Now Dasher! Dash Away!
Longitudinal Study of Fast Text Entry by Eye Gaze

Outi Tuisku, Päivi Majaranta, Poika Isokoski, and Kari-Jouko Räihä
Tampere Unit for Computer-Human Interaction (TAUCHI)
Department of Computer Sciences
FIN-33014 University of Tampere, Finland
{Outi.Tuisku, Paivi.Majaranta, Poika.Isokoski, Kari-Jouko.Raiha}@cs.uta.fi

Abstract

Dasher is one of the best known inventions in the area of text entry in recent years. It can be used with many input devices, but studies on user performance with it are still scarce. We ran a longitudinal study where 12 participants transcribed Finnish text with Dasher in ten 15-minute sessions using a Tobii 1750 eye tracker as a pointing device. The mean text entry rate was 2.5 wpm during the first session and 17.3 wpm during the tenth session. Our results show that very high text entry rates can be achieved with eye-operated Dasher, but only after several hours of training.

CR Categories: H.5.2 [Information Interfaces and Presentation]: User Interfaces – Evaluation/methodology; Input devices and strategies

Keywords: gaze writing, eye tracking, text entry, longitudinal study

1 Introduction

When the Dasher article by Ward and MacKay [2002] was published in Nature, it attracted worldwide interest in the public press. According to it, Dasher was about twice as fast and five times more accurate than any of the previous gaze writing systems. Hence it created a lot of excitement in people working with interactive eye tracking and among people with disabilities. Dasher has now been freely available for a few years in over sixty languages and it seems to be highly appreciated by users with disabilities (see comments from users quoted on the Dasher homepage [2007]). Despite all the attention, to date, no independent experiments on gaze writing with Dasher have been published to verify the results by Ward and MacKay.

In addition to the above-mentioned motivation we were interested in how easy Dasher is to learn using only an eye tracker and exactly how fast new users can become after a few hours of practice. In addition, we hoped to get insight into the pros and cons of eye-controlled Dasher, for example, to see how straining it is for the eyes and what the typical problems in learning and using it are.

ETRA 2008, Savannah, Georgia, March 26–28, 2008.
© 2008 ACM 978-1-59593-982-1/08/0003 $5.00

We will start by explaining how Dasher works and why it is so well-suited for gaze pointing with an eye tracking device. We will then briefly review related work before going into details of our own study.

2 Gaze Writing with Dasher

Dasher [Ward and MacKay 2002] is a text entry interface that is operated by continuous pointing gestures. Writing happens by zooming into a world of characters. In the initial state, all characters are located on the right side of the screen in alphabetical order (see Figure 1).

The user writes by moving the pointer towards the desired character(s). The Dasher interface zooms in, and the area around the pointed character starts to grow and move towards the center of the screen. As soon as the character crosses the central vertical line, it is selected and entered into the text box on top of the screen (see Figure 2).

While the interface zooms in towards the focused character, the language model in Dasher predicts the most probable next characters. The areas of the most probable characters start to grow within the region of the chosen character as it moves left. This brings the next most probable characters closer to the current cursor position, thus minimizing the distance and time to select them.

Cancelling entered characters is done simply by pointing left, which inverts the direction of the action. Instead of zooming in, the interface zooms out and the written characters return from left back to the right side of the screen. The central vertical line acts as a home position. All action ceases when the user holds the cursor at the center of the screen.

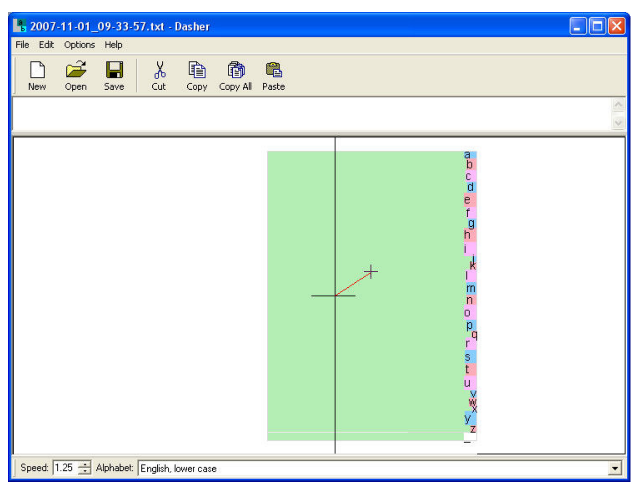

Figure 1: Dasher in its initials state.

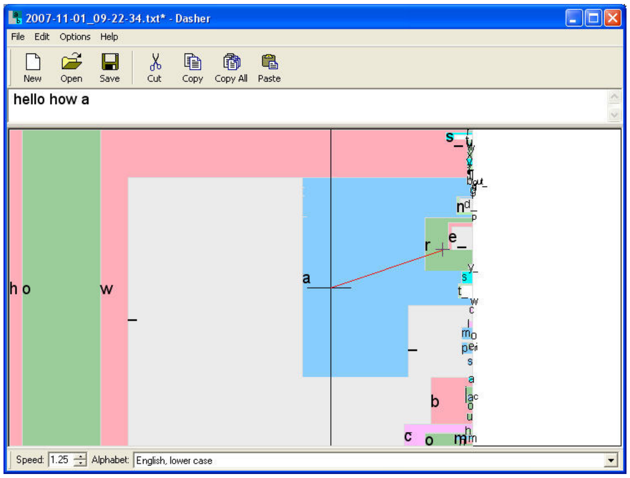

Figure 2: Dasher zooming into the world of characters. The user is in the middle of writing "hello how are you", with letter 'a' of the word 'are' just written.

Dasher can be controlled with any two-dimensional pointing device, such as a joystick, stylus, trackball, and hand, head, or foot operated mouse. There are also implementations of Dasher for other input devices, such as tilt sensors, breath-control, and buttons (switches) [Dasher homepage 2007]. However, controlling Dasher by eye movements has created the most enthusiasm because Dasher seems especially suitable for eye-control.

To appreciate Dasher's features in the context of eye tracking and gaze writing, it is important to understand the special features of eye gaze input. The eye is a perceptual organ; it is normally used for visual observation, not for control. On the other hand, gaze direction can be used for pointing. However, there is no easy way of "selecting" items by gaze alone. It needs to be combined with other modalities such as an eye blink, voice command, manual button click, frowning or any other voluntary muscle activity [Surakka et al. 2003]. If the user is only able to move the eyes, as is the case with people who are totally paralyzed, the system must infer intentions from the gaze alone. The usual method for selection by gaze alone is to use "dwell time". This means that the user fixates on a target for longer than is typically needed for observation and the system selects the fixated item. Dwell time helps to avoid making unintended selections (a.k.a. the Midas touch problem [Jacob 1991]). The downside of using dwell time is the extra time spent on selections. It sets an intrinsic limitation for the maximum text entry rate on systems that rely on discrete selecting of objects. Eye movement input has been found to be distinctly faster than other pointing methods [Jacob and Karn 2003]. Therefore it is a pity that the use of dwell time makes gaze-based user interfaces slow. Despite this problem, using dwell time to select keys (pointed by gaze direction) on a virtual keyboard is still the most common way to eye type [Majaranta and Räihä 2007].

No additional switches or dwell time are needed when using Dasher with gaze. The user simply looks at the characters. An eye tracker follows the user's gaze and moves the cursor to whatever point the user is looking at. Dasher works particularly well with gaze pointing because the desired character is the focus of the user's attention and, therefore, the user's eyes are naturally pointed at it.

Binding the cursor to eye movements has its disadvantages that are present in Dasher, too. Using mouse or any manual pointer, the user is able to look around while continuing to point at one location. Using gaze, the cursor always follows the gaze, preventing the user from looking around without moving the cursor. For example, it is difficult to review the written text without deleting parts of it; moving the gaze to the top left also moves the cursor towards left and may initiate cancelling. To prevent this from happening, it is recommended to pause Dasher before reviewing text and restart when ready.

The inaccuracy of the measured point of gaze is another noteworthy problem that sets constrains to the application of eye tracking in human-computer interaction. Most trackers can achieve an accuracy of 0.5 degrees (equivalent to a region of approximately 15 pixels on a 17" display with a resolution of 1024x768 pixels viewed from a distance of 70 cm). In practice accuracy may be nearer to 1 degree or even less. Again, this decreases the performance of gaze pointing. Generally, gaze is fast only if the targets are big enough.

Research has shown that the accuracy problems can be compensated by using fish-eye lenses [Ashmore et al. 2005] or zooming [Bates and Istance 2002]. Similarly, Dasher's zooming interface is able to alleviate the accuracy problems to a certain extent. Even if the cursor does not hit the target character's area directly, it will be selected when its region grows as the interface zooms in towards the next characters within its region. If the accuracy is way off, it will slow the writing down considerably. However, if the accuracy is reasonably good, "driving" the gaze smoothly through the characters within a relatively small area may be more comfortable to the eyes than an on-screen keyboard where the user has to constantly switch the gaze from one side of the screen to another to select each character one at a time.

Other advantages of Dasher include the built-in language model and speed control. The embedded character prediction makes separate word lists unnecessary since highly probable words will appear in the Dasher user interface automatically. This saves time and cognitive effort as the users do not need to go through a separate list to see if the word they want to write is there. It also makes writing easier and reduces error rates because probable strings get more space and are thus easier to write [Ward and MacKay 2002]. The zooming speed can be controlled easily by gaze position. The action slows towards the center and increases towards the sides of the screen. All action ceases in the center of the screen, offering a resting position for gaze and time for the user to think. This is important because while writing, Dasher requires sustained visual attention from the user. Bystanders have sometimes expressed concerns on the potential strain of the constant visual "noise" to the eyes, which, however, seems to be less disturbing to the person controlling the cursor position.

3 Related Work

An overview of gaze-based text entry systems and related research can be found in [Majaranta and Räihä 2007]. In this section, we concentrate on briefly reviewing research related to learning gaze writing and research involving Dasher.

In their original eye-controlled Dasher study Ward and MacKay [2002] report a top speed of 25 words per minute (wpm). The accompanying figure shows that after an hour of practice, the speed varied between 10 to 25 wpm depending on the user (two were novices, two experts). They compared the experts' top speed to a QWERTY on-screen keyboard (WiViK with word-completion feature enabled) in which expert users achieved 15 wpm, with error rates five times that of Dasher. The results are encouraging but it should be noted that they are based on only four participants. The authors do not report how much prior practice the experts had before the one-hour experiment. The participants reported that they felt the on-screen keyboard to be more stressful than Dasher. This was mainly because with the on-screen keyboard the users were uncertain about potential typing errors. This, in turn, prevented the word prediction program from functioning correctly. The on-screen keyboard also required extra mental effort, because the users needed to move their gaze from the keyboard to the word list that contained potential words predicted on the basis of the first few characters already written by the user. Thus, it was Dasher's inbuilt language model combined with the interface design that made it more pleasant to use.

The results with eye tracking are comparable with the previous experiment by Ward et al. [2000] with a mouse as the steering device. After 60 minutes of practice (in 12*5 min sessions) the text entry rate varied from about 12 to 25 wpm. After a few hours of practice, one expert (one of the authors) could write up to 34 wpm which is comparable to handwriting speed [Wiklund et al. 1987].

Itoh et al. [2006] compared Japanese gaze writing by Dasher and two versions of GazeTalk [Hansen et al. 2001]. Overall, there was no significant difference between the systems. Both achieved a text entry rate of 22-24 Kanji characters per minute (cpm), with performance improving from 19 to 23-25 cpm over seven short trials over three days. The metric used was cpm instead of wpm because the Japanese writing system is not directly comparable with the Latin alphabet. The two systems also elicited similar subjective responses. For example, neither of the systems induced motion sickness which was one of the questions asked. However, there was a significant difference in backspacing rate, which was significantly higher with Dasher than with the versions of GazeTalk (0.0028, 0.0029 and 0.053 backspaces per typed character for the two versions of GazeTalk and Dasher, respectively).

Urbina and Huckauf [2007] compared three new "dwell time free" eye typing approaches (Iwrite, StarWrite and pEYEdit [Huckauf and Urbina 2007]) with Dasher and a traditional QWERTY on-screen keyboard. They only report preliminary results without statistically significant differences. They did not use prediction with any of the systems. Without prediction (i.e., with a flat probability distribution where all characters had an equal share of the screen estate), Dasher lost all its speed advantage. While the participants were able to type 10 to 15 words per minute using the QWERTY keyboard (with 500 ms dwell time), the average speed using Dasher was only 4.7 wpm, and 7.4 wpm by the fastest writer.

The experiments above have been conducted with able-bodied participants. People who have no prior experience on typing or voluntary control of a computer may require a long learning time [Gips et al. 1996] and many introductory activities [Donegan and Oosthuizen 2006] before they are able to benefit from advanced gaze writing systems such as Dasher.

4 Method

4.1 Participants

Twelve able-bodied university students volunteered for the experiment (5 males, 7 females; from 21 to 30 years of age). All were native speakers of Finnish. Eleven of the participants reported normal or corrected-to-normal vision. One person reported to have poor vision which, however, did not show in any way in the experiments. One participant had seen Dasher before, but she had not used it herself. All participants were novices in gaze writing. Two participants reported that they had previously tried an eye-controlled on-screen keyboard for about five minutes.

All participants were rewarded with four movie tickets. To maintain the participants' high motivation during the ten day experiment, we informed them after the first session that the participant who would learn to use Dasher the best would receive an extra prize.

4.2 Apparatus

Dasher (version 4.4.1) was run on a personal computer with Windows XP operating system. We used the Tobii 1750 eye tracking device with its integrated 17 inch TFT color monitor (with 1280 x 1024 pixels resolution) to track the user's gaze. For mouse emulation, we used the eye mouse included in MyTobii (version 2.3.1.0). MyTobii's mouse settings were set to react faster to gaze (thus minimizing the smoothing which would slow down Dasher's reaction to gaze) and "mouse click" was set off.

Dasher alphabet "Suomalainen / Finnish with punctuation and numerals" was used. Dasher was set to "Eyetracker mode" with the "eyetracker autocalibration" option on. The "Eyetracker mode" changes the dynamics of Dasher to better suit navigation by gaze and the "Eyetracker autocalibration" option automatically detects and corrects vertical calibration errors in the gaze tracker. We placed the Dasher window so that there was a small margin above and below the Dasher canvas (the canvas size was 1025 x 640 pixels, excluding menu bars and other window elements), as suggested in the Dasher Manual [MacKay 2006]. Dasher Speed was initially set to 0.21 (information rate measured in bits per second [Ward et al. 2000]) based on pilot tests. We set the "Adapt speed automatically" option on, so that Dasher would increase its speed automatically as the participants' skills improved. With this parameter set on, the same algorithm was used to adjust the speed for each participant, thus eliminating a potentially subjective element involved in manual adjustments. After every session, the participant's end speed was saved, so that the participant could continue with the same speed at the next session. Similarly, each participant also had their own training text file with "language model adaptation" set on.

By default, Dasher is started and stopped by a mouse click. Since we wanted to study how people learn to use Dasher by eye movements alone, we enabled an option that allows starting and stopping based on the mouse cursor position. The "Start with mouse position" option with "Centre circle" attribute was set on in order to allow the participant to stop or start Dasher with gaze simply by looking at the circle (dwell time selection is applied inside the circle to prevent starting or stopping by accident). The circle is transparent when Dasher is on (as in Figure 3) and red when Dasher is stopped. We also set the attribute "Pause outside window" on so that every time the participant would look up to see the given phrase in the experimental software (described below) Dasher would pause.

Finally, Dasher was set to a direct entry mode so that everything the participant wrote could be directed to the software we used to present the stimuli phrases (see Figure 3). We used Dasher's log file to analyze the results.

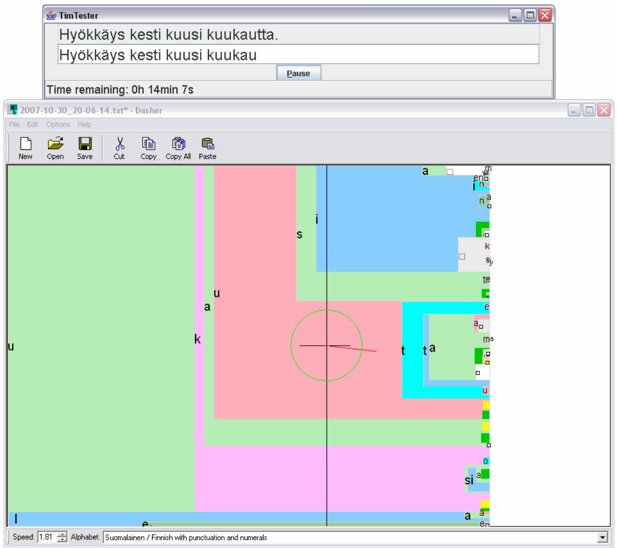

Figure 3: The text written by the user was directed to the experimental software that presented the target phrase.

4.3 Procedure and Design

Each participant was first briefed about the motivation of the study and eye-control in general. After the briefing, Dasher was briefly introduced to the participant by the experimenter using a mouse. The eye tracker was then calibrated and the participant was allowed to practice eye-control freely with an on-screen keyboard (the one included in MyTobii) to get an idea how typing by gaze is normally done and how slow it can be.

During the experiment the participant sat in front of the monitor, at a distance of 50–60 cm from the monitor. Participants were instructed to sit still. However, their movements were not restricted in any way. The eye tracker was calibrated at the beginning of every session, and sometimes also during the session, if the participant expressed the need for re-calibration. If the re-calibration occurred during a session, we tried to do it between phrases and the experimental software was set to "pause" during calibration. If re-calibration had to be done in the middle of a phrase that phrase was ignored in the analysis.

The experimental task was to write as many phrases as possible with Dasher, in Finnish, within the fifteen minutes time limit. Participants were instructed first to memorize the phrase and then to write the phrase as quickly and accurately as possible. The participants were instructed to correct errors if they detected them in the Dasher screen. If they detected errors in previous words they were to ignore them. The phrases were the Finnish translation by Isokoski and Linden [2004] of the 500 phrase set originally published by MacKenzie and Soukoreff [2003]. The phrases were easy to remember, neutral everyday sentences. Some of the phrases contained capital letters and punctuation; some only had lower case letters.

The phrases were presented one by one using Java based software (TimTester[1]) designed especially for text entry experiments by Isokoski and Raisamo [2004]. After each phrase, the participant had to enter the Enter character (included in the Dasher alphabet) to load a new phrase. The program was set to stop when fifteen minutes had passed, but if the participant was in the middle of writing a phrase, the software waited until the participant had finished the phrase before closing down. The window of the experimental program was placed above Dasher, to let the user easily see the phrases, as shown in Figure 3.

Each participant visited our eye tracking laboratory ten times in June 2007. We organized the sessions so that there never was more than 2 days between the consecutive sessions. The first and the last session took about one hour; other sessions lasted about half an hour, including preparations such as the eye tracker calibration. Each participant completed ten 15-minute writing sessions by gaze and one 15-minute session with a mouse. In total, each participant wrote two and a half hours by gaze and 15 minutes with a mouse. The mouse session was left in the very end of the trial series, because the main goal was to study how people learn to write using eye movements alone. We were interested in comparing the results with gaze to those with the mouse but we did not want to corrupt the gaze data by allowing the participants to control Dasher by any manual means. We chose mouse instead of other potential control devices because it was easily available and we assumed the participants would not require extra training in using it.

5 Results

The results are based on data from 11 participants. One participant was a clear outlier and therefore excluded from the statistics but included in the figures (marked with a red dashed line). Results from a few sessions are missing due to technical problems; the missing values were replaced with an average of the previous and the next session.

Analysis of a phrase started from entering the first character and ended to the last character that was part of the phrase. We excluded the Enter character (which ended writing of the current phrase and loaded the next) from the analysis because it took a long time for the participants to find it. It was located in the end of the punctuation marks and it was hard to find. While participants searched for the Enter, they also sometimes accidentally entered extra characters that were also excluded from the analysis.

5.1 Text Entry Rate and Dasher Speed

We measured the text entry rate in words per minute (wpm) where one word is defined as 5 characters, including space and punctuation (for more information, see e.g. [MacKenzie 2003]). The text entry rate for each participant in the 10 sessions is shown in Figure 4.

The grand mean for text entry rate was 2.49 wpm in the first session and 17.26 wpm in the tenth session, thus, significant learning happened during the experiment. The highest session average was 23.11 wpm, reached by participant 9 in session 9.

The Dasher speed was initially set to 0.21 bits per second for all participants. As seen in Figure 5 Dasher speed increased significantly for all participants during the first four sessions.

[1] Available online at http://www.cs.uta.fi/~poika/downloads.php

The average Dasher speed increased to 2.15 during the first four sessions. After that the rapid raise of the Dasher speed leveled off. The grand mean was 0.76 bps in the first session and 2.63 bps in the tenth session.

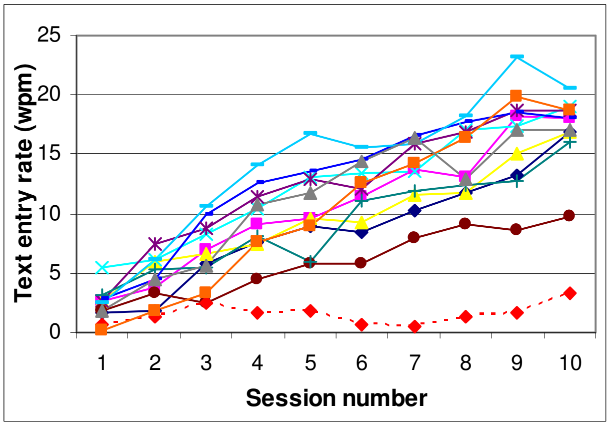

Figure 4: Text entry rate in words per minute. The lowest (red dashed) line represents the outlier who never got past 5 wpm.

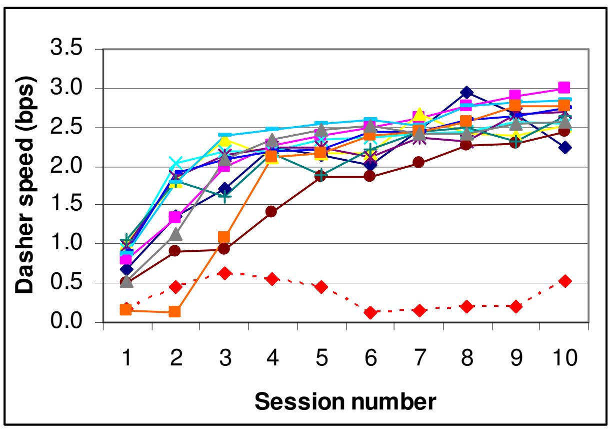

Figure 5: Dasher speed in bits per second. Again, the outlier (marked with a red dashed line) is far below others.

Figures 4 and 5 are somewhat cluttered. Because of this Figures 6 and 7 show the average text entry rate and dasher speed only. These values were computed without the slowest outlier participant.

The shape of the text entry rate vs. session number curve shown in Figure 6 is unusual. Usually longitudinal text entry experiments result in a text entry rate curve that shows rapid growth during the first few sessions and then decreasing gains with further training (see, for example, Isokoski and Raisamo [2004]). The curve in Figure 6 is almost linear. Because of this we omit the conventional fitting of a power curve to estimate the development of the text entry rate beyond our data. We have no confidence on the accuracy of such prediction in this situation.

The Dasher speed curve in Figure 7 shows the decelerating increase that we were expecting to see in the text entry rate curve. Our interpretation of this is that our experiment was not long enough for the participants to reach a level where their learning rate would start to decrease. However, the Dasher speed curve suggests that they were approaching this level. Towards the end of the experiment Dasher no longer increased its speed because the participants were barely coping with the task with the speed setting they had attained.

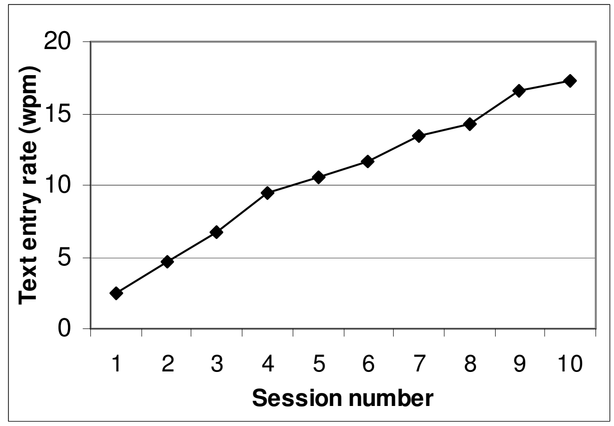

Figure 6: Mean text entry rate per session.

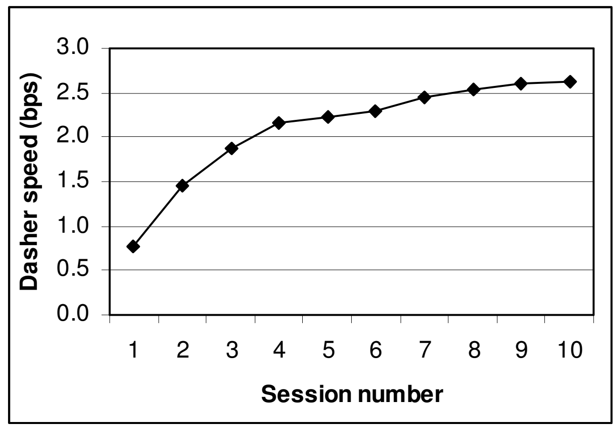

Figure 7: Mean Dasher speed in bits per second.

5.2 Error Rates

Error rates were measured in two different ways; the minimum string distance (MSD) error rate and rate of backspacing. (The over-production rate, often measured in keystrokes per character, KSPC [Soukoreff and MacKenzie 2003], is not applicable with Dasher, because it is operated with continuous navigation instead of discrete keystrokes.)

The MSD error rate was measured using the improved MSD error rate suggested by Soukoreff and MacKenzie [2003]. It is calculated by comparing the transcribed text (written by the participant) with the presented text (stimulus), using the minimum string distance (MSD). The method does not take into account errors that were corrected during the experiment. The error rates reduced significantly during the experiment. The grand mean of MSD for the first session was 10.72 and the grand mean for the tenth session was 0.57. The grand mean of MSD for the mouse session was 0.93. The grand mean of average percentage of wrong words for the first session was 33.08 % and 4.04 % for the tenth session.

The rate of backspacing indicates how often the participants cancelled characters. Thus, this measure correlates with errors to a degree. The rate of backspacing can be calculated by dividing the total number of characters erased prior to the current position by the total number of typed characters [Itoh et al. 2006]. Our participants' rate of backspacing reduced considerably during the experiment. The grand mean is 0.26 for the first and 0.13 for the tenth session.

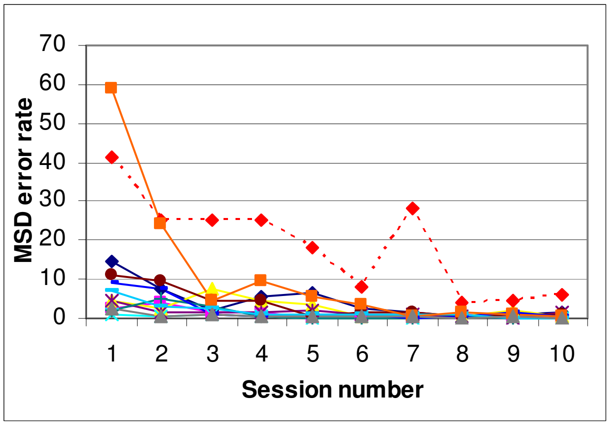

Figure 8: MSD error rate.

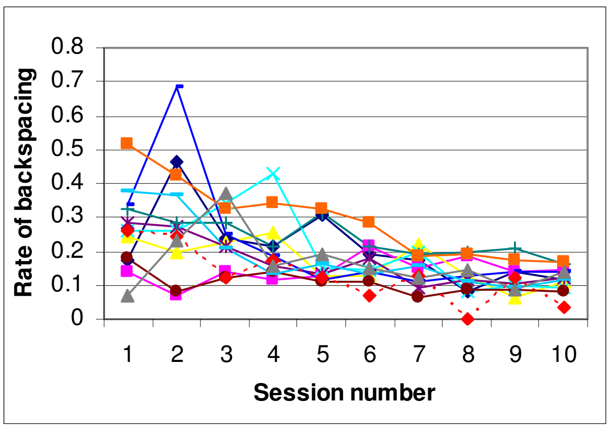

Figure 9: Rate of backspacing.

5.3 Gaze versus Mouse

Even though the participants wrote only one session with the mouse they were significantly faster with the mouse than with the eye tracker. The participants started the mouse session with the speed they had achieved after ending the tenth session by gaze, with the "Adapt speed automatically" option still on. The participants were faster with the mouse with an average of 20.69 wpm compared to the average of 17.26 wpm with gaze in the tenth session. A paired-samples t-test showed that this difference was statistically significant ($t(10) = 3.3$, $p < .01$). The participants had also significantly higher Dasher speed with the mouse (3.91 bps vs. 2.64 bps) ($t(10) = 3.01$, $p < .05$). One participant was faster with the eye tracker and one had the same speed with both devices.

The participants did slightly more errors with the mouse (with average MSD of 0.94) than with gaze (with average of 0.57 in the tenth session). The participants exhibited a lower rate of backspacing with the mouse (0.09 bps) than with the eye tracker (0.13 bps). The differences are not statistically significant.

5.4 Subjective Experience

Usability is not all about efficiency (speed) and accuracy (error rates). Learnability and subjective satisfaction, for example, are equally important. After the tenth session we interviewed the participants. They were asked about their preferences, problems that occurred during the experiment, and ideas for improvements. Participants were also free to express their opinions at any time between sessions.

Overall, the participants' comments about Dasher were positive. They felt that they had learned to use Dasher fairly well. Some commented that they had first felt Dasher to be quite hard to use but after a few sessions the writing was much easier – and even fun.

All participants were surprised and impressed by how fast the writing with Dasher by gaze actually is. They all – apart from the one outlier – thought that gaze writing with Dasher is much faster than eye typing with an on-screen keyboard because one does not have to focus one's gaze at one place for so long. Writing with Dasher was in their opinion fast also because of Dasher's embedded prediction. They gave comments such as "This is really fast", "At the end writing was faster than I could have imagined in the beginning" or "Because of the prediction writing is really fast". However, they did sometimes wonder why a word that they thought was not such a common word in Finnish was offered by Dasher more easily than the word they wanted to write.

Dasher's biggest problems from the participants' point of view were the location of punctuation (11 out of 12 participants expressed this to be a problem) and the centre circle (noted by 4 participants). They found it hard to remember where the punctuation was situated (there is no obvious order, like 'abc…', that would help in remembering the location or order). Participants had problems both in finding the punctuation marks at all, and especially, finding the more uncommon marks. They also reported that one can learn quite fast where the dot or the comma is situated but trying to find a more uncommon punctuation mark was quite difficult. The prediction gave the uncommon marks very small probabilities which also meant they got very little space compared to the more common characters with big rectangles. Participants reported that the centre circle was hard to use because it had a relatively long dwell time. One participant suggested that the centre circle should have shorter dwell time. One participant even suggested that the size of the centre circle should be bigger so that it would be easier to focus the gaze at it. Obviously the center circle issues can be easily corrected and the Dasher developers have informed us that they are already user-adjustable in the Linux version.

Participants were all novices in writing with gaze, although a couple of them had previously briefly tried an eye-controlled on-screen keyboard. It seemed that some participants had some problems in focusing their gaze (most probably due to inaccuracy in calibration). One participant reported that he had to always look above the correct letter's position. In the beginning, some of the participants commented that using Dasher was tiring to the eyes, but the effect diminished when they got used to it. Participants also reported that it was harder to write with Dasher by gaze when the Dasher speed grew faster. One participant commented that the faster the Dasher moves the more one has to concentrate on writing. Most participants felt that when the speed was faster Dasher "tossed about" more and thus writing was harder in general – and also harder for the eyes.

At the end of the whole experiment participants were asked to choose whether they preferred the mouse or the eye tracker in Dasher use. Six of them chose the eye tracker and six chose the mouse. Generally the participants reported that the speed was easier to control with the mouse. One participant who chose gaze said that it felt stupid to write with Dasher using the mouse because if one can use one's hands why not write with keyboard.

6 Discussion

None of our 12 participants learned to gaze write with the often mentioned top speed of 25 wpm, reported by Ward and MacKay [2002]. However, one of our participants came quite close with the result of 23 wpm.

One reason might be the different language. In Finnish the word endings are inflected a lot more than in English. In Finnish, not only the tense (present, past, etc.) has an effect but the inflection is also used for situations handled by prepositional constructs such as "to", "in", "for", and "from" in English. This affects the word prediction process because the number of potential continuations within a word increases. For example, let us take the word 'Finland', which is 'Suomi' in its basic, uninflected form. The Finnish analogues for 'to Finland', 'in Finland', 'for Finland', 'from Finland' are 'Suomeen', 'Suomessa', 'Suomelle', 'Suomesta', respectively. Practically all nouns are affected by the inflection, including names of persons and places.

Dasher's prediction capability is only as good (as representative or accurate) as the text corpora used to build the model. In our study we noted that Dasher's current Finnish language model (based on the modern-language novel "Pereat mundus" by Leena Krohn, WSOY, 1998) was not the best possible, because it had some words that are really not Finnish – but are still given rather high probability. For example, when the participant tried to write something that begins with 'H', for example 'Hän' ('She' or 'He' in English), Dasher always offered the word 'Håkan' which is a Swedish boy's name. Because of that, the participants had to correct many phrases that began with 'H'. Also, the word 'Jumala' ('God' in English) was a word that Dasher offered when the participant began to write a word that began with 'j' or 'J', even though 'Jumala' certainly is not the most common Finnish word to start with 'J'. If the training text had better matched the test text, the results could have been slightly better. Furthermore, had the experiment continued, the language model would also have improved automatically since the option that enables the model to adapt (learn) as the user writes was on.

Even though the participants only wrote for fifteen minutes with the mouse they were much faster with it than with the eye tracker. Of course, by that time they were already familiar with Dasher, and obviously there was a strong transfer of the learning effect from the eye-controlled Dasher to the mouse-controlled Dasher. There are several reasons why mouse performed so much better than the eye tracker. Some reasons were also noted by the participants. Most of the reasons originate from the features of the human visual system (as discussed in the beginning of this paper). First, writing with mouse is easier because the user's eyes are free to look around and search for the next characters or quickly check the written text so far. Mouse is also more accurate and does not have any calibration problems. This makes accurate pointing easier. In addition, it increases Dasher speed because the mouse's pixel-level accuracy makes it possible to keep the cursor nearer to the right edge of the screen. Dasher speed increases when the cursor is moved to the right and decreases when the cursor is moved left toward the center where all movement stops.

One of the participants was much slower than the others – so much so that we discarded his data from the analysis. The problems originated partly from the inaccuracy in the calibration. The participant mentioned that his eyes were getting tired because the cross mark (the measured point of gaze) and the line (Dasher's zooming direction) did not match. The inaccuracy does not, however, explain all of the problems he had. It seemed he had a hard time grasping (or getting used to) Dasher's working principles. For example, he complained that "it was hard to perceive [the target letters] because there were so many letters shown around [the desired letter]". He also kept constantly forgetting that the next character should always be selected inside the current character's box, which prevented the language prediction from working correctly. This may explain why he was the only one who wished Dasher would not use prediction and that the boxes around the letters should remain equally sized. Furthermore, he did not exploit the possibility to easily cancel by looking left (his backspace rating is slow compared to other participants; see Figure 9, with the outlier's line marked with dashed red line). His text entry speed was much higher with mouse than with eye tracking (mouse does not have calibration accuracy problems), but the error rate increased along the writing speed (he still did not use the cancel option to correct the errors).

It is important to acknowledge that there can be such "outliers" who will need a long learning time or who may never benefit from Dasher as much as others. It is interesting that already within a relatively small group of 12 participants there was such an outlier with an average speed of 3.24 wpm (in the tenth session). In addition, there was also another participant whose average speed was 9.82 wpm (all other participants topped at or above 15 wpm). With a larger, more heterogeneous group with varying abilities and disabilities, such outliers might be more common.

7 Conclusions

Dasher is a revolutionary concept for text entry. It has been claimed to reach text entry speeds that are almost twice as high as those obtainable by other gaze-based text entry methods. However, the controlled experiments have thus far been small in scale and the top entry rates have been obtained by the developers of the technique themselves.

We wanted to find out how long does it take to learn gaze writing at a high speed using Dasher, and also what would the top entry rate be for a fairly large group of participants (12 subjects). We carried out a longitudinal study that consisted of 10 sessions of 15 minutes of gaze writing over a period of one month. The subjects did not have prior experience with gaze writing. This resembles a situation where a disabled user suddenly loses the movement of the other muscles but the eyes, and has to learn a completely new means of communication.

The learning curve that we observed was quite exceptional: during the 10 sessions the increase in text entry rate was still almost linear. After 2.5 hours of practice, the subjects were able to enter text at an average rate of 17.26 wpm, with the top performer reaching 23.11 wpm. The numbers are lower than we expected, but this is at least partly explained by the quality of the corpus used to build the language model that Dasher used in our experiment. In real long term use the adaptation of the language model could yield further speed-ups, in addition to those obtained through the learning of the user. In our experiment the sentences were biased towards being varying, without much repetition of words, which is not likely to be the case in real life.

An interesting topic for future work would be to obtain a prediction for the top speed obtainable by an eye-controlled Dasher. This would require a considerably longer experiment with expert subjects that regularly produce large amounts of text. Judging from the linear growth of the entry rate curve, rates of 25 wpm or higher do not seem unrealistic.

Acknowledgements

We would like to thank all volunteers who participated in the experiment. We would also like to thank David MacKay for practical suggestions on Dasher's settings and Phil Cowans for tweaking Dasher's log facility for us. Thanks to I. Scott MacKenzie for the use of his Java tools in analyzing the results. This work was supported by COGAIN, Communication by Gaze Interaction (www.cogain.org), project number IST-2003-511598 (NoE) funded by the European Commission.

References

ASHMORE, M., DUCHOWSKI, A. T., and SHOEMAKER, G. 2005. Efficient eye pointing with a fisheye lens. In *Proceedings Graphics interface – GI'05,* ACM International Conference Proceeding Series, vol. 112. Canadian Human-Computer Communications Society, 203-210.

BATES, R. and ISTANCE H. 2002. Zooming interfaces! Enhancing the performance of eye controlled pointing devices. In *Proceedings of the Fifth International ACM Conference on Assistive Technologies – ASSETS 2002*, ACM Press, 119-126.

DASHER HOMEPAGE. 2007. http://www.dasher.org.uk.

DONEGAN, M., and OOSTHUIZEN, L. 2006. The 'KEE' Concept for Eye-Control and Complex Disabilities: Knowledge-based, End-user focused and Evolutionary. In *Proceedings of the 2nd Conference on Communication by Gaze Interaction – COGAIN 2006*, September 2006, Turin, Italy, 83-87. Available online at http://www.cogain.org/cogain2006/COGAIN2006_Proceedings.pdf.

GIPS, J., DIMATTIA, P., CURRAN, F. X., and OLIVIERI, P. 1996. Using EagleEyes – an electrodes based device for controlling the computer with your eyes – to help people with special needs. In J. Klaus, E. Auff, W. Kremser, & W. Zagler (Eds.), *Interdisciplinary Aspects on Computers Helping People with Special Needs*. Oldenburg, 630-635.

HANSEN, J. P., HANSEN, D. W., and JOHANSEN, A. S. 2001. Bringing Gaze-based Interaction Back to Basics. In C. Stephanidis (Ed.), *Universal Access in HCI*, Lawrence Erlbaum Associates, 325-328.

HUCKAUF, A. and URBINA, M. 2007. Gazing with pEYE: new concepts in eye typing. In *Proceedings of the 4th Symposium on Applied Perception in Graphics and Visualization (APGV '07)*, ACM International Conference Proceeding Series, vol. 253 (Tubingen, Germany, July 2007). ACM Press, 141-141.

ISOKOSKI, P., and LINDEN, T. 2004. Effect of foreign language on text transcription performance: Finns writing English. In *Proceedings of NordiCHI 2004*, ACM Press, 105-108.

ISOKOSKI, P., and RAISAMO, R. 2004. Quickwriting as a multi-device text entry method. In *Proceedings of NordiCHI 2004*, ACM Press, 109-112.

ITOH, K., AOKI, H., and HANSEN, J. P. 2006. A comparative usability study of two Japanese gaze typing systems. In *Proceedings of ETRA 2006,* ACM Press, 59-66.

JACOB, R. J. K. 1991. The use of eye movements in human-computer interaction techniques: What you look at is what you get. *ACM Transactions on Information Systems 9, 2,* 152-169.

JACOB, R.J.K and KARN, K.S. 2003 Eye tracking in human-computer interaction and usability research: Ready to deliver the promises (Section commentary). In Hyönä, J., Radach, R. & Deubel, H. (Eds.) *The Mind's Eye: Cognitive and Applied Aspects of Eye Movement Research*. Elsevier Science, 573-605.

MACKAY, D. J. C. 2006. Dasher Manual. Available online at http://www.inference.phy.cam.ac.uk/dasher/download/papers/Manual.pdf

MACKENZIE, I. S. 2003. Motor behaviour models for human-computer interaction. In J. M. Carroll (Ed.), *Toward a Multidisciplinary Science of Human-Computer Interaction*. Morgan Kaufmann, 27-54.

MACKENZIE, I. S., and SOUKOREFF, R. W. 2003. Phrase sets for evaluating text entry techniques. In *Extended Abstracts of CHI 2003*, ACM Press, 754-755.

MAJARANTA, P. and RÄIHÄ, K.-J. 2007. Text Entry by Gaze: Utilizing Eye-Tracking. In I. S. MacKenzie & K. Tanaka-Ishii (Eds.), *Text entry systems: Mobility, accessibility, universality*. Morgan Kaufmann, 175-187.

SOUKOREFF, R. W. and MACKENZIE, I. S. 2003. Metrics for text entry research: An evaluation of MSD and KSPC, and a new unified error metric. In *Proceedings of CHI 2003*, ACM Press, 113-120.

SURAKKA, V., ILLI, M., and ISOKOSKI, P. 2003. Voluntary eye movements in human–computer interaction. In J. Hyönä, R. Radach, & H. Deubel (Eds.), *The Mind's Eye: Cognitive and Applied Aspects of Eye Movement Research*. Elsevier, 473-491.

URBINA, M. H. and HUCKAUF, A. 2007. Dwell time free eye typing approaches. In *Proceedings of the 3rd Conference on Communication by Gaze Interaction – COGAIN 2007*, September 2007, Leicester, UK, 65-70. Available online at http://www.cogain.org/cogain2007/COGAIN2007Proceedings.pdf

WARD, D. J., BLACKWELL, A. F., and MACKAY, D. J. C. 2000. Dasher – a data entry interface using continuous gestures and language models. In *Proceedings of UIST 2000: The 13th Annual ACM Symposium on User Interface Software and Technology*, ACM Press,129-137.

WARD, D. J., and MACKAY, D. J. C. 2002. Fast Hands-free Writing by Gaze Direction. *Nature 414*, 838.

WIKLUND, M. E., DUMAS, J. S., and HOFFMAN, L. R. 1987. Optimizing a portable terminal keyboard for combined one-handed and two-handed use. *Proceedings of the Human Factors Society – 31st Annual Meeting – 1987*, Santa Monica, CA. Human Factors Society, 585-589.