

Combined eyetracking and keystroke-logging methods for studying cognitive processes in text production

ÅSA WENGELIN

Lund University, Lund, Sweden

MARK TORRANCE

Nottingham Trent University, Nottingham, England

KENNETH HOLMQVIST

Lund University, Lund, Sweden

SOL SIMPSON

SR Research Ltd., Osgoode, Ontario, Canada

DAVID GALBRAITH

Staffordshire University, Stoke-on-Trent, England

AND

VICTORIA JOHANSSON AND ROGER JOHANSSON

Lund University, Lund, Sweden

Writers typically spend a certain proportion of time looking back over the text that they have written. This is likely to serve a number of different functions, which are currently poorly understood. In this article, we present two systems, ScriptLog+TimeLine and EyeWrite, that adopt different and complementary approaches to exploring this activity by collecting and analyzing combined eye movement and keystroke data from writers composing extended texts. ScriptLog+TimeLine is a system that is based on an existing keystroke-logging program and uses heuristic, pattern-matching methods to identify reading episodes within eye movement data. EyeWrite is an integrated editor and analysis system that permits identification of the words that the writer fixates and their location within the developing text. We demonstrate how the methods instantiated within these systems can be used to make sense of the large amount of data generated by eyetracking and keystroke logging in order to inform understanding of the cognitive processes that underlie written text production.

Text production—the generation of multiple sentences of coherent text—involves a large and complex set of cognitive processes. Despite a number of studies, dating back to the early 1980s (reviewed in, e.g., Alamargot & Chanquoy, 2001), understanding of these processes and how they are coordinated remains limited. This contrasts with the relatively detailed understanding that we have of the cognitive processes associated with text comprehension (e.g., Crain-Thoreson, Lippman, & McClendon-Magnuson, 1997; Just & Carpenter, 1980, 2002; Kintsch & Rawson, 2005; Kintsch & Van Dijk, 1978) and the extensive literature on spoken production (e.g., Allum & Wheeldon, 2007; Clark & Fox Tree, 2002; Eklund, 2004; Goldman-Eisler, 1968; Levelt, 1989; Postma, 2000; Rapp & Goldrick, 2000; Raupach, 1980). This is partly due to the complexity of the processes being studied. It is partly also because studying

production is inherently more difficult than studying comprehension. Reading researchers can set participants clearly defined tasks with relatively easily measured outcomes. Researchers exploring production must necessarily act more as observers, studying the writing process as it unfolds in real time (e.g., Matsushashi, 1987b; Olive & C. M. Levy, 2003). They therefore require sophisticated methods for both determining and describing writers' spontaneous text-producing behavior. A particular challenge is to capture the intertwining of different writing subprocesses. Writing involves a combination of (at least) content determination, formulation, transcription, error monitoring, and reviewing. These are rarely, if ever, performed only as discrete, deliberate steps but typically occur as multiple, iterative cycles (Hayes & Flower, 1980). Text production therefore draws on mechanisms that have been explored in the con-

text of research on speech production and on mechanisms that have been explored in the study of text comprehension. However, complex interaction between these production and comprehension processes is unique to writing.

Writing researchers have used, broadly, two kinds of real-time or “online” methods for studying writing behavior. There is a tradition of studies in which participants either think aloud while writing (e.g., Hayes & Flower, 1980; van den Bergh & Rijlaarsdam, 2007) or categorize their own behavior in response to concurrent, random probes (e.g., Kellogg, 1987; Torrance, Fidalgo, & García, 2007). Self-report measures of this sort have been used fairly extensively. They are capable of providing rich information about higher level processes in text production. Alternatively or in addition to using self-report methods, researchers have simply observed and recorded text production processes without requiring any additional report or response from the writer. Observations can simply involve studying live, videotaped, or screen-capture writing sessions (e.g., Degenhardt, 2006; Matsuhashi, 1987a; Myhill & Jones, 2007) or recording writers composing by dictation (Schilperoord, 1996). They also include more sophisticated methods involving logging the keyboard activity of writers composing at the computer (Leijten & Van Waes, 2006; C. M. Levy, 1994; C. M. Levy & Ransdell, 1994; Strömquist & Karlsson, 2002; Strömquist & Malmsten, 1998; studies collected in Sullivan & Lindgren, 2006). There exist well-developed tools for keystroke capture and analysis, tailored to the needs of writing researchers—notably, Inputlog (Leijten & Van Waes, 2006) and ScriptLog (Strömquist & Karlsson, 2002). These methods provide a different kind of information from that afforded by self-report and are less reactive (Jansen, Van Waes, & van den Bergh, 1996).

Findings from keystroke and other direct observation studies suggest that writing typically proceeds in alternating bursts of inscription (handwriting/keyboard) and pauses during which nothing is written. The lengths of these pauses vary with features of the text that is being composed; where pauses occur is determined, in part, by rhetorical features of the emerging text (Foulin, 1998; Schilperoord, 1996). This indicates that pause locations and durations are not arbitrary but, instead, point to mental activity associated with writing processes. However, we do not yet have a good understanding of the nature and function of this activity (Schilperoord, 2003). Some pauses, or some parts of some pauses, are likely to be associated just with thinking about what to say next, with no additional input from the text already produced (Flower & Hayes, 1981). However, a specific feature of text production, as compared with spoken language production, is that the written text stays visible, accessible, and, hence, revisable during the entire production process. Thus, the emerging text could function as a visual external storage, which could be used to decrease the cognitive load of the writer. Despite the case that most models of writing (Alamargot & Chanquoy, 2001) include a component of reading/reviewing, we do not know how or to what extent this visual trace of the emerging text is used by the writer. However, except when visual feedback is prevented either by the writing medium (e.g., composing by dictation) or under experi-

mental conditions in which participants are prevented from seeing what they write (e.g., Olive & Piolat, 2002), pausing is also likely to be associated with writers’ reviewing what they have written. This could take place either at a local level (Britton, 1982; Matsuhashi, 1987a), with writers reading or scanning within the sentence that they are currently composing, or more globally, with writers looking further back into their text (Pianko, 1979).

Processing the text already written could serve one or more of several functions. Writers may look back to detect surface-level errors or to find errors that they have already detected but delayed correcting. It is also clearly the case that writers use visual information from the text when they explicitly review and revise what they have written. Here, we are thinking of the kind of activity that might occur when writers complete a rough draft and then go back over what they have written while making changes, an activity that appears to play an important role in the production of extended text, at least for some writers and for some kinds of writing task (Hayes, 1996). This agrees with the Hayes and Flower (1980) model of writing in which the main function of reviewing is to serve the editing processes. However, in his revision of the Hayes and Flower model, Hayes suggested that writers may also review the text produced so far in order to create “opportunities for improvement that do not necessarily stem from problems” (p. 15). This was supported by a think-aloud study of Breetvelt, van den Bergh, and Rijlaarsdam (1996), which showed that rereading was strongly correlated with text generation processes. Thus, it appears reasonable to assume that writers may look back not only to evaluate what they have written (assessing whether or not it coheres and is likely to achieve rhetorical goals), but also to help with generation of the next part of the text: to remind themselves of what has already been covered, to refresh memory about the syntactic frame that is currently being used, or to find prompts that allow search and retrieval of more ideas (e.g., Holmqvist, Holsanova, Johansson, & Strömquist, 2004). We have some understanding of the processes that underlie proofreading familiar text (Bisaillon, 2007; Hacker, Plumb, Butterfield, Quatham, & Heineken, 1994; B. A. Levy, 1983; Pilotti, Maxwell, & Chodorow, 2006), and there is a small literature of writers’ revision processes (see, e.g., the studies collected in Allal, Chanquoy, & Largy, 2004). However, empirical research to date has not directly explored the detail of how writers use visual feedback from their own texts—that is, what they look at and when, in what contexts (from a process point of view), and to what extent it can be related to “normal reading processes.” However, on the basis of the discussion above, it appears reasonable to make the following suggestions: (1) Writers might look at their emerging text to prompt content generation; (2) writers might look at their emerging text to manage reference (presumably, specifically anaphoric reference) and so maintain cohesion; (3) writers might look at their emerging text to detect and/or correct errors; (4) writers might look at their emerging text to compare it with an (internal or external) outline of intended content; and (5) writers might look at their emerging text as part of a deliberate and explicit metacognitive decision to revise what they have written.

As has already been mentioned, these issues—and specifically, the first three of these—have not previously been empirically investigated in detail for text production. One of the reasons for this has probably been the lack of available and affordable tools. The tools and methods that we describe in this article combine keystroke logging with tracking of writers' eye movements to provide information about when writers look at a text that they have already produced, where they look within this text, and how this attention to what they have already written is interleaved with ongoing production. We believe that this kind of information has the potential to provide new insights into the cognitive processes that underlie how writers compose and edit their texts.

There are established methods for tracking readers' eye movements as they process given texts, and these have yielded an extensive literature describing writers' eye movements (see the reviews by Liversedge & Findlay, 2000; Rayner, 1998; Starr & Rayner, 2001) and sophisticated models of readers' eye control (Reichle, Rayner, & Pollatsek, 2003). Reading tends to proceed through a sequence of rapid saccades (when the eye moves between and within words) and fixations (when the gaze remains relatively fixed within a single word). Reading is often identified as three or more consecutive fixations in the same direction, with a maximum time between the end point of one fixation and the start point of the succeeding fixation. Fixated words are attended to and processed (automatically and without the possibility of deliberate suppression). Some information is available from words to the side of the fixated word. These findings suggest that if we know that a familiar word is being fixated by the reader, it is reasonable to assume that the reader has available morphologic, syntactic, and semantic information associated with that word. In the context of text production, this implies that if a writer moves his or her gaze away from the point of inscription and fixates one or more words back within the text, syntactic and semantic information associated with that word becomes available to inform what he or she writes next. And the fact that he or she looked back into the text in the first place suggests that he or she (or automatic writing-related mental mechanisms) has identified informational needs that might be fulfilled by this action. Study of when and where the writer fixates his or her developing text therefore has the potential to inform theories about the complex cognitive processes that underlie written production. It is important to note, however, that reading behavior is likely to be influenced by reading goals (e.g., Lorch, van den Broek, & Pressley, 1997) and the perspective of the reader (Kaakinen & Hyönä, 2008). To our knowledge, no typology of reading goals has been established, but it should be kept in mind that the reading goals and the perspectives of the writer may, in some cases, be very different from those of any reader of a finally edited text written by someone else.

Although a small number of studies have explored typists' eye movements within the (experimenter-provided) text that they are copying (Butsch, 1932; Inhoff & Gordon, 1998), with the exception of one or two unpublished studies, writers' eye movements within text that they are producing

remains more or less unexplored. An important exception to this is work by Alamargot and coworkers using the *Eye and Pen* system (Chesnet & Alamargot, 2006), which permits detailed analysis of eye movement during handwritten composition. Participants write onto a digitizing tablet while wearing a head-mounted eyetracker. The analysis component of the system brings together pen traces and eye movement data to provide a detailed real-time description of both pen and fixation location during composition. Some early findings using *Eye and Pen* are summarized in Alamargot, Chesnet, Dansac, and Ros (2006).

This article describes two additional, and complementary, approaches to exploring writers' eye movements: *ScriptLog+TimeLine* (Andersson et al., 2006; Strömquist & Karlsson, 2002) and *EyeWrite* (Simpson & Torrance, 2007).¹ These systems differ from *Eye and Pen* in that they study writers composing by keyboarding. This provides a writing medium that, for some writers, will be more frequently used than handwriting. It also makes possible collection of precise timings for the creation of each character within the emerging text. This information can then be combined with information about the writers' eye movements to provide a synchronized log of both keyboard and eye activity.

Interpreting combined keystroke and eye movement data, however, presents technical problems that do not arise when participants are reading static text. The software that typically accompanies off-the-shelf eyetrackers provides facility for mapping eye movement information onto screen coordinates. Therefore, when participants are looking at the screen, we know, leaving issues of eyetracker accuracy and size of attentional window aside, where on the screen they are looking. When the screen displays static text, it is relatively straightforward to determine which words are being fixated. However, during composition, the text that appears at a particular location on the screen changes frequently and unpredictably. The location of a particular word or sentence in the text as it stands once the writer has finished is not necessarily the same as its location when it was first inscribed or when it was subsequently read or edited. This is particularly the case if the text fills more than one screen and scrolling becomes necessary. Analysis methods based on static areas of interest defined in terms of screen locations therefore will not work when exploring writers' eye movements. There are, broadly, two approaches to overcoming this problem. One is to display synchronized keystroke and eye movement data in a form that makes it easy for the researcher to examine them together. Alternatively, knowledge about how the text is displayed on the screen can be used to reconstruct screen states as they changed during writing.

These approaches are adopted, respectively, by the *ScriptLog+TimeLine* and *EyeWrite* systems that we describe in this article. Both systems are experimental and have been developed within specific laboratory contexts and to meet specific research needs. The remainder of this article describes these two systems and provides a summary of some preliminary findings. We describe the systems primarily to indicate what is possible in terms of collection and analysis of eye movement data from writers and to suggest different methods by which this might be achieved, rather than as

complete tools that could be imported without modification into other research contexts. We have two aims. First, we want to provide sufficient understanding of the methods that we have been using in our research to enable interested researchers to make use of them within their own studies, either by using the particular software and hardware that we describe or by developing their own software. Second, and more generally, we hope to convince readers that knowledge of where writers look within their texts has the potential to inform our understanding of the cognitive processes that underlie written production.

ScriptLog+TimeLine

In the following sections, we will describe how data from an existing eyetracking program (SensoMotoric Instruments [SMI] iView X) has been integrated with data from an existing keystroke-logging program (ScriptLog) and how the vast amount of data generated by these two programs can be merged and visualized using a third program called TimeLine. A short text production example will be used to illustrate the kinds of findings that this combination of systems makes possible.

ScriptLog. ScriptLog (Strömqvist & Karlsson, 2002; Strömqvist & Malmsten, 1998) is a keystroke-logging program designed for both laboratory and field research on the written language production process. It works in two modes: the recording mode and the analysis mode. In order to keep the number of possible variables of analysis limited and controlled, ScriptLog provides its own simple editor that is similar to NotePad in Windows. The writer can type in text and can use the mouse or cursor keys to move around in the text and to cut and paste. During the writing session, there are no formatting options, and so the writer cannot change the presentational features of his or her text. However, before the writing session starts, several formatting features can be changed in order to suit the writer and/or the purpose of the experiment. ScriptLog can be set to present different stimuli or reference materials, such as pictures, sounds, and text, which can be used to elicit composition. ScriptLog is designed for experimental and semiexperimental research, and formatting options are limited. This creates a writing environment that is not fully “natural” but avoids the situation in which the researcher has too many variables to control.

In analysis mode, ScriptLog outputs the final edited text, as well as *linear texts* and *log files*. The log file shows, for each keystroke, the exact time the keypress was made (calculated from the beginning of the recording), the category of the key (e.g., number or capital letter), the current position in the text, and which key was pressed. The log file can be used to generate detailed information on typing statistics, pauses, and edits. The corresponding linear text shows in *text format* how the text was produced, in chronological order. All keystrokes, mouse movements, and pauses are shown, but no details on the exact timing for each keystroke are given. An example of a linear text is shown below. The completed text, produced by a Swedish university student performing an expository writing task (in Swedish) in response to a short film (Berman & Verhoeven, 2002), began as follows:

Filmen jag såg nyss bestod av korta sekvenser ur några ungdomars vardag. I princip alla sekvenser visade hur ungdomar utsätts för jobbiga situationer av sina kamrater.

(The film that I just saw consisted of some short sequences from the daily life of some young people. In principle, all the sequences showed how young people are exposed to difficult situations by their peers.)

The linear file shows all of the keystroke activity associated with the production of these two sentences:

```
<5.762> Filmen jag såg nyss handla<BACKSPACE6>bestro<BACKSPACE2>od av korta klipp
<BACKSPACE4>s<BACKSPACE2>sekvenser
ur några ung<3.639>domars vardag. <9.930>
<4.027>Nästan<3.825><BACKSPACE9>.
I princip alla sekvenser ha<BACKSPACE2>
visar<BACKSPACE>de hur ungdomar utsätts
för jobbiga situationer u<BACKSPACE>av
siha<BACKSPACE2>n a<BACKSPACE2>a
kamrater.
```

The example shows that the writer paused for 5.76 sec before she started to write. Then she wrote “Filmen jag såg nyss handla” (“The film that I just saw deals with”). Then, before adding the past tense ending to the verb “handla” (“deal with”), she changed her mind, deleted it, and wrote “bestod av” (“consisted of”) instead. In the middle of the word “bestod,” she made a typo, adding an “r” that she then deleted before completing the word. Composition then continued in a similar way until the end of the sentence, where the writer made a longer pause of 9.93 sec, then typed a space, then paused again for another 4.03 sec, and finally started writing the next sentence. These pauses are marked in bold in the example.

What we would like to know is what the writer was doing during these pauses. This information is clearly not provided by the keystroke record alone.

Combining ScriptLog with SMI iView X eyetracking and Polhemus headtracking. Later versions of ScriptLog allow connection to and synchronization with an eyetracker. In our research, we have combined ScriptLog with a head-mounted eyetracker with headtracking. The eyetracker calculates a vector for the gaze direction at the eye. This, combined with information about head position, allows for a real-time calculation of the point at which the combined eye-head vector meets the monitor or the keyboard. When all the coordinate systems are in place, we are able to get precise data on where the writer looks and when. Previously or retrospectively defined areas of interest (in terms of screen coordinates) can be used to calculate total dwell time within particular areas of the writer’s visual field.

For each writer and each text, the eyetracking part of our setup outputs both a video of the visual field, with overlaid gaze cursor, time stamp, and sound, and a log file. The log file gives the following data for every 20 msec: area of interest, gaze coordinates in the coordinate system of that area, head position and head orientation, eye position, and time.




Figure 1. TimeLine screenshot showing gaze distribution across the end of the first sentence and the beginning of the second sentence (23.895–49.900 sec).

Thus, by combining an existing keystroke-logging program (ScriptLog) with a stand-alone eyetracking program, we end up with two log files: the keystroke log and the eyetracker log. However, these log files are not easy to read or directly interpret. Log files are typically files with very low readability. We therefore use a visualizing tool (Andersson et al., 2006), now named *TimeLine*, based on the multimodal time-coded score sheets method developed by Holsanova (2001). Screen coordinates could be fed back to ScriptLog. However, this has not yet been implemented.

TimeLine. TimeLine is a visualizing tool that uses the ScriptLog record in combination with eyetracking data to provide a graphical representation of verbal and visual flow. It gives an overview of how the writer's attention is distributed among the keyboard, the monitor, and elsewhere.

Figure 1 shows the end of the first sentence and the beginning of the second (*ung*<3.639>*domars vardag*. <9.930> <4.027>*Nästan*) of the fragment that was discussed in the previous section, but with the addition of a visual representation of keyboard and eye activity, including where the writer looked during the long pause between the first and second sentences. The figure is divided into three tiers, giving details of fixations on the monitor and on the keyboard and the ScriptLog keystroke record (SLog). The x-axis shows the time, and each of the thin lines in the bottom tier represents one keystroke; therefore, the faster the typing, the shorter the distance between the lines. The larger the gap, the longer the “pause” between the keystrokes. According to Figure 1, this writer spent most of the pause between the two sentences looking at the monitor and only occasionally glanced at the keyboard. She practically never looked away from both the keyboard and monitor, in which cases both the monitor and the keyboard tiers would have been blank.

An initial interpretation of these data may, therefore, be that this writer spent the whole pause reading her text. However, looking at the monitor is not the same thing as actually reading the text that is shown on the monitor, which is clearly shown in Figure 2. The figure shows how the writer first gave the monitor a quick glance, then shifted to the keyboard for an equally short glance, then went back to the monitor again, and then looked at the keyboard for about 3 sec before finally starting to write. Once she had started to write, she looked at the monitor almost all the time. It turns out that this writer was a very skilled typist, with little need to look at the keyboard while she was writing. Since the data do not tell us whether she was really a touch typist (even if we could probably assume that she was), we call her a *monitor gazer*, as distinct from a *keyboard gazer* (Johansson, Wengelin, Johansson, & Holmqvist, in press).

The figure clearly shows that we cannot interpret all monitor gazing as reading according to a reading-research-based view of reading that assumes that the main function of reading is *reading for comprehension*. However, a writer, reading his or her own text, is probably not reading mainly for comprehension. He or she may, instead, be reading for planning and/or reading for evaluation and revision. On the other hand, a writer who composes with the reader in mind probably strives toward a coherent text and may read from the perspective of a prospective reader. To differentiate between reading according to the traditional reading-for-comprehension definition and other kinds of monitor gazing, we use a reading filter developed by Kollmorgen and Holmqvist (2009). The reading filter is not based on a fixed definition but can be trained to recognize different patterns of fixation-saccade sequences and so identify, with *reasonable accuracy*, when reading occurs. The main idea behind this approach is not only to identify the patterns within the




Figure 3. TimeLine screenshot showing gaze distribution including reading patterns and some of the descriptive statistics for the end of the first sentence and the beginning of the second sentence (23.895–49.9000 sec).

duration = 0.16 sec, $SD = 0.23$) and for a total of 0.64 sec (2.47% of the time). TimeLine also generates lists of starting and ending times for each kind of gaze event—for example, looking at the monitor or looking at the keyboard. At the moment, it does not generate the fixation durations for these events, but by merging these lists with the fixation output of SMI iView, they can easily be retrieved. For example, the mean fixation duration for the whole writing session was 389 msec ($SD = 320$), whereas the mean fixation time for the 25.97 sec shown in Figures 1 and 3 was 557 msec ($SD = 338.5$) and that for the 2.86 sec of reading was 260 msec ($SD = 147.6$). This information will be directly available from future versions of TimeLine.

We are now in the position to return to our earlier question: Was the pause time between the two sentences taken up entirely by the writer's reading what she had written? Analyses suggest that she read only for about 3 sec (the first 3) (21%) of her total pause time of 13.96 sec. The video reveals that for the rest of the time (almost 10 sec), her gaze first disappeared outside the screen. Then she gazed quickly at the two last words of the text, "ungdomars vardag" ("the daily life of young people"). This was not recognized as reading, since the reading filter had not

been trained to recognize single fixations as reading. After that, her gaze went to the right side of the screen, after which it went back to the last two words of the sentence. Then she looked up at the ScriptLog menus and browsed them for a little while, perhaps looking for a menu option that exists in Microsoft Word, but not in ScriptLog. After that, her gaze went to the left corner of the monitor and then again back to the last two words of the text. Then her gaze went to the upper right corner of the text. Finally, her gaze went down to the keyboard and back to the cursor, where she immediately started writing. Note that she returned to the last two words of the sentence three times. We do not know what happened here, but one possible explanation is that she used those two words as a trigger for the next sentence. On the other hand, she could have been contemplating revising the sentence that she had just written. Despite the case that this was not categorized as reading, we can assume that she was processing information somehow, and this raises the question about the usefulness of a traditional reading-for-comprehension definition when reading during composition is investigated. According to the reading filter, during her entire writing session of about 28.5 min, the writer read parts

of her text on 78 separate occasions. Her total pause time (with pauses defined as 2 or more sec of typing inactivity) was 669.90 sec, distributed across 82 pauses, and her total reading time was 366.74 sec. The mean reading time was 4.7 sec ($SD = 5.86$). However, if there is a pause during reading, the reading filter will report the reading preceding and succeeding the pause as two instances. We analyzed the activities associated with reading and found that in 37 of the found instances, reading actually succeeded another instance of reading. The others were distributed as follows: 14 instances followed immediately on the insertion of a letter, 5 succeeded a punctuation mark, 6 succeeded a backspace, and 9 succeeded a mouse movement. Moreover, it turned out that reading did not only occur in pauses. In 6 instances, reading took place during writing, at locations other than the point of inscription. Furthermore, in 7 instances, this writer appeared to read “backward”; that is, she demonstrated sequences of fixations and saccades that are typical for reading (and were picked up by the reading filter), but from right to left.

Implementation. ScriptLog is adapted for synchronization with the SMI iView X (HED + HT 50 Hz), which is a head-mounted eyetracker. In our setup, it is mounted on a bike helmet and uses a Polhemus magnetic headtracker, which gives the participants freedom to move both body and head during the experiment. Combined eye and head information allows quite precise pinpointing of fixation location across a large visual field and in a number of pre-specified planes that can vary in size, position, and orientation. The visual field is recorded using a scene camera, also head mounted. Eye and scene camera, as well as magnetic 6DOF sensor, are all mounted on a lightweight cycle helmet that participants wear while writing. This system outputs both an MPEG-2 video of the visual field, with overlaid gaze marker and time stamp, and a data file with gaze coordinates in each of the planes. The video can also incorporate sound, which would permit collection of writers’ introspections or think-aloud protocols generated as they write.

The data file reports, alongside raw head and eye location data, the plane that is currently fixated and fixation location within that plane, reported with reference to the plane’s coordinate system. Plane coordinates are identified before data collection starts. In our research, the screen and keyboard are defined as separate planes.

ScriptLog, the keystroke-logging program, is developed for the PC. ScriptLog and eyetracking output files are synchronized by ScriptLog’s sending a start signal to the eyetracker when a ScriptLog recording is started.

TimeLine, the program used to merge eye and keystroke data, is a time-stamped data sheet coded in Python. It is flexible in that it can provide graphics representing a range of different user-defined activities, in addition to those discussed above. For example, an additional “writing” tier can be added that shows activity based on the analyst’s definition of writing versus pausing.

As was mentioned above, the reading filter is based on a hidden Markov model. The technical details of the reading filter are more thoroughly described in Kollmorgen and Holmqvist (2009).

EyeWrite

EyeWrite (Simpson & Torrance, 2007), like ScriptLog + TimeLine, aims to integrate keystroke logging with eye movement records. Its approach to achieving this is rather different, and as a result, it is able to generate information about the location of fixations relative to the text and independently of screen coordinates. By this, we mean that as well as being able to say that, at a particular time, the writer is fixating a screen location with particular x,y coordinates, EyeWrite also determines that at the point in time when that fixation was made, those coordinates placed the fixation in, for example, the middle of the word *environment* and that this word was in the third sentence of the text, 22 words from the start and 107 words and 5 sentences distant from the last-inserted character. In this section, we will provide a brief overview of the system; explore why the distinction between text- and screen-relative fixation locations is necessary; explain, again briefly, how EyeWrite generates text-relative fixation locations; and illustrate how the resulting information might be useful in exploring what happens when writers pause.

EyeWrite comprises a simple editor and an analysis program. The text editor is basic, with all cursor movements made using cursor (arrow) keys and all deletes made by backspacing. Behavior of text on the screen is as might be expected in a word-processing program, with line wrapping and vertical scrolling. The editor does not permit cut and paste. Eye movements are recorded using a head-mounted eyetracker that has sufficient accuracy to permit, at least, identification of the word that the writer is fixating. The editor program logs both keystrokes and eye movements. The analysis program interprets the combined keystroke and eye movement log file and generates text-relative fixation location information that can be played back or, for more systematic analysis, output as a data file. In playback mode, this program displays the text as it develops on the screen with a graphic representing the cursor and, when fixations are within screen coordinates, a graphic indicating the location of the current fixation. Writing sessions can be played back in (roughly) real time or stepped through by keystroke and/or fixation events. Stepping by keystroke, for example, moves the display to the point when the next keystroke was made and indicates the location of the writer’s fixation at that point in time. This is useful for fine-grained analysis of particular sets of events during production.

It is reasonably straightforward to create a playback display of this kind by combining a screen-capture movie with screen coordinates supplied by the eyetracker. Facility for doing this is a standard part of most eyetracking software. However, because of the need to generate text-relative fixation location information, EyeWrite adopts a different approach. The problem that needs to be overcome is best understood by way of an example. A writer creating text about global warming might use the word *environment* in the third sentence of the text. At the time this is created, the screen location of the center of this word (in pixels, with $1,024 \times 768$ screen resolution) might be 430,862. Fixations at 430,862 will, therefore, suggest that the writer is reading this word. However, if, some time later, the writer edits

the second sentence of the text, adding several words, the screen location of *environment* will move, with the center of the word at, maybe, 390,476. Fixations at 430,862 will now no longer represent reading *environment*, but fixations at 390,476 will. Even without editing, vertical scrolling and line wrapping means that the same text can occupy any number of different screen locations. EyeWrite accommodates the unpredictably dynamic relationship between text and screen and is, therefore, able to specify fixation locations relative to the emerging text.

EyeWrite uses the following strategy to convert fixation locations specified as screen coordinates to text-relative location variables. The editor uses a nonproportional font (a font in which all characters occupy the same screen space). The display can then be interpreted as a grid, with each character occupying one cell and cell size determined by font size, text line spacing, and screen resolution. This grid is static, and each cell can, therefore, be specified by its screen coordinates. Keystroke information allows creation, for any particular point in time during composition, of an array that specifies the character, if any, that occupies each grid square. It is then straightforward to combine this information with the screen coordinates of the current fixation, provided by the eyetracker. This then makes it possible to determine, within the limits of the accuracy of the eyetracker, which character is being fixated.

Therefore, after analysis, and if the fixation is in part of the screen that contains text, EyeWrite is able to provide information of the form, *at time t the writer was fixating character c which was n1 characters from the start of the text, n2 characters from the current end of the text, and n3 characters from the current cursor location*. By identifying relevant orthographic features (space characters, punctuation marks, paragraph markers), it is also possible to specify these distances in numbers of words, sentences, and paragraphs and also to specify the length of the currently fixated word, sentence, and paragraph. Output also includes the exact text of the fixated word and the text to either side of the point of fixation up to a user-specified number of characters. (It is, of course, also possible to extract this information manually from a playback of the writing session with a fixation location overlay. However, this is prohibitively time consuming. A 30-min writing session that we recently analyzed involved a total of 4,759 fixations, of which 1,919 [40%] were associated with fixations within the text already written.)

EyeWrite outputs this text-relative information as a data file. Each record in this file represents either a fixation or a keystroke. Full details of the information provided by the output file in the current version of EyeWrite can be found in Figure 4. For fixations, the output reports start time, duration, location relative to the screen, and, if the writer is looking at the portion of the screen that contains text, location relative to the text. It also records information about the lengths of the currently fixated word, sentence, and paragraph. For keystrokes, the output file reports key-press times (and screen update times, when the character of the key appears on the screen), the key that is pressed, and initial cursor location (relative to screen and relative to text). Records for keystrokes also report the nature of the

eye activity and the time that the key was pressed (blink, saccade, fixations). For fixations, the record then goes on to report information about that fixation (location, length of fixated word, and so forth). These form the core of the measures provided by EyeWrite. In order to answer questions about what happened prior to or following a particular kind of event (e.g., where do writers tend to look when they reach the end of a sentence?), it is also necessary to extract, for each event, information about preceding and following events. Exactly what needs to be known depends on the purpose of the analysis. Currently, some of these previous-event and next-event variables are computed by the analysis component of EyeWrite, and some by secondary analysis of the output file data within SPSS, making extensive use of the *lag* and *lead* commands. It would be straightforward to incorporate these secondary analyses into future versions of EyeWrite.

The sections that follow give some examples of the kinds of findings that EyeWrite makes possible and provide some more detailed information about how EyeWrite is currently implemented.

Illustrative findings. In this section, we will present some early findings whose aim was to explore eye activity during pauses in production that occur while writers are composing argumentative text. We present these here just to illustrate the kinds of analyses of writer behavior that are made possible by EyeWrite and will, therefore, attempt very little interpretation. Data are from eight writers who composed short persuasive essays arguing the environmental case for increasing gasoline prices. All of the writers were undergraduate psychology students, and they wrote for between 25 and 30 min. In the analyses that follow, fixations of 100 msec or less are ignored, although it would be straightforward to implement more sophisticated methods for handling short fixations, such as merging with adjacent longer fixations within the same word. We will first briefly discuss where the writers tended to look while actually inscribing text (while making character keystrokes). We will then explore, again briefly, where the writers tended to look during pauses in inscription. Finally, we will explore the extent to which in-text fixations are associated with reading, rather than, for example, with scanning for errors.

The writers that we describe were neither touch typists nor monitor gazers (in the terminology of Johansson et al., in press), and so all tended to fixate the keyboard while typing. The writers inscribed a mean of 2,397 characters ($SD = 698$) during their writing sessions, and only a very small proportion of these keystrokes were associated with fixations on or near the inserted character. On average, only 3.4% of character inserts were accompanied by fixations on the same line of text as and within 8 characters of the inserted character. This value ranged from 0.1% to 6.8% for different writers. Of the character keystrokes (ranging from 2.3% to 23%), 13.4% were associated with fixations within the portion of the screen that contained text already written. Further analysis suggested, however, that these fixations were not typically part of sequences that might be interpreted as reading but were just single, and possibly untargeted, glances at the text. We are not clear what function, if any, these might serve, but the con-

Session management
- trial label
- title text
Event
- event id (unique for each event)
- event type: fixation, insertion, deletion, cursor move
- event type for immediately previous event
- state of eye: fixation, saccade, blink
- event time (ms)
About the current fixation
- current fixation start time, end time, and duration
- current fixation location in pixels and row, column
- location from start of text at fixation start and at fixation end, in characters, words, sentences, and paragraphs
- length of fixated word / sentence / paragraph at fixation start and at fixation end
- text of fixated word
- text either side of fixation (up to a user specified maximum number of characters)
About the cursor
- initial and final cursor positions in pixels and row, column
- location of cursor in characters, words, sentences, paragraphs from start of text
- length of word, sentence, paragraph in which cursor occurs
- characters immediately preceding and following the cursor in the text (up to a user-specified maximum)
Text variables
- current length of text in characters, words, sentences, and paragraphs
- last characters appearing in text (up to a user-specified maximum)
- character inserts
- character
- time at which keyed character actually appears on screen
For previous and next insert events
- character
- event time
- location in pixels and row, column
- time elapsed since insertion of previous character
- does inserted character immediately follow previous character in the text?
- location of insert in characters, words, sentences, paragraphs from start of text

Figure 4. Summary of variables generated by EyeWrite.

siderable cross-writer variation suggests that, while typing, different writers used information from the screen in different ways. Fixations during the remaining character keystrokes ($M = 83\%$, $SD = 8.8\%$) were almost entirely off screen, with very few fixations on the blank portion of the screen below the text ($M = 3.4\%$, $SD = 2.0\%$). The eyetracker arrangement, in contrast to that used by Wengelin and coworkers, meant that for most of the participants, coordinates for fixations on the keyboard could not be recorded. However, we think it is reasonable to assume that a substantial majority of these fixations were associated with the writers' hunting for keys.

Taking the writing task as a whole, an average of 40.5% ($SD = 10.2\%$) of total time on task was associated with in-text fixations. As we have indicated, it was rare for the writers in our sample to look at or near the word that was currently being typed. The bulk of these fixations, therefore, appeared to be associated with deliberate monitoring of text already written.

For present purposes, pauses were defined as delays of greater than 2 sec between keystrokes (character inserts and deletions, but not cursor moves). Consistent with previous findings (Wengelin, 2002, 2006), pausing, defined

in this way, was much more common at sentence boundaries and paragraph boundaries than between characters and words (mean percentage of boundary locations associated with pausing were, for character, word, sentence, and paragraph boundaries, 2%, 6%, 31%, and 45%, respectively). Also consistent with previous findings (Chanquoy, Foulin, & Fayol, 1996; Spelman Miller, 2000), and as Table 1 indicates, pause lengths tended to be greater for pauses occurring at sentence and paragraph boundaries.

Figure 5 shows fixations occurring during pauses at different kinds of text location, by distance in number of words from pause location. For example, the first point on the word boundary curve indicates that when the writers paused at word boundaries, a mean of 27% of fixations lay within the 5 words prior to the word boundary where the pause occurred, the second point indicates that a mean of 43% of fixations lay within the 10 words prior to this boundary, and so forth. Only recursive fixations are shown—fixations that occur in text that appeared to the left of or above the fixation location. Distances of greater than 110 words were rare. Note that because the texts rarely comprised more than two or three paragraphs, findings related to paragraph boundaries need to be inter-

Table 1
Mean Frequencies and Percentages of Pauses Occurring at Different Locations in the Text, by Length of Pause

Pause Length (sec)	Boundary at Which Pause Occurred							
	Character		Word		Sentence		Paragraph	
	<i>M</i>	%	<i>M</i>	%	<i>M</i>	%	<i>M</i>	%
2-3	8.8	35	14.4	33	1.6	13	0.3	25
3-5	5.1	19	12.9	29	3.0	26	0.6	22
5-10	5.1	19	10.5	23	2.6	23	0.5	42
10-15	2.5	10	3.5	9	1.3	11	0	0
15-20	1.5	6	1.5	3	1.3	10	0	0
20-40	2.1	8	1.1	3	1.5	12	0.1	6
40-60	0.9	4	0.1	0	0.5	4	0.3	6
>60	0	0	0	0	0.3	2	0	0

preted with caution. The figure suggests that the extended pause times at sentence and paragraph boundaries were, as might be expected, associated with reading back to a greater depth within the text. Presenting pause-location-to-fixation distances in number of words is, to a certain extent, problematic because the depth to which a writer can read back within a text is capped by the total number of words between the pause location and the start of the text (which will often be the same as the total number of words so far written). Figure 6 controls for this by presenting distances relative to the distance between pause location and start of text. This shows a very similar pattern of results.

As was discussed above, only a proportion of fixations within the text already written are likely to be associated with “normal” (linear, left-to-right, multiple-word) reading. ScriptLog+TimeLine necessarily adopts a heuristic approach to identifying reading episodes. The text-relative fixation locations generated by EyeWrite make an algorithmic (rule-based) approach possible. We identified reading activity in two stages. First, we identified situations in which the writers made three or more consecutive fixations within the text already written (on the portion of screen containing text and at least eight characters from the last-inserted character). Of total writing time, 28.9% was associated with in-text fixation sequences (as com-

pared with 80.9% when the same writers performed a read-and-edit task on text that we provided). Then, within these in-text-fixation episodes, we identified eye activity as associated with reading when all three of the following conditions were met: (1) The writer fixated three or more words consecutively, (2) saccades between these words involved forward moves within the text (i.e., the character index for each subsequent fixation was greater than those for the previous fixations), and (3) fixated words were separated by not more than three other words and not more than 25 characters. These words and characters could be distributed across two lines of text (i.e., the analysis was blind to line breaks). Within these sequences, we permitted multiple fixations within the same word, with no constraint on the direction of within-word saccades. On the basis of this definition, a mean of 26.1% ($SD = 8.5\%$) of in-text fixation sequences were associated with normal reading. A large proportion of the remainder ($M = 37.4\%$, $SD = 7.9\%$) were associated with series of backward and forward fixations within the same sentence. This activity may be driven by a need to comprehend text, but we suspect that, in the context of writing, it is more likely to be associated with error detection or, perhaps, planning how a sentence might be improved. Exploring keystroke activity immediately following these “scanning” episodes would provide clues as to their function.

Discussion of appropriateness of different ways of operationalizing reading and more detailed analysis and interpretation of our eye movement data are beyond the present scope but are, of course, needed. In particular, we do not want to make any strong claims for the particular definition of reading that was used in the analysis above. However, the findings that we have presented serve, we believe, to illustrate ways in which text-relative fixation location information can be used to say interesting things about where writers look. The algorithmic approach to identifying reading that this information facilitates has certain advantages over the probabilistic approach adopted by the reading filter associated with ScriptLog+TimeLine: Each episode is identified as reading only if it meets tightly

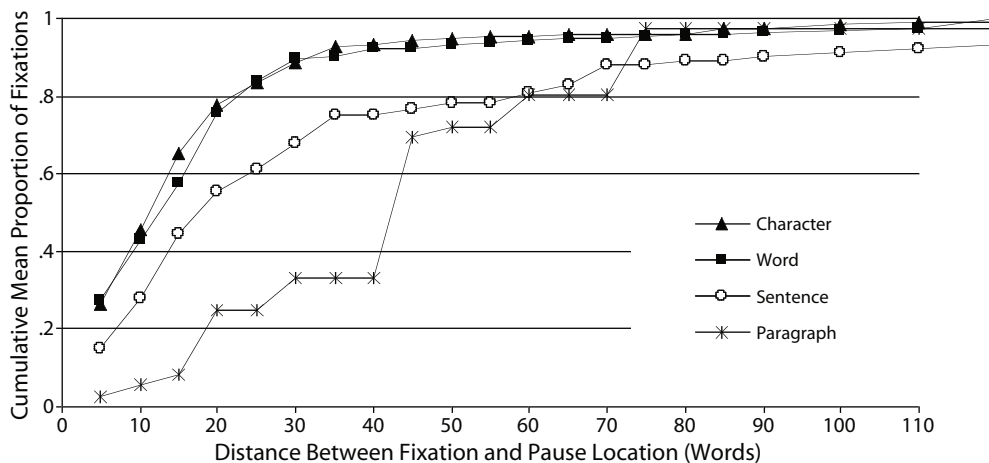


Figure 5. Fixations occurring during pauses at different kinds of text location, by distance in number of words from pause location.

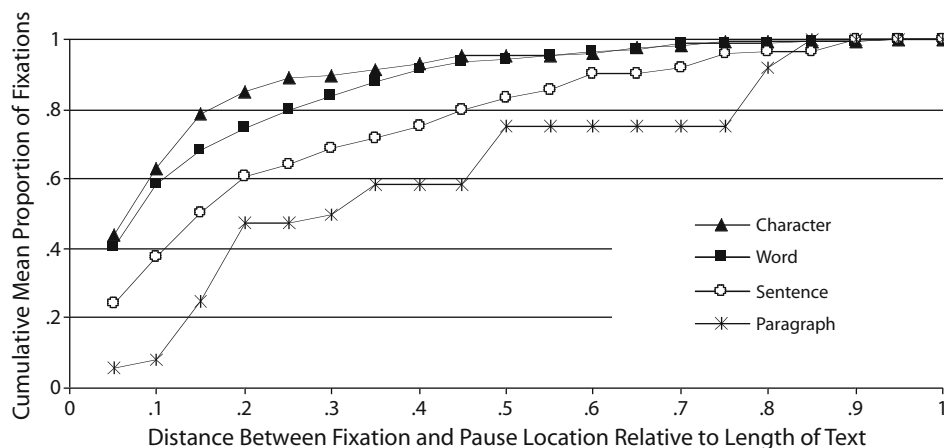


Figure 6. Same data as in Figure 5, but controlling for the number of words between the start of the text and the pause location. Fixations near the start of the existing text will have distance values approaching 1, and fixations near the pause location will have distance values approaching 0.

defined criteria, and these criteria can be theoretically driven. However, given both fuzziness in the definitions of reading and possible error in determining fixation location, ScriptLog+TimeLine's pattern-matching method is likely to yield smaller false negatives (episodes in which the writer is, in fact, reading but that are not identified as such). It is important to note, however, that writers—both monitor gazers and keyboard gazers—look at the screen and, presumably, process information to a much larger extent than what can be captured by a traditional reading-comprehension definition.

Implementation. Both editor and analysis components of EyeWrite are programmed in Python and are implemented within SR Research's Experiment Builder environment (Version 1.4 and later). The editor works as an object in an Experiment Builder project, with the option of varying several parameters, including the size and location of the editor window, font size, and line spacing. It also makes possible more than one writing task within a session, with data for all tasks collected within a single file. The editor window can start blank or can be preloaded with text. This can then be edited and developed by the writer, with full recording of keystrokes and fixations in both new and preloaded text. These features combine to give considerable flexibility in the kinds of tasks that can be presented to participants. The participants in the research presented above were simply presented with an editor window that filled most of the screen and a small panel at the top of the screen presenting the writing task. Fixations within this panel were also recorded. We also explored writers' editing behavior by preloading a text that contained several errors and asking them to find and correct these. It would also be possible to use a much smaller editor window containing a sentence stem that participants then complete or edit and complete (e.g., in an extension of the paradigm reported by Leijten, 2007, pp. 99–109). Various other paradigms—incorporating, for example, visual stimuli—could also be implemented. A copy of Experiment Builder is required to make changes

to the structure of experiments using EyeWrite. However, once these are developed, deployed versions can run on any suitable computer.

The EyeWrite editor works with SR Research's Eye-Link II eyetracker. This is a head-mounted video-based system that incorporates basic headtracking. It samples at 250 Hz in combined pupil and corneal reflection mode (the method that we have found most effective) and gives accuracy to around half a degree of visual angle. The system allows the researcher to visually monitor for calibration drift during the writing session. When this occurs, it can normally be remedied using an inline drift assessment and correction method. The writer is asked to fixate the cursor, and the experimenter then presses a key that initiates assessment and, if necessary, correction of any drift. This procedure appears minimally intrusive and, in our experience, is not required very frequently.

Our current setup, which seems to give good results, has the participant seated approximately 60 cm from a large (21-in.) screen, with his or her eyes aligned at a point about halfway down the display. The display has a resolution of $1,024 \times 768$ pixels and a 120-Hz refresh rate. The title panel and editor take up nearly all of the screen, with a small margin at the bottom and sides to avoid on-text fixations being distracted by the screen edges. Font, font size, and line spacing are defined when the text editor object is called and can potentially be varied between trials in one session or across writing sessions. We have found that we get good accuracy using 12-point Courier New font with 30-point spacing between lines (measured between line centers).

The analysis component is also implemented in Python. This makes it relatively straightforward for researchers with limited programming experience to change or add to the variables that it outputs. Data are output as tab-delimited text files that we currently manipulate quite extensively within SPSS. It would be straightforward to code the most useful of these subsequent manipulations directly into future versions of EyeWrite's analysis component.

Discussion

We have, we hope, demonstrated that eyetracking, in combination with keystroke logging, can add to our understanding of the cognitive processes that underlie written production. Traditionally, writing researchers have followed research in speech production in attributing pausing to the writer's deciding what to say next (e.g., Flower & Hayes, 1981). However, more recent research has suggested a much more complex picture (van den Bergh & Rijlaarsdam, 1996, 1999). Consistent with this, the preliminary analyses that we report in this article, made possible by combined eyetracking and keystroke-logging methods, suggest that pausing is often accompanied by directed eye movements within the text produced so far. Some of this eye activity seems to be similar to that which would be observed if a participant were reading an experimenter-provided text for comprehension. However, our preliminary analyses suggest other recurring patterns, including repeated backward and forward saccades within the sentence that is being composed, multiple well-spaced glances across several sentences or paragraphs, and a "reverse-reading" pattern involving series of fixations separated by two- or three-word right-to-left saccades. More than one of these or other patterns may occur within a single pause. They seem interpretable, in the sense that it is possible to hypothesize plausible cognitive functions that they might fulfill. Testing these hypotheses will, of course, require considerable further research in which, we suggest, combined eyetracking and keystroke-logging methods will need to play a central role.

Preliminary analyses also suggest that at least some writers fixate, and possibly process, text produced so far at some distance for the word that is currently being typed. This is consistent with findings that in speech planning, the next clause or sentence occurs while the previous one is still being spoken (Power, 1986). There is some evidence that this is also true for written production (Chanquoy, Foulin, & Fayol, 1990) and provides an explanation for why writers only sometimes pause at clause and sentence boundaries. It is also possible that fixations concurrent with but distant from the currently typed word may be associated with locating errors in the preceding sentence that the writer is aware of but delays correcting. Again, detailed exploration of eye activity that occurs during these parallel typing and planning/reviewing phases will require combined analysis of keystroke and eye movement data.

The strength of the particular systems—ScriptLog + TimeLine and EyeWrite—that we describe in this article is that they permit conversion of the large quantities of data generated by both keystroke logging and eyetracking into psychologically and linguistically interpretable variables. They permit interpretation of these data without researchers having to resort to manual coding from video, which, for a usefully sized sample of writers performing meaningfully extensive writing tasks, is likely to be prohibitively time consuming. The most significant difference between the two systems, in their current instantiations, is in terms of analysis. ScriptLog + TimeLine implements a heuristic, pattern-matching approach to identifying particular kinds of reading and reading-related

behavior within the *xy* eye movement trace. EyeWrite uses information about the editor environment to directly determine fixation location within the emerging text. Both approaches have their merits, and researchers might find it useful to implement both, using, first, heuristic methods to identify broad patterns of eye movement (with risk of false positives) and then more fine-grained algorithmic methods (with risk of false negatives). The systems currently also differ in the functionality of the editor environment in which writers compose their text. Most notably, the EyeWrite editor does not implement mouse control of the cursor or cut-and-paste actions, which, in some research contexts, might be a limitation. There is, however, no in-principle reason why these functions could not be implemented within the EyeWrite editor, or disabled within ScriptLog, to meet varying research goals.

Our intention in this article has not been, therefore, to present our particular systems as ready-made solutions to other researchers' problems. We are also not suggesting that combined keystroke and eyetracking methods are a replacement for more established ways of exploring writing processes. Even if these methods are embraced by writing researchers, there will remain an important role for collecting writers' verbal reports (both online and retrospective) and for cognitive-linguistic analysis of completed texts (e.g., Van Wijk & Sanders, 1999). We can envisage usefully combining these methods with the methods reported in this article. Both of the systems presented here allow playback of the writing sessions showing both typing and eye movements—either directly through the system (EyeWrite) or by the video generated by the eyetracking system (ScriptLog + TimeLine). These could easily be used to prompt writers' retrospective self-reports (see, e.g., Leijten, 2007; Lindgren & Sullivan, 2003). Both systems also permit audio recording of concurrent protocol, and it would be straightforward to synchronize playback of these with eyetracking and keystroke time stamps, although there are reactivity issues here: The effects of thinking aloud on pause and eye movement patterns in this context are not yet understood. Interpreting eye movements in conjunction with sophisticated text-linguistic analysis of the text's rhetorical structure is likely to prove a promising research avenue and is a logical extension of work that relates pause location to text structure (Schilperoord, 1996; Schilperoord & Sanders, 1999). Work is currently under way, as part of the development of the Inputlog suite of programs (Leijten & Van Waes, 2006), to develop a tool and standardized logfile format that would permit manual markup of keystroke files for, for example, specific linguistic features. This would extend the kinds of analysis reported in, for example, Figures 5 and 6 to features that are not orthographically marked (clause boundaries, theme change boundaries, and so forth).

We are, as will be apparent from this article, at a very early stage in research on eye movement during writing, and it is not yet clear to what extent the large body of findings relating to eye movement during reading and, more generally, to processes that underlie text comprehension can be generalized to writers reading text they have just written. Writers reading their own text have informational

needs very different from those of readers who have to extract meaning from texts provided by the researcher. This gives rise to questions relating both to eye control and, more generally, to exactly what is being processed when a writer "reads" his or her own writing. The extent to which text can be processed parafoveally or the extent to which words can be skipped while comprehension is maintained may be different when the reader is reading words that he or she has very recently produced. Writing also suggests a range of functions for reading and reading-like behavior (error checking, cuing new content, even word counting) that go beyond those typically required when participants are reading researcher-provided text. However, we believe that there is now sufficient depth in existing theory of the cognitive processes associated with writing, and sufficient understanding of how eye guidance works in other contexts, to provide a framework in which data from combined eye tracking and keystroke logging can be interpreted. The systems that we present in this article suggest two different and complementary ways in which this kind of data can be collected and managed.

AUTHOR NOTE

EyeWrite development was undertaken while M.T. was Senior Lecturer in the Psychology Department at Staffordshire University. Our thanks to Rachel Hadzic for help with initial testing. The Swedish research was funded by the Swedish Research Council (Vetenskapsrådet 2004-2687). ScriptLog was developed by Sven Strömquist and programmed by Henrik Karlsson. Our thanks to Sven for including connections to the SMI eyetracker in ScriptLog and for important input on the design of the study. TimeLine was programmed by Johan Dahl. Thanks to Jana Holsanova for contributing with her method of multimodal time-coded score sheets. Finally, thanks to Bodil Andersson and Sylvia Tufvesson for participant recruitment and data collection. Correspondence concerning this article should be addressed to Å. Wengelin, Centre for Languages and Literature, Lund University, Box 201, 221 00 Lund, Sweden (e-mail: asa.wengelin@ling.lu.se) or M. Torrance, Division of Psychology, School of Social Sciences, Nottingham Trent University, Burton Street, Nottingham NG1 4BU, England (e-mail: mark.torrance@ntu.ac.uk).

REFERENCES

- ALAMARGOT, D., & CHANQUOY, L. (2001). *Studies in writing 9: Through the models of writing*. Dordrecht: Kluwer.
- ALAMARGOT, D., CHESNET, D., DANSAC, C., & ROS, C. (2006). Eye and pen: A new device for studying reading during writing. *Behavior Research Methods, Instruments, & Computers*, *38*, 287-299.
- ALLAL, L., CHANQUOY, L., & LARGY, P. (EDS.) (2004). *Studies in writing 13: Revision. Cognitive and instructional processes*. Dordrecht: Kluwer.
- ALLUM, P. H., & WHEELDON, L. R. (2007). Planning scope in spoken sentence production: The role of grammatical units. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *33*, 791-810.
- ANDERSSON, B., DAHL, J., HOLMQUIST, K., HOLSANOVA, J., JOHANSSON, V., KARLSSON, H., ET AL. (2006). Combining keystroke logging with eye tracking. In L. Van Waes, M. Leiten, & C. M. Neuwirth (Eds.), *Writing and digital media* (pp. 166-172). Amsterdam: Elsevier.
- BERMAN, R., & VERHOEVEN, L. (2002). Cross-linguistic perspectives on the development of text-production abilities: Speech and writing. *Written Language & Literacy*, *5*, 1-43.
- BISAILLON, J. (2007). Professional editing strategies used by six editors. *Written Communication*, *24*, 295-322.
- BREETVELT, I., VAN DEN BERGH, H., & RIJLAARSDAM, G. (1996). Re-reading and generating and their relation to text quality: An application of multilevel analysis on writing process data. In G. Rijlaarsdam, H. van den Bergh, & M. Couzijn (Eds.), *Theories, models and methodology in writing research* (pp. 10-20). Amsterdam: Amsterdam University Press.
- BRITTON, J. (1982). Shaping at the point of utterance. In G. M. Pradl (Ed.), *Prospect and retrospect: Selected essays of James Britton* (pp. 139-145). London: Boynton/Cook.
- BUTSCH, R. (1932). Eye movements and the eye-hand-span in typewriting. *Journal of Educational Psychology*, *23*, 104-121.
- CHANQUOY, L., FOULIN, J.-N., & FAYOL, M. (1990). Temporal management of short text writing by children and adults. *Cahiers de Psychologie Cognitive*, *10*, 513-540.
- CHANQUOY, L., FOULIN, J.-N., & FAYOL, M. (1996). Writing in adults: A real-time approach. In G. Rijlaarsdam, H. van den Bergh, & M. Couzijn (Eds.), *Theories, models and methodology in writing research* (pp. 36-44). Amsterdam: Amsterdam University Press.
- CHESNET, D., & ALAMARGOT, D. (2006). *Eye and Pen* (Version 1.0.). Poitiers: Octares.
- CLARK, H. H., & FOX TREE, J. E. (2002). Using *uh* and *um* in spontaneous speaking. *Cognition*, *84*, 73-111.
- CRAIN-THORESON, C., LIPPMAN, M. Z., & MCCLENDON-MAGNUSON, D. (1997). Windows on comprehension: Reading processes as revealed by two think-aloud procedures. *Journal of Educational Psychology*, *89*, 579-591.
- DEGENHARDT, M. (2006). CAMTASIA and CATMOVIE: Two digital tools for observing, documenting and analysing writing processes of university students. In L. van Waes, M. Leijten, & C. M. Neuwirth (Eds.), *Writing and digital media* (pp. 180-188). Amsterdam: Elsevier.
- EKLUND, R. (2004). Disfluency in Swedish human-human and human-machine travel booking dialogues. *Linköping Studies in Science & Technology* (Dissertation No. 882). University of Linköping, Tekniska Högskolan.
- FLOWER, L. S., & HAYES, J. R. (1981). The pregnant pause: An enquiry into the nature of planning. *Research in the Teaching of English*, *15*, 229-243.
- FOULIN, J.-N. (1998). To what extent does pause location predict pause duration in adults' and children's writing? *Cahiers de Psychologie Cognitive*, *17*, 601-620.
- GOLDMAN-EISLER, F. (1968). *Psycholinguistics: Experiments in spontaneous speech*. London: Academic Press.
- HACKER, D. J., PLUMB, C., BUTTERFIELD, E. C., QUATHAMER, D., & HEINEKEN, E. (1994). Text revision: Detection and correction of errors. *Journal of Educational Psychology*, *86*, 65-78.
- HAYES, J. R. (1996). A new framework for understanding cognition and affect in writing. In C. M. Levy & S. Ransdell (Eds.), *The science of writing: Theories, methods, individual differences, and applications* (pp. 1-26). Mahwah, NJ: Erlbaum.
- HAYES, J. R., & FLOWER, L. S. (1980). Identifying the organization of writing processes. In L. Gregg & E. R. Steinberg (Eds.), *Cognitive processes in writing* (pp. 3-30). Hillsdale, NJ: Erlbaum.
- HOLMQUIST, K., HOLSANOVA, J., JOHANSSON, V., & STRÖMQUIST, S. (2004). Perceiving and producing the frog story. In D. Ravid & H. B.-Z. Shyldkrot (Eds.), *Perspectives on language and language development: Essays in honor of Ruth A. Berman* (pp. 293-306). Amsterdam: Kluwer.
- HOLSANOVA, J. (2001). *Picture viewing and picture description* (Lund University Cognitive Studies No. 83). Lund: Lund University, Department of Philosophy.
- INHOFF, A., & GORDON, A. (1998). Eye movements and eye-hand coordination during typing. *Psychological Science*, *6*, 153-157.
- JANSEN, D., VAN WAES, L., & VAN DEN BERGH, H. (1996). Effects of thinking aloud on writing processes. In M. Levy & S. Ransdell (Eds.), *The science of writing: Theories, methods, individual differences and applications* (pp. 233-250). Mahwah, NJ: Erlbaum.
- JOHANSSON, R., WENGELIN, Å., JOHANSSON, V., & HOLMQUIST, K. (in press). Looking at the keyboard or the monitor: Relationship with text production processes. *Reading & Writing*.
- JUST, P. A., & CARPENTER, A. (1980). A theory of reading: From eye fixations to comprehension. *Psychological Review*, *87*, 329-354.
- JUST, P. A., & CARPENTER, A. (2002). A capacity theory of comprehension: Individual differences in working memory. In T. A. Polk & C. M. Seifert (Eds.), *Cognitive modeling* (pp. 131-180). Cambridge, MA: MIT Press.
- KAAKINEN, J. K., & HYÖNÄ, J. (2008). Perspective-driven text comprehension. *Applied Cognitive Psychology*, *22*, 319-334.

- KELLOGG, R. T. (1987). Effects of topic knowledge on the allocation of processing time and cognitive effort to writing processes. *Memory & Cognition*, **15**, 256-266.
- KINTSCH, W., & RAWSON, K. A. (2005). Comprehension. In M. Snowling & C. Hulme (Eds.), *The science of reading: A handbook* (pp. 209-226). Oxford: Blackwell.
- KINTSCH, W., & VAN DIJK, T. A. (1978). Toward a model of text comprehension and production. *Psychological Review*, **85**, 363-394.
- KOLLMORGEN, S., & HOLMQVIST, K. (2009). *Automatically detecting reading in eye tracking data*. (Lund University Cognitive Studies, No. 144). Lund: Lund University, Department of Philosophy.
- LEIJTEN, M. (2007). *Writing and speech recognition: Observing error correction strategies of professional writers*. Utrecht: LOT.
- LEIJTEN, M., & VAN WAES, L. (2006). Inputlog: New perspectives on the logging of on-line writing. In K. P. H. Sullivan & E. Lindgren (Eds.), *Computer keystroke logging and writing: Methods and applications* (pp. 73-94). Amsterdam: Elsevier.
- LEVELT, W. J. M. (1989). *Speaking: From intention to articulation*. Cambridge, MA: MIT Press.
- LEVY, B. A. (1983). Proofreading familiar text: Constraints on visual processing. *Memory & Cognition*, **11**, 1-12.
- LEVY, C. M. (1994). FauxWord and FastFauxWord: Windows-based tools for capturing and analyzing the process of written composition. In H. Van den Bergh, T. van der Geest, D. Janssen, G. Rijlaarsdam, & S. Sengers (Eds.), *EARLI/ECWC Writing Conference abstracts*. Utrecht: Utrecht University Press.
- LEVY, C. M., & RANSDALL, S. E. (1994). Computer-aided protocol analysis of writing processes. *Behavior Research Methods, Instruments, & Computers*, **26**, 219-223.
- LINDGREN, E., & SULLIVAN, K. (2003). Stimulated recall as a trigger for increasing noticing and language awareness in the L2 writing classroom: A case study of two young female writers. *Language Awareness*, **12**, 172-186.
- LIVERSEDGE, S. P., & FINDLAY, J. M. (2000). Saccadic eye movements and cognition. *Trends in Cognitive Sciences*, **4**, 6-14.
- LORCH, R., VAN DEN BROEK, P., & PRESSLEY, M. (1997). Understanding reading comprehension: Current and future contributions of cognitive science. Commentary. *Contemporary Educational Psychology*, **22**, 213-259.
- MATSUHASHI, A. (1987a). Revising the plan and altering the text. In A. Matsuhashi (Ed.), *Writing in real time: Modelling production processes* (pp. 224-253). Norwood, NJ: Ablex.
- MATSUHASHI, A. (Ed.) (1987b). *Writing in real time: Modelling production processes*. Norwood, NJ: Ablex.
- MYHILL, D., & JONES, S. (2007). More than just error correction: Students' perspectives on their revision processes during writing. *Written Communication*, **24**, 323-343.
- OLIVE, T., & LEVY, C. M. (Eds.) (2003). *Contemporary tools and techniques for studying writing*. Dordrecht: Kluwer.
- OLIVE, T., & PIOLAT, A. (2002). Suppressing visual feedback in written composition: Effects on processing demands and coordination of the writing processes. *International Journal of Psychology*, **37**, 209-218.
- PIANKO, S. (1979). The description of the composing processes of college freshmen writers. *Research in the Teaching of English*, **13**, 5-22.
- PILOTTI, M., MAXWELL, K., & CHODOROW, M. (2006). Does the effect of familiarity on proofreading change with encoding task and time? *Journal of General Psychology*, **133**, 287-299.
- POSTMA, A. (2000). Detection of errors during speech production: A review of speech monitoring models. *Cognition*, **77**, 97-131.
- POWER, M. J. (1986). A technique for measuring processing load during speech production. *Journal of Psycholinguistic Research*, **15**, 371-382.
- RABINER, L. (1989). A tutorial on hidden Markov models and selected applications in speech recognition. In *Proceedings of the IEEE*, **77**, 257-286.
- RAPP, B., & GOLDRICK, M. (2000). Discreteness and interactivity in spoken word production. *Psychological Review*, **107**, 460-499.
- RAUPACH, M. (1980). Temporal variables in first and second language speech production. In H. W. Dechert & M. Raupach (Eds.), *Temporal variables in speech: Studies in honour of Frieda Goldman-Eisler* (pp. 263-270). The Hague: Mouton.
- RAYNER, K. (1998). Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin*, **124**, 372-422.
- REICHLE, E. D., RAYNER, K., & POLLATSEK, A. (2003). The E-Z Reader model of eye-movement control in reading: Comparisons to other models. *Behavioral & Brain Sciences*, **26**, 445-476.
- SCHILPEROORD, J. (1996). *It's about time: Temporal aspects of cognitive processes in text production*. Amsterdam: Rodopi.
- SCHILPEROORD, J. (2003). On the cognitive status of pauses in discourse production. In T. Olive & C. M. Levy (Eds.), *Contemporary tools and techniques for studying writing* (pp. 61-88). Dordrecht: Kluwer.
- SCHILPEROORD, J., & SANDERS, T. (1999). How hierarchical text structure affects retrieval processes: Implications of pause and text analysis. In M. Torrance & D. Galbraith (Eds.), *Knowing what to write: Conceptual processes in text production* (pp. 13-30). Amsterdam: Amsterdam University Press.
- SIMPSON, S., & TORRANCE, M. (2007). EyeWrite (Version 5.1). Osgoode, ON: SR Research, and Nottingham, U.K.: Nottingham Trent University.
- SPELMAN MILLER, K. (2000). Academic writers on-line: Investigating pausing in the production of text. *Language Teaching Research*, **4**, 123-148.
- STARR, M. S., & RAYNER, K. (2001). Eye movements during reading: Some current controversies. *Trends in Cognitive Sciences*, **5**, 156-163.
- STRÖMQVIST, S., & KARLSSON, H. (2002). *ScriptLog for Windows: User's manual* (Tech. rep.). Lund: Lund University, and University College of Stavanger, Centre for Reading Research, Department of Linguistics.
- STRÖMQVIST, S., & MALMSTEN, L. (1998). *ScriptLog Pro 1.04: User's manual* (Tech. rep.). Göteborg: University of Göteborg, Department of Linguistics.
- SULLIVAN, K. P. H., & LINDGREN, E. (Eds.) (2006). *Studies in writing 18: Computer key-stroke logging and writing. Methods and applications*. Oxford: Elsevier.
- TORRANCE, M., FIDALGO, R., & GARCÍA, J. N. (2007). The teachability and effectiveness of cognitive self-regulation in sixth grade writers. *Learning & Instruction*, **17**, 265-285.
- VAN DEN BERGH, H., & RIJLAARSDAM, G. (1996). The dynamics of composing: Modeling writing process data. In M. Levy & S. Ransdell (Eds.), *The science of writing* (pp. 207-232). Hillsdale, NJ: Erlbaum.
- VAN DEN BERGH, H., & RIJLAARSDAM, G. (1999). The dynamics of idea generation during writing: An online study. In M. Torrance & D. Galbraith (Eds.), *Knowing what to write: Conceptual processes in text production* (pp. 99-120). Amsterdam: Amsterdam University Press.
- VAN DEN BERGH, H., & RIJLAARSDAM, G. (2007). The dynamics of idea generation during writing: An online study. In M. Torrance, L. Van Waes, & D. Galbraith (Eds.), *Writing and cognition* (pp. 125-150). Amsterdam: Elsevier.
- VAN WIJK, C., & SANDERS, T. (1999). Identifying writing strategies through text analysis. *Written Communication*, **16**, 51-75.
- WENGELIN, Å. (2002). *Text production in adults with reading and writing difficulties* (Gothenburg Monographs of Linguistics 20). Göteborg: Göteborg University, Department of Linguistics.
- WENGELIN, Å. (2006). Examining pauses in writing: Theory, methods and empirical data. In K. Sullivan & E. Lindgren (Eds.), *Computer key-stroke logging and writing* (pp. 107-130). Amsterdam: Elsevier.

NOTE

1. Both pieces of software are free for use by other researchers. Please contact the first author for ScriptLog and the second author for EyeWrite.

(Manuscript received January 14, 2008;
revision accepted for publication November 26, 2008.)