

Haptic Feedback in Eye Typing

Päivi Majaranta
University of Tampere

Poika Isokoski
University of Tampere

Jussi Rantala
University of Tampere

Oleg Špakov
University of Tampere

Deepak Akkil
University of Tampere

Jari Kangas
University of Tampere

Roope Raisamo
University of Tampere

Proper feedback is essential in gaze based interfaces, where the same modality is used for both perception and control. We measured how vibrotactile feedback, a form of haptic feedback, compares with the commonly used visual and auditory feedback in eye typing. Haptic feedback was found to produce results that are close to those of auditory feedback; both were easy to perceive and participants liked both the auditory "click" and the tactile "tap" of the selected key. Implementation details (such as the placement of the haptic actuator) were also found important.

Keywords: Gaze input, vibrotactile feedback, haptics, eye tracking, text entry

Introduction

Real-time eye tracking enables the use of gaze as an input method in user interfaces. For people with severe physical disabilities such gaze-based interaction provides means to communicate and interact with technology and with other people through the Internet. However, using the eyes for control is missing natural feedback; it has to be provided by the application. We were interested in studying how haptic feedback – in comparison to the more commonly used visual or auditory feedback – performs in text entry by gaze.

Eye typing

There are many ways to enter text by gaze, such as using a manual switch, blinking or frowning to select the focused key, or commands based on gaze gestures (i.e. series of eye movements which are mapped to certain letters or commands), (Majaranta & Rähä, 2007). The most common method, preferred by people with disabili-

ties (Donegan, Gill, & Ellis, 2012, p. 285), is to use dwell-time to activate commands shown on the computer screen. Dwell-time was also used in the current study.

In the dwell selection method, the user first focuses his/her gaze on the desired letter on the virtual keyboard. The user then lets the gaze dwell on the letter in order to select it. The predefined dwell duration is used to separate intentional commands from casual viewing. The duration depends on the expertise and preferences of the user. Novices typically use dwell times that are between 500 and 1000 ms (Majaranta & Rähä, 2007). However, after training, most prefer a shorter dwell time. In an experiment to investigate user adjusted dwell time it was found that after about one hour of training, participants who started with a 1000 ms dwell time reduced it down to 500 ms and after some further practice, most stayed around 300-400 ms (Majaranta, Ahola, & Špakov, 2009). Expert typists may use dwell times that correspond to their normal fixation times (Rähä & Ovaska, 2012). Such short dwell times require an accurate eye tracker as there is little possibility for cancellation or correcting gaze

point inaccuracies within the very short time frame (<300 ms) between the initial focus on the key and the final selection (Majoranta, Aula, & Rähä, 2004).

Feedback

As a general rule in usability engineering practice, “the system should continuously inform the user about what it is doing and how it is interpreting the user’s input” (Nielsen, 1993, p. 134). Since the eyes are primarily used for perception, additional use as a means of controlling a computer requires careful design of the gaze interaction in order to provide adequate feedback and to avoid false activations. The false activation problem is known as the Midas touch problem in gaze interaction literature (Jacob, 1991): everywhere the user looks, something gets selected. Feedback plays an essential role in coping with the Midas touch challenge. Unless the system clearly shows which item is under focus, the user has little chance in preventing false selections.

The natural feedback that is present in manual typing is missing in eye typing. In manual typing the user feels the finger touch the key and the key go down, hears the clicking sound of the key, and is also able to see the character appear on the screen at the same time as the finger presses the key. In eye typing none of the tactile feedback is present. Also, the freedom to use eyes for observing visual feedback is limited. When using eyes to make selections, gaze is engaged in the typing process: the user needs to look at the letter to select it. In order to review the text written so far, the user needs to interrupt typing and move his or her gaze away from the virtual keyboard to the text entry field.

Previous research has shown that proper feedback can facilitate the eye typing process, by significantly improving performance and subjective satisfaction (Majoranta, MacKenzie, Aula, & Rähä, 2006). Majoranta et al. (2003) showed that compared to visual feedback alone, a simple auditory confirmation (“click” sound) on selection significantly reduced frequency of swapping the gaze between the virtual keys and the typed text. If the system gives clear feedback on focus and selection, it enables the user to concentrate on the typing instead of constantly reviewing the typed text.

Using dwell time, the user only initiates the action by dwelling on the item. The actual selection is done by the system when the dwell time exceeds the predefined duration. Animated feedback (e.g. a progress bar) indicates

the progression of time, making the dwell time process visible to the user (Lankford, 2000; Hansen, Hansen, & Johansen, 2001). Proper feedback helps the user to keep her gaze fixated on the focused item for dwell-selection (Majoranta et al. 2006). However, short dwell times do not leave much time for animated feedback. Thus, short dwell times may benefit from a sharp, distinct feedback on selection (Majoranta et al. 2006).

A meta-analysis of several different studies (not related to eye typing) indicates that complementing visual feedback with auditory or tactile feedback improves performance (Burke, et al., 2006). This is likely because human response times to multimodal feedback are faster than to unimodal feedback. Hecht, Reiner and Havelly (2006) found an average response time of 362 ms to unimodal visual, auditory and haptic feedback. For bimodal and trimodal feedback the average response times were 293 and 270 ms, respectively. Response times vary also between modalities. Results of studies on measuring perceptual latency generally agree that both auditory and haptic stimuli alone are processed faster than visual stimuli, and that auditory stimuli are processed faster than haptic stimuli (Hirsh & Sherrick, 1961; Harrar & Harris, 2008). The latency for perceiving the haptic stimuli is dependent on the location; a signal from fingers takes a different time than a signal from the head (Harrar & Harris, 2005).

Haptic feedback, i.e. feedback sensed through the sense of touch, has been used extensively in communication aids targeted for people with limited sight or the blind (Jansson, 2008). We were interested in measuring if and how haptic feedback may facilitate eye typing. To our knowledge, haptic feedback has not been studied in the context of eye typing.

Haptic feedback has been found beneficial in virtual touchscreen keyboards operated by manual touch. A good-quality tactile feedback improves typing speed and reduces error rate (Hoggan, Brewster, & Johnston, 2008; Koskinen, Kaaresoja, & Laitinen, 2008).

Haptic feedback can be used to inform the user about gaze events, provided the feedback is given within the duration of a fixation (Kangas, Rantala, Majoranta, Isokoski, & Raisamo, 2014c). If gaze is used to control a hand-held device such as a mobile phone, haptic feedback can be given via the phone. Kangas et al. (2014b) found that haptic feedback on gaze gestures improved

user performance and satisfaction on handheld devices. In addition to hand-held devices, head-worn eye trackers provide natural physical contact point for haptic feedback (Rantala, Kangas, Akkil, Isokoski, & Raisamo, 2014).

Finally, in addition to quantitative measures, feedback also has an effect on the qualitative measures. Proper feedback can ease learning and improve user satisfaction and user experience in general. This applies to visual and auditory (Majaranta et al. 2006) as well as haptic feedback (Pakkanen, Raisamo, Raisamo, Salminen, & Surakka, 2010).

Despite the earlier findings cited above, it was not clear that haptic feedback would work well in eye typing, where the task is heavily focused on rapid serial pointing with the eyes. In addition, there is a good opportunity for giving visual and auditory feedback, because the gaze stays stationary during the dwell time and audio is also easy to perceive. On the other hand, it was tempting to speculate that adding the tactile feedback that is not naturally present in eye typing could improve the user experience and possibly also the eye typing performance. Thus, we found experiments necessary to better understand how haptic feedback works in eye typing.

Experiment 1: Exploratory Study

The purpose of the first experiment was to get first impressions on using haptic feedback in eye typing and to find out how well a haptic confirmation of selection performs in comparison with visual or auditory confirmations.

Participants

Twelve university students (8 male, 4 female, aged between 20 to 35, mean 24 years) were recruited for the

experiment. Students were rewarded with extra points that counted towards passing a course for their participation in the experiment. All were novices in eye typing; six had some experience in eye tracking (e.g. having participated in a demonstration). One participant wore eye glasses. None had problems in seeing, hearing or tactile perception and all were able-bodied.

Apparatus

Eye tracking. Tobii T60 (Tobii Technology, Sweden) was used to track the eye movements, together with a Windows XP laptop. The tracker's built-in 17-inch TFT monitor (with resolution of 1280x1024) was used as the primary monitor in the experiment.

Experimental software. AltTyping, built on ETU-Driver (developed by Oleg Špakov, University of Tampere) was used to run the experiment and to log event data.

AltTyping includes a setup mode where one can adjust the layout of the virtual keyboard. The layout of the keyboard was adjusted so that it included letters and punctuation required in the experiment (see Figure 1). In addition, there were a few function keys, differentiated from the characters by green background (letter keys had blue background). The backspace key (marked with an arrow symbol '←') was used to delete the last character. A shift key (Δ) changed the whole keyboard into capitals; after selection of a key, the keyboard returned to lower case configuration. A special 'Ready' key (☺) marked the end of the typing of the current sentence and loaded the next sentence. It was located separately from the other keys to avoid accidentally ending a sentence prematurely.

Dwell time was used to select a key. Based on pilot tests, the dwell time for key selection was set to 860 ms.




Figure 1. AltTyping virtual keyboard for eye typing.

Feedback

Progression of the dwell time was indicated by a dark blue animated, closing circle (see ‘n’ in Figure 1). This animation was same for all conditions. In the end of the animation, selection was confirmed either by an audible click, a short haptic vibration, or visual flash of the key, as described below.

Visual feedback. Visual feedback for selection was shown on the screen as a “flash” of the selected key (the background color of the key changed to dark red for 100 ms, with constant intensity and tone). The confirmation feedback on selection was shown as soon as the dwell time was exceeded (immediately after the animated circle had closed).

Auditory feedback. Auditory feedback played a ‘click’ sound to confirm selection via ordinary desktop loudspeakers. The sound was similar to the default click sound used by Microsoft Windows (see Figure 2, top).

Haptic Feedback. EAI C2 Tactor vibrotactile actuator, controlled with a Gigaport HD USB sound card, was used for the haptic feedback. We chose vibrotactile actuators over other actuation technologies such as shear (Winfield, Glassmire, Colgate, & Peshkin, 2007) and indentation (Rantala, et al., 2011) because vibrotactile actuators are compact and easy to attach to different body locations. The C2 actuator converted an audio file (a 100

ms file with a sinusoidal wave, with a constant amplitude of 250 Hz) to vibration (see Figure 2, bottom).

Our aim in creating feedbacks was to make them suitable for the modality and as short as possible. The auditory feedback was easily perceived even at its 15ms length. The visual feedback was set to 100 ms long to make it clearly perceivable; with a shorter duration it could be missed e.g. due to a blink. Also the haptic feedback set to 100 ms duration would make it easy to perceive.

Since it was possible that the participant heard some sound from the vibrotactile actuators, and we wanted to isolate the effect of auditory feedback from haptic feedback, the participants wore hearing protectors. The haptic actuator was placed on the participant’s dorsal wrist. It was held in place by an adjustable Velcro strap. We considered a wrist band as a convenient way to attach the haptic actuator. Since some intelligent wrist watches already include inbuilt haptic feedback (e.g. Apple Watch with haptic alerts), this setup is also relevant to design in the wristwatch form factor.

The timing of initiation of the feedback was the same for all feedback types (they started immediately as the dwell time had run out). The experienced “strength” (intensity) of each type was set in pilot tests to be as similar as possible, although it may be considered rather subjective. In practice, for example, the same signal out-

puted as audio and haptic would cause quite different perceived “strength”. Therefore, the haptic wave had to be amplified to produce an experience that is similar in intensity with the auditory click. In other words, a sound constructed from a signal of very low amplitude is still easy to perceive but vibrotactile feedback constructed from the same low amplitude signal would be barely noticeable.




Figure 2. Signal waveforms of auditory (on the top) and haptic (on the bottom) feedback.

Task

The task that the participants completed in the experiment was to transcribe a phrase of text repeatedly. The phrases were randomly drawn from the set of 500 phrases originally published by MacKenzie and Soukoreff (2003). Because the participants were Finnish, the Finnish translation of the phrase set was used. The phrase set was translated in Finnish by Isokoski and Linden (2004). A few typos were corrected and the punctuation was unified for this experiment. The stimulus phrase was shown on top of the experimental software and the text written by the participant appeared in an input field below the target text (Figure 1).

Procedure

The participants were first briefed about the motivation of the study and the experimental procedure. Each participant was asked to read and sign an informed consent form and to fill in a demographics form.

The participant was seated in front of the eye tracker so that the distance between the participant’s eyes and the camera was about 60-70 cm. The chair position was fixed but participant’s movements were not restricted. The eye tracker was calibrated before the task started and recalibrated again before each test block. In the haptic condi-

tion, the setup was checked to make sure the haptic feedback was easily perceivable.

Prior to the experiment, the participant had a chance to practice eye typing. The setup and the task during training were similar to the actual experiment but the feedback was somewhat different: During practice, we used the default visual feedback given by AltTyping. When the participant focused on a key, it visually rose (as if the eye was pulling it up). After the 1000 ms dwell time used during the training had elapsed, the key went down (like it was pressed by an invisible finger). After 150 ms the key returned to its default state.

During the experiment, the phrases were presented one at a time. The participant was instructed to first read and memorize the phrase and then type it as fast and as accurately as possible. The presented phrase remained visible during the task, thus the participant could read it again if necessary. The participants were instructed to correct errors if they noticed them immediately but leave errors that they noticed later in the middle of the text.

After finishing the phrase, the participant activated the Ready key that loaded the next phrase. The typing time for each condition was set to five minutes. For analysis purposes, this preset typing time included only active typing time (from the first entered key to the selection of the last key, before the selection of the Ready key). However, if the preset time had elapsed during typing, the participant could finish the last phrase without interruption.

After each 5-minute block of phrases, the participant filled in a questionnaire about his/her subjective experience of the feedback used in that block. The experiment then continued with different feedback.

After all feedback conditions had been completed, the participant filled in a questionnaire comparing the feedback modes. The session ended with a short interview.

The exploratory study included one session that consisted of the training and three 5-minute blocks (one for each feedback type). The whole experiment, including instructions and post-test questionnaires took about one hour.

Design and Measurements

The experiment had a within-subject design where each participant tested all feedback conditions. The order

of the feedback types was counter-balanced between participants. The feedback type was the independent variable. Dependent variables included the performance and user experience measurements described below.

Speed. Text entry rate was measured in *words per minute (WPM)*, where a word is defined as a sequence of five characters, including letters, punctuation and space (MacKenzie, 2003). The phrase typing duration was measured as the interval between the first and last key selection. The selection of Ready key was not included, as it only served for stimuli changing, not for text entry.

Accuracy. The *error rate* percentage was calculated by comparing the transcribed text with the presented text, using the minimum string distance method. Error rate only evaluated the result but did not take into account the corrected errors. *Keystrokes per character (KSPC)* was used to calculate the overhead incurred in correcting errors (Soukoreff & MacKenzie, 2001).

Subjective experience. Participants' perceived usability and subjective satisfaction was evaluated with questionnaires. The questionnaire was given after each feedback mode, concerning perceived speed and accuracy, learning, pleasure, arousal, concentration requirements, tiredness, consistency and understandability of each feedback, with a scale from 1 (e.g. very low) to 7 (e.g. very high).

In the end (after all conditions), we also asked the participants' feedback mode preferences. They were asked to indicate which feedback mode they liked the best and which the least, which was most clear, trustworthy and pleasant. We also asked if any of the feedback modes was more "dominant (stronger)" than the others.

Statistical analysis

Repeated measures ANOVA was used to test for statistically significant differences. Pairwise T-tests were used to analyze differences between specific conditions when the ANOVA suggested a difference among more than two conditions.

Unfinished phrases and semantically incorrect phrases were left out of the analysis (8 phrases out of 287 phrases in total). Semantic errors in the typing occurred when the participant skipped, added, or mis-recalled whole words (e.g. "Never too *rich*" versus "Never too *much*"). In a few cases, calibration problems caused a premature ending of

the typing and required a re-calibration before the test could continue.

Results

Speed. The grand mean for typing speed was 8.88 wpm (0.86 standard deviation, with the range of means for individual participants was 6.49 - 10.74 wpm). Mean values for each feedback mode were: auditory 9.06 (0.83 SD), haptic 8.90 (0.76 SD), and visual 8.67 (0.98 SD) wpm. There were no statistically significant differences between the feedback modes.

Accuracy. There was variation in error rates between participants. This is no surprise as the participants were novices and they only used each feedback mode once. The grand mean of individual error rates was 0.48% (0.89 SD). The auditory feedback mode had the lowest error rate of 0.17% (0.40 SD), the haptic feedback was in the middle with 0.43% (0.81 SD) and the visual feedback had the highest error rate of 0.83% (1.46 SD). However, these differences were not statistically significant.

The participants corrected some of the errors they made. This was reflected in the keystrokes per character measure. The grand mean was 1.06 KSPC (0.06 SD). The mean KSPC for auditory, haptic and visual feedback modes were, respectively: 1.06 (0.06 SD), 1.05 (0.06 SD), and 1.06 (0.06 SD). The differences between feedback modes were small and not statistically significant.

Subjective experience. There was a lot of variation in the perceived attributes (each mode was rated as the best in at least some category) and the differences between the feedback modes were small. In general, haptic feedback was perceived positively and did not differ a lot from other feedback modes; it was experienced as a pleasurable feedback mode that was easy to learn and the participants could trust.

Auditory feedback and visual feedback were preferred by three participants. Six participants chose haptic feedback as their preferred feedback mode. This is interesting, because two participants spontaneously commented that the haptic actuator placed on their wrist felt uncomfortable, "as the second heart rate pulse" or "it tickles" (three participants marked the haptics as their least preferred option). One participant suggested to give the haptic feedback on finger, as it would "resemble the natural tactile feedback while touch typing". One participant experienced that the auditory feedback came after the

action; technically, all feedback modes were given at the same time.

Discussion

The results are perhaps best summarized by highlighting the general absence of significant differences between the feedback modes. The exploratory study was short, with a lot of variability between participants making it possible that it was not sensitive enough to detect differences. However, it was also possible that differences between the feedback modes really did not exist. There could have been at least two reasons for this: 1) all feedback modes were equivalent in terms of performance and subjective impressions. 2) None of the feedback modes were essential for accomplishing the task and therefore had no effect on the task. We suspect that reason 2 was at least partially true. This was because of the closing circle feedback shown during the dwell time. It allowed the participants to anticipate when the selection would happen. The most efficient way to utilize the system was probably to anticipate based on the circle and ignore the feedback that followed the selection. In fact, during interview in the end, some participants commented that it was hard to evaluate any differences between the feedback modes. Some of the participants commented that they mainly concentrated on the animated feedback given on dwell time progression. This supports our reasoning that the feedback after selection was not important for the completion of the task.

Still, we learned that in the tested condition of dwell time based eye typing haptic feedback did not perform any worse than the auditory feedback. While statistically non-significant, there was a slight tendency for worse performance with visual feedback, which was also reflected in the participants' preferences.

In summary, the first exploratory study showed that haptic feedback seems to work just as well as the other modes of feedback. However, we were left with the suspicion that perhaps the dwell-time feedback interfered with the measurement of the effects of the selection feedback modes. Thus, to measure the effect of the selection feedback modes, we needed an experiment that emphasized that part of the feedback.

Experiment 2: Feedback on Selection

Based on the experiences from the first exploratory study, we designed a follow-up study on the effects of haptic feedback on selection in eye typing.

Participants

Twelve volunteers (4 male, 8 female, aged between 24 to 54, mean 39 years) participated in the experiment. The volunteers received movie tickets in return for their participation. Three wore eye glasses and one had contact lenses during the experiment. The one with contact lenses experienced some problems with eye tracker calibration and required a few re-calibrations during the experiment. Participants included people with varying backgrounds and expertise levels in gaze-based interaction and use of information technology in general (some participants were total novices from outside of the academia). Five had previous experience with eye tracking but not with eye typing. Most had some experience in haptics, usually from the vibration in cell phones (eight reported using haptic feedback during input in their smart phones).

Apparatus

The eye tracker and the experimental software were the same as in the first experiment.

Feedback

Taking into account the lessons learned from the first experiment, we adjusted the feedback and modified the setup. We removed the dominating "closing circle" visual feedback on dwell. By removing the animated feedback on dwell (which also inherently indicated the selection in the end of dwell), we forced the participants to concentrate on the feedback on selection. However, some feedback on focus was necessary as the accuracy of the measured point of gaze is not always perfect and the users thus needed some indication that the gaze was pointing at the correct key. We used a simple, static visual feedback indicating focus: when the user fixated on a key, its background color changed to a slightly darker color. The color stayed darker as long as the key was fixated and returned to its default color when the gaze moved away from the key. A short delay of 100 ms was added before the feedback on focus, because the change in color was hard to detect if it happened immediately after the gaze landed on the key. The delay was included in the dwell

time (500 ms), thus the feedback on dwelling was shown for 400 ms before selection.

Visual feedback. Visual feedback for selection was the same as in the first experiment.

Auditory feedback. Auditory feedback was the same as in the first experiment.

Haptic feedback. In the first experiment, some participants did not like the haptic feedback on their wrist. Therefore, as suggested by the participants, we decided to give the feedback on the finger. The participant placed his or her right index finger on the actuator (Figure 3). One participant had an injury on the index finger and used the middle finger instead. To reduce sound from the vibration, the actuator was placed on a soft cushion. We had the option to use hearing protectors if the participant could hear the sound despite the cushioning but it was not needed. Again, we tested the setup before the actual experiment with each participant.




Figure 3. C2 Tactor haptic actuator.

Some participants complained that the 100 ms constant amplitude haptic feedback was uncomfortable because it “tickled” their skin. Considering that we moved the haptic stimulation to the finger tip, which is known to be more sensitive, we needed to reduce the intensity of the tactile signal. In informal testing we found that a signal of the same 250Hz frequency with decaying amplitude that reached 0 around 70 milliseconds (see Figure 4, bottom) was equally easy to detect, but significantly less “ticklish”.




Figure 4. Signal waveforms of auditory (on the top) and haptic (on the bottom) feedback.

Design and Measurements

We extended the experiment from one session into three so that each participant used each feedback condition three times. We hoped that a longer experiment would better reveal the potential differences between feedback types and the effect of learning. The experiment followed a repeated measures design with counterbalancing of the order of the feedback types between sessions and participants.

We ran a couple of pilot participants to test the modified setup. It was soon found that the original 860 ms dwell time felt too long already after the first session. After a few adjustments, it seemed even as short as 500 ms would work as people seemed to learn quite fast and they were eager to reduce the dwell time quite soon. This is in line with previous studies where participants have adjusted the dwell time already in the first session (Majaranta et al. 2009; Rähkä & Ovaska, 2012). However, we were aware that 500 ms may be too fast for the participants’ first encounter with eye typing. Therefore, in the practice session that preceded the three sessions analyzed below it was set to 860 ms. Although the animated circle was not a part of the second experiment, it was used as the only feedback in the practice session.

Task and Procedure. The task and the procedure were the same as in the first experiment except that the number of sessions was now three. In order to keep the time in the lab reasonable, the sessions were organized in two different days. The practice phase and the first experimental session were completed on the first day. The second and the third sessions were completed on another day within a week from the first day. Between the second and the third session, participants took a break of few

minutes. During the break the participants were instructed to stretch or walk.

Measurements. The same measurements were analyzed as in the first experiment (speed, accuracy, subjective experience). In addition, we collected two measures on gaze behavior. During eye typing, the user has to look at the virtual keys to select them. In order to review the text written so far, the user has to switch her gaze from the keyboard to the text entry field. *Read text events (RTE)* per character is a measure of the frequency of gaze switching between the virtual keyboard and the text entry field. Using dwell-selection, the user has to fixate on the key for long enough to select it. *Re-focus events (RFE)* is the measure of the number of premature exits per character, indicating the user looked away from the character too early and thus had to re-focus on it to select it. All focus-and-leave events were recorded, not only finally selected characters. A single gaze sample outside the key did not yet increase the RFE count, because the experimental software used an intelligent algorithm (described in (Räihä, 2015)) to filter out outliers: single gaze samples outside of the currently focused key did not reset the dwell time counter for that key but only decreased its accumulated time by the time spent outside, allowing the dwelling to continue when the gaze returned to the key. Thus, a real fixation on the other key was required for a focus/unfocus event. The software includes an option for the dwell-time feedback delay, which was set to 100 ms in this experiment. The feedback on focus (change of the background color of the key) was only shown after the minimum dwell time was full.

The results were analyzed using the same statistical methods as in the first experiment.

Results

19 phrases (out of 1048 in total) were left out of the analysis because they were either unfinished (due to poor calibration that forced re-calibration in the middle of the phrase) or semantically incorrect (e.g. the participant misremembered the phrase and thus replaced whole words).

Typing speed. The grand mean for typing speed was 11.20 wpm (SD 1.42), with the average of 11.80 wpm (SD 1.60) for auditory, 11.34 wpm (SD 1.31) for haptic and 10.46 wpm (SD 1.39) for visual feedback. The repeated measures ANOVA showed that the feedback mode had a statistically significant effect on the typing

speed ($F(1.73,18.99) = 10.02, p=.002$, Greenhouse-Geisser corrected). Pairwise t-tests showed that visual feedback differed significantly from auditory ($p=.003$) and haptic feedback ($p=.006$) but there was no statistical significance between auditory and haptic feedback ($p=.135$). Most participants maintained a typing speed well above ten words per minute but one participant was clearly slower than the others (reaching 5.56 wpm on average). However, since that person finished all conditions and the effect (i.e. the average speed and deviation) was similar in all conditions and sessions, the data from this participant were included in the analysis.

A small learning effect was seen in the increase of the typing speed, from an average of 10.91 wpm (SD 1.46) in the first session, 11.10 wpm (SD 1.42) in the second, and 11.59 wpm (SD 1.42) in the third (final) session (see Figure 5). However, the effect of session on text entry rate was not statistically significant ($p=.314$).




Figure 5. Typing speed in words per minute per session.

Accuracy. The mean error rate was fairly low, 1.37% (SD 1.89), but there were huge differences in the error rates between participants as shown in Figure 6. The effect of feedback type on error rate was statistically significant; $F(1.10,12.13) = 8.63, p=.011$ (Greenhouse-Geisser corrected). Visual feedback produced the highest error rates with a mean of 2.14% (SD 2.27), compared to 1.05% (SD 1.62) for auditory and 0.93% (1.77) for haptic feedback. Again, pairwise analysis showed that visual feedback differed significantly from auditory ($p=.009$) and haptic feedback ($p=.015$) with no statistical difference between auditory and haptic feedback ($p=.349$). The mean error rates also decreased slightly towards the end of the experiment (2.00%, SD 2.12, first; 1.15%, SD 1.89, second; and 0.96%, SD 1.65, for the third session) but this learning effect was not statistically significant ($p=.128$).




Figure 6. Mean error rates (%) per each feedback mode illustrate the great variability between participants.

The keystrokes per character (KSPC) revealed that participants corrected several errors while entering the phrases. The grand mean KSPC was 1.16 (SD 0.12), and the means for each feedback type were 1.11 (SD 0.10) for auditory, 1.17 (SD 0.13) for haptic and 1.20 (SD 0.14) for visual feedback. The KSPC reduced slightly towards the end of the experiment, being 1.20 (SD 0.14) in the first session, 1.16 (SD 0.12) in the second and 1.12 (SD 0.10) in the third session. These differences were not statistically significant ($p=.099$ for feedback and $p=.337$ for session).

Gaze Behavior

We counted how often people glance at the typed text field to review text written so far. Deleting errors is often linked with increased RTE, as people tend to check the effect of the correction.

One outlier with very high RTE was excluded from the calculation. There was an increase in the read text events (RTE) when visual feedback was given (average 0.17, SD 0.14) compared to auditory (average 0.08, SD 0.04) and haptic (average 0.09, SD 0.06) feedback. The effect was statistically significant $F(1.06,11.69) = 8.71$, $p=.012$ (Greenhouse-Geisser corrected). The effect of session was not statistically significant ($p=.239$).

There were slightly more re-focus events (RFE) on the key when visual feedback was used (average 2.00, SD 0.68) compared to auditory (average 1.93, SD 0.51) and haptic (average 1.89, SD 0.51) feedback but the differences were also not statistically significant ($p=.211$).

RFE decreased slightly from the first to the last session but the learning effect was not statistically significant (RFE: $p=.197$).

It should be noted that RFE includes all gaze events logged during the use of the virtual keyboard, also including data from deleted characters as well as characters focused accidentally during the search. High RFE may also reflect calibration problems that cause errors.

Subjective Satisfaction

We asked the participants to fill in a questionnaire after session one and session three, as we were interested if and how their experience changes with practice.

Participants clearly preferred auditory and haptic feedback over visual feedback. After the first session, auditory and haptic feedback got 50/50% of the votes. In the end of the third session auditory feedback got 75% of the votes with 9 participants marking it as their favorite feedback mode while haptic was preferred by 3 participants. Nobody preferred visual feedback. Participants commented that “the constant visual flashing annoys me”, or that “the visual flash is too short”. One participant even felt the flashing caused eye fatigue and ended one session a couple of minutes early. One participant commented on liking the animated feedback (a closing circle that was used in the practice phase) more than the visual flash used to confirm selection in the experiment.

Auditory feedback was also experienced as the most clear, reliable and pleasant feedback mode while haptic feedback was rated as the second best in all these categories and visual feedback as the worst.

We also measured the “dominance” of the feedback mode, meaning how “strong” the participants experienced each feedback compared to others. Auditory feedback was experienced as the most dominant feedback by 6 participants, haptic by 4, visual by nobody; and 4 participants found all equally strong.

Some participants commented that auditory and haptic feedback were easier to perceive: “I cannot always see the flash but sound and tactile feeling work even if I have my eyes closed”; “I feel more confident with auditory feedback compared to plain visual feedback”; “Hearing or feeling a click feels natural”. On the other hand, haptic feedback was felt as “weird” or “funny” at first by some participants but they commented on getting used to it by the end of the experiment.

Discussion

The second experiment did not have the visual closing circle feedback as this dwell-time feedback that was speculated to be the dominant form of feedback in the first experiment. Our goal was to tease out the possible differences between the key activation feedbacks. The results showed that haptic feedback and auditory feedback produce similar results in performance and preference metrics.

When the haptic feedback was given on the finger (instead of the wrist as in the first experiment), participants felt it “natural” to feel the “tap” of the selected key – similar to hearing the auditory “click”.

We used a haptic actuator that produced the “tap” feeling through vibration. Even if the peak of the vibration was short (in the beginning of the descending wave), the vibration could still be felt. In the first experiment, some participants described it as a tickling feeling. A smoother but still pointy single tap might work better and perhaps annoy less, especially if it was to be used on wrist or other parts of the body. It would also be possible to better differentiate the auditory and haptic feedback by using more distinct signal shape designs. Our current approach was to use amodal stimulus parameters (i.e., intensity and shape) that are common in both modalities. An alternative would be to optimize the feedback signals by using parameters such as pitch that is central in auditory feedback design (Brewster, Wright, & Edwards, 1993) but has no clear tactile equivalent.

The visual feedback produced worst performance and was also disliked by the participants. This does not mean that visual feedback in general is worse than the others. For example, the animated circle dwell-feedback has been successfully used in other experiments (e.g. by Majoranta et al. 2009) and participants in the first experiment seemed to rely on it. Thus, the results in experiment 2 apply to the “strong” flash as key activation feedback used in this experiment. For example, the color choice of the visual flash may have an effect. Instead of a strong dark red, more subtle options like different shades of the key’s normal color could be used.

With very short dwell times, there may not be enough time to show animations, thus other options should be considered. Expert typists using very short dwell times probably rely on the typing rhythm and do not wait for the feedback (Majoranta et al. 2006). On the other hand, a

short visual feedback may easily go unnoticed (e.g. due to a blink) while auditory and haptic feedback – as well as an animated circle with long enough duration – are easier to perceive.

The finding that visual feedback performed worse than auditory and haptic feedback could be partly explained by the fact that human response times are faster to auditory and haptic feedback (Hirsh & Sherrick, 1961; Hollar & Harris, 2008). However, it may also be an effect of our implementation. It is hard to create objectively similar feedback with different modalities. For example, auditory signal may appear longer than visual stimuli, especially with short signals of less than 1s (Wearden, Edwards, Fakhri, & Percival, 1998). It is possible that the results would be something else with different kind of implementation of the feedbacks. We used typical visual and auditory feedback, and created haptic feedback that resembled the auditory feedback. We then adjusted their parameters to make them as similar as possible.

Our experiments focused at feedback on selection. There is not much research on the effects of auditory or haptic feedback on dwell time progression – apart from the pilot work conducted by Zhi (2014). She studied different kinds of haptic feedback types but did not find many differences. However, one important observation and difference to the work presented here was that a constant haptic vibration during dwelling without any pause between keys (as focus changes) is not helpful and may feel annoying. In addition, constant haptic feedback is known to produce the effect of habituation (Thompson & Spencer, 1966), making the user to push attention to a continuous stimuli in the background.

One important lesson to learn in our experiments is that the implementation of the feedback required careful design and experimentation. For example, participants relied on the animated visual feedback in the first experiment but did not like the short flash used in the second experiment. Also, even if the haptic feedback given on wrist was not appreciated in the first experiment, participants did like the haptic feedback when it was given on the finger. Nobody complained about the location in the second experiment. On the contrary, the finger was seen as a natural location for a tactile ‘tap’ in a typing task.

There were large differences between our participants, shown in the performance metrics as well as in the subjective ratings. This emphasizes the importance of offer-

ing adjustable choices and supporting people with different preferences and needs.

Conclusions and Future Work

We studied the effects of haptic feedback on key selection during eye typing and found it to perform well in comparison with auditory and visual feedback. Haptic feedback produced results that are close to those of auditory feedback; both were easy to perceive and participants liked both the auditory “click” and the tactile “tap” of the selected key.

We believe there to be unexplored potential in the use of haptic feedback in gaze controlled interfaces. Haptic feedback offers a private channel that is perceived by the user only, therefore enabling feedback that may support efficient use of communication aids without disturbing other people around. The haptic feedback channel may work especially well when eye typing is done with eye tracking glasses that offer natural contact points with the user’s skin behind the ears and on top of the nose (Rantala et al. 2014; Kangas et al. 2014a).

Acknowledgments

We thank the participants who volunteered for the experiments. The work was supported by the Academy of Finland, projects HAGI (decisions numbers 260026 and 260179) and MIPI (decision 266285).

References

- Brewster, S. A., Wright, P. C., & Edwards, A. D. (1993). An evaluation of earcons for use in auditory human-computer interfaces. *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI '93)* (pp. 222-227). New York, NY, USA: ACM.
- Burke, J. L., Prewett, M. S., Gray, A. A., Yang, L., Stilson, F. R., Covert, M. D., . . . Redden, E. (2006). Comparing the effects of visual-auditory and visual-tactile feedback on user performance: a meta-analysis. *Proceedings of the 8th international conference on Multimodal interfaces (ICMI '06)* (pp. 108-117). New York, NY, USA: ACM.
- Donegan, M., Gill, L., & Ellis, L. (2012). A Client-Focused Methodology for Gaze Control Assessment, Implementation and Evaluation. In P. Majoranta, H. Aoki, M. Donegan, D. W. Hansen, J. P. Hansen, A. Hirsykyari, & K.-J. Rähkä, *Gaze Interaction and Applications of Eye Tracking: Advances in Assistive Technologies* (pp. 279-286). IGI Global.
- Hansen, J. P., Hansen, D. W., & Johansen, A. S. (2001). Bringing gaze-based interaction back to basics. *Universal Access In HCI: Towards an Information Society for All, Proceedings of HCI International 2001* (pp. 325-328). Mahwah, NJ: Lawrence Erlbaum Associates.
- Harrar, V., & Harris, L. R. (2005). Simultaneity constancy: detecting events with touch and vision. *Experimental Brain Research*, 166(3-4), 465-473.
- Harrar, V., & Harris, L. R. (2008). The effect of exposure to asynchronous audio, visual, and tactile stimulus combinations on the perception of simultaneity. *Experimental Brain Research*, 186(4), 571-524.
- Hecht, D., Reiner, M., & Halevy, G. (2006). Multimodal virtual environments: response times, attention, and presence. *Presence: Teleoperators and Virtual Environments*, 15(5), 515-523.
- Hirsh, I. J., & Sherrick, C. E. (1961). Perceived order in different sense modalities. *Journal of Experimental Psychology*, 62(5), 423-432.
- Hoggan, E., Brewster, S. A., & Johnston, J. (2008). Investigating the effectiveness of tactile feedback for mobile touchscreens. *Proceedings of the SIGCHI*

- Conference on Human Factors in Computing Systems (CHI '08)* (pp. 1573-1582). ACM.
- Isokoski, P., & Linden, T. (2004). Effect of foreign language on text transcription performance: Finns writing English. *Proceedings of the third Nordic conference on Human-computer interaction (NordiCHI '04)* (pp. 109-112). ACM.
- Jacob, R. J. (1991). The use of eye movements in human-computer interaction techniques: what you look at is what you get. *ACM Transactions on Information Systems*, pp. 152-169.
- Jansson, G. (2008). Haptics as a Substitute for Vision. In M. A. Hersh, & M. A. Johnson, *Assistive Technology for Visually Impaired and Blind People* (pp. 135-166). Springer.
- Kangas, J., Akkil, D., Rantala, J., Isokoski, P. M., & Raisamo, R. (2014a). Using Gaze Gestures with Haptic Feedback on Glasses. In *Proceedings of the 8th Nordic Conference on Human-Computer Interaction (NordiCHI EA '14)* (pp. 1047-1050). New York, NY, USA: ACM.
- Kangas, J., Akkil, D., Rantala, J., Isokoski, P., Majaranta, P., & Raisamo, R. (2014b). Gaze Gestures and Haptic Feedback in Mobile Devices. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)* (pp. 435-438). New York, NY, USA: ACM.
- Kangas, J., Rantala, J., Majaranta, P., Isokoski, P., & Raisamo, R. (2014c). Haptic Feedback to Gaze Events. *Proceedings of the Symposium on Eye Tracking Research and Applications (ETRA '14)* (pp. 11-18). New York, NY, USA: ACM.
- Koskinen, E., Kaaresoja, T., & Laitinen, P. (2008). Feel-good touch: finding the most pleasant tactile feedback for a mobile touch screen button. *Proceedings of the 10th international conference on Multimodal interfaces (ICMI '08)* (pp. 297-304). ACM.
- Lankford, C. (2000). Effective eye-gaze input into Windows. *Proceedings of the 2000 symposium on Eye tracking research & applications (ETRA '00)* (pp. 23-27). New York: ACM.
- MacKenzie, I. S. (2003). Motor behaviour models for human-computer interaction. In J. M. Carroll, *HCI models, theories, and frameworks: Toward a multidisciplinary science* (pp. 27-54). San Francisco: Morgan Kaufmann.
- MacKenzie, I. S., & Soukoreff, R. W. (2003). Phrase sets for evaluating text entry techniques. *CHI '03 Extended Abstracts on Human Factors in Computing Systems (CHI EA '03)* (pp. 754-755). New York: ACM.
- Majaranta, P., & Rähkä, K.-J. (2007). Text Entry by Gaze: Utilizing Eye-tracking. In I. S. MacKenzie, & K. Tanaka-Ishii, *Text Entry Systems: Mobility, Accessibility, Universality* (pp. 175-187). San Francisco: Morgan Kaufmann.
- Majaranta, P., Ahola, U.-K., & Špakov, O. (2009). Fast gaze typing with an adjustable dwell time. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)* (pp. 357-360). New York: ACM.
- Majaranta, P., Aula, A., & Rähkä, K.-J. (2004). Effects of feedback on eye typing with a short dwell time. *Proceedings of the 2004 symposium on Eye tracking research & applications (ETRA '04)* (pp. 139-146). New York, NY: ACM.
- Majaranta, P., MacKenzie, I. S., Aula, A., & Rähkä, K.-J. (2003). Auditory and visual feedback during eye typing. *CHI '03 Extended Abstracts on Human Factors in Computing Systems (CHI EA '03)* (pp. 766-767). New York: ACM.
- Majaranta, P., MacKenzie, I. S., Aula, A., & Rähkä, K.-J. (2006). Effects of feedback and dwell time on eye typing speed and accuracy. *Universal Access in the Information Society*, 199-208.
- Nielsen, J. (1993). *Usability Engineering*. Boston, MA : Academic Press,
- Pakkanen, T., Raisamo, R., Raisamo, J., Salminen, K., & Surakka, V. (2010). Comparison of three designs for haptic button edges on touchscreens. *Proceedings of the 2010 IEEE Haptics Symposium (HAPTIC '10)* (pp. 219-225.). Washington, DC: IEEE Computer Society.
- Rantala, J., Kangas, J., Akkil, D., Isokoski, P., & Raisamo, R. (2014). Glasses with haptic feedback of gaze gestures. *Proceedings of the extended abstracts of the 32nd annual ACM conference on Human factors in computing systems (CHI EA '14)* (pp. 1597-1602). ACM.
- Rantala, J., Myllymaa, K., Raisamo, R., Lylykangas, J., Surakka, V., Shull, P., & Cutkosky, M. (2011). Presenting spatial tactile messages with a hand-held device. *Proceedings of the 2011 IEEE World Haptics Conference (WHC' 11)* (pp. 101-106). IEEE.
- Rähkä, K.-J. (2015). Life in the Fast Lane: Effect of Language and Calibration Accuracy on the Speed of

- Text Entry by Gaze. In *Human-Computer Interaction – INTERACT 2015* (pp. 402-417). Springer.
- Räihä, K.-J., & Ovaska, S. (2012). An exploratory study of eye typing fundamentals: dwell time, text entry rate, errors, and workload. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)* (pp. 3001-3010). New York, NY: ACM.
- Soukoreff, R. W., & MacKenzie, I. S. (2001). Measuring errors in text entry tasks: An application of the Levenshtein string distance statistic. *CHI '01 Extended Abstracts on Human Factors in Computing Systems (CHI EA '01)* (pp. 319-320). New York, NY: ACM.
- Thompson, R. F., & Spencer, W. A. (1966). Habituation: a model phenomenon for the study of neuronal substrates of behavior. *Psychological Review* 73, 16-43.
- Wearden, J., Edwards, H., Fakhri, M., & Percival, A. (1998). Why "sounds are judged longer than lights": application of a model of the internal clock in humans. *The Quarterly Journal of Experimental Psychology. B, Comparative and Physiological Psychology* 51 (2), 97–120.
- Winfield, L., Glassmire, J., Colgate, J., & Peshkin, M. (2007). T-PaD: Tactile Pattern Display through Variable Friction Reduction. *Proceedings of the Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC '07)* (pp. 421-426). IEEE.
- Zhi, J. (2014). *Evaluation of tactile feedback on dwell time progression in eye typing*. Tampere, Finland: University of Tampere.