Visual Search and Reading Tasks Using ClearType and Regular Displays: Two Experiments

Andrew Dillon, Lisa Kleinman, Gil Ok Choi, Randolph Bias
University of Texas at Austin – School of Information
1 University Station, Austin, TX 78712
{adillon, kleinman, higilok, rbias}@ischool.utexas.edu

ABSTRACT
Two experiments comparing user performance on ClearType and Regular displays are reported. In the first, 26 participants scanned a series of spreadsheets for target information. Speed of performance was significantly faster with ClearType. In the second experiment, 25 users read two articles for meaning. Reading speed was significantly faster for ClearType. In both experiments no differences in accuracy of performance or visual fatigue scores were observed. The data also reveal substantial individual differences in performance suggesting ClearType may not be universally beneficial to information workers.

Author Keywords
ClearType, readability, visual displays, information tasks

ACM Classification Keywords
H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION
Improvements in image quality for display technologies are assumed to offer meaningful benefits to users performing a variety of office tasks [9]. Nielsen argues that small improvements in reading speed could save companies up to $2000 per worker who spends about 20% of their time reading emails, web pages etc. [10]. Since most information tasks involve a mix of direct reading, searching, scanning, and manipulation of documents, it can prove difficult to tease out exactly where, and to what extent image quality will have the greatest effect, but tasks involving significant serial reading where the eye makes continuous and extended contact with text are the most likely to benefit [5].

The history of research on the benefits of image quality has been dominated by reading tasks requiring users to proof or read serial text [4, 8]. While initial research on this theme pursued a single variable explanation, it has become clear to researchers over the last decade that multiple factors affect the user response to electronic text. Certainly among the key variables affecting human processing of electronic text is image quality. Reading proceeds with the visual processing of textual representations, and while the human perceptual system is capable of extracting meaning from even partial images, there can be a temporal cost associated with reduced image quality. Gould et al. [8] advanced the image quality hypothesis as the major determinant of reading speed and his and others' work indicated that reading from screens can be as much as 30% slower than reading from paper, purely as a result of differences in image quality. Subsequent studies have shown that, for highly controlled short text reading, reading speed differences between the media could be reduced with the use of high resolution screens employing black, anti-aliased text on a white background with a screen-optimized font.

However in contemporary work environments users often shift between information activities and are likely to perform a range of tasks with displays that vary considerably from the extended reading scenario of traditional experimental trials. Task variability may therefore play a significant role in determining the benefits of any image quality improvements. O'Hara and Sellen [11] for example noted that ease of manipulation was a more significant factor than image quality in determining performance and satisfaction in a note taking and summarization task. Adler et al. [1] demonstrated that work-related reading activities are often interrupted, and lengthy texts are rarely read to completion in one sitting. Their data come from a limited set of interviews and self-reports, not from direct observation of users so must be considered exploratory, but these findings may place some constraints on the impact of image quality. Nevertheless, even Adler et al. estimate that their participants are spending on average more than 80% of their working days reading. Even if such workers are not completing lengthy reading tasks in one sitting, extended reading of text is still occurring, hence even in a more diverse task environment

than has been typically studied experimentally, image quality is likely to influence performance.

Images underlie reading—they are the building blocks of the processes that lead to comprehension, and, as has been summarized extensively in recent work on the topic (see e.g. [7]), image quality limits the speed and accuracy of reading. Given the known parameters of human perceptual and cognitive processing of text, image quality will at least affect the speed with which users can process a text and therefore directly influence rate of comprehension [14].

Microsoft's ClearType technology promises significant improvement in the visual quality of text presentation on screen. ClearType is a setting that is manipulated through the Windows operating system of devices using Liquid Crystal Displays (LCDs) and works by altering the vertical color stripe within a pixel allowing for changes in how the text looks at fractional levels. These changes aim to enhance the resolution of the screen text and thus improve readability. While font design has continued to be studied as a factor in electronic document presentation [see 2, 3], ClearType can work *across* multiple fonts and may have more applicability in a world where many users like to choose their own font displays.

Experimental findings to date on this display technology are limited but early indications are that users respond favorably under some circumstances though the performance benefits are not obvious. Tyrrell et al. report that readers of LCD screens significantly preferred ClearType to regular displays and reported lower mental fatigue after 60 minutes of reading [13]. In a study of editing, scanning and reading tasks, each involving progressively greater amounts of direct eye-on-text contact, Dillon et al. observed a 5% speed improvement for reading with ClearType but no performance or preference differences for either the scanning or editing tasks [6], suggesting that the benefits of image quality improvements may be limited to only certain information tasks.

A detailed examination of user performance in that scanning study suggested a possible confound since the calculation of task speed measured both visual scanning and the physical activities of typing the task answer. As the physical components of the total reading activity are unaffected by image quality, the resulting task time was not an accurate index of scanning performance. Indeed not only was typing adding to the estimate of speed, but the requirement to specify the cell's text often necessitated several back and forth glances between target and keyboard. Similarly, for the editing task, much of the users' time was spent moving between screen and paper, and typing. In both these tasks, the impact of image quality was likely to be lessened, and its effects swamped by other factors.

The present experiments are intended to explore this issue further, particularly as it pertains to scanning, a typical information worker activity that is highly visual.

Specifically, we sought to control the user's interaction with the document so that only the visual act of interaction, rather than any physical document manipulation or typing, was involved. Unlike serial reading of text, scanning does not require users to follow a defined visual path, and the goal of target location involves different cognitive processes than text comprehension [12].

We designed two separate experiments. In the first, users scanned a set of spreadsheets for target information. The task was designed to require rapid scanning and time spent reading the task and answering did not contribute to the speed score, this overcoming one of the potential sources of confound observed in the Dillon et al. study [6]. In the second experiment, users read a 2000 word article and answered a set of comprehension questions. This task required extended eye-on-display reading. Distinct user populations were employed for each experiment to control for any task or fatigue effects which may have impacted the earlier results.

In the second experiment, the 2000 word document was presented in blocks of 400 words. By most measures, this is not a very lengthy text and no participants complained about the length or reported that the task was unusual for them. We selected this text length for reading in order to reflect typical article lengths found in popular web sites, and to overcome some of the published criticisms of Gould et al.'s [8] groundbreaking work on reading from screens where the texts of 1000 words were frequently criticized as too short.

The two experiments should be seen as part of one study, the aim of which is to engage users in different but common information tasks that vary in the amount of time their eyes spend directly on the screen. We believe this to be the first time such a manipulation has been formally conducted despite increasing evidence from years of readability and document use studies that this factor underlies much of the reported variability in performance [5].

EXPERIMENT 1: SCANNING TASK METHOD

Participants
Participants were recruited through posted advertisements at a large research university in the southwestern USA and online job web sites in the same region. All participants met the following criteria:

- age 18 or older
- English as a first language
- self-reported familiarity using Internet Explorer
- self-reported as having 20/20, standard, or corrected vision
- self-reported as having no reading disabilities or color blindness

26 participants were scheduled for 1-hour time slots and paid $10 for their involvement.

Experimental Design

The study employed a within-subjects (repeated measures) design with each participant completing ten unique scanning tasks in each of two conditions: ClearType and Regular text displays. After completing each set of tasks, the participant answered a visual fatigue questionnaire.

Performance was assessed through three dependent variables: 1) time to complete each scan, 2) accuracy of answers and 3) visual fatigue score.

Each scanning task was presented using test software, controlled by the participant, which automatically captured time data and user responses. Accuracy data was based on the number of correct answers the participant provided; the total possible score was 20 points (one point per scan). The visual fatigue scores were based on Tyrrell's six scales of mental and physical fatigue with answers ranging on a seven point scale of strongly disagree to strongly agree [12]. The wording of each of the visual fatigue questions was modified slightly to make sense for the scanning tasks (e.g. "working with computers" in the original Tyrrell questionnaire was changed to "answering these questions").

Experimental Material

For each task the subject was asked to determine the frequency of occurrence of target items in a spreadsheet display, see example in Figure 1 (note that in the test display the user did not scroll, the full spreadsheet was visible each time).

All the text content in the scanning tasks came from actual web sites with 12 pt fonts. Of the 20 test pages, 12 used Verdana, four were Courier, and four were Arial. Since all fonts were deemed representative by the experimental team no formal balancing of this variable occurred but tasks were matched for each condition so that the participant always completed a similar scan (same font, same length, similar content type) in both conditions. Order was counterbalanced for ClearType and Regular, as was the content within these conditions.
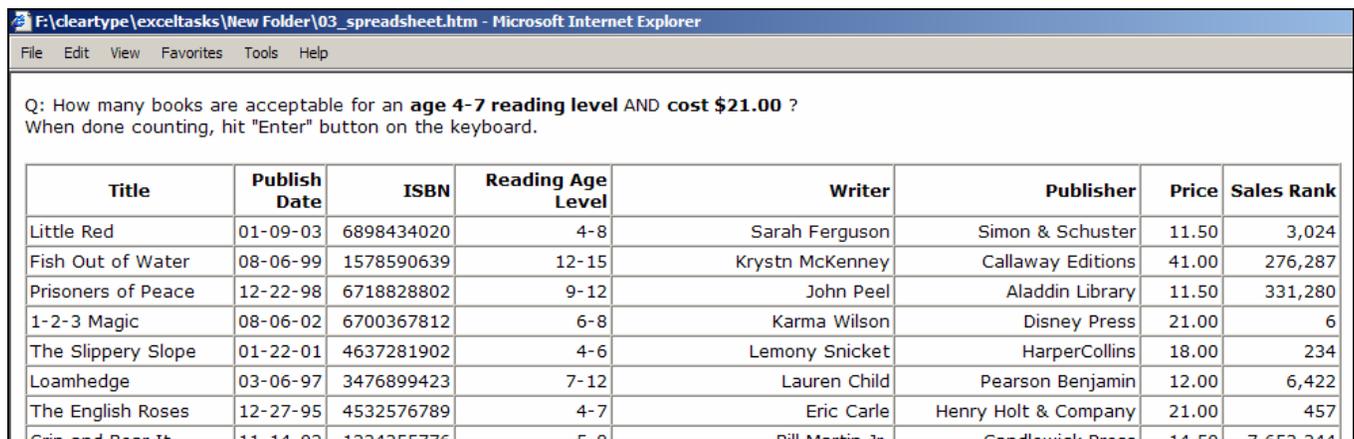
Equipment

Participants all used the same Dell Latitude C840 laptop, with the choice of using a mouse attached peripherally or the mouse-equivalents on the laptop keyboard. The screen size of the laptop was 15 inches, with a display setting of 1600x1200 pixels. ClearType was gamma tuned for the laptop to optimize its setting for this specific laptop. Gamma tuning adjusts brightness levels so that colors are correctly displayed. Internet Explorer 6.0 was used as the browser environment to present the test software and task.

Procedure

Each participant was seated in a closed room environment facing the laptop. The moderator explained the multi-part procedure for the session and then gave the participant practice scanning tasks. After completing two practice scans, participants proceeded through the first block of ten scans at their own pace. For each task, the participant initiated the activity by reading the counting question presented on the test screen (e.g. "How many books were published in 1985?") and then clicking "enter" to make visible the spreadsheet screen to be searched and thereby starting the timer. As instructed, the participant again clicked "enter" when they had determined the answer, thus stopping the timer, at which point the test software then prompted the user to type their answer. After submitting the typed answer, the next task was presented on screen and the participant proceeded as before, until the first block of ten scans was complete. In this way, the timer only recorded the duration of the task for which the participant was scanning the spreadsheet.

At the end of this block, the stimuli and the moderator prompted the participants to answer the visual fatigue questionnaire (by hand on paper). The display condition was switched automatically by the test software while the participant filled in the fatigue scales and the second block of ten scans commenced as before. Once more, at the end of this second block, the participant again completed the visual fatigue scales, after which the experiment ended.

Figure 1. Partial display of scanning task.

EXPERIMENT 1: SCANNING TASK RESULTS

Time, accuracy and visual fatigue scores were calculated for each participant and the summary data is presented in Table 1. Participants performed scans more quickly when ClearType was present, and there was noticeably greater variation in performance observed in the Regular condition. On average, participants were 24 seconds faster completing the scans with ClearType, a significant improvement of [$F_{(1, 25)} = 5.07$, p = .016], see Table 2. There were no significant differences in accuracy or visual fatigue.

	ClearType		Regular	
	Mean	**Std. Dev.**	**Mean**	**Std. Dev.**
Time (seconds)	279.54	54.46	303.81	79.42
Accuracy (max = 10)	8.46	1.07	8.73	1.22
Visual Fatigue (max = 7)	3.08	1.60	3.03	1.46

Table 1. Summary of scanning task data.

Source	Sum of Squares	df	Mean Square	F	Sig.
DISPLAY	7656.94	1	7656.94	5.07	.016
Error (DISPLAY)	37744.56	25	1509.78		

Table 2. Repeated measures 1-tailed ANOVA for scanning.

These results support the hypothesis that ClearType can improve performance in visually intensive scanning tasks. On average, where the eye-on-screen component of the task dominates, people performed 7.2% faster with ClearType than with Regular displays, with equivalent levels of accuracy and fatigue.

Mental and Physical Fatigue

The results of the perceived mental and physical fatigue scores for the scanning task are summarized in Table 3. For all scores, a greater value represents more perceived fatigue/strain (on a 7-point scale). No significant differences were observed between conditions.

Question	ClearType	Regular
1. My back and/or neck hurts from sitting in this position while answering the questions.	2.62	2.54
2. Answering these questions on a computer gives me a headache.	2.81	2.65
3. After answering these questions, my vision seems blurry when I look at distant objects.	3.73	3.69
4. I feel mentally fatigued right now.	2.50	2.54
5. My eyes feel strained right now.	4.00	4.00
6. Overall, answering these questions has made me feel fatigued.	2.81	2.73

Table 3. Mental and physical fatigue in the scanning task.

Individual Differences

While the significant effect for display points to a robust benefit for ClearType on this type of visual scanning, the individual participant scores indicate that only 16 of the 26 participants performed better in the ClearType condition. However, for these people, the advantage of ClearType was pronounced, as shown in Table 4.

Difference Score Range	# of Participants
0 to 30 sec faster in ClearType	5
31 to 100 sec faster in ClearType	7
Over 100 sec faster in ClearType	4
0 to 30 sec faster in Regular	5
31 to 100 sec faster in Regular	5
Over 100 sec faster in Regular	0

Table 4. Difference score groupings.

Of the 16 who performed better with ClearType, four performed over 100 seconds faster, on a set of tasks that on average took around 300 seconds to perform. No equivalently large differences were observed for participants who were faster with Regular text. We can examine the full distribution of these individual differences graphically in Figure 2.
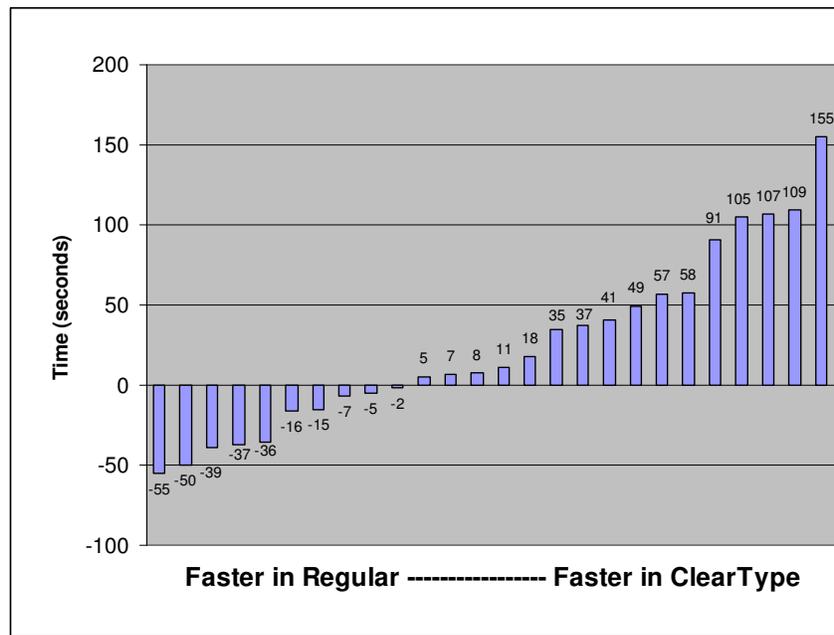
Figure 2. Individual difference scores in the scanning task.

As this shows, once you remove those who demonstrated less than a 10% difference between conditions in either direction (30 seconds) then there are only 5 participants who were much better with Regular text but 11 who performed much better with ClearType. In fact, the mean performance benefit between displays was 53.5 seconds for the "ClearType-advantaged" users compared to a 26.2 seconds benefit for the "Regular-advantaged" (or "ClearType-disadvantaged") users. These data seem to suggest that while ClearType can yield a general improvement in reading speed for many people, its benefits are not universal.

EXPERIMENT 2: READING TASK METHOD

Participants
As before, participants were recruited through posted advertisements at a large research university in the southwestern USA and online job web sites in this same region. All participants met the following criteria:

- age 18 or older
- English as a first language
- self-reported familiarity using Internet Explorer
- self-reported as having 20/20, standard, or corrected vision
- self-reported as having no reading disabilities or color blindness

25 participants were scheduled for 1-hour time slots and paid $10 for their involvement.

Experimental Design
This experiment used a within-subjects 2 x 2 design for ClearType/Regular text and content of the reading material.

Performance was assessed through four dependent variables: 1) time spent reading, 2) accuracy of answers, 3) visual fatigue score and 4) preference score.

Once more, test software was developed to enable automatic capture of time data for each task. To give us better insight into the reading process we divided the text into five screens of text, each approximately 400 words in length. To remove scrolling from the task, each screen was presented as a whole, and the participant pressed the space bar to move forward. The timer started running at the moment the participant displayed the first screen and also captured each screen's reading duration.

Accuracy data was based on the number of correct responses to questions on the material; the total possible score was 6 points. Five multiple-choice questions prompted for factual information from the article (as opposed to inferential reasoning questions); one point was given for each correct answer. Another additional point was calculated for a semi-structured question which asked the participant to list five themes or ideas that the article conveyed—1/5 of a point was given for each correct theme. The participant also responded to an open-ended question to summarize the article in three to four sentences, this was not calculated in the accuracy score but served as a check that the participant had actually read the text for comprehension.

Again, the visual fatigue scores were based on Tyrrell's six scales [12], modified as in the scanning task to suit this context. For this task we also added a preference score, calculated by presenting the participant with ClearType and Regular text side-by-side (counterbalanced across users). Participants rated the text sample which they preferred and
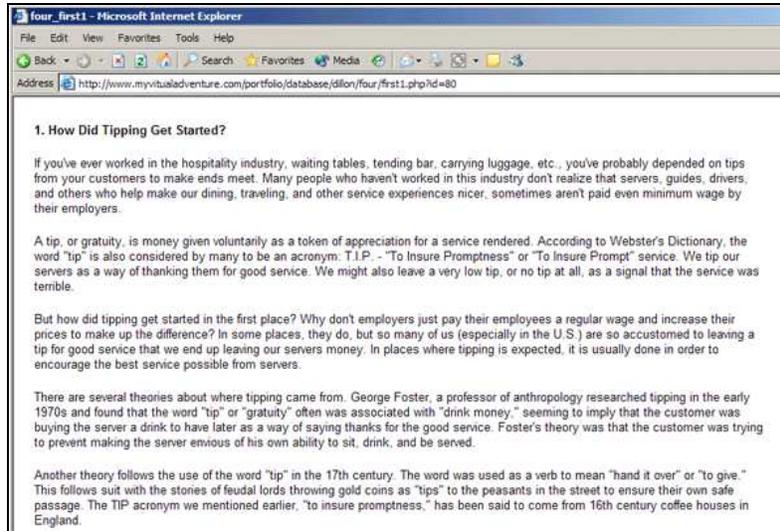
Figure 3. Reading task in Arial text.

how "blurry" each image seemed on a 7-point scale, with 0 as the mid-point expressing no preference or choice, and each of the three points either side of this (from –3 to +3) representing degree of choice (slight, fair, or strong).

Experimental Material
Each participant read two articles of general interest: one about the history of newspapers, and the other about the etiquette of tipping. These texts were selected for equivalent interest and reading levels, as judged by the experimental team. Each article was approximately 2000 words long, and was divided over five screens and presented in Arial 12 pt font. Text and display orders were counterbalanced across participants. See Figure 3 for an example screen.

Equipment
Participants all used the same Dell Latitude C840 laptop, with the choice of using a mouse attached peripherally or the mouse-equivalents on the laptop keyboard. The screen size of the laptop was 15 inches, with a display setting of 1600x1200 pixels. ClearType was gamma tuned for the laptop to optimize its setting for this specific laptop. Gamma tuning adjusts brightness levels so that colors are correctly displayed. Internet Explorer 6.0 was used as the browser environment to present the test software and task.

Procedure
Each participant was seated in a closed room environment facing the laptop. The moderator explained the multi-part procedure for the session and demonstrated an example reading task to the participant who was allowed to ask questions and become comfortable with the procedure. Participants were asked to begin the first reading block by pressing the space bar and proceeding through the reading at their own pace.

After reading five screens (one complete article) the timing stopped and the participant answered the multiple and short-answer questions presented by the test software. After answering these questions, the moderator prompted the participant to answer the visual fatigue questionnaire by hand on paper while the display condition was switched. The participant then proceeded with the second reading article and questions, and answered the visual fatigue questionnaire a second time. The moderator finished the test session by showing a comparison of the ClearType and Regular texts on one screen, asking the participant to rate preference and perceived blurriness for each. On completion of these scales the experiment ended.

EXPERIMENT 2: READING TASK RESULTS
The summary data from all participants is shown in Table 5. On average, participants read 31 seconds (5.65%) faster with ClearType. This is a significant result using a one-tailed repeated measures ANOVA [$F_{(1, 24)} = 3.12$, $p < .05$]. There were no significant effects for accuracy or visual fatigue.

	ClearType		Regular	
	Mean	**Std. Dev.**	**Mean**	**Std. Dev**
Time (seconds)	548.60	184.00	579.32	190.37
Accuracy (max = 6)	5.06	.90	5.30	.70
Visual Fatigue (max = 7)	3.11	1.77	3.01	1.72

Table 5. Summary of reading task data.

Source	Sum of Squares	df	Mean Square	F	Sig.
DISPLAY	11796.48	1	11796.48	3.12	.045
Error (DISPLAY)	90691.52	24	3778.81		

Table 6. Repeated measures 1-tailed ANOVA for reading.

Reading Speed Across Text Blocks

Figure 4 shows the comparison of task completion time between ClearType and Regular text across each of the five screens (blocks) of text. Participants consistently performed faster with ClearType across all five screens, indicating a steady performance benefit for most users.

The results of mental effort and visual fatigue score are summarized in Table 7. For all scores, a higher score represents greater perceived and physical demand or effort. Once again, no significant differences were observed between display conditions.
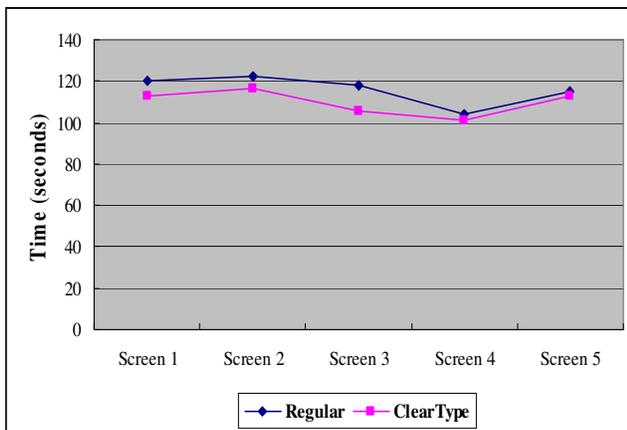

Figure 4. Reading times for each screen of text.

Question	ClearType	Regular
1. My back and/or neck hurts from sitting in this position while answering the questions.	2.69	2.58
2. Answering these questions on a computer gives me a headache.	2.88	3.15
3. After answering these questions, my vision seems blurry when I look at distant objects.	3.31	3.15
4. I feel mentally fatigued right now.	2.77	2.69
5. My eyes feel strained right now.	4.15	3.62
6. Overall, answering these questions has made me feel fatigued.	2.85	2.88

Table 7. Mental and physical fatigue in the reading task.

Preference and Blurriness Scores

Overall, more participants strongly preferred ClearType and rated it less blurry when comparing it to Regular text. Tables 8 and 9 show the number of participants for each of the scales grouped into categories of preference and blurriness. Please note that three questionnaires were discarded from the 25 participants as unusable. There is a significant negative correlation between preference and blurriness scores ($r = -.79$, $p < .01$), as expected.

Preference	ClearType	Regular
Slight	1	2
Fair	3	6
Strong	9	1
Total	13 (55%)	9 (45%)

Table 8. Preference groupings – Number of participants.

Blurriness	ClearType	Regular
Slight	2	3
Fair	7	6
Strong	0	4
Total	9 (45%)	13 (55%)

Table 9. Blurriness groupings – Number of participants.

Individual Differences

If we investigate individual differences in reading speed, as before, we find again that a proportion of participants were faster in the Regular condition. Table 10 summarizes these difference scores in global terms.

Difference Score Range	# Participants
0 to 50 sec faster in ClearType	2
51 to 100 sec faster in ClearType	10
Over 100 sec faster in ClearType	4
0 to 50 sec faster in Regular	4
51 to 100 sec faster in Regular	3
Over 100 sec faster in Regular	2

Table 10. Difference score groupings in the reading task.

Again, viewing these graphically (see Figure 5) we can see that the differences are more pronounced in one direction. Mean improvement scores for the "ClearType-advantaged" users were 85.75 seconds compared to a 65.77 seconds decrement for the "ClearType-disadvantaged" users.

These data again suggest that where ClearType helps, it does so in a more pronounced fashion than it hinders. However, once gain, the benefits are not universal.
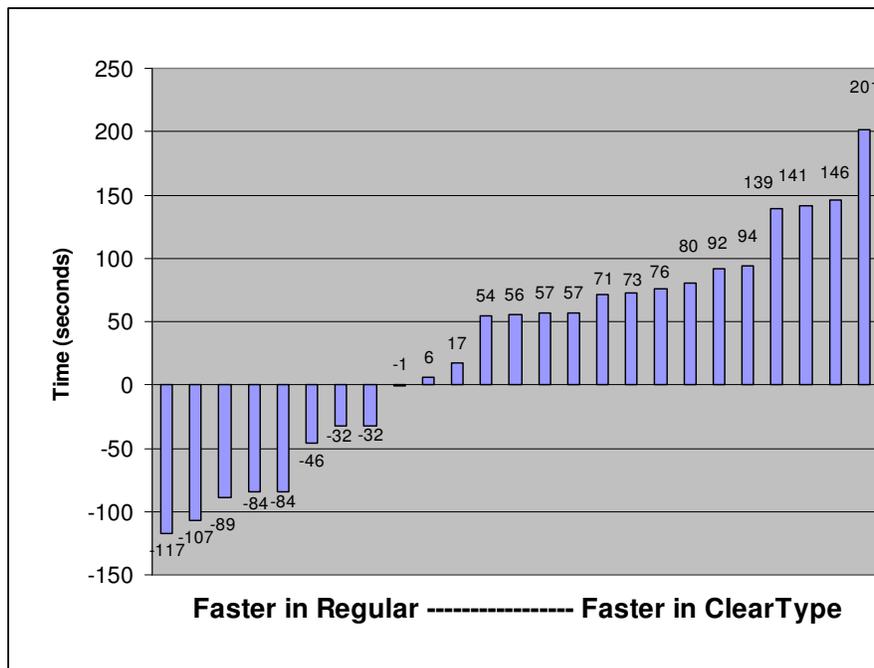
Figure 5. Difference scores in the reading task.

CONCLUSION

For most users ClearType-enhanced text produces a significant improvement in performance times on tasks involving extended eye-on-text interactions. For both reading and visual scanning tasks of the type employed here, which were designed to simulate routine office tasks for information workers, the recommendation to use ClearType is supported. Across 51 users in this study, 32 witnessed improvements in their speed compared to the Regular text display. Furthermore, the size of this performance improvement seen by these users was more pronounced than the decrement witnessed by the 19 users who were faster on Regular displays.

The data from the reading task alone confirm the earlier finding from Dillon et al., which showed a 5.1% speed advantage to ClearType over Regular displays [6]. The present study examined reading alone, not as a part of a combined task block, and it showed a 5.65% advantage to ClearType. This suggests a relatively consistent effect on reading speed for ClearType.

The performance on the visual scanning task was also improved by use of ClearType. The isolation of the visual scanning activity in the present study enabled us to gain an isolated measure of visual scanning without other task components intruding, and here we see speed was improved an average of 7.2% by ClearType.

Both these tests support the argument that image quality improvements have maximum benefit for tasks where the user is spending large proportions of the tasks with their eyes on the text. Where tasks are structured around other

activities (see e.g. [11]) such as moving between screen and paper, manipulating documents or complex typing, the effects of image quality improvements are likely lessened by other factors. Since information work consists of many interrelated physical, perceptual and cognitive activities, the impact of image quality improvements will not always be pronounced.

The issue of individual differences remains a source of interest. With 19 out of 51 users experiencing some disadvantage in a ClearType condition one cannot uniformly endorse ClearType and there is a need to determine what factors may be at work here. Since we have no demographic or perceptual processing data here by which to compare participants we can only speculate on potential acuity or related visual processing factors that may be at work. However, to check for the possibility of an unusual order effect, we conducted a post-hoc analysis with display order as a variable. Interestingly, even though order was fully counter-balanced in this experimental design, there was an interaction effect for display condition with order for the scanning task ($F_{(1, 22)} = 10.95$, $p < .01$) and a three-way interaction effect between order, display and content for the reading task ($F_{(1, 22)} = 9.20$, $p < .01$).

In both tasks, there is no evidence of a practice effect (participants were not faster in the second task block than the first) but all four of the participants who were more than 100 seconds faster with ClearType in the scanning task happened to be exposed to ClearType before the Regular display. It is impossible to determine accurately with these data if this is significant but the effect may point to different impacts resulting from the direction of the

transition to or from ClearType. In other words, it is possible that ClearType produces different effects in users when it is turned off, and they are required to read Regular text than when it is turned on after they have spent time reading Regular text. This is an issue for further study. For now the recommendation is to employ ClearType for most contexts of use but to allow users to disable it if they find it produces effects other than improved performance. Unlike Tyrrell et al. [13], we did not observe any significant fatigue effects, though we used comparable scales, a finding that may be attributable to the very different task durations involved.

Image quality still matters but its effects are subtle and highly dependent on the information task structure. To this we would add that individual differences in response to certain image enhancements are large and warrant further investigation. To develop more accurate assessments of the dollar value of image quality improvements in display technologies, it would be useful to extend Adler et al's. [1] analysis of work-related reading to determine precise estimates of time spent by the majority of information workers in visual contact with displays.

REFERENCES
1. Adler, A., Gujar, A., Harrison, B., O'Hara, K. & Sellen, A. A diary study of work-related reading: Design implications for digital reading devices. In *Proc. CHI 1998,* ACM Press (1998), 241-248.

2. Boyarski, D., Neuwirth, C., Forlizzi, J. and Regli, S. H. A study of fonts designed for screen display. In *Proc. CHI 1998*, ACM Press (1998), 87-94.

3. Bernard, M., Chaparro, B., Mills, M. and Halcomb, C. Comparing the effects of text size and format on the readability of computer-displayed Times New Roman and Arial text. *International Journal of Human Computer Studies, 59* (2003), 823-835.

4. Dillon, A. Reading from paper versus screens: A critical review of the empirical literature. *Ergonomics, 35,* 10 (1992), 1297-1326.

5. Dillon, A. *Designing Usable Electronic Text, 2nd Edition.* CRC Press, Boca Raton, FL, 2004.

6. Dillon, A., Kleinman, L., Bias, R., and Choi, G. and Turnbull, D. Reading and searching electronic documents: An experimental study of regular and enhanced screen displays. *Proc. American Society for Information Science and Technology, ASIST 2004,* Information Today (2004), 267-273.

7. Dyson, M.C. How do we read text on screen. In *Creation, Use, and Deployment of Digital Information* (H. van Oostendorp, L. Breure, & A. Dillon, Eds.), Lawrence Erlbaum Associates, Mahwah, NJ (2005), 279-306.

8. Gould, J. Reading is slower from CRT displays than paper: Attempts to isolate a single-variable explanation. *Human Factors, 29,* 3 (1987), 269-299.

9. Muter, P. Interface design and optimization of reading of continuous text. In *Cognitive Aspects of Electronic Text Processing* (H. van Oostendorp and S. de Mul, Eds.), Ablex, Norwood, NJ (1996), 161-180.

10. Nielsen, J. Avoiding commodity status. http://www.useit.com/alertbox/20020203.html.

11. O'Hara, K. & Sellen, A. A comparison of reading paper and on-line documents. . In *Proc. CHI 1997,* ACM Press (1997), 335-342.

12. Tyrrell, R. and Leibowitz, H. The relation of vergence effort to report of visual fatigue following prolonged near work. *Human Factors, 32,* 3 (1990), 341-357.

13. Tyrrell, R., Pasquale, T., Aten, T. and Francis, E. Empirical evaluation of user responses to reading text rendered using ClearType technologies. *Society for Information Display*, *Digest of Technical Papers,* (2001), 1205-1207.

14. Underwood, G. and Batt, V. *Reading and Understanding*. Blackwell, Oxford, 1996.