admirable writers as Theodore H. White, John Kenneth Galbrai[th,] Lewis Thomas, Peter Medawar, and Richard Feynman, who a[re] able to communicate their complex expertise to a wide audience [of] educated people. We will be able to achieve a just and prospero[us] society only when our schools ensure that everyone comman[ds] enough shared background knowledge to be able to communica[te] effectively with everyone else.

CHAPTER II

The Discovery of the Schema

RESEARCH ON BACKGROUND KNOWLEDGE IN READING

In 1985, Dr. Robert Glaser, president of the National Academy of Education, wrote in the foreword to a national report on reading that the past two decades of scientific research "have produced an array of information which is unparalleled in its understanding of the underlying processes in the comprehension of language."[1] What so impressed Glaser and his distinguished copanelists on the National Commission on Reading? For one thing, new research has shown that reading doesn't follow an orderly pattern, as used to be thought. We don't first identify words, then word meanings, next combine word meanings to get the meanings of sentences, and finally combine sentence meanings to get the meaning of a whole text. This model isn't wrong in all respects, but we know that it is so oversimplified and incomplete that it presents a highly misleading picture of the way we understand texts.

The new picture that is emerging from language research is more complicated and more useful. It brings to the fore the highly active mind of the reader, who is now discovered to be not only a decoder of what is written down but also a supplier of much essential information that is not written down. The reader's mind is constantly inferring meanings that are not directly stated by the words of a

text but are nonetheless part of its essential content. The expli
meanings of a piece of writing are the tip of an iceberg of mean
the larger part lies below the surface of the text and is comp
of the reader's own relevant knowledge. The past two deca
research have shown that such background knowledge is a f
important ingredient in the reading process than earlier th
accounts had supposed.

The best introduction to the new research is a ps
insight that goes back some thirty years to George M
breaking work on short-term memory, which is a sp
of our minds that lasts just a matter of millisecon
guished from long-term memory (what most of us
ory), which lasts from a few seconds to a lifetir
that our capacity for remembering briefly pres
items is severely limited. The mind cannot re
term memory more than about four to seven s
numbers, letters, points on the body, degre
of warmth, and so on. Miller's observatic
items in every domain of experience, n
language. His discovery of this univers
number of further insights into the way
understand language.

If you want to test Miller's discovery about the
term memory you can quickly do so with a series of nu
letters, or indeed with any array of items. You will not relia
remember more than four to seven discrete items presented once without rehearsal, a limit on memory experts as well as ordinary people. Memory experts have worked out indirect tricks to circumvent this universal limitation, but they can never directly overcome it, because the constraint is hard-wired into our systems. Whenever we try to perform a task that consistently forces us to exceed the capacity of short-term memory, we fail to perform it reliably.

One device that memory experts have developed to resist this built-in limitation is "chunking." We can remember telephone numbers and other series of discrete items if we chunk the disparate items into a smaller number of groupings. The number 8-0-3-2-9-6-2-6-3-2 is hard to memorize, but it becomes much easier when it is treated as 803-296-26-32. What is true of numbers is also true

are able to connect past moments with what we are no
or reading. But if short-term memories are so limited,
manage to remember meaning? Bever's explanation is
ported by the current evidence. Bever says that lang
ferred from short-term memory into long-term m
literal recollection of words but as a shorthand
gist, which normally erases from memory man
words. On the basis of this account, we co
memories for the literal words of sentences sh
our long-term memories for their gist shoul
is indeed the case, and this fact has fund
understanding the nature of reading.

From an evolutionary standpoint it
we forget the literal words we encou
ings than that we have an absolu
meanings. If we remembered all
ences, including all the words
disks in our minds would qu
have to erase them periodi
memories with detail, we a
meanings with quite rema
probably recognize or r
tences of this paragra
surface form is lost

In fact, experim
are lost to memo
reported a dem
jects narrative
tained a tar
subjects w
discover
the ori
it to
this
It

he very sentence they had heard. Subjects given the second
d that it was not the sentence they had heard. Sachs in-
these results, and later experimentation has amply con-
ference, that the original form of a sentence is rapidly
y, whereas an accurate memory for its meaning is

nbaum showed that subjects don't notice that a
ubstituted in a test sentence.[5] In one experiment,
number of original sentences, including "The
" and were then given multiple-choice rec-
this one:

sed. [correct]

rves the gist]

han any other incorrect choice,
ed an accurate memory for
by a substitution like *open*
substitution like *cold* for
t necessarily imply *cold*.
he man was not tall,"
e man was short,"
ho is not tall isn't
discovered about
not closed and
s understood,

Ev
units th
of short-term
"(1) The clause is
we assign semantic re
clause is processed, it is reco
leaving short-term memory availa

1975 by
subjects
th rec-
eners
teen
is-

This picture of the way we interpret
to the question, How do we remember
conversations and books if short-term memory
viously we manage to remember what we have hear

text but are nonetheless part of its essential content. The explicit meanings of a piece of writing are the tip of an iceberg of meaning; the larger part lies below the surface of the text and is composed of the reader's own relevant knowledge. The past two decades of research have shown that such background knowledge is a far more important ingredient in the reading process than earlier theoretical accounts had supposed.

The best introduction to the new research is a psychological insight that goes back some thirty years to George Miller's pathbreaking work on short-term memory, which is a special function of our minds that lasts just a matter of milliseconds.[2] It is distinguished from long-term memory (what most of us think of as memory), which lasts from a few seconds to a lifetime. Miller noticed that our capacity for remembering briefly presented disconnected items is severely limited. The mind cannot reliably hold in short-term memory more than about four to seven separate items, whether numbers, letters, points on the body, degrees of loudness, degrees of warmth, and so on. Miller's observation holds true for discrete items in every domain of experience, not just those presented in language. His discovery of this universal constraint has fostered a number of further insights into the way we perceive the world and understand language.

If you want to test Miller's discovery about the limits of short-term memory you can quickly do so with a series of numbers or letters, or indeed with any array of items. You will not reliably remember more than four to seven discrete items presented once without rehearsal, a limit on memory experts as well as ordinary people. Memory experts have worked out indirect tricks to circumvent this universal limitation, but they can never directly overcome it, because the constraint is hard-wired into our systems. Whenever we try to perform a task that consistently forces us to exceed the capacity of short-term memory, we fail to perform it reliably.

One device that memory experts have developed to resist this built-in limitation is "chunking." We can remember telephone numbers and other series of discrete items if we chunk the disparate items into a smaller number of groupings. The number 8-0-3-2-9-6-2-6-3-2 is hard to memorize, but it becomes much easier when it is treated as 803-296-26-32. What is true of numbers is also true

of letters. After just one reading you can remember a short sentence like "The cat is on the mat," and later you can recall in correct order all the sixteen individual letters that compose the sentence. But after just one brief presentation, without the chunking provided by the word groupings, you would not be able to remember sixteen separate letters.

These observations are connected with the new insights about language processing referred to by Dr. Glaser. Consider just one implication of the constraints of short-term memory. How did I remember "The cat is on the mat?" I could manage it for subtler reasons than the length of the sentence — fewer than seven words. I can also remember "The cat is on the road to Mandalay" and "The cat is on the verge of a severe nervous breakdown." None of this exceeds my memory capacity. Why? Because I am able to chunk words into phrasal units and apply my long-term memory capacities as well. How do I manage this?

Short-term memory is the mind's vestibule where incoming items enjoy a brief equality lasting just long enough for the mind to give them a structure. Short-term memory holds in momentary suspension items that have come in one after another, thus enabling us to convert them into a stable structure. Amidst the temporal flow of language, short-term memory allows us to form a few words into nontemporal structures. Then we transfer those structures to long-term memory, leaving short-term memory free to deal again with the temporal flow of incoming words.

The limitation imposed by short-term memory accounts for the fact that all languages must form brief bursts of words into clauses. Every known language divides its sentences into semicomplete clausal units that are small enough to be structured within the limitations of short-term memory. As the psycholinguist T. G. Bever explains: "(1) The clause is the primary perceptual unit; (2) within each clause we assign semantic relations within major phrases; (3) after each clause is processed, it is recoded into relatively abstract form, thereby leaving short-term memory available for processing the next clause."[3]

This picture of the way we interpret language provides an answer to the question, How do we remember the meanings of whole conversations and books if short-term memory is so fleeting? Obviously we manage to remember what we have heard or read and

are able to connect past moments with what we are now hear or reading. But if short-term memories are so limited, how do manage to remember meaning? Bever's explanation is strongly su ported by the current evidence. Bever says that language is tra ferred from short-term memory into long-term memory not as literal recollection of words but as a shorthand recoding of the gist, which normally erases from memory many of the individu words. On the basis of this account, we could predict that ou memories for the literal words of sentences should be poor, wherea our long-term memories for their gist should be quite reliable. Tha is indeed the case, and this fact has fundamental implications fo understanding the nature of reading.

From an evolutionary standpoint it is much more desirable that we forget the literal words we encounter and remember their mean- ings than that we have an absolute memory for both words and meanings. If we remembered all the surface details of our experi- ences, including all the words we ever heard or read, the floppy disks in our minds would quickly fill to capacity, and we would have to erase them periodically. Instead, without cluttering our memories with detail, we are nonetheless able to recall and recognize meanings with quite remarkable subtlety and accuracy. Readers can probably recognize or recall the meaning of the previous two sen- tences of this paragraph, but not their precise form. Memory of surface form is lost very quickly.

In fact, experimenters have found that surface forms of sentences are lost to memory within a few seconds. In 1967 Jacqueline Sachs reported a demonstration of this phenomenon.[4] She read her sub- jects narratives that were one paragraph long. Each paragraph con- tained a target sentence. Shortly after hearing the target sentence, subjects were presented either of two incorrect test sentences, to discover if they had accurately remembered the precise words of the original. A typical original sentence was "He sent a letter about it to Galileo, the great Italian scientist." The two test sentences for this original were "A letter about it was sent to Galileo, the great Italian scientist" and "Galileo, the great Italian scientist, sent him a letter about it."

Half the subjects were given one incorrect recognition probe and half the other. Those who were given the first probe responded that

it was the very sentence they had heard. Subjects given the second responded that it was not the sentence they had heard. Sachs in- ferred from these results, and later experimentation has amply con- firmed her inference, that the original form of a sentence is rapidly lost to memory, whereas an accurate memory for its meaning is retained.

In 1966 S. Fillenbaum showed that subjects don't notice that a synonym has been substituted in a test sentence.[5] In one experiment, subjects were given a number of original sentences, including "The window is not closed," and were then given multiple-choice rec- ognition tests, including this one:

a. The window is not closed. [correct]
b. The window is closed.
c. The window is not open.
d. The window is open. [preserves the gist]

Subjects selected (d) far more often than any other incorrect choice, showing that they had indeed preserved an accurate memory for meaning. But although they were fooled by a substitution like *open* for *not closed* they were not fooled by a substitution like *cold* for *not hot*, presumably because *not hot* doesn't necessarily imply *cold*. Similarly, when the original sentence was "The man was not tall," subjects did not assume they had heard "The man was short," presumably because they realized that a man who is not tall isn't necessarily short. Most remarkable of all, it was discovered about a decade later that inferences regarding *open* versus *not closed* and *short* versus *not tall* are made at the time a sentence is understood, not later in recognition or recall.[6]

Other tests of synonym substitution were made in 1975 by W. F. Brewer, who found it to be a robust phenomenon. His subjects changed surface forms but preserved meaning in tests of both rec- ognition and recall.[7] And in 1974 Eric Wanner showed that listeners had already lost the exact wording of a sentence when only sixteen syllables intervened between the sentence and the test.[8] Other dis- cussions and experiments regarding memory for meaning and loss of memory for surface forms of language abound in the experi- mental literature.[9] The two allied phenomena are now firmly es-

tablished: it is known that accurate memory of surface form normally disappears in a few seconds, while accurate memory for gist is quite persistent.

The implications of such phenomena became a focus of psycholinguistic research within a few years after the Fillenbaum and Sachs papers of 1966 and 1967. The next phase of work started in 1972 with an important paper by J. R. Barclay, J. D. Bransford, and J. J. Franks.[10] In a series of experiments they found that our initial understanding of a text depends on our applying relevant background knowledge that is not given in the text itself. Dividing their subjects into two paired groups, they conducted a series of recognition tests in which each group was given one of two slightly different sentences. One of these sentence pairs was as follows:

1. Three turtles rested *beside* a floating log, and a fish swam beneath them.
2. Three turtles rested *on* a floating log, and a fish swam beneath them.

Notice one difference in these sentences. In sentence 1, we do not necessarily infer that the fish swam under the log, and in fact it may not have done so. But in sentence 2, we can infer from our knowledge of logs, fish, and swimming that the fish *did* swim under the log. Subjects were later asked whether they recognized test sentences, of which this was one:

3. Three turtles rested (*beside/on*) a floating log, and a fish swam beneath *it*.

The (*beside/on*) notation means that if subjects had been given the *beside* version of the original, they got the *beside* version of the recognition test, and if they had been given the *on* original, they got the *on* recognition test.

At stake for the experimenters was a choice between two conflicting hypotheses about the way we understand language. According to one theory, we interpret the meanings of clauses and sentences and store them in long-term memory. But according to the other theory, we construct an elaborated model of what the

words *imply* and store that. Under this second theory, we always go beyond a text's literal meanings to supply important implications that were not explicitly stated by the words of the text. We then store the fuller model of the text's meaning in long-term memory. This constructive hypothesis predicts that subjects who had been read the sentence which said that turtles rested *beside* a log wouldn't mistakenly think the original had said that the fish swam under the log. By contrast, subjects who originally heard the sentence which said that turtles rested *on* a log would think the original sentence explicitly said that a fish swam under the log, even though the sentence didn't mention a relation between the fish and the log. Subjects consistently made this mistake about the explicit words of the sentence. The constructive hypothesis, in other words, proved to be correct.

The result was not altogether surprising in light of Sachs's demonstration that we have bad memory for words but good memory for meaning. In the *on*-the-log pair of sentences, the meaning of the probe sentence that substitutes *it* for *them* is the same as that of the original, if we assume that verbal meaning includes inferences based on our prior knowledge of turtles, water, fish, logs, and so on. If all these inferences belong to the basic meaning of the sentence, the meanings of the two different sentences, the original and the probe, would indeed be the same. But this conception of meaning is exactly what was at issue. Is our background knowledge itself part of the meaning of a text? Are all these extralinguistic inferences part of the meaning we initially understand? The answer given by this and other experiments is "Yes, that is the nature of verbal meaning: inferences based on prior knowledge are part of meaning from the very beginning." This constructive hypothesis of verbal meaning has been closely studied and reconfirmed many times in the past several years. According to R. J. Spiro, "The study by Bransford, Barclay and Franks served as a prototype for what was to become a near avalanche of demonstrations that inferential elaborations are a part of the process of understanding prose."[11]

From this finding we can predict that there will be mutual influences between the mental models we have in our minds on the basis of prior knowledge and the words of the text as we read them. To make sense of what we read, we must use relevant prior knowledge

to form a model of how sentence meanings hang together. T model constructed from our prior knowledge and the words of text in turn helps us make sense of further words and sentences the text. The idea was quickly tested by Bransford and his associa in 1972.[12] They used a passage written in language so general a vague that, in the absence of a context, it was difficult to constru a mental model from it. But if the passage was given a title th invoked relevant prior knowledge, subjects constructed a men model that enabled them to understand and remember the se tences. The passage began:

> The procedure is actually quite simple. First you arrange the items in different groups. Of course one pile may be sufficient depending on how much there is to do. If you have to go somewhere else due to lack of facilities that is the next step; otherwise you are pretty well set.

Some subjects were given the title "Washing Clothes" before they read the passage, some were given it after, and some were not given it at all. Only the group who had been given the title before they started to read could recognize sentences from the passage. The title enabled them to integrate the sentences into a mental model that they constructed from prior knowledge about washing clothes. This model gave the sentences meaning, and the sentences in turn adjusted the model of the passage that was finally stored in memory. Apparently such an integrated model is essential to understanding and remembering discourse. Bransford and his associates conducted further experiments, using pictures instead of titles to provide the appropriate model, and the effect on reading was the same.[13]

It's easy to imagine variations on these Bransford experiments that would illustrate the connection between background knowledge and literacy. Persons who lack cultural knowledge are in just the position of the subjects who were given the clothes-washing passage without benefit of a title to form a context for the sentences. Informationally deprived people constantly run across passages that look like the Bransford one, because the texts contain important referential clues they can't understand. Although they can read the individual sentences, they can't make sense out of the whole.

Some less formal examples will corroborate the point. In 1978 community college students were given a prose selection to read, part of which follows.

> GRANT AND LEE
>
> When Ulysses S. Grant and Robert E. Lee met in the parlor of a modest house at Appomattox Courthouse, Virginia, on April 9, 1865, to work out the terms for the surrender of Lee's army of Northern Virginia, a great chapter in American life came to a close, and a great new chapter began.
>
> These men were bringing the Civil War to its virtual finish. To be sure, other armies had yet to surrender, and for a few days the fugitive Confederate government would struggle desperately and vainly, trying to find some way to go on living now that its chief support was gone. But in effect it was all over when Grant and Lee signed the papers. And the little room where they wrote out the terms was the scene of one of the poignant, dramatic contrasts in American history.
>
> They were two strong men, these oddly different generals, and they represented strengths of two conflicting currents that, through them, had come into final collision.
>
> Lee was Tidewater Virginia; he embodied a way of life that had come down through the age of knighthood and the English country squire. . . .
>
> Grant, the son of a tanner on the western frontier, was everything Lee was not. He had come up the hard way and embodied nothing in particular except the eternal toughness and sinewy fiber of the men who grew up beyond the mountains.[14]

The community college test population found this passage difficult to understand because they were, surprisingly, ignorant of the identities of Grant and Lee. They were not stumped by the vocabulary of the text in the ordinary sense of that term. For instance, they easily understood *burdensome, reliability, availability*, and *reassurance* in the following passage, which they read with perfect facility and accurate recall.

A GOOD FRIEND IS HARD TO FIND

Troubles always seem less burdensome and joys more pleasant whe a good friend is present to share them. What makes a person good friend? The qualities are numerous and hard to define; bu two traits, availability and reliability, are essential. When I say tha good friends are "available," I mean they are always nearby, read to lend a hand. When I say that they are "reliable," I mean tha they stand by their friends no matter what happens.

Good friends are always interested in their companions' day-to day activities, always able to give reassurance. We never know wha surprises the day may bring, and this uncertainty makes us nervous. When the telephone rings, our anxiety increases. Then we hear our friend's familiar greeting; and as the warm and cheery voice dispels our tension, our fears subside.

A person often finds himself in family predicaments that call for an outsider's help.[15]

These two passages were part of some large-scale experiments that my coworkers and I conducted in 1978.[16] In the course of them we developed a technique for testing large numbers of subjects (more than 700) for variations in reading skill that can be attributed to variations in relevant background knowledge. Dividing our audiences in half, we used two versions of the same text, one well written, the other stylistically impaired according to definite principles. As long as both versions of a text concerned a subject familiar to the readers, their performances showed a definite sensitivity to the stylistic superiority of one version over the other. When the topic was familiar, the group reading the better version understood it more efficiently; when it was unfamiliar, the performances of the two groups were nearly the same.

The following is a graph of the result we obtained with a pair of texts written on a familiar topic. The dotted line represents the performance of the group that was given the superior text, the solid line the performance of the group given the stylistically degraded version.[17] (These are called quintile graphs because the audience is divided into five groups according to reading rate, with the slowest fifth of readers labeled 1 along the base of the graph, the next fastest fifth labeled 2, and so on. The vertical line marks off reading rates

in words per minute. The points through which the two lines are drawn are the average rates for each quintile taken as a group. Comprehension tests determined that all subjects had understood the passages, making reading rate an informative variable.)

QUINTILE GRAPH: "*TIME* VS. *NEWSWEEK*"

Reading Rates of Group A Compared With Rates of Group B When Reading Student Paper (A) and Stylistically Improved Version (B).

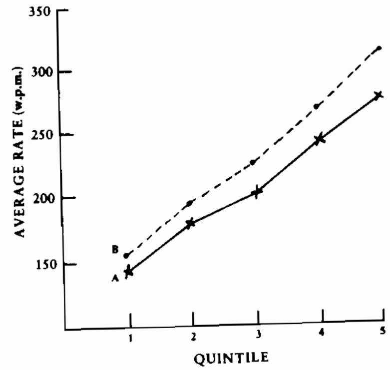

We were able to determine beforehand that our two groups were well matched, because we got the result shown in the next graph when each group was given the same text.

QUINTILE GRAPH: CONTROL ESSAY

Reading Rates of Group A Compared With Rates of Group B When Reading the Same Control Essay.

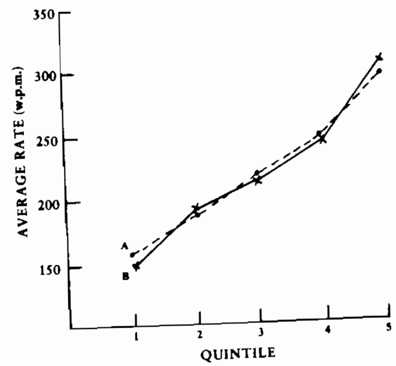

Then, some distance into our experiments, we noticed an int esting effect when we started using texts that addressed unfamil topics. Lack of familiarity not only debased the reading rates our audiences, it also erased the differential effects of good and b style. Below are two results from texts taken from *The History Civilization* by Will and Ariel Durant. The style of the Duran prose is of course clear and nontechnical, and the procedure

QUINTILE GRAPH, "EPICUREAN ROME"

Reading Rates of Group A Compared With Rates of Group B When Reading Original (A) and Stylistically Degraded Version (B).

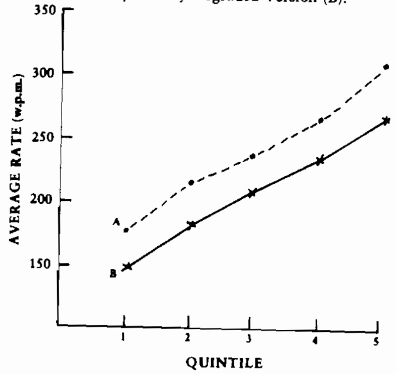

QUINTILE GRAPH: DURANT, "HEGEL"

Reading Rates of Group A Compared With Rates of Group B When Reading Original (A) and Stylistically Degraded Version (B).

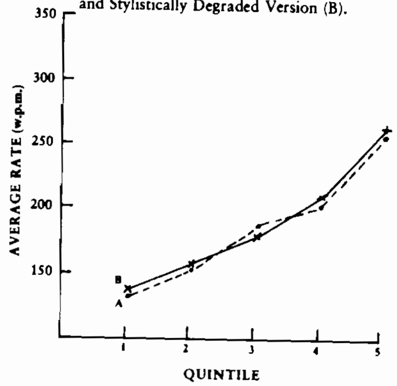

used for impairing their prose style was exactly the same in all cases. The text for the Epicurean Rome graph focused on the public baths of ancient Rome; the other graph was for an explanation of Hegel's conception of logic as metaphysics, a topic unfamiliar to our audience.

This effect was duplicated every time we tested it. We had unwittingly discovered a way to measure the variations in reading skill attributable to variations in the relevant background knowledge of audiences. Later, when we compared the results of our university presentations with the results we obtained at a community college, the contrasts were striking. The comparative results for the essay on Grant and Lee, one version of which had been stylistically impaired, are shown below.

QUINTILE GRAPH:
CATTON, "GRANT AND LEE"
UNIVERSITY AUDIENCE

Reading Rates of Group A Compared With Rates of Group B When Reading Original (A) and Stylistically Degraded Version (B).

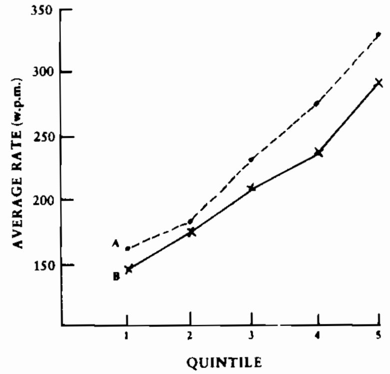

QUINTILE GRAPH:
CATTON, "GRANT AND LEE"
COMMUNITY COLLEGE AUDIENCE

Reading Rates of Group A Compared With Rates of Group B When Reading Original (A) and Stylistically Degraded Version (B).

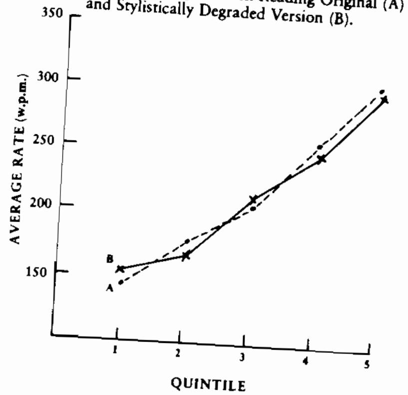

But such comparisons did not always reflect badly on the readin skills of the community college students. Here is a comparison o the two groups' comprehension of the essay on friendship.

QUINTILE GRAPH: "FRIENDSHIP"
UNIVERSITY AUDIENCE

Reading Rates of Group A Compared With Rates of Group B When Reading Student Paper (A) and Stylistically Improved Version (B).

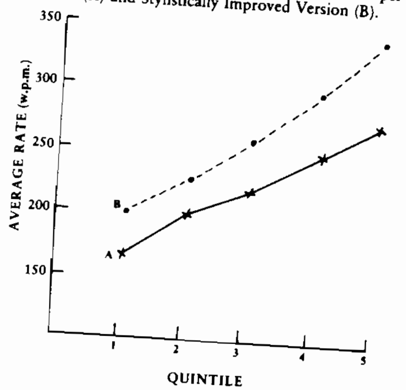

QUINTILE GRAPH: "FRIENDSHIP"
COMMUNITY COLLEGE AUDIENCE

Reading Rates of Group A Compated With Rates of Group B When Reading Student Paper (A) and Stylistically Improved Version (B).

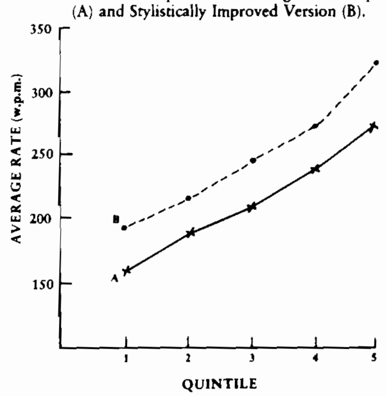

To those who know how to read them, these graphs tell a sad yet promising story — sad because the results with "Grant and Lee" show that the community college students lacked the background information needed for general literacy, but promising because the results with "Friendship" show that background information was really all they were lacking. Their memory capacities, eye movements, basic vocabularies, and reading strategies — in short, their reading skills — were all up to par. What they were missing was background information about Grant, Lee, and the Civil War (and, it can be inferred, much else that writers take for granted in their readers).

In experiments with children, researchers have drawn significant conclusions about the importance of background knowledge for general reading ability. T. Trabasso and his colleagues discovered that differences in reading ability between five-year-olds and eight-year-olds are caused primarily by the older children's possessing more knowledge, not by differences in their memory capacities, reasoning abilities, or control of eye movements.[18]

Similar work by P. D. Pearson and his associates has shown that, among seven-year-olds who score the same on reading and IQ tests, those who have greater knowledge relevant to the text at hand show superior reading skills. In one experiment, equally talented second-graders were tested to find out how much they knew about

spiders.[19] Those who knew a little about them were much be readers of a story involving spiders than were the other child even though the story contained no special concepts or diffi words. One might be tempted to say that these results are pred able and obvious. Nevertheless, educators have not acknowled or acted on the implications of such results.

THE PSYCHOLOGICAL STRUCTURE OF BACKGROUND KNOWLEDGE

As I write these words, I am glancing out my window. My view blocked by some objects a few feet away that are flat, rough oblong, and green, attached to irregular brown tubes of differe sizes. These are leaves and branches, and, together with the roo beneath them, I know them as trees. All continuing experience partial and fragmented like my view of those trees. Our cogniti life takes place through a small window of attention that is frame by short-term memory. We use past knowledge to interpret th window of experience, to place its momentary fragments withi larger wholes that give them a function and a place. The raw dat of visual or linguistic perception come in bits and pieces, and the will remain meaningless if they continue to be nothing but colore flecks or mere sequences of letters or words. In our minds these transitory fragments can acquire meaning only by being placed within larger, not presently visible wholes that are based on past knowledge.

When I ask people what they see out of my window, they invariably answer trees. They don't respond that they are seeing the out-of-doors or nature or woods. Those categories are apparently too broad. At the other extreme, they don't say that they see some ends of branches with oak leaves, although that is all they see directly. Without being able to view a single tree trunk or top or bottom, they nonetheless say that they see trees. In structural terms, this response parallels the perception by the subjects who read that a fish swam under a log, even though the sentence didn't explicitly say that. The unseen trunks and roots are parts of what we "see," just as the fish that swam under the log is part of what we understand from a sentence saying the fish swam under turtles resting on a log. We "see" trees when we see only ends of branches, because we know from past experience that the branches are part of unseen structures beyond the edge of perception.

But why should people interpret what they see from my windows as trees rather than branches or the out-of-doors, since those categories are also valid interpretations based on past knowledge? The answer is that, if no other constraints are at work, people tend to interpret their experiences through middle categories, that is, through classifications that are neither specific, like oak leaves, nor general, like flora. We normally interpret experience through the categories that have been most useful to us in the past. We try them out initially as the most efficient instruments for perceiving the world and acting in it. Research has shown that middle-level categories are the ones children learn first in acquiring language; they learn *tree* before *oak*, and they learn *dog* before *animal*.[20]

Eleanor Rosch has called these middle categories basic-level terms. We owe much of our recent insight into their primacy to work done by her and her associates. In a number of papers she has argued that middle-level categories are the basic furniture of our conceptual world.[21] She has demonstrated that, in the absence of any special context or constraint, people understand the world in terms of basic classifications. Furthermore, people understand these normal classifications through commonly encountered examples. For instance, a typical association of the category tree for North Americans is leaves, rather than fronds.

In one experiment Rosch asked her subjects to arrange some specific names under the categories *fruit*, *furniture*, *bird*, and *clothing*, ranking the names from the most to the least typical.[22] The results were quite consistent from person to person. For instance, under *fruit*, the subjects agreed on the following order of typicality:

apple	fig
plum	prune
cherry	coconut
watermelon	olive

Under *furniture*, Rosch found the following ranking:

chair	lamp
dresser	cupboard
davenport	radio
footstool	ashtray

And under *bird* she got:

robin	pheasant
swallow	goose
eagle	chicken
crow	penguin

If people in our culture tend to agree that a robin is more typic than a pheasant in the bird category, it is likely that, other thin being equal, they will think of a robin-sort-of-creature when th hear the word *bird*. With this in mind, Rosch asked subjects t compose isolated sentences using *bird*, and they produced sentenc like

I heard a bird twittering outside my window.
Three birds sat on the branch of a tree.
A bird flew down and began eating.

Then Rosch replaced the class name *bird* with various specific members of the class (e.g., *robin, ostrich, chicken*), asking other subjects to judge whether the sentences of the type shown above made good sense. She found that if you replace *bird* with *robin* in each of the above sentences, subjects always say that the sentences make sense, but if you replace *bird* with *ostrich* or *chicken*, subjects say that the sentences don't make sense. Chickens don't act the way the birds in the sentences do — they don't twitter, sit on branches, or fly. In isolated sentences, when an untypical member of a class is substituted for the class name, sentences that previously sounded reasonable often fall into absurdity.

These two findings fit together neatly — people agree about typical examples of a class and the isolated sentences they make with class names work best with typical examples. The more typical an example is for its class, the more it accords with an isolated sentence

that uses the class term. This finding suggests that we hold in our minds typical exemplars or prototypes of the category words we use and that these prototypes constitute the usual furniture of our minds. They are the classifications of things we use every day and keep at the ready in perception, conversation, reading, and writing. Rosch and others have shown that such prototypes are processed much faster than atypical examples in verifying the truth or falsity of isolated sentences. For instance, it takes a subject much longer to say true or false to the sentence "A chicken is a bird" than to "A robin is a bird."[23]

Reseachers have concluded that our usual prototypes for terms like *bird, tree,* and so on, contain definite features that we match up when we verify sentences like "A chicken is a bird." Since a chicken isn't like the usual prototype, we take rather long to compare features of birds and chickens before reaching a decision. But we can make an instant decision about robins. The consistent difference in the time it takes to reply true or false to typical and untypical examples suggests that part of our mental representation of a category is a prototype, e.g., for a bird the prototype is a vaguely robinlike creature.

Knowing about prototypes is essential for understanding how we apply past knowledge to the comprehension of speech. These mental models are the same as Hilary Putnam's "stereotypes."[24] We are able to make our present experiences take on meaning by assimilating them to prototypes formed from our past experiences.[25] Psychologists have been investigating these prototypes intensively and have given them various names such as frames, theories, concepts, models, and scripts. Researchers who have been relating these mental entities to reading, particularly R. C. Anderson and his associates, have chosen the word *schema* for them, and it is the term I use to refer to the phenomena. *Schema* and its plural *schemata* correctly suggest somewhat abstract mental entities rather than concrete images.[26]

If Rosch had tested her subjects in Australia, her results would probably have been different, and she would have found the typical bird schema to be unlike a robin in many respects. The creature nearest the Australian schema would be large and colorful, like a parakeet or a galah, with bright scarlet, green, and blue plumage,

for those are the brilliant colors of the everyday birds one see one's garden in Australia. It would be odd indeed if the basis an Australian bird schema were as drab a specimen as a robin. W this kind of qualification in mind, psychologists have ampli Rosch's work in one important respect. They have found that schemata under which we create and interpret isolated senten are not necessarily the initial schemata we apply in other contex Demonstrations of this variability have been conducted with t words *cup, eat, red,* and *held.*[27] It is clear that Rosch's work mu be further generalized so that it reflects the context sensitivity our initial schemata.

The research with the word list just referred to has shown th our minds have a remarkable ability to change initial schema according to the situations in which we find ourselves. Such ad justments are largely unconscious, but one can perceive the patter of the quick adjustments we make even with isolated sentences, a Anderson has suggested in the following passage:

1. The punter kicked the ball.
2. The baby kicked the ball.
3. The golfer kicked the ball.

A different sort of ball is, loosely speaking, implied by each sentence. The punter is kicking a football, and the golfer a golfball. Although a baby could be kicking either of these kinds of ball, this is not the inference that will be drawn by most readers. Instead, a ball a baby is likely to kick will be hypothecated — perhaps a brightly colored, inflated, plastic ball. *Kick* has different senses in the three sentences. Compare the smooth powerful kick of a punter with the hesitant, uncoordinated, possibly accidental kick of a baby. Golfers do not ordinarily kick their balls; this fact leads to the supposition that this one was angry or maybe cheating.[28]

Anderson and others have thus made an important adjustment to the principle that we interpret the world through a basic set of schemata. We are in fact much more flexible. In reading, for example, we adjust our initial schemata to the specific text. In a story

about a Thanksgiving feast, our initial schema for *bird* is *turkey, not robin.*[29] Two-way traffic takes place between our schemata and the words we read. We apply past schemata to make sense of the incoming words, but these words and other contextual clues affect our initial choices of schemata and our continuing adjustment of them.

In light of what has been described so far, we can say in general how a reader actively brings past schemata to bear upon what he or she is reading. According to the picture sketched by recent research, the reader is confined to a rather narrow window of attention that is limited by short-term memory. Through this window, the reader constantly connects a few words into clauses that have meaning and the clauses to appropriate schemata based on past experience. Thus, the reader is not just passively receiving meaning but is actively selecting the most appropriate schemata for making sense of the incoming words. Then the reader actively adjusts those schemata to the incoming words until a good fit is achieved. This process can work efficiently only if the reader has quick access to appropriate schemata. When the appropriate schemata are not quickly available, and the reader is forced to do a lot of pondering to construct them at the time of reading, the limits of short-term memory are quickly reached, and the process has to be painfully restarted and restarted.[30]

The immediately needed schemata go far beyond birds and cups to include the wider knowledge of cultural literacy. Consider, as a brief example of what is required, the opening words of the Grant and Lee passage that caused community college students so much trouble.

> When Ulysses S. Grant and Robert E. Lee met in the parlor of a modest house at Appomattox Courthouse, Virginia, on April 9, 1865, to work out the terms for the surrender of Lee's army of Northern Virginia, a great chapter in American life came to a close, and a great new chapter began. These men were bringing the Civil War to its virtual finish.

These words are like the leaves outside my window. If I do not know how they connect with larger schemata, I cannot make sense

of them. On the other hand, the schemata I need are rather abs entities that do not require profound knowledge. Consider little detail there is in the schematic information needed to read passage.

1. America fought a Civil War.
2. The two sides were the Union and the Confederacy.
3. Grant was the chief general for the Union.
4. Lee was the chief general for the Confederacy.
5. The Union won.

These ideas do not take the form of such a list in our minds b probably exist as a schematic model.[31] Clearly, these five propo tions don't exhaust the schematic information one needs to re the passage, but, on the other hand, the necessary further schema are much more elementary and were certainly known to the com munity college students. More specialized information is provid explicitly by the text.

We know that, in reading, our short-term memories allow us t attend only to a few salient features of word meanings — just thos that are needed to make the meanings fit together. In the quote passage, a reader could be attending to only a few surface idea like Lee surrenders to Grant. Nonetheless, those meanings are con nected to a lot of subsurface information, just as leaves are con nected to information about trunks and roots of trees. One must integrate the passage on Grant and Lee with schematic background information about what a general is, what Civil War entails, what surrender means, and so on, although one does not pay direct attention to the implications.

The salient surface meanings that we pay attention to in reading thus *stand for* a whole world of relationships that we are not paying conscious attention to. When we really comprehend what we are reading, we are able to supply those implied background relationships. A schema functions as a unified system of background relationships whose visible parts stand for the rest of the schema. Because our narrow windows of attention confine us to just a few elements at a time, the technique of using surface elements to stand for larger wholes is an essential feature of our mental life.[32] Schemata are our necessary instruments for making the surfaces of what we read connect significantly with the background knowledge that is withheld from immediate consciousness by the limits of short-term memory.

To pursue the practical importance of this fact, let us examine more closely how these schemata work. The evidence that schemata function roughly as described so far goes back many years. As long ago as 1922, the German psychologist Franz Wulf observed that we have a tendency to "schematize" the details of what we remember about pictures and events so that, with a lapse of time, the details we remember become more and more like the average or typical things of our experience.[33] F. C. Bartlett, in his famous book *Remembering*, showed that our memories are freshly constructed from our habitual schemata, not just reproduced from actual memory traces. Our memories are always introducing elements from our normal schemata that weren't in an original event, and, by the same token, they are always suppressing some elements of the original event that don't exist in our normal schemata.[34] Thus, according to Bartlett, memory is active and constructive rather than passive and reconstructive.

> An individual does not normally take a situation detail by detail and meticulously build up the whole. In all ordinary instances he has an overmastering tendency simply to get a general impression of the whole; and on the basis of this, he constructs the probable detail. Very little of his construction is literally observed. But it is the sort of construction which serves to justify his general impression.[35]

Bartlett backed up his observation by testing subjects' memories of an American Indian folktale. His subjects' recollections of the story were distorted in the direction of their habitual schemata. All this is evidence that schemata are essential instruments of our cognitive lives.

Further strong evidence for the importance of schemata in our understanding of language was provided by B. H. Ross and G. H. Bower in 1981.[36] They set up their experiments on this reasoning: if we understand language through schemata, then any important

element in a schema ought to call to mind other important matic elements. If we are presented the term "U. S. Grant," it sh call to mind the term "Civil War." If schemata do indeed go our associations, then our tendency to associate words that schematically connected ought to be greater than would be expe merely from the frequency of their co-occurrence in our past perience, which is the prediction of traditional learning theory associationism. Ross and Bower gave subjects eighty sets of f words to study. Some had schematic associations, as in the case *driver, trap, rough,* and *handicap,* all of which belong to the sche golf. When one word from this set (or from any other schema set) was later given as a cue, subjects showed a strong tenden to recall the other three words belonging to the schema. Th performance was much better than had been predicted stimulus-response theory or associationism. This is strong eviden that schemata function powerfully in both comprehension a memory.

Although we know in a general way what schemata are and ho they function, psychologists do not understand many details. Th latest view (that of Johnson-Laird, for instance) is that we use a least two radically different types of schemata, one analogous t static pictures and another to scripts or procedures.[37] In addition we have not only general schemata of categories like birds but als memory traces of occasions when we perceived or read about birds.[38] We know that schemata overlap and get embedded in other sche mata, that we treat them as provisional theories open to revision, and so forth. Indications are that an accurate modeling of our use of schemata would be very complex indeed. But we do not have to know the fine detail of the process to make intelligent practical advances in the teaching of literacy. We need to know how to put the principles of schemata to work to improve the performance of semiliterate people. What do we know about schemata that will enable us to do that?

We know that schemata perform two essential functions that are relevant to literacy. The first is storing knowledge in retrievable form; the second is organizing knowledge in more and more efficient ways, so that it can be applied rapidly and efficiently. Without appropriate background knowledge, people cannot adequately un-

derstand written or spoken language. And unless that knowledge is organized for rapid and efficient deployment, people cannot perform reading tasks of any complexity. Although there are sizable variations in reading rates among good readers, no good reader is a very slow reader. Slowness of reading beyond a certain point makes assimilation of complex meaning impossible. The limits of short-term memory do not allow the integration of "unchunked" material, and so crucial parts of meaning are lost to memory while other parts are being painstakingly worked out.

The evidence has shown that schemata are useful as a kind of mental shorthand that helps us work within the limits of short-term memory. One element of a schema represents many other subterranean elements, allowing us to deploy that one surface element as though the whole complex schema were a single item. In other words, a well-developed schema not only helps us make sense of incoming data but also helps us to manipulate and connect information rapidly. In this process, speed of comprehension is equivalent to quality of comprehension, because without the speed and the shorthand provided by well-organized schemata our circuits get overloaded. We can't perform the necessary manipulations before we run out of memory, and we continually have to start over again.

Since use of schemata is of course common to literate and illiterate alike, we need to ask two specific questions. What sorts of schemata must people acquire to be really literate? And how can these particular schemata be made efficient enough to create good readers? A. M. Collins and M. R. Quillian published an important piece of work on this subject in 1969.[39] Their theory is that the knowledge most often needed is also the most directly available and is, so to speak, right at the surface of the schema. Other parts of the schema are "deeper" and take longer to retrieve. One of their examples also concerns a bird. They reasoned that general information about birds would be farthest from the surface and that specific information (e.g., direct information about robins or canaries) would be most often used and closest to the surface.

They posited the following schematic network for canaries, starting with the most available information provided by a mention of the word *canary* and going down to the least available and most general level.

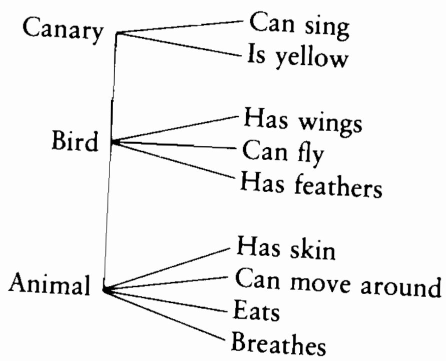

Starting from the top of the vertical line, there is an increasi long path from bird traits to animal traits. The prediction is t if knowledge is organized in this kind of pattern, the decision whe sentences are true or false will take longer the farther down vertical path the sentences require us to go. Collins and Quill gave subjects the following sentences:

Canaries are yellow.
Canaries lay eggs.
Canaries can breathe.

Subjects did in fact take longer to decide the truth of sentences th farther down the conjectured schematic path they had to go. This kind of result has been confirmed many times by researchers.[40]

What does this experimentation tell us about the rapid deploy ment of background information in reading? Collins and Quillian's observations suggest that the top portion of a schema is the im portant part to know. The schema canary can yield an indefinite number of facts and association by remote inference from knowl edge of the world: canaries have backbones; canaries have their own special pattern of DNA; canaries are descended from reptiles; canaries must drink water; canaries mate; canaries die. One could go on this way for a long time, with specifications not only from biological but from physical knowledge: canaries obey the law of gravity and the laws of motion, and so on. But this secondary information about canaries is not important in communicating with

ow human beings. What is functional in reading, writing, and nversing is the distinctive system of traits in the schemata we e — the traits that differentiate canaries from other birds: their allness, yellowness, ability to sing, use in human culture, being pt in cages, and so on. We need to know the primary traits mmonly associated with canary in our culture in order to deploy he associations rapidly when we encounter the word *canary* in reading.

There is no substitute for simply knowing these primary associations. They must be called up with lightning speed in the course of reading and conversing. We do not have the luxury of figuring out such associations one at a time. We may do this with one or two words in a paragraph — that is the way we learn new words — but we cannot pause over many words at a time. When we encounter U. S. Grant, the primary associations must be available to us in milliseconds: that he was an important Union general, that he became president, that he drank. These are some of the implicit associations needed to make meaningful what is explicitly written about Grant.

Research into the importance of primary associations thus introduces a subject of profound significance for teaching reading and writing. Successful communication depends upon shared associations. To participate in the literate national culture is to have acquired a sense of the information that is shared in that culture. No adult-level discourse retreats to the rudiments of knowledge. If assumptions about rudiments could not be made, ordinary discourse would be so lengthy and intricate as to obscure its own point.

Educators in the Rousseau-Dewey tradition, who favor less emphasis on mere fact and more emphasis on the intensive study of a few cases, encourage us to believe that students will thereby understand general principles and learn how to think critically.[41] But literacy requires us to have both intensive knowledge of relationships and extensive knowledge of specifics. We need not only a general understanding of the principles of biology (which would enable us to infer that canaries breathe and lay eggs) but also specific knowledge of facts about canaries. We need to know not only the broad social and historical significance of the American Civil War but also who U. S. Grant was, and what the word *Appomattox*

signifies. It is not enough to say that students can look these facts up. The research reviewed above shows that in order for readers to integrate phrases into comprehensible meanings, they must already possess specific, quickly available schemata. When readers constantly lack crucial information, dictionaries and encyclopedias become quite impractical tools. A consistent lack of necessary information can make the reading process so laborious and uncommunicative that it fails to convey meaning.

SKILL AS KNOWLEDGE AND KNOWLEDGE AS SKILL

Researchers in cognitive psychology and the area of computer science known as artificial intelligence (AI) have come to strikingly similar conclusions about the knowledge-bound character of all cognitive skills. AI research demonstrates that the ability of humans to exercise a skill depends on their possession of specific schemata that are sufficiently numerous and detailed to handle the many varieties of the tasks they are called on to perform. It is more accurate to speak of "reading *skills*" than of "reading skill." The graphs showing the community college students' high degree of skill in reading the essay on friendship but their lack of skill in reading about Grant and Lee accord with the recent discovery of AI and cognitive psychology that a skill is not a unified system of intellectual muscles that can be developed by calisthenics into a vigorous all-purpose ability.

Dr. Herbert A. Simon, a leading figure in AI research, once wryly remarked that saying an expert performance is caused by a "skill" is like Molière's doctor saying that the sleep-inducing properties of opium are caused by its "dormative power."[42] Simon, a Nobel laureate, has been working since the 1950s on the detailed structure of cognitive skills. The discoveries that he and his coworkers have made should induce a deep skepticism toward the belief that our schools can teach reading, writing, and critical thinking as all-purpose general skills applicable to novel problems. Simon and his colleagues have cast doubt on the idea that there are *any* general or transferable cognitive skills. All cognitive skills depend on procedural and substantive schemata that are highly specific to the task at hand.

Once the relevant knowledge has been acquired, the skill follows.[43] General programs contrived to teach general skills are ineffective. AI research shows that experts perform better than novices not because they have more powerful and better oiled intellectual machinery but because they have more relevant and quickly available information. What distinguishes good readers from poor ones is simply the possession of a lot of diverse, task-specific information.

Probably the most dramatic illustrations of the knowledge-bound character of human skills came from some remarkable experiments conducted by Adriaan de Groot, a Dutch psychologist, who described his findings in a book entitled *Het Denken van den Schaker* (literally, "the thinking of chess players").[44] De Groot discovered that chess masters are astonishingly skilled at remembering and reproducing chess positions after a very brief exposure to them. The subjects in his experiments were players of various abilities, as indicated by their official chess rankings. In one experiment, de Groot displayed for five to ten seconds a chess position from an actual game in which twenty-five pieces were left on the board. The subjects were asked to reproduce the position from memory. Grand masters performed this feat with 100 percent accuracy, masters with 90 percent accuracy. Weaker players were lucky if they could correctly place five or six pieces.

Then de Groot varied the conditions of his experiment in one respect. Instead of placing the twenty-five pieces in positions from an actual game, he placed them on the board randomly. The results were unexpected. All his subjects — grand masters, masters, class A players, and class B players — performed the same as novices did, placing only five or six pieces correctly. This experiment has been duplicated in several different laboratories, and structurally in several other fields, including algebra, physics, and medicine, always with the same striking results.[45,46]

When the configuration of a task is significantly changed, past skills are not transferred to the new problem.[47] In normal circumstances, of course, elements from past problems appear in present ones, and experts perform well with duplicated elements. But beyond similar or analogous circumstances, skill is not transferred.

Since the 1950s, when Simon introduced de Groot's work to the American research community, no convincing counterarguments or experiments have challenged these results. Researchers have consistently found that people do not develop general, transferable skills in problem solving, critical thinking, or in any other field.

What are the reasons for these astonishing results? The fullest explanations have come from those engaged in making computer models of skilled and unskilled performances. AI scientists have formulated programs that conform extremely well to actual human performances and to the models currently inferred by cognitive psychologists. The convergence of the accounts from both AI and empirical psychology increases one's confidence in their correctness.[49]

There is another reason why the AI research results should be taken seriously. AI models actually work. You can play chess against a computer program and lose. An up-to-date chess program behaves like a skilled player, not because the computer is directed to zip through every possible variation — it could not possibly do that in the few seconds it takes to respond — but because in some ways, though not in all, it imitates the knowledge-based procedures of experts.

I said that models from research in artificial intelligence and from experimental psychology are converging. Both approaches show that expert performance depends on the quick deployment of schemata. In the case of de Groot's chess experiments, the currently accepted models would explain the results as follows. When de Groot's subjects were exposed to a chessboard for about six seconds, the limits of short-term memory were quickly reached. The constraint of short-term memory explains, of course, why novices could manage to place only about six pieces — the range predicted by research. There's no difficulty, then, in explaining the poor performance of the novices. But what explains the superb performance of the experts?

In his second experiment de Groot proved that expert chess players have no better memories or general skills than novices. Like everyone else, they can remember only five or six pieces from a random pattern. But when they are asked to survey a position that is not random, they quickly match the observed patterns with spe-

cific schemata already stored in memory. That process instantly facilitates the cognitive task. Consider the arithmetic of the situation. If expert schemata for chess patterns average five or six pieces per schema, that would remove the burden of remembering twenty-five individual items. Think of the experts as perceiving the positions of five or six pieces all at once, as words rather than individual letters. Because they already possess positional schemata, experts would need to remember only four or five chunked items — well within the range of short-term memory.

It is probable that the grand masters in de Groot's experiment quickly focused on one or two variations from known patterns, and that those one or two variations were all they needed to remember. They were in the masterful position we would be in if we were asked to spell out the sixteen letters of "The cat is on the mat" after an exposure of only six seconds. Our performance would be made possible by our possession of spelling schemata for the words. When de Groot's chess masters were asked to reproduce the original position, they needed only to recall two or three variations and were able to reconstruct the rest from the schemata they already possessed.

Skill in reading is like skill in chess in many respects. Good reading, like good chess, requires the rapid deployment of schemata that have already been acquired and do not have to be worked out on the spot. Good readers, like good chess players, quickly recognize typical patterns, and, since they can ignore many small-scale features of the text, they have space in short-term memory to take in an overall structure of meaning. They are able to do all of this because, like expert chess players, they have ready access to a large number of relevant schemata. By contrast, unskilled readers lack this large store of relevant schemata and must therefore work out many small-scale meaning relationships while they are reading. These demanding tasks quickly overload their short-term memories, making their performance slow, arduous, and ineffective.

How large is the "large number of schemata" that skilled persons have acquired? It has been estimated that a chess master can recognize about 50,000 positional patterns.[50] Interestingly, that is the approximate number of words and idioms in the vocabulary of a literate person. Is there any significance in this coincidence of num-

bers? Probably, according to Simon and Barenfeld.[51] A schema for a chess pattern is the functional equivalent of a schema for a word or idiom. There must be an upper limit to the number of schemata that we can effectively utilize, because we must have immediate access to them if they are to be effective. Apparently, a number in the range of 50,000 marks the upper limit of items to which we can have rapid access in any domain of activity. Using a significantly greater number of items would cause the process of search and retrieval to occupy too much time.

There are, of course, many more than 50,000 items stored in the full text of long-term memory. A basic vocabulary of 50,000 schemata serves merely as a quickly accessible index to a much larger volume of knowledge.[52] Any of the 50,000 schemata can be related to others, and the further relationships can be stored in long-term memory. A useful illustration of the way one schema can serve as a mental shorthand for enormous complexes of associated schemata is the field of mathematics, where many decades of mathematical labor can be reduced to a single symbol that can be manipulated as though it were a mere digit like the number 5. The term "Civil War" represents not just its own schema, but also whole ranges of further, associated schemata that can be applied when needed but do not obtrude into the window of the mind unless they are called forth. As we have seen, the most effective system for the mind, whether in chess or in reading, is to keep most of its indexed schemata at a medium level of generality.[53]

But although parallels between skill in reading and skill in chess are informative, the two kinds of expertise are structurally different in one respect. Chess players do not have to communicate their stored patterns to anyone else and can organize their schemata in whatever way they like. But when we use verbal schemata, we have less freedom. The patterns of association in the verbal schemata of one person must approximate those of another, otherwise we could not use surface meanings to represent larger systems of subsurface associations. Words and idioms therefore represent systems of association that belong not just to the individual mind but to the language community as a whole. Words, idioms, and grammatical systems represent shared systems of association[54] — cultural literacy.

The last piece of research that I shall describe brings some of

these strands together and proves the importance of cultural literacy in a striking way. The work was conducted and reported by R. M. Krauss and S. Glucksberg.[55] Imagine an experiment in which some of the special constraints of reading and writing are acted out physically in the laboratory. One constraint of reading is that author and reader cannot converse with each other. Only the author is the source of the words; the reader has no access to an author's gestures or facial expressions. To reproduce this constraint, we erect a physical barrier between two people so that only one of the subjects can be heard and neither can see the other. Then we give these two subjects a communication task to perform. The speaker must explain to the listener the order in which the six following unfamiliar shapes must be listed.

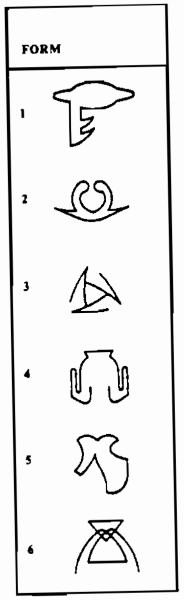

Charts on pages 65–67 from R. M. Krauss and S. Glucksberg, "Social and Nonsocial Speech."

To establish a standard of performance, we run the experime with pairs of literate adults, who find the task almost ridiculous easy. They make virtually no errors on their first attempt. Here a typical set of descriptions by a literate adult playing the role author:

FORM		
1		Looks like a motor from a motorboat. It has a thing hanging down with two teeth.
2		It looks like two worms or snakes looking at each other. The bottom part looks like the rocker from a rocking chair.
3		It's a zigzag with lines going in all different directions.
4		It's like a spaceman's helmet; it's got two things going up the sides.
5		This one looks something like a horse's head.
6		It's an upside-down cup. It's got two triangles, one on top of the other.

Now we try the experiment on six pairs of young children fro four to five years old. Not one pair of children completes a tria correctly. Their messages are short and cryptic rather than full an elaborated like those of adults. Moreover, the messages depen upon associations that are not widely shared in the speech com munity.[56] A summary chart of some results with young children i on the next page.

FORM		CHILD				
		1	2	3	4	5
1		Man's legs	Airplane	Drape-holder	Zebra	Flying saucer
2		Mother's hat	Ring	Keyhold	Lion	Snake
3		Somebody running	Eagle	Throwing sticks	Strip-stripe	Wire
4		Daddy's shirt	Milk jug	Shoe hold	Coffeepot	Dog
5		Another Daddy's shirt	Bird	Dress hold	Dress	Knife
6		Mother's dress	Ideal	Digger hold	Caterpillar	Ghost

Later the experimenters tested older children from kindergarten through fifth grade. Krauss and Glucksberg give the following account of what they found:

> The results were somewhat surprising. Kindergartners performed no better than nursery school children, and displayed the same lack of improvement with practice. Considering that adults made virtually no errors on the very first trial, the performance of children in the first, third, and fifth grades is even more surprising: they were no better than kindergartners on the first trial. The older children did show marked improvement with practice, but it seemed clear that even the fifth-graders (who were about 10 years old) did not approach the adult level. . . . We proceeded to test children from grades three through nine in the Princeton school system. . . . We were again surprised at the generally low level of performance. . . . Ninth-graders, although they showed dramatic improvement in suc-

cessive trials, still did not attain the virtually perfect accuracy that adults displayed from the very first trial.[57]

Krauss and Glucksberg inferred from this — correctly, I believe — that the older children performed poorly not because of their insufficient cognitive development but because of their lack of information.[58] They had not acquired an easy sense of the degree to which their schemata were shared with strangers in the speech community. Their lack of information strained their cognitive capacities when they tried to carry out the task.

> When the demands of the task are relatively light, children do engage in social, nonegocentric speech, and they communicate rather successfully. As the demands become heavier children may still attempt to employ social-communication strategies, but they do so less effectively than adults. Finally, when the demands of the task become heavy enough, children may not have the opportunity to bring into play the social-communication skills they possess.[59]

The performance of the young children in this experiment is like the performance of poor readers. Both are ineffective because of cognitive overload. A semiliterate person reading or a young child describing strange shapes has to figure out too many things at one time. In effective reading, one must not only call up one's own schematic associations but also monitor whether they are appropriate ones shared by the wider speech community. Literate adults have internalized these shared schemata and have made them second nature. As the Krauss-Glucksberg experiments indicate, literate adults have a sure instinct for what will and will not be shared by others in the wider culture. Young children and other semiliterates do not confidently know what other members of the speech community can be expected to know. Thus, besides their lack of substantive information, young children and semiliterates also lack something equally important: readily accessible information about what is shared by others. In reading and writing, as in communicating shapes across a physical barrier, it is essential to know the other person's unspoken systems of association.

Such cultural knowledge is not entirely a function of age. Some children acquire it rather early, and illiterate adults never acquire it fully. Experiments reported by Basil Bernstein suggest that the performances of children attempting tasks similar to those in the Krauss-Glucksberg experiments vary according to the children's degree of literacy, and their performances therefore have, under present circumstances, a strong correlation with social class.[60] In Bernstein's experiments, the most important variable was whether the children were used to talking with literate strangers and had thereby gained a sense of the wider speech community. The shared schemata necessary for reading and writing are always those of the wider community. Bernstein's experiments reconfirmed that literacy is a function of a national rather than a local culture.