Animated Dynamic Highlighting for Reading

Olga Gurevich
Linguistics Department
University of California - Berkeley
1203 Dwinelle Hall
Berkeley, CA 94720-2650
olya@socrates.berkeley.edu

William C. Janssen, Lauri Karttunen
Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304
{janssen,karttunen}@parc.com

ABSTRACT

The recent years have seen an exponential increase in the amount of information available through the Internet on any given topic. Information retrieval techniques have been steadily improving and can provide a mass of relevant results, but those results still have to be processed and digested by a human reader. Computer displays of traditional documents often simply model the way a paper rendition of the document would work. But more effective display techniques may be possible by exploiting the dynamic display properties of the computer screen. In this paper, we describe a system called Animated Dynamic Highlighting (ADH), which has been added to the ReadUp reader, part of the UpLib personal digital library system. It is an interactive, user-controlled technique that enhances the presentational aspects of the reading task. We ran a pilot study to get an initial impression of user acceptance of active display techniques, comparing ADH to two other display techniques. The results of the study are discussed here.

Categories and Subject Descriptors

H.3.7 [**Information Storage and Retrieval**]: Digital Libraries—*systems issues, user issues*

General Terms

Design,Human Factors

Keywords

Personal digital library, reading, ADH, RSVP, UpLib, ReadUp

1. INTRODUCTION

Since the invention of writing, literacy has been an important factor in individual success. The ability to read has given us access to findings, theories, and thought of many others, some of them giants in their fields. The value of this access is immense, and it may be one reason the practice of reading has changed so slowly over the centuries. The Phoenician alphabet appeared about 4000 years ago; it took another 1000 years for vowels to be invented; punctuation came 800 years after that, and the idea of putting spaces between words took another 1000 years after that; the printing press, one of the major innovations in human history, appeared in 1436, some 600 years later.

Computers have accelerated this a bit. Englebart's group at SRI invented word processing in 1965. Xerox PARC's Alto (1972) can be called the first document computer, and Ethernet (1975) the first document network. PRESS and InterPRESS (1980) were the precursors of Postscript and PDF (1991). And the World Wide Web (1992) has changed the reading habits of the whole world. Computers have changed he practices of writing and publishing, greatly increasing the amount of material available to any individual, but the practice of reading itself, and technology to help the reader absorb this new mass of material, has remained largely unchanged.

The current model of the reading interface is heavily based on the static experience of words imaged on paper. This model has been carried over directly to the presentation of text on a computer screen. Some attention has been given to using computation to modify the presentation structure of documents [4],[13], but with certain exceptions [6] these presentations are inherently static.

The major exception to this is the presentation technique commonly known as "rapid serial visual presentation" (RSVP). The overview of studies in RSVP given in [12] suggest that a dynamically altered presentation of text may be able to enhance comprehension without negatively affecting reading speeds. However, RSVP is often found to suffer from some serious disadvantages, notably eyestrain, usually attributed to the fact that the user's eyes do not move from a fixed position, and user anxiety, due to the inability to look back at previously-read text. Other studies such as [5] have demonstrated ways to alleviate some of these issues.

The work described in this paper is part of a larger effort at PARC called "Productive Reading". We are looking at ways in which computation can be applied to the reading process, in two major ways: to enhance document content, and to enhance the user experience of reading. One of the experimental reading techniques we are looking at is called "animated dynamic highlighting", or ADH.

The goal of ADH is to preserve the apparent advantages of RSVP, such as enhanced comprehension, while mitigating the apparent disadvantages. It paces the user through an electronic document, sequentially highlighting parts of the text, each a few words long, without modifying the spatial layout of the original page, so that the reader's eyes move in a normal reading fashion. Those parts of the text not highlighted are dimmed, so that they are less likely to distract from the highlighted text. The entire page is shown so that the location of the phrase can be seen by the reader.

2. HOW ADH WORKS

ADH is a mechanism that presents a low-contrast display of the page being read, and systematically highlights phrases in the text. Each phrase is emphasized for a period that varies depending on the complexity of the phrase, then dimmed, and the next phrase is highlighted. The base speed with which the highlighting moves can be adjusted by the reader; it may be somewhat faster than the reader's normal unassisted reading spead. The reader may also restart ADH from any point in the document. When the end of each page is reached, the system automatically turns to the next page in the document being read.

2.1 Phrase Identification

Several psycholinguistic and processing studies indicate that reading comprehension and speed improve when the presented text is split into smaller units. Castelhano and Mutter [5] show that putting extra pauses at punctuation breaks improves user satisfaction with RSVP displays. Cocklin et al. [7] further show that RSVP comprehension improves when sentences are broken up into coherent "idea units" by human experts. According to Pynte and Noizet [11], sentence comprehension is accelerated if text is presented in sections; for sentences with shorter words, this means separating noun phrases from verb phrases, and for sentences with longer words, subdivisions are even smaller. Abney [1] introduces the idea of "chunks" and cites some psycholinguistic support for their existence: chunks correspond to prosodic units, and are manifested in naive sentence diagramming and pause durations in reading. Chunks have fixed syntactic parses that are mostly independent of the specific words used, whereas co-occurrence restrictions between chunks may be lexically dependent. These results indicate that both structure and length are important in identifying chunks.

Based on the above evidence, a document presented with ADH is first divided into suitable phrases. We have developed a sequence of text analysis processes that accomplishes this task.

The text of a document is first annotated with part-of-speech tags using the Inxight tagger [8]. In contrast to most taggers, the Inxight tool has a large inventory of labels to distinguish between different types of determiners, adverbs, and pronouns. While the information is less detailed than a syntactic parser could produce, the markup makes it possible to divide the text into semantically coherent pieces. We have defined a large set of phrasal patterns and compiled them into finite-state transducers [3]. The transducers are applied in a cascade taking the output of one pattern matching step as input to the next one. This process splits the input text into phrases proceeding from larger constituents (sentences and clauses) to smaller constituents (NPs, VPs, PPs) and their components. Each phrase should contain between 2 and 4 content words (such as nouns, verbs, adjectives, and adverbs); the boundaries of syntactic constituents are in most cases preserved. An example of a partitioned sentence is in figure 1.

> *The Marine Corps band / played the national anthem / as Dailey unveiled a space-suited Glenn / in his new place of honor, / suspended 40 feet above the floor / of the museum's breathtaking Gallery 100.*

Figure 1: An example of phrase-breaking.

2.2 Page Display

The ADH presentation system is part of a larger software system at PARC for archiving and reading documents, called UpLib [9]. UpLib provides a secure long-term storage and retrieval system for a wide variety of personal documents such as papers, photos, bills, books, and email. It includes a document reader, called ReadUp, which normally presents a conventional page-oriented document display, along with support for search, highlighting, annotations, bookmarks, overviews, and other important elements of the reading process. ReadUp was modified to present documents in both RSVP and ADH mode, as well as the normal page-display mode.

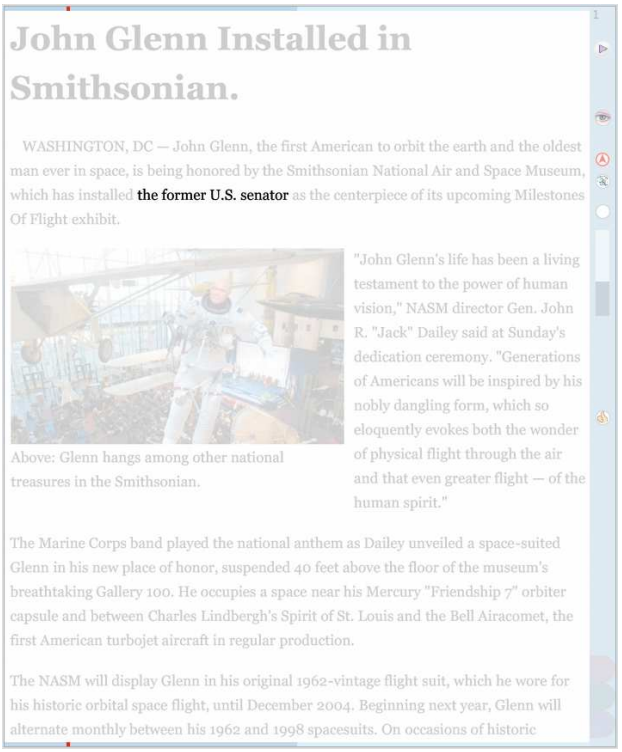

Figure 2: A document page shown with ADH highlighting.

2.3 Phrase Display Timing

Each phrase is allocated an initial display time based on the user-selected speed. This base span is then modified in a number of ways: shorter phrases get somewhat less time, longer ones more time. The timespan is further modified to reflect the findings in Just and Carpenter's theory of reading [10]: phrases ending a line, at the end of a page, at the beginning of a new line, or ending a sentence all receive varying amounts of extra time, reflecting the extra time human subjects tend to take with these kinds of phrases. Finally, the occurrence of linguistic constructs in the phrase, such as pronouns and compound nouns, is used to modify the timespan in additional ways.

3. USER STUDY

The goal of the user study was to assess the effectiveness of the ADH technique, to compare its performance to that of RSVP, and to get users' reactions and suggestions for possible improvements. The same phrase-breaking and timing were used for ADH and RSVP; the only difference was that ADH highlighted phrases sequentially in their original position on the screen while RSVP displayed one phrase at a time.

3.1 Method

3.1.1 Participants

There were 18 participants, mostly scientific researchers and interns; 13 were male, 5 were female. Age ranged from 21 to 63, with a mean of 36 years.

3.1.2 Materials and Technology

The participants read texts containing simple factual information about encyclopedia concepts that they were not likely to know well. Six texts of 2-3 pages long were used in the study. Each document was accompanied by five questions testing the accuracy of the recall.

The documents were presented in three alternative modes of presentation: plain (not modified in any way), ADH, and RSVP. All documents were presented in ReadUp document reader with modifications appropriate for each mode.

In both ADH and RSVP modes, a thin toolbar appeared at the top of the screen, where a red dot indicated the current position within the document. By clicking anywhere on the toolbar, participants were able to move quickly forwards or backwards within the document. Speed of display was controlled by a mouse scroll wheel: scrolling up resulted in faster display, scrolling down resulted in slower display. A toolbar on the side of the screen indicated current speed settings relative to a default. In addition, in ADH mode participants could click anywhere on a page, causing the ADH process to restart from that point. Pressing the space bar in either mode caused ADH or RSVP to pause; pressing the space bar again resumed it.

In ADH, phrases were sequentially highlighted on the screen while the rest of the page remained visible but greyed out. In RSVP, the phrases were displayed centered on the screen for the same amount of time as they would be displayed in ADH, and the rest of the page was not visible.

3.1.3 Design and Procedure

At the beginning of the experiment, each participant was given a baseline reading test, consisting of a short document and several questions testing the recall of factual information. Participants also filled out short questionnaires with information about their educational level and reading habits.

There were three conditions: Plain (unmodified), ADH, and RSVP. The design was factorial, cross-matching the documents, the three modes of presentation, and the order of presentation.

After reading each document, participants were asked to answer comprehension questions via a web form. After trying out a new presentation mode, participants also answered questions about their reactions to the technology. Participants were invited to come back for a shorter second stage of the experiment. Most agreed to come back one or two days later.

The second stage had the same three conditions as the first stage, but the ADH and RSVP modes used manual phrase-breaking instead of automatic.

3.2 Results

3.2.1 Timing and Accuracy

Although there were too few subjects for statistically significant results, some interesting trends emerged. Overall, ADH was found to be faster than either plain or RSVP mode; it was also somewhat less accurate.

With automatic phrase breaking, ADH timing was about the same as plain timing, but faster than RSVP. With manual phrase breaking, ADH was faster than both plain and RSVP modes. With automatic phrase breaking, ADH was less accurate than RSVP, and RSVP was less accurate than the plain mode. With manual phrase breaking, ADH and plain modes were closer together, and RSVP was more accurate than the plain mode. In general, there was a tradeoff between speed and accuracy in ADH: the faster a document was read, the less accurate was the recall. However, both the speed and accuracy results were better with manual phrase-breaking than with automatic phrase-breaking.

There was no significant effect of age, initial reading speed, or initial accuracy on ADH and RSVP timing and accuracy.

3.2.2 User Satisfaction Ratings

Users found both ADH and RSVP to be somewhat annoying, but rated RSVP worse than ADH. There was no significant difference between automatic and manual phrase breaking in user ratings.

With increasing age, there was a slight increase in ratings for RSVP and a slight decrease in ratings for ADH. Younger participants tended to like ADH more than RSVP, and older participants tended to prefer RSVP over ADH.

3.2.3 User Comments

Despite complaints about phrase timing, most users said they would use ADH again for skimming through short ar-

ticles, especially with improved phrase-breaking and timing algorithms. On the other hand, most users rejected future uses of RSVP.

3.3 Discussion

It takes time to get used to new technology, so the lower user ratings and reading speeds may be the result of novelty shock. The results are nevertheless encouraging: younger subjects in particular were very enthusiastic about ADH. The tradeoff between speed and accuracy is to be expected; the goal of reading-aid technology should be to decrease the slope of the tradeoff. Our user study produced many suggestions for future improvements of ADH technology as well as possible applications.

4. CONCLUSION

ADH is one of the many possibilities inherent in the idea of actively presented text. Interfaces that attempt to work with the user in understanding the underlying text would seem to have wide applicability for reading text of all kinds, from technical papers to email to biography, particularly in overview reading, such as Adler's "systematic skimming" and "superficial reading" [2]. They may offer special advantages to those with reading disabilities, or for specific tasks, such as proofreading. Our initial investigations into this technique seem promising, and a number of improvements in both phrase analysis and presentation timing are already being investigated.

5. REFERENCES

[1] S. Abney. *Principle-Based Parsing*, chapter Parsing By Chunks. Kluwer Academic Publishers, 1991.

[2] M. J. Adler and C. V. Doren. *How to Read a Book*. Touchstone Books, revised edition, 1972.

[3] K. Beesley and L. Karttunen. *Finite State Morphology*. Stanford: CLSI Publications, 2003.

[4] T. G. Bever, R. Burwell, S. Jandreau, R. M. Kaplan, and A. Zaenen. Spacing printed text to isolate major phrases improves readability. *Visible Language*, 25:75–87, 1990.

[5] M. S. Castelhano and P. Muter. Optimizing the reading of electronic text using rapid serial visual presentation. *Behaviour & Information Technology*, 20(4):237–247, 2001. http://cogprints.ecs.soton.ac.uk/archive/00001934/00/castelhano_muter2001.pdf.

[6] B.-W. Chang, J. Mackinlay, and P. T. Zellweger. Fluidly revealing information in Fluid Documents. In *Proceedings of Smart Graphics 2000 AAAI Spring Symposium*, 2000. http://www2.parc.com/istl/projects/fluid/fluid-docs-smartgraphics2000.pdf.

[7] T. G. Cocklin, N. J. Ward, H. C. Chen, and J. F. Juola. Factors influencing readability of rapidly presented text segments. *Memory and Cognition*, 12(5):431–442, 1984.

[8] I. S. Inc. Inxight linguistx platform, 2000. http://www.inxight.com/products/sdks/lx/.

[9] W. C. Janssen and K. Popat. UpLib: A universal personal digital library system. In *Proceedings of the 2003 ACM symposium on Document Engineering*, pages 234–242. ACM Press, November 2003.

[10] M. A. Just and P. A. Carpenter. A theory of reading: From eye fixations to comprehension. *Psychological Review*, 87(4):329–354, July 1980.

[11] J. Pynte and G. Noizet. Optimal segmentation for sentences displayed on a video screen. In *Proceedings of the 2nd Conference on Processing of Visible Language*, pages 375–385, September 1979.

[12] K. Sicheritz. *Applying the Rapid Serial Presentation Technique to Personal Digital Assistants*. PhD thesis, Uppsala University, 2000.

[13] R. C. Walker and S. D. Walker. An introduction to live ink technology, 2001. http://www.clipread.com/sbir/investigators_manual.doc.