Psychological Review

VOLUME 87 NUMBER 4 JULY 1980

A Theory of Reading: From Eye Fixations to Comprehension

Marcel Adam Just and Patricia A. Carpenter
Carnegie-Mellon University

This article presents a model of reading comprehension that accounts for the allocation of eye fixations of college students reading scientific passages. The model deals with processing at the level of words, clauses, and text units. Readers make longer pauses at points where processing loads are greater. Greater loads occur while readers are accessing infrequent words, integrating information from important clauses, and making inferences at the ends of sentences. The model accounts for the gaze duration on each word of text as a function of the involvement of the various levels of processing. The model is embedded in a theoretical framework capable of accommodating the flexibility of reading.

Although readers go through many of the same processes as listeners, there is one striking difference between reading and listening comprehension—a reader can control the rate of input. Unlike a listener, a reader can skip over portions of the text, reread sections, or pause on a particular word. A reader can take in information at a pace that matches the internal comprehension processes. By examining where a reader pauses, it is possible to learn about the comprehension processes themselves. Using this approach, a process model of reading comprehension is developed that accounts for the gaze durations of college students reading scientific passages.

The research was supported in part by Grant G-79-0119 from the National Institute of Education and Grant MH-29617 from the National Institute of Mental Health.

We thank Allen Newell and Robert Thibadeau for their very helpful discussions.

The order of authorship was decided by the toss of a coin.

Requests for reprints should be sent to Marcel Adam Just, Department of Psychology, Carnegie-Mellon University, Pittsburgh, Pennsylvania 15213.

The following display presents an excerpt from the data to illustrate some characteristics of eye fixations that motivate the model. This display presents a protocol of a college student reading the first two sentences of a passage about the properties of flywheels. The reader averages about 200 words per minute on the scientific texts. In this study, the reader was told to read a paragraph with understanding and then recall its content. Consecutive fixations on the same word have been aggregated into units called *gazes*. The gazes within each sentence have been sequentially numbered above the fixated word with the gaze durations (in msec) indicated below the sequence number.

One important aspect of the protocol is that almost every content word is fixated at least once. There is a common misconception that readers do not fixate every word, but only some small proportion of the text, perhaps one out of every two or three words. However, the data to be presented in this article (and most of our other data collected in reading experiments) show that during ordinary reading, almost all content

Eye fixations of a college student reading a scientific passage. Gazes within each sentence are sequentially numbered above the fixated words with the durations (in msec) indicated below the sequence number.

1	2	3	4	5	6	7	8	9	1	2
1566	267	400	83	267	617	767	450	450	400	616
Flywheels	are	one	of	the	oldest	mechanical	devices	known to	man.	Every internal-

3	5	4	6	7	8	9	10	11	12	13		
517	684	250	317	617	1116	367	467	483	450	383		
combustion	engine	contains	a	small	flywheel that	converts	the jerky	motion	of the	pistons into the		

14	15	16	17	18	19	20	21
284	383	317	283	533	50	366	566
smooth	flow	of	energy	that powers	the	drive	shaft.

words are fixated. This applies not only to scientific text but also to narratives written for adult readers. The current data are not novel in this regard. The eye fixation studies from the first part of the century point to the same conclusion (Buswell, 1937, chap. 4; Dearborn, 1906, chap. 4; Judd & Buswell, 1922, chap. 2). When readers are given a text that is appropriate for their age level, they average 1.2 words per fixation. The words that are not always fixated tend to be short function words, such as *the*, *of*, and *a*. The number of words per fixation is even lower if the text is especially difficult or if the reader is poorly educated. Of course, this is not the case when adults are given simple texts, such as children's stories; under such circumstances, these same studies show an increase to an average of two words per fixation. Similarly, readers skip more words if they are speed-reading or skimming (Taylor, 1962). These old results and the current results are consistent with the report of McConkie and Rayner (1975; Rayner, 1978) that readers generally cannot determine the meaning of a word that is in peripheral vision. These results have important implications for the present model; since most words of a text are fixated, we can try to account for the total duration of comprehension in terms of the gaze duration on each word.

The protocol also shows that the gaze duration varies considerably from word to word. There is a misconception that individual fixations are all about 250 msec in duration. But this is not true; there is a large variation in the duration of individual fixations as well as the total gaze duration on individual words. As the preceding display shows, some gaze durations are very long, such as the gaze on the word *Flywheels*. The model proposes that gaze durations reflect the time to execute comprehension processes. In this case the longer fixations are attributed to longer processing caused by the word's infrequency and its thematic importance. Also, the fixations at the end of each sentence tend to be long. For example, this reader had gaze durations of 450 and 566 msec on each of the last words of the first two sentences. The sentence-terminal pauses will be shown to reflect an integrative process that is evoked at the ends of sentences.

The link between eye fixation data and the theory rests on two assumptions. The first, called the immediacy assumption, is that a reader tries to interpret each content word of a text as it is encountered, even at the expense of making guesses that sometimes turn out to be wrong. Interpretation refers to processing at several levels such as encoding the word, choosing one meaning of it, assigning it to its referent, and determining its status in the sentence and in the discourse. The immediacy assumption posits that the interpretations at all levels of processing are not deferred; they occur as soon as possible, a qualification that will be clarified later.

The second assumption, the eye–mind assumption, is that the eye remains fixated on a word as long as the word is being processed. So the time it takes to process a newly fixated word is directly indicated by the gaze duration. Of course, comprehending that word often involves the use of information from preceding parts of the text, without any backward fixations. So the concepts corresponding to two different words may be compared to each other, for example, whereas only the more recently en-

countered word is fixated. The eye–mind assumption can be contrasted with an alternative view that data acquired from several successive eye fixations are internally buffered before being semantically processed (Bouma & deVoogd, 1974). This alternative view was proposed to explain a reading task in which the phrases of a text were successively presented in the same location. However, the situation was unusual in two ways. First, there were no eye movements involved, so the normal reading processes may not have been used. Second, and more telling, readers could not perform a simple comprehension test after seeing the text this way. By contrast, several studies of more natural situations support the eye–mind assumption that readers pause on words that require more processing (Just & Carpenter, 1978; Carpenter & Daneman, Note 1). The eye–mind assumption posits that there is no appreciable lag between what is being fixated and what is being processed. This assumption has also been explored in spatial problem-solving tasks and has been supported in that domain as well as in reading (Just & Carpenter, 1976). The immediacy and eye–mind assumptions are used to interpret gaze duration data in the development of the reading model.

The article has four major sections. The first briefly describes a theoretical framework for the processes and structures in reading. The second section describes the reading task and eye fixation results accounted for by the model. The third section presents the model itself, with subsections describing each component process of the model. The fourth section discusses some implications of the theory for language comprehension and relates this theory of reading to other approaches.

Theoretical Framework

Reading can be construed as the coordinated execution of a number of processing stages such as word encoding, lexical access, assigning semantic roles, and relating the information in a given sentence to previous sentences and previous knowledge. Some of the major stages of the proposed model are depicted schematically in Figure 1. The diagram depicts both processes and structures. The stages of reading in the left-hand column are shown in their usual sequence of execution. The long-term memory on the right-hand side is the storehouse of knowledge, including the procedural knowledge used in executing the

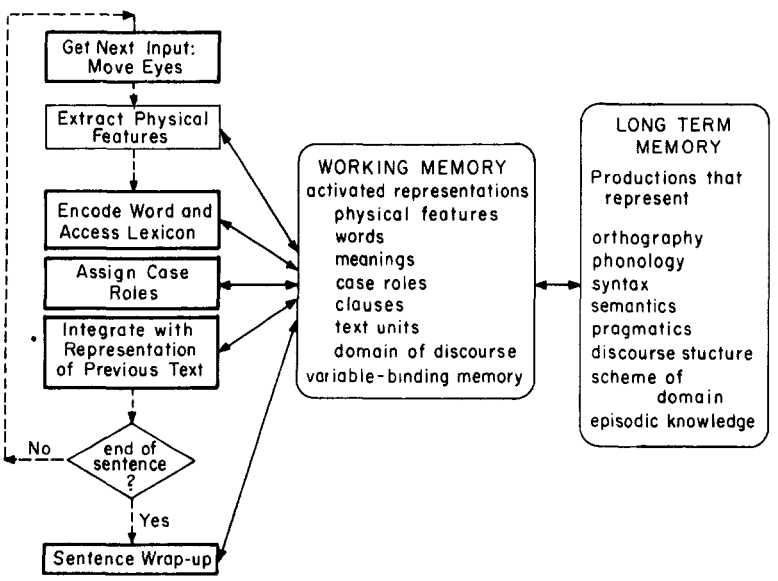

Figure 1. A schematic diagram of the major processes and structures in reading comprehension. (Solid lines denote data-flow paths, and dashed lines indicate canonical flow of control.)

stages on the left. The working memory in the middle mediates the long-term memory and the comprehension processes. Although it is easy to informally agree on the general involvement of these processes in reading, it is more difficult to specify the characteristics of the processes, their interrelations, and their effects on reading performance.

The nature of comprehension processes depends on a larger issue, namely the architecture of the processing system in which they are embedded. Although the human architecture is very far from being known, production systems have been suggested as a possible framework because they have several properties that might plausibly be shared by the human system. Detailed discussions of production systems as models of the human architecture are presented elsewhere (Anderson, 1976; Newell, 1973, 1980). The following three major properties are of particular relevance here.

1. Structural and procedural knowledge is stored in the form of condition–action rules, such that a given stimulus condition produces a given action. The productions "fire" one after the other (serially), and it is this serial processing that consumes time in comprehension and other forms of thought. In addition to the serial productions, there are also fast, automatic productions that produce spreading activation among associated concepts (Anderson, 1976; Collins & Loftus, 1975). These automatic productions operate in parallel to the serial productions and in parallel to each other (Newell, 1980). These productions are fast and automatic because they operate only on constants; that is, they directly associate an action with a particular condition (such as activating the concept *dog* on detecting *cat*). By contrast, serial productions are slow because they operate on variables as well as constants; they associate an action with a class of conditions. A serial production can fire only after the particular condition instance is bound to the variable specified in the production. It may be the binding of variables that consumes time and capacity (Newell, 1980). This architectural feature of two kinds of productions permits serial comprehension processes to operate in the foreground, whereas in the background, automatic productions activate relevant semantic and episodic knowledge.

2. Productions operate on the symbols in a limited-capacity working memory. The symbols are the activated concepts that are the inputs and outputs of productions. Items are inserted into working memory as a result of being encoded from the text or being inserted by a production. Retrieval from long-term memory occurs when a production fires and activates a concept, causing it to be inserted into working memory. Long-term memory is a collection of productions that are the repositories of both procedural and declarative knowledge. In the case of reading, this knowledge includes orthography, phonology, syntax, and semantics of the language, as well as schemas for particular topics and discourse types (Schank & Abelson, 1977). A new knowledge structure is acquired in long-term memory if a new production is created to encode that structure (Newell, 1980). This occurs if the structure participates in a large number of processing episodes.

One important property of working memory is that its capacity is limited, so that information is sometimes lost. One way in which capacity can be exceeded (causing forgetting) is that the level of activation of an item may decay to some subthreshold level through disuse over time (Collins & Loftus, 1975; Hitch, 1978; Reitman, 1974). A second forgetting mechanism allows for processes and structures to displace each other, within some limits (Case, 1978). Heavy processing requirements in a given task may decrease the amount of information that can be maintained, perhaps by generating too many competing structures or by actively inhibiting the maintenance of preceding information. There is recent evidence to suggest that working memory capacity (as opposed to passive memory span) is strongly correlated with individual differences in reading comprehension performance, presumably because readers with greater capacity can integrate more elements of the text at a given time (Daneman & Carpenter, in press).

3. Production systems have a mechanism for adaptive sequencing of processes. The items in working memory at a given time

enable a given production to fire and insert new items, which in turn enable another production, and so on. In this way, the intermediate results of the comprehension process that are placed in working memory can influence or sequence subsequent processing. There is no need for a superordinate controlling program to sequence the mental actions.

The self-sequencing nature of productions is compatible with the model depicted in Figure 1. The composition of each stage is simply a collection of productions that share a common higher level goal. The productions within a stage have similar enabling conditions and produce actions that serve as conditions for other productions in the same stage. The productions within a stage need not be bound to each other in any other way. Thus the ordering of stages with a production system is accomplished not by direct control transfer mechanisms but an indirect self-sequencing accomplished by one production helping to create the conditions that enable the "next" production to fire.

This architecture permits stages to be executed not only in canonical orders but also in noncanonical orders. There are occasions when some stages of reading seem to be partially or entirely skipped; some stages seem to be executed out of sequence, and some "later" stages sometimes seem to be able to influence "earlier" stages (Levy, in press). Stages can be executed earlier than normal if their enabling conditions exist earlier than normal. For example, if a context strongly primes a case role, then the case assignment could precede the lexical access of a word. Having read *John pounded the nail with a* _____, a reader can assign the last word to the instrumental case on the basis of cues provided by the words *pound* and *nail*, before encoding *hammer*. This organization can permit "context effects" in comprehension, where a strong preceding context shortens reading time on a given word or clause. This might occur if a processing stage that is normally intermediate between two others is partially or entirely eliminated. It could be eliminated if the preceding stage plus the context provided sufficient enabling conditions for the later stage. Analogously, a misleading context could lengthen comprehension time by providing elements that enable conflicting processes.

The production system organization can also explain how "later" stages can influence "earlier" stages, so that higher level schemas can affect word encoding, for example. If the productions of the normally later stage are enabled earlier than usual, then their outputs can serve as inputs to the normally earlier stage. The ordering of stages does not have to be entirely reversed to obtain this top-down influence. It may be sufficient for just a portion of the productions of the "later" stage to fire in order to influence the "earlier" stage.

In this view of processing stages, several stages can be executed cotemporaneously in the sense that firings of productions of two or more stages may be interleaved. Consequently, data and control can be transferred back and forth among different stages, somewhat similarly to computer programs organized into coroutines. Coroutines are two or more subprograms that have equal status (i.e., there is no master–slave relationship). When one coroutine obtains control, it executes until it detects a condition indicating it should relinquish control, and then another coroutine executes, and so on. One interesting difference between coroutines and the production system model is that coroutines generally transfer data between each other only along specified paths, used especially for this purpose. By contrast, productions "transfer" data by placing it in the working memory, so that all processes have access to it. In this sense, the working memory serves as a message center, and communication among stages is by means of the items in working memory. This is distinct from one stage feeding its output directly to another stage.

Research

Texts

This section describes the texts that were used in the reading research because their properties, both local and global, have a large influence on the processing. The global organization of a narrative text has been shown to influence how a reader recalls

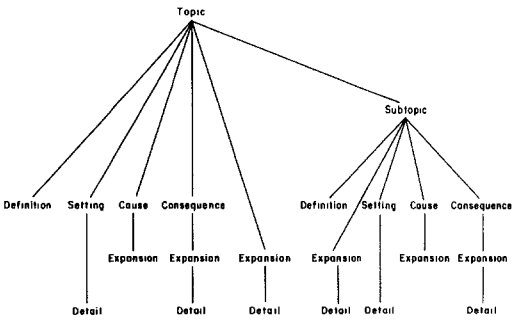

Figure 2. A schematic diagram of the major text-grammatical categories of information in the scientific paragraphs.

the text (Kintsch & van Dijk, 1978; Mandler & Johnson, 1977; Meyer, 1975; Rumelhart, 1977b; Thorndyke, 1977). The experiment reported next shows that the organization has at least part of its effect when the text is being read. Scientific texts were selected from *Newsweek* and *Time* because their content and style is typical of what students read to learn about technical topics. The passages discussed a variety of topics that were generally unfamiliar to the readers in the study. When readers were asked to rate their familiarity with the topic of each passage on a 5-point scale, the modal rating was at the "entirely unfamiliar" end of the scale. There were 15 passages, averaging 132 words each. Although the texts are moderately well written, they are on the borderline between "fairly difficult" and "difficult" on Flesch's (1951) readability scale, with 17 words per sentence and 1.6 syllables per word. The following is an example of one of the passages:

Flywheels are one of the oldest mechanical devices known to man. Every internal-combustion engine contains a small flywheel that converts the jerky motion of the pistons into the smooth flow of energy that powers the drive shaft. The greater the mass of a flywheel and the faster it spins, the more energy can be stored in it. But its maximum spinning speed is limited by the strength of the material it is made from. If it spins too fast for its mass, any flywheel will fly apart. One type of flywheel consists of round sandwiches of fiberglass and rubber providing the maximum possible storage of energy when the wheel is confined in a small space as in an automobile. Another type, the "superflywheel," consists of a series of rimless spokes. This flywheel stores the maximum energy when space is unlimited.

The content of the passages was analyzed by segmenting the text into idea units and categorizing these units by means of a simple text grammar. First, all of the 15 passages were segmented into text units called *sectors*, producing 274 sectors. The average sector length was seven words. Each sector was judged to be a single meaningful piece of information, whether it consisted of a word, phrase, clause, or sentence. The general criteria for segmentation into sectors were similar to those used by Meyer and McConkie (1973), who related such text units to recall performance.

A simplified grammar was developed to categorize the sectors of the texts. The grammar (shown schematically in Figure 2) classifies the text units into a structure that is quasi-hierarchical. This abbreviated grammar captures most of the regularities in our short passages (see Vesonder, 1979, for a more complete grammar for longer scientific passages). The initial sentences generally introduced a topic—a scientific development or event. The beginnings of the passage sometimes gave details of the time, place, and people involved with the discovery. Familiar concepts were simply named, whereas unusual concepts were accompanied by an explicit definition. The main topic itself was developed through specific examples or through subtopics that were then expanded with further descriptions, explanations, and concrete examples. Consequences, usually toward the end of the passage, stated the importance of the event for other applications. Table 1 shows how each text unit or sector in the "Flywheel" passage was classified according to these categories. Each of the 274 sectors was assigned to one of the five levels of the grammar by one of the authors. The levels of the grammar were further confirmed by a pretest involving 16 subjects who rated the importance of each sector in its passage on a 7-point scale. The mean importance ratings differed reliably among the five levels $F(4, 270) = 40.04$, $p < .01$. Specifically, the means decreased monotonically through the five postulated levels. Hence, the grammar potentially has some psychological reality, and its relevance to

Table 1
A Classification of the "Flywheel" Passage Into Text-Grammatical Categories

Category	Sector
Topic	Flywheels are one of the oldest mechanical devices
Topic	known to man
Expansion	Every internal-combustion engine contains a small flywheel
Expansion	that converts the jerky motion of the pistons into the smooth flow of energy
Expansion	that powers the drive shaft
Cause	The greater the mass of a flywheel and the faster it spins,
Consequence	the more energy can be stored in it.
Subtopic	But its maximum spinning speed is limited by the strength of the material
Subtopic	it is made from.
Expansion	If it spins too fast for its mass,
Expansion	any flywheel will fly apart.
Definition	One type of flywheel consists of round sandwiches of fiberglas and rubber
Expansion	providing the maximum possible storage of energy
Expansion	when the wheel is confined in a small space
Detail	as in an automobile.
Definition	Another type, the "superflywheel," consists of a series of rimless spokes.
Expansion	This flywheel stores the maximum energy
Detail	when space is unlimited.

reading will be demonstrated with the eye fixation data. The next section presents the data collection and analysis procedures, followed by the model and results.

Method and Data Analysis

The readers were 14 undergraduates who read 2 practice texts followed by the 15 scientific texts in random order. Although the readers were asked to recall each passage immediately after reading it, they also were told to read naturally without memorizing. They were also asked not to reread the passage or parts of it. The texts were presented on a television monitor using uppercase and lowercase letters and a conventional paragraph layout. To initiate the reading of a passage, the reader had to look at a fixation point (located where the first word of the paragraph would later appear) and press a "ready" button. If the reader's point of regard (as measured by the eye tracker) was within 1° of the fixation point, then exactly 500 msec later the passage appeared in its entirety on the screen. The passage appeared instantaneously (i.e., within one video frame) and remained there until the reader signaled that he had finished reading by pushing a response button.

The reader's pupil and corneal reflections were monitored relatively unobtrusively by a television camera that was 75 cm away. The monitoring system, manufactured by Applied Science Laboratories, computed the reader's point of regard (as opposed to eye or head position) every 16.7 msec. The accuracy of the tracker was verified before and after each passage was read by having the reader look at a fixation point and determining whether the obtained point of regard was within 1° of that point. This procedure indicated that accuracy was maintained during the reading of 195 of the 210 experimental passages in the entire experiment; the data from the 15 inaccurate trials were discarded.

Data reduction procedures converted the 60 observations per sec into fixations and then into gazes on each word. While the data were being acquired, a new "fixation" was scored as having occurred if the point of regard changed by more than 1° (the size of a three-letter syllable). The durations of blinks that were preceded and followed by fixations on the same location were attributed to the reading time on that location. Another program aggregated consecutive fixations on the same word into gazes and computed the duration of gaze on each of the 1,936 words in the 15 passages. Fixations on interword spaces were attributed to the word on the right because the perceptual span is centered to the right of the point of regard, at least for readers of left-to-right languages (McConkie & Rayner, 1976; Schiepers, 1980). The durations of saccades, blinks that occurred between words, regressions, and rereading were not included in the data analysis. Because of the instructions not to reread, these categories account for relatively little of the total reading time, approximately 12% in all. The mean duration of gaze on each word was computed by averaging over readers; these 1,936 mean gaze durations constitute the main dependent measure of interest.

The model presents a number of factors that influence various reading processes; some factors have their effect on individual words and some on larger units, such as clauses. The data were fit to the model with a multiple linear regression in which the independent variables were the factors postulated to affect reading time and the dependent variable was the mean gaze duration on each word. Since the model also applies at the level of clauses and phrases, a

second regression analysis was done at the phrase/clause level. The independent variables for the latter analysis were the factors postulated to affect reading time at the clause level, and the independent variable was the mean gaze duration on each of the 274 sectors described previously.

The psychological interpretation of the independent variables in the two regression analyses will be described in detail in the sections that follow. The equation for the analysis of the gaze duration on individual words was

$$GW_i = \sum a_m X_{im} + \epsilon_i,$$

where GW_i is the gaze duration on a word i, a_m is the regression weight in msec for independent variable X_m, X_{im} are the independent variables that code the following seven properties of word i:

(a) length, (b) the logarithm of its normative frequency, (c) whether the word occurs at the beginning of a line of text, (d) whether it is a novel word to the reader, (e) its case grammatical role (one of 11 possibilities), (f) whether it is the last word in a sentence, (g) whether it is the last word in a paragraph.

The equation for the analysis of the gaze duration on individual sectors was

$$GS_j = b_0 + \sum b_n Z_{jn} + \epsilon_j$$

where GS_j is the gaze duration on sector j, and b_n is the regression weight in msec for independent variable Z_n. The Z_{jn} are the independent variables that code the following eight properties of sector j:

(a) its text grammatical level, multiplied by the number of content words; (b) length; (c) the sum of the logarithms of the frequencies of its component words; (d) the number of line-initial words it contains; (e) the number of novel words it contains; (f) the sum of the case role regression weights of its component words; (g) whether it is the last sector in a sentence; (h) whether it is the last sector in a paragraph.

Results

The mean gaze duration on each word (239 msec) indicated reading rates that are typical for texts of this difficulty. If the 239 msec per word is incremented by 12% to allow for saccades, blinks, and occasional rereading, the reading rate is 225 words per min. The standard deviation of the 239-msec gaze mean was 168 msec, indicating considerable variability in gaze duration from word to word. The results of the regression analyses are shown in Table 2. The table is divided into three sections, corresponding to the three major processing stages postulated by the model, encoding and lexical access, case role assignment, and interclause integration. The regression weights shown in Table 2 for the word-by-word analysis (above the double line) are derived from a regression equation involving 17 independent variables (11 of which are the case role indicator variables). The standard error of estimate of this model was 88 msec, and the R^2 value was .72. The results of the interclause integration stage make use of both the word-by-word analysis and the sector-by-sector analysis. (The latter analysis will be explained in more detail in the section on interclause integration). Since the gaze durations on successive words and phrases are time-series data, it is interesting to note that there was no reliable positive serial correlation among the residuals in the word-by-word regression or the sector-by-sector regression.

The Reading Model

The next five subsections describe the major stages shown in Figure 1: get next input, encoding and lexical access, case role assignment, interclause integration, and sentence wrap-up. Each subsection describes the processes in that stage together with the factors that affect the duration of those processes, and hence the gaze durations.

Get Next Input

This is the first stage of a cycle that finds information, encodes it, and processes it. When the perceptual and semantic stages have done all of the requisite processing on a particular word, the eye is directed to land in a new place where it continues to rest until the requisite processing is done, and so forth. The specification of what constitutes "all of the requisite processing" is contained in a list of conditions that must be satisfied before the reader terminates the gaze on the current word and fixates the next one.These conditions include a specification of the goals of normal reading. For instance, one condition may be that a meaning of the word be accessed and another condition may be that a case role be assigned. These conditions can also reflect more specific reading goals. A reader who is trying to memorize a text may have as a condition that the word or phrase be transferred to long-term memory. By setting the conditions appropriately, the reader can

adjust his processes to the situation at hand. When the goal conditions for processing a word are satisfied, the resulting action is get next input.

The command to get next input usually results in a saccade to the next part of the text, one or two words forward. The process that selects the placement of the next forward fixation does not have to be very complex or intelligent. The choice of where to place the next forward fixation appears to depend primarily on the length of the next word or two to the right of the current fixation (McConkie & Rayner, 1975). The length information, which is encoded parafoveally, is then used to program a rightward saccade. However, if only the right margin is visible in the parafovea, then the eye is directed to the first word of the next line, producing a return sweep. In this case the information in peripheral vision is not adequate for accurate targeting. The return sweep is typically too short; the eye often lands on the second word of the new line for a brief amount of time (50 or 75 msec) and then makes a corrective saccade leftward to the first word of the line (Bayle, 1942). On occasion, a comprehension stage may require a review of previously read text to reencode it or process it to deeper levels. In those cases, the get next input stage results in a regressive saccade to the relevant portion of the text.

The duration of the get next input stage is short, consisting of the time for a neural signal to be transmitted to the eye muscles. In monkeys, this takes about 30 msec (Robinson, 1972). This duration must not be

Table 2
Application of the Regression Model to the Gaze Duration on Each Word (Above Double Line) and to Each Sector (Below Double Line)

Processing stage	Factor	Regression weight (msec)
Encoding and lexical access	no. of syllables	52**
	log frequency	53**
	beginning of line	30**
	novel word	802**
Case role assignment	agent (86)	51**
	instrument (110)	53**
	direct or indirect object (174)	25*
	adverb/manner (35)	29
	place or time (64)	23
	possessive (genitive) (39)	16
	verb (368)	33**
	state/adjective (451)	44**
	rhetorical word (15)	70**
	determiner (243)	26**
	connective (351)	9
Interclause integration	last word in sentence	71**
	last word in paragraph	157**
	Integration time per content word from regression analysis of data aggregated into sectors	
	topic (22)	72**
	definition/cause/consequence (23)	94**
	subtopic (48)	78**
	expansion (68)	73**
	detail (113)	60**

Note. Frequency of occurrence of case roles is in parentheses.
* $t = p < 0.05$; ** $t = p < .01$.

confused with the typical 150- to 200-msec latency of a saccade to a visual stimulus that has spatial or temporal uncertainty (Westheimer, 1954). These latencies include stimulus detection, interpretation, and selection of the next fixation target. In normal reading, there is very little uncertainty about direction of the next saccade (it is almost always rightward for forward fixations, except for the return sweeps), nor is there much uncertainty about distance. On the average, the saccade distance may be simply the mean center-to-center distance between words, a distance that does not vary much, relative to the physically possible variation in eye movements. Thus it is reasonable to suppose that the preprogramming time is very short here, consisting usually of a "go" signal and the time it takes that signal to be translated into a motor movement, about 30 msec (Robinson, 1972). The actual movements, the saccades, constitute about 5%–10% of the total reading time. Recent analyses suggest that the saccade itself may destroy the visual persistence of the information from the preceding fixation so that it does not mask the input from the new fixation (Breitmeyer, 1980). Consequently, it is reasonable to assume that stimulus encoding can commence soon after the eye arrives at a new location.

Word Encoding and Lexical Access

The reading process involves encoding a word into an internal semantic format. It is assumed that prior to this encoding, the transduction from the printed word to the visual features has already taken place, and that the features have been deposited into the working memory. Perceptual encoding productions use the visual features as conditions; their action is to activate the representation of the word. Once the representation of the word has been sufficiently activated, its corresponding concept is accessed and inserted into working memory. The concept serves as a pointer to a more complete representation of the meaning, which consists of a small semantic network realized as a set of productions. The major nodes of the network are the possible meanings of the word, the semantic and syntactic properties of the meanings, and information about the contexts in which they usually occur (see Rieger, 1979, for a related proposal).The word meanings are represented as abstract predicates, defined by their relations to other predicates.

The productions that encode a word generally trigger on orthographically based subword units such as syllables (Mewhort & Beal, 1977; Spoehr & Smith, 1973; Taft, 1979). However, there are times when alternative codes, including orthographic, phonological, and whole-word codes, are used (Baron, 1977; Kleiman, 1975; LaBerge & Samuels, 1974). Since the syllablelike encoding is believed to be the dominant mode, the data were analyzed in terms of the number of syllables in each word. Encoding time increased by 52 msec for each syllable, as shown in Table 2.

The mechanism underlying lexical access is the activation of a word's meaning representation by various sources. There are three ways that a concept's level of activation can be temporarily increased above its base level. One activation mechanism is perceptual encoding; the encoded representation of a word can activate its meaning. A second source is the parallel productions that produce spreading activation through the semantic and episodic knowledge base of the reader. The third source is activation by the serial productions that do the major computations in all of the stages of processing. When a concept has been activated above some threshold by one or more of these sources, a pointer to its meaning is inserted into working memory. The activation level gradually decays to a subthreshold level unless some process reactivates it. If the word soon reoccurs in the text while the concept is still activated, lexical access will be facilitated because the activation level will still be close to threshold. When the activation level does decrease, it decreases to an asymptote slightly higher than the old base level. In this way, the system can learn from both local and long-term word repetitions. Frequently used words will have a high base level of activation, and consequently will require relatively less additional activation to retrieve them. Thus, frequent words should

take less time to access than infrequent words (Morton, 1969). Similarly, the various possible interpretations of each word will have different base activation levels, such that the more common interpretations have higher base activation levels. For example, although the word *does* has at least two very different meanings, the "third-person-singular verb" interpretation would have a higher base activation because it is more common than the "female deer" interpretation (Carpenter & Daneman, Note 1). The more common interpretation would then be accessed faster, since less additional activation would be required to bring the activation level to threshold. This model of lexical access can account for word frequency effects, priming effects, and repetition effects in reading.

The gaze duration showed both frequency and repetition effects. Frequency was analyzed by relating gaze duration to the logarithm of the normative frequency of each word, based on the Kučera and Francis (1967) norms. It was expected that gaze duration would decrease with the logarithm of the word's frequency; that is, small differences among infrequent words would be as important as much larger differences among frequent words (Mitchell & Green, 1978). For algebraic convenience, the normative frequencies were increased by one (to eliminate the problem of taking the logarithm of zero), and the logarithm was computed and then subtracted from 4.85, the logarithm of the frequency of the most frequent English word. The analysis indicated a clear relation between this measure of frequency and gaze duration. As shown in Table 2, gaze duration increased by 53 msec for each log unit of decrease in word frequency. A moderately frequent word like *water* (with a frequency of 442) was accessed 140 msec faster than a word that did not appear in the norms.

At one extreme of the frequency dimension are words that a reader has never encountered before. In scientific passages, the novel words tend to be technical terms. To read these words, a reader cannot depend on contacting some prior perceptual and semantic representation; neither exists. The reader must construct some perceptual representation (perhaps phonological as well as orthographic), associate this with the semantic and syntactic properties of the concept that can be inferred from the passage, and then possibly construct a lexical entry. These processes seem to take a great deal more time than ordinary encoding and access processes. Two judges identified seven words in the texts (that had zero frequency) as probably entirely novel to the readers. Novelty was coded as an indicator variable, and it was found that these words took an additional 802 msec on average to process, as shown in Table 2. However, there was considerable variability among the words; their gaze durations ranged from 913 msec (for *staphyllococci*) to 2,431 msec (for *thermoluminescence*).

Once a word has been encoded and accessed once, it should be easier to access it when it occurs again. Other research has suggested that frequency and repetition have their primary effect on lexical access rather than encoding (Dixon & Rothkopf, 1979; Glanzer & Ehrenreich, 1979; Scarborough, Cortese, & Scarborough, 1977), although the possibility of some small effects on the encoding process does exist. According to the model, repetition effects should occur in reading because the first time a word meaning is accessed, it should temporarily achieve a higher activation level similar to the level of a more frequent word. This mechanism particularly predicts repetition effects for infrequent words, whose activation levels are low to start with, but not for the highly frequent words that occur in natural text. Generally, repetition effects are larger for low-frequency words (Scarborough et al., 1977). "Low frequency" in the Scarborough study was defined as less than 28 occurrences per million, the boundary of 28 emerging from a median split of the frequencies of their stimuli. So the analysis of repetition effects was limited to words with frequencies of 25 occurrences per million or less. There were 346 such instances in the text; 251 were initial occurrences and 95 were repetitions. The repetitions were words with the same morphological stem, disregarding affixes. An analysis of covariance on this subset of the data examined the effects of repe-

titions covarying out the number of syllables. The adjusted mean gaze durations were 49 msec longer on the initial appearance of these words than on the subsequent appearances, $t(343) = 2.21$, $p < .03$. Most of this effect (43 msec) was obtained on the second appearance of a word. These results indicate that once an infrequent word appears in a text, processing time on that word is decreased on subsequent appearances.

Lexical access is complicated by the fact that some words have more than one meaning, so the appropriate interpretation must be selected, or at least guessed at. When a polysemous word is accessed, the word representation that is retrieved is a pointer to a semantic network that includes the multiple representations. The interpretation that is selected is the one with the highest activation level, and several factors can affect the activation. First, some interpretations start off with a higher activation level; for instance, the "third-person-singular" interpretation of *does* has a higher base activation level than the "deer" interpretation. Second, the automatic productions that produce spreading activation can contribute selectively to the activation level of one particular interpretation. The spreading activation can emanate from the preceding semantic and syntactic context, from the reader's knowledge of the domain, and from knowledge of the discourse style. Third, the output of other stages operating on the same word may activate a particular interpretation. For example, although *hammer* can be interpreted as a noun or a verb, a sentence context that suggests an instrument to the case role assignment stage (e.g., *John hit the nail with a* _____) may help activate the noun interpretation. Fourth, when a word with many highly related meanings occurs in an impoverished context, there may be no single interpretation with higher activation than the others, and the superordinate concept may be the selected interpretation of the word. This probably occurs for words that have many closely related interpretations, such as *get* and *take*.

The selection of only one interpretation of each word, posited by the immediacy assumption, provides a measure of cognitive economy. Selecting just one interpretation allows the activation of the unselected interpretations to decay, preventing them from activating their associates. Thus, the contextual effects would remain focused in the appropriate semantic domain. This permits a limited-capacity working memory to cope with the information flow in a spreading activation environment that may activate many interpretations and associations for any lexical item. This method of processing also avoids the combinatorial explosion that results from entertaining more than one interpretation for several successive words.

This aspect of the model is consistent with some recent results on lexical access that indicate that although multiple meanings of a word are initially activated, only one meaning remains activated after a few hundred milliseconds. In one experiment, the subjects simultaneously listened to a sentence and pronounced a visually presented word. When an ambiguous word (*rose*) was presented auditorally in a syntactic context (e.g., *They all rose*), the speed of pronouncing a simultaneous visual probe related to either meaning (stood or flower) was faster than in a control condition (Tanenhaus, Leiman, & Seidenberg, 1979). In another experiment, the subjects listened to a sentence and performed a lexical decision task on visually presented stimuli. When an ambiguous word (*bug*) was presented in a semantic context (*John saw several spiders, roaches, and bugs*), the speed of a simultaneous lexical decision related to either meaning (insect or spy) was faster than a control (Swinney, 1979). In both studies, the facilitation of the inappropriate meaning was obtained only within a few hundred milliseconds of the occurrence of the ambiguous word. If the probe was delayed longer, the inappropriate interpretation was no faster than the control. These results suggest that both meanings are available when an ambiguous word is being accessed, but the inappropriate meaning is lost from working memory after a short time.

As the interpretation of the text is constructed, a corresponding representation of the extensive meaning—the things being talked about—is also being built. If the referents of the words in a passage cannot

be determined, the text will be more difficult to understand. One example of this problem is highlighted in a passage from Bransford and Johnson (1973) concerning a procedure that involved arranging "things into groups. Of course, one may be sufficient depending on how much there is. . . ." (p. 400). Subjects who were not given the title *washing clothes* though the story was incomprehensible. The referential representation helps the reader disambiguate referents, infer relations, and integrate the text.

The immediacy assumption posits that an attempt to relate each content word to its referent occurs as soon as possible. Sometimes this can be done when the word is first fixated, but sometimes more information is required. For example, although the semantic interpretation of a relative adjective like *large* can be computed immediately, the extensive meaning depends on the word it modifies (e.g., *large insect* vs. *large building*). The referent of the entire noun phrase can be computed only after both words are processed. The immediacy assumption does not state that the relating is done immediately on each content word, but rather that it occurs as soon as possible. This is an important distinction that will be made again in the discussion on integrative processes.

Assigning Case Roles

Comprehension involves determining the relations among words, the relations among clauses, and the relations among whole units of text. This section describes the first of these processes, that of determining the relations among the words in a clause (or in Schank's, 1972, terms, determining the dependencies among the concepts). These relations can be categorized into semantic cases, such as agent, recipient, location, time, manner, instrument, action, or state (Chafe, 1970; Fillmore, 1968). The case role assignment process usually takes as input a representation of the fixated word, including information about its possible case roles and syntactic properties. For example, hammers tend to be instruments rather than locations or recipients, and information about a word's usual case role can be an important contributor to the assignment process. But this normative information generally is not sufficient to assign its case role in a particular clause. Consequently, the assignment process relies on heuristics that use the word meaning together with information about the prior semantic and syntactic context, as well as language-based inferences. The output of the process is a representation of the word's semantic role with respect to the other constituents in its clause.

Just as certain meanings suggest particular case roles, so, too, can the context prime a particular case role. Consider the sentence *John was interrogated by the* _____. The semantic and syntactic cues suggest that the missing word will be an agent, such as *detective*. The strength of the context becomes evident if the primed case does not occur, for example, *John was interrogated by the window*. The prior semantic context can precede the affected case assignment by more than a few words. In the sentences *The lawyer wanted to know where in the room John had been interrogated* and *Mary told him that John was interrogated by the window*, the thematic focus of the first sentence on a location alters the interpretation of *by* and facilitates a locative case role assignment for *window*.

The specific heuristics that are used in case role assignment have received some attention (see Clark & Clark, 1977, for some examples). Many proposals contain the suggestion that readers use the verb as a pivotal source of information to establish the necessary and possible case roles and then fit the noun phrases into those slots (Schank, 1972). But the immediacy assumption posits that the case role assignment for an item preceding the verb is not postponed in anticipation of the verb. Similar to the lexical access stage, the case assignment stage makes a best guess about a word's case when the word is fixated, rather than making the decision contingent on subsequent words. So, the model would not accord any special status to verbs. Another suggested heuristic (that children appear to use) is to assign a sequence consisting of animate noun–verb–noun to the case roles of agent–action–object (Bever,

1970). Like all heuristics, this one sometimes fails, so young children sometimes misinterpret passive sentences (Fraser, Bellugi, & Brown, 1963). This heuristic may be employed by adults, but in a modified version that conforms to the immediacy assumption. Rather than waiting for the three major constituents before assigning case roles, the reader should assign an animate noun to the agent role as soon as it is encountered, in the absence of contrary prior context.

The immediate assignment of a case role implies that readers will sometimes make errors and have to revise previous decisions. For example, an adult who assigns the role of agent to an animate noun and then encounters a passive verb will have to revise the agent assignment. (Presumably, young children do not make this revision.) The immediacy of the case assignment process is evident in the reading of sentences such as *Mary loves Jonathan. . . .* The immediacy assumption suggests that a reader would assign to *Jonathan* the role of recipient; this would in turn result in an incorrect assignment if the sentence continued *Mary loves Jonathan apples.*

Because case roles are assigned within clauses, the assignment process must include a segmentation procedure to determine clause boundaries within sentences. Sentences can sometimes be segmented into clauses on the basis of explicit markers, such as a subordinating conjunction (e.g., *because*, *when*). More often, the reader cannot tell with certainty where one clause ends and another starts until beyond the clause boundary (or potential boundary). A general strategy for dealing with such cases has been suggested, namely to assign a word to the clause being processed, if possible (Frazier & Fodor, 1978). For example, the word *soil* in the sentence *When farmers are plowing the soil . . .* can continue the initial clause (*When farmers are plowing the soil, it is most fertile*) or start a new one (*When farmers are plowing the soil is most fertile*). The suggested strategy is to continue the initial clause until contrary information is encountered. Interestingly, the strategy discussed by Frazier and Fodor (1978) presupposes the immediacy assumption; the segmentation decision arises because case roles are assigned as soon as the words are encountered.

There is no direct mapping between particular case roles and the duration of the assignment process. For example, there is no a priori reason to expect that assignment of instruments takes more or less time than locations. The time for a particular assignment might depend more on the context and properties of the word than on the particular case role being assigned. Detailed specification of the process is not within the scope of this article; it probably requires a large-scale simulation model to examine the complex interactions of different levels of processing. Nevertheless, we examined whether, all things being equal, different case role assignments tend to take different amounts of time.

The analysis included the usual case roles just noted (Fillmore, 1968), as well as other categories such as determiners and adjectives that are not cases but still play a part in the parsing and assignment process. Each word was classified into 1 of 11 categories: verb, agent, instrument, indirect or direct object, location or time, adverb, adjective or state, connective (preposition or conjunction), possessive, determiner, and rhetorical word (such as *well*). Some cases were pooled (such as location and time) because they were relatively infrequent in the text and because they have some conceptual similarity. The case roles were coded as indicator variables and were all entered into the regression with the intercept forced to zero.

The results of the case role assignment analysis, shown in Table 2, indicate that there are some variations among the cases. As expected, verbs did not take particularly long (33 msec), and in fact, although the time was significantly different from 0, it was not greater than the agent or instrument cases (51 msec and 53 msec, respectively). Four cases had parameters that were not significantly different from 0, connectives (9 msec), adverb/manner (29 msec), place or time (23 msec), and possessives (16 msec). These parameters could reflect some properties of particular word classes, in addition to parsing and case role assignment processes. For example, if a connective (e.g., *and* or *but*) simply takes less

time to access than other words, the advantage should appear in this parameter. However, the parameters are not due solely to length or frequency, since these variables make a separate contribution to the regression equation. Although this analysis does not examine any of the contextual effects thought to be of some importance in the case assignment process, it does indicate roughly the relative amount of time spent assigning various categories of words to their case roles in a clause. Later theories will have to account for the precise pattern of case assignment durations in terms of specific operations that use prior context and word meanings to assign the various cases.

Interclause Integration

Clauses and sentences must be related to each other by the reader to capture the coherence in the text. As each new clause or sentence is encountered, it must be integrated with the previous information acquired from the text or with the knowledge retrieved from the reader's long-term memory. Integrating the new sentence with the old information consists of representing the relations between the new and the old structures.

Several search strategies may be used to locate old information that is related to the new information. One strategy is to check if the new information is related to the other information that is already in working memory either because it has been repeatedly referred to or because it is recent (Carpenter & Just, 1977a; Kintsch & van Dijk, 1978). Using this strategy implies that adjacency between clauses and sentences will cause a search for a possible relation. For instance, the adjacent sentences *Mary hurt herself* and *John laughed* seem related (John must be a cad) even though there is no explicit mention of the relation. This strategy also entails trying to relate new information to a topic that is active in working memory. This is a good strategy, since information in a passage should be related to the topic.

A second strategy is to search for specific connections based on cues in the new sentence itself. Sentences often contain old information as well as new. Sometimes the old information is explicitly marked (as in cleft constructions and relative clauses), but often it is simply some argument repeated from the prior text. The reader can use this old information to search his or her long-term text representation and referential representation for potential points of attachment between the new information and the old (Haviland & Clark, 1974). This second strategy may take more time than the first. In fact, it takes longer to read a sentence that refers to information introduced several sentences earlier than one that refers to recently introduced information (Carpenter & Just, 1977a).

There are two main points at which integration can occur. First, as each ensuing word of the text is encountered, there is an attempt to relate it to previous information (Just & Carpenter, 1978). Second, a running representation of the clause is maintained, with an updating as each word of the clause is read. This running clause representation consists of the configuration of clause elements arranged according to their case relations. This second type of integration involves an attempt to relate the running clause representation to previous information at each update. Integration occurs whenever a linking relation can be computed. Consider the sentence *Although he spoke softly, yesterday's speaker could hear the little boy's question.* The point of this example is not so much that the initial integration of *he* and *speaker* is incorrect, but that the integration is attempted at the earliest opportunity. This model implies that integration time may be distributed over fixations on different parts of a clause. Moreover, the duration of the process may depend on the number of concepts in the clause; as these increase, the number of potential points of contact between the new clause and previous information will increase. There is also evidence for integration triggered by the end of the sentence; this process is discussed next in more detail.

Integration results in the creation of a new structure. The symbol representing that structure is a pointer to the integrated concepts, and this superordinate symbol is then available for further processing. In this

way, integration can chunk the incoming text and allows a limited working memory to deal with large segments of prose. The macrorules proposed by Kintsch and van Dijk (1978) can be construed as productions that integrate.

Integration can also lead to forgetting in working memory. As each new chunk is formed, there is a possibility that it will displace some previous information from working memory. Particularly vulnerable are items that are only marginally activated, usually because they were processed much earlier and have not recently participated in a production. For instance, the representation of a clause will decay if it was processed early in a text and was not related to subsequent information. This mechanism can also clear working memory of "lower level" representations that are no longer necessary. For example, the verbatim representation of a previously read sentence may be displaced by the processes that integrated the sentence with other information (Jarvella, 1971). By contrast, the semantic elements that participate in an integration production obtain an increased activation level. This increases the probability that they will become a permanent part of long-term memory.

The main types of interclause relations in the scientific passages correspond to the text-grammatical categories described previously, such as definitions, causes, consequences, examples, and so forth. Text roles that are usually more important to the text and to the reader's goals, such as topics or definitions, are integrated differently than less important units, such as details. The more central units will initiate more retrievals of relevant previous knowledge of the domain (schematic knowledge) and retrievals of information acquired from the text but no longer resident in the working memory. In addition, more relations will be computed between the semantically central propositions and previous information because centrality inherently entails relations with many other units. By contrast, details are often less important to the reader's goals and to the text. Moreover, when a detail is to be integrated, the process is simpler because details are often concrete instantiations of an immediately preceding statement (at least in these scientific texts), so they can be quickly appended to information still present in the working memory. Thus, higher level units will take more time to integrate because their integration is usually essential to the reader's goals, and because integration of higher units involves more relations to be computed and more retrievals to be made.

The nature of the link relating two structures may be explicitly denoted either in the text (with connectives like *because*, *therefore*, and *for instance*) or it may have to be inferred on the basis of schematic knowledge of the domain. For example, the causal relation between the sentences *Cynthia fell off the rocking horse* and *She cried bitter tears* is inferred from the reader's knowledge about the temporal and causal relation between falling and hurting oneself (Charniak, Note 2).

The model predicts that the gaze duration on a sector depends on its text-grammatical role and on the number of concepts it contains. Because integration can occur at many points in a sector, the gaze duration associated with integration cannot be localized to a particular word. Thus, to do the clause level of analysis, the gaze durations on the individual words of a sector were cumulated, producing a total of 274 sector gaze durations as the dependent variable. The independent variables were the aggregates of the word-level variables, except for case roles. The independent variable that coded the case-role effect for a sector was the sum of the case-coefficients (obtained from the word-by-word regression analysis) for each of the words in the sector. A new independent variable coded the text-grammatical role of a sector and its number of content words; it was the interaction of the indicator variables that represented the five text-grammatical levels and the number of content words in the sector, with content words defined as in Hockett (1958).

The results indicate that the integration time for a given sector depends on its text-grammatical role. The portion of Table 2 below the double line shows the integration time per content word for each type of sec-

tor. Generally, more important or central sectors take longer to integrate. The model describes this effect in terms of the integrative processes initiated by the semantics of the different types of information and their relevance to the reader's goals. An analysis of covariance examined the effect of text roles covarying out the number of syllables. The adjusted mean gaze durations differed reliably, $F(4, 268) = 8.82, p < .01$; paired comparisons indicated that details took significantly less time than all other roles, and expansions took significantly less time than topics and definitions/causes/consequences (all ps $< .01$). These results quantitatively and qualitatively replicate those reported previously for a slightly different paradigm (Carpenter & Just, in press). The previously obtained coefficients for the five text-grammatical categories were 65, 106, 81, 76, and 47 msec per content word, respectively, corresponding to the newly obtained 72, 94, 78, 73, and 60. The model accounts very well for the sector-level data. The R^2 value was .94, and the standard error of estimate was 234. The mean gaze duration on a sector was 1,690 msec, with a standard deviation of 902 msec, and the mean sector length was 4.9 words.[1]

One cost of immediate interpretation, case role assignment, and integration is that some decisions will prove to be incorrect. There must be mechanisms to detect and recover from such errors. The detection of a misinterpretation often occurs when new information to be integrated is inconsistent with previous information. Thus, misinterpretation detection may be construed as inconsistency detection. For example, the sentence *There were tears in her brown dress* causes errors initially because the most frequent interpretation of *tears* is not the appropriate one here, and the initial interpretation is incompatible with *dress*. The eye fixations of subjects reading such garden path sentences clearly indicate that readers do detect inconsistencies, typically at the point at which the inconsistency is first evident (Carpenter & Daneman, Note 1). At that point, they use a number of error-recovery heuristics that enable them to reinterpret the text. They do not start reinterpreting the sentence from its beginning. The heuristics point them to the locus of the probable error. Readers start the backtracking with the word that first reveals the inconsistency, in this case, *dress*. If that word cannot be reinterpreted, they make regressions to the site of other words that were initially difficult to interpret, such as ambiguous words on which a best guess about word meaning had to be made. The ability to return directly to the locus of the misinterpretation and to recover from an error makes the immediacy strategy feasible.

Sentence Wrap-Up

A special computational episode occurs when a reader reaches the end of a sentence. This episode, called sentence wrap-up, is not a stage of processing defined by its function, but rather by virtue of being executed when the reader reaches the end of a sentence. The processes that occur during sentence wrap-up involve a search for referents that have not been assigned, the construction of interclause relations (with the aid of inferences, if necessary), and an attempt to handle any inconsistencies that could not be resolved within the sentence.

The ends of sentences have two important properties that make them especially good places for integration. First, within-sentence ambiguities are usually clarified by the end of the sentence. For example, if a sentence introduces a new object or person whose identity cannot be inferred from the preceding context, some cue to their identity is generally given by the end of the sentence. For that reason, if readers can-

[1] It might be argued that the variables coding the text-grammatical roles ought to be independent of the number of content words. One might argue that a definition, for example, takes a fixed amount of time to integrate, regardless of the number of content words it contains. Although the model predicts a length-sensitive duration, the analysis can also be done with five simple indicator variables to encode the five levels of the grammar. This analysis produced a fit that was almost as good ($R^2 = .93$). The weights (assuming a zero intercept) were 250, 341, 257, 214, and 118 msec for the five categories, from topics to details. Although this alternative is not ruled out by the data, we will continue to retain the view that integration time depends on the number of content words involved.

not immediately determine the referent of a particular word, then they can expect to be told the referent or given enough information to infer it by the end of the sentence. Indeed, readers do use the ends of sentences to process inconsistencies that they cannot resolve within the sentence (Carpenter & Daneman, Note 1). The second property is that the end of a sentence unambiguously signals the end of one thought and the beginning of a new one. It can be contrasted with weaker cues that signal within-sentence clause boundaries such as commas, relative pronouns, and conjunctions that can signal other things besides the end of a clause. Since ends of sentences are unambiguous, they have the same role across sentences, and they may be processed more uniformly than the cues to within-sentence clause boundaries.

There is ample empirical support for the integrative processing at the ends of sentences. Previous eye fixation studies show that when a lexically based inference must be made to relate a new sentence to some previous portion of the text, there is a strong tendency to pause at the lexical item in question and at the end of sentence that contains it (Just & Carpenter, 1978). Readres were given paragraphs containing pairs of related sentences; the first noun in the second sentence was the agent or instrument of the verb in the first sentence:

(1a) *It was dark and stormy the night the millionaire was murdered.*
(1b) *The killer left no clues for the police to trace.*

In another condition, the integrating inference was less direct:

(2a) *It was dark and stormy the night the millionaire died.*
(2b) *The killer left no clues for the police to trace.*

It took about 500 msec longer to process Sentence 2b than 1b, presumably due to the more difficult inference linking *killer* to *die*. There were two main places in which the readers paused for those 500 msec, indicating the points at which the inference was being computed. One point was on the word *killer*, and the other was on the end of the sentence containing *killer*. Another eye fixation study showed that integration linking a pronoun to its antecedent can occur either when the pronoun is first encountered or at the end of the sentence containing the pronoun (Carpenter & Just, 1977b).

Reading-time studies also have shown that there is extra processing at the end of a sentence. When subjects self-pace the word-by-word or phrase-by-phrase presentation of a text, they tend to pause longer at the word or phrase that terminates a sentence (Aaronson & Scarborough, 1976; Mitchell & Green, 1978). The pause has been attributed to contextual integration processes, similar to the proposed interclause integration process here. Yet another source of evidence for sentence wrap-up processes is that verbatim memory for recently comprehended text declines after a sentence boundary (Jarvella, 1971; Perfetti & Lesgold, 1977). The model attributes the decline to the interference between sentence wrap-up processes and the maintenance of verbatim information in working memory. Finally, another reason to expect sentence wrap-up processes is that we have observed pauses at sentence terminations in an eye fixation study similar to the one reported here (Carpenter & Just, in press). However, the current study provides stronger evidence because the text was presented all at once.

The results indicate that readers did pause longer on the last word in a sentence. As Table 2 shows, the duration of the sentence wrap-up period is 71 msec.

It is possible that wrap-up episodes could occur at the ends of text units smaller or larger than a sentence. For example, the data of Aaronson and Scarborough (1976) suggest that there are sometimes wrap-up processes at the ends of clauses. It is also possible that wrap-up could occur under some circumstances at the ends of paragraphs. The decision of when and if to do a wrap-up may be controlled by the desired depth of processing. For example, skimming may require wrap-up only at paragraph terminations, whereas understanding a legal contract may require wrap-up at clause boundaries. In fact, the clause-boundary effects obtained by Aaronson and Scarborough are sensitive to the subjects' reading goals. The current analysis indicated that the final word in the paragraph might

also be a wrap-up point; it received an additional 157 msec of fixation. However, since readers also pressed a button to indicate that they had finished reading the passage, this parameter might be influenced by their motor response.

Finally, the model included one other factor that involves a physical property of reading, namely the return sweep of the eyes from the right-hand side of one line of text to the left-hand side of the next line. Return sweeps are often inaccurate, landing to the right of the first word in a line. The inaccuracy is often corrected by a leftward saccade to the first word. As a result of this error and recovery, the first word on a line eventually receives an increased gaze duration, relative to a line-medial word. Almost all readers we have studied display the undershoot, but there are considerable individual differences in whether they compensate for it by making an extra leftward fixation to the first word. In fact, some researchers have associated these corrective leftward movements with poor readers (Bayle, 1942). To test for increased gaze durations on line-initial words, an indicator variable coded whether a word was the first one on a line. As Table 2 shows, these words received an additional 30 msec of fixation.

Fit of the Model

To see how well the model accounts for the data, one can informally compare how closely the estimated gaze durations match the observed gaze durations. The display that follows shows the estimated (in italics) and observed (in msec) gaze durations for two sentences from the "Flywheel" passage. The estimated durations can be computed by an appropriate combination of the weights given in Table 2. These estimates take into account the processes of encoding, lexical access, case-role assignment, sentence wrap-up, and the beginning of the line effect; they do not include integration time for text roles, since there is no way to distribute this time on a word-by-word basis. In spite of this, the match is satisfactory, and as mentioned earlier, the standard error of estimate was 88 msec overall.

Observed mean gaze durations (msec) on each word of a text sample and estimates (italicized, from the word-by-word regression analysis.

	169	215	165	295	290	73	196	504	29	482	0	328	431	51			
	165	*236*	*75*	*409*	*304*	*75*	*249*	*438*	*75*	*413*	*80*	*338*	*349*	*78*			
. . .	One	type	of	flywheel	consists	of	round	sandwiches	of	fiberglas	and	rubber	providing	the			
369	326	308	22	272	253	128	199	69	336	32	41	267	197	70	164	195	
354	*318*	*297*	*75*	*378*	*138*	*77*	*239*	*128*	*326*	*87*	*102*	*206*	*209*	*112*	*87*	*127*	
maximum	possible	storage	of	energy	when	the	wheel	is	confined	in	a	small	space	as	in	an	
340	323	182	72	626	276	46	21	346	60	467	519						
465	*334*	*236*	*77*	*513*	*304*	*75*	*102*	*289*	*75*	*361*	*319*						
automobile.	Another	type,	the	"superflywheel,"	consists	of	a	series	of	rimless	spokes . . .						

Table 3 presents an analogous comparison from the sector-by-sector analysis; this includes integration time. Again, the estimates from the model match the observed data quite well. The standard error of estimate was 234 msec overall.

Another way to evaluate the goodness of fit is to compare the regression results to those of another model that lacks most of the theoretically interesting independent variables and contains only the variable that codes the number of syllables. For the word-by-word analysis, this rudimentary model produces an R^2 of .46, compared to .72 for the complete model. For the sector-by-sector analysis, the rudimentary model accounts for a large portion of the variance between the gaze durations on sectors (R^2 = .87). This is not surprising, since there is considerable variation in their lengths. The complete sector-by-sector model accounts for 94% of the variance, or 54% of the variance unaccounted for by the reduced model.

The regression equations were also fit to

Table 3
Observed and Estimated Gaze Durations (msec) on Each Sector of the "Flywheel" Passage, According to the Sector-By-Sector Regression Analysis of the Group Data

Sector	Observed	Estimated
Flywheels are one of the oldest mechanical devices	1,921	1,999
known to man.	478	680
Every internal-combustion engine contains a small flywheel	2,316	2,398
that converts the jerky motion of the pistons into the smooth flow of energy	2,477	2,807
that powers the drive shaft.	1,056	1,264
The greater the mass of a flywheel and the faster it spins,	2,143	2,304
the more energy can be stored in it.	1,270	1,536
But its maximum spinning speed is limited by the strength of the material	2,440	2,553
it is made from.	615	780
If it spins too fast for its mass,	1,414	1,502
any flywheel will fly apart.	1,200	1,304
One type of flywheel consists of round sandwiches of fiberglas and rubber	2,746	3,064
providing the maximum possible storage of energy	1,799	1,870
when the wheel is confined in a small space	1,522	1,448
as in an automobile.	769	718
Another type, the "superflywheel," consists of a series of rimless spokes.	2,938	2,830
This flywheel stores the maximum energy	1,416	1,596
when space is unlimited.	1,289	1,252

the gaze durations of each of the 14 readers individually. The subjects varied in their reading skill, with self-reported Scholastic Aptitude Test scores ranging from 410 to 660, which were correlated with their reading speeds in the experiment, ranging from 186 words per min. to 377 words per min. $r(12) = .54$, $p < .05$. The mean R^2 of the 14 readers was .36 on the word-by-word analysis and .75 on the sector-by-sector analysis. This indicates substantial noise in each reader's word-by-word data. Some of the regression weights of the readers indicated considerable individual differences with respect to certain processes. For example, 4 of the 14 readers spent no extra time on the last word of a sentence. Another parameter of great variability among readers was the extra time spent on novel words, which ranged from 94 msec to 1,490 msec.

Although the sector-by-sector regression analysis uses an independent variable (the sum of the case role coefficients) that is estimated from the same data, this procedure does not do violence to the results. To estimate the effect of this procedure, the 14 subjects were divided randomly into two subgroups, and the case-role coefficients were obtained for each subgroup in a word-by-word analysis. Then these coefficients were aggregated and used as independent variables in a sector-by-sector analysis, such that one subgroup's coefficients were used in the analysis of the other subgroup's sector gaze durations. The results indicated no difference of any importance between the two subanalyses, and generally confirmed that using the case role coefficients from the word analysis in the sector analysis was an acceptable procedure.

Some of the variables that were reliable in the word-by-word analysis were not reliable in the sector analysis. For example, sectors that included a line-initial word did not have reliably longer durations, and sectors that included the end of a sentence took 57 msec longer, but the reliability of the effect was marginal ($p < .08$). The sum of the logarithms of the frequencies of the words in a sector did not reliably affect gaze duration on the sector. These differences between the two levels of analysis indicate that some effects that are word specific are not reliable or large enough to be detected when the data are aggregated over groups of words. Nevertheless, some of these effects can be detected at the sector level if the appropriate analysis is done. For example, the reason that the frequency effect was not reliable is that the aggregation of the logarithms smooths over the differences between infrequent words and

frequent words. A regression analysis of the sector data shows a reliable word frequency effect if the independent variable encodes the number of infrequent words (arbitrarily defined as less than 25 in Kučera & Francis, 1967) occurring for the first time. This latter analysis indicates 82-msec extra spent for each infrequent word, and has an R^2 of .94. (Carpenter & Just, in press, reported a 51-msec effect for this variable).

Recall Performance

The recall of a given part of a text should depend in part on what happens to the information as it is read. A clause that is thoroughly integrated with the representation of the text should tend to be stored in long-term memory, and therefore should be recalled better. There are two factors that determine how well a clause will be integrated. First, those sectors on which more integration time has been spent, like topics and definitions, should be recalled better. As predicted the integration parameter for a text role (i.e., the five weights at the bottom of Table 2) reliably affected the probability that a sector would be recalled, $t(271) = 2.01$, $p < .05$. A second factor affecting integration is the number of times an argument of a clause is referred to in the text; each repetition involving that argument may initiate another integration episode that increases its chances of being recalled (Kintsch & van Dijk, 1978). A rough index of this kind of repetition was obtained by counting the number of times the arguments of each sector were repeated in subsequent sectors. The frequency of reference to the arguments did increase the probability of recalling a sector, $t(271) = 5.90$, $p < .01$.

The recall measure just reported was the proportion of the 14 subjects that recalled each of the 274 sectors. Two independent judges assigned 100%, 50%, or 0% credit for the recall of each sector, depending on whether it had been fully, partially, or not at all correctly recalled. Synonyms and paraphrases were given full credit if they were close to the gist of the sector. If only a part of a sector was recalled, then partial credit was given. The two judges were in full agreement about 80% of the time and in partial agreement (i.e., within 50%) on 94% of the judgments; disagreements were resolved by a third judge.

Text units that were higher in the text grammar were generally recalled better, $F(4, 269) = 5.67$, $p < .01$. There was a monotonic increase in the probability of recall as a function of a sector's level in the text grammar. Recall probabilities were lowest for details (.31), then increased for expansions (.34), subtopics (.39), definitions/causes/consequences (.41), and topics (.53). This replicates previous text-role effects observed with other types of texts (Meyer, 1975; Thorndyke, 1977). The model partially explains this result in terms of the processes that occur during comprehension. In addition, retrieval processes may play a role in this effect. For example, there may be many retrieval paths from less important concepts that lead to topics, but not vice versa. Also, a complete model of recall will have to consider how the recall of particular facts is affected by the reader's previous knowledge. Although the passages were generally unfamiliar, particular facts surely differed in their familiarity, and this could have a powerful effect on recall (Spilich, Vesonder, Chiesi, & Voss, 1979). Finally, there could be response output effects in recall. In summary, the results show that a model of the comprehension processes can be used to partially account for recall performance. To totally explain recall will require a precise account of the role of prior long-term knowledge and the role of retrieval and reconstruction processes in recall.

Discussion

This section discusses three aspects of the theory: first, the implications of the immediacy assumption for language processing in general; second, how variation in reading modes can be handled by the theory; and third, the relation of the current theory to other theories of reading.

The Immediacy Assumption

The model's ability to account for fixation durations in terms of the processes that operate on words provides some valida-

tion for the immediacy and eye–mind assumptions. Readers interpret a word while they are fixating it, and they continue to fixate it until they have processed it as far as they can. As mentioned before, this kind of processing eliminates the difficulties caused by the potential ambiguity in language. It avoids the memory load and computational explosion that would result if a reader kept track of several possible meanings, case roles, and referents for each word and computed the final interpretation at the end of a clause or a sentence. This architectural feature also allows a limited-capacity processor to operate on a large semantic network without being bombarded by irrelevant associations. After a single interpretation has been selected, the activation of the unselected meanings can be dampened to their base levels so that they will not activate their semantic associates any further. This minimizes the chances that the reader will be conceptually driven in many directions at the same time.

The cost of this kind of processing is fairly low because the early decisions usually are correct. This is accomplished by taking a large amount of information into account in reaching a decision. The processes have specific heuristics to combine semantic, syntactic, and discourse information. Equally important, the processes operate on a data base that is strongly biased in favor of the common uses of words and phrases, but one that also reflects the effects of local context. The cost is also low because the reader can recover from errors. It would be devastating if there were no way to modify an incorrect interpretation at some later point. However, there are error-recovery heuristics that seem fairly efficient, although the precise mechanisms are only now being explored (Carpenter & Daneman, Note 1).

The fact that a reader's heuristics for interpreting the text are good explains why the garden path phenomenon is not the predominant experience in comprehension; it only happens occasionally. Perhaps the most common, everyday garden path experiences occur when reading newspaper headlines; for example, *Carter Views Discussed* and *Judge Admits Two Reporters*. The incorrect initial interpretations occur because headlines are stripped of the syntactic and contextual cues that guide the processing of normal text. Similarly, many jokes and puns explicitly rely on the contrast between two interpretations of an ambiguous word or phrase (Schultz & Horibe, 1974). Even garden path sentences sometimes seem funny. The humor in all of these cases resides in the incongruity between the initial interpretation and the ultimate one. Garden path sentences are also infrequent because writers usually try to avoid ambiguities that might encourage or allow incorrect interpretations. These kinds of sentences are useful tools for studying comprehension because they indicate where the usual comprehension strategies fail. But the fact that they are not frequent indicates that a reader's heuristics usually are sufficient.

Variation in Reading

There is no single mode of reading. Reading varies as a function of who is reading, what they are reading, and why they are reading it. The proposed model for the reading of scientific texts in this task is only one point in a multidimensional space of reading models. However, such variation can be accommodated within the framework presented in this article.

The reader's goals are perhaps the most important determinant of the reading process. A reader who skims a passage for the main point reads differently than someone who is trying to memorize a passage, or another person who is reading for entertainment. Goals can be represented in several aspects of the theory, but the main way is to require that each goal is satisfied or at least attempted before proceeding on to the next word, clause, or sentence. These goals correspond to the major products of each stage of comprehension and to the specific demands of a particular task. For example, an obvious goal associated with lexical access might be that one interpretation is selected. An added goal associated with the task of memorizing a passage may require

rehearsing phrases or constructing explicit mnemonics before going on to the next phrase or sentence. But goals can be deleted as well as added. A speed-reader may well eliminate goals for syntactic coherence, because the strategy of skipping over many words will destroy the syntax. Variations in goals can be detected with the current theory and analytic techniques. For example, it is possible to determine how much time is spent integrating different kinds of text roles in different tasks. When readers anticipate a recognition comprehension test, rather than recall, they spend less time integrating details (Carpenter & Just, in press).

Reading also depends on the text, the topic, and the reader's familiarity with both. A well-written paragraph on a familiar topic will be easier to process at all stages of comprehension. The lexical items will be easier to encode, the concepts will be more easily accessed, the case and text roles will be easier to infer, and the interrelations will be easier to represent. All of these dimensions of variation can be accommodated, measured, and evaluated within the theoretical framework. Moreover, any adequate theory must be sufficiently flexible to encompass such variation.

Even reading of the same text under the same circumstances will vary from person to person. There are several plausible sources of individual differences in the theory. One interesting source is the operational capacity of the working memory. Readers with a large working memory should be able to retain more of the text in the memory while processing new text, so their integration of the information may be more thorough. A promising first exploration of this hypothesis has found a very strong correlation between working memory capacity and various aspects of reading comprehension tests (Daneman & Carpenter, in press). By contrast, traditional measures of passive short-term memory capacity do not have a strong correlation with reading comprehension. Operational capacity may depend on the automaticity of basic reading processes such as encoding and lexical access. Poor readers may devote more time and attention to these processes (Hunt, Lunneborg, & Lewis, 1975; Perfetti & Lesgold, 1977) and consequently have less capacity for maintaining previous information and integrating the new information (Case, 1978).

Theories of Reading

Previous theories of reading have varied in their choice of dependent measures, the levels of information represented in the theory, and the implementation of top-down effects. It is useful to consider how the current theory compares to these alternative proposals along these three dimensions.

One important feature of the current theory is its attempt to account for reading time on individual words, clauses, and sentences. This approach can be distinguished from research that is more centrally concerned with recall, question answering, and summarizing (e.g., Rumelhart, 1977b). The dependent measure is not an incidental aspect of a theory; it has important implications for which issues the theory addresses. The present focus on processing time has resulted in a theory that accounts for the moment-by-moment, real-time characteristics of reading. By contrast, the theory pays less attention to retrieval and reconstruction, two later occurring processes that are important to an account of summarization.

Another feature of the theory is the attempt to account for performance at several levels of processing. Previous theories have tended to neglect certain stages. For example, the reading models of LaBerge and Samuels (1974) and Gough (1972) focus on the word-encoding processes, whereas the model of Kintsch and van Dijk (1978) focuses on integration. This is not to say that these models do not acknowledge other aspects of processing, but simply that they describe detailed mechanisms for one aspect of reading and no comparable mechanisms for other stages. The current theory has attempted to span the stages of reading by describing mechanisms for the word-encoding and lexical-access stages, as well as the parsing and text integration stages. Moreover, it has

attempted to describe some formal similarities by placing them all within the architecture of a production system.

A final but important distinction among reading theories is the manner in which they accommodate top-down and bottom-up factors in reading (see Rumelhart, 1977a). Some reading theories, particularly those addressed to word encoding, omit mechanisms to account for top-down or contextual effects (e.g., Gough, 1972). At the other extreme, there have been some theories that appear to place a major burden of comprehension on contextual effects. Some of these are recent schema-based theories of language comprehension (Schank & Abelson, 1977). Others are the older top-down models, developed out of analysis-by-synthesis theory; these models suggested that readers form explicit predictions about the next word and fixate it merely to confirm the hypothesis (Goodman & Niles, 1970). The current model falls somewhere between the extremes. It allows for contextual influences and for the interaction among comprehension processes. Knowledge about a topic, syntactic constraints, and semantic associates can all play a role in activating and selecting the appropriate concepts. However, the printed words themselves are usually the best information source that the reader has, and they can seldom be entirely replaced by guesses from the preceding context. Thus the top-down processes can influence the bottom-up ones, but their role is to participate in selecting interpretations rather than to dominate the bottom-up processes. Finally, the production system architecture permits a degree of coordination among different processes, so that any stage can be influenced by any cotemporaneously or previously executed stage.

Future Directions

The current theory suggests two major avenues of reading research. One direction is to construct computer simulations that are driven by reading performance data. The postulated human heuristics can be implemented in a computer program to examine the resulting complex interactions among knowledge sources. Reading-time data may be sufficiently constraining to select among various alternative heuristics. We are currently implementing aspects of the model presented here as a production system in collaboration with a colleague, Robert Thibadeau, to develop greater specification and more stringent tests of the model.

Although the production system framework is not essential for the interpretation of the empirical results in the present study, it has other benefits. First, it provides an architecture that can accommodate the flexibility and interaction that has been observed among the processes in reading and still express typical or canonical processing. Even though this theoretical framework is minimally specified, it seems sensible to start at this point and allow successive generations of data to constrain it, as Newell (1980) suggests. Finally, when expressed as a computer simulation, the model retains correspondence to postulated human processes and structures. Collections of serial productions may correspond to heuristic processes employed in comprehension. The firing of parallel productions can be identified with spreading activation in long-term memory. The production system's working memory can be identified with the reader's working memory. Thus, the production system can be viewed as a useful theoretical vehicle, or excess baggage, depending on one's intended destination.

The second avenue includes further empirical research on the real-time characteristics of reading. Eye movement and reading-time methodologies can reveal reading characteristics with other types of texts, tasks, and readers. The useful property of these methodologies is that they can measure reading time on successive units of text. One method is to present the successive words of a sentence one at a time, allowing the reader to control the interword interval (Aaronson & Scarborough, 1976). This procedure is only one end of a continuum defined by what units are presented. Rather than single words, they could be phrases, clauses, sentences, or entire passages (Carpenter & Just, 1977a; Mitchell & Green, 1978; Kieras, Note 3). In this way, it will be possible to gain more information about

human performance characteristics and then use these data to develop a more complete theory of reading.

Reference Notes

1. Carpenter, P. A., & Daneman, M. *Lexical access and error recovery in reading: A model based on eye fixations.* Unpublished manuscript. Carnegie-Mellon University, 1980.
2. Charniak, E. *Toward a model of children's story comprehension* (Tech. Rep. 266). Cambridge, Mass.: MIT Artificial Intelligence Laboratory, 1972.
3. Kieras, D. E. *Modelling reading times in different reading tasks with a simulation model of comprehension* (Tech. Rep. 2). Tucson: University of Arizona, 1979.

References

Aaronson, D., & Scarborough, H. S. Performance theories for sentence coding: Some quantitative evidence. *Journal of Experimental Psychology: Human Perception and Performance*, 1976, *2*, 56–70.

Anderson, J. R. *Language, memory, and thought.* Hillsdale, N.J.: Erlbaum, 1976.

Bayle, E. The nature and causes of regressive movements in reading. *Journal of Experimental Education*, 1942, *11*, 16–36.

Baron, J. Mechanisms for pronouncing printed words: Use and acquisition. In D. LaBerge & S. J. Samuels (Eds.), *Basic processes in reading: Perception and comprehension.* Hillsdale, N.J.: Erlbaum, 1977.

Bever, T. G. The cognitive basis for linguistic structures. In J. R. Hayes (Ed.), *Cognition and the development of language.* New York: Wiley, 1970.

Bouma, H., & deVoogd, A. H. On the control of eye saccades in reading. *Vision Research*, 1974, *14*, 273–284.

Bransford, J. D., & Johnson, M. K. Considerations of some problems of comprehension. In W. G. Chase (Ed.), *Visual information processing.* New York: Academic Press, 1973.

Breitmeyer, B. G. Unmasking visual masking: A look at the "why" behind the veil of "how." *Psychological Review*, 1980, *87*, 52–69.

Buswell, G. T. How adults read. *Supplementary educational monographs (45).* Chicago, Ill.: University of Chicago Press, 1937.

Carpenter, P. A., & Just, M. A. Integrative processes in comprehension. In D. LaBerge & S. J. Samuels (Eds.), *Basic processes in reading: Perception and comprehension.* Hillsdale, N.J.: Erlbaum, 1977. (a)

Carpenter, P. A., & Just, M. A. Reading comprehension as eyes see it. In M. A. Just & P. A. Carpenter (Eds.), *Cognitive processes in comprehension.* Hillsdale, N.J.: Erlbaum, 1977. (b)

Carpenter, P. A., & Just, M. A. Cognitive processes in reading: Models based on readers' eye fixations. In A. M. Lesgold & C. A. Perfetti (Eds.), *Interactive processes in reading.* Hillsdale, N.J.: Erlbaum, in press.

Case, R. Intellectual development from birth to adulthood: A neo-Piagetian interpretation. In R. Siegler (Ed.), *Children's thinking: What develops?* Hillsdale, N.J.: Erlbaum, 1978.

Chafe, W. L. *Meaning and the structure of language.* Chicago, Ill.: University of Chicago Press, 1970.

Clark, H. H., & Clark, E. V. *Psychology and language.* New York: Harcourt Brace Jovanovich, 1977.

Collins, A. M., & Loftus, E. F. A spreading activation theory of semantic processing. *Psychological Review*, 1975, *82*, 407–428.

Daneman, M., & Carpenter, P. A. Individual differences in working memory and reading. *Journal of Verbal Learning and Verbal Behavior*, in press.

Dearborn, W. *The psychology of reading* (Columbia University contributions to philosophy and psychology). New York: Science Press, 1906.

Dixon, P., & Rothkopf, E. Z. Word repetition, lexical access, and the process of searching words and sentences. *Journal of Verbal Learning and Verbal Behavior*, 1979, *18*, 629–644.

Fillmore, C. J. The case for case. In E. Bach & R. T. Harms (Eds.), *Universals in linguistic theory.* New York: Holt, Rinehart & Winston, 1968.

Flesch, R. F. *How to test readability.* New York: Harper, 1951.

Fraser, C., Bellugi, U., & Brown, R. Control of grammar in imitation, comprehension, and production. *Journal of Verbal Learning and Verbal Behavior*, 1963, *2*, 121–135.

Frazier, L., & Fodor, J. The sausage machine: A new two-stage parsing model. *Cognition*, 1978, *6*, 291–325.

Glanzer, M., & Ehrenreich, S. L. Structure and search for the internal lexicon. *Journal of Verbal Learning and Verbal Behavior*, 1979, *18*, 381–398.

Goodman, K. S., & Niles, O. S. *Reading process and program.* Urbana, Ill.: National Council of Teachers of English, 1970.

Gough, P. B. One second of reading. In J. F. Kavanagh & I. G. Mattingly (Eds.), *Language by eye and ear.* Cambridge, Mass.: MIT Press, 1972.

Haviland, S. E., & Clark, H. H. What's new? Acquiring new information as a process in comprehension. *Journal of Verbal Learning and Verbal Behavior*, 1974, *13*, 512–521.

Hitch, G. J. The role of short-term working memory in mental arithmetic. *Cognitive Psychology*, 1978, *10*, 302–323.

Hockett, C. F. *A course in modern linguistics.* New York: Macmillan, 1958.

Hunt, E., Lunneborg, C., & Lewis, J. What does it mean to be high verbal? *Cognitive Psychology*, 1975, *2*, 194–227.

Jarvella, R. J. Syntactic processing of connected speech. *Journal of Verbal Learning and Verbal Behavior*, 1971, *10*, 409–416.

Judd, C. H., & Buswell, G. T. Silent reading: A study of the various types. *Supplementary Educational Monographs (23).* Chicago, Ill.: University of Chicago Press, 1922.

Just, M. A., & Carpenter, P. A. Eye fixations and cognitive processes. *Cognitive Psychology*, 1976, *8*, 441–480.

Just, M. A., & Carpenter, P. A. Inference processes during reading: Reflections from eye fixations. In J. W. Senders, D. F. Fisher, & R. A. Monty (Eds.), *Eye movements and the higher psychological functions*. Hillsdale, N.J.: Erlbaum, 1978.

Kintsch, W., & van Dijk, T. A. Toward a model of text comprehension and production. *Psychological Review*, 1978, *85*, 363–394.

Kleiman, G. M. Speech recoding in reading. *Journal of Verbal Learning and Verbal Behavior*, 1975, *14*, 323–339.

Kucěra, H., & Francis, W. N. *Computational analysis of present-day American English*. Providence, R.I.: Brown University Press, 1967.

LaBerge, D., & Samuels, S. J. Toward a theory of automatic information processing in reading. *Cognitive Psychology*, 1974, *6*, 293–323.

Levy, B. A. Interactive processes during reading. In A. M. Lesgold & C. A. Perfetti (Eds.), *Interactive processes in reading*. Hillsdale, N.J.: Erlbaum, in press.

Mandler, J. M., & Johnson, N. S. Remembrance of things parsed: Story structure and recall. *Cognitive Psychology*, 1977, *9*, 111–151.

McConkie, G. W., & Rayner, K. The span of the effective stimulus during a fixation in reading. *Perception & Psychophysics*, 1975, *17*, 578–586.

McConkie, G. W., & Rayner, K. Asymmetry of the perceptual span in reading. *Bulletin of the Psychonomic Society*, 1976, *8*, 365–368.

Mewhort, D., & Beal, A. L. Mechanisms of word identification. *Journal of Experimental Psychology: Human Perception and Performance*, 1977, *3*, 629–640.

Meyer, B. *The organization of prose and its effect on recall*. Amsterdam: North-Holland, 1975.

Meyer, B., & McConkie, G. W. What is recalled after hearing a passage? *Journal of Educational Psychology*, 1973, *65*, 109–117.

Mitchell, D. C., & Green, D. W. The effects of context and content on immediate processing in reading. *Quarterly Journal of Experimental Psychology*, 1978, *30*, 609–636.

Morton, J. Interaction of information in word recognition. *Psychological Review*, 1969, *76*, 165–178.

Newell, A. Production systems: Models of control structures. In W. G. Chase (Ed.), *Visual information processing*. New York: Academic Press, 1973.

Newell, A. Harpy, production systems and human cognition. In R. Cole (Ed.), *Perception and production of fluent speech*. Hillsdale, N.J.: Erlbaum, 1980.

Perfetti, C. A., & Lesgold, A. M. Discourse comprehension and sources of individual differences. In M. A. Just & P. A. Carpenter (Eds.), *Cognitive processes in comprehension*. Hillsdale, N.J.: Erlbaum, 1977.

Rayner, K. Eye movements in reading and information processing. *Psychological Bulletin*, 1978, *85*, 618–660.

Reitman, J. S. Without surreptitious rehearsal, information in short-term memory decays. *Journal of Verbal Learning and Verbal Behavior*, 1974, *13*, 365–377.

Rieger, C. J. Five aspects of a full-scale story comprehension model. In N. V. Findler (Ed.), *Associative networks*. New York: Academic Press, 1979.

Robinson, D. A. Eye movements evoked by collicular stimulation in the alert monkey. *Vision Research*, 1972, *12*, 1795–1808.

Rumelhart, D. E. Toward an interactive model of reading. In S. Dornic (Ed.), *Attention and performance VI*. Hillsdale, N.J.: Erlbaum, 1977. (a)

Rumelhart, D. E., Understanding and summarizing brief stories. In D. LaBerge & S. J. Samuels (Eds.), *Basic processes in reading: Perception and comprehension*. Hillsdale, N.J.: Erlbaum, 1977. (b)

Scarborough, D. L., Cortese, C., & Scarborough, H. S. Frequency and repetition effects in lexical memory. *Journal of Experimental Psychology: Human Perception and Performance*, 1977, *3*, 1–17.

Schank, R. C. Conceptual dependency: A theory of natural language understanding. *Cognitive Psychology*, 1972, *3*, 552–631.

Schank, R. C., & Abelson, R. P. *Scripts, plans, goals and understanding: An inquiry into human knowledge structures*. Hillsdale, N.J.: Erlbaum, 1977.

Schiepers, C. Response latency and accuracy in visual word recognition. *Perception & Psychophysics*, 1980, *27*, 71–81.

Shultz, T., & Horibe, F. Development of the appreciation of verbal jokes. *Developmental Psychology*, 1974, *10*, 13–20.

Spilich, G. J., Vesonder, G. T., Chiesi, H. L., & Voss, J. F. Text processing of domain-related information for individuals with high and low domain knowledge. *Journal of Verbal Learning and Verbal Behavior*, 1979, *18*, 275–290.

Spoehr, K. T., & Smith, E. The role of syllables in perceptual processing. *Cognitive Psychology*, 1973, *5*, 71–89.

Swinney, D. A. Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior*, 1979, *18*, 645–659.

Taft, M. Recognition of affixed words and the word frequency effect. *Memory & Cognition*, 1979, *7*, 263–272.

Tanenhaus, M. K., Leiman, J. M., & Seidenberg, M. S. Evidence for multiple stages in the processing of ambiguous words in syntactic contexts. *Journal of Verbal Learning and Verbal Behavior*, 1979, *18*, 427–440.

Taylor, S. E. An evaluation of forty-one trainees who had recently completed the "reading dynamics" program. *Eleventh yearbook of the national reading conference*, 1962, 41–55.

Thorndyke, P. W. Cognitive structures in comprehension and memory of narrative discourse. *Cognitive Psychology*, 1977, *9*, 77–110.

Vesonder, G. T. *The role of knowledge in the processing of experimental reports*. Unpublished doctoral dissertation, University of Pittsburgh, 1979.

Westheimer, G. H. Eye movement responses to a horizontally moving visual stimulus. *Archives of Ophthalmology*, 1954, *52*, 932–943.

Received May 14, 1979
Revision received February 29, 1980 ■