IDEAL READERS AND REAL READERS

By

Charles J. Fillmore

Cognitive Science Program

Institute of Cognitive Studies

University of California at Berkeley

January 1983

In: *Georgetown University Round Table on Languages and Linguistics 1981*. Edited by Deborah Tannen. Washington, D.C.: Georgetown University Press. 248-270.

The production of the Berkeley Cognitive Science Report Series is supported by a grant in cognitive science to the University of California at Berkeley from the Alfred P. Sloan Foundation.

1. For the past year and a half I have been working on a research project which investigates the ways in which school children interact with standardized tests of reading comprehension.* The children in the study are skilled and medium-skilled readers chosen from third and fifth grade classes in two schools in Berkeley and Oakland, California.[1] The tests we have been examining are selected from those currently given in American schools to children at our subjects' grade level.[2]

An example of the kind of material the team is working with is the following passage, taken from the Metropolitan Achievement Tests and intended to be administered to third-grade students.[3]

> The carpenter was astonished that such a weird, weak-looking creature as Nasrudin was applying for a job. At dusk Nasrudin returned. 'Okay, I'll give you a chance,' said the doubtful carpenter finally. 'Take this ax and chop as much lumber as you can.'
>
> 'How many trees have you felled?' questioned the carpenter.
>
> 'All the timber in the forest,' Nasrudin replied.
>
> Shocked, the carpenter glanced out his window. There were no trees standing on the hillside. Nasrudin had destroyed the entire forest. 'Where did you learn to chop lumber?' asked the astonished carpenter.
>
> 'In the Sahara Desert,' answered Nasrudin.
>
> 'That's ridiculous!' shrieked the carpenter. 'There aren't any trees in the desert!'
>
> 'There aren't any, NOW,' said Nasrudin calmly.

The passage is followed, of course, by a series of test questions. Item (1) gives the first question.

(1) The carpenter told Nasrudin to
(a) look for another job
(b) cut down as many trees as he could
(c) go back to the Sahara Desert
(d) plant as many seeds as he could

The children confronting this question, presented as it is immediately after the reading passage, are expected to understand that they are not here being asked to continue the narrative. That is, they

---

* The work reported here has been supported by the National Institute of Education under Grant No. G-790121 Rev. 1, 'Text semantic analysis of reading comprehension tests.' Co-PIs are C. J. Fillmore and Paul Kay; others on the project are reading specialist Judith Langer and a team of Berkeley graduate students from Anthropology, Education, and Linguistics (Karen Carroll, Linda Coleman, Katharine Kovacic, Thomas Larsen, Mary Catherine O'Connor); the project has received good advice from Patrizia Violi, Mary Sue Ammon, and Haj Ross.

1 The schools in which we have been allowed to do our work are Cragmont Elementary School in Berkeley and Sequoia Elementary School in Oakland. The principals of the two schools, Dr. Benton Ng and Mr. Alfred C. Valdix, have our gratitude.

2 These have included CTBS Comprehensive Tests of Basic Skills, Level 1, Form S, Complete Battery, Expanded Edition, CTB/McGraw-Hill, Inc., copyrights 1968, 1969, and 1973; and the Metropolitan Achievement Tests: Reading Instructional Tests, Form JI, Elementary, The Psychological Corporation, Division of Harcourt, Brace, Jovanovich, copyright 1978; and (not cited in this report) the California Achievement Tests, the Stanford Achievement Tests, and the Gates-McGinnitie Reading Tests.

3 Form JI, page 15.

have to sense that the story they have just read has been finished, as far as they are concerned, and that they are now being asked to show how well they understood it. Should they by mistake construe their task as one of advancing the narrative beyond the point where Nasrudin made the boastful claim about creating the Sahara Desert, they might find it quite reasonable that the carpenter should advise Nasrudin to look for another job (since he was no longer needed here), tell him to go back to the Sahara Desert (as a kind of 'get-out-of-my-life' remark), or indeed, order him to plant as many seeds as he could (to make sure that something could get growing on the bared hillsides). The test-takers must first keep in mind the test-taking maxim, that if two answers appear to be equally good, both are probably wrong, but they must then realize that they are probably not being asked to advance the narrative. What they must remember is that in Nasrudin's probationary period, he had been given the ax and told to 'chop as much lumber' as he could. They had to figure out that 'chop lumber' is our author's unusual way of saying 'cut down trees,' and they must sense that it was the early conversation between the carpenter and Nasrudin with which the question is concerned.

The second question is shown in (2).

(2) How long did it take Nasrudin to complete the job?
(a) one day
(b) three days
(c) thirty days
(d) three years

In order to answer this question correctly, the children are required to realize that in the sentence 'At dusk, Nasrudin returned,' the phrase 'at dusk' refers to the dusk of that same day, and they must also realize that there is nothing in the story that could back up any answer with the number 'three' or 'thirty' in it. Those children who are uncommonly sensitive to language will wonder what it might mean to 'complete the job' under these circumstances, since the only task Nasrudin had been given was to 'chop as much lumber' as he could.

The third question is stated in (3).

(3) Nasrudin suggested that there were no trees in the Sahara Desert because
(a) trees can't grow in the desert
(b) no one had ever planted any there
(c) they had been destroyed by fire
(d) he had chopped them all down

The answer is that Nasrudin had chopped them all down. This was 'suggested,' to use the question's word, by Nasrudin's answer, 'There aren't any, NOW,' said after Nasrudin had explained that the Sahara Desert was the place where he had 'learned' to 'chop lumber.'

The fourth question is given in (4).

(4) After Nasrudin finished work, he
(a) left for the Sahara Desert
(b) told the carpenter what he had done
(c) applied for a new job
(d) yelled at the carpenter

The expected answer is that Nasrudin told the carpenter what he had done. The ordinary scene a reader might construct based on that description, however, is probably a bit different from what we saw in the story, so a certain amount of construing is necessary. The carpenter, it will be recalled, asked Nasrudin, ‘How many trees have you felled?’, to which Nasrudin replied, ‘All the timber in the forest.’ This utterance, an elliptical answer to a question which presupposed an understanding of what he had done, has to be construed as an instance of Nasrudin telling the carpenter what he had done. The tempting possibility that the correct answer is ‘yelled at the carpenter’ is presumably introduced to take advantage of the printer’s decision to put the word ‘now’ in capital letters, in the sentence, ‘There aren’t any, NOW.’ The capitalized word suggests shouting, so what is being tested with this foil is the child’s ability to notice that what followed ‘There aren’t any, NOW’ in the text is ‘said Nasrudin calmly.’

The fifth question appears in (5).

(5) The carpenter had not expected that Nasrudin
(a) had ever seen the Sahara Desert
(b) really needed a job
(c) would be so rude
(d) could do the job so quickly

In this item the pluperfect form, ‘had not expected,’ plays an important role. The sentence has to be situated in the text at some time point where it serves an explaining role. The text reveals that something was unexpected in the place where it shows the carpenter surprised. The text describes the carpenter as ‘shocked,’ through a grammatical device whose function may not be transparent to most third grade readers, and backs this up by showing that the carpenter spoke from then only in sentences ending in exclamation points. He had been shocked when he learned--and hence he had not expected that it had been possible--that at the end of the first day on the job Nasrudin had conquered the entire forest.

2. It is of interest to our group how young readers construct an understanding of reading passages of the kind we have just seen, and how well that understanding can be appealed to in finding ‘best answers’ to test questions about those passages.

There are two intended goals in our research, and a third that we have taken on against our will. Our first goal is to analyze reading-test items, both the passages and the test questions, in such a way as to be able to isolate and describe the kinds of background knowledge and the kinds of interpreting and integrating skills which a reader must bring to the passages in order to get out of them what their creators intended. The second goal is to find out, by interviewing our young subjects, whether they have that knowledge and those skills. We begin by analyzing tests and devising a system of annotations for them which can serve three purposes: it represents our view of the comprehension process of someone who understands the passage with no difficulty, it provides the material for our choice of the interview probes, and it gives us a framework or checklist against which we can evaluate the children’s performance with the texts. The second part of our work involves close observation of the children’s experiences with the texts, with free retellings, interviews, and metacognitive probes, looking for the presence or absence of the kinds of knowledge and skills which seem to us to be necessary for understanding them.

The third and unwelcome task which fell to us is due to the nature of the corpus we had chosen. Oddly enough, in view of the time and expense that goes into the construction of reading test items, we are dealing in this research with seriously flawed texts, texts which frequently require of their readers an uncommon degree of tolerance and cooperation. The testing industry, we have come to realize, has created a new genre for English written language, a genre whose characteristics are determined by very unnatural requirements of lexical choice, grammatical structuring, and synonym alterations, these dictated, I presume, by the intention to test knowledge of particular vocabulary items, the need to produce something which fits accepted readability formulas, which avoids gender or ethnic stereotyping, and which satisfies copyright laws with respect to the percentage of material that needs to be modified in the case of passages taken from the trade literature. This third goal of ours, then, is to understand some of the consequences of the development of this special genre, to show why it introduces complications that interfere with the ability of these tests to measure what they are designed to measure, and to see how well young children are able to master this particular genre through their repeated experiences with it.[4]

3. We need, in our work, to be able to compare real readers in their experience of these texts with the kind of reader who gets out of them everything that is needed. For this comparison we had to develop a particular kind of abstraction, something we call the ‘Ideal Reader.’ It is this invented Ideal Reader who sees the connections, creates the expectations, performs the inferences, and asks the questions which our annotation is designed to represent. It is with this Ideal Reader that our flesh-and-blood readers are to be compared.

My purpose in this paper is to explain the nature of this idealization and to show how we think we can decide on its characteristics.

Our notion of the Ideal Reader is localized to given texts and to given interpretations of those texts. That is, we do not speak of ‘an ideal reader’ in the abstract, but of ‘the Ideal Reader’ of a given text on a given interpretation. We see before us a written text; we determine what we take to be a ‘correct’ or somehow necessary interpretation of it; and we then project, from the text and the interpretation, to an invented Ideal Reader: that individual who has exactly what it takes to get from the text to the interpretation via the usual principles of compositional semantics, schema building, inferencing, goal and plan detection, and so on. (I should point out that the notion is a technical one, designating an abstraction that is of limited utility. It is not necessarily a desirable thing to be an Ideal Reader in our sense.)[5]

Once we have decided on the characteristics of an Ideal Reader (for a given text on a given interpretation), we can then ask certain important empirical questions about the text and its living readers. If, for example, we find ourselves convinced that there are in the world no real readers who

---

[4] One part of each interview is the genre question, which generally takes some such form as ‘Where do you think you might read a passage like this? In a story book? In a letter from a friend? In the Weekly Reader?’ One answer we received, in connection with an intendedly humorous narrative passage, was ‘in a CTBS test.’

[5] Research applications of the Ideal Reader concept must make use of a large and unarticulatable common sense component. We all know that there are situations in which texts communicate new information by presupposition, and very many situations in which a text develops its argument by reminding its readers of things they already know. The Ideal Reader who has the subtlety we need it to have will be able to convert, by Peircean abduction, information that is formally presupposed into information intended to be derived from the text, and will be able to recognize situations in which old information is being introduced mainly as a link in the chain of an argument or a narrative.

match certain essential characteristics of our invented Ideal Reader, then we can believe that we have on our hands a useless text. It is not written for anybody. If we find among our young subjects that the only real readers who match the conditions we have specified for the Ideal Reader are those children who have had experiences not shared by the others, especially if those experiences are distributed among the children in patterns that follow social, economic, or ethnic differences, then we have learned something about the fairness of the item or the representativeness of the test results as these are affected by the item. With items that we believe are reasonable and fair, we can examine in detail the nature of the gaps between particular real readers and our projected Ideal Reader, and ask how such gaps can be filled. Discovering such gaps is a matter of reading failure diagnosis; instituting measures for bridging the gaps by providing children with new experiences, facts, or skills, or by helping them become aware of the gaps, is a matter of improving children's reading abilities.

4. Simplifying a little, the Ideal Reader is someone who knows, at each point in a text, everything that the text presupposes at that point, and who does not know, but is prepared to receive and understand, what the text introduces at that point. Real readers, then, can differ from the text's Ideal Reader in two directions. With respect to any given point in the text, they may be underqualified, in that they do not know what the text assumes they know at that point, or they may be overqualified, in that they already know what the text introduces.[6]

(A passage in a linguistic textbook which uses the term 'contrast' without explanation but which defines, explains, and exemplifies the notion 'neutralization,' is addressed to a reader who already understands 'contrast' in its technical linguistic sense, but who does not yet know how linguists use the term 'neutralization.' The Ideal Reader for this text, then, is in part described as someone who knows, at this point, about contrast, but who does not know about neutralization. In this respect, at least, such a text appears to be perfectly reasonable. There are, we can be sure, any number of readers who meet perfectly this pair of conditions. A text in which the opposite choice had been made, however--one which explained 'contrast' but presupposed knowledge of 'neutralization'--would be a pedagogically defective text, since only readers with very unusual educational histories would be likely to satisfy just those conditions).

5. Linguists will have noticed that I have been talking about readers and written texts, whereas it is much more common in linguistic circles to talk about hearers and spoken language. Our idealization is possible for written language more than for spoken language because of the fact that written texts are more characteristically monologic and closed: in a word, they tend to be 'composed.' By contrast, we find in conversational language situations in which the interlocutors need to negotiate a common background as the conversation progresses, we find texts in which the linguistic part of an interaction is insufficient for constructing anything that could be called the intended interpretation, and we find, in dialogue between people communicating at cross purposes, texts that lack any

[6] It will seem, to many readers of this text, that my colleagues and I have a hopelessly naive view of the nature of the reading experience. People find pleasure in reading texts they only minimally understand, there are often rhythms and patterns in texts that can only be detected on repeated re-reading, readers can experience suspense and surprise in a narrative even if they know, as in a thoroughly familiar text, exactly how it's all going to end, and quite frequently texts are designed precisely to invite idiosyncratic imaginative responses in their readers--a situation for which the notion of a 'correct' interpretation is thoroughly unwelcome. It should be remembered that the simplistic view of the reading experience which we adopt for our research purposes is aimed at a level of understanding for which readers can be reasonably expected to choose 'best' answers to multiple-choice questions.

discernible point or structure.

6. While I have said that the Ideal Reader is defined with respect to a text that is composed--that is, to a text that is fixed and in some sense complete--I do not intend to imply that composed written texts can be interpreted statically. The interpreter's experience always has a clear dynamic aspect, to which our work has to pay close attention. We need to show, for example, that a text can create expectations in the reader's mind at one point which it then satisfies or subverts at a later point. The recognition of structure, development, or point; of suspense, surprise, or closure; of the interruption and resumption of a 'thread,' etc., all make up part of a reader's experience with a written text.

It is, in fact, just this dynamic aspect of the reading experiences which is the most important part of our analysis, and the most difficult part of our interview process with the children. What I mean by that is that in the interview process, we have had to be particularly sensitive to the difference between asking questions which will reveal the dynamics of the text-understanding process as it occurs naturally, and questions which themselves guide and advance that process.

We are pleased if we find a child, who, on reading the seventh sentence in a text, spontaneously says, 'Oh, now I see what's going on!', because we have just had revealed to us the workings of the comprehension process in that child. We have learned not to be pleased when a child says something like that in response to one of our interviewers' questions. Our job is to track the child's coherence-creating process, not to shape it for him.

In order to present the dynamic aspect of a reader's experience with a text, we have developed a method of text analysis which takes the text one segment at a time, asking ourselves at each point in the unrolling of the text something like, 'Having read this far, what would it have figured out, or be puzzled by, or be expecting?'. In order to make our analytic task reasonably finite, we have found it useful to keep the analyst's point of view distinct from that of both the Ideal Reader and the real readers. In particular, we, the analysts, have the advantage of knowing the whole text, knowing the point of the text, knowing how everything comes out in the end. We might know that the point of a particular passage is that it introduces a surprise. For example, if we know that the point of a particular line in a story is to reveal, for the first time, that the characters we have been reading about live in a tree, and if we feel sure that the reader is expected at this point to be surprised, that assumption imposes specific requirements on our annotation of earlier parts of the text. We have to make the Ideal Reader assume, throughout some earlier part of the text, that the people in the story do not live in trees, and at the same time we have to make sure that this assumption is known not to be explicitly bound to the actual material of the text. That is, we need an annotation which will show, at the point when the surprisng information is introduced, that the reader is supposed to feel surprised but not cheated. If we ourselves, in preparing the annotation for the text, had to examine it only one segment at a time, ignorant at each point of what was going to come up next, we would have no basis for including or excluding any part of the probably boundless number of things that ordinary readers would assume in the world of the unrolling text. That is, knowing the nature of the surprise that is coming up gives us a reason to ignore many aspects of the normal reader's envisionment of what's going on.[7]

[7] It is this criterion, in conjunction with the decision to define the Ideal Reader as relativized to given texts, which simultaneously makes our work do-able (by defining a stopping place) and guarantees that the Ideal Reader offers only limited predictability to 'good' readers. There is obviously no reason whatever for a reader to create only those inferences and embellishments which the current text will build upon.

7. The analyst sees the text as a whole. The Ideal Reader and the real reader see it one segment at a time. It is necessary for us, then, if we wish to explore the real readers' experience of the text, to present it to them one segment at a time. The general method we use works like this: we show the readers the first segment, then ask our first batch of questions; then we show them the second segment, asking them our next batch of questions; and so on. At each point, we make sure that the previously exposed part of the text is still available for scanning, rereading, consultation, etc.

The mechanics of this kind of presenting and interviewing proved to be interestingly difficult. Since some of these difficulties shed light on the process of text understanding, I believe they are worth discussing here.

At first our method was to decide on a particular segmentation of a text, type it out with each segment on a separate line, and have the subjects slide a piece of cardboard down over the text one line at a time during the interviewing process. This method introduced two major difficulties. First, as the card got lower on the page, the exposed piece of text at the top ended up looking more like verse than prose, and that, we felt, could affect some people's interpretations of what they were reading. Second, the method did not make it obvious to the subject how much of the text was left, and we felt that subjects might use different strategies for interpreting a sentence if they thought it was the closing sentence in a text than if they thought more was coming. This uncertainty made the experience quite unlike normal reading; in normal reading, we almost always know how close we are to the end, and that knowledge plays a large part in shaping our expectations and putting our interpretative faculties to work.

Our current method is more expensive and troublesome, but it has eliminated both of the difficulties presented by the sliding card method. We type the text on a sheet of paper, in the normal way, double-spaced. If the text had been given n segmentations, we make n+1 xerographed copies of it, and construct a booklet. Page 1 of the booklet has the entire text blocked out with a marking pen. Page 2 has everything blocked out but the first segment. Page 3 has everything blocked out but the first two segments. And so on. (Having found that some children have trouble finding quickly the place where they left off, we have begun using a small red dot at the beginning of the increment to help them.) With this method, then, as the text gets exposed, it looks like ordinary printed prose, and it is very clear to the subject how much of the passage is left before the end is reached. (The passages we use all fit on one page.)

The interview works something like this. The child who turns the page sees

> Once upon a time xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
> xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
> etc.

exposed a the top of the second page. Our interviewer asks something like, 'What can you tell me about this passage so far?' The subject answers, 'Well, it's going to be a story. Most likely a fairy tale.' 'How do you know that?' 'You only say "Once upon a time" when you're telling a fairy tale.' 'Do you have any guesses about what we're going to read when we get to see more of the story?' 'Well, maybe something about a poor old lady who lived in the woods, or maybe a rich king with a beautiful daughter. I don't know.'

The child then turns the page to expose the next increment, and sees

> Once upon a time there was a rich king xxxxxxxxxxxxxxxx
> xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
> etc.

The interviewer says, ‘Say, you were right, weren’t you? The sentence isn’t finished yet, is it? Do you want to stick to your guess about the beautiful daughter?’ ‘Yeah.’ The child turns the page and sees

> Once upon a time there was a rich king who had three sons.
> xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
> etc.

And so on through the text.

The new method of presentation solved some of our problems in exploring the reading experiences of real readers, but there remain a number of very serious problems connected with the piecemeal presentation, no matter how it is adapted. One difficulty with our method comes from the fact that the talk produced by the interviewer is itself a text which the children have the right to believe has a point and a direction. If, in connection with the text about the people who lived in trees, we were to ask a question like, ‘Do you think these people live in houses like everyone else?’, in order to be able to predict whether our young subjects were going to be surprised when they found out the truth, the children would quickly catch on to our purposes. If we tried to balance this by asking dozens of questions about normal readers’ default assumptions (‘Do you think they grow hair on their teeth?’, ‘Do you think they sleep at night?’), our questions would be intolerably distracting.

A second problem produced by the interviewer’s own text shows up when we are exploring the reasons for the choice of particular test question answers. With some young children, the question, ‘Why did you choose this answer?’ suggests powerfully that the answer was wrong.

One source of difficulty, then, resides in the character of the conversation with the interviewer. Another is the pacing. A paragraph that takes less than a minute to read in normal circumstances can take half an hour to go through in the segment-by-segment presentation. The interview brings so many things into the subjects’ consciousness that the simple thread that is our text can get completely lost. Using the method with adult readers on adult texts, we learned that passages that are humorous when read all at once are not humorous when given out piece by piece. I am not merely saying that the passages do not seem funny; sometimes their humorous intent is not even discerned. Furthermore, if a text takes a digression and then returns to the main theme, the return to the theme can seem very striking. In natural fast reading, by contrast, the digression itself would hardly be noticed.

With young readers, the factor of interest works against us both ways. If the story is interesting, the reader wants to get to the end to see how things turn out, and becomes impatient with all these boring questions. If the story is uninteresting, the reader wants to get the whole experience over with as soon as possible, and becomes impatient with all these boring questions. A lot depends, in short, on the warmth and charm of the interviewers, and on whatever rewards children might feel in knowing that an adult is paying very close attention to their words and thoughts.

The upshot is that while we are leading children through our kind of micro-analysis of a text, we may indeed be learning a lot about whether they are prepared to understand the text we are examining--which is, after all, our purpose--but the method itself might be preventing them from actually understanding it. A control is obviously needed, and for our control we have a second group of children who are exposed to the text in a different way. These children read the passages, answer the test questions, retell what they have read in their own words, and only then submit to the piecemeal presentation. The questions for this group are somewhat differently formulated: ‘What were you thinking when you read this?’, ‘Do you remember what the next sentence is?’, and so on.

8. As we see it, the main dynamic aspect of the reading experience is that of constructing and revising an Envisionment of the ‘world of the text,’ some coherent ‘image’ or understanding of the states of affairs that exist in the set of possible worlds compatible with the language of the text. We do not intend the word Envisionment, which we have borrowed from John Seely Brown (personal communication), to suggest too strongly the visual aspects of a text world, but at the same time, we recognize that in the kinds of texts we are dealing with, the visual aspects do, in fact, predominate.

We have found it useful to distinguish various confidence levels of material in the Envisionment, according to whether such material is explicitly justified by the linguistic material of the text, whether it came into being by inferences which the text is seen as clearly inviting, or whether they represent interpretations which result from schematizations brought to the text to situate its events in common experience, but which do not follow necessarily from anything the text has provided. In our discussions of real readers, we need still another level of Envisionment: ways in which the world of the text has been shaped by the idiosyncratic experiences and imaginings of individual readers.

To illustrate the theory of levels of Envisionment, we can consider the following three-line text, borrowed by Marcelo Dascal (personal communication).

> The princess ate some jam.
> The queen slapped her.
> The princess cried.

At what we call the E^0 level of Envisionment, there are only those states of affairs that have to hold in the world of the text for the individual sentences to be separately true. In our case, the Envisionment at E^0 has three disjoint parts: somebody who is a princess eats some jam; somebody who is a queen slaps a female being; and somebody who is a princess cries.

At the second level in the Envisionment, which we call E^1, it is assumed that we are dealing with a cohesive text, rather than with three independent statements, and so the princess in Sentence 1 is the same as the princess in Sentence 3, and is the *her*, the one who is slapped, in Sentence 2. The queen in Sentence 2 is furthermore taken to be the princess’s mother.

We have here the workings of a kind of Parsimony Principle in text comprehension. If, in constructing an understanding of the first sentence, we had to instantiate a Royal Family schema, we find it parsimonious to use that same schema instantiation for identifying the queen, too. Hence the queen and the princess are mother and daughter. If at the beginning of the second sentence (‘The queen slapped her.’) we have two people ‘on stage,’ the princess and her mother the queen, then the *her* of this sentence must refer to the princess. The Parsimony Principle is a text-interpretation maxim that says something like: Don’t bring more people or props into the text world than are needed to make the text cohere.

In the example just given, the difference between E^0 and E^1 amounted to a difference between individual sentences and sentence sequences. It was when we put the sentences together that we began needing interpretations at the E^1 level. The difference between E^0 and E^1 can also be illustrated with single sentences.

The favored kinds of sentences used in psycholinguistic experiments typically have only an E^0 interpretation. In a sentence like ‘The actress sold a fish to the carpenter,’ the kind of sentence you might be expected to learn or shadow or associate something with in a psycholinguistic experiment, we find that there is nothing we know about either actresses or carpenters that we can call on to motivate the scene of the one selling a fish to the other. If we compare that with a sentence like ‘The cobbler sold a pair of boots to the mountain climber,’ we find that the latter invites more interpretative work. The Parsimony Principle would induce us to make the hypotheses that the cobbler sold boots that he had made in his shop, and that the mountain climber was buying boots that he would use in mountain climbing. That is, we tend to assume that it is as a cobbler and as a mountain climber that the two participants were engaged in this particular commercial event.

These assumptions are made at the E^1 level, rather than at the E^0 level, because we could, without contradiction, immediately find out that these hypotheses were completely wrong. It could turn out that the cobbler was sitting in for his wife, a shoe store clerk, while she was at the dentist’s, and that the person described here as the mountain climber was buying the boots to give to his mother for Mother’s Day.

The sort of ‘default’ interpretations that we assign to the Envisionment at the E^1 level remain tentative for a very brief time. Knowing that we are dealing with a composed text, we somehow believe that if we were not meant to give this interpretation, the author owes us an explanation very soon. If the text does not very quickly turn us away from the E^1 interpretation, the normal reader quickly converts it to an E^0 confidence level.[8]

At the third level of the Envisionment, which we call E^2, we ‘situate’ elements of the passage in terms of our knowledge of goals and institutions and folk theories of human nature. With the story of the princess and the jam, we make sense of what is going on by assuming that the queen’s act of slapping the princess was in punishment for the princess’s having eaten the jam, and we assume that the princess’s tears are in response to the queen’s slap, showing pain, remorse, or shame.

The next higher level, E^3, contains particular real readers’ embellishments of the text world, the filling in of details not motivated by the text itself. For our story we might find, in particular readers’ experiences of it, the assumption that the queen had wanted the jam for herself and was very selfish, or maybe that the princess had recently been made to promise to stay on a diet until she found a husband. In our annotation for the Ideal Reader, we obviously do not need such a level; but in our representations of the subjects’ interpretations of our texts, we definitely do.

My co-workers and I have no stake in the proposed number of levels, and we will not be surprised to find out that the phenomena we are concerned with do not at all lend themselves to a description in terms of discrete ranked tentativity levels. What we are sure of--and this is the main practical purpose of the levels theory--is that a child who is not able to do the kind of inferencing or ‘reading between the lines’ that is necessary for filling out material at what we are calling E^1 and E

[8] The Parsimony Principle is found in essentially the form suggested here in such works as Harvey Sacks (1972), ‘On the analyzability of stories by children,’ in J. Gumperz and D. Hymes (eds.), *Directions in Sociolinguistics*, Holt, Rinehart and Winston, Inc.; and Yorick Wilks (1973), ‘Preference semantics,’ in E. Keenan (ed.), *Formal Semantics of Natural Language*, Cambridge University Press, and (1973), ‘Understanding without proofs,’ in *Proceedings of the Third International Joint Conference on Artificial Intelligence*, Stanford Research Institute, Menlo Park, California. The Wilks references were brought to my attention by Robert Kirsner.

^{2} is not a good reader, nor is the child who fills in lots of E ^{3} material and immediately assumes that it was in the text. A child who reads very slowly, figures out the meanings of the sentences in the text one at a time but never puts them together into a cohesive whole, is a poor reader in the first sense. Such a reader cannot go beyond E ^{0} Poor in a different way would be the reader who, in reading our story about the princess with the sweet tooth, would unhesitatingly answer ‘Strawberry’ to the question, ‘What kind of jam did the princess eat?’ Such readers overtrust their own E ^{3} embellishments of a text.

9. The lexical and grammatical material of the unfolding text creates or identifies for the reader the conceptual tools needed for constructing the Envisionment. The main kind of tool I have in mind is the construct (or set of constructs) variously known as ‘scripts,’ ‘frames,’ ‘schemata,’ ‘folk theories,’ ‘cognitive models,’ and the like. These constructs are intended to represent the knowledge structures with which our experiences with the world are held together. Sometimes in a given portion of the text such tools are merely activated, merely made available for later use. At other times they are actively used in constructing and maintaining the Envisionment. In our annotation we make a great deal of use of labeled schemata, with which we appeal, informally, to structured knowledge that the Ideal Reader is believed to have and which real readers may or may not have, may or may not be able to use, or may or may not know when to use.

Two applications of such schemata have already been illustrated: the Royal Family schema, which allowed us to see instantly how the princess and the queen were related to each other, and the schemata we used in bringing coherence to the sentence about the cobbler and the mountain climber. A slightly more complicated one-sentence example is: *The defendant had forged the will*. Somebody called the *defendant* is seen as a person who plays a particular role in a Criminal Trial schema, that of the Accused; the object called the *will* in our example is the Legal Instrument in an Estate Inheritance schema; and the act of forging is the central act in a Forgery schema. The character of each of these schemata could be spelled out in considerable detail. The Parsimony Principle would invite us, at perhaps an E ^{2} level (the level at which we seek to bring in outside explanations) to imagine that the act of committing the forgery was one of the acts for which the defendant was currently involved in the judicial process, and to imagine that the will that got forged had the defendant named as a beneficiary.

The schemata just mentioned have all been brought into the Envisionment by virtue of particular lexical items that are keyed to particular roles or steps in their associated schemata. In the comprehension process, the interpreter activates the schema connected with such lexical items and works at building an Envisionment out of knowledge derived from these schemata. This is done partly by combining primary schemata into larger assemblies or networks, sometimes, as we have seen, by bringing in schemata not explicitly indicated by material in the text but needed for holding the other schemata together.

There are various ways in which schemata can be linked together. Some are linked in semantic memory by what can be called Knowledge Links (K-links), connections between schemata provided by general knowledge, independently of any information provided by the actual present text. For our forgery sentence, K-links bring to play our knowledge that all of the schemata lexically introduced by that sentence somehow fit into a larger schema that could be called Judicial Process, and make us aware that a defendant in the criminal trial has been accused of a crime, and that forgery is a crime.

The text itself, which tells us about some ‘world,’ links together instantiations of schemata by anchoring them to each other in that world. These can be called Text Links (T-links). An E ^{0} T-link in our forgery sentence established, through the grammar of the sentence, that the person who is the defendant in a criminal trial is also the perpetrator in an act of forgery, and that the product of that act of forgery was (or has been put forth as being) a legal instrument in a matter of estate inheritance.

The schemata that we have been considering so far are schemata that operate in what with many others we have come to call the Content domain. We distinguish three domains for the Ideal Reader's speculations, puzzlements, and conclusions: the domain of 'Content,' by which we have in mind the properties and events in the world of the text; the domain of 'Text,' with its schemata of grammatical structure and text structure; and the domain of 'Genre,' where we have in mind those structures of expectation that come with knowing that one is dealing with a folk tale, a detective story, an obituary, a reading test, or the like.

Within the Text domain, our forgery sentence has, in addition to the grammatical structurings which aided us in providing the T-links, a number of features which mark it as 'text-internal.' We recognize that a sentence like *The defendant had forged the will* is not likely to be the first, and certainly not likely to be the only sentence in a text. This judgment we base on the use of the definite article (*the defendant*, *the will*) and the use of the pluperfect (*had forged*). The text-schematizing necessarily associated with them shows that there must be a presupposed temporal reference point in the part of the larger text where this sentence can occur (the point on which the pluperfect is semantically anchored) and a pre-established setting within which the descriptions *the defendant* and *the will* are uniquely identifying.

Genre schemata arise from structured expectations created by familiarity with particular genres. If we read in a folk tale that the king's two older sons have both failed to slay the dragon, we are filled with hope when we learn that the king's third and youngest son has set out to try his hand. Were the story to end with the third son being slain by the dragon, we would feel that we had just been exposed to a new and cynical derivative genre, not that we were wrong in forming the expectations we had formed. When we read questions in a reading test, we know that we are not being asked to figure out a clever way to finish the passage we have just been reading, but rather that we're being asked what we remember, or what we can now figure out, about what the passage told us.

10. There is one more device that we need for our complete annotation, and that is a representation of Point of View. It has played a relatively small role in the work we have done so far, so I will say little about it. For some texts, especially texts involving histories, descriptions of objects or terrains, and narratives, the text world not only has the properties it has in some objectively describable way, but its properties are presented to us from a particular point of view: there is an observer, an experiencer, or a camera's eye, and perhaps always a temporal reference point.

I illustrate the notion merely by showing that some sequences of sentences in a text can be seen to cohere only when they are taken as sharing a single point of view. The sequence

> He was coming up the steps. There was a broad smile on his face.

is coherent in the sense I have in mind, since the position from which his ascending the stairs can be described as 'coming' is also the position from which it would be natural to see his smile. Analogously, a sequence like

> He was going up the steps. There was a wad of bubble-gum on the seat on his pants.

is also coherent, given a constant viewpoint. By contrast, a sequence like

> ??He was coming up the steps. There was a wad of bubble-gum on the seat on his pants.

seem quite bizarre if taken as a visual description, and the sequence

> He was going up the steps. There was a broad smile on his face.

seems to require us to imagine a side view. The recognition of Point of View figures in any complete account of lexical semantics (with respect, for example, to such words as *come* and *go*), the semantic role of grammatical categories (involving, for example, tense and definiteness), and the cooccurrence or sequencing of sentences in a text.

11. The descriptive framework we have ended up with includes: a streamlined record of the grammatical parsing of the sentence; the schemata of Content, Text, and Genre which are introduced by lexical and grammatical form, in the first instance, and from the interpreter's repertory of schematizing devices, in the second instance; the questions and expectations raised in the reader's mind concerning the still unrevealed portions of the text; the growing and changing Envisionment of the world of the text, with changes in the levels of tentativeness as the text develops; and point of view.

Given such machinery and the goals of our project, we seek to annotate texts in a way which represents the Ideal Reader's processing of them and which provides the basis for the construction of probing devices designed to assess a real reader's experience of the text.

I would like to return to the passage I started out with, the story about the uncommonly gifted woodcutter, to show something about how this annotation is used in our work.

The method requires an initial segmentation of the text, since we are assuming, for theoretical purposes and contrary to fact, that readers bring their full armament of interpreting skills to each piece of text without peeking at the next one. The segmentation we provide is, of course, arbitrary; since we are not equipped to carry out the kind of research it would take to find out what natural segmentations real readers would make, we have to be arbitrary. In general, our segmentation procedure amounts to taking one clause at a time.

In my demonstration I am going to do the first segment of the Nasrudin story with some degree of completeness, and then just briefly touch on a few other aspects of the passage. The first segment is

> The carpenter was astonished XXXXXXXXXXXXXXXXX

The Ideal Reader will take note that this is the first segment of the text, that it is not a complete sentence (since there is no following punctuation), that it has a definite noun phrase as its subject, a passive factive complement verb as its predicating element, and the simple past as its tense.

The lexical item *carpenter* introduces a Carpentry schema, which is merely another way of saying that when we understand the noun *carpenter* we do so by knowing something about what carpenters do. (It is a word for which one would be hard put to distinguish semantic information about what the word means and practical information about what carpenters do). The person designated as the carpenter is the practitioner in a work schema, the other elements of which include the tools, the materials, and some notion of the products of a carpenter's work. Our schematic knowledge about carpentering could be spelled out in great detail, with, for example, different types of activities associated with particular tools; but since we, the analysts, know how the story develops, we know that what we need for our Ideal Reader is a mere skeleton of the whole schema, with only certain aspects drawn out in detail, these involving the carpenter's need for wood. From that we get the K-links that connect wood with trees and the cutting and milling of lumber.

The next major lexical item in the segment is *astonished*. This word, we allow ourselves to say, introduces an Astonishment schema. To understand what it means to be astonished is to have an outline knowledge of the kind of scenario that could lead a person to have the experience which this word describes. In this schema there is at least the experiencer of the emotion, a perception on the experiencer's part which triggers the emotion, the event which was responsible for that perception, and a history or set of expectations against which this causing event stood out in sharp conflict. (One is astonished when one notices that something very unexpected has occurred).

There are no K-links between Carpentry and Astonishment. There is, of course, a T-link: it is the workman in an instantiation of the Carpentry schema who is the experiencer in an instantiation of the Astonishment schema.

In addition to the Content schemata of Carpentry and Astonishment, we recognize in this segment certain schemata in the Text domain. These include the grammatical knowledge that the verb *astonish* is a factive complement verb, giving rise to the expectation that what follows is going to be a description or mention of the event which, in the world of the text, caused the astonishment experience. When the next segment is exposed, the Ideal Reader is prepared to integrate the information it provides into the Content domain of the text-world as the cause of the carpenter's astonishment. More superficially, knowledge of the kind of word *astonish* is primes the Ideal Reader to expect that the very next segment will be *at*, *to*, or *that*, followed by a nominal, verbal, or clausal description of the causing event.

The expectation that what follows is to be an account of the causing event comes not simply from the form of our segment, but is further informed by the knowledge that the segment is not a complete sentence. There could be a text which read

> The house collapsed. The carpenter was astonished.

where our segment was a complete sentence. But that is not what we have here.

There are also Text-domain schemata associated with the definite article and the simple past tense, and their activation invites the Ideal Reader to hypothesize that what we have here is the beginning of a particular type of narrative, one in which characters and props can be introduced with definite noun phrases or proper nouns, and in which a presupposed time point can be presented without explanation. These schemata, and the conditions in which they are introduced, invite the Ideal Reader to activate a particular schema from the Genre domain. There are kinds of third-person narrative in which such *in medias res* textual features are common and proper. (I am only pretending, of course, to have access to a theory of genres. Since our collection of texts is made up mostly of narratives, expository prose, personal letters, and simple poems, there are not really many distinctions of genre that we have had to worry about. If we are to choose from that limited inventory of genres, we are already fairly safe, after even this first short segment, to guess that we are in a narrative.)

After exposure to this simple segment, *The carpenter was astonished...*, the Ideal Reader has two Content schemata instantiated, has T-links between them provided by the grammatical relations present in the segment, has expectations about both the content and the form of what the next piece of the sentence is likely to be, and has an active hypothesis about what kind of a text it is dealing with. Questions the interviewer might ask of a real reader, on the presentation of this first segment, to see how closely this real reader matches the accomplishments of the Ideal Reader, include the following:

> Do you know what a carpenter does?
> What does it mean to be astonished?
> When we turn the page and look at the next part of the paragraph, what do you think we're going to find out about?
> Can you guess what the next word is going to be?
> That's a pretty good guess. Can you think of anything else?

So much for the first segment. With third graders we tend to postpone questions about the genre until at least one or two complete sentences have been exposed.

In the next segment, Nasrudin is described as *weak-looking*. A Physical Strength scheme, or scale, has to be introduced into the Ideal Reader's awareness, with the knowledge that weak and strong are the two extremes, and the associated knowledge that somebody who is strong can do more work and heavier work than somebody who is weak. This, of course, turns out to be related to the carpenter's surprise, since Nasrudin looked weak but was asking for work that required strength. The second part of the word *weak-looking* raises in the Ideal Reader's mind the question of whether somebody who is described as weak-looking really is weak. A predication of appearance naturally invites a question about reality. That question, you will recall, gets answered very soon in the story.

We can clearly put the Parsimony Principle to work with the first full sentence, which is: 'The carpenter was astonished that such a weird, weak-looking creature as Nasrudin was applying for a job.' Without the Parsimony Principle, we could imagine the carpenter peeking in the door of a personnel office, or an employment bureau, and seeing Nasrudin standing in line. With it, we try to use the characters in our scene maximally for filling out the introduced schemata. In the applying-for-a-job scene, we make Nasrudin the applicant and the carpenter the interviewer. In the work for which Nasrudin is applying, we make Nasrudin the potential employee and the carpenter the potential employer.

These T-linking assumptions can be wrong, so we will put them in E^1 and keep them there until we see whether or not they get immediately corrected. Not only do they not get corrected: the next sentences cannot be made intelligible unless they are true. The next words of the carpenter are, 'Okay, I'll give you a chance. Take this ax and chop as much lumber as you can.' The relationship between interviewer and interviewee is suddenly transformed into a relationship between employer and employee.

The text offers the Ideal Reader a number of places to 'read between the lines.' The first sentence, we have seen, is apparently about the carpenter's inner life: it reports that he was astonished at what he saw. The very next sentence is a record of the carpenter's words: 'Okay, I'll give you a chance.' The first word, 'Okay,' can count as a signal of assent; but a signal of assent does not follow naturally a statement about that person's mental experiences. The first sentence has to be understood as introducing, by presupposition, the first move in the narrative. We then construe Nasrudin's 'applying for a job' as an act of asking the carpenter something like 'Will you give me work?' It is to that appeal, not explicitly present in the text, to which the utterance, 'Okay, I'll give you a chance,' is a response.

Similar cooperation is required of the reader in the transition between the following two sentences, the first one in quotes:

> 'Take this ax and chop as much lumber as you can.'
> At dusk Nasrudin returned.

To make this sequence cohere, the Ideal Reader has to build into its picture of the story the information that the carpenter held in his hands or otherwise indicated the ax, that he spoke the quoted sentence to Nasrudin, that Nasrudin took the ax and went off to the forest to cut down trees with it, and that it was at dusk of the same day that Nasrudin returned to the place where the carpenter was. The text tells us none of that but we have to believe it. The Parsimony Principle would have us believe that the carpenter's order was obeyed, since the author does not tell us that it was ignored or defied. We are, in short, not free to believe that Nasrudin might have stared at the ax all afternoon, giving up his ambitions as a carpenter's helper and returning at dusk to Istanbul.

The patterns of inferencing in this story reach great complexity, especially in the place where the reader figures out that Nasrudin has claimed responsibility for creating the Sahara Desert, and where the reader realizes in the domain of Genre that our narrative is a joke.

The reader of my text will have figured out (with relief or disappointment, I cannot guess which) that the few remaining pages are not going to provide a complete account of the Nasrudin story.

12. I explained earlier that one of our project's unwelcome problems was that of coping with flawed texts. The Ideal Reader abstraction works out most satisfactorily and most straightforwardly with well-constructed texts, texts in which the author's plans are discernible and in which the author's assumptions about the text's readers are reasonable. Almost all of the texts we have examined in our work are seriously flawed in one way or another.

Here is one, picked almost at random:

> In l877 a machine appeared which surprised many people. Can you guess the name of this strange new machine? As you spoke into the mouthpiece and turned the handle, a tube covered with a thin piece of tin moved around. As the tube moved a needle pressed deep lines into the tin. As you turned the handle once more, the needle touched against the same lines and played back your words. This was the first phonograph. How different from the hi-fi of today![9]

Not only do our third-grade subjects have difficulties with this text; so do many very literate adults. Our subjects had a much worse time of it, however. Although these children know the words *recordplayer* and *stereo,* not one of them knew the word *phonograph*. Where the text tells us that

---

[9] From Comprehensive Tests of Basic Skills, Form S, Level 1, page 8. Reproduced by permission of the publisher, CTB/McGraw-Hill, Del Monte Research Park, Monterey, California 93940.

the tube moved, those adults who know what the antique roll-type phonograph looked like are able to picture a cylinder rotating. The children knew about toothpaste tubes and video tubes, but that knowledge did not help them at all. Not one child was able to create anything remotely resembling the intended image. The pictures they were induced to draw of the machines they imagined ranged from tractors to microphones to Coca-Cola machines.

One attitude we could take toward such a passage, is that it is merely difficult, and that nobody has a right to complain about finding difficult items in a test. Our interview protocols on this item convince us, however, that it simply does not belong in a reading test. The one child (out of 30 subjects) who figured out that the machine was a record-player could answer very few of the questions correctly, largely because she became confused by the word *phonograph*, which she was unfamiliar to her. (She pictured it as a record player until she came upon the word *phonograph*, and then she became uncertain.) By contrast, a boy who misread the word *machine* as *magician* and who performed the wildest sort of mental acrobatics in order to preserve that part of his Envisionment, was able, by absolutely absurd reasoning plus a certain amount of test-taking know-how, to choose mainly correct answers to the test questions.

The Ideal Reader for such a test has to be, quite simply, a person who happens to know what the oldest phonograph looked like and who can be cooperative enough with the text, once the information that it is a phonograph is finally provided, to realize that the description of the machine is some author's poor attempt to describe what the Ideal Reader already knows it should look like. For a text like this, the concept of the Ideal Reader is workable, but pointless. In the usual case, we as readers have to know something about the real world in order to build on that to construct an Envisionment of the world of the current text; in a case like this, however, what we have to know in order to understand the text exhausts what the text tells us.[10] That, in my mind, is a clear case of a bad text, most assuredly a bad test item.

13. Forthcoming from the project on which I am reporting are an analysis of the interview data from our subjects and a critical study of the worst of the test items. Here I have merely tried to give an informal account of the goals and methods of the project and a survey of the problems one has to face when trying to monitor closely the reading experience, and to show the relevance to this enterprise of the Ideal Reader abstraction.

[10] The case has been slightly overstated. One does learn in this passage that the first phonograph appeared in 1877.