A new technique for measuring word processing time in reading

PETER DIXON

University of Alberta, Edmonton, Alberta, Canada

This paper reports some preliminary data on a "window-pointer" technique for measuring word processing time in reading. In this technique, almost all of the letters are masked on a page of text shown on a video monitor. The letters can be seen only through a "window" showing one word on one line. Subjects control the horizontal location of the window by moving a pointer left and right. When subjects want to move the window to the next line, the window is moved off the right end of the current line, and the pointer is moved back to the left. The window then reappears at the beginning of the next line. After some practice, subjects can move the window in approximate synchrony with their eyes, and can read the passage without any difficulty. The length of time the window spends on each word can then be used as a measure of word processing time. An analysis of these window durations indicates that they are similar in many respects to the gaze durations obtained with eye-movement monitoring techniques.

The duration of eye fixations during reading is an important source of information about the reading process, and has been used by a number of researchers to test specific hypotheses about comprehension (e.g., Carpenter & Just, 1977; Rayner, 1977; Scinto, 1978). Even more importantly, the duration of fixations on a word may represent the total time readers spend processing that word. Just and Carpenter (1980) refer to this notion as the "immediacy" hypothesis: Readers may process a word as fully as possible while they are looking at it. To the extent that the immediacy hypothesis is true, fixation duration is a measure of word processing time, and can be used to construct models of the reading process.

Unfortunately, the measurement of fixation duration is technically difficult. It requires expensive equipment, often entails a serious commitment in laboratory and software development, and usually does not approach the data-collection efficiency found in other cognitive psychology paradigms. In this paper, I will describe an alternative way of measuring word processing time in reading that eliminates much of this expense. The technique will be referred to as the "window-pointer" paradigm. Essentially, it depends on the ability of subjects to indicate, while reading, where they are looking in a passage. Instead of measuring eye position directly, the paradigm allows subjects to provide that information themselves.

In the window-pointer task, a page of text is shown on a video monitor, but with all of the letters blanked

out. The letters are visible only through a small "window" on a single line (see Figure 1). The subject controls the location of this window by moving a pointer: When the pointer is moved to the right, the window moves to the right, and when the pointer is moved to the left, the window moves to the left. To move the window to a succeeding line, the subject moves the window off the right end of the current line and then moves the pointer all the way back to the left. The window then reappears at the beginning of the next line.

The goal of this technique is to get an inexpensive measure of eye position. Hopefully, practiced subjects would be able to move the window quickly and automatically, using only a minimum of processing capacity. In that event, they may be able to move the window in close synchrony with their eyes and process the text just as they would without the window. Thus, the position of the window at any point in time would provide a simple measure of eye position. In order to approach this ideal, the pointer was made as natural and as easy to use as possible, and subjects were given ample practice

Figure 1. Upper left-hand portion of a typical passage display.

This research was supported by Grant A8263 from the Natural Sciences and Engineering Research Council of Canada. I would like to thank Renée Elio for comments on this manuscript. Requests for reprints should be sent to Peter Dixon, Department of Psychology, University of Alberta, Edmonton, Alberta, Canada T6G 2E9.

in using it. The results below suggest that actual use of the window was close enough to ideal to yield useful data on the reading process.

This paradigm is an extension of the "moving-window" paradigm studied by Just, Carpenter, and Woolley (1982). They presented a page of text with all the letters replaced with dashes. Only the letters in a single word were actually shown. Subjects pressed a button to move this window on the text to succeeding words. The principal difference between the Just et al. technique and the present one is the method of control. In Just et al., subjects indicated *when* they wanted the window to move, whereas subjects here indicated *where* they wanted the window to move. The latter may be more natural, and the data to be presented below suggest that it may be closer to the usual reading process when unconstrained by a window. Moreover, the window-pointer paradigm allows readers to make regressions by returning the window to a previous word on a line, unlike the Just et al. technique. Although the number of such regressions was small, they may be critical in allowing subjects to read in a natural way.

Other self-paced reading-time paradigms have been studied by Mitchell and Green (1978) and Aaronson and Scarborough (1976). Mitchell and Green presented text three words at a time in the center of the screen. The next group of words was shown whenever the reader pressed a button. The technique used by Aaronson and Scarborough was similar, except that they presented only a single word at a time. Although these methods can generate useful data about reading (e.g., Aaronson & Ferres, 1983), reading with the present method is probably more like unconstrained reading. For instance, Just et al. (1982) found slower reading time and less resemblance to eye-fixation data with Aaronson and Scarborough's method than with their own moving-window method. One might conjecture that Aaronson and Scarborough's method suffers because it does not provide information about the length of upcoming words.

This paper is organized as follows: First, a detailed description is provided of the paradigm and its implementation. Then some representative results are examined in terms of three measures: overall reading speed, the word-length effect, and the end-of-sentence effect. These are compared with similar measures found with eye-fixation and buttonpressing paradigms. Finally, the advantages and disadvantages of this paradigm and several variations are considered.

METHOD

Apparatus and Implementation

The pointer apparatus is shown in Figure 2. Subjects sat in front of a standard video monitor with their preferred hand and forearm resting on a movable 18 x 40 cm platform. Subjects lightly grasped a knob near the end of the platform, and could pivot the platform back and forth around a point near their elbow. The total extent of the movement was about 45 deg, which meant that the knob moved through a distance of about 22 cm. The platform was mounted on low-friction bearings, and

Figure 2. Subject using the pointer apparatus.

was attached at the pivot point to a voltage potentiometer. The voltage signal was processed by an Apple II microcomputer system using an analog-to-digital interface manufactured by John Bell Electronics. A separate timer allowed the computer to sample the signal every 16 msec.

The computer used this signal to decide where the window should be positioned. The window was always one word wide, so this amounted to deciding which word should have its letters displayed on the video monitor. The total distance over which the pointer could move was divided into 16 equal sectors. When the pointer was in sector N, the letters in the Nth word were shown on the screen; all other letters were masked. The window moved to a new word only when the pointer was moved more than halfway into a new sector; this kept the window from alternating back and forth between two words when the pointer was close to a sector boundary. When the window was moved past the last word on a line, the window disappeared, and all of the letters on the line were masked. The window reappeared at the beginning of the next line when the pointer had been moved all the way back to the left. In order for this algorithm to work, stimulus passages had to be edited so that they had no more than 15 words per line.

The paradigm was implemented on an Apple II microcomputer. The Apple's limited speed and memory posed a problem for how data was to be collected and stored. A simple solution would have been to calculate the length of time the window spent on each word right after the window moved off that word, and then store that time along with an indication of which word had been shown. This solution was judged to be too slow; it was anticipated that it would introduce an excessive delay in moving the window when the pointer was moved quickly. Another solution would have been to simply store the location of the pointer every 16 msec when its position was sampled. However, this technique would have required an excessive amount of memory space. The solution used in the present implementation was to store the times at which the window moved from one word to another. Two lists of times were kept: one for forward and one for regressive movements. A subsequent analysis program collated the lists and calculated the length of time each word in the passage was shown.

Passages were shown on a standard 20-cm video monitor at a distance of 50-70 cm. At that distance, letter spaces subtended about 15 min of horizontal visual angle. The display showed 24 80-character lines, each separated by about 35 min of visual angle. Letters were shown only within the window. Outside the window, a mask was shown in place of each letter (see Figure 1). For most letters, the mask was a square the size of one

letter space. A similarly sized square was added above the first for ascenders, and below the first for descenders. Thus, the shape of the mask indicated the size and overall shape of each word. Punctuation outside the window was left intact.

The software used to control the window was divided into three parts: an interrupt-driven device handler, an experiment control program written in Pascal, and assembly language subroutines. The device handler was a set of general-purpose routines to do response input and timing in experiments. It sampled the analog-to-digital interface every 16 msec and returned its value to the Pascal program. The Pascal program performed file input/output and some program initialization, but the actual control of the window was performed by several assembly language subroutines. This was necessary because the window moved too slowly when it was controlled directly by the Pascal program, and a noticeable lag between pointer and window movement occurred if the pointer was moved quickly. The assembly language implementation was substantially more efficient, and it was difficult to discern any such lag. Altogether, the program contained about 250 lines of Pascal code and about 400 lines of assembly language code for the subroutines.

Materials

Forty-six stimulus passages were prepared to use with the window. These were 236 words long on the average, with a range of 193 to 280 words. Passages contained an average of 11 sentences and occupied from 18 to 23 lines on the video screen. They were selected from a wide range of magazine articles and chosen to be moderately difficult for the subject population. The passages were edited so that they would be of similar lengths, have a well-formed discourse structure, and avoid very infrequent or technical words. Two comprehension questions were prepared for each passage. The first concerned a specific detail in the passage and could be answered by recalling a single phrase or sentence. The second question required subjects to select a title for the passage. A booklet of questions was prepared, with one page for the two questions on each passage.

Procedure

Subjects first read five randomly selected passages without the window in order to obtain a baseline unconstrained reading rate. The procedure was as follows: When the computer was ready to begin, the word "Ready" was displayed in the center of the screen. The passage was displayed 1.8 sec after the subject pressed a button. When the subject finished reading the passage, he or she pressed the button again and the screen went blank. Immediately after reading each passage, subjects turned over a page in the question booklet and wrote out the answers to the two questions.

After the 5 unconstrained passages, subjects read 10 randomly selected passages using the window pointer. The procedure followed was similar to the unconstrained reading task. The word "Ready" appeared as before, and subjects initiated passage presentation by pressing a button. However, when the passage was shown, all of the letters were masked. The window appeared at the beginning of the first line when the pointer was moved all the way to the left. The passage automatically disappeared when the window was moved off the end of the last line.

Ten undergraduates at the University of Alberta participated in the experiment for course credit. The entire session took about 50-60 min.

RESULTS

Three aspects of reading performance with the window were analyzed: overall reading speed, the word-length effect, and the end-of-sentence effect. The first aspect bears on the issue of how natural the window

was to use and how quickly subjects learned to use it. An analysis of the word-length effect indicates whether the length of time subjects spent on a word was sensitive to the local properties of that word. Finally, the end-of-sentence effect is the finding that subjects pause at the end of sentences. If subjects slow the window down at the end of sentences, it would indicate that the paradigm is sensitive to the kinds of sentence-level effects found with eye fixations.

Reading Speed

The average time per word for the 10 window passages is shown in Figure 3. Subjects improved substantially over the first 5 passages, continued to improve somewhat in the last 5, and reached a speed close to 75% of their unconstrained reading speed. The time to move the pointer back to the left between lines was not included in these times. Subjects' speed in unconstrained reading was 165 words per minute and is indicated by a dashed line in Figure 3.

Although subjects were clearly reading slower with the window than without it, this paradigm seems to be an improvement over buttonpressing paradigms. For instance, in the best paradigm studied by Just et al.

Figure 3. Reading time per word as a function of serial position. The dashed line indicates the average time per word for the five passages read without the window.

(1982), subjects averaged 54% of their unconstrained speed. Although differences in reading goals and materials make comparisons difficult, the subjects in this task seem to read faster than those of Just et al. without any practice at all. Thus, the window pointer may be more natural and easy to use.

Word-Length Effect

Figure 4 shows the length of time spent on each word as a function of its length and the length of the previous word. These times were computed using only the last five passages. As found with eye fixations and with buttonpressing paradigms, people spent more time on longer words. This effect may be due in part to word frequency: Longer words tend to be less frequent, and less frequent words require more lexical access time (e.g., Gordon, 1983). Most of this effect (68%) occurred the first time the word was in the window, rather than on regressive movements of the window. The slope of the increase was 16 msec per character, which is quite similar to the slope of 15 msec found by Just et al. (1982) in their moving-window paradigm, but less than the 32 msec per character found by Just and Carpenter (1980) with eye fixations.

However, an analysis of the window duration as a function of the length of the previous word (the dashed line in Figure 4) showed a similar increase with word size. This means that some of the additional processing necessary for each letter in a word carried over and affected the duration the window spent on the next word as well. An estimate of the total time needed for each additional letter can be found by adding the slope of the word N+1 effect to that for the word N effect. This total is 27 msec per letter. Twenty-seven milliseconds is much larger than the 15-msec effect found by Just et al. (1982) and approaches the 32-msec effect found using eye movements. Thus, a tentative conclusion is that the window-pointer technique captures more of the variance attributable to word length than a buttonpressing paradigm. However, some of the extra time per letter occurs after the window has actually moved to the following word.

Figure 4 shows that window durations following words of length one are exceptionally long. This one data point does not fit with the general trend of longer durations following longer words. The only word of length one in these passages was the indefinite article "a." One account of the long times following "a" is that an indefinite article often precedes the first occurrence of significant nouns in a passage. Such words may require more processing than others, and it would be reasonable for them to have longer window durations.

An analysis of variance with subjects as a random factor was performed on the word N and N+1 effects. (Since the passages were randomly selected for each subject, any variance due to materials was purposely confounded with the subject variance.) There was a

Figure 4. Effect of word N length on window duration for word N and word N+1.

strong effect of the length of word N on both the time spent on word N [$F(11,99) = 9.24, p < .001$] and the time spent on word N+1 [$F(11,99) = 3.95, p < .001$]. Ninety-five percent of the word N effect was due to a linear trend [$F(1,9) = 27.47$]; and, with the exception of the first point, 80% of the effect on word N+1 was also due to a linear trend [$F(1,9) = 33.29, p < .001$].

End-of-Sentence Effect

The end-of-sentence effect for the last five passages is shown in Figure 5. As can be seen, there was a strong tendency to slow down on the last word in a sentence. However, subjects also spent more time on the first word of the next sentence. This trend is similar to the word N+1 effect found for word length. That is, the processing that occurs on one word seems to carry over and affect the time spent on the following word. In the present case, the following word is the first word in the next sentence. An analysis of variance similar to that used for the word-length effects showed a significant effect of word position relative to the end of the sentence [$F(7,63) = 6.03, p < .001$].

DISCUSSION

The results indicate that subjects can adapt quickly to the window-pointer paradigm, and seem to require only about 15-20 min of practice to become fairly proficient. Their overall reading speed is less than that without the window, but seems to be better than those found in other word-by-word paradigms. Moreover, the

Figure 5. Window durations near sentence breaks.

same kinds of effects that are found with eye fixations can be obtained here. For instance, window durations increased 27 msec for each letter in a word, and a total of 208 msec near the ends of sentences. The only difference is that, in the present task, such effects seemed to carry over to the following word. This kind of carry-over often does not occur with eye fixations (e.g., Carpenter & Just, 1983).

The carryover may occur here (but not with eye movements) because the pointer has more inertia than the eye does. That is, any decision to slow down or speed up the window may take somewhat longer to execute than with eye movements. For example, suppose that a reader is taking longer than usual to process word N in a passage. A signal may be generated to pause the pointer on that particular word. However, if it takes a relatively long time to physically pause the pointer, the window may actually be on word N+1 when the pause occurs. Thus, both the window duration on word N and that on word N+1 would reflect processes related to word N.

This processing carryover may or may not reflect any differences in the reading process itself. On one hand, it may be that, as long as no additional visual information is required from the text, it is immaterial whether the window pauses on word N or on word N+1. In this view, only the total duration of the pause (including both word N and word N+1) would be important for models of reading comprehension. On the other hand, it is possible that the additional inertia involved

in moving the window may significantly affect some components of the reading process. For instance, readers may read at a more uniform rate, or may fixate words they would otherwise skip. The extent to which the window does affect the reading process can be determined only by a careful consideration of the patterns of word processing times with and without the window.

An important advantage of the window-pointer paradigm is that subjects can make regressions within a line, unlike buttonpressing paradigms such as Just et al.'s (1982) moving-window paradigm or Aaronson and Scarborough's (1976) word-by-word technique. When they need to, readers can easily move the window back to any previous word on the current line. Such movements were relatively rare in these data: Only 10% of the total reading time was spent on second looks at words. Nevertheless, the mere fact that readers could back up when they wanted to may have been important. For instance, readers may have been more willing to move the window at a fast rate because going too fast would not cause any difficulties; they could simply go back and take a second look at any words that needed more processing. In contrast, buttonpressing paradigms do not have a provision for backing up, and consequently may engender a more conservative strategy. Because a word can be looked at only once, readers may delay until they are sure all the processing for that word is completed before going on to the next. As a consequence, more time may be taken with some words than is really necessary, and the overall reading speed may be slowed.

Two variations of the window-pointer paradigm have also been implemented and tested. For some hypotheses about the reading process, these variations may be more appropriate than the one discussed so far. The first variation changes the way the pointer position is mapped onto window position. In the original version, the number of the word being shown on a line was determined by the number of the sector the pointer was in. For example, if the pointer was in the third sector, the third word on the line would be shown. This means that if the pointer was moved uniformly from left to right, each word would be shown an equal amount of time. In the new version, the position of the pointer directly determined the position of an imaginary cursor on the current line, and the word at that cursor position was the one shown in the window. This means that if the pointer was moved uniformly from left to right, the amount of time each word was shown would vary, and longer words would be shown for longer periods of time.

This version of the paradigm may be easier for subjects to use. Since the length of time needed to process each word generally increases with the word's length, the pointer may end up being moved at a more consistent rate in this version, and less effort would be required to speed the pointer up or slow it down. The

disadvantage of this version is that it makes it difficult to study effects at the lexical level. If more time was spent on a particular word, it may be because the word required more processing or because the pointer was moved uniformly and the word happened to be longer. However, if the effects of interest are at the level of clauses or phrases, this version may be entirely appropriate.

The second variation uses a window that is a fixed number of letters wide instead of always being one word wide. Subjectively, window movement with this version seems to be much smoother and less disruptive than in the original. Instead of jumping from word to word across the line, the window appears to move almost continuously. The fixed-width version also provides more peripheral information, since the first few letters of a word could become visible while the preceding word is still being processed. Some preliminary pilot data suggest that overall reading speed with this window may be substantially faster than with the previous two versions, perhaps because of this extra peripheral information. However, this version has the same problem as the previous one with respect to lexical effects: Longer words may be displayed for longer periods of time for reasons that are unrelated to processing difficulty. Another problem is that since the window is a fixed number of characters wide, several words may be visible or partially visible at the same time, and it is unclear which is being processed. Both of these problems may not be serious if the unit of reading-time measurement is a phrase or clause. Under such circumstances, this version may be the most appropriate of the three.

In summary, the window-pointer paradigm has been demonstrated to be a useful technique for studying word processing time in reading. Although subjects read somewhat slower using the window than they do without it,

the reading process may be substantially unaltered. In particular, a word-length effect and an end-of-sentence effect were found, both similar in size to the effects found with eye-fixation measurements. Two variations of the window-pointer task were also discussed; these variations may be appropriate when the processing unit of interest is larger than the word.

REFERENCES

- AARONSON, D., & FERRES, S. (1983, November). *A self-paced word-by-word reading procedure*. Paper presented at the meeting of the Society for Computers in Psychology, San Diego, CA.
- AARONSON, D., & SCARBOROUGH, H. S. (1976). Performance theories for sentence coding: Some quantitative evidence. *Journal of Experimental Psychology: Human Perception and Performance*, 2, 56-70.
- CARPENTER, P. A., & JUST, M. A. (1977). Reading comprehension as eyes see it. In M. A. Just & P. A. Carpenter (Eds.), *Cognitive processes in comprehension*. Hillsdale, NJ: Erlbaum.
- CARPENTER, P. A., & JUST, M. A. (1983). What your eyes do while your mind is reading. In K. Rayner (Ed.), *Eye movements in reading: Perceptual and language processes*. New York: Academic Press.
- GORDON, B. (1983). Lexical access and lexical decision: Mechanisms of frequency sensitivity. *Journal of Verbal Learning and Verbal Behavior*, 22, 24-44.
- JUST, M. A., CARPENTER, P. A., & WOOLLEY, J. D. (1982). Paradigms and processes in reading comprehension. *Journal of Experimental Psychology*, 87, 329-354.
- JUST, M. A., CARPENTER, P. A., & WOOLEY, J. D. (1982). Paradigms and processes in reading comprehension. *Journal of Experimental Psychology: General*, 111, 228-238.
- MITCHELL, D. C., & GREEN, D. W. (1978). The effects of context and content on immediate processing in reading. *Quarterly Journal of Experimental Psychology*, 30, 609-636.
- RAYNER, K. (1977). Visual attention in reading: Eye movements reflect cognitive processes. *Memory & Cognition*, 4, 443-448.
- SCINTO, L. F. (1978). Relation of eye fixations to old-new information in texts. In J. W. Senders, D. F. Fisher, & R. A. Monty (Eds.), *Eye movements and the higher psychological functions*. Hillsdale, NJ: Erlbaum.