On Measuring the Impact of Hyperlinks on Reading

Gemma Fitzsimmons, Mark Weal, Denis Drieghe
University of Southampton
England, UK
{G.Fitzsimmons, mjw, D.Drieghe}@soton.ac.uk

ABSTRACT

We increasingly spend a vast amount of time on the Web and much of that time is spent reading. One of the main differences between reading non-Web based text and reading on the Web is the presence of hyperlinks within the text, linking various related Web content and Web pages together. Some researchers and commentators have claimed that hyperlinks hinder reading because they are a distraction that may have a negative effect on the reader's ability to process the text. However, very few controlled experiments have been conducted to verify these claims.

In the experiments documented here we utilise eye tracking as a new methodology for examining how we read hyperlinked text. An eye tracker was used to observe participant's behaviour while reading. The results showed that hyperlinked text did not generally have a negative impact upon reading behaviour. However, participants did show a tendency to re-read sentences that contained hyperlinked uncommon (low frequency) words. This suggests that hyperlinks highlight important information to the reader and the hyperlinks add additional content which for more difficult concepts, invites rereading of the preceding text.

Author Keywords
Hyperlinks, Reading, Web Science, Psychology, Human Computer Interaction, Eye movements, Visual cognition

ACM Classification Keywords
H.1.2. User/Machine Systems: Human Information Processing

General Terms
Design; Experimentation; Human Factors; Theory

INTRODUCTION
One of the main differences between reading non-Web based text and reading on the Web is the presence of hyperlinks. There is an on-going debate about hyperlinks and whether they have a negative influence on reading behaviour [3, 14, 16, 17]. In this paper, we report two experiments that examined reading behaviour by recording the eye movement behaviour of participants when they read text presented to them. We will begin by describing and detailing how eye movement recording as a methodology can be utilized to explore human information processing when interacting/reading pages on the Web. We will then describe the previous research regarding hyperlinks and Web reading, and discuss how eye tracking can be used to explore the issues that have arisen.

WebSci'13, May 2–4, 2013, Paris, France.

Culture of the Blue Hyperlink
Displaying hyperlinks in blue has become part of the online culture and most people would recognise a blue word on a Web page as a hyperlink. Hall [8] argues for "ending the tyranny of the button" and giving the user more control over their interaction by using a query-based system where links themselves are virtual entities. Blue hyperlinks are just buttons indicating the presence of links, but this is not necessary in a query-based system where the user can ask for all available links if they wish to query them. However, it may be difficult to move away from the culture of blue text denoting a hyperlink. Without the blue text the users may assume there are no hyperlinks. Nielsen [14] claimed that "the mother of bad design conventions is the decision to make hypertext links blue". In support of this claim, Nielson cites the fact that a sharp blue image is difficult to bring into focus due to the fact that only 2% of the cones on the retina are sensitive to blue and that blue therefore makes for a poor colour choice in terms of usability [6]. Nielsen [14] admits that a move away from blue hyperlinks would improve usability, but he recommends that the convention of the blue hyperlink should remain because most users know that blue text denotes a hyperlink, making it easier for users to recognise which words are hyperlinks more rapidly. This idea is supported by research on automatic attention that suggests that when a user consistently searches the same environment for similar information that is always represented in a similar way, the processing becomes automatic [24, 25]. This could mean that in the case of hyperlinks, blue text is automatically processed as being a hyperlink because blue text tends to always be a hyperlink. This research all suggests that hyperlinks are processed differently to normal text and suggests that reading hypertext on the Web is different from normal reading. Denoting a hyperlink with blue text is the known convention for hyperlinks and this has an influence on how people perceive hyperlinks and interact with them.

There is on-going debate about whether in-text hyperlinks hinder reading. Carr [3] suggested that hyperlinks within the text are a distraction and therefore hinder comprehension of the text. Carr suggests that having to evaluate hyperlinks and navigating a path through them is mentally demanding and is an extraneous task to the act of reading itself.

Another suggestion as to why hyperlinks may cause a disruption to reading may be due to saliency of objects on the display attracting the users' attention. Simola, Kuisma, Oorni, Uusitalo and Hyönä [26] found that salient advertisements on a Web page, such as those with motion, can attract the user's attention and disrupt reading.

There has also been research into learning from electronic texts that suggest that hyperlinks do attract attention to them and that it does not disrupt reading, but actually assists in the retention of the hyperlinked word [16].

Therefore, the hyperlinks and reading research is divided on whether hyperlinks are beneficial or a hindrance to reading. Although reading could be said to be one of the most fundamental tasks we engage in while using the Web, up to this point, very little research has investigated how the presence of hyperlinks influence reading behaviour on the Web. This area is important to investigate because it has an impact on how hyperlinks are generated on Web pages. If hyperlinks are automatically generated, for example by cross-referencing documents in Wikipedia [12], with no human authored intentionality, it is important to understand how we read hyperlinked text in order to make sure that the function of reading is not unnecessarily disrupted.

There have been a number of studies exploring browsing behaviour on Web pages which have found that users browse in an 'F' shaped pattern as observed by Nielsen [15] whereby users focused on the left hand side of the Web page making longer vertical scans at the top of the page and longer vertical scans down the page on the left hand side later (similar in shape to the letter F). However, this cannot be generalised to reading for comprehension where readers need to read the majority of text on the screen, not just browse it. The same issue arises in studies exploring how users search for information on Web pages. Pearson and van Schaik [17] found in a visual search task that participants responded more rapidly to blue hyperlinks than red hyperlinks. We cannot generalise from these findings how reading behaviour may be affected by hyperlinks. With that in mind, we need a new methodology to explore this problem.

EYE TRACKING METHODOLOGY

Eye Movement Research

When we create experiments to explore how a user interacts with a task and processes the information, it is difficult to determine exactly what the user is doing during the task. One way to find out would be to ask the individual. However, there are a number of problems with this method. Firstly, the experimenter would have to wait until after the task had finished or risk interrupting the task and confounding the results. Another problem is that individuals have a tendency to rationalise their behaviours, making self-report unreliable [5]. Recording eye movements is a more objective way of collecting data about behaviour and a number of studies have shown that eye movements provide an unobtrusive, real-time behavioural index of visual and cognitive processing [10, 20, 21].

The Biology of the Eye

Many of the studies described above used indirect measures to explore the influence of hyperlinks on reading. For example, they use response accuracy to artificial tasks and reaction times to infer how hyperlinks affect user behaviour when engaging with the Web. However, the recording of eye movement behaviour enables the researcher to explore the cognitive processes of the online reader in detail.

When we make eye movements, the movements are called saccades. In between these saccades our eyes are relatively still, which is called a fixation. We take in visual information during fixations and vision is said to be suppressed during saccades to avoid seeing a blur or smear [11].

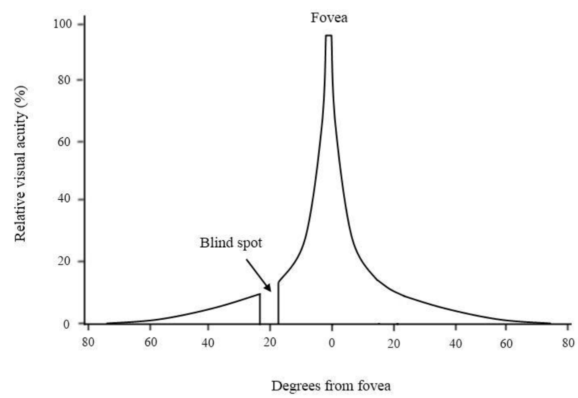

Figure 1. Visual acuity as a function of position on the retina.

Saccades are necessary due to the anatomy of the eye and the retina. The retina contains many photoreceptors called rods and cones. Cones are necessary for high visual acuity. The majority of the cones in the retina are in an area called the fovea which covers about 1° degree of visual angle on either side from the fixation point. Moving beyond the fovea, there is a large reduction in the number of cones and therefore a high acuity drop-off (see Figure 1).

In order to read, we need the words to be mostly positioned on the high acuity fovea to be able to process them. As a result, we need to move our eyes so that the fovea can be utilised to gain the most visual information while reading. There is considerable variability in how far a saccade moves (average 7-9 characters, range 1-15 characters) and how long a fixation lasts (average 200-250ms, range typically 100 - over 500 ms). This variability does not only differ between individuals, but also within an individual's own eye movements. Much of this variability can be related to how difficult the currently fixated text is to process (for reviews see Rayner [20, 21]). Rayner and Pollatsek [22] found that the more difficult the text the longer the fixations and the shorter the saccades and more backward-directed eye movements (regressions) are made to re-read information.

Eye Tracking Technology

The eye tracker used for the current experiment was the SR-Research EyeLink 1000, which records eye position every millisecond and has an accuracy of 0.25º - 0.5º of visual angle. By using infrared light directed at the participant and a camera to record the light reflected back, the eye position can be calculated. The eye position is calculated by asking the participants to look directly at points on the screen and recording the position of the pupil and the position of the corneal reflection (the first Purkinje image, which is the reflection from the outer surface of the cornea) when looking at these points. An algorithm is then used to calculate the eye position across the entire display.

Eye Movement Measures

There are a number of eye movement measures that have been developed to investigate and understand reading behaviour [20]. Only those relevant to the current experiments are listed here.

- *Skipping probability* is the probability that the target word was skipped on first-pass reading. Skipping rates show the ease of processing a word. If a word is very easy to process than it may be skipped completely in the first pass of reading.
- *First fixation duration* is the duration of the first fixation on a word. It does not matter if the target word was fixated once or more than once, only the first fixation is used. This measure is used to examine the early effects of the target word on eye movements. If the first fixation durations are short then the word is easier to process than if the first fixation durations are long [20].
- *Single fixation duration* is where the reader made exactly one first-pass fixation on the target word. Single fixation durations are similar to first fixation duration in that they show the ease of processing the target word [20]. However they only include occasions where the target word was fixated once in the first pass of reading, giving a good estimate of the difficulty or ease in processing that particular word.
- *Gaze duration* is the sum of all first-pass fixations on the target word before moving to another word. Longer gaze durations indicate that the participant is having difficulty processing the information and may make more than one fixation of the target word if needed [20].
- *Go-past times*, which are the accumulated time from when a reader first fixated on the target word until their first fixation to the right of the target word. This measure includes any regressions the reader made before moving forward past the word. Go-past times are used to examine when the reader has difficulty integrating a word and has to go back and re-read the preceding content thus increases the fixation times in the Go-past times measure.
- *Total reading time* is the sum of all fixation durations on the target word regardless of whether this reading happed during first pass or later. This measure examines any disruption seen in reading the target word. Again, if the total reading time is longer in one condition compared to another it suggests that the text in that condition was more difficult to process.

First fixation duration, single fixation duration, and gaze duration are considered early measures of processing because they indicate any initial processing difficulties the reader may have when they fixate a word because these measures only look at fixations on the first-pass of reading. Go-past times reflect difficulties in integrating a word when it is fixated and include regressive fixations on the prior content before the reader goes past the target word. Go-past times show the cost of overcoming difficulty integrating a word and they are considered a late measure of processing.

Eye Movements and Reading

The eye movement methodology has been used extensively to investigate how individuals process text during reading. As of yet, there has been very little investigation into how we read hyperlinked text, but the eye movement methodology, coupled with this vast amount of previous research into reading and eye movements, can be used to help. There are a number of factors to consider that influence when and where we move our eyes and investigating these factors can tell us a great deal about how text is processed during reading.

Where we move our eyes

The decision of where to move the eyes is largely driven by low-level cues such as word length and space information. If the word to the right of fixation is a long word of for instance 9 characters, the next saccade will be longer than if it were a medium sized 5-character word to the right of fixation [19]. In addition, spaces between words are used as cues to where the word ends and can be used to target saccades. When spaces are removed, reading speeds are reduced substantially by as much as 50% [13].

When we move our eyes

The decision of when to move the eyes is mostly driven by lexical factors and ease of processing [10]. Fixation times can be influenced by a range of lexical variables. These include word frequency, with high frequency words fixated for less time than low frequency words, (e.g., “flower” vs. “orchid”; [9]) and predictability, where highly predictable words have shorter fixation durations than unpredictable words [1]. For example, in the sentence “The baker rushed the wedding cake to the reception.” the word “cake” will have shorter fixation durations than the word “pies” because the former is more predictable from the preceding context. However, when completing a visual search task looking for a target word in a passage of text these lexical effects disappear. This indicates that eye movements during reading are very different to those observed during visual search [23].

Using Eye Movement Methodology to Explore Reading Hyperlinks

The research presented above is only an overview of how the eye movement methodology has been used to explore how we read; for a more in-depth review, see Rayner [20, 21]. The present paper is focused on how we are using eye tracking to explore how we read hyperlinked text and whether hyperlinks influence reading behaviour. In order to examine how hyperlinks affect reading behaviour, we first need to find out if any disruption of reading occurs and decide whether it is due to the hyperlink being a salient colour, or because the blue text of the hyperlink is perceived as being important due to the additional information that it can link to.

Two experiments were conducted to explore this issue. The first experiment explored whether salient, coloured words negatively impact reading behaviour. Before examining what the impact is of hyperlinks coloured in blue in a hypertext environment, it is important to examine the impact of reading a word in blue in a sentence outside of a hypertext context.

If any disruption to reading occurred during the reading of a blue coloured word embedded in a sentence with black coloured words, then it is the saliency of a coloured word that disrupted reading. Therefore blue hyperlinks in text will also disrupt reading due to the saliency of the blue word by itself. In line with previous research which have suggested that hyperlinks disrupt reading behaviour [3, 14], we predicted that the coloured words would be fixated for longer because of the saliency of the coloured word. We also predicted that grey words would be fixated for longer due to their reduced contrast, thereby making them more difficult to process [4].

The second experiment explored whether perceiving the words as hyperlinks impairs reading behaviour. Again based on previous suggestions [3, 14], we predicted that the hyperlinked words would be fixated for longer due to the saliency of the blue words. However, because blue hyperlinks are commonplace in Web pages, the processing may become automatic [24, 25]. This could mean that in the case of hyperlinks, blue text is automatically processed as being a hyperlink without an additional cost because blue text tends to always be a hyperlink. We also implemented a word frequency manipulation in order to explore whether common lexical effects are present in hyperlinked text and to investigate if they are modulated by the word being hyperlinked. We predicted that a frequency effect would still be present for the hyperlinked words. However, because high frequency words are skipped more often than low frequency words [9], the saliency of a hyperlinked word can draw attention to the word regardless of the frequency and increase the chance of fixating the hyperlinked word, reducing skipping. Together, the experiments assessed whether there is a difference between reading coloured words and reading hyperlinks and how this affects reading behaviour.

EXPERIMENT ONE

Method

Thirty native English speakers with normal or corrected-to-normal vision took part in the experiment. Eye movements were measured with an SR-Research Eyelink 1000 system running at 1000Hz (i.e. one sample every millisecond). Viewing was binocular, but eye movements were only recorded from the right eye. The stimuli consisted of single sentences displayed on a single line. The participants’ head was stabilised in a head/chin rest to reduce head movements that could affect the quality of the calibration of the eye tracker.

All characters were lowercase (except when capitals were appropriate) and presented in a monospaced Courier font. The display was 73 cm from the participant's eye, at this distance three characters equal about 1° of visual angle.

Thirty sentences were used and a single target word in each sentence would appear in one of five colours, which correspond to the five experimental conditions (black, blue, green, red or grey; see Figure 2). A counterbalanced design was used in which each participant read all thirty sentences with an equal number from each condition. Participants were instructed to read for comprehension and told that they would occasionally have to answer comprehension questions about the sentences. Comprehension questions were presented after 25% of trials and the accuracy of answering these was high (97.5% accuracy), indicating that participants were reading the text correctly.

At the beginning of each trial the participant had to look at a fixation point on the screen. When the eye tracker registered a stable fixation on the fixation point, the

The weather forecast said it would be sunny in England this weekend.

The weather forecast said it would be sunny in England this weekend.

The weather forecast said it would be sunny in England this weekend.

The weather forecast said it would be sunny in England this weekend.

The weather forecast said it would be sunny in England this weekend.

Figure 2. Different versions of one stimulus from Experiment One. Only one version of the sentence was shown.

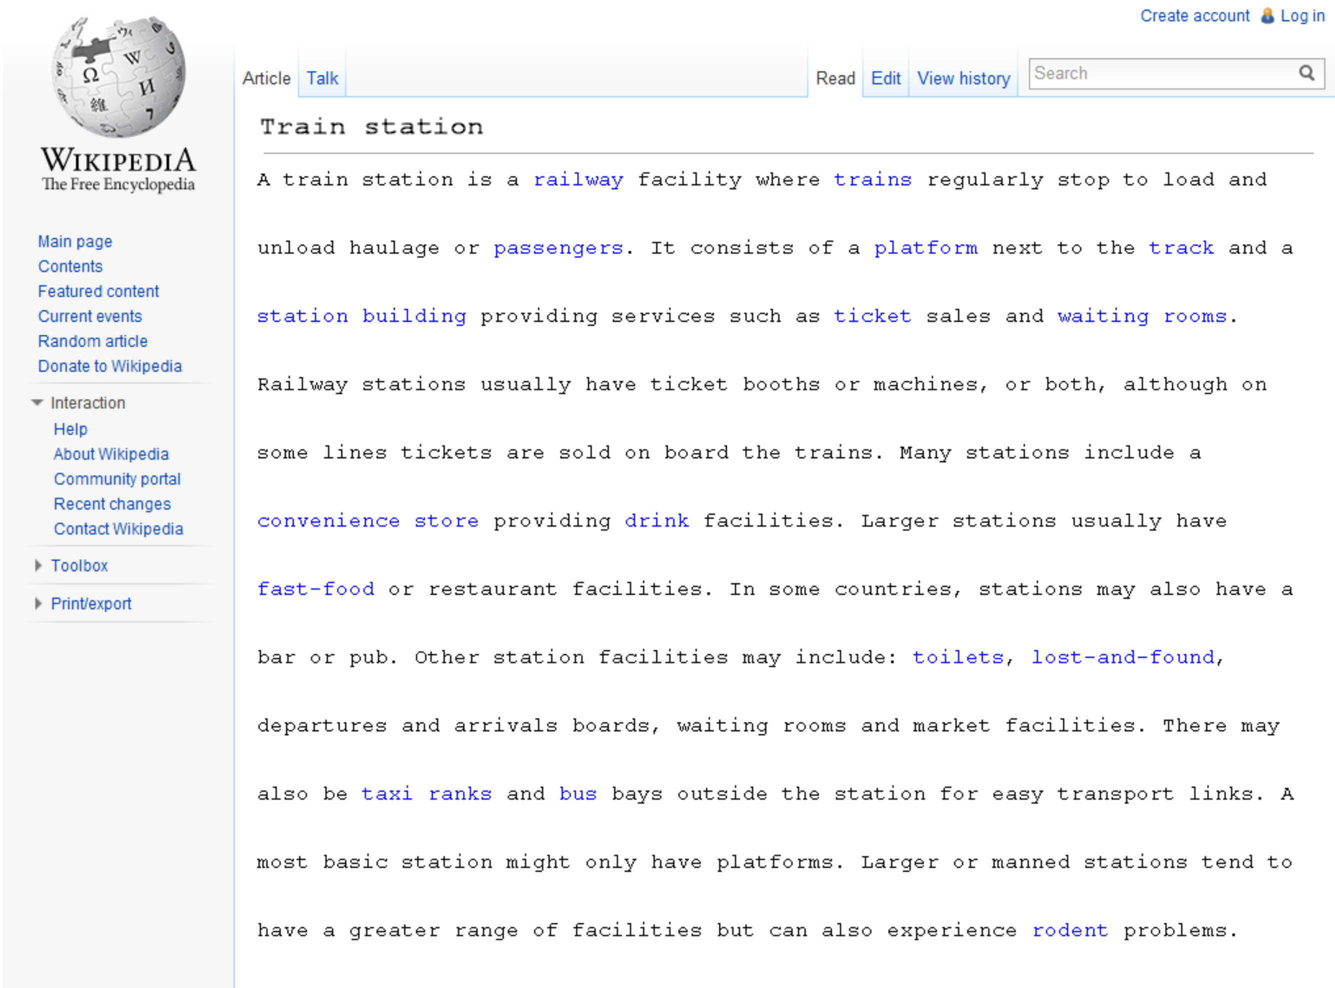

Figure 3. Example stimuli from Experiment Two.

sentence was displayed ensuring that the first fixation fell at the beginning of the sentence. This is to be certain that the reader is starting at the beginning of the sentence and not picking up information from later in the sentence if they do not start at the beginning of the sentence. When participants finished reading the sentence they moved on to the next trial by pressing a button on the response box in front of them. Each participant was given three practice trials to become familiar with the procedure. The experiment lasted approximately 15 minutes.

Table 1. Means of eye movement measures for Experiment One (in ms). Standard deviation in parentheses.

Condition (Colour of Target Word)	Skipping Probability Percentage	First Fixation Duration	Single Fixation Duration	Gaze Duration	Go-Past Times
Black	27 (19)	222 (36)	229 (54)	242 (47)	282 (106)
Blue	14 (15)	219 (49)	225 (80)	244 (55)	280 (95)
Green	11 (14)	234 (41)	243 (46)	260 (57)	339 (115)
Red	8 (10)	209 (36)	210 (44	229 (47)	279 (59)
Grey	13 (13)	244 (40)	247 (51)	279 (50)	337 (96)

Table 2. Statistical tests for Experiment One. ANOVA's and t-tests.

	Skipping Probability Percentage	First Fixation Duration	Single Fixation Duration	Gaze Duration	Go-Past Times
	F (df)	**F (df)**	**F (df)**	**F (df)**	**F (df)**
Main effect of colour	7.06 (4,116)***	4.66 (4,116)**	2.30 (4,96)	6.10 (4,116)***	3.68 (4,116)**
	t (df)	**t (df)**	**t (df)**	**t (df)**	**t (df)**
Black-Blue	3.19 (29)**	0.26 (29)	0.23 (27)	-0.23 (29)	0.09 (29)
Black-Green	3.42 (29)**	-1.65 (29)	-1.47 (27)	-1.73 (29)	-2.13 (29)*
Black-Red	5.18 (29)***	1.71 (29)	0.93 (29)	1.13 (29)	0.13 (29)
Black-Grey	2.80 (29)**	-2.61 (29)*	-2.56 (28)*	-3.43 (29)**	-2.79 (29)**

* p<.05, ** p<.01,*** p<.001

Results

Data cleaning

Eye trackers record a large amount of data (one sample every millisecond) and this data can contain erroneous fixations that are not representative of the dataset. In some cases, these erroneous or outlier fixations will be caused by errors in the eye-tracker (due to the algorithms used to track the eyes); in other cases, the participant may have had a lapse of concentration, leading to very long fixation durations. Regardless of the cause, and because we were interested only in when the participants were paying attention and reading the text, we cleaned the dataset before conducting our statistical analyses. In the current experiments we followed the standard procedures for cleaning our data that have been adopted by the reading research community. Trials where there was tracking loss were removed prior to analysis. Fixations shorter than 80 ms that were within one character of the previous or following fixation were merged and all other fixations shorter than 80 ms or longer than 800 ms were removed to eliminate outliers (4.87% of total dataset). When calculating the eye movement measures data that were more than 2.5 standard deviations from the mean for a participant within a specific condition were removed (<1% of dataset). Data loss affected all conditions similarly.

How does a coloured word effect the participant's reading behaviour?

The means for all of the eye movement measures for Experiment One are listed in Table 1. Participants were significantly less likely to skip a target word when it was any other colour except black (see Table 2). This suggests that the saliency of the coloured target word draws attention to it, making it more likely that participants will fixate it.

There was no statistical difference in fixation time across the conditions except when the target word was grey. The reduced contrast of the target word increased the fixation time on that word compared to any other condition because it was more difficult to process and read [4].

These results suggest that salient words are less likely to be skipped because the saliency draws attention to the word making it more likely to be fixated. There was no difference in fixation times across the different colours except from when the target word was grey which suggests that having a salient coloured word does not affect reading behaviour in terms of processing the coloured word, unless the contrast is reduced making it more difficult to process. When the target word was black, blue, green or red there was no difference in fixation times suggesting that a coloured word does not hinder or help the reading of that word.

Table 3. Means of eye movement measures for Experiment Two (in ms). Standard deviation in parentheses.

	Skipping Probability Percentage	First Fixation Duration	Single Fixation Duration	Gaze Duration	Go-Past Times	Total Reading Time
High/Hyperlinked	43 (23)	216 (34)	214 (35)	227 (36)	298 (123)	261 (54)
High/Unlinked	46 (24)	215 (32)	219 (40)	228 (38)	291 (84)	266 (67)
Low/Hyperlinked	40 (24)	232 (45)	249 (47)	258 (47)	364 (112)	320 (68)
Low/Unlinked	43 (24)	231 (38)	238 (44)	251 (47)	306 (88)	297 (60)

Table 4. Fixed effect estimates for all eye movement measures in Experiment Two.

	Skipping Probability Percentage	First Fixation Duration	Single Fixation Duration	Gaze Duration	Go-Past Times	Total Reading Time
Intercept	-0.32	212***	213***	221***	285***	258***
High/Low Word Frequency	-0.14	22.92***	36.77***	39.36***	67.34***	64.23***
Hyperlinked/Unlinked	0.14	2.76	4.26	6.1	9.68	10.04
Word Frequency x Hyperlinked/Unlinked	-0.03	-5.66	-14.58	-13.73	-51.35*	-33.13*

* $p<.05$, ** $p<.01$, *** $p<.001$

EXPERIMENT TWO

Method

Thirty-two native English speakers with normal or corrected to normal vision took part in the experiment. The procedure was identical to Experiment One, only the stimuli differed. The stimuli in Experiment Two consisted of twenty edited Wikipedia articles (see Figure 3). Eighty target words were embedded in sentences (one target word per sentence) and four sentences were inserted into each Wikipedia article. The target words within these articles were either displayed in blue or black to denote if the word was a hyperlink or not. There was also a word frequency manipulation where the target word is either high or low frequent. The word frequencies were taken from the Hyperspace Analogue to Language (HAL) corpus, which consists of approximately 131 million words gathered across 3,000 Usenet newsgroups. The frequency norms were used to extract both high and low frequency words to create the experimental stimuli. The high frequency words had an average log transformed HAL frequency of 9.91 and the low frequency words has an average log transformed HAL frequency of 5.75 (according to the norms collected in the HAL corpus [2]). In total there are 4 conditions in a 2 x 2 design within participants.

Participants were again instructed to read for comprehension and told that they would be have to answer comprehension questions about the sentences, these appeared after all trials and the accuracy answering them was 94.7%. The experiment lasted approximately 30 minutes.

Results

Data cleaning

The data cleaning procedure and eye movement measures used were identical to that used in Experiment One. Data loss affected all conditions similarly.

How does a hyperlinked word affect the participant's reading behaviour?

The means for all of the eye movement measures for Experiment Two are listed in Table 3.

A linear mixed-effects model (lme) using R [18] was used. The model specified participants and items as crossed random effects.

The significance values and standard errors reported reflect both participant and item variability. These analyses have the advantage that they result in considerably less loss of statistical power in unbalanced designs due to missing values than traditional ANOVA's. This is especially important for fixation times when the target word was skipped often. The *p*-values were estimated using posterior distributions for model parameters obtained by Markov-Chain Monte Carlo (MCMC) sampling. All the patterns observed in the models were identical whether they were run on log-transformed or untransformed fixation durations,

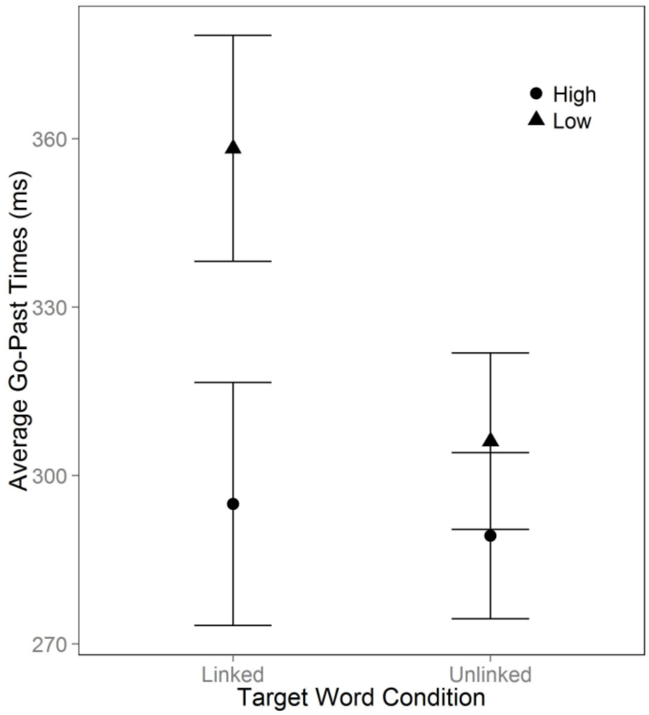

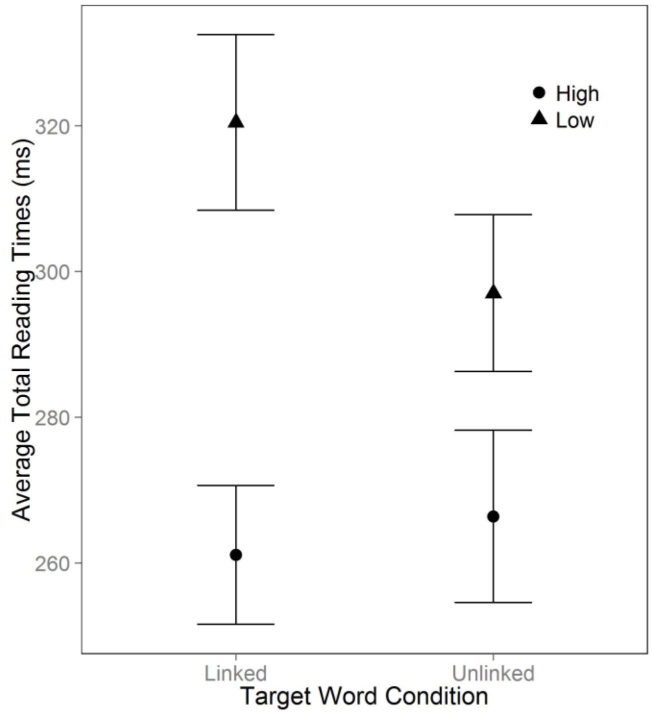

Figure 4. Word Frequency x Hyperlinked/Unlinked interaction for Go-past times and Total Reading Time

allowing us to present the data run on the untransformed fixation durations in order to increase transparency.

As fixed factors, we included the two independent variables: word frequency and whether the target word was hyperlinked or not. All the fixed effects estimates are shown in Table 4.

The effect of word frequency was present in all eye movement measures. The high frequency words were skipped significantly more often than the low frequency words and also had significantly shorter fixation times when they were fixated. This replicates previous experiments that have demonstrated that low frequency words are skipped less often and have longer fixations times because they are more difficult to process than highly frequent words [9].

Replicating the results of Experiment One there were no significant differences between whether the target word was hyperlinked or not in the early eye movement measures (first fixation durations, single fixation durations and gaze durations). This suggests that having a word hyperlinked does not make it any more difficult to process in first-pass reading. However, there was no difference in the amount of skipping observed in Experiment Two. This is especially interesting because in Experiment One when a word was a coloured it was less likely to be skipped. This suggests that there is a difference between the visual appearance of a coloured word and when it is perceived as a hyperlink.

There was a significant difference in eye movement measures between whether the target word was hyperlinked or not, qualified by an interaction with frequency in the later measures of go-past times and total time (see Figure 4). As Figure 4 shows, the low frequency hyperlinked words had significantly longer go-past times and total reading times on the target word compared to when the word was a high frequency linked word or both high and low frequency unlinked words. This interaction shows that when a word is hyperlinked and low frequency it has significantly longer fixation times. This interaction was only present in the later measures which suggest that the low frequency hyperlinked word causes regressive eye movements due to difficulty processing. This means that participants are reaching the low frequency hyperlinked words and rereading the preceding content to re-evaluate it.

CONCLUSION

By running these two experiments together we can be certain that there is a difference between reading text with a coloured word and reading text with a hyperlinked word. Experiment One demonstrated that a coloured word is less likely to be skipped, but that making a word coloured does not negatively impact reading behaviour unless the colour has reduced contrast making it difficult to read as seen when the target word was grey.

Experiment Two demonstrated that there is a difference between a word that is coloured and a hyperlinked word. In Experiment Two there was no reduced skipping of the hyperlinked word that was observed in Experiment One. There was also a significant difference in go-past times and total reading times between whether the target word was hyperlinked or not, qualified by an interaction with frequency. The hyperlinked, low frequency words had longer fixation times in these measures which indicate that the reader had difficulty integrating and processing the low frequency word when it was hyperlinked. As a result, participants were more likely to reread the preceding content to re-evaluate it. A hyperlink on a word indicates that that word is important and implies there is more information behind the hyperlink regarding the word or topic that hyperlinked word corresponds to. When the hyperlinked word is a low frequency word the reader may wonder why that word is hyperlinked and want to re-evaluate the preceding content to make sure that they understood it, or try to decide why it is important.

These experiments have shown that coloured text does not hinder reading, but also that hyperlinks can cause us to reread previous content if the word is a low frequent/difficult word in order to re-evaluate the content. In terms of Web design and layouts, the present results highlight the importance of carefully considering which words are hyperlinked in Web pages. The key lesson here is that Web designers should only hyperlink important words in pages, taking extra caution with words that are uncommon or ones that may be difficult to process. Specifically, Web designers should consider using common words in language in order to facilitate understanding and reduce the need for Web users to reread preceding text surrounding hyperlinks. Another lesson for Web designers that we have found here is that coloured words do not have any negative impact on reading behaviour. This is the case no matter the colour, unless the contrast between the text colour and the background colour is low, as seen in the longer fixation times on the low-contrast grey words in Experiment 1. Therefore efforts made in Web development to avoid using blue as the hyperlink colour and instead using a different colour may have no positive influence for the reader reading the text, but instead make it more difficult for the reader to know what is a hyperlink when they are expecting it to conform to the convention of hyperlinks being denoted in blue.

The present set of experiments represent the first steps in understanding how we read hyperlinked text. A hyperlink is not just a salient word in a passage of text, it denotes that more information that may be relevant lies behind that hyperlink. Enriching hypertext documents with large numbers of links that are automatically generated (as seen in Milne & Written [12]) may cause disruption to reading behaviour. Hyperlinks do not necessarily cause a disruption to reading, but our research here has shown that if you hyperlink a low frequent/difficult word there is disruption to reading behaviour, which if the word does not necessarily need a hyperlink, this disruption could be avoided.

Even though in the current experiments participants only engaged in reading behaviour and did not have to make decisions and click any hyperlinks, there was still a significant difference between reading a coloured word and a hyperlinked word. For the current experiments the aim was to pull apart the process of reading hyperlinked text. By taking away the decision making and clicking ability we could observe the visual effect of how hyperlinks are read.

The results presented as part of these two experiments serve as the foundation for future experimentation using the eye movement methodology to examine reading behaviour on the Web. By basing our future research on the vast amount of research already conducted on eye movements and reading we can build an understanding of how we read hyperlinked text. Future experimentation will expand our experimental task, which is a simplification of live Web behaviour. In future research we aim to explore reading behaviour alongside the navigation and decision making elements that hyperlinked text entails.

ACKNOWLEDGMENTS

This research was funded by the Research Councils UK Digital Economy Programme, Web Science Doctoral Training Centre, University of Southampton. EP/G036926/1.

REFERENCES

1. Balota, D., Pollatsek, A., & Rayner, K. The interaction of contextual constraints and parafoveal visual information in reading. *Cognitive psychology*, 17, 3 (1985), 364–90.
2. Burgess, C., & Livesay, K. The effect of corpus size in predicting reaction time in a basic word recognition task: Moving on from Kucera and Francis. Behavior Research Methods, Instruments, & Computers, 30 (1998), 272-277.
3. Carr, N. G. *The Shallows*. W. W. Norton, New York, NY, 2010.
4. Drieghe, D. Foveal processing and word skipping during reading, *Psychonomic Bulletin & Review*, 15 (2008), 856-860.
5. Ericsson, K. A., & Simon, H. A. *Protocol analysis: Verbal reports as data*. (H. A. Simon, Ed.)*Rev ed* , MIT Press, Bradford Cambridge, MA, 1993.
6. Galitz, W. O. *The essential guide to user interface design: An introduction to GUI design principles and techniques*. John Wiley & Sons, New York, NY, 1997.
7. Gilchrist, I. D., Brown, V., & Findlay, J. M. Saccades without eye movements. *Nature, 390,* 6656 (1997), 130–131.
8. Hall, W. Ending the Tyranny of the Button. *IEEE Multimedia*, 1, 1 (1994), 60–68.
9. Inhoff, W. A., & Rayner, K. Parafoveal word processing during eye fixations in reading: Effects of

word frequency. *Perception psychophysics*, 40, 6 (1986) 431–439.

10. Liversedge, S., & Findlay, J. Saccadic eye movements and cognition. *Trends in cognitive sciences*, 4, 1 (2000), 6–14.
11. Matin, E. Saccadic suppression: A review and an analysis. *Psychological bulletin*, 81, 12 (1974), 899–917.
12. Milne, D., & Witten, I. Learning to link with wikipedia. *Proc. CIKM '08*, ACM Press, (2008), 509-518.
13. Morris, R. K., Rayner, K., & Pollatsek, A. Eye movement guidance in reading: The role of parafoveal letter and space information. *Journal of experimental psychology Human perception and performance*, 16, 2 (1990), 268–281.
14. Nielsen, J. When bad design elements become the standard. *Nielsen Norman Group* (1999), http://www.nngroup.com/articles/when-bad-design-elements-become-the-standard/
15. Nielsen, J. F-shaped pattern for reading Web content. *Nielsen Norman Group* (2006), http://www.nngroup.com/articles/f-shaped-pattern-reading-web-content/
16. Nikolova, O. R. Effects of Visible and Invisible Hyperlinks on Vocabulary Acquisition and Reading Comprehension for High- and Average-Foreign. *Apprentissage des langues et systèmes d'information et de communication*, 07 (2004), 29–53.
17. Pearson, R., & van Schaik, P. The effect of spatial layout of and link colour in web pages on performance in a visual search task and an interactive search task. *International Journal of Human-Computer Studies*, 59, 3 (2003), 327–353.
18. R Development Core Team (2009). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria: R Foundation for Statistical Computing.
19. Rayner, K. Eye guidance in reading: Fixation locations within words. *Perception*, 8, 1 (1979), 21–30.
20. Rayner, K. Eye movements in reading and information processing: 20 years of research. *Psychological bulletin*, 124, 3 (1998), 372–422.
21. Rayner, K. Eye movements and attention in reading, scene perception, and visual search. *Quarterly journal of experimental psychology*, 62, 8 (2009), 1457–506.
22. Rayner, K., & Pollatsek, A. *The psychology of reading*. Prentice Hall, New York, NY, 1989.
23. Rayner, K., & Raney, G. Eye movement control in reading and visual search: Effects of word frequency. *Psychonomic Bulletin & Review*, 3, 2 (1996), 245–248.
24. Schneider, W., & Shiffrin, R. M. Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review*, 84, 1 (1977), 1-66.
25. Shiffrin, R. M., & Schneider, W. Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review*, 84, 2 (1977), 127–190.
26. Simola, J., Kuisma, J., Oörni, A., Uusitalo, L., & Hyönä, J. The impact of salient advertisements on reading and attention on web pages. *Journal of Experimental Psychology: Applied*, 17, 2 (2011), 174–90.