CHAPTER ONE

INTRODUCTION AND PRELIMINARY INFORMATION

Reading is a complex skill that is pretty much taken for granted by those who can do it. In the early-seventies (when cognitive psychologists became very interested in studying reading), one of the authors (then a graduate student) got into an elevator in the engineering department at a famous university in the United States with a copy of Smith's book *Understanding Reading* (1971) under his arm. A bright young freshman engineering student upon seeing the book was quick to remark, "Oh, reading, I learned how to do that 15 years ago." That remark is pretty consistent with most people's attitudes about reading. While those who can do it fluently take it for granted, its complexity is more apparent to those who are having trouble. Reading is sometimes difficult for children to learn (particularly in comparison to the ease with which they learn to speak), and illiterate adults find learning to read agonizingly frustrating.

Anyone reading this book is likely to be familiar with 30,000 or more words, and can generally recognize most of them within a fraction of a second. A skilled reader can do this despite the fact that the letters which make up the words are often printed in different typefonts. In the case of handwritten letters, a reader can still read and comprehend despite rather dramatic differences in style and legibility. In being able to read and identify words in spite of all this variability, a skilled reader is able to perform a feat that is well beyond the capability of the most powerful computer programs available today. But this is not all. Skilled readers can identify words that have different meanings in different contexts. Consider the use of the word *boxer* in the following two sentences:

John knew the boxer was angry when he started barking at him. (1.1)
John knew the boxer was angry when he started yelling at him. (1.2)

These two sentences are identical except for a single word which makes clear the appropriate meaning of the word *boxer*. The less common meaning for *boxer* is a dog. Since dogs bark and people don't, however, *boxer* in Sentence 1.1 clearly refers to a dog. Likewise, in Sentence 1.2 the fact that the *boxer* is yelling leads us to believe that the sentence is referring to a person. If you are observant, you may have noticed that there are actually two ambiguities in sentences 1.1 and 1.2. Not only is the word *boxer* ambiguous but also the pronoun *he*. In 1.1, *he* would be interpreted as the *boxer* because of *barking,* but in 1.2, *he* could either be *John* or *the boxer*. There is some bias, however, in sentences like 1.1 and 1.2 to associate the pronoun with the most recent antecedent. Other factors can change the bias; in 1.3, *he* is most likely *John,* but in 1.4, *he* is most likely *the boxer*.

The boxer hit John because he started yelling at him. (1.3)
The boxer hit John and then he started yelling at him. (1.4)

The point of this discussion is that we can easily understand the meaning of these different sentences despite the fact that individual words have more than one meaning and pronouns occasionally have unclear antecedents. Coupled with this fact is the observation that we can easily understand puns, idioms, and metaphors. For example, in Sentence 1.5

John thought the billboard was a wart on the landscape. (1.5)

none of us would believe the literal meaning of the word *wart* was intended. We quite easily understand the sentence to mean that the billboard was ugly and spoiled the scene. Just as we can easily comprehend the metaphor in sentence 1.5, the idiomatic nature of sentence 1.6

John hit the nail on the head with his answer. (1.6)

presents a difficulty only for nonnative readers of English who attempt a literal interpretation of the sentence and find it nonsensical. Thus, skilled readers are very good at combining the meanings of individual words to derive the meaning of sentences and paragraphs, as well as short passages and books. Readers can draw inferences by relying upon what they already

mplex task? And how is the skill acquired?
this book. For the most part, we will focus
g to explain the process of reading. Our
understand the skill itself before we can
l our primary orientation in this book is a
-processing point of view (i.e., under-
ms underlying reading). In the remainder
lace the rest of the book into perspective.
how researchers have historically viewed
n overview of the human information-
t types of processing mechanisms may be

ADING RESEARCH

, the branch of experimental psychology
s, can be traced to the establishment of
879. Workers in Wundt's laboratory were
ted to memory and to language process-
s considerable interest in the process of
ith the publication of Huey's (1908) *The*
ding. A perusal of the chapters in the first
with the psychology of reading) will reveal
le similarity to the topics covered in the
contemporary books dealing with the
his contemporaries were interested in eye
re of the perceptual span (how much
g a fixation of the eye), word-recognition
comprehension, and reading rate. Huey's
escription of his findings and those of his
l a joy to read. Many of the basic facts we
g reading were discovered by Huey and
e and seemingly archaic techniques in
devices currently available to record eye
ir discoveries have stood the test of time
using more accurate recording systems. A
val, the French oculist, first noted that

know to help understand text, and from reading words they can form images of scenes and appreciate poetry.

We have been arguing that the feats of a skilled reader are truly impressive. Very powerful computers, despite tremendous memory capacity, cannot do what a skilled reader can do; such machines (or more specifically the programs that run them), would fail on many of the tasks we have mentioned that a skilled reader handles almost effortlessly. How do skilled readers accomplish this complex task? And how is the skill acquired? These are the central questions of this book. For the most part, we will focus on the skilled reader in attempting to explain the process of reading. Our primary rationale is that we must understand the skill itself before we can understand how it is acquired, and our primary orientation in this book is a cognitive-psychology–information-processing point of view (i.e., understanding the component mechanisms underlying reading). In the remainder of this chapter, we will attempt to place the rest of the book into perspective. We will do this by first discussing how researchers have historically viewed reading. Then we will present an overview of the human information-processing system, discussing what types of processing mechanisms may be involved in reading.

HISTORICAL OVERVIEW OF READING RESEARCH

The roots of cognitive psychology, the branch of experimental psychology that studies how the mind works, can be traced to the establishment of Wundt's laboratory in Leipzig in 1879. Workers in Wundt's laboratory were keenly interested in questions related to memory and to language processing. Shortly thereafter, there was considerable interest in the process of reading which reached its apex with the publication of Huey's (1908) *The Psychology and Pedagogy of Reading*. A perusal of the chapters in the first part of his book (that part dealing with the psychology of reading) will reveal that the chapters bear a remarkable similarity to the topics covered in the present volume and most other contemporary books dealing with the psychology of reading. Huey and his contemporaries were interested in eye movements in reading, the nature of the perceptual span (how much information can be perceived during a fixation of the eye), word-recognition processes, inner speech, reading comprehension, and reading rate. Huey's marvelously cogent and concise description of his findings and those of his contemporaries prior to 1908 is still a joy to read. Many of the basic facts we know about eye movements during reading were discovered by Huey and contemporaries using cumbersome and seemingly archaic techniques in comparison to the sophisticated devices currently available to record eye movements during reading. Yet their discoveries have stood the test of time and have held up when replicated using more accurate recording systems. A contemporary of Huey, Emile Javal, the French oculist, first noted that during reading our eyes do not move smoothly across the page as our phenomenological impressions would imply. Rather our eyes make a series of jumps (or *saccades* in French) along the line. Between the jumps the eyes remain relatively still, for about a quarter of a second, in what is referred to as a *fixation*. A large number of experiments were carried out by Huey and his contemporaries to understand the work of the eyes in reading.

In order to study how much information can be perceived in a single eye fixation, the tachistoscope was devised. The t-scope (as it is often called) is a device that allows an experimenter to control how much information is presented to a subject, as well as the duration of the exposure. By varying the amount of information available in the t-scope and by presenting it for a duration brief enough to preclude any eye movement, early researchers hoped to infer the size of the *perceptual span* or the area of effective vision during a fixation. Huey's book also describes classic experiments by Cattell (1886) and by Erdmann and Dodge (1898) on word recognition, and two full chapters in the book are devoted to the role of inner speech in reading. Huey's observations on inner speech and word-recognition processes are lucid, and amazingly relevant to current issues.

Work related to the cognitive processes involved in reading continued for a few years after the publication of Huey's book. However, serious work by psychologists on the reading process pretty much came to a halt a few years after 1913. In that year, the behaviorist revolution in experimental psychology began. According to behaviorist doctrine, the only things worthy of study by experimental psychologists were activities that could be seen, observed, and measured. Since cognitive processes involved in skilled reading cannot be observed and directly measured, interest in reading waned between 1920 and 1960. While Buswell and Tinker carried out some well-known investigations of eye movements during reading, their work, for the most part, dealt with purely peripheral components of reading. Attempts to relate the activity of the eye to the activity of the mind were virtually nonexistent.

In essence, work on the cognitive processes associated with reading came to a standstill in the 1920s and did not begin again until the 1960s. Small wonder that when Huey's book was republished in 1968 it seemed so relevant. Not much had been learned about reading in the 60 years between the initial publication of the work and its second appearance. We hasten to point out that in addition to the work on eye movements during reading by researchers such as Buswell and Tinker, some work on reading did continue during the interval in question. But most of it was conducted in education schools where the primary focus is generally on more applied aspects of reading. Thus, there was work on the most appropriate method to teach reading, and many of the standardized reading tests still in existence today were developed during that period. However, work on the mental processes associated with reading was almost nonexistent.

Today, we find many psychologists interested in reading. Why has this change taken place? The primary reason appears to have been the failure of behaviorism to account for language processing in any reasonable way. The promise of behaviorism was always that if psychologists could understand the laws of learning and behavior in simple tasks (like knee jerks and eye blinks), those laws could be generalized to more complex tasks like language processing. In 1957, B. F. Skinner decided it was high time that the behaviorists produced on this promise, and he published *Verbal Behavior* which was an account of language from a behaviorist viewpoint. The linguist Noam Chomsky (1959) wrote a scathing review not only of the book but of behaviorism in general.

In essence, Chomsky argued that behaviorist principles could not account for language learning or language processes in general. Around that same time, he also published *Syntactic Structures* (1957), which was a radical departure from traditional linguistic theory. In that work, he suggested that the study of language and the mind are intimately related and presented an elegant theory of grammar. Many psychologists, disillusioned with behaviorism, became very interested in Chomsky's theories of language processing. After a hiatus of more than 40 years, work on cognitive processes was underway. (There were a number of other factors that contributed to the reemergence of the study of cognitive processes around 1960, but they are beyond the scope of our current discussion.) Out of the burgeoning interest in language processes in general, interest in the reading process began once again around 1970. Since the mid-1960s a number of scholarly journals dealing with cognitive processes and human experimental psychology have been founded, and each issue of these journals generally contains at least one article related to reading. In addition, a number of textbooks dealing with reading have appeared in the last 5 or 6 years. Clearly, there is now considerable interest among cognitive psychologists in studying reading.

It is important to note that cognitive psychologists studying reading approach the issue from slightly different perspectives. Some have a background rooted in perception research and see the study of word recognition, for example, as a means to study perceptual processes or pattern recognition using well-defined stimuli. Others approach the study of reading with a background in memory processes and verbal learning theory. They tend to approach the study of reading by examining comprehension processes. Still others are interested in reading in and of itself, because they believe, as Huey pointed out 80 years ago, that to understand what the mind does during reading would be "the acme of a psychologist's achievements, since it would be to describe very many of the most intricate workings of the human mind, as well as to unravel the tangled story of the most remarkable specific performance that civilization has learned."

It is our contention that this diversity of interests and backgrounds is healthy and can easily be accommodated within the information processing approach because it views reading as a highly complex process relying on a number of subprocesses. Indeed, most of the breakthroughs have come from researchers working on different subcomponents of the reading process and there have been few global insights that have been the key to answering many of the questions about a complex skill like reading. While it is clear that the information obtained by cognitive psychologists needs to be put together into a unified framework, the present state of the art justifies an emphasis on the different component processes of a complex skill.

Critics of the information-processing approach often argue that attempts to isolate component processes of reading result in tasks very much unlike reading. For example, to study word-recognition processes, cognitive psychologists often present a word for a very brief duration (say 50 milliseconds, which is one-twentieth of a second). A subject in such an experiment may be asked to pronounce the word or make some type of decision about it (Is it a word? Does it belong to a certain category of things? Is it larger than a breadbox?) In making decisions about the word, subjects push one button for a ''yes'' response and another for a ''no'' response. Admittedly, these tasks are unlike reading. Yet to respond appropriately, subjects may well be using the same processing mechanisms that they use during reading. Perhaps an analogy will help. Suppose we're interested in studying walking. If we study the motor responses that people make when they take two steps, critics may say, ''But that's not walking. When you walk you go a long way.'' True, but are the motor responses any different when you take two steps? Undoubtedly not. What cognitive psychologists strive to do is set up experiments in which the same processing mechanisms are used in the task derived as in reading. Sometimes we're more successful than other times. In this book, we will place the greatest weight on those experiments that most closely match the task of reading.

It is important to point out that the primary methodology of the cognitive psychologist is empirical experimentation. Theories and models of processes such as reading are also critically important because they help to formulate the kinds of research questions to be asked. Theories often arise from informal observations and intuitions as well as experiments. The ultimate test of a theory, however, is an experiment in which contrasting theoretical positions are tested against each other. With these points in mind, we now turn to a description of the human information processing system.

OVERVIEW OF THE HUMAN INFORMATION-PROCESSING SYSTEM

In this section, we will present an overview of the human information-processing system. However, we caution you that not all cognitive psychologists would agree with the notion of distinct stages. Indeed, the best way to describe the human information-processing system is quite controversial (see Craik and Lockhart 1972; Broadbent 1984). Consisting of three distinct stages (sensory store, short-term memory, and long-term memory), the type of model we will present is rather controversial at the moment, as is the assumption that the system is more or less passive. We will overlook such controversies because our primary intention is to give the flavor of the information-processing approach and to introduce terminology that virtually all cognitive psychologists use. In addition, some of the controversy over details is peripheral to understanding reading.

Figure 1.1 shows an example of a typical stage model of the human information-processing system. It consists of three stages and each stage has distinct functions and characteristics. However, prior to examining such a system, we must discuss something about the initial sensory registration of printed words by the eyes and subsequent pattern recognition processes.

The Retina and Visual Acuity

Vision depends on a pair of specialized organs (the eyes) whose neural receptors can be thought of as being part of the brain which has extended outside of the cortex. Patterns of light falling on the sensory neurons in the retina result in the sensation of seeing. When you look at a page of text (like the one you are currently reading), you are not able to see all of the words on the page equally well. This is because of *acuity limitations*. In terms of acuity, a horizontal line of text falling on the retina can be divided into three regions: foveal, parafoveal, and peripheral. The foveal area subtends about 2 degrees of visual angle around your fixation point; the parafoveal area subtends about 10 degrees of visual angle around fixation (4 degrees to the left and to the right beyond the foveal region); the peripheral area includes everything on the line of text beyond the parafoveal region. Acuity is greatest in the center of vision (the fovea) and drops off markedly in the parafovea and even more so in the periphery. This is because of the anatomical structure of the retina.

The retina is composed of two types of receptors called *rods* and *cones*. The fovea consists almost entirely of cones, and with distance from the fovea, the density of cones decreases and the density of rods increases.

FIGURE 1-1 An overview of the human information-processing system.

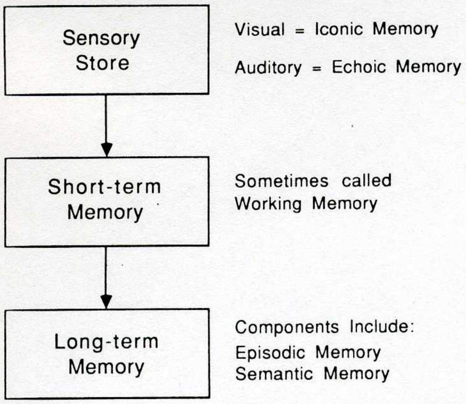

Thus, the peripheral region of the retina is composed entirely of rods. The parafovea contains a mixture of rods and cones. These two types of receptors serve dramatically different functions. The cones are specialized for processing detail and for acuity. In addition to permitting fine discrimination of detail, cones also serve in the discrimination of wavelengths or hue. The rods, on the other hand, are specialized for detecting movement and permit discrimination of brightness or shades of gray. The rods are particularly important for night vision; when you enter a dark room, at first you cannot "see" anything. However, after a short while (unless the room is totally dark) your rods adapt and you can see.

The most important point to be gleaned from the description of rods and cones is that the acuity necessary for discriminating fine detail (as is necessary in reading) is available only in the center of vision. A simple experiment can demonstrate this quite clearly. If you were asked to look into a tachistoscope and tell us the word (or letter) that appeared there, your accuracy in doing so would decrease as the stimulus was presented further from your point of fixation. The stimuli are presented briefly enough (about 150 milliseconds or less) so that it is virtually impossible for you to move your eyes to look directly at it. Figure 1.2 shows how performance in such a task would depend upon how close to fixation the stimulus was presented. In the figure, we have also plotted the relative distribution of rods and cones in the retina. Note that the accuracy function in our experiment is very similar to the distribution of cones in the retina. The purpose of this demonstration is to convince you that in order to discriminate the fine details of letters and words as we read, we must move our eyes to place the fovea over that part of the text we want to read.

FIGURE 1-2 **Relative frequency of the density of cones (solid line) and rods (dashed line) across the visual field. Dotted line shows the accuracy of identifying a target word exposed briefly to the left or right of fixation.**

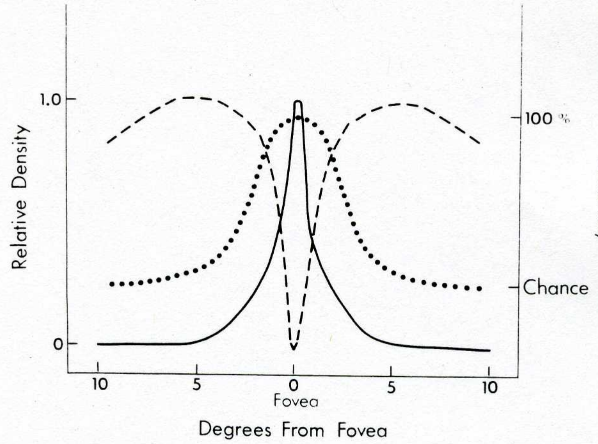

Pattern Recognition Processes

After we move our eyes to place the fovea on the word or words that we want to read, pattern recognition begins. Actually, the pattern-recognition process for a word may have begun on the prior fixation when the word was in parafoveal vision, as we shall see in Chapter 4. What we are concerned with in this section is how the brain goes about recognizing the letters and words which must be processed for us to read. To take a simple example, how do we recognize the printed letter *A*? Two major theories of pattern recognition have been proposed. The first *template-matching theory,* suggests that we have stored in our brains a representation of every pattern that we can recognize. Thus, we recognize the letter *A* by comparing the pattern of excitation from the cells in the retina to a template stored in memory. If there is a match between the representation and the stimulus, the letter *A* is perceived.

While template-matching theory works quite well in computer pattern-recognition devices that read letters and digits in highly constrained contexts, such as the digits that specify the code number for your checking account, it is well known (Neisser 1967; Crowder 1982) that such a system would fail to recognize instances of the letter *A* that were slightly deviant with respect to shape, size, or orientation as Figure 1.3 shows. The major problem for the theory in its most rigid form is that it suggests that we have a template for every variation of every pattern we are able to recognize. The simple pattern for the letter *A*, for example, can appear in a number of different typefonts and handwriting variations. Yet we are able to recognize it quite easily. It seems unwieldy to think that we have so many patterns stored in our head.

One way to make the template-matching theory more workable is to assume that before comparisons of new input to those stored take place, the input is "cleaned up" or normalized. The normalization process would separate essential information in the stimulus from most nonessential information. For example, variations in size could be taken care of before comparisons to a template occurred by transforming all images to a standard size. Accounting for variations in orientation prior to the matching process also seems quite reasonable. It is somewhat harder to understand how the normalization process would fill in missing gaps and eliminate fuzzy squiggles (as in handwriting) that are irrelevant to recognizing the pattern as the letter *A*.

While the normalization process gives plausibility to the template-matching theory, the second theory, *feature-detection theory,* is more parsimonious in accounting for how the pattern-recognition process deals with such variation and is generally considered to be a more viable account. The starting point for feature-detection theory is the idea that there are many common elements for letters (consisting of horizontal, vertical oblique, and curved lines) and that we analyze these component elements in recognizing a pattern. The letters *C* and *G* and the letters *O* and *Q,* for

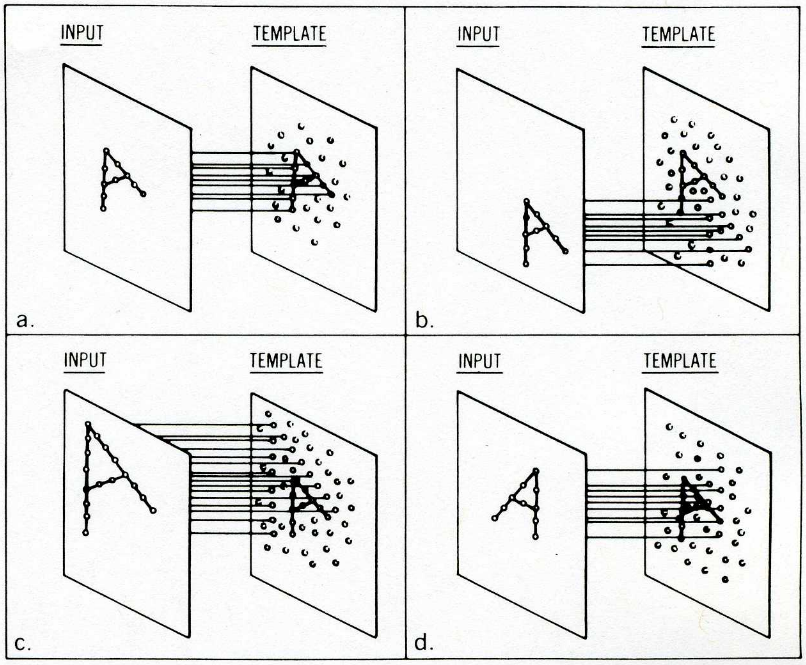

FIGURE 1-3 **Illustration of the difficulties with a simple template-matching system for character recognition. (After Neisser 1967, with permission of Prentice-Hall.)**

example, have a great deal of similarity and featural overlap. The distinguishing feature between the *C* and *G* is the horizontal line that is present in the *G*, but not in the *C*. The distinguishing feature between the *O* and *Q* is the oblique line present in the *Q* but absent in the *O*. According to feature-detection theory, when a letter is analyzed the first step is to prepare a list of its features and this list is compared with the list stored in memory. This process then is *analytical* in that we put together the different elements of the pattern until recognition occurs. Template-matching theory, on the other hand, is more of a *holistic* process. As we implied, the feature-detection theory is considerably less cumbersome because the analytic processes rely on a small number of features that are common to all typefonts (the distinguishing feature for *C* versus *G* remains invariant over different typefonts and handwriting styles), as opposed to having a different template for each style of print.

What type of evidence is there for feature-detection theory? Three types of evidence are consistent with the theory: (1) physiological data from animals, (2) stabilized image research, and (3) visual search data.

The best known physiological evidence in favor of feature detection theory comes from work by Hubel and Wiesel (1962) on the visual system of the cat. Using electrical recordings, Hubel and Wiesel were able to examine the rate of firing of individual cells in the visual system as a function of what the cat was looking at. The most important finding from their work for our purposes is that they demonstrated that cortical cells in the visual system fired differentially depending on the stimulus. Line, edge, and slit detectors, as well as more complex detectors, were all discovered. It is easy to generalize from these results (although it should be done cautiously) and suggest that there are also feature detectors in humans specialized for firing when the various horizontal, vertical, oblique, and curved lines making up the letters of our alphabet fall in front of our eyes.

The stabilized image experiments use sophisticated technology to hold an image fixed or stabilized on the retina; to whatever extent the eyes move, the presented stimulus moves a corresponding amount in the same direction. In contrast, even when we are asked to hold our eyes very still in normal vision, there is a slight movement or tremor of the eyes (called *nystagmus*). Apparently, the nystagmus is important for perception because under stabilized image conditions perception gradually blanks out so that the observer no longer sees the stimulus. What is interesting is that perception of the stabilized image does not blank out instantaneously. Instead, it is gradual, and lines sharing the same orientation disappear at the same time as depicted in Figure 1.4. The manner in which the stimulus fades is further support for feature-detection theory, since the units fading out appear to be features of the objects.

FIGURE 1-4 **Perceptual fragmentation in "stopped images." The figures at the left are stimulus patterns; the others are typical products of fragmentation. (From Pritchard 1961, with permission of Scientific American).**

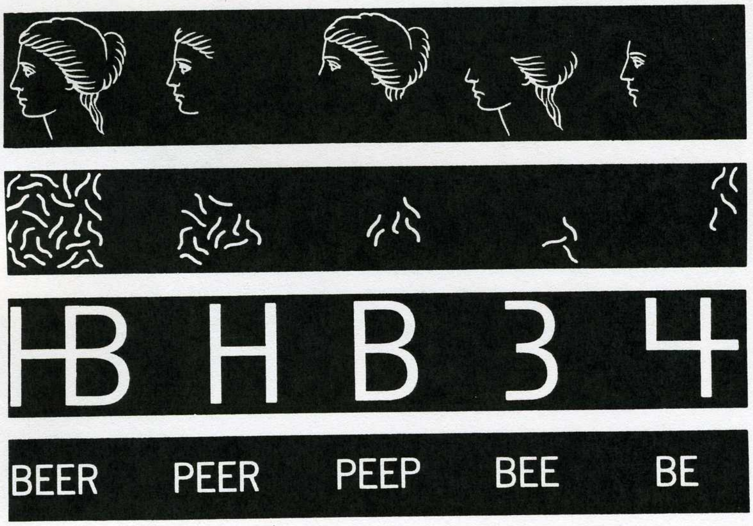

STANDARD CONDITION TARGET = *Z*	NONCONFUSABLE CONDITION TARGET = *Z*	CONFUSABLE CONDITION TARGET = *Z*
RYVMKF	CBSOGS	VMWNMW
PTHSHG	UBSQQQ	WYLKWV
GTVCBH	BQOUDG	XMWLLY
HUIRYD	SCDOBC	YXZWXL
KIREGD	CZOQUS	NMWYMN
GBZTBN	DUBCCD	YNLXLI
POLKRF	DOQUCB	MNWXMH
FTIEWR	CCOQOU	LYXWLT
. . .	. . .	. . .
. . .	. . .	. . .

FIGURE 1-5 **Neisser's visual search task. A display is exposed and the subject must scan vertically, from top to bottom, until finding a target item. Time to reach the target is recorded as a function of the position of the target.**

The third line of evidence for feature-detection theory is the work on visual search originated by Neisser (1967). Figure 1.5 shows sample stimuli in the task. You are asked to find the target letter *Z*. It turns out that the letter *Z* is much easier to find (subjects are over twice as fast) in the middle column than in either of the other columns. The reason, of course, is that the distractor letters in the middle column are all letters that are quite dissimilar to the target and do not share many features with it. Template-matching theory would suggest that there should be no difference between the columns because the number of matches that must be made would be equivalent in each. However, feature-detection theory can account for the result, since virtually the same set of feature-detector cells would be firing when the distractors are confusing whether or not the target letter was present. On the other hand, when the target is embedded in dissimilar letters, the cells that fire for the letter *Z* will only fire when the target is present.

While there are criticisms of feature detection as a theory of object perception in general, it provides a reasonably satisfactory model of how letters and words are processed. One issue that we have not discussed, which will become relevant in a later chapter, is whether pattern recognition processes occur in a *serial* or *parallel* manner. If information is analyzed serially, the pattern-recognition processes occur one at a time. Thus, in Neisser's visual-search task, feature analysis would occur for one letter and when it was complete would begin on the next, and so on. If information is analyzed in parallel, the sensory information from various spatial locations is analyzed simultaneously. Thus, within acuity limits, all of the letters in each row in Figure 1.5 would be analyzed at once. The experiments on word identification that we will discuss in Chapter 3 also argue for parallel processing.

In this section we have argued that the pattern-recognition processes that operate on print can best be conceptualized in terms of feature-analytic processes. We return now to a more detailed examination of the information-processing system depicted in Figure 1.1.

The Sensory Store

The lowest level in the information-processing system is generally referred to as the *sensory store*. For auditory information, the store is referred to as *echoic memory* (see Cowan 1984 for a discussion). For visual information, the store is referred to as *iconic memory*. Iconic memory is considered to be a temporary memory store in which much of the information physically available in the stimulus is still available after the display has gone off.

In one of the most widely cited experiments in cognitive psychology, Sperling (1960) demonstrated the existence of a visual sensory store. He presented subjects with three rows of 4 letters each for a brief duration (say 50 milliseconds or one-twentieth of a second). In a control condition, subjects were instructed to report as many of the letters as possible, and the average number of letters reported was 4.5. In the experimental conditions, Sperling used a partial report technique in which he cued subjects as to which row they should report. Thus if after the display went off, subjects heard a high pitch tone they were to report the top row, a medium tone the middle row, and a low pitch the bottom row. The tone was presented at various delays after the disappearance of the letters from the screen. Sometimes the tone occurred simultaneously with the offset of the letters and sometimes it occurred up to a second after the letters disappeared.

Sperling's reasoning was that if we have an iconic memory (though he did not call it such), we should be able to use the cue to "read out" the letters from the cued row. If the tone actually occurred prior to the onset of the letters, then we would be able to use that information to focus on the relevant row. Sperling's question was whether or not we can use the cue to examine information after it has physically disappeared from before our eyes. If so, then we must have an iconic memory in which all or most of the information present in the stimulus is still available. In fact, what Sperling found was that subjects could report on average 3.3 letters from the array, which means they must have had approximately 10 letters available in their memory system. Since the subject had no way of knowing beforehand which row was going to be signaled, it must be the case that subjects could get roughly 3.3 letters from any of the rows.

Keep in mind that when subjects in the control condition were asked to report as many letters from the display as they could, they averaged 4.5 letters. How do we account for the discrepancy between the cued partial report trials and the whole report condition? If we assume that iconic memory has a large capacity (say 10 letters from a 3-by-4 array) and that the rate at which people can "read out" letters is slow (say 3 to 5 letters in a quarter of a second), then the discrepancy can be explained. If the letters are taken from the whole display, performance is limited by the read-out process. On the other hand, if the tone cues the subject that only one row is relevant, then the subject has time to read out most of that row.

Since Sperling's demonstration, there have been literally hundreds of experiments investigating the characteristics of iconic memory. This research has revealed that the primary characteristics of iconic memory are that (1) the memory has a large capacity, (2) it has a duration of roughly a quarter of a second, (3) it is precategorical, (4) it is interfered with by new information, and (5) the read-out rate is relatively slow.

The current status of iconic memory is highly debatable (see Coltheart 1980a; Haber 1983; Turvey 1977). Some workers in the field are now arguing that it is an *epiphenomenon* (i.e., a phenomenon that occurs but is of no real functional significance) of the highly sterile and controlled experimental laboratory; after all, when does a stimulus appear before our eyes for a fraction of a second only to disappear completely? Such individuals tend to argue that iconic memory, like our appendix, has no functional utility. Others argue that all biological mechanisms are adaptive, and just because we do not know for certain what function the icon serves, does not mean it does not have a role in processing.

With respect to reading, it is not at all clear what function iconic memory might serve. Clearly, subjects in iconic-memory experiments are not reading in the sense that we would normally think of reading. In reading, the stimulus does not disappear from in front of our eyes after only a brief exposure (unless perhaps we try reading in a lightning storm). At one time, it was thought that something similar to an iconic memory played a role in integrating visual information across the eye movements we make during reading. We will discuss that idea in some detail in Chapter 4, but to anticipate that discussion a bit, the available evidence argues against such a conclusion. Indeed, the fact that we make eye movements so frequently is a problem for the utility of iconic memory in reading. Recall that the duration of iconic memory is roughly a quarter of a second, which is about the rate at which we make eye movements. Given that information in iconic memory is disrupted by new information and that eye movements occur at the rate they do, plus the fact that the information we want to read is available to us continuously (we can always look back with our eyes), it does not seem that iconic memory plays any role in reading.

At this point, you may be asking yourself: if iconic memory plays no role in reading why was it discussed in such detail? Indeed, why was it discussed at all? There are two reasons why we troubled to present the details of iconic memory. First, our primary purpose in Part 1 is to present an overview of the human information-processing system. Numerous information-processing models of the reading process (Mackworth 1972; Massaro

1975; Mitchell 1982; Gough 1972) have used the concept of iconic memory as the initial stage of registration of visual information during reading. Recall that we argued that at this point we are not so concerned about the extent to which the different stages are accurate and useful in understanding reading; they are presented to give you a flavor for the approach. Which brings us to our second reason for discussing iconic memory in such detail: the notion of *buffers* (or temporary storage units) in information processing turns out to be highly useful. Actually, the icon is little more than a buffer in which information is held for some later processing activity. The short life of the icon in fact suggests that the visual features of the print are of little use once the eye is no longer looking at them. As we shall see later, the notion of a buffer has been very useful in various types of research related to the reading process. By discussing iconic memory in such detail, we hope you will have a sense for how such a concept may be useful in designing experiments and theorizing about reading.

To summarize, iconic memory is the initial stage in an information-processing model. Although highly transient, it has a large capacity. We have also argued that its usefulness for understanding reading is limited since the stimulus is always available to us in reading. However, the concept of a bufferlike store has been useful in experiments related to the reading process.

Short-Term Memory

According to the standard view of the information-procesing system, due to the transient nature of iconic memory we need to get information registered by the sense organs into a more permanent structure. The structure is *short-term memory* (STM). Considerable information is lost before it can be transferred to STM because the read-out rate from iconic memory is quite slow. A certain amount of information, however, is transferred to STM, which it turns out has problems of its own. First, and most importantly, it has a limited capacity. The capacity of STM is about 7 plus or minus 2 items (Miller 1956). Notice that we said "items," not letters, words, or digits. Indeed, we can learn to short-circuit to some extent the limited capacity of short-term memory by various types of *chunking strategies*. If the number 967835241 is presented to you orally and you treat each individual digit as an item you will have a difficult time recalling it in the same order. Quite simply, for most people, STM will be overloaded. You will, however, be able to recall the number with 100 percent accuracy if you treat it as three 3-digit numbers (967-835-241). Another way that we deal with the capacity limitation of STM is through a process called *rehearsal*. When you look up a telephone number in the phone book, you often find yourself repeating it (often silently, but sometimes aloud) so that you won't forget it. In other words, you rehearse the number over and over. Such a strategy is another way to hold information in short-term memory.

Notice we said that you often repeat the number over to yourself silently. For a long time, it was considered that STM was exclusively an acoustic store. That is, even information coming in the visual modality was assumed to be recoded into acoustic or auditory information. The reason for this was that the kinds of errors that subjects made in recalling information in STM tended to be acoustically, not visually, related to the information actually presented. We now know that there are visual and semantic codes in STM. Still, for linguistic stimuli, STM is primarily acoustic as evidenced by the fact that we try and remember phone numbers from the telephone book by rehearsing them subvocally. This aspect of STM turns out to be particularly important for understanding the role of subvocal or inner speech in reading.

The fact that we engage in various strategies (some of them unconscious) to short-circuit the limited capacity of STM has led some workers to refer to it as *working memory* (Baddeley and Hitch 1974). That is, STM can be considered a flexible workspace whose limited capacity can be allocated to either storage or processing. Information in STM can remain there as long as it is being worked on. Working memory, in the sense of a flexible workspace for processing, is also heavily involved in reading. Words are * integrated in this memory, and as we shall see later, comprehension processes are initiated here.

To summarize, STM has a limited capacity. Using rehearsal, however, we can hold items there for long periods of time. We also develop efficient strategies for dealing with the limited capacity. Short-term memory is also primarily acoustic. Whereas iconic memory was argued to have limited usefulness in understanding reading processes, the characteristics of STM are important in understanding inner speech in reading and comprehension processes.

Long-Term Memory

The rate at which we can set up programs to transfer information from STM to long-term memory is relatively slow in relation to the rate at which new information enters STM, so that considerable information is lost. However, it is generally believed that once information enters *long-term memory* (LTM) it is stored there permanently. Patients under a local anesthetic whose brains have been electrically stimulated can remember things they long since thought they had forgotten and even relive memories of events that occurred a long while in the past (Penfield and Roberts 1959). Information in LTM is not organized in a haphazard fashion. Indeed, LTM is highly organized and much of the material that we cannot retrieve has been mislaid, not lost. The major problem with LTM is getting the appropriate retrieval key to access information stored there. This is not surprising given the vast amount of new information we process and store in LTM each day. In addition, there is evidence that the new information we learn interferes with our ability to retrieve previously stored information. Conversely, information already stored in LTM can interfere with retrieving newly learned information.

Most cognitive psychologists now believe that it is appropriate to think of two types of long-term memory: *episodic memory* and *semantic memory* (Tulving 1972). Episodic memory is the memory for sequences of events in your life. Semantic memory, which is more important for understanding reading, contains general knowledge you have. A part of semantic memory that is important for reading is the *lexicon*. The lexicon, which like LTM itself is highly organized, contains the meanings of the 30,000 or more words that you know. The goal of most reading is to understand something new and to store the gist of it in LTM. To do so involves processing the meanings of the words we know, or accessing our lexicon in LTM. Further, to understand idioms, metaphors, and the like, we have to use general world knowledge that we have stored there. And when authors are a bit vague, we must make inferences based on what we already know to understand their point.

Selection of Information

An issue that relates to the conceptual framework which we have presented is how information is selected to be passed on to the next stage of processing. In vision, of course, the eyes are a major device for selection. You point your eyes at those stimuli you want to process and ignore others. As we discussed earlier, this overt selection process is not all-or-none. While stimuli seen in extrafoveal vision are processed less well than those in foveal vision, they are processed.

However, pointing the eyes is not the only selectional mechanism in vision. In our discussion of Sperling's experiments, we tacitly assumed that there was a selection process which could help "read out" the appropriate row of letters. A great deal of recent research has documented the reality of such a covert mechanism of spatial attention. In essence, even when they do not move their eyes, human and animal subjects can respond to visual stimuli more quickly and accurately when they are cued as to where in extrafoveal vision these stimuli will be (Posner 1980). Furthermore, the locus of this attentional system in the brain is becoming increasingly well understood. We will return to these selectional issues in Chapters 4 and 5.

In contrast to the above attentional mechanisms, the processes by which information is selected to go to LTM (or is made more retrievable from LTM) are less well understood. Clearly, factors such as your motivation to remember the material, the length of time it is in STM, and the meaningfulness of the material all affect how well you will remember it later. We will touch on these issues again when we discuss memory for discourse in Chapter 8.

The Concept of Processing Stages

An assumption of the model we have outlined above is that there are discrete processing stages. That is, it is assumed that information is processed in one stage and, only when that processing is completed, is the information shipped on to the next stage. This assumption underlies a great deal of cognitive psychology because processing will be much easier to study if it can be broken down into simpler components.

Sternberg (1969) proposed a widely used test to determine whether such a stage notion is valid. The test assumes that the dependent variable used to study mental processing is response time, or the time between the presentation of a stimulus and the time to execute a response. We can best explain the idea with an example from Sternberg's experiment. He used a memory search task, in which subjects were initially given a *memory set* (1 to 6 digits) to hold in memory, and then they had to indicate if a *probe* digit (which was presented visually) was in the memory set. His finding was that the time to determine whether the probe digit matched one of the items in the memory set increased as the memory set got bigger. In fact, the time increased by about 40 milliseconds (msec) with each additional memory set item, suggesting that it took subjects 40 msec to compare each additional memory item with the probe.

Thus, a measurable search process appears to be occurring in STM. What would happen if the probe digit was presented in "visual noise" (i.e., it was embedded in a lot of random dots so that it was harder to see)? If we view the process of identifying the digit as a stage prior to searching STM, then the digit should take longer to identify when presented in noise, but the rate of search in STM should be unaffected since the identification stage would be complete regardless of whether the digit was "clean" or "noisy." In contrast, if identification and search are not discrete stages so that visual noise is still part of the item being compared to the memory set items, one would expect that the search time per item would increase. In fact, Sternberg obtained the former result—overall times increased with a "noisy" probe digit but memory search times did not—and concluded that identification of the digit was a stage prior to STM search. The basic logic of this experiment has been used widely in cognitive psychology and will come up again at several points in the book.

The Reality of Information-Processing Models

If this is your first exposure to information processing, you may be asking yourself to what extent the different structures presented in the model have been localized in the brain. The answer is that, for the most part, they haven't. Neurophysiologists working on brain functions have found chemical changes in the brains of animals during learning stages that could correspond to STM functions. In addition, studies of brain localization have revealed different functions of different parts of the cortex (especially language functions). However, there is not likely to be an anatomical division between STM and LTM. Nor have we localized the lexicon in the cortex.

In fact, the concept of the mind is a rather abstract entity. The mind is the concept that we can think of as being the executor responsible for cognitive activity, which presumably can be ultimately explained in terms of the structure and function of the brain. The task of the cognitive psychologist is to learn how the mind is structured and functions. If it were possible, perhaps an ideal way to study reading would be to open up a reader's brain and observe what kinds of activities and changes occurred during reading. But we cannot do this. Thus, cognitive psychologists are forced to infer characteristics of how the mind works in skilled cognitive tasks, like reading, on the basis of various types of evidence that we can accumulate. In a sense then, a cognitive psychologist is like a detective searching for clues to how the mind works. The type of structures presented in Figure 1.1 and elsewhere in this book serve as a convenient way of hypothesizing about how the mind works and then summarizing what we know. Throughout this book, we will present evidence accumulated by cognitive psychologists about how the mind works in reading, and we will frequently use diagrams to present the information-processing flow, such as in Figure 1.1, to summarize what we know. But it would be a mistake to think of these structures as necessarily mapping directly onto parts of the brain.

Brain Function and Reading

In the prior section we differentiated between the brain and the mind. While much of our focus will be on how the mind works in reading, we also know that there are specific brain functions related to reading. In this section, we briefly review some of them.

Information registered on the retina is transmitted to the visual cortex. The cerebral cortex of human beings is divided into two hemispheres that have different but complementary functions. Some functions are bilaterally represented in the brain. For example, there is a visual area at the back of both of the hemispheres. However, some areas, particularly those associated with language processing, seem to be localized almost exclusively in only one hemisphere. For most people, regions of the left hemisphere are responsible for language processing. Regions of the right hemisphere, on the other hand, are specialized for non-verbal, spatial processing. In some left-handed people the functions of the two hemispheres of the brain may be reversed. The two hemispheres are connected by a bundle of nerve fibers called the corpus callosum.

We know about the different functions of the two hemispheres from two types of evidence. First, experiments on normal subjects often take advantage of the fact that information seen in the left half of the visual field of either eye arrives initially in the right hemisphere and things seen in the right half of the visual field arrive initially in the left hemisphere. Material presented to the center of vision is simultaneously available to both hemispheres of the brain. (With respect to the left and right ears, a similar pattern holds as for the presentation of visual information.) From experiments in which stimuli are briefly presented in the left or right visual field, we know that words are processed more efficiently when presented in the right visual field (implying left-hemisphere processing), whereas faces and pictures are processed more efficiently when presented in the left visual field. From such experiments, it is often argued that the left hemisphere operates in a serial and analytic fashion and the right hemisphere operates in a parallel and holistic fashion, although there is no compelling reason for this conclusion. The second way we know about the functions of the two hemispheres is from research on both brain-damaged patients, who have one of the two hemispheres missing (from birth or because of brain injury), and "split-brain" patients, who have had the corpus callosum severed as a treatment for epilepsy. With respect to language, the basic evidence is that if certain regions of the left hemisphere are damaged language functions are impaired, but right-hemispheric damage does not produce language impairment. For the "split-brain" patients, the evidence is that linguistic information put into the left hemisphere is processed normally, whereas there is little comprehension of linguistic information put into the right hemisphere. Recently, a great deal about the reading process has been learned by examination of patients with brain damage (and known lesion sites). In Chapter 11, we will review the evidence obtained from such patients. We will omit the details of the physiology (such as the location of language-specific sites) here and in Chapter 11 since such information adds little to our understanding of reading.

WHAT IS READING?

In this chapter so far, we have presented preliminary information necessary to understand how cognitive psychologists think about reading. This brings us to a critical question. What do we mean by "reading"? It is obvious that to many people, reading is an all-encompassing activity that can take on different forms. For example, when you look at a map, are you reading? When you proofread a paper, are you reading? When you look at a computer program to find your programming error, are you reading? When you scan a newspaper for the latest results of the stock market, are you reading? We will take the conservative view that none of these activities are what we have in mind as reading. It is also obvious that when you read a novel on a 5-hour airplane trip, you may at times be reading slightly differently than when you read this book. Four hours into your trip you find that you are only half way through the book, so you start skipping over sections which seem redundant looking only for relevant and new information so you can finish it before reaching your destination. You would have a difficult time understanding a textbook if you read in such a fashion, yet we can generally skim most novels and still understand the story. In our chapter on speed reading, we will discuss skimming and the adjustments the reader makes under such conditions. However, apart from that chapter, we will focus on the rather careful type of skilled reading that occurs when you read to comprehend a textbook, a newspaper article, or a narrative.

It would be easy at this point to get into a lengthy argument about what is and is not reading. We do not wish to do so. We hope it is clear what we have in mind by reading. If forced to provide a definition of reading, we would probably say *reading is the ability to extract visual information from the page and comprehend the meaning of the text*. By focusing on the careful reading of a newspaper article, for example, we do not wish to imply that the other activities mentioned are not interesting. Our bias is that activities such as proofreading and skimming probably involve strategies and processes that are different from normal silent reading. At places, we will examine such tasks. However, our central concern is how people read during normal silent reading.

This brings us to a second critical question. What is the best way to study reading? The answer to the question depends on which aspect of the reading process you are interested in studying. We mentioned earlier that cognitive psychologists interested in word recognition generally present isolated words to people in a tachistoscope and ask them to make some kind of judgment about, or response, to that word. We shall discuss such tasks in detail in Chapter 3. Other researchers interested in the role of inner speech in reading have devised clever techniques to determine its role in understanding written language. If researchers are interested in how much of the text the reader comprehends, then they would want to examine how well readers can answer questions about the content of the text. Techniques used to study inner speech and comprehension will be discussed in chapters 6 and 8, respectively.

If the goal is to study the cognitive processes that occur during (normal silent) reading of text on a moment-to-moment basis, then any technique that has readers do something different, such as read words in isolation or read text out loud, may significantly distort the component process in silent reading one wishes to study, such as word identification or the role of acoustic codes in reading. While it is plausible that the components of reading do not change radically from task to task, there is no guarantee. Thus, the relevance of any technique is an open question if we don't know how the processes work during silent reading.

This brings us to eye-movement recording, the primary technique used to study cognitive processes during actual silent reading. Recording of eye movements has a long history in experimental psychology as noted earlier in this chapter. However, it has only been of late that eye-movement data have been widely used to infer moment-to-moment cognitive processes during reading. It is now fairly clear that where readers look and how long they look there provides valuable information about the mental processes associated with understanding a given word or set of words (Just and Carpenter 1980; Rayner 1978a). Eye-movement recording can be accomplished in a variety of ways, but often involves shining onto the eye a beam of invisible (infrared) light that is reflected back from the cornea or retina to a sensing device. With this methodology, readers are free to look at any part of the text for as long as they wish. As mentioned above, the technique also has a great deal of ecological validity in that subjects in eye-movement experiments are actually engaged in the task that we wish to study, namely reading.

This is not to say that eye-movement recording is free from criticism. In order to distinguish movements of the eyes from movements of the head, it is often necessary to stabilize the head. This is often done by using a bitebar (which consists of dental compound that is very soft when a subject bites into it but quickly hardens to provide a dental impression that keeps the head still). In other cases, forehead rests are used, and subjects generally read from a display placed directly in front of them. Some critics have suggested that the rigid constraints on head movement, plus the fact that in reading—outside of the eye-movement laboratory—we often look down at the text (rather than straight ahead), will lead to different reading strategies. It has even been suggested that the mere fact that our eye movements are being recorded will make us conscious of them and lead us to do something different when we read under such circumstances. Our impression is that these concerns are all ill founded. Indeed, Tinker (1939) demonstrated quite some time ago that the reading rate and comprehension of subjects in a soft easy chair with a book did not differ from the reading rate obtained in the eye-movement laboratory.

Both of the authors have been subjects in experiments using eye-movement recordings. Our firm impression is that reading in the eye-movement laboratory is virtually the same as reading outside of it, and it is definitely our sense and intuition that this latter technique provides a much better approximation of reading, itself, than any other technique. But we do not want to argue that eye-movement recording is the only way to study skilled reading. Many of the techniques mentioned throughout this book provide useful information, but the best type of evidence will be converging data, in which information obtained from a number of the techniques converge on the same answer to a given question. Our intention is to use converging evidence from a number of sources to understand reading, but our greatest emphasis will be on the data obtained while the subject is reading connected text, rather than simply being engaged in one of the clever tasks cognitive psychologists have devised.

MODELS OF READING

While there are many facts about reading that have been learned by cognitive psychologists, many people often find cognitive psychology somewhat frustrating because there is often conflicting evidence on a single issue. There are many reasons why this may be the case, including the fact that our experiments are sometimes not very good. But another reason is that cognitive psychologists often have different *models* or *theories* of how some mental process works. What are models and theories? Let's borrow from Carr (1982) in defining these two concepts. A *theory* is a set of principles (assumptions or rules or laws) that together constitutes a verbal or mathematical description of an interesting phenomenon, and an explanation of how or why the phenomenon happens. A theory defines the important characteristics of a phenomenon that are then included in a model of the phenomenon. A *model* represents a description of the major working parts of a real-life process (such as reading). The description captures the most important characteristics of each part's operation, though it might leave out large amounts of detail. Currently, there are a number of models of the reading process that, in our opinion, vary in the extent to which they capture important aspects of the skill.

We shall not attempt to describe various models of reading here. Rather, let us simply characterize these models as being primarily (1) *bottom up,* (2) *top down,* or (3) *interactive.* Incidentally, these three types of models are characteristic not only of the reading process but of the descriptions of most of the tasks and phenomena that cognitive psychologists typically investigate. Some books on reading (Just and Carpenter 1987; Smith 1971) present their audience with a model of reading and then interpret relevant evidence within the framework of that model. Some books on reading (Crowder 1982; Downing and Leong 1982; Gibson and Levin 1975) manage to avoid presenting a model of reading altogether and present only the facts as interpreted by the authors (in some cases the rationale is that a single model cannot capture the complexities of reading or the varieties of types of reading). Other books (Mitchell 1982; Perfetti 1985; Taylor and Taylor 1983) present evidence first, and then on the basis of the evidence describe a model of the reading process. We will adopt this latter strategy and present you with the facts as we see them. There is a danger in presenting the model first and then fitting the facts to the model because such a strategy often makes it sound as though we know more than we really do. We also suspect that researchers often become committed to a particular model of the reading process and that the model itself then becomes more important than the data collected.

We feel that most of these models are little more than general frameworks which provide some biases about which aspects of reading are really important. Our discussion of models will indicate many of our biases and provide the "bare bones" of a general framework. This framework will acquire more detail as we progress through the book. In the final chapter, we will briefly summarize the framework that has evolved throughout the book.

Bottom-up models (Gough 1972) stress that most information flows in a passive manner through the human information-processing system. The major idea is that this flow of information is very fast and that knowledge we have stored in memory has little impact on how the processing takes place. In contrast, proponents of top-down models (Goodman 1970; Smith 1971) feel that the passive flow of information through the processing system is relatively slow because there are numerous bottlenecks (places where the architecture of the system forces us to slow down). Accordingly, to short-circuit the bottlenecks these models stress that we rely heavily on information stored in memory (general information that we have about the world) to help speed up our processing. The primary way in which readers short-circuit the information-processing bottlenecks is to formulate hypotheses about what they will next read. This view of reading, often referred to as the *hypothesis-testing model,* was once very popular. However, evidence now suggests that the visual processing of text is very fast and that the extent to which readers engage in hypothesis testing or guessing behaviors seems to play a minimal role in the process of reading. We will return to this issue at various points throughout the chapters that follow. For now, let us simply state that a bottom-up view of reading more accurately characterizes the available evidence. That is not to say we do not think that top-down processes play no role in reading. They clearly do. Perhaps our model of the reading process can best be described as a bottom-up model in which the reader gets some help from top-down processes.

We have told you briefly what bottom-up and top-down models are, but we have not yet mentioned interactive models. Interactive models (Just and Carpenter 1980; McClelland 1986; Rumelhart 1976) allow for all sorts of communications between top-down and bottom-up processes. Proponents claim that these models are very good in accounting for the data on reading processes. Critics argue that while these models may be able to account for lots of data, they are unconstrained and hence do not predict very well what the outcome of any particular experiment might be. In contrast, the major virtue of most bottom-up models is that they are very good at making clear predictions about performance measures.

The view of reading that we will be presenting will largely be a bottom-up view, but with some influences from top-down processes. Notice that we have used the word *process* a number of times in this discussion. Elsewhere in this book, we will make a distinction between the *process* of reading and the *product* of reading. The product of reading is the information that gets stored in memory; it is what gets comprehended during reading. The major emphasis in this book is on the process rather than the product of reading (though the latter will be discussed in Chapter 8) because, from our point of view, the most important thing to understand about reading is the process. This is a bias that not everyone would agree with. For example, educators would undoubtedly argue that knowing the best way to teach children to read is more important than understanding the process of skilled reading. While we appreciate their opinion, our sense is that if we can understand the process of skilled reading, we may well be able to provide useful information to educators about what they are trying to teach. In essence, we believe that understanding the process of skilled reading should provide firm conclusions about how to instruct novices to become skilled in the task. Our discussion in chapters 9 and 10 will highlight some of the ways that we believe research has made clear how children should be instructed to learn to read.

Some cognitive psychologists who study the product of reading would also want to argue with us concerning our bias towards understanding the process of reading. To their way of thinking, what people remember from what they read may be more important than how they go about the chore of reading. However, our response to such a point is that understanding the process by which some mental structure is created almost logically entails understanding that structure. In contrast, understanding what gets stored in memory may not reveal much about the processes that created the structure. Thus, understanding what is in memory as a result of reading discourse may not be unique to reading; essentially the same structures may be created when people listen to discourse. We are not saying that understanding the product of reading and how that product gets remembered is not important. It's just that reading is a remarkable skill that must be understood—quite apart from issues like general comprehension skills and intelligence.