INSTITUTE FOR THE
STUDY OF
INTELLECTUAL
BEHAVIOR

Cognitive Processes in Skimming Stories

Michael E. J. Masson
Department of Psychology
University of Colorado

Technical Report No. 84-ONR

Institute for the Study
of Intellectual Behavior
University of Colorado
Boulder, Colorado 80309

September 1979

This research was sponsored by the Personnel and Training Research Programs, Psychological Science Division, Office of Naval Research, under contract No. N00014-78-C-0433, Contract Authority Identification Number NR 157-422

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER: 84 - ONR	2. GOVT ACCESSION NO.
3. RECIPIENT'S CATALOG NUMBER	
4. TITLE (and Subtitle): Cognitive Processes in Skimming	5. TYPE OF REPORT & PERIOD COVERED: Occasional report 10/1/78 - 9/30/79
	6. PERFORMING ORG. REPORT NUMBER: ONR
7. AUTHOR(s): Michael E. J. Masson	8. CONTRACT OR GRANT NUMBER(s): N00014-78-C-0433
9. PERFORMING ORGANIZATION NAME AND ADDRESS: ISIB, University of Colorado, Boulder, Colorado 80309	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS: NR 157-422
11. CONTROLLING OFFICE NAME AND ADDRESS: Personnel and Training Research Programs, Office of Naval Research (Code 458), Arlington, VA 22217	12. REPORT DATE: September 1979
	13. NUMBER OF PAGES
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)	15. SECURITY CLASS. (of this report): Unclassified
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Reading Rate, Text Memory, Processing Resources, Rapid Reading, Skimming

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

The research in this report investigated comprehension processes and memory representations involved in skimming stories. Experiments 1 and 2 used a timed recognition test to study the formation of macrostructure representations of narrative and newspaper stories while reading at rates ranging from 225-600 wpm. In some cases subjects were required to read stories typed in alternating cases in which every other letter was capitalized. Alternating case was meant to disrupt whole word visual identification processe which are hypothesized to play an important role in rapid reading tasks such as

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102-014-6601

skimming. Even while skimming at 600 wpm and while reading alternating cas subjects were able to form macrostructure representations of stories that included much of the important information in a story and few details. Decreased performance on the surface memory recognition test as reading rate increased and when alternating case was used was consistent with the hypothesis that rapid reading and alternating case both reduce the degree of detailed surface processing of text. In Experiment 3 subjects read stories from specific perspectives and were better able to recognize test statement relevant to the assigned perspectives than they were able to recognize irrelevant statements, even when skimming as fast as 600 wpm. Subjects who read alternating case, however, were not able to selectively process and represent in memory relevant information when they skimmed the stories. It appears that the demands imposed on cognitive processing resources by alternating case and by the selection strategy based on a given perspective exceeded readers' capabilities. All subjects were, however, able to form macrostructures based on statement importance independent of assigned perspectives. In Experiment 4, recall protocols taken from subjects who read stories at various rates (225-600 wpm) were successfully simulated by a text processing model developed for normal reading tasks. The results of these four studies argue that the basic processes in skimming are similar to those of normal reading, especially in terms of the development of macrostructure representations, although surface structure processing is limited and readers must rely to a greater extent on the use of general knowledge an predictive reading processes while skimming.

ABSTRACT

The research in this report investigated comprehension processes and memory representations involved in skimming stories. Experiments 1 and 2 used a timed recognition test to study the formation of macrostructure representations of narrative and newspaper stories while reading at rates ranging from 225-600 wpm. In some cases, subjects were required to read stories typed in alternating cases in which every other letter was capitalized. Alternating case was meant to disrupt whole word visual identification processe which are hypothesized to play an important role in rapid reading tasks such as skimming. Even while skimming at 600 wpm and while reading alternating case subjects were able to form macrostructure representations of stories that included much of the important information in a story and few details. Decreased performance on the surface memory recognition test as reading rate increased and when alternating case was used was consistent with the hypothesis that rapid reading and alternating case both reduce the degree of detailed surface processing of text. In Experiment 3 subjects read stories from specific perspectives and were better able to recognize test statements relevant to the assigned perspectives than they were able to recognize irrelevant statements, even when skimming as fast as 600 wpm. Subjects who read alternating case, however, were not able to selectively process and represent in memory relevant information when they skimmed the stories. It appears that the demands imposed on cognitive processing resources by alternating case and by the selection strategy based on a given perspective exceeded readers' capabilities. All subjects were, however, able to form macrostructures based on statement importance independent of assigned perspectives. In Experiment 4, recall protocols taken from subjects who read

stories at various rates (225-600 wpm) were successfully simulated by a text processing model developed for normal reading tasks. The results of these four studies argue that the basic processes in skimming are similar to those of normal reading, especially in terms of the development of macrostructure representations, although surface structure processing is limited and readers must rely to a greater extent on the use of general knowledge and predictive reading processes while skimming.

CHAPTER I

INTRODUCTION

The process of reading is one of the most complex cerebral functions of which human beings are capable. The cognitive apparatus required to carry out the steps involved in reading the written word appears to be correspondingly complex and intricate. Given the degree of influence of cognitive processes upon reading and the fundamental importance of these processes for cognitive theory in general, it is highly appropriate that reading has become a prominent area of concern for cognitive psychology.

This thesis will focus on a particular type of reading that is commonly referred to as *skimming*. In the initial part of this chapter I will provide a justification for studying skimming through a discussion of objectives underlying the research to be described in later chapters. Next, the relationship between skimming and speed reading will be considered in some detail in the context of previous research on skimming, speed reading, and reading improvement. The chapter will conclude with a discussion of the rationale for the experimental work reported in Chapters II-V. The final chapter will summarize these empirical results and will present a discussion of implications for skimming and speed reading as well as a plan for future research on skimming.

Objectives in Studying Skimming

Reading has long been a primary area of concern among researchers in education. Only in the comparatively recent past have cognitive psychologists returned to a serious consideration of reading processes. A long hiatus followed Huey's insightful treatise written in 1908. Current psychological research and models pertaining to reading have focused on various aspects of the process, from word recognition and decoding to comprehension of texts. A wide variety of methodological techniques have been employed in these efforts, some more closely related to natural reading processes than others (e.g., tachistoscopic word recognition).

One of the fundamental conclusions that can be drawn from much of cognitive research on reading is that the process of reading should not be described as a simple "bottom up" flow of information from the visual system to more complex areas of the brain suggested in models proposed by Gough (1972) and LaBerge and Samuels (1974). A more realistic account of reading, and natural language processing in general, should include consideration of "top down" or <u>conceptually driven</u> processes (e.g., Bobrow & Norman, 1975).

The construct of conceptually driven processing is consistent with theories that claim that reading is a predictive, hypothesis testing type of process (Gibson & Levin, 1975; Goodman, 1967; Hochberg, 1970; Hochberg & Brooks, 1976; F. Smith, 1971). Some examples of research that support ideas about the

predictive nature of the reading process include Marcel's (1974) finding that increasing contextual constraint on sentences by using G. A. Miller and Selfridge's (1950) method of successive approximations to English has the result of enlarging the effective visual field during reading. The apparent cause of the effect is reduced sampling of visual information for word recognition due to increased contextual constraint. This reduction in sampling, in turn, allows more capacity for visual processing of information in the periphery. Steinheiser and Guthrie (1974) found that scanning prose or word strings for a particular word takes less time than scanning for a single letter. Scanning for a single phoneme required the longest search time. The unit of analysis therefore appears to be at least as large as a whole word and specific graphemic or phonemic information must be extracted after encoding of the whole word has at least begun. Drewnowski and Healy (1977) have provided evidence that subjects may even read in units larger than single words. Subjects were assigned the task of detecting the words "the" and "and" in isolation and when embedded in larger words (e.g., another). Failures to detect the target words were more common when the targets occurred in isolation as individual words, indicating that they had been processed as part of a larger unit with little attention paid to the target words themselves. The highly predictable nature of the occurrence of "the" and "and" as function words leads readers to take less notice of their occurrence.

Of course people do not confine themselves to predicting what is written before them on a page without consulting their visual perceptions of the actual markings. There must be some way in which a reader's predictions and hypotheses about what is being read can be confirmed and elaborated, rejected, or revised. A theoretical solution to the problem of how predictive reading can be constrained and made accurate that is gaining wider acceptance is the construct of interactive processes (Adams & Collins, 1977; Bobrow & Norman, 1975; J. Frederiksen, 1978; Goldstein & Papert, 1977; Masson & Sala, 1978; Rumelhart, 1977a). Within the framework of interactive processes, the task of reading is conceptualized as a set of subskills that interact with one another. The "bottom up" or data driven skills and processes, which are responsible for deciphering and bringing into the information processing system the actual symbols being read, do not operate independently of the predictions and hypotheses generated by conceptually driven processes. Rather, conceptually and data driven processes interact with and influence one another. For example, conceptually driven predictions about what to expect next while reading can be verified or rejected by the products of a data driven analysis of what is actually printed. If a prediction is verified reading can proceed rapidly and further predictions may be more precise; if not, new predictions can be generated, hopefully within the constraints defined by the new information provided by data driven processes. Similarly, data driven processes can be

influenced by conceptually driven processes, as when a reader is searching for or expecting a certain kind of information. This type of influence can allow speeded perception of what is written with a high degree of accuracy, and it can also lead one to misread what is actually written, making one's perception of the written words consonant with expectations even though the true message may contain contradictory information.

Rumelhart (1977a) has reviewed a number of empirical studies that support the notion that knowledge of the world and the language act to influence perceptions in a conceptually driven manner. Beyond the studies discussed by Rumelhart there exist classes of other supportive findings. From a developmental perspective, Biemiller (1970-71) studied oral reading errors of children of various ages. Biemiller detected three developmental stages, the first of which was characterized by contextually constrained reading errors. That is, reading errors were most commonly of the sort that conformed to semantic constraints of the surrounding passage. The second stage was defined primarily by graphically constrained reading errors, and the final stage was characterized by errors constrained by graphics and context (both semantic and syntactic). Thus, in the final stage of reading, constraints on predictive reading are both conceptually and data driven. A developmental study by Mason (1977) found that word decoding was dependent on letter properties as well as familiarity of whole words and that the decoding process is best viewed as holistic and analytic.

Dependence of decoding on individual letter properties and on familiarity of whole word patterns supports the influential role of data and conceptually driven processes in decoding. Flood (1978) constructed passages that contained information that conformed or did not conform to expectations based on the first sentences of the passages. After reading each passage, high and low level eighth grade readers answered questions about the passage. At least among the poor readers, question answering was least accurate when the passage violated expectations based on the first sentence. Using samples of good and poor readers as well, Steiner, Wiener, and Cromer (1971) found that good readers could extract and use syntactic and semantic cues in reading while poor readers were unable to extract such cues or even make use of them if they were explicitly provided by the experimenter. Success in reading appears to depend on the ability to make use of various contextual constraints while reading. It is not sufficient to plod through a complete letter by letter decoding process. In fact, Samuels, Begy, and Chen (1975-76) found that fluent readers had faster word recognition performance and were better at generating a target word given context and minimal cues about constituent target letters than slower readers. Furthermore, Samuels, Dahl, and Archwamety (1974) demonstrated that training retarded and normal children on hypothesis testing word recognition subskills led to improved tachistoscopic word recognition and cloze comprehension performance in both groups. These results imply that as the reading process develops, the

role of conceptually driven constraints on data driven processes becomes increasingly important, and consideration of the interaction of these processes is critical.

Evidence for the important role played by conceptually driven processes in reading is also available from cognitive research. Knowledge of syntactic structure of sentences decreases both the amount of time and cognitive resources required to read the sentences (Wisher, 1976). Semantic constraints also serve to decrease the amount of cognitive effort required in reading. Cutler and Fodor (1979) used a phoneme detection task to examine efficiency of reading words related or unrelated to the goal in reading a passage. If the target phoneme occurred in a word relevant to the reader's purpose, detection time was lower than when the phoneme occurred in a word not related to the purpose. Sentence comprehension, then, can be facilitated by the rapid identification of relevant information. Masson and Sala (1978) found that knowledge about semantic and surface aspects of sentences that subjects have previously read can positively influence the reading time and recognition of those sentences when presented a second time.

The role of conceptually driven processes is also an important one from a perceptual point of view. Research on eye movements during reading has provided ample evidence to conclude that eye movements are strongly influenced by cognitive processes occurring during fixations (Rayner, 1978). This conclusion implies an active, conceptually driven processing mechanism that searches out particular information from relatively specific locations. Fixations during reading that involve processing of both position and text information take longer than fixations involving processing of position information alone (Abrams & Zuber, 1972-73). The latter type of fixations primarily occur just prior to return sweeps of the eyes after the end of a line has been reached and just before experimentally imposed elongated spaces between words. During a fixation certain types of information are available for processing and can influence the reader's expectations. McConkie and Rayner (1975; Rayner, 1975) have shown that subjects acquire enough information within ten character spaces to the right of the fixation location to semantically interpret a word. Further into the periphery, up to about 15 characters to the right of the fixation point, the subject can perceive only general characteristics of word shape. But even this vague information is used to guide eye movement patterns such as saccade length. In fact, oral reading accuracy depends heavily upon presence of information in the periphery (Poulton, 1962). P. Carpenter and Just (1977a, 1977b) and Shebilske, Reid, and Wright (1977) have demonstrated that fixations are longer in duration at those points when subjects are integrating fixated information into higher level memory structures. In their recent development of a model of eye movements during reading Just and Carpenter (1979) found empirical support for longer fixation durations corresponding to the occurrence of more complex stages of the reading process.

Subjects pause to complete at least preliminary comprehension of one part of a text before going on to obtain further information.

The data on eye movements during reading are consistent with the theoretical concept of interaction between perceptually oriented data driven processes and conceptually driven cognitive processes. This is mutual influence and responsiveness on the part of both types of process. Developmental and cognitive data on reading are also consonant with interactive processing theories. The apparent importance of interaction between processes precludes the feasibility of research completely devoted to the study of processes in isolation. A more fruitful approach would include study of the whole reading process active in different situations. Study of a particular subprocess or subskill could be accomplished by selecting reading tasks that maintain the basic nature of the reading process and that emphasize the specific subprocess of concern. A major objective of this thesis is to investigate the role of conceptually driven processes in reading and their influence on and responsiveness to data driven processes. Consequently, it is essential to develop a reading task that emphasizes the role of conceptually driven processes. Masson and Sala (1978) have used transformed typography of sentences as a method of increasing the importance of conceptually driven processes with some success. This type of methodology may be adequate for determining basic influences of conceptually driven processes, but a more detailed investigation of the subtle effects of conceptually driven processes probably requires a more naturalistic reading situation. Less natural reading tasks may distort subtle effects to a great extent and, thus, provide an inaccurate portrayal of conceptually driven processes.

A natural reading situation that presents itself as a prime candidate for an in-depth study of conceptually driven processes is skimming. When reading certain types of material, one is not always concerned with reading and fully comprehending every detail of what was written. Often one reads for a particular kind of information "ignoring" irrelevant information in the text. In these cases the reader moves more swiftly through the text than when reading for full comprehension. Since the reading process is speeded the reader probably does not fully decode many words or sentences in the text. Rather, skimming likely involves an important interaction between data driven processes which provide barely adequate information about what is actually on the page and conceptually driven processes which use that information in developing a comprehensible memory representation of the text by confirming or rejecting expectations related to the reader's goal. The evaluative role played by the products of data driven processes and the highly active nature of conceptually driven processes in their use of the decoded information places relatively small demands on the quality of that information and allows (or even requires) a very rapid flow of unintegrated information. Thus, the process of skimming was selected as a task that would emphasize the role of conceptually driven processes in

reading.

A related objective is, of course, to study the skimming process itself. Skimming is a form of reading that almost everyone has done in various situations and, therefore, is a basic part of a reader's skill. It is important to know how people skim and the extent to which skimming can be effective. That is, if a person skims a text for its gist or for a specific kind of information, can that person actually extract the desired information while processing the text at a relatively rapid rate? To the degree that this is possible, skimming will be defined in this thesis as effective. The two objectives in studying skimming are closely related and to the extent that the research described here is successful in meeting one of the objectives, the other objective will have been met to the same general degree.

Skimming, Reading Improvement, and Speed Reading

Skimming represents a speeded form of reading and, while there has been little research directed at skimming processes themselves, there are two branches of empirical work that relate to the general issue of reading speed. One of these branches of research is concerned with improvement of reading speed and the other branch pertains to speed reading.

Educational research on improvement of reading speed has generally conformed to the basic paradigm of training some subjects in a reading improvement program and evaluating their improvement either on the basis of their reading performance prior to training or relative to reading performance of a control group. A basic and unresolved issue in this field is the measurement of reading speed relative to comprehension. Improvements in reading speed are useful only if comprehension is not sacrificed. A trade-off probably exists in which, beyond certain reading rates, increased reading speed will result in decreased comprehension. Researchers have not used uniform measures of comprehension and as a result conflicting results have often been found. In fact, a similar problem exists in research on speed reading. A resolution of this troublesome issue will be offered at the end of this section. In the meantime, evaluation of experiments will be made in the context of the types of comprehension measures used in the hope that a consistent picture of the relationship between reading speed and comprehension will emerge.

Evaluation of the effectiveness of the reading speed improvement methods discussed below should be done in light of Huey's (1968) claim that some increment in speed without comprehension loss is possible simply by forcing a faster rate and avoiding the tendency to "plod" while reading. Poulton (1963) has also suggested that most readers operate below their maximum potential rate of information processing. Furthermore, Rankin (1963) has shown that reading speed improves more reliably if training first emphasizes speed then comprehension, rather than comprehension first. Subjects can profit from some emphasis on speed even before comprehension training begins. Reports of

increased reading rates may simply be due to subjects realizing their maximum potential rather than representing the result of some major change in the reading process due to a training program.

Reading rate is related to visual span of perception (Buswell, 1957; Gilbert, 1959; Jackson & McClelland, 1975) and a number of studies have reported attempts to improve reading speed by providing practice in improving the visual span of perception (Amble, 1967; Brim, 1968; Lloyd, 1978). While these studies were successful in demonstrating some reading rate improvement, the quality of comprehension was in doubt. Only in Amble's (1967) review of developmental studies was comprehension shown to acutally improve as rate increased. These improvements, however, were attributable to accelerated development of reading skills that likely would have occurred by adulthood without special training. Olsen, Harlow, and Williams (1977) have used a parallel method of rate training in the realm of braille reading. Blind subjects entered a reading improvement program that emphasized the use of more than one finger, working on multiple lines simultaneously, and integration of the obtained bits of information into comprehensible ideas. Reading rate increased from a beginning level of 85 words per minute (wpm) to 120 wpm, with no reliable drop in comprehension. It was found, however, that the largest rate gains were associated with the lowest comprehension scores. On the other hand, Marcel (1974) found that after receiving training in advanced reading subjects demonstrated a greater use of visual periphery. Providing subjects peripheral vision training can improve reading rate (Sailor & Ball, 1974) but the quality of comprehension at the improved rate is in doubt. In general, then, it may be that improved span of perception, whether it be visual or tactile, may provide moderate increases in reading speed. Beyond modest increases expected by Huey (1968) and Poulton (1963), however, comprehension is not maintained.

Improvements in reading speed that are concomitant with improved comprehension are the most powerful demonstrations of all. T. Carpenter and Jones (1975) trained subjects to improve reading rate and comprehension by seeking main ideas and developing critical reading ability. In this way, reading speed increased to an average of 515 wpm with improved comprehension on the Nelson-Denny test. The standardized Nelson-Denny test is constructed so that the comprehension questions are typed next to the text to be read and there is no guarantee that subjects do not adopt strategies to search for answers to particular questions rather than reading the text as rapidly as 515 wpm then attempting to answer the questions. After all, training in rapid reading can lead to rate and comprehension improvement when comprehension questions are presented prior to reading a passage (Flynn, 1977). Moreover, presenting questions before reading a text will improve performance on those questions but will decrease performance on questions not presented prior to reading (Anderson & Biddle, 1975). These results cast doubt on the

validity of comprehension scores based on the Nelson-Denny test as a pure measure of ability to read rapidly and then answer comprehension questions on virtually any part of the text. Comprehension tests of the Nelson-Denny variety suffer from another problem in the context of measuring reading improvement. If pre- and posttest scores are used to determine whether reading rate has improved without comprehension loss, it is quite possible that subjects' posttest scores will be inflated due to prior practice on the test and the development of more efficient strategies for taking the test. Consequently, the results of studies using the Nelson-Denny test, especially in a pre- and posttest paradigm, must be highly suspect.

Another method for improving reading speed involves giving subjects practice at scanning paperbacks and training of efficient eye fixations, recognition span, and reduction of subvocalization (Berger, 1972). Baer (1974), Berger (1968), and Brown (1976) have successfully increased reading rates by training subjects to scan. No apparent loss in comprehension was found, but the specific nature of the comprehension tests were unclear. Similar results were reported by Gluck (1969), but comprehension scores based on the Nelson-Denny test showed declines from base scores in many cases, although there was no overall change in comprehension.

A number of methods of pacing subjects at rapid rates of speed have been used in an effort to increase natural reading rates. One class of methods involves the use of compressed speech which requires subjects to listen to tape recordings of speeded oral readings of passages. With extended practice in listening to compressed speech Berg (1977) and Thames and Rossiter (1972) have shown that subjects can increase reading speeds somewhat with no loss of comprehension. The amount of improvement in speed, however, was not great and may have been due to subjects' approaching realization of their maximum potential rather than being due to a real improvement in reading skills. Without extended practice, listening to compressed speech does not appear to be helpful (Stamper, 1976).

Attempts to improve reading rate by training subjects to read at fast, paced rates of speed have had limited success (Brown, 1976; Gluck, 1969; Himelstein & Greenberg, 1974; Laffitte, 1964; Maxwell & Mueller, 1967; Poulton, 1961). In all cases either the amount of reading rate improvement was no larger than might be expected by Huey's (1968) suggestion about self-improved reading speed, or if the improvements were reasonably high, quality of comprehension was questionable or clearly worsened (Buswell, 1957; C. Smith, 1976). In his review of research on the use of mechanical devices to pace rapid reading Tinker (1967) warned that resulting gains in reading speed are no greater than those that could be accomplished by increased motivation.

Efforts to increase reading speed in children have commonly resulted in decreased comprehension (Bonsall & Dornbush, 1969; Neville, 1975; Swalm & Kling, 1973) and have led researchers to

recommend concentration of training on comprehension and development of basic learning skills before proceeding to an emphasis on speed. In fact, there is some concern that emphasis on speed in reading at too early an age may retard the development of flexibility in reading strategies or even basic skills.

A rather interesting alternative approach to improving reading speed consists of the use of reduced text (Bassin & Martin, 1976; Martin & Bassin, 1977; Martin & Sheffield, 1976). Passages are reduced to a fraction of the number of original words by eliminating unimportant or redundant words. Reduction can be based on a number of different criteria, such as word frequency, rated subjective importance, or grammatical importance. For the most part, both sighted and blind subjects were able to read the reduced text versions without major losses in comprehension. But reading rate declined as the amount of reduction increased so that the total time taken to read full and reduced texts were the same. Subjects seemed to be able to cope with a limited rate of information flow (cf. Carver, 1977; Poulton, 1958, 1963), and these limitations appear to be based at the level of higher order processing rather than low level perception (Keen, 1973).

Given that methods of training subjects to read at rates of speed greater than their normal reading rates has not proven very successful, it would be worthwhile to examine possible causes of this lack of success. Research on the effects of increasing reading rate without prior training on comprehension and memory has provided insight into the information processing constraints imposed on a reader. I argue that it is these constraints that are responsible for the failure of training in speeded reading to produce dramatic increases in reading rate. In an early study on the effects of forcing subjects to read beyond their normal rates, Jester and Travers (1966) paced subjects at rates of up to 350 wpm. Pacing was accomplished by the use of controlled visual presentation of segments of the text, compressed speech, or both. Comprehension was measured by a multiple-choice test. It was found that comprehension decreased in a linear fashion as reading rate inceased. At 300 wpm greatest comprehension was observed for those subjects who were paced both visually and aurally. Comprehension was lowest among subjects paced with compressed speech alone. Speeded reading comprehension probably was best when compressed speech accompanied visual pacing because the auditory information could encourage efficient eye movements, especially by reducing regressions, and could provide text information not completely processed by the visual system. Lower comprehension of compressed speech presented in isolation indicates that reading may have greater potential for rapid information input than listening. The extent of this potential will be examined in later sections of the thesis.

Kieras (1974) has also investigated the consequences of speeded reading. Subjects read sets of related sentences either at their own rates or at a rapidly paced rate. Recognition of the

sentences in a later test was lower for those subjects forced to read at a fast pace. Kieras suggested that the effects of the speeded pace were to cause sentence information to be stored in a less complete and less accessible form. It appears that readers are able to process information for optimal comprehension at some maximum rate of speed. Beyond that rate comprehension is not as complete as when lower rates are used. If subjects are forced to read beyond the maximum rate they will necessarily process less of what is "read" and must, therefore, learn to select the appropriate information for full memory representation (cf. Poulton, 1963). In a supporting experiment by Poulton (1958), subjects studied all or some fraction of a set of test sentences for a fixed time period. On a test of the whole set of sentences only those subjects who studied one quarter of the set did poorly. Subjects who studied more than one quarter, including those who studied the whole set, did equally well on the test of total recall. A maximum information input rate had apparently been reached by those subjects who studied more than one quarter of the sentences. Those who were given only one quarter of the sentences were operating well below their optimal input rate.

Sticht (1977) has presented a theory of reading literacy in which he has claimed that the optimal level of reading that one can achieve is constrained by one's ability to comprehend speech. Even with the most efficient decoding processes available, a reader's comprehension rate will be limited to his or her speech comprehension rate. This view of the relationship between reading and speech comprehension hints at some fundamental constraints on natural language processing and is shared by Carver (1977). In his work on pacing readers at rates of up to 100,000 wpm Carver has found that the maximum rate of comprehension of written material and of speech for college students is in the vicinity of 300 wpm. The optimal rate of 300 wpm was found across all levels of difficulty of material tested (elementary school to college). Rates greater than this led to decreased comprehension, and beyond 1000 wpm comprehension test scores were at chance level. Furthermore, a test of free reading speeds of college students on materials of differing levels of difficulty showed that the average rate was about 300 wpm.

The results of research on forcing subjects to read at increased rates of speed are consistent in their finding of decreased comprehension. Even with training at faster reading rates people are not able to dramatically improve their reading speed and maintain full comprehension. Contrasting with these results are claims made by advocates of speed reading who have argued that people can be taught to read at rates of thousands of words per minute without loss of comprehension. In evaluating these claims the issue of what is meant by "comprehension" will once again play a crucial role. In resolving the commercial enthusiasm of speed reading advocates with experimental fact it will become apparent that the single most important underlying issue is that of comprehension.

Wood (1960) touted speed reading as a breakthrough in

improving reading ability and implied that people who read very fast become very well informed. Of course, it is also possible that being well informed is a prerequisite of being able to read quickly. In fact, G. Stevens and Oren (1963) have argued that speed reading requires adequate conceptual background knowledge. They also claimed that speed reading can be accomplished through a procedure of reading parts of several different lines of text at once. It has been shown that success in speed reading is linked to alterations in certain aspects of the reading process. In particular, being able to avoid serial subvocalization of the text's words is a factor in developing speeded reading (Hignerey, 1975; Voluse, 1973). Avoiding serial subvocalization would be related to the process of reading different parts of the text simultaneously as it is difficult to carefully attend to multiple "conversations" at the same time. Apparently there are some basic changes in reading strategy that accompany training in speed reading. The question remains as to whether speed reading can provide adequate comprehension.

One of the most convincing demonstrations of successful speed reading to date was performed by Barrus, Brown, and Iouye (1978). They argued that it is not necessary to see every word to comprehend all the information in a text, and that speed reading can provide full comprehension. They selected three groups of subjects, one of which was composed of select graduates of a speed reading course. Graduate students in social psychology and undergraduate honors students all without speed reading training formed the other two groups. The material used in the testing were passages from a social psychology text. The procedure for reading always involved reading the passage for an allotted time period, making notes about the passage from memory, postviewing the passage at a rapid rate, making more notes from memory, then providing a detailed verbal summary. One passage was read at 3000 wpm and postviewed at an average of 4500 wpm. Another passage was read for a fixed time corresponding to a rate of 200 wpm. Subjects could read the passage multiple times and, of course, the speed readers did just that. Postviewing time corresponded to 650 wpm. One passage was read at the readers' preferred rates: 1800 wpm for speed readers and about 320 wpm for the others. Scoring of the summaries was based on an outline of the text consisting of major and minor concepts and supporting details. Judges blindly rated the quality of the protocols. Subjects who had received speed reading training showed superior recall on the first two reading tasks, and were not significantly different from the other groups on the free reading passage. This seemingly convincing demonstration of the effectiveness of speed reading training suffers from a number of serious problems. Perhaps the most important flaw is that there is a differential selection process involved in the classification of subjects as speed readers or normal readers. Although all subjects had reasonably similar undergraduate grade point averages, there is no guarantee that the selected group of speed reading trainees did not differ from the other subjects in important ways, some of which may

involve reading abilities quite unrelated to speed reading training per se. A related criticism is that the experiment was designed so that no valid comparison of recall after reading at different rates could be made within-subjects. Since the between-subject comparisons were questionable at best, it would be very important to determine whether speed readers recall more information when forced to read at slower speeds than when speed reading. Such a result would call into question the claim that one need not see every word to comprehend all the information in the text. Also, the reading and scoring methods were both biased in favor of the speed reading trainees who have been specially instructed in methods such as postviewing and forming outlines. These procedures could have provided the speed reading group with an extremely unfair advantage. Finally, comprehension was estimated in the experiment by comparing recall protocols to an outline of the passage being read. Quality of comprehension, then, must be defined within the context of general information included in an outline. Given these considerations, the Barrus et al. (1978) experiment does not seem very convincing.

Even less promising results of speed reading training have been reported in other experiments. In some cases, training has led to improved speed but significantly lower comprehension (Graf, 1973; D. Stevens & Adams, 1968). In a study by Liddle (1966) training subjects to speed read led to decreased comprehension on three different measures. Speed readers were not able to obtain sufficient information or details to recognize basic facts or draw related inferences or necessary conclusions.

The measurement of comprehension in experiments on speed reading is crucial, for it is this measure that will shape the conclusions that one can draw about the degree of success that a speed reader has achieved. A number of studies are very revealing in this regard. Hansen (1975) carried out a discourse analysis of recall protocols of speed reading trainees and untrained subjects all of whom were allowed to read an unlimited amount of material in a fixed time period. A striking effect was obtained in that control subjects recalled fewer idea clusters than the trainees (probably because they didn't read as far in the passage), but the control subjects recalled more information per idea cluster. A study of recall protocols of speed readers done by McLaughlin (1969) revealed a good deal of reconstructive recall implying that subjects pieced together bits of information based on general knowledge. Arguments have been made by B. Smith (1975) and by Witty (1969) that the standard comprehension tests used in speed reading courses may be limited in the scope of evaluation of comprehension and that there is also some question about the ability to transfer learned speed reading skills to materials quite different from those used in training. Demonstrations of the validity of these concerns are readily available. Rauch (1971) gave a sample training multiple-choice comprehension test to a group of teachers after they had read only the title of the relevant article. Nearly half of the teachers scored 70% or better on the test. Similar effects were reported by Carver

(1971). In another demonstration Ehrlich (1963) gave graduates of a speed reading course a page of typewritten material to speed read. The subjects read the material over three times for an overall speed of 1700 wpm and were satisfied that they had understood the article. Ehrlich had constructed the article by writing out two lines at a time from two different articles, alternating the pairs of lines to form the final passage. The subjects had "understood" a rather incomprehensible passage! This serves as a rather strong example of how much the reader must contribute to what is being read and of just how little external guidance of thinking (Neisser, 1967) some forms of reading provide.

Further doubts about the feasibility of speed reading stem from research on eye movements during reading which has led to the conclusion (Spache, 1962; Taylor, 1965) that reading at speeds beyond 800-900 wpm is impossible due to fixation durations and the number of fixations required to perceive all words in a text. McLaughlin (1969) has refuted these claims and argued that speed readers have patterns of eye movements that are characterized by straight sweeps down the page sometimes marked by small horizontal movements. These patterns are quite different from those of normal readers which are characterized by general left to right horizontal movements and regressions. According to McLaughlin, speed readers are able to fully comprehend material because they can process different parts of a text in parallel and can perceive parts of different sentences in a single fixation. Both of these claims are highly questionable. The work of McConkie and Rayner (1975; Rayner, 1975) has shown that the visual system is relatively limited in the degree of resolution of words seen in the periphery. Without the necessary eye movements it seems impossible to clearly perceive widely separated parts of a text simultaneously. It is also very doubtful that a reader would be able to adequately integrate into a memory structure of a text information from a number of different sentences at once. The cognitive processing resources required for such a task probably would exceed the reader's capacity (Kahneman, 1973; Norman & Bobrow, 1975).

A stringent test of the hypothesis that speed readers are able to smoothly move their eyes down a page while comprehending a text would consist of a comparison of eye movement data of normal and speed readers. Taylor (1962) has provided just such a comparison. He collected eye movement protocols of speed reading trainees while reading normally and while speed reading. He found that eye movement patterns during speed reading resembled those that occur during skimming or scanning and consisted of arhythmic, small left to right saccades while generally moving down the page. Similar patterns have been observed among subjects who skim or scan a text (Hultgren, 1968; Spache, 1962; Spragins, 1974). When Taylor pressed his subjects to read at a greater rate of speed with smooth, almost purely vertical eye movements, comprehension suffered badly. Data such as these have led to the conclusion that speed reading should be considered a form of

skimming (Taylor, 1965). Taylor has argued that people can be taught to skim and scan at thousands of words per minute and that the corresponding eye movements involve looking and reading, with the reading phase similar to eye movements during normal reading. Tinker (1962) has made the point that speed reading should be considered a form of skimming because when one reads at rates over 1000 wpm many words can't be seen and are, therefore, skipped as in skimming. Schale (1965) and Carver (1971) have also argued that speed reading should be considered a form of skimming due to similarity of eye movements and resulting patterns of comprehension. Hansen's (1975) result on recall of texts in which it was found that speed readers recalled more idea clusters than normal readers but recalled less about each idea is consistent with the notion that speed reading results in comprehension of something akin to the gist of a story. Often the goal in skimming is to obtain the gist of a story.

The similarity between skimming and speed reading eye movements and consequent patterns of comprehension suggests that both methods of reading should be considered as part of a repertoire of reading strategies available to a sophistocated reader. Rather than defining reading as an inflexible process in which all aspects of texts are perceived and comprehended at a uniform rate (Carter, 1977, 1977-78; Coke, 1974; G. R. Miller & Coleman, 1972), reading should also be defined as including processes of skimming, scanning, and slow study rates (Hoffman, 1978). One's reading rate and strategy should be flexible enough to fit one's purpose (Gibson & Levin, 1975).

In a series of studies using payoff schemes to alter readers' goals McConkie and his coworkers (McConkie & Meyer, 1974; McConkie & Rayner, 1974; McConkie, Rayner, & Wilson, 1973) have shown that reading rate is sensitive to payoff contingencies in which rate and comprehension are related. Subjects increased their reading rates up to about 300 wpm and were able to retain information relevant to their purpose in reading while retaining less irrelevant information. Coke (1976) and Samuels and Dahl (1975) have shown that establishing different purposes in reading can affect reading processes and, consequently, reading rate. For example, reading for general information can be done nearly 200 wpm faster than reading for details (Samuels & Dahl, 1975). Readers can adjust their reading rates as they become more familiar with a task (Kershner, 1964) and also can adjust their rates as different parts of the same text become more or less difficult (Rankin, 1971). Flexibility is important in the development of speeded reading as more success in increasing reading rate is achieved by readers who have greater flexibility (Thompson & Whitehill, 1970).

Part of the ability to develop flexible reading strategies depends on being able to rapidly cover a large body of material for particular or gist information. These processes of skimming or speed reading are dependent on the reader's success in suppressing subvocalization (Hardyck & Petrinovich, 1969; Moore, 1962; Wheeler & Wheeler, 1962). On the other hand, reading

difficult material often involves an increase in subvocalization, as if this process were an additional resource for comprehension (Hardyck & Petrinovich, 1969).

From a general point of view, it has been argued that comprehension should be evaluated in light of a reader's purpose (Stauffer, 1962). Reading efficiency, defined as reading rate per unit of comprehension, has been advocated as an alternative to the usual measures which treat rate and comprehension as separate quantities (L. Miller, 1973). While such a measure would encourage flexibility due to inflation from possibly incredible reading rates (despite low comprehension), it masks the basic trade-off of speed and completeness of comprehension. It would seem more suitable to speak in terms of reading rate and strategy and quality or type of comprehension (detailed, gist, etc.). One could then begin to judge the adequacy of certain reading strategies (e.g., speed reading, skimming) with regard to more reasonable criteria of comprehension. Through consideration of a good deal of research it has become clear that reading at incredible rates is, in fact, highly credible if one defines comprehension in the appropriate manner. Skimming and scanning should be considered no more effective reading strategies than speed reading, as long as one realizes that effectiveness is determined by the degree to which a reading strategy is successful in providing the reader with the desired information. The development of clearly defined goals in reading is a requirement of flexible reading (Steinacher, 1971), and it is these goals that specify the type of comprehension that should be achieved. If a reader is interested in obtaining detailed information from a text, then skimming or speed reading will not be useful unless done as part of a more complex reading strategy that uses rapid reading as a prelude to careful reading (Pauk, 1964).

In conclusion, the evaluation of speed reading, skimming, reading improvement programs, and other reading strategies must be done within the specific limits of the pattern of comprehension that is sought by the reader. It is futile to continually argue about the success of a reading strategy until a criterion of comprehension is clearly defined and agreed upon. It is equally futile to argue about whether activities such as speed reading and skimming should be called "reading" (Carver, 1971). Given that these processes are forms of reading strategies, they must be considered within the realm of reading and we must broaden the definition of and scope of concern about reading processes accordingly. The research described below is part of a program designed to extend the range of cognitive research on reading to include issues related to rapid reading strategies.

Rationale for Experiments

Theoretical Issues

The experiments to be reported in this thesis focus on a number of issues relevant to cognitive psychology, artificial intelligence, and education. They are generally concerned with

processes in skimming stories and the relationship of these processes to reading at normal rates. Included as a major objective in studying skimming is the goal of investigating conceptually driven processes. The specific approach to be taken in accomplishing this objective is to consider the role of knowledge structures in guiding conceptually driven processes. These knowledge structures can be thought of as frames (Charniak, 1975; Minsky, 1975) or schemata (Bobrow & Norman, 1975). A number of different schemata are required to successfully read even a relatively simple story and to acquire the new knowledge contained therein (Rumelhart, 1977b; Thorndyke & Hayes-Roth, 1979). These schemata include knowledge of many aspects of the world and the language.

Schank and Abelson (1977) have implemented notions of schemata-based comprehension in the development of natural language processing programs. Not only do they suggest that knowledge structures for common events (scripts) are required for comprehension, but knowledge about plans and goals of actors is also required for understanding and inferences about events in a story. Bower, Black, and Turner (1979) have obtained a number of results that confirm the importance of the part played by scripts in comprehending stories. Subjects exhibit confusion between actions stated in a text and unstated script default actions, indicating that events expected on the basis of prior knowledge are likely to enter into the comprehension process. Script actions will be recalled in script order even if a scrambled presentation order is used. In support of ideas about predictive reading processes, reading time of statements occurring later in a story (presumably after a script has been instantiated) is less than for earlier statements. With respect to the importance of actors' goals, Bower et al. have shown that goal-relevant script deviations are better remembered than standard script actions.

The schema or set of schemata that is of particular interest in the research reported below is related to story structure. Proceeding from Bartlett's (1932) work on story structure schemata a number of current researchers have conducted experiments testing the influence of story structure on recall (C. Frederiksen, 1975a; Kintsch, 1974; Kintsch & van Dijk, 1975; Mandler & Johnson, 1977; Meyer, 1975, 1977; Rumelhart, 1975; Thorndyke, 1977). A general finding is that the importance of a statement or proposition to the gist or theme of a text is predictive of that statement's probability of being included in a recall protocol or summary, with more important statements having higher probabilities. Another experimental result is that if paragraphs of a story are presented in scrambled order (Kintsch, Mandel, & Kozminsky, 1977) or if multiple episodes of a story are interleaved (Mandler, 1978) subjects will recall the story in canonical form, reordering the paragraphs or separating episodes to produce a reasonably coherent protocol which adheres to basic principles of story structure. It is these principles that are thought to be contained in story structure schemata (Anderson, 1977; Crothers, 1972; Rumelhart & Ortony, 1977).

Knowledge about story structure can serve to guide the reader's expectations about what kind of information is likely to be provided at the next point in the story. Such knowledge can also be used to aid in comprehending a poorly (or highly stylistically) structured story by suggesting possible structural categories (e.g., setting, plot event) for otherwise disconnected pieces of information. Finally, story schemata can be used to guide recall of a story and improve coherence of a recall protocol, particularly in those instances where the story was structured in some unusual way (e.g., scrambling paragraph order) but the requisite story information is available from the memory representation. In fact, recalling a poorly organized story may help to change a muddled memory representation that is highly susceptible to forgetting into a more coherent, stable comprehension structure. A process such as this may account for the success of note taking and outlining procedures used in speed reading courses: a rather piecemeal set of information can be organized into a more coherent, though not completely accurate, whole.

With regard to processes in skimming, the availability of a schema for story structure could guide the information selection and comprehension processes. By taking advantage of a story's structure a reader can efficiently locate certain kinds of information. And by knowing what aspects of story structure are most important, the reader can selectively elaborate reading and inferential processes at appropriate points in the story so that the most important and relevant information in the story will earn an appropriate place in the resulting memory representation. In skimming there is not sufficient time for perceptual or cognitive processes to operate as completely as they ordinarily might, and so these processes must be made to work more selectively. I claim that it is knowledge of story structure that plays a large role in this selectivity of processing operations. Therefore, a basic issue to be pursued in the research discussed in the following chapters is the use of story structure schemata in skimming.

Another issue of importance is the question of the kind of memorial representation that is produced by skimming. While the pattern of memory performance will be used to make inferences about the operation of various processes, the nature of the representation itself is of interest. One aspect of memory representations of text that will be investigated is the extent to which the representations include reliable information about the surface structure of the material that was skimmed. The degree of surface memory observed would have implications for how memory of skimmed material is structured and also for how the material was processed. For example, if skimming involves perception of bits of information and higher order inference processes to integrate them, as suggested earlier, one would expect to observe rather poor surface memory. On the other hand, if subjects predominantly restrict themselves to careful reading of select sentences one might expect to observe rather good

memory for surface structure.

The issue of surface memory is one that has been of concern to researchers in natural language processing for some time. Sachs (1967) tested surface memory in what I will refer to as a <u>direct</u> manner by presenting verbatim, paraphrase, and false test sentences drawn from previously processed textual material and requesting subjects to judge whether or not the statements were different from the statements that had actually occurred in the text. Subjects were far more likely to reject false statements than they were to reject paraphrase statements, indicating that memory for meaning was far superior to memory for surface structure. Recent evidence has been obtained by a number of people that suggests that surface memory may be more accurate and durable than originally believed (see Hayes-Roth & Hayes-Roth, 1977). In an experiment by Hayes-Roth and Hayes-Roth (1977) subjects studied sets of unconnected sentences and were tested for recognition of those sentences when presented in original or synonymous form. Using what I will refer to as <u>indirect</u> measures of surface memory, it was found that surface memory was quite reliable. The measures consisted of confidence in recognition judgments, which was higher for verbatim than for paraphrase test statements, and of time required to verify test statements, which was lower for verbatim statements. Kintsch and Bates (1977) found evidence for verbatim memory of lecture material even after two days in an indirect test. Verbatim test statements were recognized as being consistent with material presented in the lecture more reliably than were paraphrase statements. When they used a five-day delay interval and a direct test, Kintsch and Bates found that surface memory was reduced. Thus, it appears that aspects of surface memory may be retained over some period of time. In fact, using an indirect test Kolers (1976) has demonstrated that subjects remember aspects of the typography of text material even a year after the original reading. His test involved measurement of the time required to reread passages originally typed in normal or transformed typography.

In studies of surface memory a distinction has arisen in methods of testing surface memory. It is important to realize the consequences of direct and indirect methods of testing. These consequences have recently been elucidated in a study by Sala and Masson (1978) in which subjects read sets of unrelated sentences then were given a recognition test in which original sentences were presented in verbatim or paraphrase form. If a subject recognized the meaning of a statement he or she was then asked to indicate whether the wording was the same as the original sentence's wording. The indirect test of surface memory consisted of the difference in meaning recognition of verbatim versus paraphrase test statements. The degree of superiority of recognition of the verbatim statements was maintained across test delays of up to seven days and showed no sign of decreasing. The direct test, in which subjects overtly indicated whether wording had changed, showed that memory performance approached chance by the seventh day. The general lesson here is that very different

results may be obtained depending on the method used to test surface memory. It appears that surface memory may be more stable than originally supposed and can exert some influence on processing operations such as recognition of meaning and reading time while explicit access to the surface information itself may become less feasible over time. Thus, the extent to which surface memory is accurate must be evaluated not just in terms of the actual score on a test but also with respect to the method of testing.

In skimming a story processing is not only constrained by elements of time but also by limited resources and limited data (cf. Norman & Bobrow, 1975). With respect to resource limitations, a reader who is skimming a story is faced with a rapid flow of information that often is incomplete with regard to coherence. One piece of information may introduce an event's setting and the next may suddenly describe that event's conclusion, leaving the reader with an enormous inferential computation. This computation requires cognitive resources that draw upon prior knowledge about the topic in an attempt to construct a reasonably likely chain of events making the obtained information comprehensible. At the same time, more pieces of information about other events may be made available and require attention. To the extent that the reader is able to construct a meaningful memory representation of a story, more resources will be required as meaningfulness increases (Britton, Holdredge, & Westbrook, 1979). All these demands on cognitive resources place a heavy burden on a reader while skimming and it is important to determine how these resources are distributed among various of the processing tasks.

Related to the concern about cognitive resources is the issue of how perceptual processes operate in skimming to provide data for higher-order cognitive processes. A relevant hypothesis, which is related to the research on speed reading and reading improvement (Hardyck & Petrinovich, 1969; Moore, 1962; Wheeler & Wheeler, 1962), is that as reading rate increases the reader may no longer be able to rely on phonological decoding processes as these appear to demand time and cognitive resources (Kleiman, 1975; Levy, 1977), neither of which are abundant during skimming. A more efficient strategy would be to use whole word visual identification (Baron, 1977) which involves rapid identification of words based on their visual characteristics rather than phonological decoding of syllables. In fact, fast readers can be distinguished from slow readers on the basis of their speed of accessing overlearned memory codes for visually presented letters (Jackson & McClelland, 1979). Thus, when skimming the reader may be provided with products of data driven processes that are not based on phonological decoding, but a primarily visually accessed code. Some of the consequences of this situation may be very subtle but a more obvious one would be reduced demands on cognitive processing resouces. It seems reasonable to assume, contrary to Bobrow and Norman (1975), that under some circumstances data driven processes require some conscious

attention (Masson & Sala, 1978) and that reduction in the amoun of attention or resources required at the data driven processing level would make available more resources at the conceptually driven processing level. More resource-efficient data driven processing probably is a prerequisite for effective skimming. With a decrease in phonological decoding the reader can rely on more rapid visual perception of fragments of information and can adjust more readily to a rapidly forming comprehension structure.

The results of the research reported here will have implications for text processing models, and there are two such models that will be considered with respect to this research. One of the models is an artificial intelligence program called FRUMP (Fast Reading Understanding and Memory Program) developed by DeJong (1979). FRUMP was designed specifically to skim certain types of newspaper stories for important information. The program processes stories rapidly and will output a summary containing the critical information. There are two issues relevant to the operation of the program and to how people skim: how important information is defined as important and how it is selected. FRUMP uses _sketchy scripts_ which are specially condensed scripts (Schank & Ableson, 1977). The sketchy script contains requests for various types of information and it is this information that is defined as important. The requested information can be provided by the active part of the program which parses the story. Only statements or conceptualizations (Schank & Abelson, 1977, Chapter 1) relevant to an activated sketchy script will be fully processed and incorporated into memory for later summarization. Conceptualizations not directly relevant to the requests of a sketchy script will be parsed only as far as is necessary to determine that they will not satisfy any requests. Not all sketchy scripts in FRUMP's memory are active at once and an important aspect of the program is how relevant sketchy scripts are selected for processing a story. There are three ways in which a sketchy script may be instantiated, depending on how the script is referenced in the article and on any possible information currently being processed. A script may be activated by explicit reference to the script, by an inference implying that the defining action of a script has occurred, or by explicit reference to an action that is part of the script's set of requests. Once a sketchy script has been activated, the program seeks information to fill the script's requests. In this way FRUMP operates in a conceptually driven manner, expecting particular kinds of information relevant to the instantiated sketchy script and readily incorporating that information into memory.

At a general level, the FRUMP program as a model of skimming implies that aspects of the text that are relevant to the goal of reading the text (as defined by sketchy script requests) are far more likely to be incorporated into memory than are statements not related to the goal. As skimming rate increases, FRUMP's parser is used less frequently and important conceptualizations may be missed or incompletely processed. The implication for

human readers is that higher skimming rates should result in virtually exclusive selection of relevant conceptualizations as well as occasional omission of some relevant conceptualizations. Omission of important statements from text processing should increase with skimming rate.

A second model of text processing that is relevant to the research described in the following chapters concerns the extraction of macropropositions and micropropositions from stories (Kintsch & van Dijk, 1978). The model consists of a cyclic processing of text constrained by limitations of working memory. Macro-operators infer and extract from the text macropropositions that together represent the gist or macrostructure of the story which conforms to story structure. These macropropositions are defined by the reader's goals. Recall primarily involves retrieval of macropropositions from memory, as well as construction of relevant inferences. In the case where one is reading for the gist of a text, one is more likely to extract and infer macropropositions than micropropositions as the latter contain details not essential to comprehension of the story's gist. The model is particularly concerned with establishing referential coherence in the comprehension structure. Thus, as the model cycles through the text, a number of propositions are entered into the processing at each cycle. As a new cycle begins some number of propositions are held over from the last cycle in order to make a referential connection with some argument(s) of the newly processed propositions. If no such connection can be made more complex processes such as inference and reinstatement of earlier propositions must be called into play. Consequently, a poorly connected text will be difficult to process.

When a reader skims a text it is likely that the distinction between macropropositions and micropropositions will become more prominent, with macropropositions being more fully processed both in terms of information provided by data driven processes and at the level of verification or rejection of conceptually driven conjectures about what the text says. On the other hand, it is also possible that skimming does not allow one to fully exploit the distinction between macropropositions and micropropositions. That is, the ability to develop an accurate representation of a text's gist or macrostructure may deteriorate to a state in which one has only a mixed collection of very important statements and trivial details. Another interesting issue is whether the cyclic processing of text proposed in the model is maintained when skimming or if some other processing strategy for construction of a coherent representation is used.

The experiments reported below were designed to deal with the general issue of how people skim stories and with the theoretical expectations based on interpretations of DeJong's (1979) FRUMP program and Kintsch and van Dijk's (1978) text processing model. I should emphasize, however, that the experiments are not meant to be direct tests of either of these theoretical formulations. FRUMP and the macrostructure model were

discussed here in order to provide an indication of how skimming processes might be formalized in processing models. FRUMP was not meant to be a psychological model; it is an artificial intelligence program. People may skim stories in ways totally unrelated to FRUMP's behavior, or there may be striking similarities between human processing and FRUMP's operation. The extent of these similarities and differences will determine FRUMP's potential for being a psychological model. Similarly, Kintsch and van Dijk's model was not devised as a description of skimming processes, but the model does have potential ramifications for ideas about what kinds of comprehension processes might be involved when a reader skims. In fact, the model may be so general as to be made to simulate skimming behavior simply by making appropriate adjustments in its parameters. If not, it may be feasible to suggest changes in the model that would result in a related processing model more specifically designed to deal with skimming processes.

Experimental Methodology

The primary questions addressed in this research are concerned with how a text is processed when it is skimmed and with the characteristics of the resulting memory representation. Processing details and memory representations can be assessed by tests of what subjects remember about a text: the characteristics of what is remembered have implications both for what is in memory and for how that information was processed. The use of memory performance as an indicator of comprehension processes must include an acknowledgement that comprehension processes have been tapped only indirectly. Ortony (1978) has argued that separate theories of comprehension and memory should be developed, and has pointed out that not everything that is understood is remembered and that not everything that is remembered is understood. I would argue, however, that while investigations of comprehension processes by direct means (e.g., Shebilske et al., 1977) is highly desirable, development of comprehension theories probably will not be completely successful without consideration of realistic and compatible memory theories. As was discussed earlier, processes that interact in fundamental ways with other processes, as comprehension and memory clearly do (Rumelhart, 1977b; Rumelhart & Ortony, 1977), must be understood in the context of those processes with which they interact. Study of a process occurring in isolation may lead to conclusions that badly misrepresent the true state of affairs. In the case of memory and comprehension, while it may be that not all that is remembered is understood, it is apparent that memory and comprehension are interdependent (Bransford & McCarrell, 1974).

Observations about aspects of memory representations, then, do have promise for contributing to the understanding of comprehension processes. An example of this contribution that is used in the studies reported later in this thesis involves testing memory for surface form. It was suggested that skimming may involve a process of obtaining partial information from

sentences and then drawing inferences to relate those pieces of knowledge into a coherent structure. To the extent that comprehension during skimming relies on inferential processes, little processing of surface structure is likely. A consequence for memory of a text is that little knowledge about surface structure will be represented. Surface memory, then, was tested in a number of experiments in a direct manner by asking subjects to judge whether test statements drawn from texts that they had read were worded in the same way as when they had occurred in the texts. A direct test of surface memory allows measurement of the extent to which subjects can consciously access their knowledge of surface structure. If subjects primarily process surface structure without conscious attention it is probable that they will remember very little of the surface features of a text (Kolers, 1975a; Masson & Sala, 1978; Schneider & Shiffrin, 1977; Shiffrin & Schneider, 1977). But if attention is devoted to the processing of surface structure when reading normally or skimming, subjects should display some ability to detect changes in surface form of test statements.

A memory for meaning test was also used in the experiments described in the following chapters. In the first two studies a recognition test was used to provide a sensitive test of information available from normally reading and skimming stories. One of the most important issues considered in designing the meaning test items was the use of story structure schemata in reading. Test statements were derived that represented different levels of importance with respect to a story's macrostructure. Importance ratings were used as a basis of determining which statements contained propositions that would correspond to macropropositions and which statements contained information corresponding to micropropositions. Johnson (1970) and Caccamise and Kintsch (1978) have shown that statements rated as more important to a text's gist are more likely to be recalled and recognized. Test statements that were rated as important in the experiments reported here will be called <u>macrostatements</u> and unimportant statements will be called <u>microstatements</u> so that the correspondence with models of story macrostructure (Kintsch, 1977; Kintsch & van Dijk, 1978; van Dijk, 1977) will be preserved. Other test statements were inferences that could be made on the basis of information provided in the text. Inferences were chosen to be of high importance in order to increase the likelihood that subjects would have made the inferences during reading (Goetz, 1977). Inferences that are too easily predicted by explicit information are likely to be processed superficially and represented in memory in an unstable fashion (Spiro & Esposito, 1977), so the selected inferences also required the availability of a number of pieces of information and cognitive effort in connecting those pieces into a coherent knowledge structure. The ability to recognize such inferences would indicate that during reading (normal or skimming) subjects were able to comprehend and possibly combine the information necessary to form important conclusions. To the extent that subjects are

able to recognize macrostatements and inferences more often than microstatements, evidence for the role of story structure schemata or macrostructures will be obtained. For example, if subjects are unable to make good use of story structure in skimming and do not selectively process macrostatements nor make important inferences, then recognition performance on the different types of statements should not differ.

Stories of two different types were included in the first two experiments in order to observe effects of skimming on the operation of two kinds of story schemata. One set of stories was composed of narratives, a commonly used type of text in cognitive research. The other type was newspaper stories. This type was selected so as to provide a comparison of obtained data with FRUMP and because newspaper stories have a rather different structure from narratives. In newspaper stories the important information ordinarily occurs early in the story and this may affect the strategy or effectiveness of skimming. Narratives may place important information at any point in the story, although information given early is important in establishing the basic nature of the narrative. Moreover, newspaper stories and narratives contain different kinds of information with narratives emphasizing any number of themes such as character development, while newspaper stories concentrate on presentation of newsworthy information. Readers are usually familiar with newspaper story structure and are likely to recall important more than unimportant information in a newspaper story (Thorndyke, 1979).

In the third experiment a recognition test was used to investigate skimming processes that occur when subjects are instructed to read for information relevant to a specified goal. Rather than reading to obtain the gist of a story, subjects were requested to read from a particular perspective. The degree to which relevant and irrelevant information is recognized provides an index of how selectively subjects are able to process a text on the basis of the assigned perspective.

A variable included in the second and third experiments was the way in which stories were typed. Normal typography was used for half the subjects and an alternating case typography was used for the remaining subjects. It was hypothesized that reading alternating case would be difficult if whole word visual identification processes were in effect, due to the unusual nature of each word's appearance (Baron, 1977). Thus, if skimming involves the use of whole word visual identification more so than does normal reading, then one would expect to see a greater detrimental effect of alternating case for subjects who skim than for those reading at a normal rate.

The final experiment in this report involved collection of recall protocols from subjects who read stories at different rates of speed. These protocols were simulated using a version of Kintsch and van Dijk's (1978) text comprehension and production model that has been developed by J. Miller and Kintsch (1979). Success in adequately simulating recall protocols taken from subjects who skimmed the stories would support the general

validity of the model and would point out those aspects of processing that are common to normal and skimming rates of reading. The parameter values associated with successful simulations would also provide useful information about how reading processes change as reading moves from normal rates of speed to skimming.

CHAPTER II

EXPERIMENT 1

The major purpose of Experiment 1 was to establish the validity of methods of studying skimming behavior. In this first experiment a naturalistic setting for skimming stories was sought so that test performance could be compared to performance in experiments that included more stringent experimental control during reading. In particular, subjects were provided with stories to read and were requested to read either at the rate they would normally use for full comprehension or at a rate that they would use to skim a story for important information. Two drawbacks to this method are that different subjects can, and no doubt will, read at widely different rates when reading for any particular purpose (e.g., full comprehension), and that in order to keep track of these differences and to be sure that subjects who are asked to skim do, in fact, read faster than those asked to read for full comprehension it is necessary to obtain reading time measures for each subject. In Experiment 1 reading times were collected by having subjects start and stop a stopwatch at points coinciding with the beginning and completion of reading a story. In this way each subject's reading time for each story could be measured and a subject would be allowed to read in a manner and at a rate to which he or she was accustomed. Two practice stories were read before the stories that were to be tested so that subjects could adjust to the requirements of the

reading and timing tasks. The use of the stopwatch might have led to some distraction on the subject's part and in order to determine subjects' attitudes toward the experimental situation, including use of the stopwatch, a questionnaire was administered at the end of the session which inquired as to subjects' reactions to the experimental setting. The questionnaire also requested information about subjects' skimming behavior outside the laboratory with the expectation that such information could be informative with respect to usual skimming habits and the degree to which the sample population of subjects uses skimming as a reading technique.

Another aspect of the experimental methodology that Experiment 1 was meant to validate was the recognition testing procedure. Before moving on to more complex experimental designs it was important to establish that the recognition tests of meaning and surface structure memory were sensitive measurement devices. It was also highly desirable to demonstrate that the various types of test statements (inference, macrostatement, and microstatement) did, in fact, reflect different levels of macrostructure information (Kintsch & van Dijk, 1978). Moreover, Experiment 1 was intended to establish a basic pattern of test performance against which performance in more highly controlled experiments could be compared.

The recognition test consisted of two parts, the first of which tested surface memory and the second tested memory for meaning. The surface structure test included only statements that had appeared in the stories and their paraphrase versions and subjects were required to indicate whether a test statement's wording was the same as that used in the story. The first and last critical stories read by subjects were tested, with the last story tested first, so that the effects of a delay interval on surface memory could be assessed. The meaning test included all three statement types drawn from the stories not tested in the surface test. In both tests, added sensitivity was obtained by measuring the time taken to respond to each test statement.

A critical experimental outcome that would have implications for the process and effectiveness of skimming is the nature of the interaction between rate and statement type in the meaning test. Previous research (Caccamise & Kintsch, 1978; Johnson, 1970; Kintsch & Keenan, 1973; McKoon, 1977) suggests that the more important test statements (inferences and macrostatements) should be recognized more often than the unimportant ones (microstatements). Differential recognition would imply that a story has been processed with respect to a story structure schema and that a macrostructure (Kintsch & van Dijk, 1978) has been formed and represented in memory. A key question is whether or not the processes involved in forming a structured memory representation of a story are operative during skimming, and, if so, to what extent their operations are similar to those used in normal reading. Assuming that differential recognition of statements due to differing degrees of importance is observed among subjects reading at normal rates, indicating that a story

structure schema has guided reading, there are three possible effects that could be observed among subjects who skim stories. First, if the processes responsible for forming a macrostructure do not operate effectively during skimming, then no effect of statement types would be observed in the recognition of meaning test. Subjects would be equally likely to recognize inferences and macro- and microstatements. A second possibility is that reading processes operate in skimming similar to the way they do in normal reading, and that one should observe better recognition of important statements than of unimportant ones, although overall recognition performance may decline (i.e., no statistical interaction between rate and statement type. That is, enough of the essence of macro-operations is preserved so that important information is selectively processed but sampling of information from the text is generally reduced. Finally, when skimming the process of extracting important information may be emphasized to such a degree as to cause subjects to form a macrostructure that contains an unusual proportion of macrostatements relative to detail information. The corresponding result on the recognition of meaning test would be an increased difference between important and unimportant statements relative to the difference observed among subjects reading at their normal rates. This is the type of interaction that McConkie et al. (1973) found among subjects reading up to 300 wpm. It is not certain that such an effect would occur when faster rates such as skimming are used, although a common impression of the purpose of skimming is that skimming is used to derive the important information from a story and skip over unimportant details. If this selective processing strategy is feasible and highly accurate in its selection decisions, then an interaction between rate and statement type should occur in which the importance effect is greater for those subjects who skim than for those who read at normal rates.

A result related to the importance effect is recognition of (important) inferences relative to macrostatements. While both are important to a story's macrostructure, only the macrostatements are explicitly stated in a single location in the story. Nevertheless, there is abundant evidence arguing in favor of the belief that readers do not process only explicit information but also draw and represent in memory relevant inferences based on explicitly presented information and general knowledge (C. Frederiksen, 1975b; Kintsch 1974; Reder, 1979; Thorndyke, 1976). For example, McKoon and Keenan (reported in Kintsch, 1974, Chapter 8) found that time required to verify inferential statements was not reliably greater than verification time for explicitly presented statements once the effectiveness of surface structure information in aiding verification declined. Not only are inferences represented in text memory, they are apparently integrated with explicitly presented information (Bransford & Franks, 1971; Moeser, 1976, 1977) and can be used to aid in the verification (Reder, 1979) or recall (Masson, 1979) of explicitly presented information. If an inference is an important aspect of a story's macrostructure it should be included in that

macrostructure representation and recognized just as well as macrostatements. In order to verify an inference, however, a subject must have processed the relevant pieces of information (and possibly have already drawn the inference) as part of the macrostructure while reading (Hayes-Roth & Walker, 1979; Reder, 1979). The process of skimming must allow the reader to efficiently select and organize sometimes physically separated but semantically related information to form an important conclusion. To the extent that such processing is practicable, recognition of inferences should be about as reliable as recognition of macrostatements and should be almost as rapid if subjects have formed adequate configurations of the underlying facts of an inference (Hayes-Roth & Walker, 1979).

When skimming, the processing of surface structure is hypothesized to play a different role than when reading normally. Subjects are less likely to fully process a sentence's complete surface representation when skimming. It is more likely that subjects will select certain relevant aspects of a sentence for further semantic processing with little attention paid to the surface features of even the selected portions of a sentence. Less processing of surface information is expected on the basis of highly active conceptually driven processes which can make accurate predictions at least at general levels. Complete checking of the surface details is not necessary, nor is it feasible, when skimming. Decreased processing of surface structure is expected to produce a general decline in recognition of surface information. Whereas subjects reading at normal rates may remember more about the surface structure of more important statements, a much weaker effect, if any, would be expected for subjects who skimmed the stories.

Two types of stories were used in an attempt to obtain more generalizable results and to determine which aspects of comprehension and/or memory representation are susceptible to effects of processing in the context of different story structure schemata. Narrative and newspaper stories were chosen as two distinctly different and familiar styles. In his recent work, Thorndyke (1979) has suggested that different story schemata may exist for different types of stories. If this is so, it is especially important to discover those aspects of normal reading and skimming processes that are generalizable across different kinds of stories.

Method

Subjects

Subjects in the experiment were 72 students at the University of Colorado participating in partial fulfillment of an introductory psychology course requirement. Thirty-six subjects were assigned to each of the two reading rate conditions according to their order of appearance at the laboratory.

Materials

A 500-word narrative story was selected for use in obtaining a preliminary reading rate score for each of the subjects. Four

other report style narratives of about 1000 words each and four newspaper stories, each of about 400 words in length, were selected for use in the experiment proper. Three stories of each type were used as critical stories and a set of test statements was developed for each critical story. A set of test statements was developed by selecting from a story, statements (whole or partial sentences) that intuitively appeared to vary in their degree of importance for the general meaning of the story. Some statements were meant to be important and others unimportant. A third type of statement was developed by drawing reasonable and apparently important inferences from specific points in the story and writing these inferences in proper English. Six to eight statements of each type were developed for each story.

Intuitions about the relative degree of importance of the selected statements were evaluated by having a separate group of subjects from the introductory psychology subject pool provide importance ratings. Twenty-seven subjects read each critical narrative and then evaluated each statement in its corresponding set of selected statements for importance to the general meaning of the story. A six-point rating scale was used where a rating of one meant a statement was unimportant and a rating of six meant a statement was very important. Similarly, a set of 26 new subjects read each newspaper story and evaluated the corresponding selected statements.

Based on the mean importance rating for each statement in a set of selected statements, a set of test statements was chosen.

The three most highly rated statements that explicitly appeared in the story were chosen as true macrostatements for the test set. The three explicit statements receiving the lowest ratings were chosen as true microstatements. The three most highly rated inferences were included in the test set as true inferences. Selections sometimes were not solely determined by ratings but also by degree of similarity between statements. For example, if two highly rated statements represented very similar information, only one was selected for the set of true statements.

From the remaining statements, the three explicit statements with the highest ratings were selected for the false macrostatement test set. The three explicit statements with the lowest ratings were chosen for the false microstatement test set, and the three most highly rated inferences were selected for the false inference test set. Each of the six explicit statements selected for the false test set was rewritten in two new versions. One version was a paraphrase of the original statement that altered the lexical and syntactic structure of the original statement. Statements with this version were used in the surface memory test. The second version was syntactically similar to the original statement but was altered in some way so as to make the new statement semantically false with respect to the original statement. Statements with this version were used in the memory for meaning test. Each of the three inferences selected for the false test set was rewritten in a new version that semantically differed from the original inference and implied something that

would be incorrect with respect to what had been stated in or could be inferred from the story. These inference versions were used in the memory for meaning test.

Thus, 24 test sentences were developed for each critical story. Three macrostatements, three microstatements, and three inferences formed the true set. Three statements of each type were used as false statements in the meaning test and three statements of each explicit type were used as paraphrases in the surface test. An example of a narrative story and the corresponding set of test sentences and an example of a newspaper story and its test statements are presented in Appendix A. The mean importance rating and mean number of words for each type of test statement are shown in Table I. Note that inferences are almost as highly rated as macrostatements and that newspaper test statements are generally shorter than the narrative test statements. Importance ratings for the two kinds of false statements (new meaning and paraphrase) are the same since the ratings are based on the original explicit version. No paraphrase versions of inferences were developed as inferences were not included in the surface test.

The six critical and two practice stories were typed single spaced, right and left justified, on separate pages and headed by appropriate titles. The stories were arranged in booklets so that the two practice stories occurred first and the critical stories next. For three of the booklets the practice narrative occurred first, followed by the practice newspaper story. The three

TABLE I

MEAN RATED IMPORTANCE OF AND NUMBER OF WORDS IN TEST STATEMENTS USED IN EXPERIMENTS 1 AND 2

Story	Test Version	Importance Ratings			Number of Words		
		Statement Type[a]					
		Inf.	Mac.	Mic.	Inf.	Mac.	Mic.
Narrative	Old/Verbatim	4.69	4.81	3.05	16.8	15.9	15.4
Narrative	New	4.00	4.26	3.46	15.6	18.0	16.0
Narrative	Paraphrase	--	4.26	3.46	--	16.4	16.9
Newspaper	Old/Verbatim	4.87	4.97	2.68	12.2	13.0	13.2
Newspaper	New	4.11	4.21	3.21	14.0	13.6	14.3
Newspaper	Paraphrase	--	4.21	3.21	--	13.4	14.8

[a]Inf. = Inference; Mac. = Macrostatement; Mic. = Microstatement

critical narratives, then the three critical newspaper stories. Each of these three booklets represented one ordering of the critical stories determined by a 3 X 3 Latin square yoked counterbalancing of the critical narratives and newspaper stories. The other three booklets were similarly arranged except that the narratives and newspaper stories switched positions so that the newspaper stories occurred first. A blank sheet of colored paper was inserted before each story.

<u>Design</u>

For the meaning test a 2 X 2 X 3 mixed factorial design was used. Reading rate (normal and skimming) was a between-subject variable, while story type (narrative and newspaper) and statement type (inference, macrostatement, and microstatement) were within-subject variables. For the surface test a 2 X 2 X 2 mixed factorial design was used. All three variables were the same as in the meaning test except that the statement type variable included only two levels (macrostatement and microstatement).

<u>Procedure</u>

Subjects participated in small groups ranging in size from one to six. As each subject entered the laboratory he or she was instructed to read the 500-word practice narrative story at the rate he or she would use in reading for full comprehension. Subjects were provided a copy of the story and a stopwatch. They were instructed to start the watch when they began reading and to stop the watch when they finished reading. The reading time for each subject was recorded by the experimenter who then reset the stopwatch for each subject and collected the stories. Once all subjects in the group had finished the preliminary reading task, each subject was given a booklet containing instructions and a set of practice and critical stories. For subjects in the normal reading rate condition the experimenter instructed subjects to read at their individual normal rates of speed and to try to comprehend the meaning of the stories. Subjects were asked to try to maintain a constant reading rate. Subjects in the skimming condition were similarly instructed and were asked to use a skimming rate that they might ordinarily use on their own and to try to extract the important information from each story. The experimenter then had the subjects read the instructions on the cover of the story booklet. These instructions explained that the subject would be asked to read a set of stories at the specified rate of speed and would then be tested on the material contained in the stories. Reading for meaning rather than memorization of material was stressed. The instructions also described the types of stories contained in the booklet and explained that the experimenter would record the time taken to read each story as was done for the preliminary story.

After reading the instructions in the booklet and having any questions answered, subjects went on to read the stories, starting and stopping a stopwatch in accordance with the onset and completion of the reading of each story. After reading each story the subject turned to the next blank page. The experimenter

recorded the reading time for each subject and story, reset the watch, and the subject proceeded to the next story.

Once all subjects completed reading the whole set of stories, they were told that they would next be given a short test on their knowledge of the stories they had read. Each subject was assigned to a CRT terminal in a room adjoining the main laboratory room. There were two terminals facing opposite walls in each room. The instructions for the test and the test items were presented and the subjects' responses and reaction times were recorded using the DTES system developed by the CLIPR laboratory at the University of Colorado (Spear, Overgard, & Christian, 1975). The subject began the test by viewing the first part of the instructions on the terminal, then advanced through the instructions by pressing a button on a button box located in front of the terminal. The first part of the test involved a test of surface memory. The instructions on the screen informed the subject that the first part of the test would be concerned with how well he or she remembered how statements from the stories were worded. The subject was told that some statements would be presented in their original verions, while others would be worded differently. It was also explained that for some statements pronouns would be replaced by the original proper nouns so that the subject would have a better idea of what the statement was about, but that if this was the only change the statement should be considered as being presented in its original form. The subject was instructed to keep his or her left index finger on the button marked "SAME," which was to the left of center on the button box, and his or her right index finger on the button marked "DIFF" for "different," which was to the right of center. If a test statement was worded in what the subject thought was the original version, the subject was told to respond by pressing the button marked "SAME." If a test statement appeared to be worded differently, the subject was to press the button marked "DIFF." The instructions requested the subject to respond to each item as rapidly as possible but cautioned that it was very important to be reasonably sure of the correctness of a response before making it. The subject also was told that the test would have two sections, one devoted to statements from each of two different stories.

Each set of test statements was preceded by a display that presented the title of the appropriate story. When the subject was ready to proceed he or she initiated the test by a button press. The test statements appeared individually in random order and remained on the screen until the subject responded. The screen was blank for about 1 sec between each test statement. The first story tested was the story that appeared in the last position in the story booklet, and the second story tested was the first critical story read by the subject. Over the course of the experiment each of the six orderings of stories was read by a different set of six subjects in each of the reading rate conditions. Thus, each subject was given the surface memory test on one narrative and on one newspaper story. Also, for half of

the subjects the narrative story was tested relatively soon after reading, while the newspaper story was tested after a number of intervening stories had been read. The reverse was true for the other half of the subjects. The ordering of stories in the booklets also meant that the specific stories tested for surface memory were counterbalanced across subjects.

After the subject finished the surface memory test, instructions for the meaning test appeared on the screen. The instructions indicated that the next part of the test phase would be concerned with how well the subject could remember the meaning of the stories that were read. The subject was told that some of the statements would represent information contained in or that could be inferred from a critical story, while other statements would contain information that would not be consistent with what was said in a story. The subject was to respond to the former type of statement by pressing the button marked "SAME" and to the latter by pressing the button marked "DIFF." Instructions about speed, accuracy, and placement of fingers were the same as those given for the surface test. The subject was also told that the test would contain four sections, one for each of four different stories, with the title of the appropriate story appearing before each set of test statements. The procedure then followed the same steps as those described for the surface test. The four stories tested were the second, third, fourth, and fifth critical stories, none of which were represented in the surface test for that subject. The four stories were tested in the order in which they had been read. Thus, each subject was tested for memory of meaning on two stories of each type (narrative and newspaper).

When the subject finished the meaning test a message appeared on the screen requesting him or her to return to the main laboratory room. The subject then filled out a questionnaire concerning skimming habits outside the laboratory, a comparison of usual reading speeds and those used in the laboratory, an estimate of how successful the subject was in extracting important information from the critical stories, and history of possible training in speed reading.

Results and Discussion

In all experiments to be reported, a statistical significance level of .05 has been adopted. Except where noted, all significant effects are reliable at least at that level.

Reading Time

In order to demonstrate that subjects in the skimming condition were, in fact, reading the critical stories at a faster rate than subjects in the normal reading rate condition, an analysis of variance of reading times was calculated. Each subject's mean reading time for each story type was entered into the analysis. The reading rate for narratives and newspaper stories did not differ ($F < 1$), while there was a reliable effect of reading rate condition with subjects in the skimming condition reading faster (382 wpm) than subjects in the normal condition (232 wpm), $F(1, 70) = 67.62$, $MS_e = 11932$.

It is clear that instructing some subjects to skim resulted in faster reading rates than that exhibited by subjects instructed to read at a normal rate of speed. These mean reading speeds should be interpreted cautiously, however. A good deal of variability in reading speed was observed as the skimming rate for one subject might have been slower than the normal rate for another. What is important is how a subject's reading rate in the experiment (normal or skimming) compares to what the subject does when reading on his or her own. The questionnaire administered to each subject at the end of the experimental session was designed to deal with this and other issues.

Questionnaire

The questionnaire that each subject filled out at the end of his or her session addressed two basic issues. The first issue was concerned with subjects' skimming behavior outside the laboratory and the second was concerned with subjects' reading performance in the experiment. With respect to skimming outside the laboratory, subjects were first asked to estimate the percentage of their total reading time that is taken up by skimming. The mean estimate for time spent skimming was 32%, with a range of 0-90%. Next, the subjects were asked to indicate what type of material they skim on their own. The most commonly mentioned materials were newspapers and magazines, and 44% of the subjects claimed that they skimmed at least one of these two kinds of material. Interestingly, 12% of the subjects stated that they skimmed some parts of their college work materials (notes, texts, etc.). Other materials mentioned idiosyncratically included lists, numbers, familiar material, and material read for personal enjoyment.

To obtain an estimate of how successful subjects believe they are when they skim material on their own, they were asked if they felt that they are generally able to extract the important information in a story when they skim. Replies were interpreted on a four-point scale where a rating of zero meant that the subject felt that he or she was not able to select important information while skimming, and a rating of three meant that the subject felt that he or she was quite successful. A mean rating of 1.96 was observed for the subjects in Experiment 1, indicating a relatively high degree of confidence.

Subjects also were asked if they had had any training in speed reading and 15% of them indicated that they had received some form of training. The nature of this training ranged from high school courses in reading improvement to current commercial speed reading programs. The subjects who had received training displayed a greater degree of confidence in their skimming abilities as their mean evaluation rating was 2.36 compared to the overall mean of only 1.96 for all subjects in Experiment 1.

With respect to the experiment and the reading rate that the subject was instructed to use, each subject was asked if he or she was able to read the experimental materials as he or she ordinarily would read such stories outside the laboratory. For subjects in the normal reading rate condition, 61% indicated that

they were able to read as they normally would, 28% indicated that they read at least slightly faster than they ordinarily would, and 11% claimed that conditions forced them to read at a somewhat slower pace than normal. For subjects in the skimming condition, 42% stated that they were able to skim as they would on their own, 47% indicated that they read faster than they ordinarily would skim, and 11% stated that they read at a rate below their usual skimming rate. Subjects in the skimming condition also were asked if they felt that they were able to extract the important information from the stories read in the experiment. Responses were interpreted using the same scale as that used for the question about success in extracting important information when skimming outside the laboratory. The mean rating assigned to these responses was 2.14.

In general, the questionnaire revealed that subjects spend a significant portion (almost one third) of the time they devote to reading on skimming through material. Skimmed material ordinarily consists of newspapers, magazines, and other information for which an individual is not formally held responsible. It is also interesting to note that a nontrivial proportion of the subjects use skimming when dealing with some aspect of their college work. This fact, taken in conjunction with a relatively high degree of confidence demonstrated by subjects in their ability to select important information while skimming, indicates that college students believe skimming to be a viable and practiced form of reading. Few subjects have received any formal training in speed reading and, therefore, appear to have developed the ability to skim through experience with normal reading.

The subjects appeared to have reacted well to having their reading processes placed under experimental observation. For the most part, subjects claimed to be able to maintain their usual reading paces (normal or skimming) with a tendency, especially among subjects in the skimming condition, to read faster than usual. Relatively few subjects claimed to have had their reading rates slowed by being in the experiment. Subjects in the skimming condition maintained their relatively high degree of confidence in their ability to skim for important information when questioned about how successful they felt in their attempts to skim the experimental materials. The subjective reports about being able to maintain usual reading habits is important in two respects. First, they make generalization of experimental results to usual reading conditions more credible. Second, the reports indicate that a reasonable degree of success may have been achieved in placing skimming processes under experimental observation, supporting the validity of the reading time results described above. Not only did the subjects in the skimming condition read faster than those in the normal condition, but they claimed that they were, in fact, skimming as they would on their own. It is important to be able to establish with some degree of confidence that subjects were reading normally or skimming as they ordinarily might before moving on to a discussion of the more theoretically meaningful results of the

experiment. We are now in a position to interpret the remaining empirical results in a context that is applicable beyond a laboratory setting and that is representative of cognitive processes involved in normal reading and skimming behavior of college students.

Memory for Meaning

Recognition performance. Subjects' ability to recognize the meaning of test statements that explicitly appeared in or could be inferred from one of the critical stories was characterized in two ways. First, the proportion of hits and false alarms for each statement type for both kinds of story were calculated for each subject. The mean proportion of hits and false alarms are presented in Table II. The second characterization of recognition performance involved calculating d' scores for every pair of hit and false alarm rates for each subject. In calculating the d' scores, when a hit rate of 1.0 was encountered it was transformed into a more realistic value by using the estimate $1-(1/(2N))$, where N was the total number of hits possible. False alarm rates of 0.0 were transformed into realistic values by using the estimate $1/(2N)$, where N was the total number of false alarms possible. These estimates have the desirable property of providing a probability value that lies between 1.0 for hits or 0.0 for false alarms and the probability obtained when just one miss or false alarm, respectively, occurs. Furthermore, a hit rate of 1.0 or a false alarm rate of 0.0 should be considered a more reliable result when more items are involved (i.e., when N

TABLE II

MEAN HIT AND FALSE ALARM RATES ON MEANING TEST IN EXPERIMENT 1

Story Type	Reading Rate	Statement Type[a]					
		Inf.		Mac.		Mic.	
		Hit	FA[b]	Hit	FA	Hit	FA
Narrative	Normal	.89	.12	.90	.12	.87	.24
Narrative	Skim	.85	.19	.86	.16	.78	.30
Newspaper	Normal	.88	.12	.87	.12	.78	.23
Newspaper	Skim	.76	.21	.86	.20	.71	.33

[a]Inf. = Inference; Mac. = Macrostatement; Mic. = Microstatement

[b]FA = False Alarm

is large). The term 1/(2N) reflects this consideration, assigning higher hit rate estimates (approaching 1.0) and lower false alarm estimates (approaching 0.0) as N increases. The d' scores for each subject were used in the formal analyses of recognition performance as they reflect a consideration of both hit and false alarm rates.

The mean d' scores for recognition of statements in the meaning test are presented in Figure 1. An analysis of variance of the d' scores was carried out, including story and statement type as within-subject variables and reading rate as a between-subject variable. The analysis indicated that subjects in the normal reading rate condition had higher d' scores (1.97) than subjects in the skimming condition (1.60), F (1, 70) = 14.40, MS_e = .870. Subjects were better able to recognize statements from the narratives (1.87) than from the newspaper stories (1.70), F (1, 70) = 8.39, MS_e = .830. The effect of statement type was also reliable, F (2, 140) = 41.84, MS_e = .339. A Neuman-Keuls test indicated that recognition of inferences (1.91) and of macrostatements (2.01) did not reliably differ, while recognition of either of these two statement types was reliably greater than recognition of microstatements (1.43). None of the interactions were significant (Fs < 1.5).

The reading time results indicated that subjects instructed to skim did, in fact, read faster than subjects instructed to read normally and this effect was further reflected by reduced recognition performance among subjects in the skimming condition.

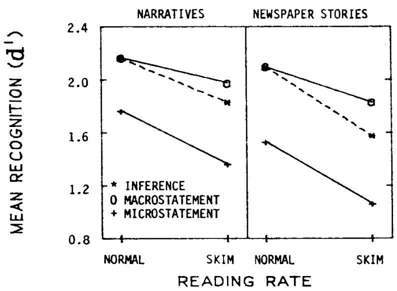

Figure 1. Mean recognition performance on meaning test in Experiment 1.

The faster reading rates led to reduced comprehension, as measured by the recognition test used here, and as expected from the review of research presented in Chapter I. More interesting than the quantity of comprehension are the qualitative aspects of comprehension. One of the results that represents the quality of comprehension is the statement type effect. As predicted by theories of macrostructure (Kintsch & van Dijk, 1978; Mandler & Johnson, 1977; Thondyke, 1977), the important inferences and macrostatements were more reliably recognized than microstatements. The failure to find a significant difference between recognition of inferences and macrostatements attests to the fact that development of a macrostructure is not completely stimulus-bound, but relies on the reader's background knowledge and ability to draw relevant conclusions (C. Frederiksen, 1975b; Hayes-Roth & Walker, 1979; Kintsch, 1974; Kintsch & van Dijk, 1978).

No interaction between rate and statement type was observed, indicating that the processes responsible for producing a macrostructure representation of a story when reading at normal rates were not completely abandoned when skimming. Some essential aspects of macro-operations were in effect during skimming, but subjects were not able to strictly adhere to the usual macroprocessing operations as recognition performance on both important and unimportant statements declined significantly and to the same extent. Were macro-operations in full effect, one would expect to observe no decline in recognition of macrostatements and inferences while recognition of microstatements would decrease severely. Nor, however, was the distinction between important and unimportant information diminished to any extent. It is clear that subjects who basically have taught themselves to skim can effectively select important information from stories. Memory representations of stories that are skimmed are not comprised of equal amounts of important and unimportant information. At least a rough, incomplete macrostructure is preserved when subjects skim at moderate rates: processing does not deteriorate to a random selection of pieces of information. The reliably greater recognition of inferences than of microstatements is maintained even when stories are skimmed, suggesting that while skimming subjects are able to draw together enough partial information found in different parts of a story to form at least some of the important conclusions.

Even though recognition of statements drawn from narratives was greater than recognition of newspaper story statements, exactly the same pattern of reading rate and statement type effects were observed for both kinds of story. Thus, while subjects may be more interested in or have more background knowledge about narratives, their ability to form story macrostructures when skimming or reading at normal rates is quite good when reading newspaper stories as well. Whatever differences exist between newspaper and narrative story schemata, subjects are able to use these schemata to shape a reasonably accurate

macrostructure representation of either kind of story. The processes involved in forming a macrostructure appear, therefore, to be general across a number of different kinds of stories.

__Reaction time for hits__. The time taken for a subject to respond to each test statement was recorded by the computer system used for the test phase of the experiment. For each subject a mean reaction time was calculated for each combination of response and statement type. These means were adjusted in the following way to account for unduly long reaction times. If a subject's mean reaction time was based on one, three, or more responses (depending on how many times the subject responded in a certain way to statements of a particular type) and if that mean exceeded 15 sec, then the mean was changed to 15 sec. If a mean was based on exactly two responses and if one or both values exceede 15 sec, the mean was recalculated with the excessive value(s) adjusted to 15 sec. These adjustments were made for all reaction time categories in Experiment 1. In cases where a subject made no responses of a particular type that subject was assigned the mean reaction time of the appropriate type for his or her condition.

The mean time taken to correctly verify each statement type in Experiment 1 is presented in Figure 2. An analysis of variance of the reaction times for hits was carried out with reading rate (normal and skimming) as a between-subject factor and story type (narrative and newspaper) and statement type (inference, macrostatement, and microstatement) as within-subject factors.

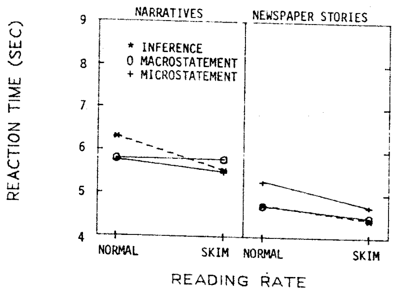

Figure 2. Mean reaction time on hits on meaning test in Experiment 1.

Only two effects were reliable. Newspaper story statements were verified more quickly (4.728 sec) than narrative statements (5.792), F (1, 70) = 48.72, MS_e = 2.506. Story and statement type interacted, F (2, 140) = 3.39, MS_e = 1.529, indicating that the pattern of reaction time to the different statement types varied as a function of story type.

The interaction of story and statment type probably was due to differences in materials and no reliable conclusions about processing differences can be drawn from this effect. The same can be said for the main effect of story type, especially since statement lengths for the two types of stories differed so markedly. The effects of greatest interest in the analysis of the reaction time data are reading rate and its interactions with other variables. None of these effects approached significance in the reaction time data for hits. Subjects who skimmed the stories were just as swift in verifying true statements as were those who read normally. Different types of statements did not require greatly different response times and the pattern of what small differences there were was maintained across reading rate. These results suggest that the memory representation resulting from skimming is about as well organized as that resulting from normal reading. Of course, a representation based on skimming stories contains less accurate information as witnessed by reduced recognition performance.

<u>Reaction time for correct rejections</u>. Reaction time data were prepared for correct rejections in the same way as they were

for hits. Mean reaction times for correct rejection of different statement types are shown in Figure 3. An analysis of variance of these data indicated that correct rejection of newspaper statements took reliably less time (5.055) than did correct rejection of narrative statements (6.211), F (1, 70) = 65.96, MS_e = 2.190. The three statement types differed in time taken for correct rejection, F (2, 140) = 10.33, MS_e = 1.634. A Newman-Keuls test found reaction time for microstatements (5.796) and macrostatements (5.863) did not reliably differ, but both were significantly greater than reaction time for inferences (5.240). There was also a significant interaction between story type and statement type, F (2, 140) = 5.14, MS_e = 1.233. In order to interpret this interaction and its implications for the main effect of statement type, separate analyses of variance were conducted of the data from narratives and newspaper stories. The analysis of reaction times for narratives found only a main effect of statement type, F (2, 140) = 5.63, MS_e = 1.831. A Newman-Keuls test found that time taken to reject false microstatements (6.143) and inferences (5.872) did not reliably differ while both took less time than rejection of macrostatements (6.619). The analysis of rejection time for newspaper statements showed that subjects in the normal reading rate condition tended to take more time to reject statements (5.382) than did subjects in the skimming condition (4.728), F (1, 70) = 3.94, MS_e = 5.879, p < .06. Rejection times for the different statement types were reliably different, F (2, 140) =

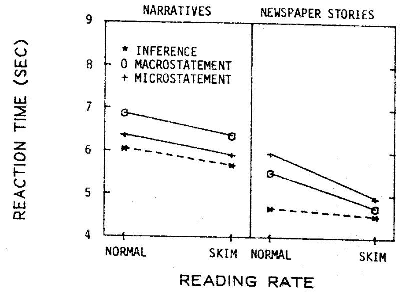

Figure 3. Mean reaction time on correct rejections on meaning test in Experiment 1.

12.46, MS_e = 1.036, with a subsequent Newman-Keuls test indicating that microstatements required the longest time (5.450), macrostatements required significantly less time (5.108), and inferences were rejected reliably more quickly (4.608) than either of the two explicit types of statements. A reliable interaction between rate and statement type, F (2, 140) = 3.46, MS_e = 1.036, indicated that the effects of statement type diminished when subjects skimmed as compared to when stories were read at normal rates.

A general finding that was consistent across both story types was that inferences were rejected more rapidly than either of the explicit statement types. This finding may be due to uninteresting differences in materials (e.g., statement length), but its consistency suggests that subjects are able to reject inferences rapidly because the critical information represented in these statements is quite discrepant from anything stated in a story. False statements based on explicitly stated information contain enough accurate information (semantic and surface) to force a more careful memory search before rejection can be made with confidence. Another general trend is for correct rejection reaction times to be longer than reaction times for hits, probably due to the more exhaustive examination of memory representations required in the former case. The data from newspaper stories provided some additional information. The tendecy for subjects in the skimming condition to more rapidly reject false statements may stem from the possibility that memory

representations that result from skimming a story are not as detailed or elaborate as those that result from normal reading and, hence, can be searched more quickly. It is also possible that the subjects who skimmed stories acquired a speeded reading set and had a tendency to respond faster in some cases because of that set (Kieras, 1974). Finally, the decrease in differences between statement types as reading rate increased revealed that there may be somewhat less distinction of important, unimportant, and inferred information when skimming. This explanation is a very tenuous one, however, as the effect did not occur for reaction times for hits nor for correct rejection of narrative statements. Also, the effect may have been due to a speeded reading set since it occurred under the same conditions that indicated the possible existence of a speeded reading set during the test.

Reaction time for false alarms. Time taken to make false alarms was analyzed in the same way as other types of reaction time data. Mean false alarm reaction times are presented in Figure 4. An analysis of variance found that subjects in the normal rate condition took more time in making false alarms (7.139) than subjects in the skimming condition (5.690), F (1, 70) = 21.15, MS_e = 10.724. False alarms to newspaper statements were made more swiftly (5.535) than to narrative statements (7.295), F (1, 70) = 92.03, MS_e = 3.638. The main effect of statement type was reliable, F (2, 140) = 3.57, MS_e = 3.794, and a Newman-Keuls test indicated that only

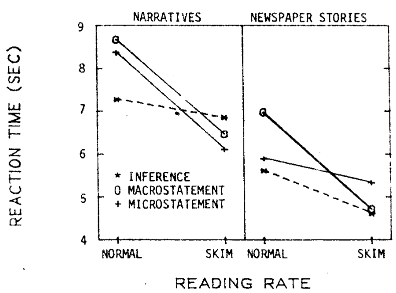

Figure 4. Mean reaction time on false alarms on meaning test in Experiment 1.

macrostatements (6.710) and inferences (6.098) differed reliably, while microstatements (6.436) were not significantly different from either of the other two types of statement. All three main effects must be interpreted in light of significant interactions. Rate and statement type interacted, F (2, 140) = 5.49, MS_e = 3.794, and the three-way interaction involving rate, story and statement type was also significant, F (2, 140) = 3.76, MS_e = 3.394. The general interpretation of these interactions, supported by subisidiary analyses of each story type, is that statement types differ in false alarm reaction times among subjects in the normal rate condition and these differences are reduced and overall reaction time is lower among subjects in the skimming condition. Story type enters the interaction because the ordering of statement types differs across the two stories.

False alarms on macrostatements consistently take longer than on inferences for subjects in the normal rate condition. Assuming that this effect is not an artifact of materials, it suggests that the subjects may have more knowledge about the explicitly stated macrostatements that must be checked before acceptance or rejection can occur, while for inferences a false alarm may occur more rapidly because the test statement represents an inference that the subject must draw and find credible because some piece of contradictory information is not available in memory or because the subject had previously drawn a related incorrect inference. The reaction time differences are virtually eliminated when subjects skim and, as with correct rejections, this may by due to a less differentiated, less surface structure oriented representation, or to the establishment of a speeded reading set that obscures the subtle differences among statement types.

Reaction time for misses. The reaction time data for subjects' incorrect rejection of true statements was prepared in the same way as the other reaction time data. The mean time taken for misses is presented in Figure 5. An analysis of variance was done and it was found that the normal rate condition had longer reaction time (6.790) than the skimming condition (5.913), F (1, 70) = 10.13, MS_e = 8.199, and that newspaper statements were responded to faster (5.637) than narrative statements (7.067), F (1, 70) = 71.61, MS_e = 3.084. There were also two significant interactions. Rate and statement type interacted, F (2, 140) = 7.34, MS_e = 4.377, and the interaction involving rate, story and statement type was reliable, F (2, 140) = 4.09, MS_e = 4.493.

The main effects observed in the analysis of reaction time for misses were consistent with those found in the analyses of other reaction time measures. The interactions, however, are less interpretable. Because of rather high hit rates, very few misses were made and for a particular statement type about half the subjects missed no statements and, therefore, did not contribute data to the reaction time analysis. The replacement of the empty cells with the cell mean caused the pattern of results to be heavily influenced by a few subjects' reaction times, especially

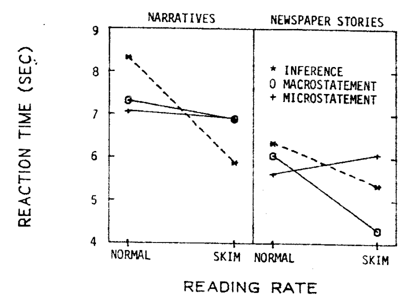

Figure 5. Mean reaction time on misses on meaning test in Experiment 1.

extreme times. The statistical effect was to increase within-subject variance (notice the mean-squared error values involving within-subject variables in the correct rejection and hit analyses compared to values for within-subject effects in the analysis of misses) and to greatly exaggerate differences between within-subject means, such as statement type. That the significant interactions may be due to statistical artifact must be considered a strong possibility. The uninterpretability of the interactions increases suspicions about their validity.

Memory for Surface Structure

Recognition performance. Recognition of the surface structure of statements contained in the critical stories was analyzed by calculating the hit and false alarm rate for each subject's performance on each type of statement. The mean hit and false alarm rates are presented in Table III. A d' score for each subject's performance on each statement type was calculated as was done for recognition of meaning data. The mean d' scores are shown in Figure 6. Also shown in the figure are confidence intervals for significance of difference from zero. The dashed line represents the 99% confidence level and the solid line represents the 95% confidence level. Means falling below these cutoff points should be considered not significant from zero (chance performance) with the appropriate degree of confidence. These confidence criteria were based on the error estimates obtained from an analysis of variance of the d' scores. The analysis was originally carried out including delay interval as a

TABLE III

MEAN HIT AND FALSE ALARM RATES ON SURFACE TEST IN EXPERIMENT 1

Story Type	Reading Rate	Statement Type			
		Macrostatement		Microstatement	
		Hit	FA[a]	Hit	FA
Narrative	Normal	.68	.48	.70	.50
Narrative	Skim	.61	.56	.64	.49
Newspaper	Normal	.76	.54	.72	.42
Newspaper	Skim	.74	.58	.53	.49

[a]FA = False Alarm

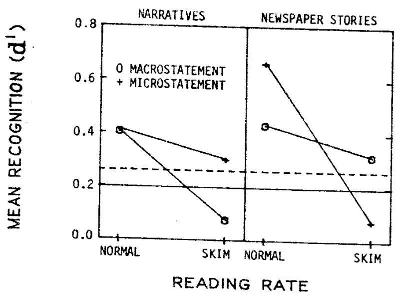

Figure 6. Mean recognition performance on surface test in experiment 1. Significance of difference from zero at 99% and 95% confidence levels are indicated by dashed and solid lines, respectively.

variable but no consistent or interpretable effects due to this factor were found and so it was not included in any of the reported analyses of surface memory.

The analysis found a main effect of reading rate with better surface memory d' scores for subjects in the normal rate condition (0.48) than for subjects in the skimming condition (0.19), F (1, 70) = 8.52, MS_e = 0.686. None of the other effects were reliable, although the three-way interaction involving rate, story and statement type approached significance, F (1, 70) = 3.13, MS_e = 0/696, p < .08.

The fact that surface memory deteriorated when reading rate was increased supports the hypothesis that subjects who skim stories do not fully process surface structure, but apparently rely on organization of partial knowledge obtained through highly predictive processing of a story. While the statement type differences observed at normal reading rates did not diminish (if anything, they increased) when subjects skimmed stories, the differences observed after skimming were primarily due to near chance performance on one statement type and chance performance on the other. The relative degree of surface memory for macro- and microstatements varied as a function of story type and it is unclear whether these variations reflect real and general processing differences or artifacts of test stimulus selection. The lack of high reliability of the three-way interaction discourages elaborate speculation.

Reaction time for hits. Reaction time data were prepared and analyzed for the surface test data in the same way as the data from the meaning test. The mean hit reaction times for the surface test are presented in Figure 7. An analysis of variance indicated that the effect of story type approached significance, F (1, 70) = 3.21, MS_e = 4.680, p < .08, with reaction time to newspaper statements (5.739) less than narrative statements (6.196). Microstatements were responded to more rapidly (5.751) than were macrostatements (6.184), F (1, 70) = 4.64, MS_e = 2.915, and there was a tendency for reaction time among normal rate subjects to be higher (6.350) than among subjects who skimmed (5.584), F (1, 70) = 3.53, MS_e = 11.990, p < .07. These last two main effects must be interpreted with respect to a reliable interaction between rate and statement type, F (1, 70) = 9.45, MS_e = 2.915. As indicated in Figure 7, the longer reaction time to macrostatements decreased as reading rate increased from normal to skimming.

The effect of decreased reaction time during skimming could reflect the same kind of speeded reading set as that implied by the recognition of meaning reaction time data. The interaction, however, implies that the surface memory distinction between macro- and microstatements may have diminished as a result of skimming. Certainly, the actual recognition performance approached chance when subjects skimmed and the reaction times for correct verification of surface structure are correspondingly not differentiated on the basis of importance. It is as though subjects process surface structure to a very limited extent,

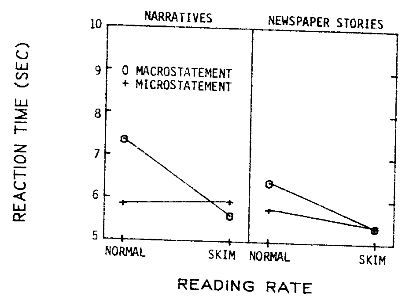

Figure 7. Mean reaction time on hits on surface test in Experiment 1.

regardless of importance. Once surface and semantic processing have determined that a statement is important the semantic gist is incorporated into the comprehension process with little additional surface processing, and if the statement is unimportant the gist is not further elaborated.

Reaction time for correct rejections. The mean reaction time for correct rejections is shown in Figure 8. An analysis of these times found that newspaper statements were responded to reliably faster (6.773) than were narrative statements (8.030), F (1, 70) = 14.85, MS_e = 7.671. Story and statement type were involved in an interaction, F (1, 70) = 7.78, MS_e = 3.161, and these two variables further interacted with reading rate, F (1, 70) = 30.37, MS_e = 3.161. Once again, the differential statement type effects may have been due to artifacts of differences in materials or to some, as yet, obscure processing difference between narratives and newspaper stories. In either case, the differences that manifested themselves at the normal reading rate were eliminated at the skimming rate, suggesting that skimming processes do not strongly differentiate surface aspects of story information in the resulting memory representation.

Reaction time for false alarms. The mean reaction times for false alarms made on different statement types are presented in Figure 9. The analysis of these data revealed that normal rate subjects took longer (7.149) in making false alarms than did skimming subjects (6.227), F (1, 70) = 4.82, MS_e = 12.703, and that newspaper statements were responded to faster (6.116) than

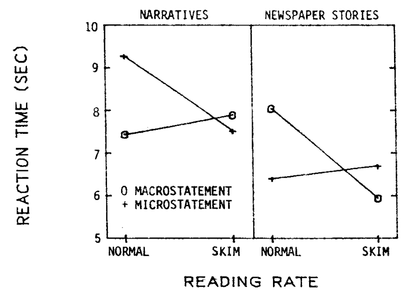

Figure 8. Mean reaction time on correct rejections on surface test in Experiment 1.

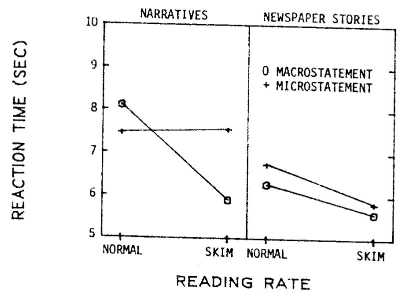

Figure 9. Mean reaction time on false alarms on surface test in Experiment 1.

narrative statements (7.261), F (1, 70) = 17.70, MS_e = 5.336. The main effect of statement type approached significance, F (1, 70) = 3.64, MS_e = 3.764, p < .06, with reaction time to microstatements (6.906) longer than macrostatements (6.470). Rate and statement type interacted reliably, F (1, 70) = 5.08, MS_e = 3.764, and these variables also interacted with story type, F (1, 70) = 6.12, MS_e = 4.594. Separate analyses supported the impression obtained from Figure 9 that the three-way interaction was due to the differential effect of skimming on reaction time for macro- and microstatements. The general decrease in reaction time observed for all other statement types was not found for narrative microstatements. It may be that subjects who skimmed were, for some reason, more cautious when responding to the statements from narratives that represented less important information and a speed-accuracy trade-off effect (Pachella, 1974) resulted. In fact, a review of Table III reveals that false alarm rates increased as reading speed increased for all statement types except narrative microstatements, which showed a slight <u>decrease</u>. Except for this trade-off effect, a general decrease in reaction time was apparent for false alarms just as with other reaction time measures. The generality of the decrease of reaction times when subjects skimmed is consistent with the hypothesis that subjects who skimmed adopted a speeded reading set. The set could be altered, however, by speed-accuracy concerns and a more cautious strategy led to longer reaction times.

__Reaction time for misses__. The mean time taken to incorrectly reject a verbatim statement is shown in Figure 10. An analysis of these reaction times found that normal rate subjects reacted more slowly (7.132) than skimming subjects (6.319), F (1, 70) = 4.60, MS_e = 10.333. Newspaper statements were responded to more quickly (6.325) than narrative statements (7.216), F (1, 70) = 12.61, MS_e = 5.488. Rate and statement type interacted, F (1, 70) = 6.21, MS_e = 4.198, and the interaction of rate, statement and story type was reliable as well, F (1, 70) = 7.89, MS_e = 3.341. The observed main effects were consistent with those found in analyses of other reaction time measures. The interactions were due to the lack of a decrease in reaction time of misses on newspaper macrostatements as reading rate increased. As with false alarm reaction times, a speed-accuracy trade-off appeared to be operating. Table III shows that the hit rate for newspaper macrostatements did not decrease with increased reading rate to the same extent as for other statement types, particularly newspaper microstatements. The cost of maintaining a relatively high hit rate for newspaper macrostatements was a relatively high reaction time. The reason why subjects selected for special consideration the specific type of statement they did is a matter of speculation.

Summary and Conclusions

The time the subjects required to read the critical stories indicated that the instructions to read normally or to skim

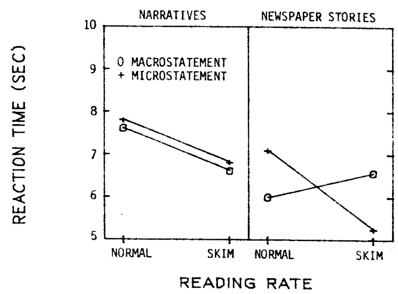

Figure 10. Mean reaction time on misses on surface test in Experiment 1.

produced reliably different reading rates. Although a good deal of variability between individuals was observed, within-individual reading rates were remarkably consistent. The results of the questionnaire suggested that usual strategies of normal reading and skimming were, for the most part, maintained in the experiment. Both of these reading strategies appear to be centered around the formation of a macrostructure as a memory representation of information explicitly presented in and inferred from a story. The formation of a macrostructure is common to both narratives and newspaper stories, and was evidenced by superior recognition performance on macrostatements and inferences relative to microstatements. Memory for meaning and for surface structure generally decreased as a result of skimming. Surface memory, however, was not consistently superior for one kind of statement and any differences observed among subjects reading at normal rates tended to diminish when subjects skimmed the stories. Reaction times in the meaning and surface tests primarily reflected theoretically uninteresting characteristics of test statements such as length, but also tended to decrease as reading rate increased, suggesting that subjects in the skimming condition adopted some form of a speeded reading set during the test phase.

The results of Experiment 1 were very encouraging with respect to the validity of experimental procedures used to study the reading process and resulting memory representations. The effect of statement importance was maintained across reading

rates, indicating that some essential macrostructure processes were operative. These processes did not completely deteriorate when skimming rates were in effect. It was hypothesized that macro-operations (Kintsch & van Dijk, 1978) might even be more strongly emphasized during skimming and as a result the importance effect in recognition would increase in the skimming condition. While this effect was not obtained using the present measures, it is theoretically possible that macro-operators were especially emphasized during skimming but that certain memory characteristics obscured the effect. Specifically, suppose we assume that memory for macrostatements and inferences is related to reading rate by a function that has the same basic shape as that for microstatements (see Figure 11), but that the latter function is, in fact, steeper and is displaced further toward lower reading rates. That is, the steepest drop in memory performance on microstatements would be associated with lower reading rates than would that of important statements. Suppose we also assume that the tested reading rates sampled these two functions at points beyond that associated with the steepest drop in memory performance on microstatements. Then the sampling of the microstatement function would include points with relatively little difference in performance, while the sampling of the important statement function would include points of performance that were just as differentiated as those sampled from the microstatement function (see Figure 11). The result is a general impression of equal effects of rate on recognition of statements produced reliably different reading rates. Although a good deal of variability between individuals was observed, within-individual reading rates were remarkably consistent. The results of the questionnaire suggested that usual strategies of normal reading and skimming were, for the most part, maintained in the experiment. Both of these reading strategies appear to be centered around the formation of a macrostructure as a memory representation of information explicitly presented in and inferred from a story. The formation of a macrostructure is common to both narratives and newspaper stories, and was evidenced by superior recognition performance on macrostatements and inferences relative to microstatements. Memory for meaning and for surface structure generally decreased as a result of skimming. Surface memory, however, was not consistently superior for one kind of statement and any differences observed among subjects reading at normal rates tended to diminish when subjects skimmed the stories. Reaction times in the meaning and surface tests primarily reflected theoretically uninteresting characteristics of test statements such as length, but also tended to decrease as reading rate increased, suggesting that subjects in the skimming condition adopted some form of a speeded reading set during the test phase.

The results of Experiment 1 were very encouraging with respect to the validity of experimental procedures used to study the reading process and resulting memory representations. The effect of statement importance was maintained across reading

rates, indicating that some essential macrostructure processes were operative. These processes did not completely deteriorate when skimming rates were in effect. It was hypothesized that macro-operations (Kintsch & van Dijk, 1978) might even be more strongly emphasized during skimming and as a result the importance effect in recognition would increase in the skimming condition. While this effect was not obtained using the present measures, it is theoretically possible that macro-operators were especially emphasized during skimming but that certain memory characteristics obscured the effect. Specifically, suppose we assume that memory for macrostatements and inferences is related to reading rate by a function that has the same basic shape as that for microstatements (see Figure 11), but that the latter function is, in fact, steeper and is displaced further toward lower reading rates. That is, the steepest drop in memory performance on microstatements would be associated with lower reading rates than would that of important statements. Suppose we also assume that the tested reading rates sampled these two functions at points beyond that associated with the steepest drop in memory performance on microstatements. Then the sampling of the microstatement function would include points with relatively little difference in performance, while the sampling of the important statement function would include points of performance that were just as differentiated as those sampled from the microstatement function (see Figure 11). The result is a general impression of equal effects of rate on recognition of statements

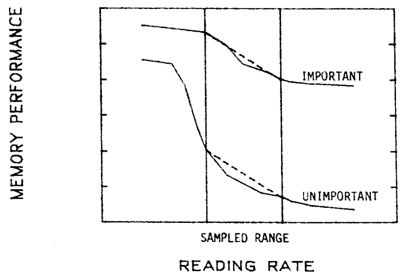

Figure 11. Theoretical functions relating memory to reading rate for important statements and unimportant statements.

of different levels of importance. But sampling of a greater range of rates might reveal more subtle differences between the two functions, particularly a greater maximum decrease for microstatement performance than for performance on important statements. Because of the complexity of this issue and the nature and probable instability of the functions relating reading rate and memory performance, a clear distinction between the theoretical alternatives of maintaining versus increasing macro-operations during skimming will not be a primary goal in this thesis. Emphasis will instead be placed on the more testable (and more critical) issue of whether or not macro-operations function well enough to honor the relative significance of different facts and form at least a general macrostructure when a reader skims. The results of Experiment 1 strongly suggest that they do, though an incomplete macrostructure results, as not all information can be adequately processed by the macro-operations. The experiments reported in the remaining chapters were designed to further investigate the characteristics of these macro-operations and their relationships with data driven and conceptually driven processes in varied reading situations.

CHAPTER III

EXPERIMENT 2

In Experiment 1 there was not a great deal of control over the reading rates adopted by subjects. The result was a wide range of reading rates operative in each reading rate condition. Experiment 2 was designed to establish strict control over reading rates and to increase rates beyond those used by subjects in a free skimming situation. A number of procedures could be used to control reading rates such as presenting limited portions of the text for fixed time intervals (Carver, 1977). Controlled presentation of text has commonly been used in speeded reading research (see Chapter I) but suffers from the major disadvantage of being very unnatural. Subjects are not able to make decisions about which parts of a text they wish to process more elaborately nor are regressive eye movements to earlier parts of a text possible. Skimming processes probably would be very different in this situation compared to more realistic reading situations.

The primary objective in controlling reading rates is to allow the subject an amount of time to process a text that corresponds to a particular rate of speed. This can be accomplished in a relatively natural way by making the whole text available to the subject for the fixed period of time. Rather than having the subject read as much as possible in the allotted time at his or her own rate (Carver, 1977), however, the subject should be encouraged to read at rates consistent with the rate

established by the time interval allowed for reading. That is, depending on the rate in effect subjects would be told they will have enough time to read at a rate similar to their normal reading rates or at a rate that corresponds to skimming, and subsequently encouraged to operate at those rates. Practice texts are necessary to allow subjects to accommodate their reading style to the rate in effect. Assistance in reading at the established rate can be provided by keeping the subjects informed of elapsed time (Wright, 1971). Subjects need not read the text at a uniform rate under this procedure, but can allocate reading time according to whatever criteria they would ordinarily use in skimming or normal reading. The overall reading rate would conform to the established rate. Information about elapsed time is presented only to keep the subjects informed of time constraints and through experience on practice texts a general sense of the reading rate requirements can be obtained. This procedure has the advantage of being reasonably natural with only time constraints imposed upon subjects. Otherwise, subjects are free to use any reading strategy and time allocation priorities they choose. No other pacing procedure offers this degree of freedom. Anything more restrictive may have severe consequences for normal reading processes such as formation of macrostructures of text information. For these reasons the pacing procedure used in Experiment 2 allowed subjects access to the whole text during a fixed time interval corresponding to the established reading rate. The extent and nature of the influence of the pacing procedure used in Experiment 2 can be gauged by comparing the experimental results with those found in Experiment 1 where subjects read at self-paced rates.

Three reading rates were selected for Experiment 2. The first two were based on the subject-paced reading rates observed in Experiment 1. The normal and skimming rates were based on the median rates of the corresponding reading conditions of the first experiment. The selected paced rates were 225 wpm and 375 wpm for normal reading and skimming, respectively. A third rate of 600 wpm was chosen to push subjects to skim stories very rapidly and will be referred to as the _fast skimming_ rate. The reason for including such a fast rate was that while subjects may skim at less than 400 wpm when they know they will be tested, it may be that this type of skimming is unlike the kind of skimming that is done at very fast rates. Also, it was important to determine whether subjects could successfully form macrostructures of story information when skimming at a rate that exceeds more careful skimming speeds.

Experiment 2 involved the same stories and test materials as those used in Experiment 1 so that direct comparisons of experimental results could be made in evaluating the effects of pacing subjects' reading rates. The two different story types, narrative and newspaper, represent basically different story structures and their use allowed observation of the effects of various reading rates on the processing of stories with diverse structures. Use of the different statement types provided

information about the formation of a macrostructure representation in memory and the ability to draw inferences from information obtained while reading at widely varying rates. The surface memory test was included to provide further evidence about the nature of surface processing during normal reading and skimming.

The final variable included in Experiment 2 was that of typographical case. Half of the subjects read stories typed in normal case while the other half read stories typed in alternating case (every other letter capitalized). The logic underlying the use of different typographies was related to the hypothesis that when readers skim they must rely more on whole word visual identification of words because they lack the time and cognitive resources necessary to make full use of phonological decoding processes. Baron (1977) has shown that use of alternating case sufficiently disrupts visual patterns so as to make whole word visual identification very difficult. The alternating case typography must be prepared with capital letters larger than normal letters in order to obtain the effect. If capital letters are no larger than the lower case letters no disruption occurs since subjects are able to integrate alternative sets of features into whole word visual identification when size is controlled (F. Smith, Lott, & Cronnell, 1969). When alternating case is used for normal reading rates a relatively small detrimental effect occurs (Fisher, 1975; Fisher & Lefton, 1976) but when subjects must scan material at rapid rates as in a search task, the effects of alternating case are accentuated (Fisher & Lefton, 1976; F. Smith et al., 1969). These results are consistent with the hypothesis that whole word visual identification is particularly important during skimming.

One possible effect of alternating case on meaning recognition performance would be a general decrease due to reduced efficiency or completeness of processing text information. This effect might also be observed in the surface memory test if subjects are unable to devote much attention to processing of surface structure since full surface processing requires time consuming decoding of visually unfamiliar word shapes. Perhaps the most interesting effect of all would involve the effect of alternating case on the ability to develop a macrostructure representation while reading normally or skimming. If decoding alternating case is a sufficiently difficult task with high demands on processing time and resources it may be that subjects would not have adequate time or cognitive resources (Kahneman, 1973; Norman & Bobrow, 1975) for the selective processing required to form a macrostructure from a story being read. Formation of macrostructures does appear to require some amount of cognitive resources (Britton et al., 1979), but Experiment 1 demonstrated that macrostructures can be formed even while skimming. The important question, then, is whether macrostructures can be formed during normal reading or skimming when perceptual processes cannot operate optimally and when further demands are placed on cognitive resources as a result of

having to decode alternating case. Those reading situations susceptible to alternating case effects would be those that rely primarily on whole word visual identification processes for rapid and resource-efficient text comprehension.

Method

Subjects

The subjects in Experiment 2 were 180 students recruited from the same source as that used in Experiment 1. There were 30 subjects assigned to each of the six combinations of reading rate and typographical case. Assignment was based on the order of the subjects' appearance at the laboratory.

Materials

The reading and test materials used in Experiment 1 were also used in Experiment 2. In addition, a second set of story booklets identical to that used in Experiment 1 was prepared using alternating case typography in which alternating letters were capitalized and all other letters were lower case, regardless of normal rules for capitalization.

Design

For the meaning test a 3 X 2 X 2 X 3 mixed factorial design was used. Reading rate (normal, skimming, and fast skimming) and typographical case (normal and alternating) were between-subject variables. Story type (narrative and newspaper) and statement type (inference, macrostatement, and microstatement) were within-subject variables. For the surface test a 3 X 2 X 2 X 2 mixed factorial design was used. All four variables were the same as in the meaning test except that the statement type variable included only two levels (macrostatement and microstatement).

<u>Procedure</u>

Subjects participated in small groups ranging in size from two to six. Each subject was first required to read the practice narrative as in Experiment 1. Once all subjects in the session had completed this task and the materials were collected, the experimenter explained that the subjects would next be asked to read a set of stories. Subjects were told that they would be asked to read at a particular rate of speed and that that rate would be similar to their normal reading rate, to the rate they might use when skimming, or to a rate they might use when skimming very fast, depending on the reading rate in effect during the session. The method of pacing the subjects through each story was then described and subjects were reminded that their task would be to obtain from each story the important information that would contribute to the story's general meaning. The story booklets were distributed and the subjects read the instructions on the cover page. These instructions were basically the same as those that appeared on the cover page of the story booklets used in Experiment 1, except that the instructions for Experiment 2 also included a description of the pacing procedure. Also, for those subjects receiving stories typed in alternating case, a paragraph was included that introduced the characteristics of the typography. After reading the instructions

and having any questions answered, subjects began the paced reading procedure.

Reading rate was paced using a tape recorded message for each story. The message began with the word "Start," indicating that the subjects were to open the booklets to the first story and begin reading. A series of numbers, "one" through "four," was heard with each number occurring at equally spaced time intervals. Thus, when subjects heard the number "one" it meant that one quarter of the reading time had elapsed, and when they heard the number "four" it meant that the total reading time had elapsed. At that point subjects were to be finished reading and were required to turn to the next blank page in the booklet. Six sets of messages were recorded, taking into account the two different orderings of narrative and newspaper stories (the two story types differed in length), and the three different reading rates. Two sets of messages were prepared for each of the three rates: 225 wpm, 375 wpm, and 600 wpm.

When the pacing procedure was described to the subjects it was emphasized that when they heard a number, for example "one," it did not mean that they necessarily had to be finished reading the corresponding proportion of material in the story (1/4). Subjects were told that they could distribute their reading time over the story in any way they desired, as long as they completed their reading by the end of the allotted reading time. Consequently, subjects were asked to adjust their general reading rate to cover the whole story in the time allowed. The numbers were present only to keep subjects informed of elapsed and remaining time. Once the time for reading a story had elapsed, subjects turned to a blank page while waiting for the signal to start reading the next story. There was a pause of about 5 sec between stories. At the end of the reading phase subjects began the test phase which involved procedures and materials identical to those used in Experiment 1. Once the test phase had been completed subjects filled out the same questionnaire as that used in Experiment 1.

Results and Discussion

Questionnaire

The subjects in Experiment 2 were not given the opportunity to read the critical stories at their preferred reading rates. Instead, they were required to read at paced rates. Therefore, in order to be able to claim that the reading processes in Experiment 2 were reasonably representative of natural reading, it is very important to evaluate subjects' comments on how well their experimental reading experiences corresponded to their normal reading experiences. Subjects' replies to the questionnaire item regarding a comparison of usual reading rates with those used in the experiment are relevant to this issue.

The percentages of subjects in each experimental condition claiming that the paced reading rate was the same as, faster than, or slower than usual reading rates are shown in Table IV. In general, subjects reading at the normal rate felt that they

TABLE IV

SUMMARY OF RESPONSES TO QUESTIONNAIRE ITEMS REGARDING REACTION TO PACING PROCEDURE IN EXPERIMENT 2

Case	Reading Rate	Paced Reading Rate Relative to Own Rate (Percent of Subjects)			Experimental Skimming Success Ratings
		Same	Fast	Slow	
Normal	Normal	56	17	27	--
Normal	Skim	3	97	0	2.03
Normal	Fast Skim	10	90	0	1.40
Alternating	Normal	40	40	20	--
Alternating	Skim	7	93	0	2.03
Alternating	Fast Skim	3	94	3	1.23

were reading at or slightly below their normal reading rates. There was a tendency for a greater percentage of subjects reading alternating case at the normal rate to claim that the pace was too fast. Given that all subjects were forced to read at one speed, however, it appears that a representative speed was chosen. For subjects in the skimming and fast skimming conditions there was almost complete agreement that the rates were too fast. Since these responses by subjects in both skimming conditions appeared to be based on comparisons with personal normal reading rates rather than personal skimming rates (as in Experiment 1), it is difficult to interpret them in the context of evaluating paced skimming rates with respect to usual skimming rates. A more reliable impression can be obtained by evaluating subjects' responses to the questionnaire item asking how successful subjects believed they were in extracting important information from the stories while skimming. These responses were evaluated using the same four-point scale as that used for this item in Experiment 1. The mean rated subjective success in extracting important information for subjects in both skimming conditions is presented in Table IV. It is apparent from these means that subjects in the skimming condition were more confident (2.03) that they had successfully gleaned the important information from the stories than were the subjects in the fast skimming condition (1.32), F (1, 116) = 12.54, MS_e = 1.228. Typographical case did not have a reliable effect on these ratings, nor was the interaction between reading rate and case significant (Fs < 1).

Thus, subjects in the skimming condition were able to maintain a relatively high degree of confidence (very close to that observed for subjects in the skimming condition in Experiment 1) in their ability to skim effectively at the paced rate of 375 wpm. The rate of 375 wpm seems to be a representative skimming rate for these subjects. When subjects were forced to read at 600 wpm, however, the degree of confidence in what they learned while skimming declined significantly, indicating that the rate of 600 wpm was probably beyond these subjects' preferred rates of skimming. These are exactly the roles that the rates of 375 wpm and 600 wpm were meant to play.

With respect to the subjects' responses to items on the questionnaire that dealt with skimming habits outside the laboratory, the mean estimate of percentage of total reading time that is spent skimming was 30%, with a full range of 0-100%. As in Experiment 1, the most common types of materials that are skimmed are newspapers and magazines as 37% of the subjects claimed to skim at least one of these two types of material. Also, 17% of the subjects stated that they skim at least some part of their college materials. A number of other types of material were mentioned idiosyncratically as in Experiment 1. A general estimate of how effectively subjects feel they skim on their own was obtained by evaluating subjects' responses to the question concerning how successful subjects' felt they are at extracting important information while skimming on their own. Responses were evaluated using the four-point scale described in

Experiment 1. The mean rating for the subjects in Experiment 2 was 1.86, indicating that these subjects felt that they can skim reasonably well on their own. Finally, 12% of the subjects indicated that they had taken some course of training in improving reading speed. Among those subjects who have had training in increasing reading speed, the mean evaluation of confidence in success while skimming outside the laboratory was a very high 2.33. With respect to these subjects' confidence in their success while skimming stories in Experiment 2, those who were paced at 375 wpm (N = 9) had a mean evaluation rating of 2.78, while those who were paced at 600 wpm (N = 7) had a mean rating of only 2.0. The difference between these mean ratings was reliable, F (1, 14) = 6.00, MS_e = 0.397. Apparently, even those subjects who had received training in improving reading speed found some difficulty in skimming at 600 wpm.

In summary, the results of the questionnaire provide two general conclusions about the skimming behavior of subjects in Experiment 2. First, the characteristics of subjects' skimming behavior outside the laboratory are similar to those described by subjects in Experiment 1. Second, although subjects in Experiment 2 read stories at paced rates of speed, it appears that they were, for the most part, able to maintain their usual reading strategies whether reading at the normal or at the skimming rates. At the fast skimming rate, however, subjects felt that they were less able to successfully extract important information from the stories. Consequently, it appears safe to conclude that

the pacing method did not seriously disrupt usual reading processes and that the selected reading rates were representative of normal, skimming, and fast skimming rates ordinarily used by subjects.

<u>Memory for Meaning</u>

<u>Recognition performance</u>. Ability to recognize the meaning of test statements that appeared in or could be inferred from one of the critical stories was characterized in the same manner as for Experiment 1. The mean proportions of hits and false alarms are presented in Table V. Each subject's performance was also characterized by a set of d' scores combining appropriate pairs of hit and false alarm rates in the same way as for Experiment 1. The mean d' scores are presented in Figure 12.

In order to further establish the validity of the pacing procedure as a method of ensuring particular rates of reading while not seriously disrupting usual reading processes, an analysis was carried out comparing the d' recognition of meaning scores of subjects in Experiment 1 to those of subjects in Experiment 2. To equalize the number of subjects in each condition, six subjects from each reading rate condition (one subject from each of the six counterbalancing conditions) in Experiment 1 were randomly chosen for exclusion and all subjects in the fast skimming and in the alternating case conditions in Experiment 2 were excluded from the analysis. The analysis, then, included experiment and reading rate as between-subject factors, and story and statement type as within-subject factors. The only

TABLE V

MEAN HIT AND FALSE ALARM RATES ON MEANING TEST IN EXPERIMENT 1

Case	Story Type	Reading Rate	Statement Type[a]					
			Inf.		Mac.		Mic.	
			Hit	FA[b]	Hit	FA	Hit	FA
Normal	Narrative	Normal	.88	.13	.91	.12	.92	.26
Normal	Narrative	Skim	.89	.19	.91	.16	.82	.25
Normal	Narrative	Fast Skim	.81	.27	.85	.33	.72	.38
Normal	Newspaper	Normal	.84	.13	.93	.17	.77	.28
Normal	Newspaper	Skim	.78	.14	.86	.14	.78	.31
Normal	Newspaper	Fast Skim	.79	.27	.82	.26	.70	.32
Alternating	Narrative	Normal	.86	.19	.88	.17	.87	.28
Alternating	Narrative	Skim	.83	.25	.84	.24	.70	.40
Alternating	Narrative	Fast Skim	.81	.24	.79	.24	.65	.40
Alternating	Newspaper	Normal	.83	.19	.84	.22	.81	.28
Alternating	Newspaper	Skim	.76	.19	.83	.25	.72	.35
Alternating	Newspaper	Fast Skim	.73	.27	.75	.28	.62	.28

[a]Inf. = Inference; Mac. = Macrostatement; Mic. = Microstatement

[b]FA = False Alarm

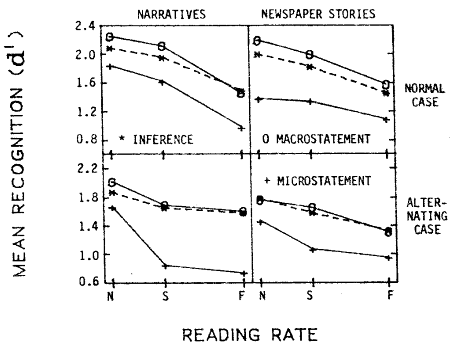

Figure 12. Mean recognition performance on meaning test in Experiment 2 (N = Normal; S = Skim; F = Fast Skim).

effects of interest in this analysis were those involving the main effect of experiment and any interactions involving this variable. None of these effects were reliable and the only one that approached significance was the interaction between experiment and reading rate, F (1, 116) = 3.16, MS_e = 0.791, p < .08. This interaction effect was due to a greater decrease in performance due to skimming in Experiment 1. Thus, it seems clear that the use of the pacing procedure did not produce sufficient aberration of the reading process to seriously affect performance on the recognition of meaning task. If anything, the pacing procedure involved a slightly more moderate skimming rate than that adopted voluntarily by subjects in Experiment 1.

An analysis of variance of the d' scores for the whole set of subjects in Experiment 2 was carried out, including reading rate and typographical case as between-subject variables and story and statement type as within-subject variables. All four main effects were significant. As reading rate increased, recognition performance decreased, F (2, 174) = 23.30, MS_e = 1.232. The mean recognition performance for subjects in the normal (1.85), skimming (1.61), and fast skimming (1.29) conditions all reliably differed from one another according to a Newman-Keuls test.

Subjects who read stories typed in normal case had reliably higher d' scores (1.70) than did subjects who read stories typed in alternating case (1.47), F (1, 174) = 11.24, MS_e = 1.232. Subjects were better able to recognize statements taken from

narratives (1.63) than statements taken from newspaper stories (1.54), F (1, 174) = 4.33, MS_e = 0.556. The different statement types differed in the degree to which they were recognized, F (2, 348) = 86.31, MS_e = 0.373. A Newman-Keuls test indicated that recognition of inferences (1.71) and of macrostatements (1.79) did not reliably differ, while recognition of either of these two statement types was significantly superior to that of microstatements (1.24).

Only one interaction was significant in the analysis. This was a three-way interaction involving case, story type, and statement type, F (2, 348) = 3.26, MS_e = 0.398. In order to interpret this interaction two separate analyses of variance were computed, one for each story type. Each analysis included reading rate and case as between-subject factors and statement type as a within-subject factor. In both analyses all three main effects were reliable with patterns of significant differences between means almost identical to those found in the original analysis. Additionally, the analysis of recognition of statements taken from narratives revealed three interactions that were at least close to being reliable. The interaction between rate and statement type was relaible, F (4, 348) = 2.82, MS_e = 0.319, indicating that ability to recognize inferences and macrostatements declined less rapidly than did ability to recognize microstatements as reading rate increased. The interaction between rate and case approached significance, F (2, 174) = 2.87, MS_e = 0.984, p < .06. This interaction implied that the negative effects of reading alternating case increased as reading rate increased from normal to skimming, but disappeared when the rate was set for fast skimming. Alternatively, one could interpret the interaction as indicating that for subjects reading normally typed stories, a small drop in d' scores was observed as reading rate went from normal (2.05) to skimming (1.89) and a rather large drop occurred when the rate was set at fast skimming (1.30). On the other hand, for subjects reading the stories typed in alternating case a large drop in d' scores was observed as soon as reading rate went from normal (1.84) to skimming (1.39) and a lesser decrease occurred when the rate was fast skimming (1.30). Finally, the interaction between case and statement type approached significance, $\underline{F}$ (2, 348) = 2.68, $\underline{MS}_e$ = 0.319, $\underline{p}$ < .07, indicating that reading alternating case had a more detrimental effect on recognition of microstatements than of macrostatements or inferences.

These interactions taken together with Figure 12 imply that the effect of reading rate on recognition of different statement types depends both on the type of story and the case in which the story is typed. In order to further investigate this possibility and to delineate those situations in which interactions between reading rate and statement type are likely to occur, four more subsidiary analyses of variance were done. One analysis was done for each combination of story type and case. In each analysis the variables were rate (between-subject) and statement type (within-subject), and the question of interest was whether these

two variables interacted. In each analysis both main effects were highly reliable as they had been in the earlier analyses, but, as expected from viewing Figure 12, only the analysis involving narratives read in alternating case revealed a reliable interaction between rate and statement type, $\underline{F}$ (4, 174) = 3.40, $\underline{MS}_e$ = 0.318. This interaction did not approach significance in the other analyses ($\underline{F}$s < 1). The significant interaction clearly indicates that the condition of reading narratives typed in alternating case leads to a more rapid decline in recognition of microstatements than of macrostatements or inferences as reading rate increases.

The comparison of recognition of meaning in Experiments 1 and 2 was very encouraging with respect to the validity of the pacing procedure used to control reading rates. The basic pattern of results found when subjects read at self-paced rates were replicated when a controlled pacing procedure was used. In conjunction with the results of the questionnaire, these findings support the argument that subjects can be paced in reading at normal and skimming rates without serious disruption of regular reading processes. The use of a third reading rate provided an indication of the effects of skimming at a rate beyond that of careful skimming. Even at 600 wpm subjects successfully formed macrostructure representations of the stories.

The effects of alternating case were generally detrimental as recognition performance was higher for subjects who read normally typed stories than for subjects who read stories typed in alternating case. There was a tendency for the effect of alternating case on recognition of narrative statements to be greater in the skimming condition than in the normal rate condition, but there was no effect at all in the fast skimming condition. It appears that when subjects are pressed to read at 600 wpm the ability to carefully process perceptual information is sufficiently compromised that use of alternating case did not add a significant burden. Subjects could revert to a rapid but impoverished perceptual processing strategy or a more selective one in which certain parts of a story were read carefully. In the latter instance alternating case would not have a great effect (Fisher, 1975; Fisher & Lefton, 1976). Whatever the general processing strategy, subjects were capable of forming a representation of a story that was characteristic of a macrostructure, though not complete, even when a visually unfamiliar typography was used. Clearly, even under rather extreme circumstances subjects are well-equipped to use macro-operations to regulate processing of information extracted from stories.

A particularly interesting consequence of reading alternating case was to amplify the effect of statement importance for narratives as reading rate increased. Unlike recognition of newspaper statements, the narrative microstatements showed a steeper decline as a function of reading rate than did inferences or macrostatements in the alternating case condition. For narratives, subjects were not as successful

at processing detailed information as they were for newspaper stories. The cause of this difference may lie in the different length and structural properties of the two story types and in the flexibility of allocating processing resources during reading. In newspaper stores the more important information is located early in the stories and the length of the stories used in the experiment was relatively short. Subjects could more easily select and process important information in this situation then go on to deal with the unimportant details later in the story. Even when the use of alternating case created excessive demands on processing capacity subjects could still be confident that the critically important information had been processed once the early part of the story had been read. Less concentrated processing would then be devoted to the remainder of the story without fear of missing critical information. In the narratives, however, the length was greater and the important information in narratives can be scattered throughout. In the situation of skimming narratives in alternating case, subjects may have had time and processing resources enough only to survey the whole story, concentrating on important information and allowing unimportant details only minimal processing. Thus, as rate increased, differential processing was maximized and processing of microstatements decreased markedly. The specific effects of alternating case on processing of narratives and newspaper stories point up the flexibility of resource allocation and the different underlying story structures and awareness of these differences appears to be part of the subjects' story schemata.

Reaction time for hits. Reaction time data were prepared and analyzed in the same way as they were for Experiment 1. The mean time taken to verify true statements is shown in Figure 13. An analysis of these reaction times found that the shorter newspaper statements were verified faster (4.171) than the narrative statements (5.269), F (1, 174) = 157.75, MS_e = 2.064. The effect of story type became more pronounced as reading rate increased, as rate and story type interacted, F (2, 174) = 5.01, MS_e = 2.064. The different statement types differed in time required for verification, F (2, 348) = 11.25, MS_e = 0.804. A Newman-Keuls test found that reaction time for microstatements (4.838) and inferences (4.782) did not reliably differ while both were significantly longer than reaction time for macrostatements (4.540). The effect of statement type interacted with reading rate and case, F (4, 348) = 2.77, MS_e = 0.804. This interaction can be interpreted in terms of deviations from the pattern of statement effects just described. For the normal case condition, reaction time to microstatements for fast skimming subjects was especially high, and for the alternating case condition, normal rate subjects responded especially slowly to inferences while skimming subjects were almost as fast on inferences as macrostatements. Two other interactions involving rate, story and statement type approached significance but both represented elaborations of the significant interaction without contributing further insight and will, therefore, not be described.

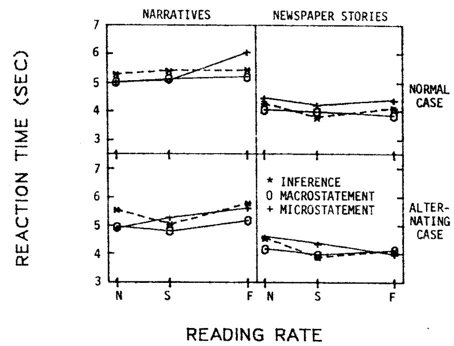

Figure 13. Mean reaction time on hits on meaning test in Experiment 2 (N = Normal; S = Skim; F = Fast Skim).

The time taken to verify statements in the meaning test did not consistently vary with reading rate. In general, even when subjects skim they can respond to test statements swiftly on the basis of the general macrostructure representation they have formed. The increase in the story type effect may be tied to statement length, but also to the greater recognition of narrative statements. Subjects may be more careful in their responses to narrative statements due to the greater amount of information and interrelationship of story elements found in the longer stories. The statement type effect may partially reflect differences in materials, but the finding that macrostatements are responded to more quickly is consistent with their critical role in the macrostructure, while the slower response times for inferences probably reflects lack of prior processing of any surface representation of the inferred information. The slower response to inferences was accentuated in the normal reading rate condition when alternating case was used, suggesting that subjects may have been particularly reliant on surface structure processing when dealing with alternating case at a slow reading rate (cf. Masson & Sala, 1978). In the skimming condition, however, reaction to any important information was rapid and inferences were responded to quickly. Part of this effect could be due to a speeded reading set, as more careful consideration of test statements among subjects in the fast skimming condition once again elevated reaction time to inferences. In reading normal case, fast skimming subjects were able to process

microstatements to a limited extent as these subjects' reaction time to microstatements was particularly high.

__Reaction time for correct rejections__. The mean time taken to correctly reject false statements is presented in Figure 14. The analysis of these reaction times found the usual effect of faster responding to newspaper statements (4.627) than to narrative statements (5.812), $\underline{F}$ (1, 174) = 208.75, $\underline{MS}_e$ = 1.818. The effect of statement type also was reliable, $\underline{F}$ (2, 348) = 35.87, $\underline{MS}_e$ = 1.185, and a Newman-Keuls test showed that macrostatements required significantly more time for rejection (5.495) than microstatements (5.329) and both explicit statement types required far more time than inferences (4.834). Statement and story type interacted, $\underline{F}$ (2, 348) = 7.05, $\underline{MS}_e$ = 1.057, and subsidiary analyses of each story type indicated that while narrative statements followed the described pattern, newspaper micro- and macrostatements switched positions.

There are two general outcomes in the correct rejection reaction time data that are consistent with findings in Experiment 1. First, time taken to reject statements is greater than verification time, suggesting a more exhaustive search before a rejection can be made. Second, false inferences are rejected more rapidly than any other statement type. This result is large and consistent enought to support arguments about processing and memory effects rather than stimulus selection artifact explanations. The more rapid rejection of false inferences probably was due to the lack of prior experience with

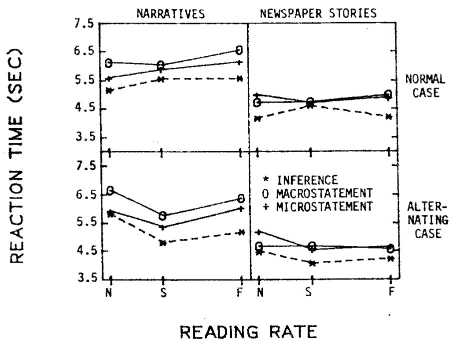

Figure 14. Mean reaction time on correct rejections on meaning test in Experiment 2 (N = Normal; S = Skim; F = Fast Skim).

a test inference's surface structure and the statement's lack of accurate information relevant to the story did not encourage very close checking of the inference with the memory representation of the story.

__Reaction time for false alarms__. The mean reaction time for false alarms is shown in Figure 15. An analysis of variance of these data showed that as reading rate increased time taken to make a false alarm decreased, F (2, 174) = 3.52, MS_e = 8.562. A Newman-Keuls test showed that reaction time in the normal reading rate condition was slower (6.040) than in the skimming (5.546) and fast skimming (5.531) conditions which did not differ reliably from each other. Newspaper statements were responded to more rapidly (5.149) than narrative statements (6.262), F (1, 174) = 108.32, MS_e = 3.088. Reading rate and case interacted, F (2, 174) = 4.60, MS_e = 8.562, indicating that when the normal reading rate was in effect subjects who read alternating case responded more slowly than subjects who read normal case, but at the fast skimming rate and especially the skimming rate subjects who read alternating case responded more rapidly. While the statement type main effect was not reliable this factor interacted with reading rate, F (4, 348) = 4.53, MS_e = 3.156, and with case, F (2, 348) = 3.06, MS_e = 3.156. The interaction of statement type and reading rate resulted from a faster response to macrostatements at the normal rate and a decrease in this effect as rate increased up to fast skimming where little, if any, statement differences were found. The

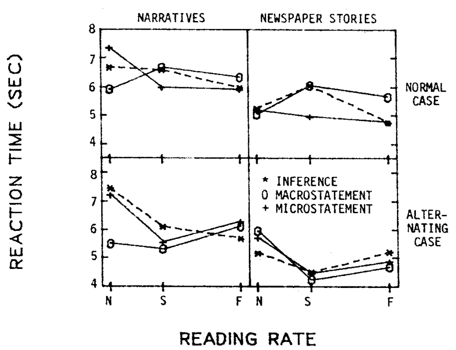

Figure 15. Mean reaction time on false alarms on meaning test in Experiment 2 (N = Normal; S = Skim; F = Fast Skim).

statement type interaction with case showed that the effect of alternating case on false alarm reaction time was largest in reducing response time for macrostatements and almost nonexistent for microstatements. The interaction involving story and statement type was reliable, F (2, 348) = 3.50, MS_e = 3.379, indicating that the idiosyncratic pattern of reaction time for different statement types varied across the two types of stories. Story and statement type further interacted with reading rate, F (4, 348) = 2.49, MS_e = 3.379. Subsidiary analyses of each story type showed that for narratives, macrostatements were responded to more rapidly than other statements but that this effect diminished as reading rate increased. No main effect or interaction involving statement type was significant in the analysis of newspaper statement responses.

The time taken to respond when subjects made false alarms was affected primarily by reading rate and case. Unlike Experiment 1, there was no consistent effect of statement type suggesting that interactions involving this variable were the result of material selection artifacts or the problem of inflated effects due to many subjects not making any false alarms to certain statement types and being assigned the cell mean instead. Any interactions involving statement type were characterized by reduced statement type effects as reading rate increased, concomitant with a smaller proportion of subjects failing to make at least one false alarm. The interaction between rate and case, however, suggests that subjects who read at normal speeds were able to establish a memory representation with confidence and any false alarms were the result of relatively long deliberation. When subjects skimmed the stories less certainty about representations in memory was possible, particularly when alternating case was used, and reasonable (but actually false) statements were accepted as valid rather than rejected on the basis of incomplete knowledge. A review of Table V confirms this general trend as increases in reading rate and the use of alternating case had proportionally greater effects on false alarms than on hits.

<u>Reaction time for misses</u>. The means for time taken in failing to verify a true statement are shown in Figure 16. The analysis of these reaction times showed that the only reliable main effect was story type, with newspaper statements responded to more rapidly (4.869) than narrative statements (6.349), $\underline{F}\ (1, 174) = 185.76$, $\underline{MS}_e = 3.186$. There was a reliable interaction involving reading rate and case, $\underline{F}\ (2, 174) = 7.09$, $\underline{MS}_e = 5.883$. Reaction time for subjects who read alternating case was greater than for subjects who read normal case when the normal reading rate was in effect, but when the skimming rate was used the reverse was true, and at the fast skimming rate there was no large difference due to case. The interaction between story and statement type was significant, $\underline{F}\ (2, 348) = 9.77$, $\underline{MS}_e = 2.692$, representing what is probably an effect due to selection of materials rather than processing differences between narratives o and newspaper stories. Four other interactions

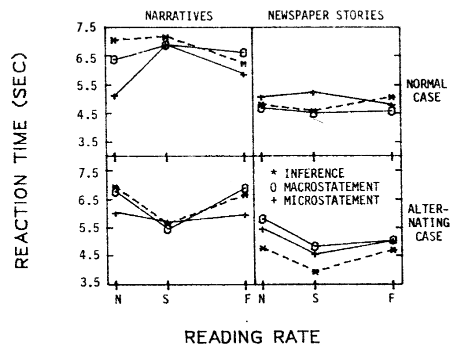

Figure 16. Mean reaction time on misses on meaning test in Experiment 2 (N = Normal; S = Skim; F = Fast Skim).

approached significance, but once again the high incidence of subjects not missing any true statements of a particular type suggests that these effects may well be artifactual.

As with reaction times for false alarms, the effect of alternating case varied with reading rate and was strongest in the skimming condition. In that condition, subjects were rather fast in rejecting true statements compared to subjects who had skimmed normally typed stories. Skimming appeared to lead to less certainty about resulting memory representations, especially when skimming alternating case, and subjects rapidly rejected unfamiliar statements even though they had occurred in a story that was read. Part of this effect could be due to the speeded reading set established during the paced reading procedure. The effect could be particularly strong for subjects who read stories in alternating case and on their test had their first opportunity to read experimental materials in normal case.

__Memory for Surface Structure__

__Recognition performance__. Hit and false alarm rates were calculated for each subject's performance on each statement type on the surface test. The mean hit and false alarm rates are presented in Table VI. A set of d' scores was also calculated for each subject and the mean d' scores are shown in Figure 17. An analysis of the d' scores showed that, in general, time of test (immediate or delay) had no consistent or interpretable effect on performance and was, therefore, not included in the analyses reported here. As reading rate increased, surface memory

TABLE VI

MEAN HIT AND FALSE ALARM RATES ON SURFACE TEST IN EXPERIMENT 2

Case	Story Type	Reading Rate	Statement Type[a]			
			Mac.		Mic.	
			Hit	FA[b]	Hit	FA
Normal	Narrative	Normal	.78	.58	.83	.56
Normal	Narrative	Skim	.73	.53	.69	.53
Normal	Narrative	Fast Skim	.69	.54	.57	.52
Normal	Newspaper	Normal	.77	.48	.64	.48
Normal	Newspaper	Skim	.67	.56	.53	.52
Normal	Newspaper	Fast Skim	.71	.59	.51	.52
Alternating	Narrative	Normal	.74	.62	.73	.52
Alternating	Narrative	Skim	.64	.62	.53	.47
Alternating	Narrative	Fast Skim	.69	.55	.53	.47
Alternating	Newspaper	Normal	.73	.54	.62	.43
Alternating	Newspaper	Skim	.61	.58	.54	.61
Alternating	Newspaper	Fast Skim	.70	.62	.54	.48

[a]Mac. = Macrostatement; Mic. = Microstatement

[b]FA = False Alarm

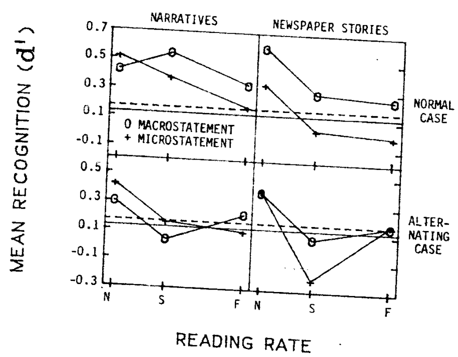

Figure 17. Mean recognition performance on surface test in Experiment 2 (N = Normal; S = Skim; F = Fast Skim). Significance of difference from zero at 99% and 95% confidence levels are indicated by dashed and solid lines, respectively.

performance decreased, F (2, 174) = 6.72, MS_e = 0.751. A Newman-Keuls test found that subjects in the normal reading rate condition had reliably greater recognition performance (0.42) than subjects in the skimming (0.16) or fast skimming (0.18) conditions, which did not reliably differ from one another. Subjects reading normal case had higher recognition performance (0.32) than subjects who read alternating case (0.18), F (1, 174) = 4.50, MS_e = 0.751. No other effects approached significance.

As in Experiment 1, surface memory deteriorated as reading rate increased, adding support to the hypothesis that when subjects skim they are not able to fully process surface structure. In fact, peformance on the surface memory test was rarely above chance, especially in the fast skimming condition. The effect of alternating case was to further decrease ability to fully process surface structure, as evidenced by lower recognition performance among subjects who read alternating case. In those situations where surface structure was not processed to any great extent, subjects probably had to rely to a greater extent on their own knowledge about the story's content and were required to form a comprehensible representation on the basis of partial information and predictive reading processes. Skimming generally led to decreased surface processing because of lack of available time. Particularly in the fast skimming conditions, the feasibility of devoting time and cognitive resources to the complete processing of surface structure would be extremely low. Rather, subjects would need to develop a more selective processing strategy, causing performance on a test for memory of detailed surface information to be quite low. The further effect of alternating case was to reduce the efficiency of surface processing by interfering with whole word visual identification processes. Less efficient decoding processes would be more costly in terms of time, and full decoding of surface structure would be less common. Only enough information for development of a macrostructure-like representation would be extracted with the details of a sentence's surface structure receiving minimal attention.

Reaction time for hits. The mean reaction time in verifying true statements is shown in Figure 18. The analysis of these data showed only that newspaper statements were verified more rapidly (4.991) than narrative statements (5.775), F (1, 174) = 24.21, MS_e = 4.575. No other effects approached significance. No effect of reading rate was observed, as subjects were able to verify correctness of surface form just as rapidly after they had skimmed stories as after they had read them at a normal rate. Whatever knowledge about surface structure that achieved representation in memory could be rapidly accessed regardless of skimming rate or typography used during reading.

Reaction time for correct rejections. The mean reaction times for correct rejections are shown in Figure 19. An analysis of variance of these data found that newspaper statements were responded to more quickly (6.064) than narrative statements (7.195), F (1, 174) = 62.31, MS_e = 3.689. This effect varied with

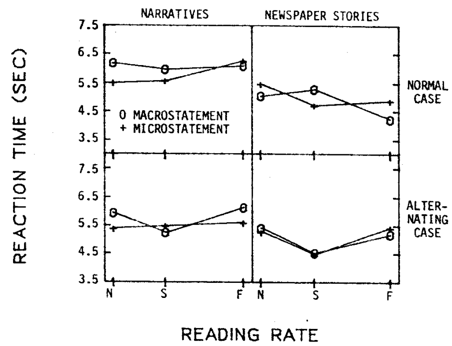

Figure 18. Mean reaction time on hits on surface test in Experiment 2 (N = Normal; S = Skim; F = Fast Skim).

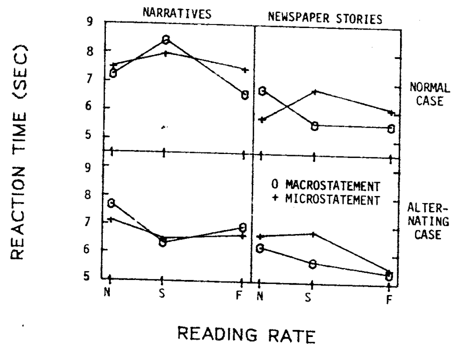

Figure 19. Mean reaction time on correct rejections on surface test in Experiment 2 (N = Normal; S = Skim; F = Fast Skim).

case, *F* (1, 174) = 4.80, $\underline{MS}_e$ = 3.689, and with reading rate and case, *F* (2, 174) = 4.73, $\underline{MS}_e$ = 3.669. Separate analyses of reaction time for narrative and newspaper statements were carried out to aid in the interpretation of these interactions. The analysis of narratives found that responses in the normal case condition took longer (7.526) than responses in the alternating case condition (6.863), *F* (1, 174) = 5.19, $\underline{MS}_e$ = 7.607. This effect appeared, however, only in the skimming condition as rate and case interacted, *F* (2, 174) = 3.75, $\underline{MS}_e$ = 7.607. In the analysis of newspaper statements it was found that as reading rate increased, reaction time decreased, *F* (2, 174) = 3.15, $\underline{MS}_e$ = 5.228. A Newman-Keuls test found that only the normal and fast skimming conditions differed reliably. There was also a significant effect of statement type as macrostatements were responded to more rapidly (5.865) than microstatements (6.264), *F* (1, 174) = 3.86, $\underline{MS}_e$ = 3.697. This effect held only in the skimming condition as the interaction between reading rate and statement type was reliable, *F* (2, 174) = 4.00, $\underline{MS}_e$ = 3.697.

The reaction time data for correct rejections must be interpreted with respect to those conditions that yielded above versus below chance recognition performance. In general, the observed interactions were due to effects in the skimming condition where speeded reading sets have been observed previously. Moreover, these effects involve comparison of reaction time to a type of statement that was below chance recognition performance with that of a statement type that was above chance on the recognition test. The fact that differences were observed could easily be due to strategic decisions on the part of subjects which are dependent on the existence or nonexistence of a reliable memory representation of surface structure information. As such, the correct rejection reaction time effects reflect the recognition performance results rather than provide further direct information about processing and memory.

<u>Reaction time for false alarms</u>. The mean reaction time for false alarms is shown in Figure 20. The analysis of these reaction times revealed the usual effect of faster responses to newspaper statements (5.520) than to narrative statements (6.473), F (1, 174) = 37.78, MS_e = 4.328. Macrostatements were responded to more rapidly (5.805) than microstatements (6.188), F (1, 174) = 9.71, MS_e = 2.719. There was a reliable interaction of story type and case, F (1, 174) = 8.40, MS_e = 4.328. These effects were offset by a four-way interaction involving reading rate, case, story and statement type, F (2, 174) = 3.95, MS_e = 3.741. In order to interpret the four-way interaction, separate analyses of each of the four combinations of story type and case were conducted. As expected from Figure 20, the combination of newspaper statements and normal case and of narrative statements and alternating case produced no significant effects. The analysis of narrative statements and normal case found a tendency for statement type and rate to interact, F (2, 87) = 2.83, MS_e = 3.374, p < .07, and the analysis of newspaper statements in the

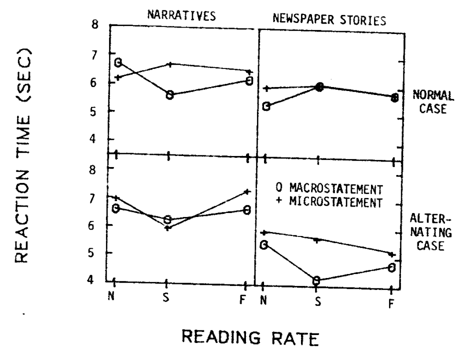

Figure 20. Mean reaction time on false alarms on surface test in Experiment 2 (N = Normal; S = Skim; F = Fast Skim).

alternating case condition found that macrostatements were responded to more quickly than microstatements, F (1, 87) = 8.33, MS_e = 3.268.

The effects of false alarm reaction times in the surface test must also be interpreted in terms of level of recognition performance. The effect in the narrative-normal case condition is clearly due to an advantage of macrostatements on speed when the skimming rate was used. In this case, the subjects were rather quick in making false alarms when presented paraphrase macrostatements. The reaction time advantage of macrostatements held across all reading rates in the newspaper-alternating case condition. These effects indicate a greater readiness to accept as true, information about statements more central to a story's gist.

Reaction time for misses. The mean time taken when subjects failed to verify a true statement is shown in Figure 21. The analysis of these data found only two reliable effects. As reading rate increased, reaction time decreased, F (2, 174) = 8.65, MS_e = 6.674. A Newman-Keuls test showed that the reaction time in the normal reading rate condition (6.742) was reliably longer than in the skimming (5.948) or fast skimming (5.846) conditions which did not reliably differ from each other. Once again, newspaper statements were responded to more swiftly (5.394) than narrative statements (6.974), F (1, 174) = 112.49, MS_e = 3.948.

The effect of reading rate on reaction time is consistent

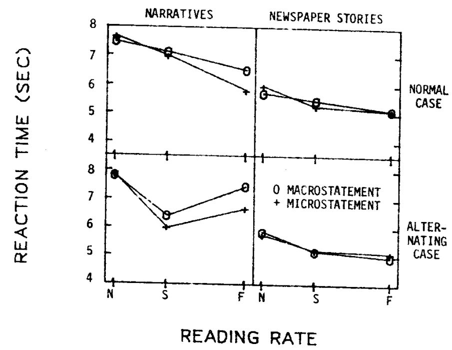

Figure 21. Mean reaction time on misses on surface test in Experiment 2 (N = Normal; S = Skim; F = Fast Skim).

with the idea that subjects in the skimming and fast skimming conditions had a less embellished memory representation of the stories. As a result, these subjects could more rapidly review their knowledge about a story and reject a statement that was not comprehended and represented with respect to its original surface structure.

Summary and Conclusions

The results of Experiment 2 served to verify two methodological aspects of the research reported here. First, subjects' evaluations of their ability to effectively skim stories in the experiment were reasonably consistent with memory performance in that fast skimming reduced confidence in skimming effectiveness and reduced actual memory performance (cf. Carver, 1973). Second, the overall degree of consistency between Experiments 1 and 2 in recognition performance validates the pacing procedure as a method of controlling reading rate without serious consequences for natural reading processes. Beyond these results the data demonstrated a number of other important points. Subjects were able to successfully form macrostructure memory representations of stories even when skimming at 600 wpm and even when reading alternating case. The use of alternating case, however, did not go unnoticed as recognition performance was lower when subjects read alternating case. The effect of alternating case appeared to be most pronounced in the skimming condition. Furthermore, when subjects read narratives in

alternating case, reading rate and statement type interacted: a greater reduction in recognition of microstatements than of macrostatements or inferences was observed as reading rate increased. True macrostatements were generally verified faster than other statement types while false inferences were rejected faster than other statements. Recognition of surface structure suffered badly as reading rate increased and dropped to chance levels in some cases even in the skimming condition. Alternating case also had a detrimental effect on surface memory.

The formation of a story macrostructure in reading appears to be a general goal in reading. Subjects in Experiment 2 continued to show patterns of recognition that indicated a macrostructure, albeit incomplete, had been formed even under adverse circumstances of rapid skimming and alternating case. In fact, when these two elements were combined in reading narratives, the distinction of important and unimportant information was accentuated: recognition of microstatements declined more rapidly than recognition of important statements as reading rate increased. That subjects did not show this effect in their reading of newspaper stories suggests they were able to make more efficient use of newspaper story structure to swiftly ascertain the important points by concentrating processing resources on information in the initial part of the story, and to then read more detailed information with greater speed and a smaller investment of processing resources. In reading the narratives, subjects had to process the whole story for important information rather than just the first part, as in newspaper stories, and had little time for processing of details. In this case, allocation of processing resources and time could not be concentrated on a selected portion of the text but had to be available at any time that an apparently important fact was encountered. It would be a rare event for a significant amount of resources to be allocated to processing of detail information. The differential effects of narratives and newspaper stories supports the suggestion (Thorndyke, 1979) that readers may possess different story schemata for different kinds of stories. At least subjects have knowledge about the characteristics of different story types and can use this knowledge to streamline processing and determine policies of resource allocation. Whether the knowledge forms a single schema or multiple schemata is another issue.

The effects of alternating case on comprehension and memory were most pronounced in the skimming condition, consistent with the hypothesis that alternating case interferes with whole word visual identification and that faster reading depends on whole word processing to a greater degree than normal reading. At the fast skimming rate, however, alternating case had no effect implying either that whole word processing does not play an important role in skimming at such high rates or that readers were so pressured at the fast skimming rate that only processing at a global level could be done accurately and allotment of significant resources to detailed processing of surface structure

was not feasible (alternating case would be expected to affect detailed processing of surface structure). In fact, memory for surface structure was quite poor in the fast skimming condition, and was at or near chance level for almost all statement types and experimental conditions. It is particularly interesting to note that surface memory approached chance performance levels at the skimming rate in the alternating case condition, while more knowledge about surface structure was apparent in the normal case condition. This result is consistent with the hypothesis that controlled processing (Shiffrin & Schneider, 1977) of surface structure of statements during skimming was not feasible when alternating case was used since decoding processes were less efficient. Whole word visual identification was disrupted and subjects were less frequently able to completely process the surface structure of statements. A partial processing scheme likely was in effect, in which some part of a statement's surface structure was processed in order to comprehend the important information contained in it. This processing would be highly predictive and dependent on minimal input from data driven processes. As a result, processing of semantic information may have been incomplete or incorrect as the reader relied heavily on prior knowledge about the story topic. The overall consequence was lower performance on memory for meaning and memory for surface structure tests. Despite the frequency of inaccurate or incomplete processing, though, subjects were able to form a basic macrostructure of stories they skimmed and this representation reflected at least part of the fundamental content of the stories.

The reaction time data from the meaning test revealed two interesting effects. The time taken to verify macrostatements was less than that required for inferences or microstatements, indicating that macrostatements represented a critical part of macrostructure representations of stories. While inferences also represented important information, the longer reaction time suggests that they may not have been directly represented in memory but rather the pieces of information required to draw an inference might have needed to be combined to verify the inferential statement presented during the test (cf. Hayes-Roth & Walker, 1979; Reder, 1979). The other reaction time effect pertains to the fact that false inferences were rejected more rapidly than false statements based on explicitly presented information. False inferences could be more quickly rejected because they dealt with important facts that likely would be a basic part of a macrostructure and because an inference required subjects to combine these facts or to alter a fact in an unacceptable way. Thus, the discrepancy between the false inference and the macrostructure could be detected even as the subject was first reading the statement. False statements based on explicitly presented information were fundamentally correct except for one aspect of each statement and the correctness of the majority of a statement's information probably prompted more careful comparison with the memory representation of the story.

The general conclusion from Experiment 2 is that subjects are able to form limited macrostructure representations as they read at very rapid speeds and even when reading a typography that interferes with efficient decoding processes. The problem of reading under such circumstances appears to be solved by implementation of a processing resource allocation policy of which processing surface structure to a lesser degree and increasing reliance on conceptually driven processes for predictive reading and development of a macrostructure are characteristic. To the extent that the reader can read predictively, sentence processing will be more efficient and the long process and high resource requirements of detailed data driven analysis can be de-emphasized. The reader can then form a general macrostructure that is centered around important information as defined by a story schema or knowledge of the story's topic, although a price has been paid in terms of completeness and accuracy of the macrostructure as a representation of a story. Unlike a summary given after normal reading of a story, a macrostructure that results from skimming would be more likely to contain incorrect information even in macropropositions due to the highly predictive nature of the reading process and the slackening of constraints usually imposed by data driven processes. Despite these deficiencies, a reasonably accurate macrostructure representation can be formed while skimming.

CHAPTER IV

EXPERIMENT 3

In Experiments 1 and 2 subjects were assigned the general task of extracting important information from stories while reading at normal or skimming rates. The task was, apparently, not vaguely defined as subjects consistently formed macrostructures that represented key information contained in the stories. In Experiment 3 a more specific reading goal was introduced, however, in order to observe how subjects normally read and skim for particular kinds of information. When required to read a story for a certain type of information subjects need to use not only their knowledge about story structure but also their knowledge about the target information so that selective processing is a realistic possibility. A reader must be able to establish expectations about the sorts of information to be encountered while reading so that this information can rapidly be processed and incorporated into a memory representation. If a reader knows little about a topic it will be more difficult both to select and to form a representation of information relevant to the topic.

Reading stories from a particular perspective affects comprehension and memory processes in a way that favors information relevant to the chosen perspective. Processing of information relevant to a consistently maintained perspective is faster (Black, Turner, & Bower, 1979) and less resource-demanding

(Cutler & Fodor, 1979) than processing of irrelevant sentences. Relevant information is more likely to be retrieved from memory after reading than is irrelevant information (Anderson & Biddle, 1975; Anderson & Pichert, 1978; Pichert & Anderson, 1977; Kozminsky, 1977; Poulton, 1958). In fact, when subjects attempt to recall statements that are inconsistent with the emphasized perspective, errors are made that make the recalled information more consistent with the perspective (Black et al., 1979). Although retrieval from memory is biased in favor of relevant information, at least part of the effect is due to a problem of information access rather than availability in the memory representation. Anderson and Pichert (1978) demonstrated that subjects could be prompted to recall more irrelevant information from a story that had been read if they were given a descriptive title that was relevant to that information. Thus, while irrelevant information may not occupy prime locations in a memory representation of a text, that information appears to be processed and represented to at least some degree.

Experiment 3 was designed to further investigate the extent to which relevant and irrelevant information is processed and represented in memory when subjects read stories from specific perspectives. It was especially important to determine how, or even if, differential processing of relevant information is achieved while skimming. Carver (1977) has claimed that at a variety of reading rates little effect on the comprehension process is to be found as a function of the reader's purpose.

Certainly the research discussed in Chapter I concerning the flexibility of the reading process argues against Carver's position as do the results concerning effects of relevance on recall. The issue of whether selective processing occurs during increased reading rates remains, however, as previous work on effects of relevance has not included reading rate variables. It is possible, for example, that when subjects skim material they cannot effectively process relevant and irrelevant information differentially. Time pressures may force inadequate processing of relevant information as subjects must move on to find new relevant information once a relevant statement has received minimal processing. Minimal processing would also be devoted to irrelevant statements in order to reject them as not pertinent to the reader's goal. Unless the reader is able to process relevant statements more elaborately and incorporate them into a memory structure based on his or her goal, effects of relevance on comprehension and memory will be modest.

The manipulation of typographical case was also included in Experiment 3. It was found in Experiment 2 that reading alternating case had a detrimental effect on memory for meaning and surface structure, especially when subjects skimmed the stories at a moderate rate. The effect appeared to be due to interference with efficient decoding processes such as whole word visual identification. Subjects were able to form macrostructures, however, even when reading alternating case. The issue in Experiment 3 was whether the control processes and

schemata involved in selective processing of relevant information can operate as well as macro-operations under the demanding conditions of alternating case. The process of forming macrostructures may be inherent in the reading system and relatively robust with respect to reasonable reading rate and decoding constraints. On the other hand, processes involved in selecting relevant information to form the basis of a comprehensible memory representation biased in favor of a particular perspective may be less well practiced. To the extent that a reader is not practiced at reading particular kinds of stories from specific perspectives, each new combination of story and perspective is a new situation requiring possibly infrequently occuring combinations of schemata in comprehension and formation of a memory representation. A greater degree of attention would be required to selectively process the relevant information than would be required to carry out a reading task that merely involved formation of a macrostructure. Consequently, the added demands on processing resources and attention introduced by the use of alternating case might have the effect of causing selective processing of relevant information to fail. When skimming, subjects may be forced to attend to processing of surface structure to such a degree that not sufficient attention or cognitive resources (Kahneman, 1973; Norman & Bobrow, 1975) would be available for use in guiding the interaction of schemata controlling selection of relevant information and its incorporation into a developing memory structure. The consequence for resulting memory representations would be that subjects would include both relevant and irrelevant information in their representations. A recognition test was used to examine subjects' memory representations of stories so that a sensitive estimate of the relative availability of relevant and irrelevant information could be obtained without the strong constraints of retrieval biases that likely would be operating in a recall test (cf. Anderson & Pichert, 1978).

Finally, an effort was made in Experiment 3 to obtain information about reading strategies that subjects use when their task is to skim rather than carefully read a story. The questionnaire administered to subjects upon completion of the memory tests asked for a description of the subjects' skimming strategies that were normally used outside the laboratory and those subjects in the skimming conditions were also asked whether and how their usual skimming strategies differed from those used in the experiment. It was hoped that these descriptions of skimming strategies would provide a general indication of how subjects believe they go about successful skimming and of how skimming strategies are adjusted to meet various experimental demands.

The fact that relevant information was defined by an arbitrary perspective and that the statements relevant to the assigned perspective were embedded among irrelevant statements may serve to reduce relevance effects relative to what might be obtained if more distinctive definitions of relevance and textual

cues such as headings were used to guide selection of relevant information. The results of Experiment 3 should be interpreted in light of this consideration.

Method

<u>Subjects</u>

The subjects in Experiment 3 were 192 students recruited and assigned to experimental conditions in the same manner as in Experiments 1 and 2. There were 32 subjects assigned to each of the six combinations of reading rate and typographical case.

<u>Materials</u>

The 500-word narrative used in the preliminary reading task in Experiments 1 and 2 was used in Experiment 3 as well. Four other narratives of about 355 words each were developed for use in the experiment proper. Two of the stories were used by Pichert and Anderson (1977) and the other two were used by Kozminsky (1977). These stories were selected because they could be read from at least two different perspectives depending on the specific title that preceded the story. The four stories were altered by adding or deleting material to make them conform to the length restriction. In addition, two practice stories were selected and prepared so as to meet the length restriction imposed on the critical stories. An arbitrary title was assigned to each practice story.

A set of test statements was developed for each of the critical stories, with half of the statements in each set selected to be relevant to one possible perspective of the story and the other half relevant to the other possible perspective. For the Pichert and Anderson stories relevance was determined by fairly obvious intuitions about which statements would be closely related to specific perspectives. For the Kozminsky stories relevance was determined by his the recall protocols generated by his subjects. Statements recalled most often when a particular perspective was suggested before reading the story were taken as relevant to that perspective. In this way 16 statements were selected from each story, with eight relevant to one of a story's two perspectives. Each subset of eight statements was further divided into two groups of four statements. One group of four statements comprised the true statement test set and the other group of four formed the false statement test set. Each of the statements assigned to false statement test sets was rewritten in two versions as in Experiment 1, with one version semantically the same as the original but altered in surface form and the other version similar to the original in surface form but altered so as to be false with respect to the original statement's meaning. Statements written in paraphrase version were used in the surface memory test and those written with basic semantic alterations were used in the memory for meaning test.

Thus, 24 test sentences were developed for each critical story. Eight statements formed the true test set with each half relevant to a different one of the two possible perspectives. Eight statements formed the paraphrase set for the surface test

with each half relevant to a different perspective. The last eight statements formed the false set for the meaning test with each half relevant to a different perspective. An example of one of the critical stories, its two perspective-inducing descriptions, and the corresponding set of test statements are presented in Appendix B. The mean number of words in each of the true statements was 10.5, the mean for the paraphrase statements was 9.6, and the mean for the semantically false statements was 9.5.

The four critical and two practice stories were typed single spaced, right and left justified, on separate pages. A general, neutral title was typed centered at the top of each story page. Two title pages for each critical story were prepared with the title and one of the perspective-inducing descriptions typed centered on each. One title page was similarly prepared for each practice story. Enough copies of title and story pages were produced to form 16 story booklets. Eight booklets contained one set of perspectives while the other eight contained the other set. Each group of eight booklets was arranged to represent the eight different orderings of stories generated by the following counterbalancing scheme. Four booklets represented all four possible orderings of stories with the restrictions that the Kozminsky stories occurred in the first and fourth positions and the Pichert and Anderson stories occurred in the second and third positions. Four booklets represented all four possible orderings with the reverse set of restrictions. Thus, eight counterbalanced orderings of stories were produced and assigning the two different sets of perspectives to each ordering produced 16 story booklets. The booklets were arranged so that the two practice stories appeared first, then the four critical stories in proper order. The appropriate title page appeared before each story and a blank colored page was inserted in front of and behind each title page. Another set of 16 story booklets was prepared in the same way, except that the text of each story was typed in alternating case as in Experiment 2. In total, 32 story booklets were prepared. A page summarizing the instructions for reading the stories was appended to the front of each booklet.

Design

A 3 X 2 X 2 mixed factorial design was used for the meaning and surface tests. Reading rate (normal, skimming, and fast skimming) and typographical case (normal and alternating) were between-subject factors, and statement type (relevant and irrelevant) was a within-subject factor.

Procedure

Subjects participated in groups ranging in size from one to six. As in Experiments 1 and 2 each subject began the experimental session by reading the practice narrative. When this task was completed the materials were collected and the subjects were told that in the next part of the experiment they would be required to read a set of stories at a fixed rate of speed. The instructions and procedure for the reading phase were similar to those used in Experiment 2, except that subjects were told to

read each story from a particular perspective. It was emphasized that subjects should try to derive from each story information relevant to the given perspective. The recorded pacing tapes were constructed as in Experiment 2. They were set for reading rates of 225 wpm, 375 wpm, and 600 wpm. The cue word "Next" prompted subjects to turn to and read the title page of the next story. After 5 sec the cue word "Start" prompted the subjects to turn to the story page and begin reading. A period of 5 sec was inserted at the end of each story.

At the end of the reading phase subjects were tested in the same manner as in Experiments 1 and 2. The test instructions did not mention anything about inferences as none were tested. Otherwise, the instructions basically were the same as those used in earlier experiments. Test statements from the fourth and then the first critical stories the subject read were tested in the surface test then the statements from the second and then the third stories were tested in the meaning test. As in Experiments 1 and 2, all test statements for a particular story were tested consecutively in a random order and preceded by a message containing the story title.

When the subject finished the test phase he or she filled out the questionnaire used in the first two experiments. In addition, this questionnaire included two final questions. The first asked the subject to briefly describe the method he or she normally uses to skim material outside the laboratory. The other question was relevant to subjects in the skimming and fast skimming conditions and asked whether subjects were able to use their normal skimming strategies in the experiment. If subjects had altered their strategies they were also asked to describe how their experimental strategy differed from their usual strategy.

Results and Discussion

Questionnaire

As in Experiment 2, subjects in Experiment 3 were required to read the critical stories at paced rates of speed rather than at preferred reading rates. To check the representativeness of subjects' reading behavior in Experiment 3 the answers to a number of questionnaire items were analyzed. First, subjects' comparisons of usual reading rates to rates used in the experiment were examined. The percentage of subjects in each experimental condition claiming that the paced reading rate was the same as, faster than, or slower than usual reading rates are shown in Table VII. As in Experiment 2, subjects reading at the normal rate generally felt they were reading at or slightly above their normal rates, with a greater tendency among those subjects who read alternating case to claim that the pace was too fast. It does appear, therefore, that the selected rate for normal reading speed represents a reasonable estimate of the sampled population's normal reading rate. For subjects in the skimming and fast skimming conditions there was nearly uniform agreement that the paced rates were too fast. This agreement seems to stem from the same misinterpretation of the questionnaire item as that

TABLE VII

SUMMARY OF RESPONSES TO QUESTIONNAIRE ITEMS REGARDING REACTION TO PACING PROCEDURE IN EXPERIMENT 3

Case	Reading Rate	Paced Reading Rate Relative to Own Rate (Percent of Subjects)			Experimental Skimming Success Ratings
		Same	Fast	Slow	
Normal	Normal	50	9	41	--
Normal	Skim	6	91	3	1.81
Normal	Fast Skim	0	100	0	1.34
Alternating	Normal	37	38	25	--
Alternating	Skim	6	91	3	1.59
Alternating	Fast Skim	0	100	0	0.41

which occurred in Experiment 2. That is, most subjects in the skimming and fast skimming conditions probably compared the paced rates to their usual normal rates of speed. Thus, these subjects almost always claimed that the paced rates were too fast.

A more valid estimate of the representativeness of the skimming and fast skimming rates of natural skimming rates can be obtained by evaluating subjects' responses to the questionnaire item asking how successful subjects believed they were in extracting important information from the stories while skimming during the experiment. These responses were evaluated using the same four-point scale as that used in the first two experiments. The mean rated subjective success in selecting important information while skimming in the experiment is presented for subjects in the skimming and fast skimming conditions in Table VII. The mean for subjects in the skimming condition who read normally typed stories was 1.81, which is very close to the overall mean rating of success while skimming outside the laboratory (1.90) based on all subjects in the experiment. The subjects who read alternating case, however, did not show the same degree of confidence in their success at selecting important information while skimming in the experiment. The overall mean for subjects skimming or fast skimming normal case was 1.58, while for subjects skimming or fast skimming alternating case the mean was 1.00. This was a reliable difference, $\underline{F}$ (1, 124) = 9.24, $\underline{MS}_{\underline{e}}$ = 1.158. Furthermore, subjects in the skimming condition were more confident in their skimming success (1.70) than were

subjects in the fast skimming condition (0.88), $\underline{F}$ (1, 124) = 18.96, $\underline{MS}_{\underline{e}}$ = 1.158. The interaction between rate and case approached significance, $\underline{F}$ (1, 124) = 3.57, $\underline{MS}_{\underline{e}}$ = 1.158, $\underline{p}$ < .06. The interaction indicates that there was a tendency for the detrimental effect of reading alternating case to be stronger among subjects in the fast skimming condition. Thus, the subjective ratings of skimming success indicate that at least for normally typed stories the rate of 375 wpm is a representative skimming rate for subjects in Experiment 3, while the rate of 600 wpm was apparently beyond usual effective skimming rates. As in Experiment 2, these are the roles that the skimming and fast skimming rates were meant to play.

Regarding responses to questionnaire items pertaining to subjects' skimming habits outside the laboratory, the mean estimate of percentage of total reading time that is spent skimming was 33%, with a range of 0-95%. The most commonly mentioned materials that subjects skim are newspaper and magazines, as 30% of the subjects claimed to skim one or both of these types of material. Parts of college work were mentioned by 16% of the subjects as materials for skimming. As in the first two experiments, a number of other types of material were mentioned idiosyncratically. As mentioned earlier, the mean estimate of these subjects' confidence in their success in selecting important information while skimming on their own was 1.90 on the four-point scale used in the earlier experiments. This indicates a reasonably high level of confidence in effectiveness of skimming.

Of the subjects in Experiment 3, 20% indicated that they had received some form of training in improving reading speed. Among these subjects the mean rating of confidence in success while skimming outside the laboratory was 1.86, very close to the overall mean rating. With respect to these subjects' confidence in the effectiveness of their skimming stories in the experiment, those who were paced at 375 wpm ($N = 12$) had a mean rating of 2.67, and those paced at 600 wpm ($N = 15$) had a mean rating of 0.93. This difference was highly reliable, $F\ (1, 25) = 25.55$, $MS_e = 0.784$, indicating that even for those subjects who have received training in speeded reading, the rate of 600 wpm was too fast to allow effective selection of important information.

Finally, subjects' responses to the questionnaire items regarding skimming strategies were very revealing. Descriptions of skimming methods used while skimming outside the laboratory fell into three general categories. The most commonly mentioned strategy involved some form of selective skimming such as reading particular paragraphs or sentences, searching for key words, reading important points carefully, or reading headings. Of the 188 subjects responding to the questions on skimming methods, 79% stated that they used some form of selective strategy. Glancing over all words, or simply reading everything faster than usual was a strategy mentioned by 10% of the subjects. Another 7% of the subjects claimed to use the method of skipping over various parts (words, sentences, etc.) of a text indiscriminantly or in a

structural pattern without regard to semantic content in an effort to obtain the gist. Only 4% of the subjects stated that they use a method that involves perceiving chunks of information (sentences or phrases) simultaneously. Two subjects specifically mentioned that they did not subvocalize while skimming. Of those subjects claiming to use a selective skimming strategy, 53% based selection on certain words that are important or relevant to goals in skimming and then read contiguous words, 18% selected only specific sentences or paragraphs in certain locations in the text (e.g., first and last) to read carefully, and 29% mixed the two strategies by reading particular sentences or paragraphs carefully and skimming the remaining material for important points. When subjects in the skimming and fast skimming conditions were asked if the skimming strategies they used in the experiment differed from their usual methods, almost all subjects gave responses indicating that there was no fundamental change, that the rate was faster, or that the alternating case made the task more difficult. None of the subjects who read normal case in the skimming condition stated that their skimming strategy had changed fundamentally. Three subjects in the skimming condition who read alternating case described fundamental changes in their skimming strategies. One subject who selected specific sentences or paragraphs for careful reading and one subject who skipped material without regard to content changed to the method of selecting key words. One subject who read chunks simultaneously changed to the method of swiftly glancing over the whole text.

Five fast skimming subjects who read normal case claimed that their skimming methods changed. Two changed from a mixture of selecting key words and specific sentences or paragraphs to a strategy of selecting only key words. One changed from the mixed selection strategy to selecting only specific sentences or paragraphs. One subject who ordinarily selected key words changed to rapidly glancing over the whole text, and one subject made the reverse switching of strategies. In the fast skimming of alternating case condition, only one subject claimed to have altered skimming strategies and this was a change from selecting specific sentences to selecting key words.

Four general conclusions can be drawn from the results of the questionnaire: (a) the extra-experimental skimming practices of subjects in Experiment 3 are very similar to those of subjects in Experiments 1 and 2, (b) the paced normal and skimming rates appeared to be representative of the subjects' usual rates and subjects, for the most part, were able to maintain their basic reading strategies, while the fast skimming rate was beyond comfortable skimming rates, (c) the use of alternating case reduced subjects' confidence in their success at skimming for relevant information during the experiment, and (d) the most common skimming strategies involve selection of certain aspects of the text such as key words or sentences in particular locations in the text and when pressured by time or the use of alternating case those subjects who alter their skimming strategies most often change to the method of selecting key words

and reading contiguous information. These conclusions encourage the assumption that subjects reading normal case at the normal or skimming rates were able to maintain their usual reading strategies and were reading at rates close to their usual speeds for normal reading and skimming. The pacing procedure, then, did not appear to severly alter usual reading processes and the selected rates were representative of normal, skimming, and fast skimming rates ordinarily used by these subjects.

<u>Memory for Meaning</u>

<u>Recognition performance</u>. The mean proportion of hits and false alarms for relevant and irrelevant statements in the meaning test are presented in Table VIII. Performance also was characterized by calculating d' scores for each subject as in Experiments 1 and 2. The mean d' scores are presented in Figure 22.

An analysis of variance of d' scores was carried out with reading rate (normal, skimming, and fast skimming) and typographical case (normal and alternating) as between-subject factors and statement type (relevant and irrelevant) as a within-subject factor. The main effect of reading rate was reliable, $\underline{F}$ (2, 186) = 28.93, $\underline{MS}_e$ = 0.697, indicating that as reading rate increased, recognition performance decreased. A Newman-Keuls test indicated that the mean d' score for the normal reading rate condition (1.62) was significantly higher than for the skimming condition (1.21) and each of these was reliably higher than the mean for the fast skimming condition (0.83). The

TABLE VIII

MEAN HIT AND FALSE ALARM RATES ON MEANING TEST IN EXPERIMENT 3

Case	Reading Rate	Statement Type			
		Relevant		Irrelevant	
		Hit	FA[a]	Hit	FA
Normal	Normal	.83	.29	.76	.28
Normal	Skim	.75	.30	.75	.32
Normal	Fast Skim	.70	.36	.60	.39
Alternating	Normal	.86	.19	.80	.27
Alternating	Skim	.71	.32	.75	.35
Alternating	Fast Skim	.72	.41	.69	.39

[a]FA = False Alarm

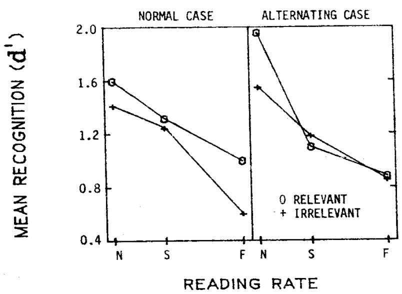

Figure 22. Mean recognition performance on meaning test as a function of relevance in Experiment 3 (N = Normal; S = Skim; F = Fast Skim).

only other reliable effect was the main effect of statement type, with recognition of relevant statements (1.30) significantly higher than recognition of irrelevant statements (1.14), F (1, 186) = 4.78, MS_e = 0.551.

Despite the failure of any interaction involving typographical case to reach significance, Figure 22 suggests that the effects of reading rate on the differential recognition of relevant and irrelevant statements may depend on whether the stories were typed normally or in alternating case. Two further analyses of variance were calculated so that the possibility of these differential effects could be more thoroughly investigated. Separate analyses were carried out for subjects who read normally typed stories and for those who read stories in alternating case. In each analysis reading rate (normal, skimming, and fast skimming) was a between-subject factor and statement type (relevant and irrelevant) was a within-subject factor. Both analyses found a reliable effect of reading rate and the patterns of differences between means based on Newman-Keuls tests basically conformed to the pattern observed in the full analysis. In addition, the analysis involving subjects in the normal case condition indicated that the effect of statement type approached significance, F (1, 93) = 3.29, MS_e = 0.677, p < .07, with recognition of relevant statements (1.30) greater than recognition of irrelevant statements (1.08). The interaction between rate and statement type was not reliable (F < 1). On the other hand, the analysis involving subjects in the alternating

case condition did not reveal a main effect of statement type, while the rate by statement type interaction approached significance, F (2, 93) = 2.50, MS_e = 0.424, p < .09. This interaction implies that subjects who read alternating case better recognized relevant statements than irrelevant ones when reading at normal speed, but when required to skim at any rate reasonably beyond the normal rate the superiority of relevant statements was lost. The patterns of reading rate effects on the recognition of relevant versus irrelevant statements apparent in Figure 22 are, therefore, at least moderately supported by the subsidiary analyses.

In Experiment 2 when subjects were required to read narratives typed in alternating case it was found that as reading rate increased, the superiority of recognition of inferences and macrostatements over microstatements <u>increased</u>. In contrast, when subjects in Experiment 3 were required to read a different set of narratives in alternating case and from specific perspectives, the superiority of recognition of relevant statements over irrelevant statements was <u>eliminated</u>. These disparate effects could be due to the use of different materials or subjects, or to the introduction of a new task, namely, skimming for information relevant to a given perspective. In order to separate these possibilities and to determine which underlies the differential effect of rate on importance or relevance, a post hoc analysis was carried out. The objective of this analysis was to select macro- and microstatements from among those statements used as test items in Experiment 3. Recognition of these two statement types as a function of reading rate and typographical case would then be analyzed. If subjects show higher recognition of macrostatements than of microstatements across all reading rates, then it would be reasonable to conclude that the elimination of the relevance effect among subjects reading alternating case was not due to subjects or story materials. Rather, the elimination of the relevance effect would be closely related to the differences in task demands when instructed to read for the gist of a story as opposed to reading for information relevant to a particular perspective.

The task of selecting macro- and microstatements from the body of test statements used in Experiment 3 was accomplished by obtaining importance ratings for each statement. Paraphrase and false statements were rewritten in their original forms for the purpose of obtaining ratings. Ten sophistocated subjects (most of whom were graduate students in psychology) read each story and rated its corresponding set of test statements using the same scale as that described in Experiment 1. Three sets of ratings were obtained for each story. The first was based on a story's gist and the other two were based on each of experimental perspectives. The latter two sets of ratings basically agreed with original relevance classifications without providing strong distinctions between important and unimportant relevant statements, for example, and will not be discussed further. The mean rated general importance of each statement was computed, and

on the basis of the six-point rating scale (1 = unimportant, 6 = very important) statements were labeled as macrostatements, microstatements, or neither. Statements with mean ratings of 4.5 or greater were classified as macrostatements and those with a mean rating of 3.5 or lower were classified as microstatements. It was found that two of the stories, one taken from Pichert and Anderson (1977) and one from Kozminsky (1977), did not have any statements that were classifiable as macrostatements. These stories appeared, therefore, to exhibit less than acceptable amounts of narrative structure and no statements from these stories were used in the post hoc analysis. From the two remaining stories six true macrostatements, five true microstatements, and five false statements of each type were obtained. Since each of the two remaining stories provided at least two statements in each cell resulting from the combination of statement type and truth value, and since each subject was tested on one of the two stories in the meaning test, it was possible to include all subjects in the post hoc analysis. A control analysis of recognition of relevant and irrelevant statements based only on the two selected stories was carried out, and the same pattern of results as those shown in Figure 22 was obtained.

The mean hit and false alarm rates on the selected macro- and microstatements are shown in Table IX. Also, d' scores based on hit and false alarm rates were calculated for each subject in the same manner as described for earlier experiments. The mean d'

TABLE IX

MEAN HIT AND FALSE ALARM RATES ON MEANING TEST AS A FUNCTION OF POST HOC IMPORTANCE RATINGS IN EXPERIMENT 3

Case	Reading Rate	Statement Type			
		Macrostatement		Microstatement	
		Hit	FA[a]	Hit	FA
Normal	Normal	.90	.32	.66	.25
Normal	Skim	.86	.29	.65	.24
Normal	Fast Skim	.79	.38	.55	.44
Alternating	Normal	.95	.25	.72	.25
Alternating	Skim	.88	.39	.58	.34
Alternating	Fast Skim	.98	.46	.52	.41

[a]FA = False Alarm

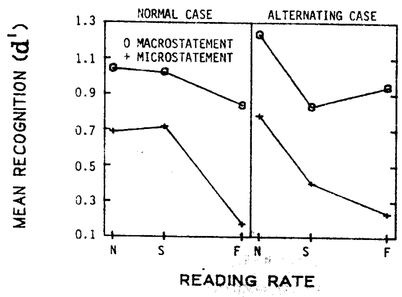

Figure 23. Mean recognition performance on meaning test as a function of post hoc importance ratings in Experiment 3 (N = Normal; S = Skim; F = Fast Skim).

scores are shown in Figure 23. An analysis of variance of the d' scores was carried out including reading rate (normal, skimming, and fast skimming) and typographical case (normal and alternating) as between-subject factors and statement type (macrostatement and microstatement) as a within-subject factor. In the analysis the main effect of reading rate was reliable, $F(2, 186) = 10.99$, $MS_e = 0.441$. A Newman-Keuls test indicated that subjects reading at the normal rate had reliably higher recognition (0.93) than did subjects in the skimming condition (0.74), and subjects in each of these two reading rate conditions had reliably higher recognition than subjects in the fast skimming condition (0.54). More important, recognition of macrostatements (0.98) was significantly higher than recognition of microstatements (0.50), $F(1, 186) = 43.87$, $MS_e = 0.518$. The interaction involving rate and statement type was not reliable. Finally, there was a significant interaction between reading rate and case, $F(2, 186) = 3.27$, $MS_e = 0.441$. This interaction followed the same pattern as the rate by case interaction described in Experiment 2. At the normal reading rate, case had a mildly facilitative effect (cf. Kolers, 1975b) on recognition (normal case = 0.86, alternating case = 1.00). Alternating case was detrimental at the skimming rate (normal case = 0.87, alternating case = 0.62) but had no effect at the fast skimming rate (normal case = 0.50, alternating case = 0.59). Another way to interpret this interaction is to consider the amount of decrement in recognition as reading rate increases. For normal

case there was virtually no decrement as the rate increased from normal to skimming, but a rather large decrease in recognition occurred when the rate reached fast skimming. For alternating case, however, a large decrease in recognition was found when the rate changed from normal to skimming and only a small decrease was observed when the rate was moved to fast skimming. The tendency for the decrease in recognition of meaning performance as a function of reading rate to decline markedly at a lower reading rate when alternating case was used was also found in Experiment 2 and was apparent (though not statistically supported) in the recognition of relevant and irrelevant statements in Experiment 3. The consistency of this effect suggests that a fundamental change in processing occurs when subjects begin to skim alternating case. Specifically, it appears that as soon as subjects skim, the role of whole word visual identification processes is highlighted and any manipulation (e.g., use of alternating case) that interferes with these processes will cause a general decrease in the completeness and accuracy of a memory representation.

The results of the two sets of analyses of recognition of meaning in Experiment 3 strongly argue in favor of the suggestion that readers can more readily extract macrostructure information from stories than they can selectively process statements relevant to a given perspective. One probable cause of this effect is the relative degree of resource requirements in reading for gist compared to reading for particular kinds of information.

The processing demands appeared to be greater in the task of reading for particular kinds of information, as selective processing was prevented when the added demands of alternating case were imposed. Notice that comprehension and recognition degraded gracefully (Norman & Bobrow, 1975) and did not suffer a complete breakdown as demands on resources exceeded readers' processing capacities. It is as if the subjects found that the task of decoding alternating case and reading from an assigned perspective could not both be adequately executed and they therefore allocated processing resources (Navon & Gopher, 1979) in a way that would allow at least some information to be decoded. That is, rather than attending to relevance to a great extent, subjects processed information in a general way, forming macrostructures and permitting selective processing of relevant information to lapse. Resources for selection of relevant information and development of a special comprehension structure based on that information were not available due to the added demands of reading alternating case. In keeping with the alternatives discussed by Navon and Gopher (1979), however, it is also possible that decoding and comprehension processes do not draw on a single common pool of cognitive resources but, instead, overlap in their demands on processing mechanisms that possess their own resource capacities. This possibility, in fact, is more favorable to the general theory of interactive processes which emphasizes interdependence of subskills such as decoding and comprehension and mutual dependence on various other processing

mechanisms.

Irrespective of the problems caused by reading alternating case, subjects who read normally typed stories were able to maintain a strategy of selective processing of information even when skimming at 600 wpm. The size of the relevance effect was not as large or as stable as the importance effect observed in comparisons of macro- and microstatements, but it does provide evidence that readers are capable of selecting and elaborately processing statements pertaining to a specific goal while skimming. Subjects were able to determine which aspects of a text were relevant to the perspective and were able to elaborate the processing of those statements to form a memory representation centered around relevant information. Selective processing of relevant information occurred in the context of macrostructure development (note the sizable importance effect in the post hoc analysis) and, while there is no direct evidence, this finding is consistent with the possibility that selective processing of information relevant to a specific goal is done in conjunction with and possibly under the guidance of processes responsible for developing macrostructure representations.

<u>Reaction time for hits</u>. Reaction time data in Experiment 3 were prepared in the same way as data in the previous experiments, except that any reaction time value in excess of 15 sec was reduced to 15 sec before a subject's mean reaction time for a specific statement type was calculated. The mean time taken to verify true statements is shown in Figure 24. The analysis of

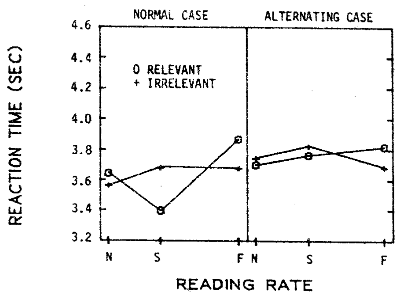

Figure 24. Mean reaction time on hits on meaning test in Experiment 3 (N = Normal; S = Skim; F = Fast Skim).

variance of these data found no significant effects. Of interest here is the fact that regardless of reading rate or case, once subjects had established a memory representation of a text's meaning any aspect of the representation that was accessible could be recognized quickly. This result is unlike the tendency in Experiment 2 for subjects to verify macrostatements more rapidly than microstatements. The decrease in reaction time to relevant statements in the normal case skimming condition apparent in Figure 24 was not reliable, and probably represents a speed-accuracy trade-off effect (Pachella, 1974) in that the relevance effect in that condition was lower than for the normal rate and fast skimming conditions. It is possible that memory representations based partly on relevance do not draw as strong a distinction between relevant and irrelevant information as do representations based only on macrostructures.

Reaction time for correct rejections. The mean reaction time for correct rejections is presented in Figure 25. No reliable effects were found in the analysis of variance. If a statement was not consistent with information in a story, regardless of its relevance to the assigned perspective, it was rejected after about 4 sec of processing.

Reaction time for false alarms. The mean time taken by subjects to incorrectly accept a false statement is presented in Figure 26. An analysis of variance found that only one effect approached significance as reading rate and case tended to interact, $F(2, 186) = 2.75$, $MS_e = 2.766$, $p < .07$. As reading

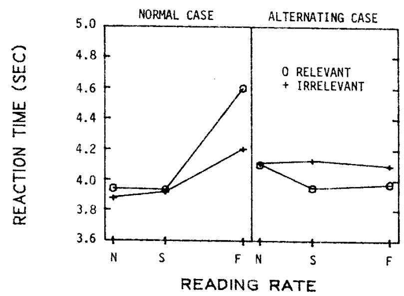

Figure 25. Mean reaction time on correct rejections on meaning test in Experiment 3 (N = Normal; S = Skim; F = Fast Skim).

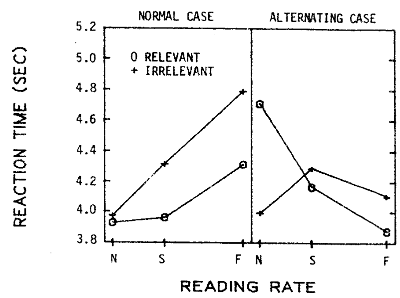

Figure 26. Mean reaction time on false alarms on meaning test in Experiment 3 (N = Normal; S = Skim; F = Fast Skim).

rate increased reaction time in the normal case condition increased, indicating a greater degee of caution and uncertainty about memory representations. In the alternating case condition, though, increased reading rate led to lower reaction times, implying that subjects had sufficiently low confidence in their memory representations that they would more rapidly accept statements that appeared at least partially correct.

Reaction time for misses. The mean reaction time for failing to verify a true statement is shown in Figure 27. The only reliable effect in the analysis of these data was the statement type effect, F (1, 186) = 6.06, MS_e = 2.201. Subjects considered true relevant statements longer (4.484) before rejecting them than they did true irrelevant statements (4.112). This effect may represent a facet of the memory representation of a story after reading for particular information or a response bias to consider apparently relevant test statements more carefully than irrelevant ones.

Memory for Surface Structure

Recognition performance. Hit and false alarm rates were calculated for each subject as in earlier experiments, and the mean rates are presented in Table X. A d' score for each statement type was calculated for each subject as in earlier experiments and the mean d' scores are presented in Figure 28. The analysis of these scores found only a main effect of reading rate to be reliable, F (2, 186) = 4.95, MS_e = 0.398. As reading rate increased, recognition performance decreased. A Newman-Keuls

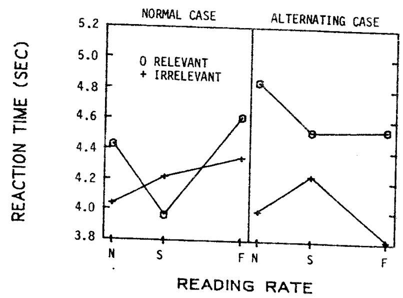

Figure 27. Mean reaction time on misses on meaning test in Experiment 3 (N = Normal; S = Skim; F = Fast Skim).

TABLE X

MEAN HIT AND FALSE ALARM RATES ON SURFACE TEST IN EXPERIMENT 3

Case	Reading Rate	Statement Type			
		Relevant		Irrelevant	
		Hit	FA[a]	Hit	FA
Normal	Normal	.69	.52	.64	.48
Normal	Skim	.60	.52	.64	.48
Normal	Fast Skim	.61	.56	.54	.46
Alternating	Normal	.68	.50	.65	.49
Alternating	Skim	.61	.47	.61	.46
Alternating	Fast Skim	.63	.58	.58	.47

[a]FA = False Alarm

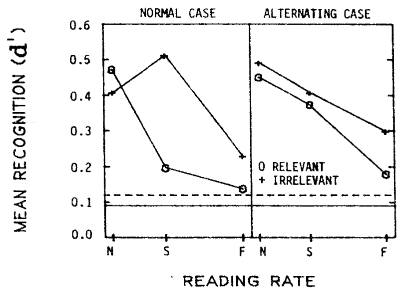

Figure 28. Mean recognition performance on surface test in Experiment 3 (N = Normal; S = Skim; F = Fast Skim). Significance of difference from zero at 99% and 95% confidence levels are indicated by dashed and solid lines, respectively.

test found that while recognition scores in the normal (0.45) and skimming (0.37) conditions did not differ reliably, both were significantly greater than recognition scores in the fast skimming condition (0.21).

The deterioration of surface memory as reading rate increased is consistent with the findings in Experiments 1 and 2 and adds further support to the hypothesis that when subjects skim they are not able to fully process surface structure. The use of alternating case did not contribute significantly to the decline of surface memory, possibly because subjects opted to devote more attention to decoding processes and allow selective processing of relevant information to lapse. The combination of reading rapidly and attempting to maintain selective processing may have prevented subjects in the normal case condition from devoting more attention to complete execution of decoding processes than subjects in the alternating case condition.

Reaction time for hits. The mean reaction times for hits are shown in Figure 29. An analysis of variance failed to reveal any reliable effects, again indicating that regardless of reading conditions subjects were equally fast in correctly recognizing surface structure of statements they had read, even though subjects possessed differing degrees of knowledge about surface structure.

Reaction time for correct rejections. The mean reaction times for correct rejections are presented in Figure 30. The analysis did not find any reliable effects, as reading conditions

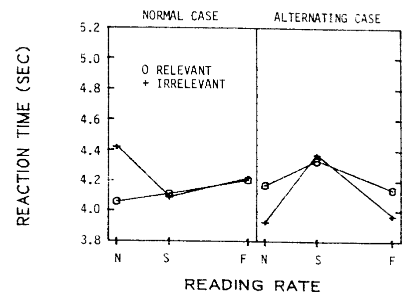

Figure 29. Mean reaction time on hits on surface test in Experiment 3 (N = Normal; S = Skim; F = Fast Skim).

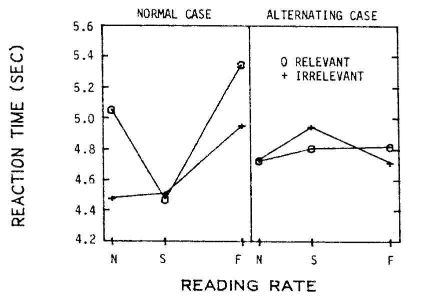

Figure 30. Mean reaction time on correct rejections on surface test in Experiment 3 (N = Normal; S = Skim; F = Fast Skim).

did not affect the time required by subjects to reject paraphrase statements.

Reaction time for false alarms. Mean time taken by subjects to incorrectly accept a paraphrase statement is shown in Figure 31. An analysis of these data showed that there was a reliable interaction of reading rate and case, F (2, 186) = 3.23, MS_e = 3.262, indicating that the reaction time for subjects in the skimming condition who read alternating case was particularly high. There was also a reliable interaction between case and statement type, F (1, 186) = 4.77, MS_e = 1.347. According to subsidiary analyses, subjects in the normal case condition tended to be quicker to accept relevant paraphrase statements as verbatim than irrelevant ones, F (1, 93) = 3.52, MS_e = 1.227, p < .07, while no effect of relevance was apparent for subjects in the alternating case condition.

The slower reaction time for subjects in the alternating case condition who skimmed the stories probably was due to a speed-accuracy trade-off effect as these subjects also had the lowest false alarm rates (see Table X). For those subjects who read normally typed stories the greater familiarity with the meaning of relevant statements evidently led to more ready acceptance of surface structure of paraphrase statements.

Reaction time for misses. The mean time taken by subjects before rejecting a verbatim statement is presented in Figure 32. None of the effects in the analysis of these reaction times reached significance. Subjects were apparently able to

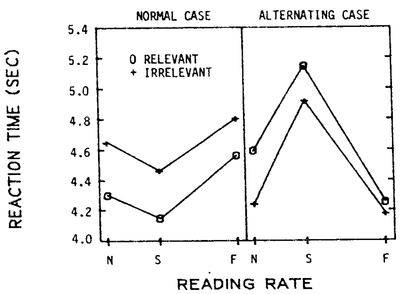

Figure 31. Mean reaction time on false alarms on surface test in Experiment 3 (N = Normal; S = Skim; F = Fast Skim).

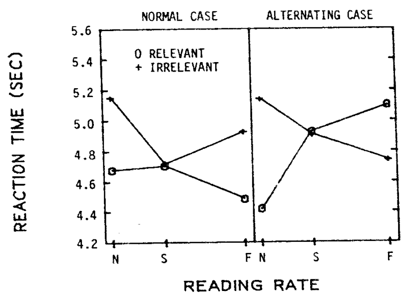

Figure 32. Mean reaction time on misses on surface test in Experiment 3 (N = Normal; S = Skim; F = Fast Skim).

exhaustively search their representations of surface structure in about the same amount of time, though not necessarily at the same rate, regardless of reading conditions.

Summary and Conclusions

Responses to the questionnaire administered to subjects at the end of the experimental session provided information about skimming strategies that subjects claim they use. For the most part subjects used some kind of selection strategy with more elaborate processing of selected information. The fast skimming and alternating case conditions caused some subjects to alter their skimming strategies and among these subjects most switched to some form of selection strategy from a nonselection strategy. The selection strategies described by subjects were similar to the key elements in Maxwell's (1973) method of teaching skimming. This method involved selection of critical words and construction of major ideas from them based on prior knowledge of the text's topic. Questionnaire responses also reliably foreshadowed the effects of reading rate and alternating case on the subjects' success in extracting relevant information from the stories. Self-ratings of effectiveness in skimming were sensitive both to the use of alternating case and fast skimming manipulations. Recognition of meaning performance provided evidence that, in fact, when forced to skim at 600 wpm subjects' memory performance declined below that observed when subjects read at normal or moderate skimming rates. Evidence with a limited degree of

reliability was also provided showing that alternating case interfered with subjects' ability to select and elaborately process relevant information. A post hoc analysis of macrostructure formation, however, showed that subjects were able to form representative macrostructures under any of the reading conditions. Surface memory suffered under speeded reading conditions but not under alternating case conditions. Note that in Experiment 2 when success in skimming was defined by formation of a macrostructure, self-ratings in skimming success were not decreased and eventual formation of macrostructures was not prevented by the use of alternating case.

The effect of statement relevance on recognition memory was not as strong, either in terms of accuracy or reaction time, as was the effect of statement importance in Experiments 1 and 2. Part of the cause of a diminished effect of statement type in Experiment 3 probably is related to the fact that relevant information in the critical stories was embedded in other irrelevant information and not readily discernible. As a result, subjects had to process irrelevant information more carefully before reliable judgments about relevance could be made and relevant information could be more elaborately processed in development of the memory structure. Subjects were able to maintain a reasonable degree of selectivity in their processing as reading rate increased, so long as stories were typed normally. Use of alternating case led to a breakdown of selective processing as soon as subjects began to skim. While the variables of rate and case were manifesting their effects on subjects' ability to selectively process relevant information, subjects were proceeding without setback in their development of macrostructures. The formation of macrostructures seems to be a fundamental part of reading and successfully went on in conjunction with (sometimes unsuccessful) attempts to selectively process information relevant to a perspective.

The fact that subjects in the alternating case condition could not selectively process relevant information when skimming, even though that was their primary goal, while they were able to form a representative macrostructure, suggests a number of theoretical possibilities. One possibility that will be elaborated here is that the requirements of decoding alternating case and maintaining a system of selective processing of information based on criteria imposed by an assigned perspective exceeded subjects' processing resource limits (Norman & Bobrow, 1975) when skimming. Decoding can be done efficiently when reading normal case, possibly through the use of whole word visual identification mechanisms. More controlled processing (Shiffrin & Schneider, 1977) of surface structure must be done, however, when decoding alternating case, causing too great a demand on cognitive processing resources for both decoding and selective processing to occur during skimming. Consequently, subjects elected to focus on decoding and slackened selectivity criteria so that at least some information could be accurately processed. The necessity of accurate surface processing,

particularly in decoding alternating case, when searching for perspective-relevant information is reflected in rather high (above chance) surface memory recognition performance and in the finding that surface memory was not negatively affected by use of alternating case.

A related theoretical explanation of the loss of the relevance effect when subjects skimmed alternating case is based on Navon and Gopher's (1979) recent suggestion that two primary processing tasks may show interference effects if they both rely on a processing mechanism or set of mechanisms that requires processing capacity and allocation of that mechanism's or set of mechanism's capacity to both primary processing tasks is not feasible. With respect to the effects observed in Experiment 3, it may be argued that decoding and selective processing of perspective-relevant information are two processes that are dependent on a set of mechanisms responsible for formation of a macrostructure while reading. These mechanisms require processing capacity and are hypothesized to involve the use of conceptually and data driven processing of a text in an effort to extract information relevant to a basic story schema so that a macrostructure can be formed. Processing resources associated with macrostructure formation can be allocated to such tasks as decoding to aid in efficient, conceptually driven reading, and to the maintenance of selectivity criteria that are responsible for selection of perspective-relevant information. If the combination of decoding and selectivity place excessive demands on the resources of the macrostructure formation process, an austere policy of allocation must be put into effect in which one of the primary tasks of decoding or selectivity must suffer. In the case of Experiment 3 it is apparent that a greater allocation of processing capacity was provided to the decoding process so that subjects would be able to extract at least some information and successfully form a macrostructure.

This alternative theoretical formulation is quite similar to the one based on the notion of a single, common pool of processing resources but is preferred over that account for two reasons. First, it takes into account the elegant arguments developed by Navon and Gopher (1979) in favor of a multiple resource theory. Second, it helps to elucidate the functions of conceptually and data driven processes in decoding and macrostructure formation. It also points up the interdependence of processes, as it probably is no coincidence that macrostructures were formed when decoding of textual information was successful even if a price was paid in terms of lack of selective processing of perspective-relevant information. The pattern of results is consistent with the notion that developing a memory structure that favors perspective-relevant information is dependent on the concurrent formation of a macrostructure. On the other hand, processing mechanisms responsible for development of a macrostructure are not reliant on selectivity of processing of perspective-relevant information. Macrostructure formation

appears to be the process most closely associated with the fundamental, resource-rich mechanisms of reading.

CHAPTER V

EXPERIMENT 4

The experiments reported thus far have all been concerned with recognition tests of subjects' memories for stories they have skimmed. Experiment 4 was an extension of the evaluation of memory representations that involved the examination of recall protocols. The rationale for studying recall data was that in recalling a story subjects are in a position to indicate which aspects of text information they have represented in memory and how those apsects are organized. Clearly, subjects are more likely to recall centrally important information than unimportant details, and while details may be recognized in a verification task, the probability of their being output in a recall test is low. Thus, recognition tests may not reveal the full distinction between important and unimportant information which characterizes macrostructure representations.

Another procedural difference between Experiment 4 and the earlier studies was that reading rate was manipulated as a within-subject variable using the pacing procedure. It was hoped that even more stable comparisons of the effects of different reading rates on comprehension and memory could be obtained if each subject was observed at each reading rate. It was also expected that with the questionnaire items regarding skimming strategies and changes in strategies it would be possible to determine how skimming strategies operate and change

within-individuals as reading rate increases.

While recall protocols can provide information about which aspects of a story subjects are likely to remember, generalizations about recall protocols cannot be concise unless some definition of text structure can be applied. In the first three experiments text structure was characterized by degree of importance of test statements. Greater recognition of important than of unimportant statements was taken as evidence that a macrostructure-like representation had been formed. In Experiment 4 a more strict definition of text structure was imposed using the text processing model developed by Kintsch and van Dijk (1978). Recall protocols were scored using Kintsch's (1974) propositional scoring method to obtain a recall profile across propositions of each story. The observed profiles or patterns of recall were then compared with predictions based on the text comprehension and production model described by Kintsch and van Dijk. The specific, computer implemented version of the model that was used was described by J. Miller and Kintsch (1979), and a summary of that version of the model is provided in the next section. Comparison of obtained data with predictions based on the text processing model will serve not only to provide the recall data with a reasonable structural organization, but will also, importantly, provide an interesting test of the model. Specifically, recall protocols obtained from subjects when reading at a normal rate should be successfully simulated by the model, given that the model's basis rests on normal reading processes. Successful simulation would also indicate that the procedures used to establish reading rates did not seriously alter the usual macro-operations predicted by the model. Moreover, claims about the use of macro-operations during skimming could be tested by evaluating the success of the model in simulating recall protocols given after skimming stories. If the model can predict these recall protocols, strong evidence for the use of macro-operations during skimming will have been obtained, and the processing model itself will have gained claim to a greater degree of generality than originally supposed.

Text Comprehension and Production Model

The version of Kintsch and van Dijk's (1978) text processing model that has been implemented by J. Miller and Kintsch (1979) preserves the fundamental concepts of macrostructure representation, cyclic processing, and limitations of working memory. In general, the programmed implementation of the model works toward a representation of a text that is based on the relative importance of key propositions in the text structure (Kintsch & Keenan, 1973; Kintsch, Kozminsky, Streby, McKoon, & Keenan, 1975; McKoon, 1977) and the coherence relations between propositions in the text (Kintsch & van Dijk, 1978). These criteria are used in selective processing of key propositions for the development of a memory representation.

The critical assumptions of the model are as follows. Chunks of the text (sets of propositions) are processed simultaneously

in a single cycle of the program. Only a limited number of propositions can be included in a chunk, and this number is determined by one free parameter and aspects of text structure. Propositions will be selected consecutively from the proposition list representing the text for inclusion in a processing chunk until either a sentence boundary is reached or all of a number of text-dependent criteria are met. These criteria, generally stated, are: (a) at least two propositions have been selected, (b) the next unselected proposition in the list is not associated with the most recently selected proposition in either an embedding, modifier, or shared argument relationship (Kintsch, 1974), and (c) at least some minimum number, *i*, of words in the text have been included in the chunk. These criteria depend both on aspects of the surface structure of the text itself and on the text's propositionalized representation. Once a chunk has been established, each constituent proposition is then given some probability, *p*, of being processed into long-term memory such that it will be reproduced in a recall or summarization task. In this way, *p* captures a combined probability of successful encoding and production. The constituent propositions are also arranged in a continually growing coherence graph (Kintsch & van Dijk, 1978), that indicates the pattern of argument overlap and reference among propositions. In order to maintain coherence across processing cycles, the next cycle begins with a selection of a subset of propositions from among those present in the previous cycle for inclusion in the current cycle. The coherence graph can then be expanded to include the new chunk of propositions and proper coherence relations will be drawn to previously processed propositions. Two aspects of the selection of propositions to be held over for the next cycle are important. First, only some maximum number, *s*, can be held over due to limitations of working memory. After the first processing cycle only, *s* can be expanded to include one additional proposition to account for additional allocation of resources at the beginning of a story as the reader makes an effort to initiate some kind of organizational structure for the story's content. The value of *s* may also be expanded by one in exceptional cases after later processing cycles as noted below. The second critical aspect of the selection process is the heuristic for deciding which propositions to include. The heuristic found to be most successful in simulations (Kintsch & van Dijk, 1978) is the leading edge strategy (Kintsch & Vipond, 1979). This strategy makes use of the coherence graph as it develops after each processing cycle and consists of selection of the most recently occuring propositions at each level of the graph. Selection begins at the highest level and ends when *s* propositions have been selected. If the *s*th proposition is involved in an embedding relationship with a proposition not yet selected, then *s* will temporarily be expanded by one to include that proposition. Furthermore, if the *s*-1th selected proposition is involved in an embedding relation with a proposition at the next level and the latter proposition is not the most recent at its level, it will

be selected as the *s*th entry into working memory rather than the most recent proposition and no expansion of *s* will occur.

Inclusion in multiple processing cycles through the selection process is what determines the relative probabilities of encoding and production of the different propositions. Probability, *p*, of encoding and production is applied to a proposition each time it is included in a processing cycle. Thus, the more important a proposition's role in the coherence graph, the more often it will be included in processing cycles and the higher its likelihood of encoding and subsequent production. The macrostructure representation of a story, then, is defined by the nature of the coherence graph and the parameters *s* and *i*, which control the number of processing cycles and the frequency with which propositions are carried over in multiple processing cycles.

Of interest in Experiment 4 was whether or not the text processing model would successfully predict recall protocols in skimming conditions and, if so, what parameter changes would be necessary in obtaining reasonable predictions. Changes in *s* and *i* would indicate alterations in the amount of information held in working memory and the amount of information processed in each cycle. The nature of the macrostructure would also be altered. The quality of comprehension, especially with respect to later recallability, would be captured by the parameter *p*.

Method

Subjects

The subjects were 72 students recruited from the same population as that used in the previous experiments.

Materials

The 500-word narrative used in the preliminary reading task in the first three experiments also was used in Experiment 4. Seven other narratives of about 500 words each were developed: three test and four practice stories. The three test stories contained a particular critical paragraph located near the middle of the story. These critical paragraphs were among those used by J. Miller and Kintsch (1979) in their simulation of recall protocols and, therefore, an external criterion of comparison was available. The critical paragraphs, the number of propositions in each, and the title of the corresponding test stories are presented in Appendix C.

The stories were typed single spaced, right and left justified on separate pages. An appropriate title was typed centered at the top of each story page. Three story orders were arranged representing a Latin square counterbalancing of order of the test stories. The practice stories always appeared in the same location, one at the beginning and another before each test story. A blank colored page was inserted before each story and a page that contained a summary of recall instructions was appended to the front to complete the construction of each story booklet. Three booklets of each story order were made.

Recall booklets were composed of a cue and a recall page for the first practice story and for each test story. The cue page gave a brief description of the critical paragraph that the subject was to recall from the story that had just been read, and the recall page contained blank lines. These pages were inserted into a folder so that when the subject opened the folder to the materials for a specific story, the cue page was on the left and the recall page was on the right. The cue used for each critical paragraph is shown in Appendix C. Three sets of recall folders, each corresponding to one of the counterbalanced orderings of test stories, were prepared.

Three sets of pacing tapes were constructed with the same rates (225, 375, and 600 wpm) and general design as that used in Experiment 2. Each set represented a Latin square counterbalanced ordering of the three reading rates. The first paced rate was always 225 wpm and the other six were ordered in pairs such that each pair corresponded to one of the three rates and followed the counterbalanced ordering of rates.

Design

The counterbalancing of story order and order of reading rate were combined to form nine combinations, so that each test story was read at each position in the story order at each reading rate. Eight subjects were assigned to each of the nine combinations according to order of appearance at the laboratory.

Procedure

Subjects participated in groups ranging in size from one to six. An experimental session began with the preliminary reading task used in the earlier experiments. Afterward, subjects were given the story and recall booklets and they read the recall instructions which stated that they would be required to recall, as completely and accurately as possible, one paragraph cued from each test story. The pacing procedure was explained as in Experiments 2 and 3, then subjects read the first practice story at the paced rate of 225 wpm and attempted to recall the appropriate paragraph. A period of 5 min was allowed for recall, after which the experimenter reviewed the procedure to be sure that subjects understood the requirements of the reading and recall tasks. The remaining stories were read and tested in three cycles. Each cycle began with an explanation of the reading rate that was in effect, followed by the paced reading of two stories in a row, first the practice then the test story. The practice story was meant to allow subjects to adapt to the rate in effect. After reading the test story, subjects turned to the next blank page in their story booklet, then opened their recall folders to the next section and were allowed up to 5 min to recall what they could of the cued paragraph that was present in the test story they had just read. Time allowed for recall was ample in all cases. After completion of the reading and recall cycles subjects completed the form of the reading questionnaire used in Experiment 3. Items relevant to rate adaptation were to be answered with respect to each of the three reading rates used in the experiment.

Results

Questionnaire

Subjects in Experiment 4 seemed to have skimming habits similar to those described by subjects in the earlier studies. The mean estimate of percentage of total reading time spent skimming was 29%, with a range of 0.05-75%. Newspapers and magazines were the most commonly mentioned topics for skimming, as 44% of the subjects indicated that they skim one or both types of material. Aspects of college work were mentioned by 22% of the subjects as topics for skimming and 15% of the subjects stated that they had received some form of speeded reading training. The mean estimate of subjects' confidence in their success in selecting important information while skimming outside the laboratory was 1.72, based on the four-point scale of zero to three used in the earlier experiments. Subjects who claimed to have had training in speeded reading expressed higher confidence in their skimming abilities, with a mean rating of 2.15. With respect to mean estimates of skimming success in the experiment, at the skimming rate the mean rating was 1.69, very close to the value for independent skimming, and the mean rating for fast skimming was 1.57. The rather high rating for fast skimming may have been due to failure on the part of some subjects to distinguish between the two skimming rates in providing their responses and the subsequent use of the same rating value in calculation of the means for both skimming and fast skimming. Another indication that subjects adapted well to the reading rates used in Experiment 4 was found in comparison of experimenter-paced and self-paced rates. For the normal reading rate, 51% of the subjects claimed the paced rate matched their own normal rates, while only 21% claimed the paced rate was too fast. The paced skimming rate was also representative of self-paced rates as 53% of the subjects felt that the paced rate was close to their own skimming rates, and 43% stated that it was too fast. For the fast skimming rate, however, only 18% of the subjects claimed that the paced rate was the same as their own fast skimming rates, while 82% felt the paced rate was too fast.

The description of skimming strategies followed the pattern set by subjects in Experiment 3 rather closely. A selective skimming strategy of one sort or another was described by 79% of the subjects, 13% claimed they use the procedure of glancing over all words or reading everything faster than usual, 6% said they try to perceive chunks of a story simultaneously, and 3% stated that they skip various parts of a text indiscriminantly or without regard to semantic content. Of those subjects who ordinarily use a selection strategy, 54% base selection on key words, 24% base selection on specific sentences or paragraphs in certain locations in the text, and 22% mix these two strategies. Fundamental deviations from these described strategies due to the pacing procedure were mentioned by only five subjects. Two subjects changed from a selection strategy to a nonselection strategy, two switched in the opposite direction, and one altered the form of the selection strategy ordinarily in effect.

Simulation of Recall Protocols

Recall protocols were scored using Kintsch's (1974) system of propositional analysis and scoring. The mean proportion of propositions recalled from each critical paragraph at each reading rate is shown in Table XI. It is clear that while the paragraphs differ in general level of recall, all three exhibited decline in recall peformance as reading rate increased.

A more detailed analysis of the characteristics of recall protocols was provided by attempts to simulate these data with the text processing model described above. Possible values of *i* were constrained by text structure and maximum sentence length. On some occasions a range of *i* values produced the same pattern of cyclic processing, so the range of values was combined and only the minimum value will be reported. Values of *s* ranging from one to five were tested. For each combination of *i* and *s* a hill climbing parameter estimation routine was used to find the value of *p* that minimized the chi-square goodness-of-fit comparison of observed and predicted frequency of recall of each proposition.

In calculating chi-square goodness-of-fit values a convention used by J. Miller and Kintsch (1979) was adopted. It was assumed that a few propositions will not be output during recall because of their low degree of importance in the text. Even if subjects remembered such propositions they probably would not write them down in a recall test. These propositions generally are modifiers and an example from the *Panda* paragraph is the fact that panda bears are excellent climbers. While

TABLE XI

MEAN PROPORTION OF PROPOSITIONS RECALLED IN EXPERIMENT 4

Paragraph	Reading Rate		
	Normal	Skim	Fast Skim
Hitchcock	.450	.398	.266
Panda	.433	.325	.168
Roses	.533	.425	.204

subjects are likely to write down the fact that pandas can climb, they may well omit the fact that they are excellent climbers. Propositions such as these were not included in the calculation of chi-square values. An objective criterion for excluding or pruning propositions involved two steps. First, any proposition with a recall frequency across all subjects that was more than one standard deviation below the mean frequency was considered for pruning. Usually propositions recalled by fewer than three or four subjects of the 24 who read the test story at the rate in question met this part of the criterion. Next, propositions that were redundant with respect to information already in the developing coherence graph were pruned while other propositions were retained under the assumption that they were poorly recalled for other, unaccountable reasons. Pruned propositions usually were modifiers or details related to such things as the setting of the story.

The best fitting parameter estimates and corresponding chi-square values for each story and each reading rate are shown in Table XII. Two bases for pruning were used. The rate-specific basis allowed pruning to be based on the recall data specific to the reading rate being simulated. The general result was a greater number of pruned propositions at the skimming and fast skimming rates. Propositions in these two reading rate conditions were pruned just as long as they met the criterion of one standard deviation below mean recall frequency, as it was felt that skimming processes may be responsible for as type of

TABLE XII

PARAMETER ESTIMATES AND CHI-SQUARE GOODNESS-OF-FIT VALUES

Paragraph	Parameter, Pruning Index, or Chi-Square Value	Basis of Pruning				
		Rate-Specific			Normal Rate	
		Reading Rate[a]				
		N	S	F	S	F
Hitchcock	i	6	6	6	6	6
	s	3	3	3	3	3
	p	.346	.326	.206	.302	.216
	PP[b]	2	5	3	2	2
	χ^2	28.60	15.24	44.33[c]	32.60	46.23[c]
	df	21	18	20	21	21
Panda	i	14	13	13	13	13
	s	2	1	1	1	1
	p	.413	.339	.162	.335	.176
	PP	5	6	3	5	5
	χ^2	38.50	16.74	28.86	18.67	25.34
	df	23	22	25	23	23
Roses	i	6	16	6	16	6
	s	5	4	5	4	5
	p	.401	.441	.142	.419	.131
	PP	4	8	8	4	4
	χ^2	30.33	40.40[c]	20.69	56.30[c]	36.61
	df	24	20	20	24	24

[a] N = Normal; S = Skim; F = Fast Skim

[b] PP = Number of propositions pruned

[c] $p < .01$

selectivity not found at normal reading rates. Nevertheless, pruned propositions generally were modifiers or similar details and not centrally important propositions. The normal rate basis for pruning simply involved allowing exactly those propositions pruned in the normal reading rate condition to be pruned in either skimming or fast skimming conditions.

Once the pruning criteria had been applied, chi-square goodness-of-fit values were calculated on the basis of the observed and predicted frequency of recall of each proposition in a paragraph. Low chi-square values indicated that the pattern of recall frequencies across propositions was replicated by the model with a reasonable degree of accuracy. In such cases, the recall data and model agree on which propositions are most central to the paragraph and, hence, are more likely to be recalled.

With respect to the model's success in fitting the recall protocols in the normal reading rate condition, all three paragraphs were successfully simulated. The chi-square values were quite acceptable and similar to those obtained by J. Miller and Kintsch (1979). They obtained the same optimal values of i and s for the *Hitchcock* and *Panda* paragraphs, while for *Roses* they found that $i = 13$ and $s = 1$ were optimal. Their optimal values of p were generally higher than those obtained here as their subjects read the critical paragraphs in isolation rather than embedded in longer stories, and so recall was much higher.

The chi-square values obtained in simulation of the skimming and fast skimming conditions were very good, particularly when rate-specific pruning was used. The model had some difficulty with *Hitchcock* at the fast skimming rate and with *Roses* at the skimming rate, but otherwise the model fared very well, with chi-square values as good as or better than those obtained in the normal reading rate condition. The parameter changes necessary to obtain the successful simulation of protocols in the skimming and fast skimming conditions can easily be summarized. The values of i and s changed very little, except in the relatively poor fit of *Roses* in the skimming condition. The most noteworthy parameter change was in p, which generally decreased as reading rate increased.

An example of the model's success in predicting the recallability of selected important and unimportant propositions as reading rate increases is presented in Table XIII. Obtained and predicted data are in general agreement as to which propositions will maintain a relatively high probability of recall even after skimming. A further example of the model's predictions is provided in Appendix C. For each reading rate condition, propositions from the *Panda* paragraph that were predicted to be recalled by at least 10 subjects were written out in English to provide protocols that simulated recall of facts actually contained in the text.

TABLE XIII

OBSERVED AND PREDICTED PROBABILITIES OF RECALL OF SELECTED IMPORTANT AND UNIMPORTANT PROPOSITIONS FROM "PANDA" PARAGRAPH

Sample Propositions and Reference Numbers		Reading Rate					
		Normal		Skim		Fast Skim	
No.	Proposition	O[a]	P[bc]	O	P	O	P
2	(THINK SCIENTISTS 3)	.79	.80	.75	.71	.46	.41
3	(RELATIVE-OF PANDA BEAR)	.92	.66	.83	.56	.54	.30
9	(CLASSIFY SCIENTISTS 10)	.67	.66	.62	.56	.33	.30
10	(RELATIVE-OF PANDA RACCOON)	.92	.96	.83	.92	.54	.65
24	(ISA PANDA CLIMBER)	.42	.66	.25	.34	.29	.16
8	(TIME: TODAY 9)	.46	.41	.38	.34	.08	.16
20	(WALK PANDA SOLE)	.62	.41	.33	.34	.08	.16
21	(OF SOLE FEET)	.25	.41	.25	.34	.04	.16
22	(LIKE 20 23)	.50	.41	.29	.34	.08	.16
23	(WALK BEAR)	.50	.41	.29	.34	.08	.16

[a]O = Observed

[b]P = Predicted

[c]Predictions based on rate-specific pruning

Discussion

The questionnaire data indicated that subjects in Experiment 4 were similar to subjects in the other experiments, both in terms of their usual skimming habits and strategies and in terms of their ability to adapt to the pacing procedures used in the experiment. The reading behavior of subjects in the experiment appeared to be representative of their usual reading processes. This fact is important in evaluating the validity of the recall protocols obtained in Experiment 4. Given that the protocols are valid, success in simulating them provides valuable evidence in favor of the underlying simulation model.

In general, the text processing model developed by J. Miller and Kintsch (1979) was highly successful in simulating recall protocols taken after normal reading or skimming of stories. Embedding of the critical paragraphs in longer texts did not seriously affect performance of the model. This fact is somewhat surprising given that the model's predictions are based on structural coherence of propositions rather than the relationship of currently processed propositions with semantic content occurring much earlier in the text. The similarity of optimal parameter values found here and by J. Miller and Kintsch (where paragraphs were read in isolation) points up the relative importance of the role of local coherence relations in predicting recall.

Moreover, the success of the model in predicting recall protocols in the skimming and fast skimming conditions was

striking. The fundamental assumptions of the model pertaining to cyclic processing, holding over selected propositions for the next cycle, and limitations of working memory in the process of holding over propositions were shown to be valid possibilities not only for normal reading but also for skimming. In fact, the basic aspects of these processes, captured by the parameters *i* and *s*, remained virtually the same across reading rates. As mentioned above, these are the parameters that control the nature of the macrostructure and their constancy over reading rates is very revealing of the similarity of the macrostructure to be expected when reading normally and when skimming. Only the parameter *p* showed a large change and the number of rate-specific pruned propositions increased with reading rate. The lower values of *p* could be interpreted as poorer comprehension, possibly due to incomplete processing of surface structure, or less elaborate or extensive encoding into long-term memory due to time constraints. The increase in the number of detail propositions that were pruned probably reflects the tendency for subjects to skip over certain details while skimming.

A description of skimming processes from the point of view of the text processing model might be as follows. A number of propositions are selected for processing as a group. The number selected is about the same or possibly less than that for normal reading, but the extent or quality of processing of each is diminished (reduced value of *p*). Depending on current coherence relations a number of propositions are selected to be held over for the next processing cycle, and this number and selection process are approximately the same as for normal reading.

There are two aspects of this description of skimming that may not be consistent with reality. First, the notion that each proposition in a chunk is processed, but to a lesser degree than is true when reading normally, may not be correct. It is possible that of the propositions engulfed in a chunk only a few are selected (or even perceived), on whatever basis, for further processing. The nature of the propositions that were pruned suggest that this hypothetical selection process would be sensitive to the general importance of propositions. The model might be more accurate in its predictions if some sort of selective processing parameter or criterion were included at this stage. Second, the selection of propositions for inclusion in the next processing cycle was based on the leading edge strategy. It is possible that some other heuristic might provide an even more accurate account of skimming processes. These possibilities will be elaborated in the concluding chapter.

Whether the issues of selective processing within a chunk or of selection heuristics are handled in exactly the right manner by the version of the text processing model described here is not certain. But it is very clear that even if changes in these two processes are warranted they will result in a version of the model that is not very different from the current one. The success in predicting recall obtained with the current version was quite remarkable and serves to emphasize the importance of

the processes detailed in the model. These processes are very general with repsect to a broad range of reading rates and this generality suggests the existence of fundamental commonalities of reading for full comprehension and skimming.

A further suggestion of common dependence on macro-operations at different reading rates is provided by the degree of consistency between Experiments 1-3 which involved recognition tests and the recall data of Experiment 4. In all cases strong evidence for the formation of macrostructures was obtained, regardless of the type of memory test that was used. Experiments 1-3 indicated that macrostructures were formed even during skimming. Although these macrostructures were not as complete or accurate as macrostructures formed during normal reading, they still reflected a strong bias for inclusion of important over unimportant information. In particular, the relative proportions of important and unimportant information contained in macrostructures was maintained over different reading rates. No fundamental change in selective processing of macrostructure-relevant information was observed when resource demands were reasonable. In Experiment 4 the pattern of parameter changes required to simulate recall data obtained in skimming conditions relative to the normal reading rate condition, also failed to reflect any fundamental change in the nature of macrostructure formation. Only the encoding parameter, $\underline{p}$, was altered, suggesting that the major effect of skimming was to reduce the likelihood of encoding any particular proposition.

within the ranges of p found in Experiment 4, the overall decrease in frequency of recall of propositions that participated in one processing cycle was not very different from the decrease observed for propositions that participated in multiple cycles. If p is reduced drastically (as should occur with extremely high reading rates), however, the largest decrease in recall probability would be observed in the important propositions, and eventually the distinction between important and unimportant propositions would be very small. Thus, if the model's predictions are correct, pushing subjects to extreme reading rates would result in a large decrease in the distinction between important and unimportant information in the memory representation, rather than an increase in the distinction. The breakdown apparently would be at the encoding level, while selective processing still would occur with fundamental macro-operations in effect. This extreme case would require a very small value of p, as even values smaller than .2 led to a macrostructure that exhibited distinction of important information. If speed reading is successful in developing accurate, representative macrostructures, then it may be that the fundamental processes described by the model are not appropriate for speed reading. The degree of distinction between important and unimportant information would have to be determined before the model's validity for speed reading could be assessed.

Taken as a whole, the four experiments provide converging evidence for the conclusion that fundamental macro-operations are

in effect at a range of reading rates when resource demands are not increased by use of a typography that is difficult to decode. In skimming, resources are not sufficient for full processing of all propositions and, consequently, there is a general decrease in recall and recognition performance. The resource requirements for selecting macropropositions for more extensive processing (e.g., inclusion in multiple processing cycles) are consistently met and a basic, though incomplete, macrostructure is formed.

CHAPTER VI

CONCLUSION

The first section of this chapter summarizes the major findings of the four experiments reported here, and highlights the significance of some of these findings for hypotheses about cognitive processes in skimming stories. A consideration of skimming as a form of reading and implications for other rapid reading methods will be discussed in the next section. A third section will be devoted to a discussion of the implications of the research reported here for the Kintsch and van Dijk (1978) text processing model and for DeJong's (1979) newspaper skimming program, FRUMP. The final section will provide a plan for further investigations of skimming and other forms of rapid reading, including recommendations for a variety of experimental paradigms and a discussion of related implications.

Summary of Experiments

A number of interesting facts concerning skimming were obtained from the 516 subjects whose recall or recognition data were analyzed in the four experiments discussed in earlier chapters. For these subjects, an average of nearly one third of their reading time is spent skimming. Although the most common material skimmed by these subjects is newspaper and magazine articles, about 17% of them are confident enough in their skimming abilities to apply them to various aspects of their college work. For the most part, skimming skills are developed independently, and probably are derived from normal reading processes, since skimming is a very common practice and only 16% of the subjects claimed to have received any formal training in speeded reading. It would not seem surprising, then, to find strong relationships between normal reading and skimming processes. In fact, Experiments 1-4 all found evidence for a good deal of overlap in processes involved in reading for full comprehension and skimming.

In Experiment 1 subjects were allowed to read narratives and newspaper stories at their own normal reading and skimming rates. A subsequent recognition test showed that, while performance generally declined when subjects skimmed, greater recognition of important than of unimportant statements was maintained. This effect indicated that when skimming, subjects form macrostructure representations similar to, but less complete than, those formed when reading normally. Story schemata both for narratives and newspaper stories were effectively used and were adhered to at various reading speeds. Increased reliance on conceptually driven processes was evidenced by the maintenance of at least the basic macro-operations while skimming and by decreased surface memory, which suggested that subjects did not fully process surface features but, rather, relied on conceptually driven processes to provide accurate hypotheses about what actually was written. In this way less surface processing was required and skimming could proceed rapidly.

In Experiment 2 a replication of the first study was provided with an extension to more extreme reading rates and the use of a method of pacing reading rates. It was also found that reading alternating case typography while skimming apparently interferes with whole word visual identification processes which play a key role in the rapid decoding of words. Moreover, different effects of alternating case on the macrostructure representations of narratives and newspaper stories pointed to differences in basic story schemata and allocation of processing resources. Subjects were able to make use of the story structure of newspaper stories and could glean the important information from the early part of those stories even when skimming alternating case. Then subjects could process details later in the stories. In skimming narratives typed in alternating case, however, it was necessary for subjects to search for important information throughout the stories with little attention allocated to processing of details. The ability to process surface structure features was generally reduced by the use of alternating case, forcing the judicious use of processing resources outlined here.

An even more revealing demonstration of the role of processing resource allocation in skimming was provided in Experiment 3. When assigned a goal in reading normally or skimming subjects were successful in extracting more goal-relevant than irrelevant information from stories. If the reading task also required subjects to read alternating case,

though, selective processing of goal-relevant information was possible only when reading at the normal rate. The requirements of skimming alternating case and maintaining a selective processing criterion for a particular goal were not compatible. As rate increased from 225 wpm to 375 wpm, differential recognition of goal-relevant and irrelevant statements was lost. Subjects were able, however, to maintain some basic macro-operations as adequate macrostructures of the stories, regardless of the goal in reading, were formed under all reading conditions. These results suggest that decoding and selective processing of goal-relevant information may both rely on processing resources associated with the mechanisms (cf. Navon & Gopher, 1979) responsible for forming a macrostructure representation.

In the final experiment, a version of Kintsch and van Dijk's (1978) text processing model was used to successfully simulate protocols provided by subjects who had read stories at various rates. The best fitting set of parameter values indicated that while processes involved in determining the nature of the macrostructure representation were similar across reading rates, the quality of processing of each proposition was drastically reduced as reading rate increased. The fact that particularly important propositions were retained in working memory and participated in multiple processing cycles served to enhance the likelihood, as far as possible under the circumstances, of their inclusion in the memory representation of a story. The success of the model in predicting protocols based on such a range of reading rates with relatively few parameter changes stongly suggests that there are fundamental macro-operations that underlie reading at a wide range of rates.

Skimming as a Form of Reading

The similarity of memory representations and inferred comprehension processes observed in Experiments 1-4 for normal reading rate and skimming conditions strongly argues against Carver's (1977) position that processes such as speed reading and skimming should not be considered reading, and that reading involves the encoding of information at a relatively uniform rate of about 300 wpm. While the classification of "what is not what" in reading could be done in absence of data on reading processes, arguments for inclusion or noninclusion of possible candidates could not properly be evaluated. In the face of compelling empirical evidence, however, the question of defining reading takes on a new aspect involving delineation of processes and evaluation of the degree of overlap in those processes associated with different forms of reading. The question of whether or not something is "reading" then becomes less important. The experiments reported here provided evidence for a high degree of overlap of processes of reading for full comprehension and skimming.

As indicated by the questionnaires a number of different skimming strategies are used by subjects. In spite of these

individual differences, a general description of the basic processes underlying most of the strategies is still possible. It appears that as a reader skims a story he or she probably misses portions of the text, but through selection of key words or phrases or by chance will locate at least some of the potentially important information in a text. Once a statement has been selected as potentially important, the reader can process the statement in more detail, making more use of data driven processes. If further processing indicates that the statement is, indeed, important for the macrostructure it will be processed further or more elaborately (Reder, 1979) and incorporated into the memory structure. Reading or searching for important information are processes that are aided by knowledge structures such as story schemata and by emphasis on conceptually driven processes that can operate on minimal information from data driven processes. In this situation the reader must draw on a good deal of knowledge about the background of a text in order to establish reliable expectations about what is being read so that data driven processing can be curtailed, and in order to draw accurate inferences integrating piecemeal aspects of a story into a reasonably comprehensible whole.

The selective reading and processing of statements is based on some goal that the reader has in skimming. We have seen that a fundamental goal in reading, to which other goals appear to be closely tied, is the formation of a macrostructure representation. Even when the reader's primary goal in skimming fails, as in Experiment 3, a basic macrostructure still is formed. It is as though the mechanisms involved in forming a macrostructure are responsible for a large amount of the processing resources that can be allocated to subskills such as decoding and selection of goal-relevant information (cf. Navon & Gopher, 1979). If the combined sources of processing resources are insufficient for decoding and selective processing of goal-relevant information, the available resources must be allocated so that a reasonable level of reading performance can be achieved. In Experiment 3 it was observed that subjects chose to maintain a minimum standard for decoding processes but allowed selectivity of processing to lapse. Nevertheless, the subjects processed information with respect to a general macrostructure, suggesting that this may be a fundamental goal underlying any reading task.

The selective reading and processing of information discussed in the foregoing description of skimming implies a basic conclusion about studying reading processes. On almost all occasions researchers use the term reading rate as a unitary measure of how a text was read. This rate probably is not based on a constant reading rate, but rather a fluctuating rate of information processing that can vary widely within a single reading of a text (cf. Rankin, 1971). As a result, classification of reading speeds in words per minute can be very misleading unless accompanied by an explication of goals in reading and possible variation in rate across a text. Readers can adjust to

different requirements in reading, exhibiting a high degree of flexibility. Depending upon the type of information the reader wants from a text, the reader can read very slowly or skim and even scan at incredible "rates" for a specific piece of information. Thus, very rapid reading, even speed reading or scanning, can be "effective" if the appropriately defined degree of comprehension is attained. The truly interesting questions, then, deal with the kind of comprehension associated with different reading methods and how information relevant to goals of comprehension are selectively processed. To what extent is a memory representation that is formed after, say, speed reading a story based on information contained in the story versus information already in the reader's knowledge base? Given the remarkable speeds at which some speed readers process stories, it is apparent that they must contribute a good deal of their own knowledge of the world and of the specific topic in particular to the formation of a coherent memory representation. Reading at thousands of words per minute renders full surface processing physiologically impossible (Spache, 1962; Taylor, 1965). After all, "no amount of practice can make us perceive what we do not look at" (Gibson & Levin, 1975, p. 549). The speed reader must be able to make highly efficient use of the small amount of accurate surface processing that is done, calling into play conceptually driven processes to close many of the gaps between bits of comprehended information. The key to speed reading seems to be taking in as much information as possible through a pattern of eye movements similar to those used in skimming (Taylor, 1962), and being able to devote sufficient processing resources to fleshing out otherwise incoherent pieces of data. The process of coalescing a subset of a story's information into a macrostructure is critical. It is clear that speed readers cannot possibly fully process all the information contained in a story. Only systematic investigation of memory for text information and degree of background knowledge can suggest the extent to which a speed reader's memory representation of a story is based on fact or fabrication.

Implications for Models of Text Comprehension and Skimming

In Experiment 4 it was found that Kintsch and van Dijk's (1978) text processing model provided a very good account of skimming processes. It was suggested, however, that the model might be made more specifically appropriate for simulation of skimming processes if two possible adjustments were examined. First, the parameter p, which represents the probability of encoding into long-term memory and later recalling a proposition, may not sufficiently represent the selective processing characteristics of skimming. While the cyclic processing of a text in conjunction with a sophistocated strategy of selecting propositions for inclusion in multiple processing cycles is representative of macrostructure formation, there is no provision for selective processing of specific goal-relevant information. For example, in Experiment 3, subjects reading normally typed

stories were capable of selectively processing perspective-relevant information while skimming and at the same time formed appropriate macrostructures. This result indicates that the macrostructure formation processes represented in the model are well-motivated, but that the use of a single encoding parameter is not sufficient. In order to account for skimming tasks that involve specific goals a selection criterion probably should be established in place of or in conjunction with the encoding parameter. A system such as this would allow the basic macrostructure to be formed as usual, while propositions relevant to the goal in skimming would be especially well represented in memory.

A possible goal-relevant selection process that could be developed would involve one additional parameter, q. The parameter q would behave just like p, except that one of these two parameters, say, p, would be applied to goal-relevant propositions while the other, q, would be applied to irrelevant ones. A priori decisions about the relevance of each proposition for the goal would need to be made, but such decisions should ordinarily be intuitively apparent. Under any reasonable circumstances, of course, p would be greater than q, representing selective processing of propositions relevant to the goal. Selection of propositions for inclusion in multiple processing cycles would go on as usual, leading to the formation of a macrostructure, just as shown in Experiment 3. The application of p or q on multiple processing cylces of a proposition represents, at least minimally, the interaction of goal-relevant and macrostructure-relevant selective processing criteria. Even in cases where selective processing is not feasible (i.e., $p = g$) for such reasons as resource allocation priorities, macrostructure formation still would occur but without influence from the specific goal. With the addition of the parameter g the model might do quite well in predicting the results of an experiment such as Experiment 3 in which recall protocols would be taken after skimming normal or alternating case.

The second possible modification of the text processing model pertains to use of the leading edge strategy. It may be that in skimming, selection of propositions for inclusion in later processing cycles may not be based on the recency plus levels criteria that comprise the leading edge strategy. A number of possible alternative strategies are highly unlikely on the basis of their failure on occasions when the leading edge strategy was quite successful. If selection is based on levels plus primacy, recency only, or upon a random selection strategy, the model does poorly (Kintsch & van Dijk, 1978). The joint influence of levels and recency combine to form a selective strategy that is relevant to macrostructure formation and makes optimal use of coherence of a text as more recent propositions are most likely to be related to currently processed propositions. Whatever alternative selection strategy is developed, these criteria probably will need to be included, implying that any alternative would be rather complex. The

leading edge strategy accounts well for data based on skimming for gist. If it could also be used successfully in conjunction with selective encoding parameters in simulating data based on skimming for goal-related information, its validity for skimming processes would be greatly enhanced. On the other hand, if the inclusion of selective encoding parameters is not sufficient, then a selection strategy sensitive to goal-relevant propositions might be required. Such a selection strategy would further emphasize the importance of the interaction between goal-relevant processing and macrostructure formation.

With respect to DeJong's (1979) newspaper skimming program, FRUMP, the results of the experiments reported here have some general implications. As an artificial intelligence program, FRUMP is not meant specifically as a model of human processing, but still it can be evaluated with regard to its consistency with those processes that we believe contribute to human skimming operations. The general predictive nature of FRUMP's processing is consistent with the conceptually driven processing hypothesized to underlie skimming processes in humans. In FRUMP, predictive processing is based on knowledge structures called <u>sketchy scripts</u>. These scripts determine which aspects of a story are important. Within the sketchy scripts is the key to macrostructure formation or goal-relevant selective processing. If a certain kind of information is sought, the relevant sketchy scripts must be modified to include requests for that information. If such requests are not included in the relevant sketchy scripts, the goal-relevant information will not be extracted for inclusion in the program's output, nor is it even likely to be processed at preliminary stages, because of FRUMP's highly conceptually driven nature. Sections of an input text are parsed only when a specific piece of information is requested by a sketchy script and when no script has yet been instantiated. Information satisfying requests can be obtained by direct text processing of relevant conceptualizations (cf. Schank & Abelson, 1977) which add pieces of conceptualizations to memory in order to satisfy requests, or by inferences that may add partial or complete conceptualizations. The use of text processing is more costly, as it would be in humans, in terms of efficiency. Faster moving versions of FRUMP rely less on text processing and more on inferences, making for a less accurate representation but a more rapid processor. The consistency of this design with the hypothesized human processes of text processing and inference are remarkable.

The processing based on sketchy scripts depends on selection of the appropriate sketchy script. Selection mechanisms were briefly mentioned in Chapter I, but will be repeated here. The text may contain explicit reference to a script, activate an inference implying that the defining action of a script has occurred, or contain explicit reference to an action that is part of a script's set of requests. In the first two cases scripts can be activated directly, but in the third instance there must be a way to conclude which script is referenced by the action. FRUMP

uses a discrimination tree which is given the conceptualization containing the action. The conceptualization is filtered through the tree until the appropriate script is found. These script selection processes are dependent on a mixture of text information and prior knowledge, as would be true of readers who are attempting to initially determine the topic of a text. FRUMP confines these processes to the first paragraph of a story, however, which means that it is dependent on a newspaper story structure. This is very much like the hypothesized differences in processing of narrative and newspaper stories described in Chapter III. In reading narratives, a successful program would have to be ready to activate scripts throughout the story. This capability may prove as costly to machine resources as it did to human resources in Experiment 2.

Finally, FRUMP's objective is to provide a story summary based on the important information represented as sketchy script requests. Only this information will be fully processed, represented, and output by FRUMP. As a result, FRUMP cannot be considered a complete model of human processing and comprehension as subjects apparently are capable of comprehending and remembering irrelevant or unimportant information while skimming. To the extent that comprehension of unimportant information is part of the selection process, models of skimming should include at least limited processing and representation in memory of unimportant information. In FRUMP, the processing is so conceptually driven that this does not appear feasible, as only requested information receives adequate processing for inclusion in memory. In this respect, FRUMP may be an ideal conceptually driven processor, reaching beyond the constraints of human processing limitations. The final output of FRUMP is a summary of important information and if subjects were in agreement with or instructed to search for the important information as defined by sketchy scripts, their summaries of stories most likely would be highly similar to FRUMP's. If subjects were asked to recall a story, though, recall protocols likely would contain some detail information as evidence of at least superficial processing of unimportant information. On the other hand, if given instructions to process stories for highly specific information at high rates of speed, such as number of people killed or an earthquake's rating on the Richter scale, subjects may be able to process stories more selectively in a manner quite similar to FRUMP with little or no processing of irrelevant information. This kind of processing approaches what might be considered scanning for specific information, rather than skimming for the gist of a story.

Given these considerations, FRUMP might be interpreted as a scanning program with powerful capabilities regarding determination of story topics. From a more general perspective, FRUMP should be regarded as a skimming program with strictly pre-defined macrostructures for different topics and powerful conceptually driven mechanisms for comprehension. The existence of these possible interpretations attests to FRUMP's potential

generality as an artificial intelligence program and as a model of human information processing.

Future Considerations

In this concluding section I would like to propose a general plan for further research on rapid reading processes. The types of rapid reading that, at this time, appear to be most worthy of study are skimming, speed reading, and scanning. The latter process is broadly defined here to mean rapid searching through a text for some small number of specific pieces of information that might easily be detected due to uniqueness of visual patterns, such as numbers. In scanning, as opposed to skimming, the objective is not to form a macrostructure but to locate and read a particular set of information that represents a very small part of a text. Even in this case the basic macrostructure formation operations may be in effect. Or the reader may not process information that does not match some predefined visual pattern. In speed reading, the process of training and extremely rapid reading rates are the defining characteristics. For all three rapid reading methods a number of theoretical questions can be posed, and various empirical paradigms developed. Whatever the reading method, the procedure of presenting story materials should be as natural as possible so that usual reading processes are not seriously disrupted. Self-paced and experimenter-paced reading rates were used here, and both provided reasonable and consistent results.