Using Gaze Patterns to Study and Predict Reading Struggles due to Distraction

Vidhya Navalpakkam
Yahoo! Research
Santa Clara, CA
nvidhya@yahoo-inc.com

Justin M. Rao
Yahoo! Research
Santa Clara, CA
jmrao@yahoo-inc.com

Malcolm Slaney
Yahoo! Research
Santa Clara, CA
malcolm@ieee.org

CHI 2011, May 7–12, 2011, Vancouver, BC, Canada.
ACM 978-1-4503-0268-5/11/05.

Abstract
We analyze gaze patterns to study how users in online reading environments cope with visual distraction, and we report gaze markers that identify reading difficulties due to distraction. The amount of visual distraction is varied from none, medium to high by presenting irrelevant graphics beside the reading content in one of 3 conditions: no graphic, static or animated graphics. We find that under highly-distracting conditions, a struggling reader puts more effort into the text — she takes a longer time to comprehend the text, performs more fixations on the text and frequently revisits previously read content. Furthermore, she reports an unpleasant reading experience. Interestingly, we find that whether the user is distracted and struggles or not can be predicted from gaze patterns alone with up to 80% accuracy and up to 15% better than with non-gaze based features. This suggests that gaze patterns can be used to detect key events such as user struggle/frustration while reading.

Keywords
Visual distraction, animated graphics, eye-tracking, reading, user experience, gaze patterns

ACM Classification Keywords
H.5.2 Information Interfaces and Presentation: User interfaces Evaluation/ methodology

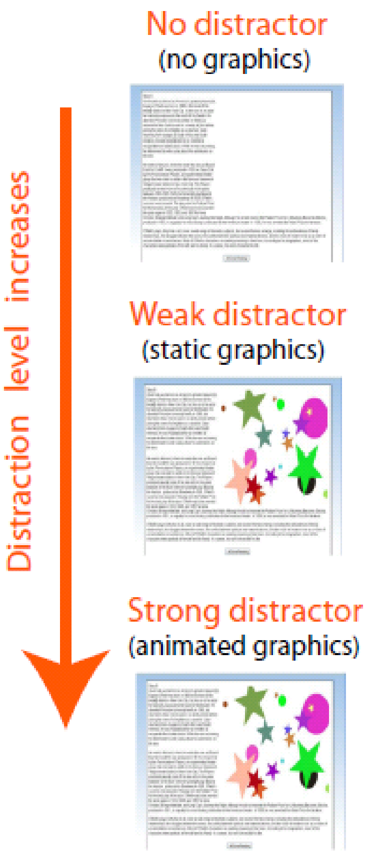

Figure 1: Distraction conditions and sample displays used in the experiment. The animated graphic was similar to the static graphic in color, shape etc., but was blinking and moving randomly.

Introduction

Online environments present several obstacles to readers in the form of visual distraction, including pop-up windows, email alerts, IM, ads, pictures, videos etc. This raises the important question of how visual distraction affects a reader's attention and mental state, and bears consequences for the web and other online reading environments. Indeed, the method in which many websites are monetized, through presenting ads *alongside* content, fundamentally distracts the user. In this eye-tracking study we address the question of how such distraction impacts the user. To do so, we vary the amount of visual distraction from none, medium to high by presenting no graphics, static or animated irrelevant graphics alongside the reading content (see Figure 1). We study the extent to which distraction affects reading time, comprehension, gaze patterns and the reader's experience, measured by the reported interestingness and pleasantness of reading experience. Our analysis shows *how* users are distracted by animated graphics, and which gaze markers are important predictors of a user struggling to read.

Previous studies have analyzed gaze patterns during reading (for a review, see [12]). They show that gaze pauses at some words for 100-500ms (fixations) and jumps between a few words (saccades). While most of these saccades are forward (eye moves left to right on a line, and down to the next line), occasionally, there are backward saccades (regressions). These are thought to reflect confusion and difficulty in comprehension [11, 13]. Gaze patterns are known to vary depending on the age [1], skill-level [10, 13] and reading task difficulty [6, 7]. We extend these studies by testing how visual distraction due to graphics — a widely prevalent factor in online reading environments — affects gaze patterns as well as the reader's mental state.

Most previous work on the effect of distraction on gaze patterns has used the reading-with-text-distraction paradigm [5], using distractors of different color, text style (e.g., italic vs. upright) [8], or semantic meaning than the words in the text [9]. These studies differ from ours in three important ways — they did not study distraction due to graphics, did not vary the distraction-level systematically, nor did they study how distraction affects the reader's mental state.

In a design similar to ours, [2] measured how pictures affect gaze patterns while reading. The study used relevant pictures and ads (both static) and found that the total dwell time is longer and regressions are higher for relevant pictures (presumably in trying to relate the picture to the text). As we are interested in studying the effects of visual distraction, we use irrelevant graphics that are static (less distracting) or animated (more distracting).

An innovation of our study is that we use gaze patterns as a behavioral measurement tool *and* a predictive tool of subjective user experience. Most previous studies measure the effect of age, skill, difficulty level etc. on gaze patterns while reading. Here, we attempt to go beyond measurement of reading behavior to ask whether we can *predict* the reader's mental state, such as whether she struggles and finds the reading experience unpleasant. Identifying these gaze-based markers allows for one to build a user-interface that is informed by objective signals of user experience. Examples of such a corrective approach include gaze-based reading assistants that infer reading difficulty from long pauses on words, for example, and offer remedies such as presenting the meaning of the word to facilitate reading [3, 14]. To our knowledge, no one has used gaze patterns to measure and predict the user's distraction level due to irrelevant graphics, such as widely prevalent on the web.

Here, we use gaze patterns to show how people cope with visual distraction while reading: they struggle to ignore the distraction and work hard to focus on the text. This is reflected in more time on text, increased

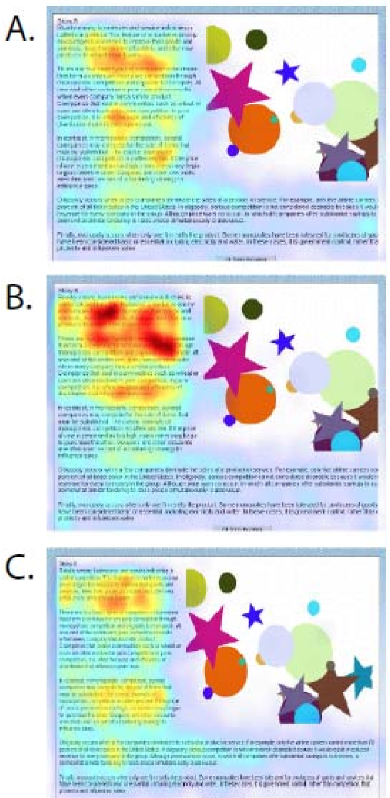

Figure 2: Panels A-C show the heatmaps for the different conditions. Guess which is what? The answer is provided at the end of this paragraph. The hotspots (page regions that are most viewed by users) are shown in red, and the coolspots are in blue.

Answer: the conditions are static, animated and no graphics respectively.

fixation duration, more fixations, more revisits to previously read content and poor pleasantness ratings. We show that the user's mental state — the pleasantness ratings that indicate whether the user struggles or not — can be predicted reasonably well by looking at the gaze patterns alone, demonstrating the utility of gaze as a measurement and predictive tool for understanding reading behavior.

Experiment

Design: The experiment consisted of a within-subject randomized design with 3 factors: no, medium and high distraction created by presenting no, static and animated graphics. Each subject saw 3 essays (from the Test of English Language Fluency), that were randomly paired with one of the graphic types. Each essay consisted of 300-400 words, followed by 5 factual/theme-based multiple-choice questions, and 2 subjective questions where subjects were asked to rate their user experience on a scale of 1-5 for interestingness and pleasantness (e.g., 1 for least pleasant and 5 for most pleasant).

Apparatus: We recorded subjects' gaze patterns during task performance using a Tobii 1750 eye tracker (50Hz sampling frequency), with a 17" LCD monitor, set at resolution 1024x768, at roughly 85 cm viewing distance. We collected a log of eye and mouse movement.

Participants: There were 20 subjects aged 19-60, with normal or corrected vision. Subjects were paid $1 for every correct answer (5 questions per essay x 3 essays). During data cleaning, 3 subjects were excluded for the following reasons: poor calibration (2 subjects did not maintain their head in the correct position), or outliers in fixation duration or number of fixations (3 standard deviations, 1 subject).

Procedure: The study began with a 5-point calibration procedure followed by the task-instruction screen. This was followed by one practice essay paired with animated graphics to help subjects familiarize themselves with the task, types of graphics and the format of the questions in the reading-comprehension test. Next, subjects saw the 3 essays randomly paired with no, static or animated graphics. At the conclusion of the study, subjects were paid based on their performance.

Results

Gaze-based measurement of distraction

We considered 14 metrics in total. The non-gaze metrics are the reading time and comprehension score. The gaze-based metrics are the number of fixations on the graphic, fixation duration on the graphic, time to first visit the graphic, total time on the graphic, number of fixations on the text, fixation duration on the text, saccade length, re-reading, saccade length of re-reading and the total time on text. Apart from the above 12 objective metrics, we also considered the user's self-reported mental state in the form of pleasantness ratings and interestingness ratings of reading experience (on a rating scale of 1-5 for 1 being very bad and 5 being very good).

Figure 2 shows the user interface and heat maps of fixations. It is immediately evident that the conditions differed significantly, with the difference mainly occurring on the text regions of the UI. Figures 3 and 4 show the metrics that differed significantly due to distraction (as computed by one-way repeated measures ANOVA). These are detailed below.

We first report the results from the non-gaze metrics. Subjects reported all the essays to be equally interesting (mean rating of 4.0 out of 5, all stories within a half standard deviation of each other), regardless of the amount of distraction. There were no significant differences in accuracy of reading comprehension; subjects got an average of nearly 4 correct answers per essay across conditions. But there were differences in the reading time. When the distraction is absent (no graphics) or mild (static

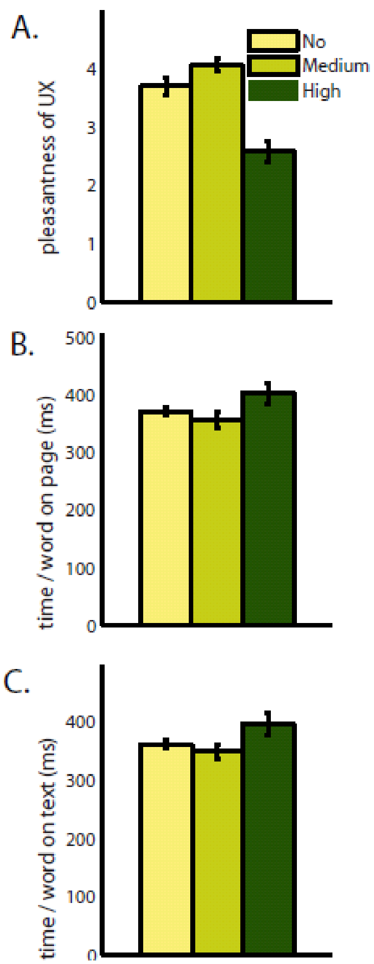

Figure 3: Panels A-C show some metrics that differed significantly when the distraction level increased from none to medium to high.

graphics), users report higher pleasantness levels and spend less time comprehending the text (Figures 3A and 3B). These differences are significant as given by a one-way repeated measures ANOVA test ($F(2,48)=23.08$, $p=0.00$ and $F(2,48)=3.47$, $p=0.04$ respectively).

The results from gaze patterns offer valuable insights on why the reading performance and experience lag in the animated graphics condition. First, the poor reading performance is *not* due to overt attention or fixations on the distractor. Although the animated distractors were always noticed and reported as distracting (based on verbal reports by subjects at the end of the study), they were barely fixated (<5% fixations on the graphics) and overt attention was mainly deployed on the text. There was little time spent on the graphics (<3 seconds) regardless of its type; total dwell time on the graphic did not significantly differ ($F(2,48)=1.15$, $p=0.32$, one-way ANOVA). The differences in total time on page are entirely driven by the time on text.

Second, the animated graphics are seen much earlier than static or no graphics (within 10 seconds of page onset, compared to 30 seconds for the no and static graphic; Figure 4B). This is consistent with studies on attention from cognitive psychology which show that moving objects appear more salient and capture attention more rapidly than static objects [4, 15]. However, they are quickly rejected as irrelevant, and this is reflected by the small fixation duration on animated graphics (250ms compared to 450ms for no and static graphics, $F(2,48)=3.27$, $p=0.05$, one-way ANOVA; Figure 4C). Thus, the animated distractors capture overt attention, but do not sustain it.

The reason for the poor reading performance and experience could be due to covert attention on the distractors using peripheral vision. This is reflected in increased cognitive cost of processing the text, as reflected in longer time on text (and page) and more fixations on text. In particular, the "regressions", revisits to previously read content, which we define as "re-reading," are 30% more frequent with animated graphics (Figure 4A). The difference is significant ($F(2,48)=3.61$, $p=0.03$, one-way ANOVA). No differences were observed in other metrics such as saccade length and saccade length of regressions. Our results show that users struggle to ignore the animated distractors and work hard to focus on the text.

Gaze-based prediction of distraction

Can gaze patterns be used to predict whether the user struggles or not while reading? We fed the 12 metrics discussed earlier as input features to a decision-tree classifier, with the pleasantness of user experience (not pleasant: ratings<2.5, pleasant: ratings>=2.5) as the output variable. Specifically, there were 9 gaze-based metrics and 3 non-gaze metrics (reading comprehension accuracy, time on page, graphics type). To avoid overfitting, we computed the leave-one-out cross validation error as a function of tree height, and selected the tree height that minimized the cross-validation error.

The prediction accuracy using all 12 metrics was 87%. The most predictive metrics were the graphics type (82% accuracy), followed by two gaze-based metrics: the amount of re-reading (75%) and time to first fixation on the graphic (70%). The remaining metrics were less predictive (<70%). Combining the two gaze-based metrics lead to 80% prediction accuracy. In comparison, the non-gaze metrics (time on page and accuracy) were less predictive (74% accuracy).

To test whether gaze patterns are predictive simply because they co-vary with a strong predictor like distractor type, we repeated the above analysis for the strong distractor condition (we focused on this condition due to the higher variability in subjects' experience; in the no/weak distractor condition, the experience was always pleasant, hence little or no variability). We find that non-gaze metrics are no longer informative about reading experience (60%

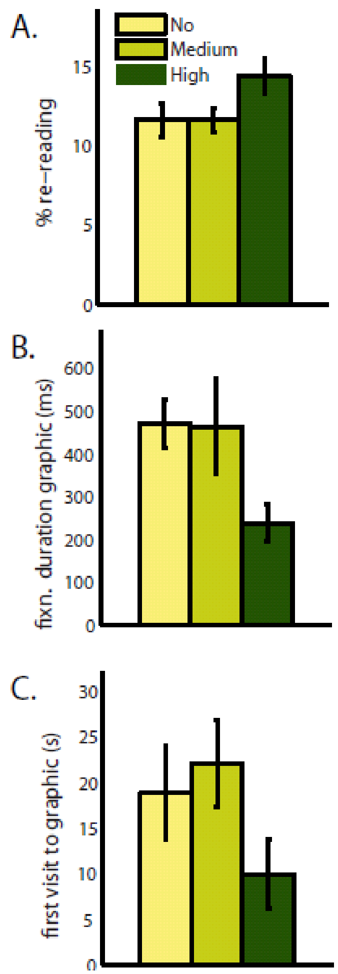

Figure 4: Panels A-C show some gaze-based metrics that differed significantly when the distraction level increased from none to medium to high.

accuracy), scarcely superior to chance. In contrast, gaze patterns continue to predict experience well (72% accuracy). Confirming the above findings, the most predictive gaze-based metrics were the amount of re-reading and time to first fixation on the graphic; these were sufficient to predict UX as well as all 9 gaze-metrics combined. These gaze based metrics can be used in future usability studies as objective quantifiers of subjective user experience. We view this as a promising advance, especially given that technological progress is making eye-tracking more accessible.

Discussion and Conclusion

We use gaze patterns to measure and predict distraction-induced struggles in reading. When the distraction is absent or mild (such as due to static graphics), we find that readers take less time comprehending the text and report a pleasant reading experience. In contrast, under high levels of distraction (such as due to salient, animated graphics), users behave differently — they take much longer to comprehend the text, report unpleasant reading experience, and their gaze patterns reveal important differences. For example, a strong visual distractor causes eye capture. Although users quickly reject it as irrelevant and deploy further eye fixations on the text, it continues to attract attention covertly through peripheral vision, which leads to an increased cost of processing the text. This is reflected in longer time to comprehend the text, higher number of eye fixations on the text, significantly higher amount of revisits to previously read content and poor reading satisfaction levels reported by the user. In other words, the user struggles to cope with the strong visual distraction.

An important contribution of this study is that we show that gaze patterns can be used as a predictive tool to detect distraction-induced struggles in reading. We trained a decision-tree classifier to predict pleasantness of reading experience from various gaze- and non-gaze metrics. We find that gaze-based prediction of user pleasantness level is up to 80% accurate and up to 15% better than non-gaze based prediction. The most predictive signals are the amount of revisits to previously read content and the time to first fixate the distractor. To summarize, we use gaze patterns to measure the effects of distraction, and more importantly predict distraction levels while reading. This suggests that gaze-tracking is useful beyond usability studies to predicting key events such as whether the user is frustrated or not.

Potential applications of gaze-based prediction include developing reading assistants that take corrective measures to decrease visual distraction upon detecting reading struggle. On the web, this would translate to updating page and graphics properties (e.g., ads, pictures, videos) to enable better reading experience and avoid page abandonment due to reading struggle.

Acknowledgements
We thank Robert Dougherty, and the IE group at Yahoo! Research for their valuable feedback, especially Elizabeth Churchill. We thank Bob Moore and Prasad Kantamneni for helping setup the eye tracking study.

References

[1] L. Abel, B. Troost, and L. Dell'Osso. The effects of age on normal saccadic characteristics and their variability. *Vision Research*, 23(1):33–37, 1983.

[2] D. Beymer, P. Orton, and D. Russell. An eye tracking study of how pictures influence online reading. In *Proceedings of the 11th IFIP TC 13 international Conference on Human-computer Interaction-Volume Part II*, pages 456–460. Springer-Verlag, 2007.

[3] R. Biedert, G. Buscher, and A. Dengel. The eyebook–using eye tracking to enhance the reading experience. *Informatik-Spektrum*, 33(3):272–281, 2010.

[4] R. Carmi and L. Itti. Visual causes versus correlates of attentional selection in dynamic scenes. *Vision Research*, 46(26):4333–4345, 2006.

[5] S. Connelly, L. Hasher, and R. Zacks. Age and reading: The impact of distraction. *Psychology and Aging*, 6(4):533–541, 1991.

[6] J. Henderson and F. Ferreira. Effects of foveal processing difficulty on the perceptual span in reading: Implications for attention and eye movement control. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16(3):417–429, 1990.

[7] J. Jacobson and P. Dodwell. Saccadic eye movements during reading. *Brain and Language*, 1979.

[8] S. Kemper and J. McDowd. Eye movements of young and older adults while reading with distraction. *Psychology and aging*, 21(1):32, 2006.

[9] I. Mund, R. Bell, and A. Buchner. Age differences in reading with distraction: Sensory or inhibitory deficits? *Psychology and aging*, 2010.

[10] W. Murray and A. Kennedy. Spatial coding in the processing of anaphor by good and poor readers: Evidence from eye movement analyses. *The Quarterly Journal of Experimental Psychology Section A*, 40(4):693–718, 1988.

[11] A. Poole. Differences in Reading Strategy Use Among ESL College Students. *Journal of College Reading and Learning*, 36:1, 2005.

[12] K. Rayner. Eye movements in reading and information processing: 20 years of research. *Psychological bulletin*, 124:372–422, 1998.

[13] K. Rayner and A. Well. Effects of contextual constraint on eye movements in reading: A further examination. *Psychonomic Bulletin & Review*, 1996.

[14] J. Sibert, M. Gokturk, and R. Lavine. The reading assistant: eye gaze triggered auditory prompting for reading remediation. In *Proceedings of the 13th annual ACM symposium on User interface software and technology*, pages 101–107. ACM, 2000.

[15] J. Theeuwes. Exogenous and endogenous control of attention: The effect of visual onsets and offsets. *Perception & Psychophysics*, 49(1):83, 1991.