Citation for published version:

Duggan, GB & Payne, SJ 2009, 'Text skimming: the process and effectiveness of foraging through text under time pressure', *Journal of Experimental Psychology: Applied*, vol. 15, no. 3, pp. 228-242.
<https://doi.org/10.1037/a0016995>

DOI:

[10.1037/a0016995](https://doi.org/10.1037/a0016995)

Publication date:

2009

Document Version

Peer reviewed version

[Link to publication](#)

This article may not exactly replicate the final version published in the APA journal. It is not the copy of record

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Running head: TEXT SKIMMING

Text Skimming: The Process and Effectiveness of Foraging Through Text under Time
Pressure.

Geoffrey B. Duggan and Stephen J. Payne

University of Bath

Bath, United Kingdom

Correspondence concerning this article should be sent to: Geoffrey Duggan,
Department of Computer Science, University of Bath, Bath, BA2 7AY, UK. Email:
gbduggan@yahoo.co.uk.

Abstract

Is skim reading effective? How do readers allocate their attention selectively? We report three experiments that use expository texts and allow readers only enough time to read half of each document. Experiment 1 found that, relative to reading half the text, skimming improved memory for important ideas from a text but did not improve memory of less important details or of inferences made from information within the text. Experiment 2 found no advantage of skimming over reading the first or second half of every paragraph. Two final experiments using a hierarchical, website-like layout of documents showed that advantage of skimming found in Experiment 1 was dependent on the linkages between pages and, thus, the ease with which participants could navigate through the text. Data on page-by-page reading times, and eye-tracking analyses from Experiment 2 indicated that skim readers spent more time reading text that was earlier in the paragraph, towards the top of the page and in an earlier page of the document. These findings were interpreted as evidence in support of a satisficing account of skimming process.

Text Skimming: The Process and Effectiveness of Foraging Through Text under Time Pressure.

How do readers cope when there is too much text to read in the time available? The rise of the Web and the ready availability of source documents have made this problem more common than ever. The literature on web browsing behaviour (e.g. Liu, 2005; Morkes & Nielsen, 1997) indicates that readers often skim; that is, they employ some form of rapid, selective reading strategy such as omitting words, paragraphs or pages. Presumably, if skimming is effective, it is because it enables readers to ignore less valuable parts of the text and to focus attention on the more useful parts.

But is this presumption correct? Are readers able to allocate attention effectively when faced with a longer document than they can read in the time available? And if so, by the use of what strategy? In fact, the little empirical literature that directly addresses the primary question is inconclusive at best.

Experimental Studies of Skim Reading

Perhaps the most extensive investigation of the impact of skimming upon comprehension and memory was conducted by Masson (1982). In four experiments he tested participants' recognition memory after reading through short texts (between 400-1000 words) at different rates. The different reading rates were manipulated by instruction in one experiment and by restricting reading time in the remaining experiments. Memory for the meaning of the text was subsequently tested by asking participants to judge whether test statements contained information mentioned in, or consistent with, the text. The statements referred to either "important" or "unimportant" information from the text or required participants to make an inference using information contained within the text. Important versus unimportant was a

judgment made by a separate group of readers; Masson assumed it was related to the distinction by Kintsch and colleagues (Kintsch, 1998) between macropropositions (important) and micropropositions (relatively unimportant). False statements were constructed by changing the meaning of true statements.

Results showed that recognition accuracy (d') for all statement types deteriorated roughly equally as reading rate increased (i.e. the *difference* in d' scores was more or less the same) across all statement types. There was also evidence that faster reading produced slower response times when verifying unimportant statements but did not affect response times to important statements. Masson suggested that skim readers were unable to perceptually select and focus on more important passages, but they had devoted more conceptual processing to the most relevant information and that this additional processing enabled faster responses. (This is a kind of selection, of course, but not the kind of selection that our intuition supposes underlies the adaptive character of skimming.)

Studies since Masson have similarly found little evidence for effective skimming. Carver (1984) presented 100-word passages to readers for different lengths of time and found shorter reading times reduced performance on tests of macrostructure and microstructure to similar extents. Dyson and Haselgrove (2000) asked participants to read texts of approximately 500 words at their normal reading speed or at a self-paced faster rate – on average about twice as fast. Again, there was no interaction between the type of information retained and reading speed. Performance on multi-choice recall questions was worse at higher reading speeds, and at both reading speeds, general information was recalled better than specific details. This finding was also reported by Kiwan, Ahmed and Pollitt (2000) who presented a

text of approximately 500 words text to children one paragraph at a time varying the presentation time between participants.

In summary, there is little experimental evidence that skim readers are able to devote their limited time to the most valuable parts of a single document. But perhaps the reasons for this are to be found in some consistent aspects of the experimental designs. First, in all the above experiments, the texts that readers were asked to skim were relatively short and unstructured. We regard it as an important experimental design constraint that all parts of the text should be, *a priori*, relevant – it would not be very interesting to show that readers of experimental reports who are instructed to learn about the method (rather than the results, say) are able to focus in on the Method section. Nevertheless, if skimming is to be effective at discovering “important” aspects of a text, the reader must be able to make inferences about the content of parts of the text before they are read (so as to know whether to skip over them). Assumptions of continuity and coherence may allow such inferences, but it seems to us possible that these processes may be more effective at a coarser grain, and therefore within a longer document. Our reasoning here is supported by our proposal of *how* skimming might operate in longer texts and in the absence of distinct sections of variable relevance marked by headings – namely, by a patch-leaving satisficing strategy, as discussed below.

Second, the logic of all the above studies is the same. For skimming to be effective, the argument goes, readers must identify and devote more attention to more important parts of the text, and so any reduction in comprehension due to reduced time should be less severe for more important content. This logic seems compelling, but statistical issues may make it difficult to measure the predicted effect, leading to the reported null effects. Statistically the prediction in all these cases was an

interaction effect (greater reduction in performance scores under time pressure for one type of question). However, interaction effects (or their absence) are hard to interpret when baseline scores are at different levels. For example, in Masson's Experiment 1, the d' score for important statements after normal reading of narratives was approximately 2.2, whereas for unimportant items it was approximately 1.8 (Masson, 1982, Figure 1). Because the underlying psychological scale may not be linear, it is not clear whether the same psychological decrement in sensitivity would lead to the same reduction in the scores (e.g. Anderson, 2004).

Our response to this analysis is to use longer texts, and a different experimental logic. We will discuss our logic shortly, but first, we review a study that did provide some limited evidence for effective skimming, and at the same time introduced a sketch of the process by which skimming may proceed.

Satisficing as a Strategy for Skimming in Multiple-Texts

A recent experiment on skimming of multiple on-line texts (Reader & Payne, 2007) provides some evidence of adaptive allocation of attention. Reader and Payne presented participants with a browser interface to four separate texts on the human heart (each around 1500 words long), and asked them to study the texts for 15 minutes in order subsequently to write a brief essay. The texts varied considerably in terms of their difficulty level. Participants' time allocation across the texts was reliably affected by their prior expertise (as measured by a pre-test on the heart). Readers that are more expert spent more of their time on higher-level texts, despite the absence of any cues to difficulty other than the content of the texts themselves.

On the basis of the readers' behavioural protocols – the timing of their switches from one page to another, Reader and Payne suggested that readers were able to achieve some adaptive allocation of attention by using a simple satisficing

strategy. The strategy they outline derives from an information foraging approach to browsing (Pirolli & Card, 1999) in that it assumes that readers are sensitive to some proxy for “information gain” and use this as a basis for their choices. Suppose a reader monitors the rate of information gain while reading. Further, suppose that they set a threshold of acceptable gain. If readers begin to read linearly, they may continue until the rate of information gain drops below the acceptable threshold. At this point they will leave the current “patch” of text and skip to the beginning of the next patch.

As an account of skimming, this version of satisficing is very unspecified, but nevertheless provides a framework, and given some additional assumptions, makes some predictions for how skimming might proceed. First, although we do not know how information gain may be monitored or measured, we might assume that when the time available to read is low given the quantity of text available, the satisficing threshold might be set rather high, so that the reader will only dwell on very informative parts of the text.

Second, although we do not know how the reader may divide a text into patches, we might assume that when all major sections of a document might be relevant, the patch-size will be small, perhaps a paragraph. Reader and Payne (2007) reported that most visits to a document in their experiment incorporated sequential visits to all four pages of the document. On this basis, they concluded that the patch-size must be a page or lower (because when a page was abandoned, the next page in the document was visited). Because paragraphs are conventional topic-units, and because paragraph-beginnings are perceptually easy to scan to, paragraphs may be a plausible “patch” in many situations.

If this satisficing strategy is to allow successful skimming, then patches of the text must be differentially valuable to the reader, and a patch that begins well must be

more likely to contain important information. If, on the other hand, “important” facts are randomly or evenly distributed among patches, then skimming will not be successful. With very short articles or stories, such as those used in the literature on skimming reviewed above, this may be the case.

The Current Experiments

In the light of the above arguments, the experiments reported below use a new experimental logic to test whether skimming is effective for longer, on-line expository texts organized as traditional linear documents (Experiments 1 and 2) or hierarchical websites (Experiment 3). Further, we test predictions from the satisficing framework about how skimming might be controlled.

All the experiments we report use texts of around 3000 words in length, organized into thematic sections (all of which are relevant to the readers’ goals). The texts used in all experiments were divided into 11 pages and each page had a separate heading and covered a different subtopic. We argue that this increases the scope for participants to allocate their time adaptively, especially if they are using a satisficing strategy (as we predict they will). The texts were adapted from Scientific American articles.

We tested understanding and memory in the same way as Masson, by asking participants to make recognition judgements about statements from the text. The Important and Unimportant statements may be taken to refer to knowledge about the macrostructure and microstructure of the text respectively and the Inference statements to provide an indication of the situation model (using the general framework for text comprehension due to Kintsch and colleagues; Kintsch, 1998).

With regard to experimental logic, we suggest that a more powerful test of the advantages of skimming is to hold constant the amount of time available and to

manipulate the discretion of the reader to allocate time across the text. If a reader only has half of the time they would require to read a text, do they nevertheless benefit from having the whole text available for skimming, compared with linear reading through an arbitrary half of the text? This seems to us to be an experimental test of the effectiveness of skimming that represents the everyday dilemma of time being fixed at too low a level to allow complete reading.

Different experiments below use different versions of a control condition, in which participants only have access to part of the text, and read this part linearly, with no, or little, discretionary skimming. We should not call these control conditions “normal reading” – they involve straightforward linear reading at the usual pace, but through only a part of a document. In the first experiment, this part is either the first or the second half of the document. In the second experiment, it is the first or second half of every paragraph. In the third experiment, it is a half of the document defined according to a hierarchical structuring of the document in a web-page layout.

In the first experiment, we included an additional test of the readers’ representation of the text. Payne and Reader (2006) demonstrated that after reading a text participants had an improved “map” of the structure of a text and could locate the position of target sentences within the text more quickly. This structure map becomes particularly valuable when readers are faced with a large amount of text that they must navigate back and forth between as often happens with hypertext and on the Web generally. To test the structure map formed through skimming, participants were presented with a subset of the Important and Unimportant sentences and required to identify which page they were on.

Whatever strategy they pursue, skim readers have to choose to read some parts of the text, in advance of knowing how important they are. What parts of the text

should they choose? A simple answer, in alignment with the satisficing hypothesis, is that they should read beginnings - of documents, of sections, of paragraphs. After all, documents are written to be read from beginning to end, and advice to writers often includes injunctions to begin every sub-section with important orienting material (e.g. Strunk & White, 1979). Therefore, beginnings are not only likely to contain important information, but likely reliable indicators of the value of middles and ends.

Our documents were organized into pages, with pages corresponding to meaningful sections. Consequently, to gain insight into skimming *process* we consider the time-allocation across pages and additionally, in the second experiment, using eye-tracking, within pages.

Experiment 1

In this experiment, when participants were in the half-text condition, they received either the first or the second half of the text and were asked to read through at their normal speed. The time available in both cases was set to an underestimate of the time required to read half the text at normal rates – we did not want any participants in the “Half text” condition to have any spare time they might not use.

Method

Participants

The participants were 10 male and 22 female students from the University of Manchester. The mean age was 21.75 years ($SD = 2.66$). They received £5 (U.S. \$9.79) for taking part in the study. There was no overlap between participants across any of the experiments reported.

Design

All participants read one of two texts in the Skim condition and the other text in the Half-text condition. In the Half-text condition, one group of 16 participants

could only see the first half of the text (Half 1 condition) and the remaining 16 participants (Half 2 condition) could only see the second half of the text. Presentation order was counterbalanced across participants such that each combination of text, order and condition appeared an equal number of times.

The time spent reading each individual page and the number of times participants switched between the pages of the text was recorded. Participants' memory for the meaning of each text was assessed using the proportion of true and false sentences correctly categorised during a recognition memory test. Participants' structure maps were tested using the number of target sentences that were correctly located within the text.

Materials

Texts. Two texts were adapted from Scientific American articles on Hypnosis (3134 words) and Attention-Deficit Hyperactivity Disorder (3252 words). Word frequency profiles of the texts were extracted from the "Edict" text analyser. These found that for the ADHD text 72.08% of words were in the 2000 most frequent list and 8.21% of words were in the 2000-5000 most frequent lists. For the Hypnosis text, 71.93% of words were in the 2000 most frequent list and 8.93% of words were in the 2000-5000 most frequent lists. Readability of the texts was tested using the Gunning-Fog Index. The ADHD text scored 11.7 and the Hypnosis text scored 12.6 (6 = easy, 20 = hard).

The two texts were restructured into 11 pages of approximately equal length (Page length in words: ADHD, $M = 292.64$, $SD = 75.42$; Hypnosis, $M = 281.73$, $SD = 87.20$). Page boundaries were selected so that the content within each page addressed a distinct macroproposition that could be summarised using a single heading.

Section headings from the original text were removed and for each page a separate heading was devised. The goal of these headings was to provide an indication of the content without explicitly mentioning information from any target sentences.

For the Half-text condition the texts were divided in half at the nearest end of page to the midway point of the text. Thus, pages 1-5 of the ADHD text formed the first half (1623 words) and pages 6-11 formed the second half (1629 words). The Hypnosis text was divided into pages 1-6 for the first half (1600 words) and pages 7-11 for the second half (1534 words).

Memory-for-meaning test sentences. From each text 36 sentences were selected for the test of readers' memory for meaning. These sentences were categorised following the procedure used in Masson (1982). A group of 20 undergraduates who did not participate in the main study rated the sentences for their importance to the general meaning of the text. The nine most highly rated sentences were used as the True Important sentences and the nine lowest rated sentences were used as the True Unimportant sentences. Of the remaining sentences the nine most highly rated were used as False Important sentences and the others as False Unimportant sentences.

An additional group of 18 statements for each article was constructed by drawing intuitively important inferences from specific information mentioned in the text and then writing these inferences in Standard English. For example, the inference statement "Similarities in the incidence rates of ADHD between identical twins do not seem to be purely due to environmental factors." was based upon specific information in the text about a higher rate of ADHD in identical twins relative to non-twin siblings. These statements were rated in the same way as above; the nine sentences rated most important were included as True Inference sentences and the remainder

were used to generate False Inference sentences. Each of the 27 false sentences for each text was altered so as to be semantically incongruent with the original statement. Thus, altering the earlier example would produce the following statement “Similarities in the incidence rates of ADHD between identical twins are purely due to environmental factors.”

The different sentence types were distributed approximately evenly across the different pages in the text with all pages containing at least two target sentences. The standard deviation of the number of sentences per page was 1.85. The mean number of words between the start of the text and the target sentence was 1467.47 ($SD = 958.27$) for the Important sentences and 1570.89 ($SD = 898.23$) for the Unimportant sentences. This difference was not significant ($t < 1$). Across both texts, there were 36 examples of each different sentence type and there were 54 true sentences and 54 false sentences. These sentences were distributed such that in total in the first half of a text there were 21 Important, 19 Unimportant, 18 Inference, 29 True and 28 False sentences.

For the test of sentence location a group of 12 target sentences was formed using six of the True Important sentences and six of the True Unimportant sentences from the recognition test. No more than two of these sentences occurred on any one page and both the Important and Unimportant sentences were equally split between the first and second half of both texts.

Task interface. A Microsoft Visual Basic 6 program presented the texts and the subsequent test of memory for meaning. The program recorded and timestamped all mouse clicks and keyboard entries during the experiment.

The texts were presented one page at a time in a large box that filled most of the screen. Participants could navigate back and forth through the text by clicking on

one of 11 buttons from a column down the right hand side of the screen (top to bottom in page order). Each button was labelled with the heading from the corresponding page. The heading for the currently selected page was displayed at the top of the screen with the main body of text for that page presented directly below the heading. The texts were presented in Microsoft Sans Serif font size 12, with up to 106 characters fitting on each line. There was a single line space between each line of text and a double line space between each paragraph of text. In the bottom right hand corner of the screen a clock counted down the remaining number of seconds before the text was removed.

In the Half conditions participants could only view half of the text. Buttons corresponding to pages in the other half of the text were displayed in a lighter shade of grey and were disabled.

Procedure

Participants were instructed that they would be presented with a text on either Hypnosis or ADHD for a limited period of time before being tested on their understanding of the whole text. In the Skim condition, they were told that there was not enough time to read through the text at normal reading speed but that they should allocate their time while reading so as to maximise performance on the subsequent test of memory for meaning. In the Half-text condition, participants were instructed to read through the text unidirectionally at their normal reading speed being aware that they would subsequently be tested on their understanding of the text. In the Half-text condition, it was emphasised that they did not have to finish reading all the visible text.

To ensure that in the Skim condition participants could not read all the text in the time available, the time limit was set so that participants would have to read at 600

words per minute to read through the complete text. (Masson reported that normal reading speed was approximately 225 words/minute.) This meant that, in all conditions, the ADHD text was presented for 325 seconds and the Hypnosis text was presented for 313 seconds.

When the time for reading had elapsed participants were required to complete the test of memory for meaning. After one practice sentence all 54 target sentences were presented serially in the centre of the screen in a randomised order. Each sentence remained on screen until the participant made a response. Participants were instructed to press a key labelled “True” if they thought the sentence was contained in the text they had just read or if the sentence contained information that was consistent with the text. If they thought the sentence was inconsistent with the text they were told to press a key labelled “False”. Participants were told to prioritise accuracy over speed when making their responses.

After completing the memory-for-meaning test, participants were required to indicate the specific page within the text that contained a target sentence. The interface was the same as for the reading phase except that the page of text was no longer displayed and was replaced by a target sentence. Participants read the sentence then selected one of the page headings that labelled the buttons on the right hand side of the screen. All 12 target sentences were presented in this way in a different random order for each participant.

After completing the sentence location test participants repeated the whole procedure in a different reading condition using the other text.

Results

Analysis Strategy

Some features of our analysis strategy were common across all three experiments. Our main experimental hypothesis relates to memory-for-meaning, which we predict will be greater when skimming than when presented with a half-text, especially for Important sentences. Our experimental design allowed analysis of these effects in a 5-way ANOVA: 2 (Text Assignments, i.e., the group of each participant according to which text was skimmed and which was half-read: ADHD-skimmed or Hypnosis-skimmed) x 2 (Condition Order: Skim first or Skim second) x 2 (Text Half: First or Second half read in the Half-text condition) x 2 (Reading Condition; Skim or Half-text) x 3 (Sentence Type: Important or Unimportant or Inference), with the first three factors between-participants and the final two factors within-participants.

This design did not allow Text (i.e. ADHD or Hypnosis irrespective of whether it was skimmed or not) as a factor, yet exploratory analysis revealed reliable differences in sensitivity scores according to text. Consequently, recognition memory indices of sensitivity (d') and bias (C) were converted into z -scores (using the means and standard deviations across all conditions for each text) prior to the 5-way ANOVA. For all experiments the difference in d' score between the two texts is reported and thereafter “sensitivity” and “bias” is used to refer to the z -scores computed from the d' and C scores.

In view of the complexity of the ANOVA, we report all, but only significant effects. In addition to the memory-for-meaning results, some indications of reading process are reported, as they relate to skimming strategy. These analyses were done separately for Skim-reading and Half-text conditions, or only for skimming. Consequently, in these analyses Text was treated as a between-participants factor.

Memory for Meaning

The d' scores were higher for the ADHD text ($M = 1.37$, $SD = .58$) than for the Hypnosis text ($M = 1.02$, $SD = .58$), $t(31) = 3.39$, $p < .01$. There was no effect of Presentation Order ($t < 1$). Means and standard deviations for hit rate, false alarm rate, d' and C scores are given in Table 1.

Sensitivity. Table 1 shows that for all 3 sentence types d' scores were higher for the Skim condition than the Half-text condition. The main effect of Reading Condition across all sentence types was reliable, $F(1, 24) = 7.78$, $MSE = .60$, $p < .05$, $\eta^2_p = .25$.

Inspection of Table 1 also suggests that the difference between Skim and Half-text was larger for the Important sentences than for the Unimportant and Inference sentences. Planned comparisons supported this observation as there was improved sensitivity for Important sentences after skimming the text than reading half the text, $F(1, 24) = 10.99$, $p < .01$, $\eta^2_p = .31$. There was no significant effect of Reading Condition upon sensitivity for the other two sentence types ($F_s < 2$).

There was also a main effect of Sentence Type, $F(2, 48) = 10.80$, $MSE = .64$, $p < .001$, $\eta^2_p = .31$. This was orthogonal to our hypotheses and indicated that memory for meaning score was lower for the Unimportant sentences.

Bias. The C scores in Table 1 indicate that across both True and False sentences in the Skim condition there was a greater bias to respond “True” after skimming a text than after reading half the text. This bias was reliable across all sentence types, $F(1, 24) = 17.09$, $MSE = .66$, $p < .001$, $\eta^2_p = .42$. Planned comparisons found the effect was significant for the Important sentences, $F(1, 24) = 8.78$, $p < .01$, $\eta^2_p = .27$ and the Unimportant sentences, $F(1, 24) = 6.28$, $p < .05$, $\eta^2_p = .21$.

Reading Times per Page

Reading time per word was calculated for each page, to assess whether pages that were earlier in the text were read more slowly. For each condition times were averaged across texts and presented in Figure 1.

Figure 1 indicates that reading time per page decreased over the course of the text in all three conditions. To test this, bearing in mind the non-independence of times across individual pages given the total time limit, the texts were divided into rough thirds and the time spent reading the first and last third of each text was compared. For the Skim condition mean time per page was calculated for the first and last three pages and in the Half conditions the reading time per page for the first and last pages were used. For each condition a three-way ANOVA was conducted with Text (ADHD or Hypnosis) and Presentation Order (First or Second) as between participants variables and Text Part (First or Last) as a within participants variable. Technical difficulties meant that reading times were not collected for four participants in the Skim condition, one participant in the Half 1 condition and two participants in the Half 2 condition.

The analyses indicated that in all three conditions the first part of the text was read for longer than the last part of the text. Skim condition, $F(1, 24) = 50.13$, $MSE = 3824.86$, $p < .001$, $\eta^2_p = .68$; Half 1 condition, $F(1, 11) = 26.40$, $MSE = 12757.65$, $p < .001$, $\eta^2_p = .71$; Half 2 condition, $F(1, 10) = 28.36$, $MSE = 13949.90$, $p < .001$, $\eta^2_p = .74$. No other effects or interactions were significant.

The mean number of pages visited was 9.22 ($SD = 2.34$) in the Skim condition, 5 ($SD = .94$) in the Half 1 condition and 4.69 ($SD = 1.31$) in the Half 2 condition. In the Skim condition the mean number of times participants revisited a page was 1 ($SD = 1.76$). When these revisits were excluded, 31 out of 32 participants read the pages in the order they were presented.

Memory for Sentence Location

We used two measures of the accuracy of participants' judgments of sentence location. These were the proportion of sentences for which the correct page was selected and the distance of the selected page from the correct page (a correct answer would have a score of 0, and selecting page 5 when page 3 was the correct answer would produce a score of 2). The mean scores on both measures for the Skim and Half-text conditions are given in Table 2.

These scores were analysed using the same 5-way ANOVA as for the memory for meaning results. The analyses found no reliable differences between the Skim and Half-Text conditions for location accuracy of either Important or Unimportant sentences. There were no reliable interactions between Reading Condition and any other variables.

Discussion

The memory-for-meaning results provided clear support for our predictions. Readers correctly recognised more of the Important sentences after skimming the text than after reading linearly through half the text at normal pace. Therefore, readers successfully gleaned more information about the macrostructure of the text when skimming than when reading at normal speed. The satisficing model would suggest skimmers achieved this by successfully leaving a section of text when the rate of information gain was low.

As anticipated, skimming, relative to reading at normal speed, did not improve the understanding of unimportant sentences from the text suggesting that skimmers successfully skipped over less important parts of the text. Nor did the reading conditions differentially support a situation model of the text as assessed by understanding of Inference statements. This is consistent with our hypothesis that

skimming meant there was often insufficient time to develop a deeper, more elaborated understanding of the text. There was no evidence from the sentence location task that readers had an improved structure map after skimming than after reading at normal speed.

Relative to reading half the text, skimming also led to an increased tendency to respond “True” to both true and false sentences. This bias is consistent with an earlier pilot study (Duggan & Payne, 2006) and may indicate a discrepancy between the readers’ perceptions and the incomplete situation model developed during skimming.

Experiment 2

In this experiment the “Half text” conditions gave participants access to either the first half or the second half of every paragraph. Because they read paragraph-halves in sequence, readers in these conditions very likely lose some of the text’s coherence cues. But, we would argue, so do skim readers, who presumably select which parts of text to skip without knowing what they contain. Indeed, our reason for using half-paragraph as a comparison condition is to test our intuition that reading the first part of every paragraph would be quite an effective and simple-to-apply heuristic for skimming. We still predict an advantage for skim readers who have discretion over how they allocate their attention, but we predict that this advantage will be considerably attenuated in the case of the first-half of paragraph condition.

This experiment also uses eye-tracking to record finer-grained information about skim readers’ strategies. Following Reader and Payne (2007), we suggest the paragraph may be the patch-size evaluated by many readers when satisficing. In this case, skim readers should spend more time in the first half of paragraphs than the second half (because if the rate of information gain from the paragraph is judged as

below threshold, it will be abandoned, and the reader will leap to the beginning of the next).

Method

Participants

The participants were 12 male and 20 female students from the University of Manchester. The mean age was 20.97 years ($SD = 2.25$). They received £5 (U.S. \$9.79) for taking part in the study.

Apparatus

A Tobii 1750 binocular eyetracker monitored participants' eye movements while they read. The eyetracker consisted of a high resolution camera embedded in a 17-inch monitor. A PC was connected to the eyetracker and the stimulus materials were presented on the eyetracker monitor. The fixed wide-angle camera allowed data to be recorded from a freely moving person with approximately 20 cm of freedom on each side. The screen resolution on the monitor was set to 1024 x 768 pixels and the sampling rate was 20 ms. Participants were seated 61 cm from the monitor.

The application ClearView was used to export the eye-tracking data. The fixation filter from this package determined the start and end of a fixation. Thus, a fixation was deemed to start when the distance between each successive gaze point remained within a specified number of pixels for a specified period of time. A fixation ended when the distance between successive gaze points exceeded the threshold. The fixation filter size was set to 20 pixels and the minimal fixation duration was 40 ms.

To maximise ecological validity, relatively large amounts of text were displayed on screen and the font size of the text was smaller than in standard eye-tracking studies of sentence comprehension. This, combined with the technical specification of the eyetracker, prevented a detailed analysis of gaze location as a

function of word location. Therefore, the half-paragraph was adopted as the smallest unit of analysis.

We analysed two standard eye-movement measures. First-pass time is the sum of all fixation durations beginning with the first fixation in a region until the reader's gaze leaves the region. Total time is the sum of all fixation durations in a region, regardless of order.

Materials

Each paragraph within the texts was divided in half. The boundary point between the two halves was selected by finding the midway point of the paragraph in words and then inserting the boundary at the nearest end of sentence to this halfway point.

For 6 paragraphs in the ADHD text and 4 paragraphs in the Hypnosis text, the halfway boundary was moved by one sentence to counterbalance the sentences across each half of paragraph. This meant that, within each text, there were an equal number of Important and Unimportant, and True and False sentences in the first and second half of a paragraph. Typically, the Inference sentences could not be easily categorised as relying purely on information from one half or another of a paragraph. After these alterations the mean length of the first half of the paragraph was 53.34 words ($SD = 21.93$) and the mean length of the second half of the paragraph was 53.75 words ($SD = 24.41$).

Text presentation for the Skim condition was identical to Experiment 1. In the Half-text condition the text was laid out in exactly the same way as for the Skim condition but the half paragraph not presented was occluded with a string of "X"s. Thus, the visible half of the paragraph was presented in the same position as for the Skim condition.

All other aspects of the method were the same as in Experiment 1.

Results

Memory for Meaning

The d' scores were higher for the ADHD text ($M = 1.35$, $SD = .52$) than for the Hypnosis text ($M = .96$, $SD = .49$), $t(31) = 3.91$, $p < .001$. There was no effect of Presentation Order ($t < 1$). The data were analysed using the same 5-way ANOVA as in Experiment 1. Table 3 shows the means and standard deviations for hit rate, false alarm rate, d' and C scores.

Sensitivity. Sensitivity scores were lower for Unimportant sentences than for Important or Inference sentences, and the main effect of Sentence Type on sensitivity was reliable, $F(2, 48) = 22.56$, $MSE = .57$, $p < .001$, $\eta^2_p = .49$. Table 3 suggests that across all 3 sentence types, memory for meaning was slightly better in the Skim condition than the Half condition, however, this difference did not approach significance ($F < 1$). Planned comparisons between the Skim and Half-text conditions for each sentence type revealed no significant differences.

Bias. Table 3 indicates a greater tendency to respond “True” to the Important sentences and the Inference sentences than the Unimportant sentences. This was supported by a main effect of Sentence Type on bias, $F(2, 48) = 3.74$, $MSE = .51$, $p < .05$, $\eta^2_p = .23$.

Eye-tracking Results

Analyses of eye-tracking data focused on the Skim condition. In the Half-paragraph conditions readers did not have to be selective - and could NOT be selective between half-paragraphs (although they did perhaps have to search for paragraph beginnings in perceptually awkward pages containing blocks of Xs as well

as text). We had no particular hypotheses about reading strategy under these conditions.

To analyse the data from the Skim condition each paragraph in the Skim condition was divided into two halves in the same way as for the presentation of the text in the Half 1 and Half 2 conditions. The total time each participant spent fixating within each half-paragraph was computed and divided by the number of words in that half-paragraph. This measure of “total fixation duration per word” was then aggregated across all eleven pages to produce a mean time for half-paragraphs within the first, second and third paragraphs on each page. These are given in Figure 2. Only one page from each of the texts had a fourth or fifth paragraph therefore data from these paragraphs were not included in any analysis. Technical difficulties achieving accurate calibration meant that the eye-tracking data from four participants were not useable.

The total fixation duration per word data in the Skim condition were analysed using a 2 (Text: ADHD or Hypnosis) $\times$ 2 (Presentation Order: First or Second) $\times$ 3 (Paragraph: 1, 2 or 3) $\times$ 2 (Paragraph Half: 1st or 2nd) mixed ANOVA with repeated measures on the last two variables. Figure 2 indicates that paragraphs higher in the page were read for longer [$F(2, 48) = 14.97$, $MSE = 544.89$, $p < .001$, $\eta^2_p = .34$] and the first half of the paragraph was read for longer than the second half of the paragraph [$F(1, 24) = 115.74$, $MSE = 640.56$, $p < .001$, $\eta^2_p = .83$]. There was also a significant Paragraph $\times$ Paragraph Half interaction, $F(2, 48) = 3.94$, $MSE = 231.88$, $p < .05$, $\eta^2_p = .14$. Simple effects analysis found significant effects of Paragraph at both Half 1 and Half 2 and significant effects of Paragraph Half at all 3 paragraphs (for space reasons, and because these simple main effects simply confirm the effects of both paragraph-on-page and half-within-paragraph, we will not report the statistics).

In both texts there were 5 pages with only 2 paragraphs, therefore, the reading times for the third paragraph were calculated using fewer pages than the times for the first and second paragraph. To ensure this uneven weighting did not bias the results all analyses were also conducted using only pages with 3 paragraphs or more. This produced an identical pattern of results and significance.

First pass time data were very short indeed, reflecting a tendency among skim readers to glance across an entire page before settling to read. First pass times were longer for the first half of a paragraph ($M = 3.67$ ms/word, $SD = 3.17$) than the second half of a paragraph [$M = 1.30$ ms/word, $SD = 1.15$, $F(1, 24) = 24.13$, $MSE = 9.69$, $p < .001$, $\eta^2_p = .50$]. No other effects or interactions were significant.

Reading Times per Page

Reading time per word was calculated for each page using the number of *visible* words. For each condition times were averaged across text and presented in Figure 3. There was a clear trend in all 3 conditions to spend more time on the earlier pages in the text. For each condition a 2 (Text: ADHD or Hypnosis) $\times$ 2 (Presentation Order: First or Second) $\times$ 2 (Text Part: First three pages or Last three pages) mixed ANOVA with repeated measures on the last factor was conducted on participants' reading time per word.

In the Skim condition the first part of the text was read for longer than the last part of the text, $F(1, 28) = 60.27$, $MSE = 3799.41$, $p < .001$, $\eta^2_p = .68$. There was also a main effect of Text Part in the Half 1 condition, $F(1, 12) = 15.31$, $MSE = 3498.16$, $p < .01$, $\eta^2_p = .56$, but there was also an interaction between Text Part and Text, $F(1, 12) = 19.80$, $MSE = 3498.16$, $p < .01$, $\eta^2_p = .62$. Simple effects indicated that more time was spent reading the first part of the Hypnosis text ($M = 277.87$, $SD = 50.55$) than the first part of the ADHD text [$M = 177.78$, $SD = 57.81$; $F(1, 12) = 13.01$, $p < .01$,

$\eta^2_p = .52$]. For the last part of the text more time was spent reading the ADHD text ($M = 188.99$, $SD = 49.66$) than the Hypnosis text [$M = 103$, $SD = 71.47$; $F(1, 12) = 15.57$, $p < .01$, $\eta^2_p = .57$]. The difference between the first and last part of the text was significant for the Hypnosis text, $F(1, 14) = 36.75$, $MSE = 679.74$, $p < .001$, $\eta^2_p = .72$. In the Half 2 condition the first part of the text was read for longer than the last part of the text, $F(1, 12) = 13.40$, $MSE = 5489.48$, $p < .01$, $\eta^2_p = .53$.

The mean number of pages visited was 8.97 ($SD = 2.19$) in the Skim condition, 10.56 ($SD = .81$) in the Half 1 condition and 10.38 ($SD = 1.15$) in the Half 2 condition. Number of page revisits in the Skim condition was $M = 1.81$ ($SD = 3.74$). After excluding these revisits 28 out of 32 participants read the pages in the same order they were presented on the screen.

Discussion

We found no support for our hypothesis that reading the first halves of paragraphs would be more effective than reading the second halves. Instead, in both Half-paragraph conditions, memory for meaning was not significantly worse than in the Skim condition (although there was a non-significant tendency in that direction). Instead of supporting our hypothesis that skim readers will benefit by having discretion over the areas of text to which they allocate attention, these data support the idea of an upper limit on this benefit. Furthermore, in conjunction with the findings of Experiment 1, these data suggest that a strategy of spreading time allocation across the extent of a text, by sampling each paragraph, will be an effective strategy.

The eye-tracking data and measures of task performance while skimming were in line with our predictions from the satisficing model. In the Skim condition participants spent more time reading the first half of each paragraph than the second

half. This supports the idea that skimmers satisficed when reading a paragraph and moved to the next paragraph when the information gain was not sufficiently high.

Eye-tracking also showed that the text at the top of the page was more likely to be read than at the bottom of the page. Further, as in Experiment 1, there were also longer reading times for pages earlier rather than later in the text. This pattern of findings is consistent with our supposition that skim readers will prefer beginnings to ends of documents and document-parts. How to explain the preference for first-parts at the level of document, page and paragraph is something to which we return in the general discussion.

Experiments 3A and 3B

Experiments 3A and 3B aimed firstly to replicate the improved memory for the meaning of the text after skimming that we discovered in Experiment 1 (but not Experiment 2). The second goal was to investigate aspects of website design, where some of the most important applied issues in design-for-skimming surely reside. In particular, this next experiment studied whether alterations in the ease of navigation and the linkages between the pages influenced the effectiveness of skimming. Finally, because access to separate pages on the web is often organized hierarchically, we feel it allows more natural, applicable versions of our “Half-text” comparison conditions.

Duggan and Payne (2001) demonstrated that a small cost to access a text could have large implications for the representation formed. They found that when participants were required to click on a button to access a page of instructions they chose to read a larger amount of text per access than when the text was always visible. (This finding mirrors the effect of travel time on patch residence in foraging research related to Charnov’s marginal value theorem; Charnov, 1976). The decision to read larger chunks improved the situation model - see Salmerón, Cañas, Kintsch and

Fajardo (2005) for further evidence that the linkages between pages can affect the comprehension of a text.

On the Web, pages are linked in unpredictable ways (Park & Thelwall, 2003). To test whether the links between pages could influence the extent to which skimming facilitated memory for text meaning we constructed two alternative Website structures. In Experiment 3A the pages of the text were organized in a hierarchy and readers could only navigate up and down through the levels of the hierarchy rather than within the individual levels. The organization of pages in Experiment 3B was also hierarchical and very similar, however, crucially we increased the number of links between pages enabling readers to navigate within as well as between each level and, thus, move between the pages more readily. We predict that the comprehension benefits from skimming will be dependent on the ease with which readers can navigate through the text.

Experiment 3A Method

Participants

The participants were 12 male and 20 female students from the University of Manchester. The mean age was 22.19 years ($SD = 2.18$). They received £5 (U.S. \$9.79) for taking part in the study.

Materials and Procedure

Excluding the first page, the content of each text was subdivided into 4 distinct topics. A topic was at least 2 pages in length and included only consecutive pages. Each topic was then given a heading and the pages were organized into a hierarchical structure. Some of the page headings were also rewritten slightly to increase the extent to which they described the contents of the page. Figure 4 shows the tree structure and page headings from the ADHD text.

The first page of each text acted as the homepage at the top level of the menu structure. The next level of the tree structure contained the first page of each topic, labelled using the four topic headings. The bottom level contained the remaining pages. The top page linked to the 4 pages on the middle level and each page on the middle level linked to the other pages within that topic (all on the bottom level). There were no links between the top and bottom levels nor were there any links between pages on the same level. All links were bidirectional.

In the previous experiments there were 11 buttons that referred to the different pages. In this experiment these buttons were removed from the interface and replaced with buttons that linked to pages in the level below in the tree structure. These buttons were labelled with the corresponding page heading (for the first page in a topic this was the same as the topic heading). Clicking on a button caused the appropriate page to be displayed on screen, all buttons were then removed and replaced with buttons for any pages lower in the hierarchy that were linked to the newly displayed page. To navigate up the menu structure an additional button labelled “Back” was included. Clicking on this took the participant up one level in the tree structure.

Prior to reading, the hierarchical structure of the pages was explained to the participants, as was the functionality of the back button. Participants in the Skim condition were instructed that there were 11 pages in total and that they could navigate freely between the pages as they chose. As in the previous experiments all participants were told they would be tested on their understanding of all 11 pages. In the Half-text condition at the homepage the link buttons to two of the topics were disabled. In the Half 1 condition participants were only allowed access to the first and third topic (ADHD text pages 1, 2, 3, 7, 8, 9; 1669 words; Hypnosis text pages 1, 2, 3, 4, 8, 9; 1585 words). In the Half 2 condition they were only allowed access to the

second and fourth topic (ADHD text pages 1, 4, 5, 6, 10, 11; 1831 words; Hypnosis text pages 1, 5, 6, 7, 10, 11; 1716 words). Because the first page was included in both halves of the text, a few target sentences also appeared in both halves. Across both texts in the first half there were 18 Important, 18 Unimportant, 18 Inference, 26 True and 28 False sentences. In the second half there were 19 Important, 21 Unimportant, 19 Inference, 32 True and 27 False sentences.

All other aspects of the method were the same as in Experiment 1.

Experiment 3A Results

Memory for Meaning

The d' scores were higher for the ADHD text ($M = 1.16$, $SD = .48$) than for the Hypnosis text ($M = .86$, $SD = .53$), $t(31) = 3.27$, $p < .01$. There was no effect of Presentation Order ($t < 1$). The data were analysed using the same 5-way ANOVA as in Experiments 1 and 2. Table 4 shows the means and standard deviations for hit rate, false alarm rate, d' and C scores.

Sensitivity. Sensitivity was lower for Unimportant sentences than for Important or Inference sentences, thus the main effect of Sentence Type was reliable, $F(2, 48) = 8.56$, $MSE = .91$, $p < .01$, $\eta^2_p = .26$. Table 4 shows that sensitivity was marginally better in the Half-text condition than the Skim condition for all 3 sentence types. There was also a 3-way interaction between Text Assignments, Condition Order and Text Half, $F(1, 24) = 9.96$, $MSE = 1.70$, $p < .01$, $\eta^2_p = .29$.

Analysis of simple effects found an effect of Condition Order such that the text skimmed second had a higher score than the text skimmed first. This effect was present when the text was the first half of the hypnosis text, $F(1, 24) = 6.33$, $p < .05$, $\eta^2_p = .21$, and when the text was the second half of the ADHD text, $F(1, 24) = 4.99$, p

$< .05$, $\eta^2_p = .17$. Sensitivity was better for Half 1 than Half 2 when the ADHD text was skimmed first, $F(1, 24) = 5.28$, $p < .05$, $\eta^2_p = .18$.

Bias. Participants were more likely to respond “True” after skimming the text than after reading half the text [Reading Condition, $F(1, 24) = 5.24$, $MSE = .37$, $p < .05$, $\eta^2_p = .18$].

The Reading Condition $\times$ Condition Order interaction was significant, $F(1, 24) = 18.59$, $MSE = .35$, $p < .001$, $\eta^2_p = .44$. Analysis of simple effects found that the C scores were higher for the Half-text condition relative to the Skim condition, $F(1, 24) = 21.79$, $p < .001$, $\eta^2_p = .48$, when the Skim condition was second but not when the Skim condition was first.

Reading Times per Page

Reading time per word for each page is given in Figure 5. Each text was divided in half using different pages, therefore the times for the ADHD and Hypnosis texts are presented separately. Figure 5 shows that in all conditions and for both texts reading was slower for earlier pages. The data were analysed using the same 3-way ANOVA as in the previous experiments.

The first part of the text was read for longer than the last part of the text in the Skim condition, $F(1, 28) = 52.99$, $MSE = 6433.55$, $p < .001$, $\eta^2_p = .65$, the Half 1 condition, $F(1, 12) = 13.30$, $MSE = 12040.83$, $p < .01$, $\eta^2_p = .53$, and in the Half 2 condition, $F(1, 12) = 36.47$, $MSE = 10032.60$, $p < .001$, $\eta^2_p = .75$.

The mean number of pages visited was 9.22 ($SD = 2.34$) in the Skim condition, 5.5 ($SD = .89$) in the Half 1 condition and 5.06 ($SD = .93$) in the Half 2 condition. Number of page revisits in the Skim condition was $M = 7.23$ ($SD = 3.87$). Participants were presented with the pages in a hierarchy rather than a linear structure thus the sequence in which they read the pages was not compared to the original text.

Experiment 3B Method

Participants

The participants were 15 male and 17 female students from the University of Manchester. The mean age was 22.40 years ($SD = 2.64$). They received £5 (U.S. \$9.79) for taking part in the study.

Materials

All pages within a particular topic were linked to each other and to the first page of the next topic in the text. This meant that participants could navigate forward through the topics without returning to the homepage. (In the Half-text condition the pages linked to the next accessible topic.) The “Back” button was replaced with a button labelled “Home” which returned participants to the Home page (the first page of the text).

All other aspects of the method were the same as in Experiment 3A.

Experiment 3B Results

Memory for Meaning

The d' scores were higher for the ADHD text ($M = 1.25$, $SD = .46$) than for the Hypnosis text ($M = .77$, $SD = .62$), $t(31) = 4.86$, $p < .001$. There was no effect of Presentation Order ($t < 1$). The data were analysed using the same 5-way ANOVA as in the previous experiments. Table 5 shows the means and standard deviations for hit rate, false alarm rate, d' and C scores.

Sensitivity. Table 5 shows that sensitivity in the Skim condition was better than in the Half-text condition for the Important and Inference sentences but not for the Unimportant sentences. The main effect of Reading Condition was not reliable ($F < 2$). There was a main effect of Sentence Type, $F(2, 48) = 8.69$, $MSE = .67$, $p < .01$, $\eta^2_p = .27$, this was partly due to the better memory for the Important sentences relative

to the Unimportant sentences. The Reading Condition $\times$ Sentence Type interaction was significant, $F(2, 48) = 4.42$, $MSE = .63$, $p < .05$, $\eta^2_p = .16$. Simple effects indicated that the effect of Sentence Type was significant for the Skim condition, $F(2, 23) = 9.11$, $p < .01$, $\eta^2_p = .44$, but not for the Half condition ($F < 1$).

Planned comparisons found that skimming improved sensitivity for Important sentences relative to reading half the text, $F(1, 24) = 8.57$, $p < .01$, $\eta^2_p = .26$. There was no significant effect of Reading Condition upon sensitivity for the other two sentence types ($F_s < 2$).

Bias. Participants were more likely to respond “True” to the Important sentences and the Inference sentences than to the Unimportant sentences [Sentence Type, $F(2, 48) = 8.50$, $MSE = .38$, $p < .01$, $\eta^2_p = .26$]. Planned comparisons found no significant differences in *C* scores between Skim and Half-text conditions for any of the 3 sentences types.

Reading Times per Page

Reading time per word for each page is given in Figure 6. As in the previous experiments, there was a general trend towards slower reading on earlier pages. The reading time data were analysed in the same way as Experiment 3A.

The first part of the text was read for longer than the last part in the Skim condition, $F(1, 28) = 40.27$, $MSE = 4674.06$, $p < .001$, $\eta^2_p = .59$, the Half 1 condition, $F(1, 12) = 10.61$, $MSE = 19447.18$, $p < .01$, $\eta^2_p = .47$, and in the Half 2 condition, $F(1, 12) = 25.38$, $MSE = 12100.07$, $p < .001$, $\eta^2_p = .68$.

The mean number of pages visited was 8.84 ($SD = 2.05$) in the Skim condition, 5.06 ($SD = 1.00$) in the Half 1 condition and 4.75 ($SD = .77$) in the Half 2 condition. Number of page revisits in the Skim condition was $M = 4.94$ ($SD = 10.93$).

Experiments 3A and 3B Discussion

The results of Experiment 3B successfully replicated the pattern of results for memory for meaning found in Experiment 1. Readers demonstrated a greater understanding of important statements from the text after skimming than after reading through an imposed half of the text. As in Experiment 1, skimmers were no better than half-text readers on understanding of unimportant statements or in ability to make inferences using information from the text.

However, in Experiment 3A there were no differences between the two reading conditions on any of the measures of memory for meaning. After skimming, the d' scores for the Important sentences were lower than in any of the other three experiments. Thus, the lack of an effect of Reading Condition was presumably due to deterioration in the representation constructed by the skimmers rather than any improvement in the Half-text reading condition. This suggests that the reduction in the number of linkages between the pages in Experiment 3A limited the effectiveness of skimming.

General Discussion

This article addresses two main questions. Can skim readers effectively allocate their limited time? And if so, by what strategy?

The reported experiments have demonstrated that when readers have a limited amount of time they are indeed able to gain a greater understanding of a text by skimming rather than by reading linearly through an imposed half of the text. Skimming did not reliably aid memory for less important details, nor did it facilitate inferences about information from the text. However, skimming did allow improved memory for the text's most important ideas, relative to reading part-texts.

To our knowledge, this is the first demonstration of such an advantage for skimming under controlled laboratory conditions. The selective advantage in the

memory for important ideas has not been observed in previous studies of skimming (e.g. Carver, 1984; Dyson & Haselgrove, 2000; Kiwan, Ahmed & Pollitt, 2000; Masson, 1982). The result does build on one published trend, reported by Masson (1982), for skimmers to categorize less important statements more slowly during tests of memory for the meaning of a text. This pattern of results was not replicated in the response times for the experiments here and they are not included for brevity. But, note, we emphasised to participants that they should prioritise accuracy over response times during the recognition tests.

Previous studies have typically compared the decline in performance on different components of comprehension as reading rate increased. By contrast, our study has investigated the *improvement* in comprehension that readers can achieve when allowed to allocate time as they choose rather than reading linearly through the text. The higher memory-for-meaning scores for important sentences in the skim conditions indicate that skim readers successfully allocated their time to more important material within the text.

However, the results from Experiment 2 suggest that this may be primarily due to reading a more distributed selection of the text rather than specifically identifying more important elements. In that experiment the comparison condition read through half of each paragraph at normal speed and, for important sentences, the memory for meaning was equivalent to that in the skimming condition.

Nevertheless, this finding is consistent with our claim that skimming is adaptive. We suggest below that the rate of information gain within a section of text may be subject to diminishing returns. If this is true, leaving a section of text before it has been completely read should, on average, improve the rate of information gain when under time pressure. Ultimately, when skimming, readers must trade off the cost

of exploring the text to find the most important elements against the gain from exploiting (reading) the text (Reader & Payne, 2007).

Turning, then, to the question of skim readers' strategy, we suggest that time-pressured readers achieve this improved understanding of the main elements of the text by *satisficing*. That is, they stop reading one section of text and move forward to the next when the rate at which they are gaining information drops below a threshold. In this way, a gist representation of the main ideas within the text is constructed. In all the experiments we collected reading-process data that support the use of this satisficing strategy.

First, consider the page-to-page data. Typically, in Experiments 1 and 2 participants very rarely revisited pages. Moreover, when these revisits were excluded from the data almost all the participants worked their way linearly through the text reading each page in sequence. This suggests that a simple strategy of reading linearly then skipping over some text accurately characterises the skimming behaviour of readers in these experiments.

Reading time per page shows that skim readers on average spend more time on earlier than later pages. Furthermore, the eye-tracking data in Experiment 2 confirm that this pattern of preference for beginnings is recursively true at the level of pages and paragraphs – skim readers allocate more time to the first paragraph on each page, and to the first half of each paragraph.

How might one explain this simultaneous preference for beginnings at different levels? Preference for the first half of paragraphs emerges naturally from a satisficing strategy if one assumes that the text patch about which a stay-or-go decision is being made is the paragraph: all paragraphs are begun, and most will fall below a demanding threshold imposed by severe time pressure.

The allocation of more time early in pages and early in documents suggests to us that pages and documents suffer from diminishing returns, in terms of rate of gain of valuable information. We propose two factors that might be responsible for this characteristic gain curve. First, there will be some redundancy of information. The gist of a section of text may be inferable from any paragraph within it, therefore, this information will be extracted from the first paragraph read and not from later paragraphs. In addition, well-written text includes coherence cues to help the reader predict subsequent paragraphs (Strunk & White, 1979). Any repetition of information whether inferential or explicit will increase the likelihood of diminishing returns.

Second the value of information may decline, because of writers' deliberate placing of more important information early in paragraphs (Kieras, 1980), early in sections and early in documents. (The distinction between important and less-important information points to an interesting aspect of information foraging theory – to understand the gains in utility it is sometimes necessary to consider the quality of information rather than its mere quantity.)

Finally, the phenomenon of more time spent early rather than late in pages and documents as well as paragraphs might be explained as information foragers judging patches at more than one level – i.e. nested patches – simultaneously. Such an assumption would complicate our satisficing heuristic considerably, but it seems attractive, intuitively. Whenever a paragraph is judged uninteresting, perhaps the section that contains it, and the document that contains that section are also being judged. Mechanistically, this could be made to work in several different ways – for example, by raising the threshold of acceptability for all subsequent paragraphs within a section and document.

Such extensions of the satisficing account of skimming await further research. However, progress has already been made in trying to uncover the mechanisms that may generate patch-leaving decisions in problem solving tasks, where the rate of gain of currency is much easier to measure and track. Payne, Duggan and Neth (2007) investigated problem solvers generating words from sets of letters and choosing to switch from one set to another. They showed that a version of Green's rule (Green, 1984), in which judgment of patches' potential increases with every success and decreases with time between successes could explain most decisions to switch tasks (see also Hutchinson, Wilke & Todd, 2008). Such a paradigm could be extended to test the possibility of foragers judging nested patches.

This article has focussed on skim reading as a means to understand a text in limited time. We contend that this is an important kind of skim reading, from paper and electronic sources. However, it is not the only purpose to which skimming might be put. For example, websites may be skimmed just to make a first judgement about whether they contain anything of interest at all – a “sampling” or exploration phase in advance of any exploitation. Nielsen (2006) contends that when browsing the Web readers typically skim a page in an “F-shape” – a few horizontal eye movements are made before scanning vertically along the left hand side of the page. Although this phenomenon was not observed in Experiment 2 and awaits controlled study, it highlights the potential for readers to adapt their strategy according to the environment. For example, the sheer amount of text available on the Web, the greater variability in text content or the richness of the medium (graphics, adverts, multimedia) could all have made sampling an adaptive strategy for the users observed by Nielsen.

Experiments 3A and 3B demonstrated the potential impact of another aspect of the Web upon skimming behaviour. These experiments showed that the better memory for important sentences after skimming was dependent on the linkages between pages and, thus, the ease with which participants could navigate through the text. One practical implication from this finding is that Website designers should be careful to maximise the flexibility of navigation within the site. Links between pages should allow users to either maintain their current narrative or skip to the beginning of other sections within the site. Although users did typically open most pages in the limited time available, in Experiment 3A the cost of this navigation appeared to impair the representation formed from the content of the text. Therefore, being able to navigate a site is not enough, time-pressured readers need to be able to *readily* skim between pages.

These experiments have demonstrated that skimming can be an effective method to grasp the main points of a text. However, given the increasing temptation to skip text that is afforded by the Web, it is also important to recognize the advantages of reading text in its entirety. Where a deeper understanding of the text is required, it will sometimes be necessary not only to read the most important elements but also the micropropositions that set the context and provide coherence. The value of slower, more extensive reading can be seen not only in controlled studies of skimming (e.g. Masson, 1982) but also within anecdotal accounts of online reading patterns (Carr, 2008).

References

- Anderson, N.H. (2004). Methodology and statistics in single-subject experiments. In H. Pasher, S. Yantis, D. Medin, R. Gallistel, & J. Wixted (Eds.) *Stevens Handbook of Experimental Psychology, 3rd Edition*, Vol. IV. New York: Wiley.
- Carr, N. (2008). *Is Google making us stupid?* Retrieved July 2008 from <http://www.theatlantic.com/doc/200807/google>.
- Carver, R. P. (1984). Rauding theory predictions of amount comprehended under different purposes and speed reading conditions. *Reading Research Quarterly*, 19 (2), 205-218.
- Charnov, E. L. (1976). Optimal foraging: the marginal value theorem. *Theoretical Population Biology*, 9, 129-136.
- Duggan, G. B. & Payne, S. J. (2001). Interleaving reading and acting while following procedural instructions. *Journal of Experimental Psychology: Applied*, 7, 297-307.
- Duggan, G. B. & Payne, S. J. (2006). How much do we understand when skim reading? *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems, CHI '06, Montreal, Canada*, 730-735. New York: ACM Press.
- Dyson, M. C., & Haselgrove, M. (2000). The effects of reading speed and reading patterns on the understanding of text read from screen. *Journal of Research in Reading*, 23 (2), 210-223.
- Edict text analyser. Retrieved February, 2009, from (<http://www.edict.com.hk/textanalyser/>)

- Green, R. F. (1984). Stopping rules for optimal foragers. *American Naturalist*, 123, 30-43.
- Hutchinson, J. M. C., Wilke, A., & Todd, P. M. (2008). Patch leaving in humans: Can a generalist adapt its rules to dispersal of items across patches? *Animal Behaviour*, 75, 1331-1349.
- Kieras, D. E. (1980). Initial mention as a signal to thematic content in technical passages. *Memory and Cognition*, 8 (4), 345-353.
- Kintsch, W. (1998). *Comprehension: A paradigm for cognition*. New York: Cambridge University Press.
- Kiwan, D., Ahmed, A., & Pollitt, A. (2000). The effects of time-induced stress on making inferences during text comprehension. *European Conference on Educational Research, Edinburgh, September, 2000*.
- Liu, Z. (2005). Reading behavior in the digital environment. *Journal of Documentation*, 61 (6), 700-712.
- Masson, M. E. J. (1982). Cognitive processes in skimming stories. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 8, 400-417.
- Morkes, J., & Nielsen, J. (1997). *Concise, SCANNABLE, and Objective: How to write for the Web*. Retrieved July 2008 from <http://www.useit.com/papers/webwriting/writing.html>.
- Nielsen, J. (2006). *F-shaped pattern for reading Web content*. *Alertbox*, April 2006. Retrieved July 2008 from http://www.useit.com/alertbox/reading_pattern.html.
- Park, H. W., & Thelwall, M. (2003). Hyperlink analyses of the world wide web: A review. *Journal of Computer-Mediated Communication*, 7 (4).

- Payne, S. J., Duggan, G. B., & Neth, H. (2007). Discretionary task interleaving: Heuristics for time allocation in cognitive foraging. *Journal of Experimental Psychology: General*, 136 (3), 370-388.
- Payne, S. J., & Reader, W. R. (2006). Constructing structure maps of multiple on-line texts. *International Journal of Human-Computer Studies*, 64, 461-474.
- Pirolli, P., & Card, S. K. (1999). Information foraging. *Psychological Review*, 106 (4), 643-675.
- Reader, W. R., & Payne, S. J. (2007). Allocating time across multiple texts: Sampling and satisficing. *Human-Computer Interaction*, 22, 263-298.
- Salmerón, L., Cañas, J. J., Kintsch, W., Fajardo, I. (2005). Reading strategies and hypertext comprehension. *Discourse Processes*, 40 (3), 171-191.
- Strunk, W., & White, E. B. (1979). *The Elements of Style*. New York, NJ: Macmillan.

Table 1.

Experiment 1: Hits, false alarms and measures of sensitivity (d') and bias (C) for all 3 sentence types after skim reading or after reading half the text.

Sentence	Measure	Skim		Half-text	
Type		M	SD	M	SD
Important	Hit Rate	.83	.16	.67	.21
	False Alarm Rate	.28	.17	.29	.21
	C	-.18	.35	.06	.44
	d'	1.67	.74	1.13	.90
Unimportant	Hit Rate	.70	.20	.61	.21
	False Alarm Rate	.40	.21	.32	.20
	C	-.14	.47	.10	.52
	d'	.92	.80	.86	.66
Inference	Hit Rate	.77	.22	.68	.20
	False Alarm Rate	.31	.19	.27	.23
	C	-.12	.49	.08	.50
	d'	1.38	.82	1.21	.86

Table 2.

Experiment 1. Location accuracy of Important and Unimportant sentences in document as assessed by proportion correct and mean distance in pages of answers from correct page.

Sentence Type	Measure	Skim		Half-paragraph	
		<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
Important	Proportion correct	.22	.17	.23	.21
	Distance to answer	2.29	.89	2.42	1.01
Unimportant	Proportion correct	.26	.17	.29	.21
	Distance to answer	2.38	1.04	2.22	1.18

Table 3.

Experiment 2: Hits, false alarms and measures of sensitivity (d') and bias (C) for all 3 sentence types after skim reading or after reading half the text.

Sentence Type	Measure	Skim		Paragraph		Paragraph	
				Half 1 Visible		Half 2 Visible	
		<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
Important	Hits	.81	.11	.79	.12	.74	.19
	False Alarms	.33	.21	.35	.17	.33	.16
	C	-.21	.35	-.21	.35	-.10	.39
	d'	1.42	.78	1.27	.61	1.21	.71
Unimportant	Hits	.65	.19	.67	.17	.56	.22
	False Alarms	.36	.17	.38	.19	.36	.19
	C	-.02	.39	-.07	.43	.11	.51
	d'	.85	.70	.83	.64	.61	.69
Inference	Hits	.77	.16	.80	.14	.76	.13
	False Alarms	.31	.18	.35	.16	.31	.18
	C	-.12	.38	-.22	.29	-.10	.39
	d'	1.36	.72	1.32	.71	1.31	.57

Table 4.

Experiment 3A: Hits, false alarms and measures of sensitivity (d') and bias (C) for all 3 sentence types after skim reading or after reading half the text.

Sentence Type	Measure	Skim		Half-text	
		M	SD	M	SD
Important	Hit Rate	.79	.19	.77	.15
	False Alarm Rate	.39	.16	.36	.20
	C	-.27	.38	-.20	.43
	d'	1.19	.76	1.22	.68
Unimportant	Hit Rate	.65	.20	.61	.20
	False Alarm Rate	.43	.21	.34	.23
	C	-.11	.43	.08	.51
	d'	.67	.87	.76	.80
Inference	Hit Rate	.75	.15	.78	.13
	False Alarm Rate	.37	.20	.38	.21
	C	-.21	.42	-.23	.43
	d'	1.06	.63	1.19	.68

Table 5.

Experiment 3B: Hits, false alarms and measures of sensitivity (d') and bias (C) for all 3 sentence types after skim reading or after reading half the text.

Sentence Type	Measure	Skim		Half-text	
		M	SD	M	SD
Important	Hit Rate	.77	.15	.65	.24
	False Alarm Rate	.28	.16	.31	.20
	C	-.09	.36	.06	.55
	d'	1.44	.68	1.05	.79
Unimportant	Hit Rate	.56	.23	.59	.22
	False Alarm Rate	.34	.22	.29	.22
	C	.15	.54	.19	.52
	d'	.65	.85	.89	.87
Inference	Hit Rate	.73	.20	.65	.22
	False Alarm Rate	.36	.23	.33	.22
	C	-.13	.51	.03	.55
	d'	1.09	.84	.96	.76

Figure Captions

Figure 1. Experiment 1: Mean time per word spent viewing each page in Skim, Half 1 and Half 2 conditions. For page 6 in Half 1 condition mean calculated from Hypnosis text only and in Half 2 condition mean calculated from ADHD text only. All other means averaged across text.

Figure 2. Means and Standard Errors calculated across participants for time spent viewing each half of a paragraph when skim reading during Experiment 2.

Figure 3. Experiment 2: Mean time per word spent viewing each page in Skim, Half 1 and Half 2 conditions.

Figure 4. Hierarchical menu structure for ADHD text from Experiment 3A. Arrows signify a link between the pages all links were bidirectional.

Figure 5. Experiment 3A: Mean time per word spent viewing each page for ADHD text (top) and Hypnosis text (bottom) in Skim, Half 1 and Half 2 conditions.

Figure 6. Experiment 3B: Mean time per word spent viewing each page for ADHD text (top) and Hypnosis text (bottom) in Skim, Half 1 and Half 2 conditions.