Evaluating Reading and Analysis Tasks on Mobile Devices: A Case Study of Tilt and Flick Scrolling

Stephen Fitchett
Department of Computer Science
University of Canterbury
Christchurch, New Zealand
saf75@cosc.canterbury.ac.nz

Andy Cockburn
Department of Computer Science
University of Canterbury
Christchurch, New Zealand
andy@cosc.canterbury.ac.nz

ABSTRACT
Flick scrolling is a natural scrolling method for mobile touch devices such as the iPhone™. It is useful not only for its performance but perhaps even more so for its ease of use and user experience. Tilt scrolling instead uses the device's tilt to determine the rate of scrolling, which offers several potential interaction advantages over touch sensitive alternatives: scrolling can be achieved without occluding a large proportion of the screen with a hand, finger, or thumb; it frees drag input events for other important actions such as text selection and drag-and-drop; and it works regardless of the hand's state (e.g. moist or gloved). Although previously described, the performance of tilt scrolling has not been compared to flick scrolling, which is now the state of the art. Furthermore, it is unclear how such an empirical comparison should be conducted. To better understand interaction with mobile scrolling, we propose a new method of evaluating scrolling interfaces in the context of reading or analysis tasks. These activities typically involve slow subtle scroll movements rather than large movements typical investigated in most scrolling evaluations. We use this method to thoroughly compare flick scrolling and tilt scrolling. We show that tilt scrolling results in better performance for tasks performed while stationary while there is no significant difference while moving. However, we find that participants prefer flick scrolling and walk faster when completing moving tasks with flick scrolling than tilt scrolling.

Author Keywords
Accelerometer, flick scrolling, mobile device, tilt scrolling

ACM Classification Keywords
H5.2 [User Interfaces]: Input devices and strategies (e.g., mouse, touchscreen).

INTRODUCTION
Scrolling is a fundamental activity in most user interfaces, and consequently it has been extensively studied. Most scrolling evaluations focus on navigating across large regions of a document to either visually search for clearly marked targets (e.g. Andersen (2005), Appert & Fekete (2006), and Cockburn, Savage & Wallace (2005)) or to return to previously viewed sections (e.g. Hinckley, Cutrell, Bathiche & Muss (2002), Alexander, Cockburn, Fitchett, Gutwin & Greenberg (2009)). These rapid long range scrolling activities ignore an important application of scrolling: slowly moving through the document to continually extract information. This is typical for *analysis tasks* such as reading or visual inspection. This type of scrolling is perhaps even more common on mobile devices such as the iPhone™where documents are typically shorter than on the desktop and navigating across large regions is rare; reading a 50 page PDF document, for example, would be uncommon, while reading an email or web page would be more typical.

Scrolling for this purpose has some subtle differences to other types of scrolling which may influence the relative efficacy of scrolling interfaces. Movements are typically small as the user is interested in what is directly below the viewport, not several pages ahead. It is also important that scrolling does not interrupt or distract the user – Kaptelinin, Mäntylä & Åström (2002) found that scrolling and paging caused almost 30 percent of disruptions when reading aloud text displayed in a window. Interaction techniques that easily produce slow, smooth and consistent scrolling motions may be more advantageous for these kinds of tasks.

We therefore propose an evaluation methodology for understanding the efficiency of mobile scrolling interfaces for analysis tasks. We then demonstrate its use by comparing two state of the art mobile scrolling interfaces (flick scrolling and tilt scrolling) implemented on an iPod touch™. Flick scrolling is widely used, provides a good user experience, and has been demonstrated to be as good as the traditional scrollbar for short documents (Aliakseyeu, Irani, Lucero & Subramanian 2008). It is also increasingly used commercially, with our evaluation using the standard flick scrolling implementation on the iPod touch™(other implementations are also available (Aliakseyeu et al. 2008)).

Tilt scrolling, first introduced by Rekimoto (1996), provides a radically different scrolling method that uses accelerometer input to calculate the scrolling direction and speed, based on the device's rotation relative to a neutral angle. It offers several potential benefits over flick scrolling, described in the following section. Although several researchers have adapted and extended tilt scrolling methods, there is little em-

OZCHI 2009, November 23-27, 2009, Melbourne, Australia.

OZCHI 2009 Proceedings ISBN: 978-1-60558-854-4

pirical evidence of its relative efficiency, and none (that we know of) that compares it to flick scrolling.

BACKGROUND

Tilt Scrolling

Over the last 15 years accelerometers have become increasingly widely available and interfaces that make use of them have emerged. Accelerometer based input is particularly interesting in a mobile context for several reasons:

Zero-space input modality. Mobile devices are necessarily small, meaning that there is little space to display interactive controls. Hard or soft buttons consume precious device and screen space, while accelerometer input provides an additional input modality that is independent of any display requirements.

Occlusion-free input. Hard and soft buttons are typically activated with fingers or thumbs, and these occlude the underlying input controls. These problems can be severe when the device is small or functionally rich (necessitating small controls). Although these problems can be ameliorated in a variety of ways, most solutions have associated limitations: using a stylus requires retrieving and holding a device, and necessitates two handed use; 'pointing lenses' (Ramos, Cockburn, Balakrishnan & Beaudouin-Lafon 2007), as used in the iPod touch™, show an offset magnified region, but consume additional screen space and require a two-part target acquisition process; and finally back-of-device interaction techniques (Baudisch & Chu 2009) are in the very early stages of research investigation, and are unsupported by commercial devices. Tilt based input, in contrast, avoids occlusion.

Tilt scrolling frees drag events for other actions. Flick scrolling is achieved by a dragging action on the device surface, meaning that dragging actions cannot be used for traditional important interaction activities such as text selection and drag-and-drop. This limitation of flick scrolling can be worked around through additional input modalities, such as swiping across the screen edge to trigger selections (Roth & Turner 2009), but such solutions reduce the simplicity and elegance of the direct physical metaphors provided by dragging and tilting. Tilt input leaves the drag event unbound, so it can be used for the most appropriate input activity.

Hand-state independence. Touch sensitive screens are sensitive to the hand state. Input fails or is error prone if the user's hand is sweaty or wet, or if it is covered in a glove in cold weather. Tilt input, in contrast, can be achieved whenever the device can be held.

More generally, tilt input typically fall into two classes: gesture-based, which interpret discrete device movements and map them to interface functions, and rotation-based, which use the angle of the device to perform some functionality. Other techniques involve a certain amount of overlap between the two groups. There has been a significant amount of research in tilt-based scrolling interfaces in a variety of contexts in both categories. In the former group, techniques such as TiltText (Wigdor & Balakrishnan 2003) use tilt gestures to disambiguate text entry on mobile phones. The Rock 'n' Scroll input method demonstrates several gestures for scrolling, selecting and rotating photos (Bartlett 2000). Cho, Choi, Sung, Lee, Kim & Murray-Smith (2007) also investigate tilt scrolling for photos and found its performance to be inferior to button based navigation for search tasks using both discrete and continuous tilt input.

Rekimoto (1996) was perhaps the first to investigate tilt-based scrolling and implemented prototypes for scrolling menus and maps, noting that rotations of 10 to 15 degrees are typical during operation. Harrison, Fishkin, Gujar, Mochon & Want (1998) describe tilt-based scrolling for lists. They map the device's rotation relative to a predefined neutral angle onto one of 6 predefined scrolling rates. Tilt operations are started and stopped by physically squeezing the device.

Speed-dependent automatic zooming (Igarashi & Hinckley 2000) has also been integrated into a version of tilt scrolling (Eslambolchilar & Murray-smith 2004) and later enhanced with audio feedback (Eslambolchilar, Williamson & Murray-Smith 2004). Others have also investigated audio or tactile feedback to improve tilt scrolling performance (Cho et al. 2007, Poupyrev, Maruyama & Rekimoto 2002).

In summary, although there has been a large amount of development work on tilt scrolling, particularly in combination with other input and feedback modalities (haptics and zoom), there is relatively little evidence of its comparative efficiency for everyday scrolling activities.

Flick Scrolling

Flick scrolling was first described in the context of virtual walls (Geißler 1998) and tabletops (Reetz, Gutwin, Stach, Nacenta & Subramanian 2006). More recently, several variants have been evaluated on tablets (Aliakseyeu et al. 2008). Aliakseyeu et al. noted that flick scrolling was intuitive and enjoyable to use, although scrollbars performed much better for navigating large distances in long documents. Apple also use flick scrolling extensively throughout the iPod touch™and iPhone™interfaces.

Evaluations of both flick scrolling and tilt scrolling have been sparse and have tended to focus on search tasks. To the best of our knowledge, the two interfaces have never been directly compared in a formal evaluation. Comparing them in the context of reading and analysis tasks therefore provides useful insights not seen in previous work.

DESIGN OF FLICK AND TILT SCROLLING INTERFACES

Two one-dimensional scrolling methods were evaluated for scrolling analysis tasks: flick scrolling and tilt scrolling. The flick scrolling implementation used was identical to Apple's implementation on the iPod touch™and iPhone™. Touching the screen with one or more fingers and subsequently moving them scrolls the view as if physically grabbing it; for example, moving a finger down two centimetres will scroll the view *up* by two centimetres. Additionally, a flick operation can be initiated by very quickly touching the display, moving a finger and releasing it. This causes the view to scroll even after releasing the finger, with velocity decreasing over time. The initial velocity is determined by the speed at which the finger is moved when it is in contact with the display.

Tilt Scrolling Implementation

The tilt scrolling method is controlled entirely by accelerometer input. The device's pitch (rotation in the YZ plane) is constantly monitored and θ, its angle relative to a neutral angle, is used to calculate the scrolling speed. The view is scrolled in the direction that the device is tilted, for example tilting the device forward scrolls down, although users are split about which way is preferred (Bartlett 2000).

The neutral angle is originally the device's initial rotation, however two events reset it to the current location: first, an acceleration of less than 0.2g in every axis for at least a tenth of a second following an acceleration exceeding 0.5g in any axis (e.g., shaking the device or placing it on a surface); second, touching or continuing to touch the device's screen (this was implemented for our experiment, but is not a practical solution). Note that in the former case, scrolling is paused in the time between the initial acceleration greater than 0.5g and the time the neutral angle is reset. This allows for actions such as placing the device down on a desk without the scroll position changing.

An adequate angle-to-speed conversion is then needed to allow both small subtle movements and rapid scrolling. A linear conversion would achieve at most one of these, much like moving a mouse cursor with no acceleration. Other research has used a polynomial mapping (Hinckley, Pierce, Sinclair & Horvitz 2000). We chose to use a circle geometry based mapping, summarised in Figure 1 and explained below, which shows the device's rotation relative to the neutral angle represented as a semicircle. This mapping more closely approximates the physical geometry when the device is rotated.

Let p_1 be the point on the outside of the semicircle at a rotation of θ from the neutral angle. Note that we restrict θ to a maximum magnitude of $\pm 60°$. Let α be a threshold angle and t the point on the outside of the semicircle at a rotation of α from the neutral angle. Let L_1 be the line tangent to the semicircle at the neutral angle and let L_2 be the line parallel to L_1 which passes through t. Then let p_2 be the point on L_2 closest to p_1. The scrolling speed is proportional to the distance d between p_1 and p_2, scaled to be in the range of ± 0 to 5000 pixels per second. We chose to set α to 6.5° to give a reasonably sized safe region of 13° where no scrolling is performed while also allowing scrolling to be performed easily when required.

The $\pm 60°$ of angular rotation is a fairly conservative range as NASA anthropometry and biomechanics measures (NASA 1991) show fifth percentile wrist flexion and supination values of 40.1 and 83.4° respectively, and 95th percentile values of 78.0 and 125.4° (applicable if the device is held in the hand with the fingers pointing away from the face), and forearm supination values of 83.4 to 125.8° for the 5th and 95th percentiles (applicable if the device is held with the fingers pointing across the body. Thus, the range of motion is attainable without need for finger-based manipulation to achieve most tilt angles. The screen is also reasonably easily viewed within the $\pm 60°$ of angular rotation; less so at the extremities, but at very high scroll speeds users are less likely to be concerned about precise visual details.

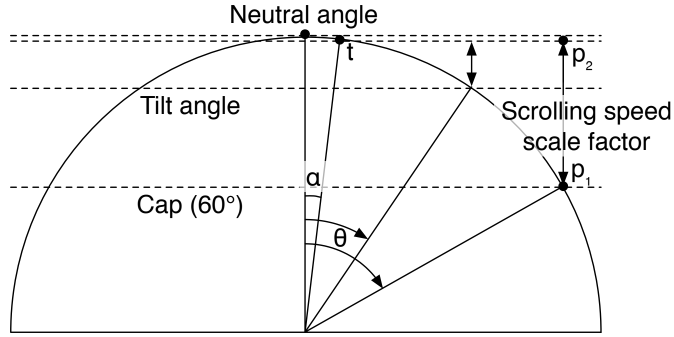

Figure 1: Tilt speed calculations

EXPERIMENTAL METHODOLOGY

Our methodology for empirically comparing mobile interfaces for scroll-based analysis activities uses two different task types. The first, called 'text', is based on skim reading a passage of text, but operationalises the activity by having participants count the number of occurrences of one specific target word. The aim is to be similar to typical day to day tasks, such as reading through email subject headers, skim reading a webpage, or scanning through a text file for a salient keyword. While the exact cognitive processes involved in each of these scenarios differ, the scrolling behaviour is essentially the same and our tasks provide useful insights into real world usage.

The second task type, called 'grid', aims to eliminate all extraneous factors from the interaction essence of the scrolling interfaces. For example, the text based task is somewhat similar to real user activities, but it suffers problems in that some words are more recognisable and familiar than others. The grid task is based on counting differently coloured dots in a grid.

Both tasks begin with a piece of text stating "Scroll down to begin". The task content begins immediately below the initial viewport, requiring users to initiate scrolling before any task data is displayed. Timing begins when the first scroll action is initiated. Tasks are completed when the participant scrolls to the bottom of the view and presses a "Done" button to end the task timing. Having done so, they are asked to answer a question concerning a count of the data items displayed during the task. The method therefore provides data on both the speed and accuracy of the tasks, as described further below.

Text Tasks

Each text task consists of a block of text rendered in left aligned 22 point Helvetica font, as shown in Figure 2a. The blocks of text contain 120 five letter English words and are typically 25 or 26 lines long. One target word occurs multiple times, with the number of occurrences following a Binomial distribution with $p = 0.05$ and $n = 120$, meaning that when generating each word there was a 5% probability that it will be the target word and a 95% chance that it will be different, resulting in an expected count of 6 target word occurrences (on average) per trial. This procedure results in random placement of the target words, and allows the possibility of clustering (although unlikely). The words generated for each trial are independent (except that the target word is the same across all trials for each participant). Participants are prompted to count the number of times this word appears, with a reminder of the target word displayed before each trial together with the “Scroll down to begin” prompt. Following the task participants enter how many times they counted the target word. Participants are instructed to complete the trials as quickly and accurately as possible.

Grid Tasks

Grid tasks display a grid of circles on a white background, as shown in Figure 2b – six circles per row, and 25 rows in total. Circles are randomly either grey or black, each with a 50% probability. We ensured that all grids had between 45% and 55% of their circles coloured black.

Our initial method involved simply counting the number of black circles. However, a small pilot study indicated that the speed at which participants could count in their heads was the major determinant in performance. We then modified the task to counting the number of rows with even parity, which reduced the cognitive load – rather than counting relatively high numbers (up to eighty or so), the participants make a quick judgement about each row, counting up to thirteen (on average).

On completing each trial participants were asked to enter how many rows had an even number of black circles (including zero). Like text tasks, they were asked to complete the trials as quickly and accurately as possible.

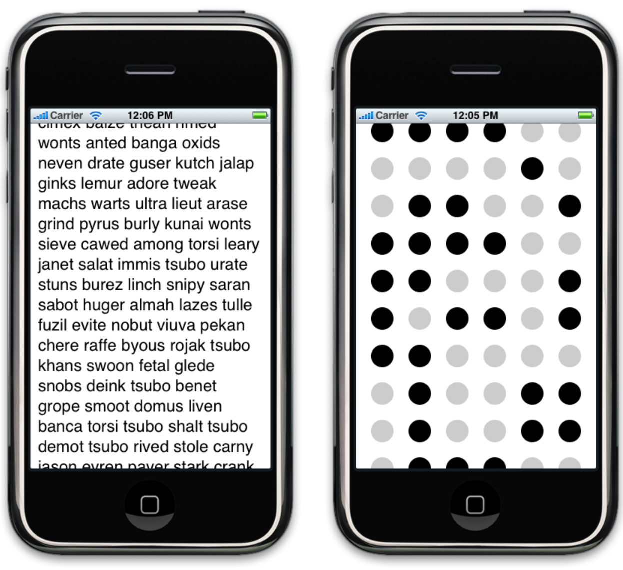

(a) Text tasks **(b) Grid tasks**

Figure 2: Tasks given to evaluation participants

Walking and Stationary Conditions

Mobile devices are (obviously) used in both mobile and stationary settings, and it is important to understand their performance in both. Our methodology inspects both with text and grid tasks. Stationary tasks are conducted while participants are seated. Mobile tasks are conducted while walking.

For walking tasks, we add an additional dependent variable – percentage preferred walking speed (PPWS) (Petrie, Furner & Strothotte 1998, Pirhonen, Brewster & Holguin 2002) – which has been previously used to evaluate electronic travel aids as well as mobile devices. PPWS gives an indication of the effect of the device’s interface on walking speed: as Pirhonen et al. state, *the further below their normal walking speed that they walked the more negative the effect the device was having on them*.

We use a similar method to Pirhonen et al. to find participants’ preferred walking speeds. Two chairs are placed eight metres apart with a third in the centre. Participants are asked to walk 3 laps weaving around the chairs. Each lap is roughly 19 metres taking into account the turns around the chairs. The times of their second and third laps are averaged to calculate each participant’s preferred walking speed.

EXPERIMENT COMPARING FLICK AND TILT SCROLLING

We used the experimental methodology described above to compare the state of the art flick scrolling interface with tilt scrolling for analysis tasks.

Participants and Apparatus

14 postgraduate computer science students (one female) participated in an experiment comparing flick scrolling and tilt scrolling. Their mean age was 26 and all had normal or corrected to normal vision. Five had previous experience with an iPod touch™or iPhone™.

The evaluation was performed on a second generation iPod touch™running iPhone OS 2.2. The display’s resolution was 480×320 pixels and it was always oriented in portrait.

Procedure and Design

At the start of the experiment, participants were given a brief introduction and asked to provide basic demographic information. They then carried out the preferred walking speed calibration, as described above. Participants were then given a quick demonstration of flick scrolling and tilt scrolling and given one minute to practice both methods on a sample view 1000 pixels high which allowed only vertical scrolling. This view allowed both forms of scrolling while later views would only allow a single method of scrolling. They were then shown an example of a text task and a grid task. They then completed the tasks, completing all tasks for one scrolling interface before continuing to the other. After completing tasks for each interface participants completed NASA Task Load Index (TLX) worksheets (Hart & Staveland 1988). Once all tasks were complete, participants answered several general questions and gave comments.

All factors were counterbalanced, including movement type (walking and stationary). Each condition had three trials, with participants being informed that the first in each condition was a practice trial which was not included in statistical analysis. Participants were prompted before each tilt scrolling task to reset the device to a neutral angle, and the neutral angle was then reset to the current rotation at the beginning of each tilt scrolling trial.

The dependent variables are analysed using two 2×2 repeated measures analyses of variance (ANOVA) for factors *interface type* (flick scrolling and tilt scrolling) and *task type* (counting dots and reading a passage). This analysis is applied separately to stationary and walking conditions. The dependent variables are task time and error rate (for walking and stationary), and also PPWS for the walking condition.

RESULTS

During the experiment, six tasks were repeated due to outside interruptions or participants not properly following instructions. The data for the original attempts was not included in the analysis. Additionally, one outlier trial was removed which had a trial time of greater than three standard deviations above the global mean.

Task Times

For stationary tasks (shown in Figure 3a), there was a significant main effect of *interface* ($F_{1,13} = 5.58, p = 0.03$), showing that tilt scrolling (mean: 17.1 seconds) is faster than flick scrolling (mean: 18.6 seconds). This difference was similar for both grid and text tasks, resulting in no significant *interface*×*task type* interaction ($F_{1,13} < 0.1$). There was no significant main effect of *task type* ($F_{1,13} < 1$), but this is an unimportant methodological coincidence of similarity between text and grid tasks.

For walking tasks (shown in Figure 3b), there were no significant effects (e.g. main effect of *task type*, $F_{1,13} = 0.070, p = 0.796$), suggesting similar interface performance for both task types.

Error Rates

We analysed the difference between actual counts and participants' counts for each task. There was no significant effect of interface type for either stationary ($F_{1,13} = 1.677, p = 0.218$) or moving tasks ($F_{1,13} = 0.052, p = 0.824$). Mean errors are summarised in Figure 4. This suggests that the performance benefits of tilt scrolling shown above are *not* due to participants differently addressing the speed-accuracy trade off when tilt scrolling; indeed, the trends shown in Figure 3a (significant) and Figure 4a (not significant) suggest both faster and more accurate performance with tilt scrolling while stationary.

There was a significant difference between task type for both stationary tasks ($F_{1,13} = 9.782, p < 0.01$) and moving tasks ($F_{1,13} = 10.367, p < 0.01$), with grid tasks having significantly fewer errors. As above, this is best explained as a methodological coincidence.

Percentage Preferred Walking Speeds (PPWS)

We found a significant difference ($F_{1,13} = 6.438$, $p = 0.025$) between PPWS for flick scrolling (mean of 67.2% of normal walking speed) and tilt scrolling (64.2%) – participants walked slightly slower when using tilt scrolling. *Task type* also showed a significant main effect, with participants walking slightly more slowly for text tasks ($F_{1,13} = 10.449$, $p < 0.01$), presumably due to the additional difficulty of scanning words rather than filled and unfilled circles. Finally, there was a marginal *interface*×*task type* interaction ($F_{1,13} = 3.246$, $p = 0.095$), with little difference between interfaces for grid tasks but larger difference in text tasks. These results are summarised in Figure 5.

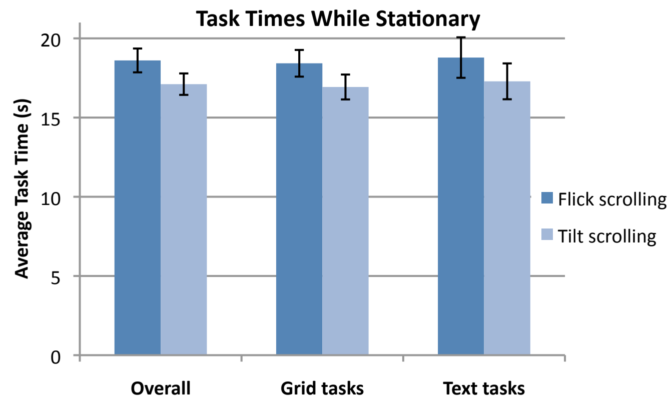

(a) Stationary tasks

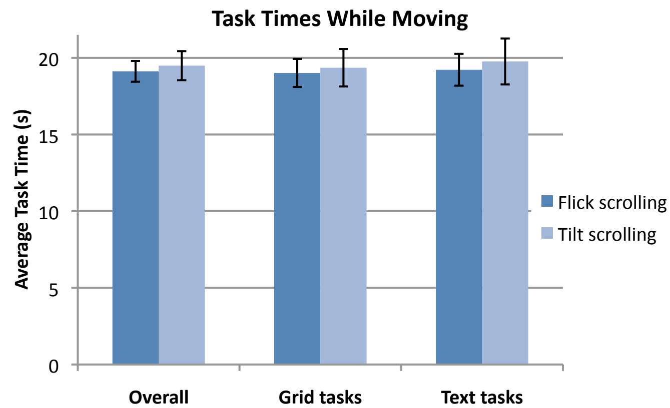

(b) Moving tasks

Figure 3: Task times for stationary and moving tasks. Error bars show standard error.

Preferences

Of the 14 participants, six preferred tilt scrolling while stationary, four preferred it while moving, but only two preferred it overall. Most of the participants who preferred tilt scrolling while stationary but flick scrolling while moving stated that this was because their gait added noise to the device's tilt, causing the scrolling to become jerky. Interestingly, three participants preferred the exact opposite: flick scrolling while stationary and tilt scrolling while moving. These participants explained that one handed control was preferable while walking and that this was much easier with tilt scrolling. Several participants also mentioned that a disadvantage of flick scrolling is that the need of a finger on the screen means that there is effectively less screen real estate. This diversity in preferences should be carefully considered by future interface designers.

The mean NASA-TLX responses were better for flick scrolling than tilt scrolling in all measures, significantly so for mental demand (Wilcoxon $z = 1.64$, $p = 0.05$) and frustration level (Wilcoxon $z = 1.64$, $p = 0.05$). A summary of mean responses is shown in Figure 6.

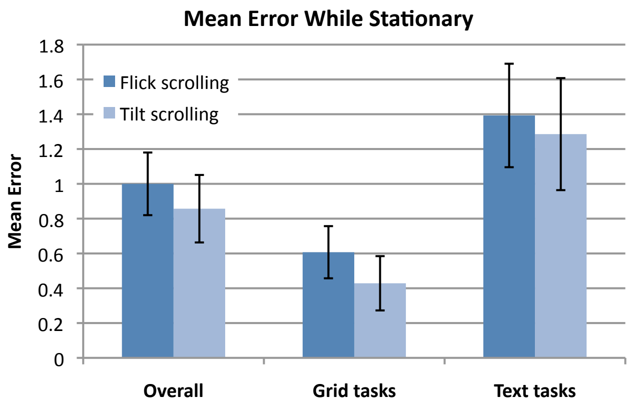

(a) Stationary tasks

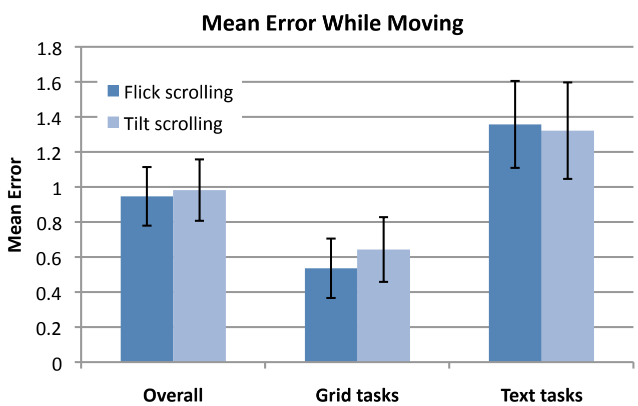

(b) Moving tasks

Figure 4: Mean difference between correct counts and user counts. Error bars show standard error.

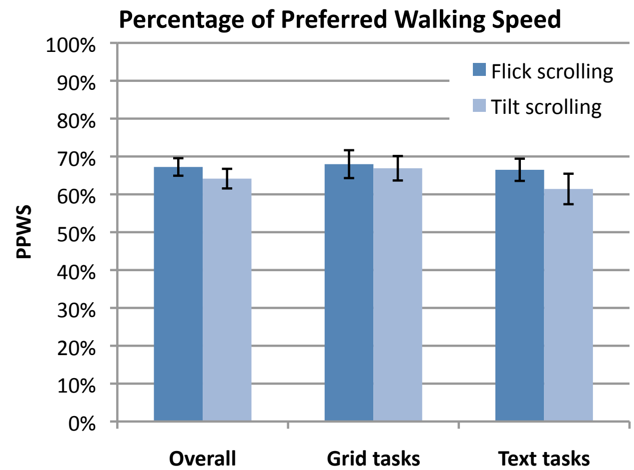

Figure 5: Percentage of preferred walking speed (PPWS). Error bars show standard error.

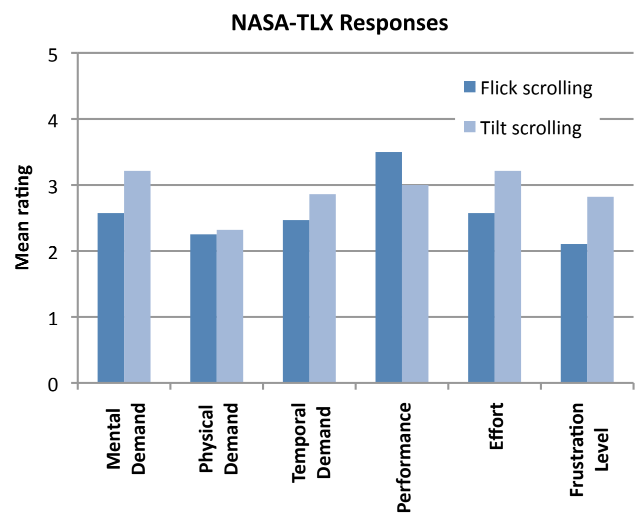

Figure 6: Mean NASA-TLX responses. Lower numbers are better except for performance.

DISCUSSION AND FURTHER WORK

Summarising the main results, tilt scrolling significantly outperformed flick scrolling when stationary, with reliably faster task completion times and an apparent trend to fewer errors. Both interfaces performed similarly while moving, although participants walked more slowly with tilt scrolling. Most participants preferred the flick scrolling interface overall.

As well as providing new empirical evidence on the effectiveness of these two important mobile scrolling interfaces, the results also demonstrate that our methodology is successful in revealing nuanced performance characteristics.

Limitations of the Study

Several issues related to tilt scrolling may have negatively impacted the tilt scrolling results:

1. Several participants mentioned that with practice, they may have viewed tilt scrolling more favourably. It may indeed be the case that tilt scrolling has a longer learning curve than other methods of scrolling as its input method is so different to conventional scrolling interfaces and users must also learn to compensate for accidental device movement. Additionally, five of the participants had previous experience with an iPod touch™ or iPhone™ so would have been previously exposed to flick scrolling.

2. Drifting was observed for some participants where the device would be rotated for scrolling down at the end of one task and would not be moved back sufficiently before starting the next task, despite instruction to do so. This caused some participants to start certain tasks at uncomfortable viewing angles.

3. Several participants expected the device to scroll in the opposite direction than it did for tilt scrolling and this may have impacted their performance. This is consistent with previous research showing split views about which direction is preferred (Bartlett 2000).

4. Subjective responses were taken after doing all tasks for an interface, both stationary and moving. Tilt scrolling responses may have therefore been negatively influenced by bad experiences with one movement type. In hindsight,

it would have been useful to gather separate NASA-TLX responses for stationary and moving tasks.

5. As a relatively new technique, there will inevitably be many future improvements to tilt scrolling that improve user performance, particularly with the angle to scrolling velocity mapping. To perhaps a lesser extent, this also applies to flick scrolling; for example, Aliakseyeu et al. (2008) found that the flick scrolling implementation on the iPhone™was not the best for either performance or user preferences compared to other flick scrolling implementations.

A key assumption of this research is that users spend significant time scrolling for reading or analysis tasks rather than to navigate large distances. To back up this premise, further research into the use of scrolling on mobile devices is needed to prove that this type of interaction is indeed common.

It would also be beneficial to compare flick scrolling and tilt scrolling with other scrolling techniques for reading and analysis tasks.

Insights

Perhaps the most interesting finding of the evaluation is that while most participants preferred flick scrolling and thought they had performed better while using it, they actually tended to perform better with tilt scrolling. One participant in particular commented that he thought he had performed very poorly with tilt scrolling due to a lack of familiarity. He was later shocked to learn that he had actually performed better with tilt scrolling. This example illustrates that given a choice, users will not always choose the most efficient interface.

It was also surprising to see some participants prefer tilt scrolling only while stationary and others to prefer it only while moving. Based on participants' preferences and comments, we recommend that interface designers consider offering tilt scrolling as an alternative to existing methods, but that they not enforce its use. A suitable method of toggling tilt scrolling should be available. Additionally, users should be able to select the direction of the tilt to scrolling velocity mapping.

CONCLUSIONS

Tilt scrolling has not been implemented in any mainstream interfaces to date, although it is extensively used for analogous two dimensional movement effects in many mobile games. We have demonstrated that it performs at least as well, and in some cases better, than flick scrolling, providing a strong motivation for implementing it as an additional scrolling mechanism in future interfaces. Its main barrier is user perceptions that it does not perform as well, although it remains to be seen if these perceptions persist with more practice.

We have also contributed a methodology for comparing scrolling interfaces for reading tasks. We hypothesise that this context accounts for a large percentage of navigation tasks and hope that this inspires other researchers to investigate ideal navigation methods for such tasks.

ACKNOWLEDGEMENTS

Thanks to the anonymous reviewers for their comments and to the evaluation participants for their time.

This work was partially funded by New Zealand Royal Society Marsden Grant 07-UOC-013.

REFERENCES

Alexander, J., Cockburn, A., Fitchett, S., Gutwin, C. & Greenberg, S. (2009), Revisiting read wear: analysis, design, and evaluation of a footprints scrollbar, *in* 'CHI '09: Proceedings of the 27th international conference on Human factors in computing systems', ACM, New York, NY, USA, pp. 1665–1674.

Aliakseyeu, D., Irani, P., Lucero, A. & Subramanian, S. (2008), Multi-flick: an evaluation of flick-based scrolling techniques for pen interfaces, *in* 'CHI '08: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems', ACM, New York, NY, USA, pp. 1689–1698.

Andersen, T. H. (2005), A simple movement time model for scrolling, *in* 'CHI '05: CHI '05 extended abstracts on Human factors in computing systems', ACM, New York, NY, USA, pp. 1180–1183.

Appert, C. & Fekete, J.-D. (2006), Orthozoom scroller: 1d multi-scale navigation, *in* 'CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems', ACM, New York, NY, USA, pp. 21–30.

Bartlett, J. F. (2000), 'Rock 'n' scroll is here to stay', *IEEE Comput. Graph. Appl.* **20**(3), 40–45.

Baudisch, P. & Chu, G. (2009), Back-of-device interaction allows creating very small touch devices, *in* 'CHI '09: Proceedings of the 27th international conference on Human factors in computing systems', ACM, New York, NY, USA, pp. 1923–1932.

Cho, S.-J., Choi, C., Sung, Y., Lee, K., Kim, Y.-B. & Murray-Smith, R. (2007), Dynamics of tilt-based browsing on mobile devices, *in* 'CHI '07: CHI '07 extended abstracts on Human factors in computing systems', ACM, New York, NY, USA, pp. 1947–1952.

Cockburn, A., Savage, J. & Wallace, A. (2005), Tuning and testing scrolling interfaces that automatically zoom, *in* 'CHI '05: Proceedings of the SIGCHI conference on Human factors in computing systems', ACM, New York, NY, USA, pp. 71–80.

Eslambolchilar, P. & Murray-smith, R. (2004), Tilt-based automatic zooming and scaling in mobile devices - a state-space implementation, *in* 'Proceedings of Mobile HCI 2004', Springer, pp. 120–131.

Eslambolchilar, P., Williamson, J. & Murray-Smith, R. (2004), Multimodal feedback for tilt controlled speed dependent automatic zooming, *in* 'UIST '04: Proceedings of the 17th annual ACM symposium on User interface software and technology'.

Geißler, J. (1998), Shuffle, throw or take it! working efficiently with an interactive wall, *in* 'CHI '98: CHI 98 conference summary on Human factors in computing systems', ACM, New York, NY, USA, pp. 265–266.

Harrison, B. L., Fishkin, K. P., Gujar, A., Mochon, C. & Want, R. (1998), Squeeze me, hold me, tilt me! an exploration of manipulative user interfaces, *in* 'CHI '98: Proceedings of the SIGCHI conference on Human factors in computing systems', ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, pp. 17–24.

Hart, S. & Staveland, L. (1988), Development of nasa-tlx (task load index): Results of empirical and theoretical research., *in* P. Hancock, ed., 'Human Mental Workload', pp. 139–183.

Hinckley, K., Cutrell, E., Bathiche, S. & Muss, T. (2002), Quantitative analysis of scrolling techniques, *in* 'CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems', ACM, New York, NY, USA, pp. 65–72.

Hinckley, K., Pierce, J., Sinclair, M. & Horvitz, E. (2000), Sensing techniques for mobile interaction, *in* 'UIST '00: Proceedings of the 13th annual ACM symposium on User interface software and technology', ACM, New York, NY, USA, pp. 91–100.

Igarashi, T. & Hinckley, K. (2000), Speed-dependent automatic zooming for browsing large documents, *in* 'UIST '00: Proceedings of the 13th annual ACM symposium on User interface software and technology', ACM, New York, NY, USA, pp. 139–148.

Kaptelinin, V., Mäntylä, T. & Åström, J. (2002), Transient visual cues for scrolling: an empirical study, *in* 'CHI '02: CHI '02 extended abstracts on Human factors in computing systems', ACM, New York, NY, USA, pp. 620–621.

NASA (1991), Nasa man-systems integration standards: 3 anthoropometry and biomechanics, Technical Report NASA-MSIS Volume 1, part 3, National Aeronautics and Space Administration.

Petrie, H., Furner, S. & Strothotte, T. (1998), 'Design lifecycles and wearable computers for users with disabilities', Presented at The First Workshop on Human-Computer Interaction With Mobile Devices.

Pirhonen, A., Brewster, S. & Holguin, C. (2002), Gestural and audio metaphors as a means of control for mobile devices, *in* 'CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems', ACM, New York, NY, USA, pp. 291–298.

Poupyrev, I., Maruyama, S. & Rekimoto, J. (2002), Ambient touch: designing tactile interfaces for handheld devices, *in* 'UIST '02: Proceedings of the 15th annual ACM symposium on User interface software and technology', ACM, New York, NY, USA, pp. 51–60.

Ramos, G., Cockburn, A., Balakrishnan, R. & Beaudouin-Lafon, M. (2007), Pointing lenses: facilitating stylus input through visual-and motor-space magnification, *in* 'CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems', ACM, New York, NY, USA, pp. 757–766.

Reetz, A., Gutwin, C., Stach, T., Nacenta, M. & Subramanian, S. (2006), Superflick: a natural and efficient technique for long-distance object placement on digital tables, *in* 'GI '06: Proceedings of Graphics Interface 2006', Canadian Information Processing Society, Toronto, Ont., Canada, Canada, pp. 163–170.

Rekimoto, J. (1996), Tilting operations for small screen interfaces, *in* 'UIST '96: Proceedings of the 9th annual ACM symposium on User interface software and technology', ACM, New York, NY, USA, pp. 167–168.

Roth, V. & Turner, T. (2009), Bezel swipe: conflict-free scrolling and multiple selection on mobile touch screen devices, *in* 'CHI '09: Proceedings of the 27th international conference on Human factors in computing systems', ACM, New York, NY, USA, pp. 1523–1526.

Wigdor, D. & Balakrishnan, R. (2003), Tilttext: using tilt for text input to mobile phones, *in* 'UIST '03: Proceedings of the 16th annual ACM symposium on User interface software and technology', ACM, New York, NY, USA, pp. 81–90.