

TiltText: Using Tilt for Text Input to Mobile Phones

Daniel Wigdor, Ravin Balakrishnan
Department of Computer Science
University of Toronto
dwigdor | ravin @dgp.toronto.edu
www.dgp.toronto.edu

ABSTRACT

TiltText, a new technique for entering text into a mobile phone is described. The standard 12-button text entry keypad of a mobile phone forces ambiguity when the 26-letter Roman alphabet is mapped in the traditional manner onto keys 2-9. The *TiltText* technique uses the orientation of the phone to resolve this ambiguity, by tilting the phone in one of four directions to choose which character on a particular key to enter. We first discuss implementation strategies, and then present the results of a controlled experiment comparing *TiltText* to *MultiTap*, the most common text entry technique. The experiment included 10 participants who each entered a total of 640 phrases of text chosen from a standard corpus, over a period of about five hours. The results show that text entry speed including correction for errors using *TiltText* was 23% faster than *MultiTap* by the end of the experiment, despite a higher error rate for *TiltText*. *TiltText* is thus amongst the fastest known language-independent techniques for entering text into mobile phones.

Keywords: Text entry, mobile phones, tilt input

INTRODUCTION

Most mobile phones are equipped with a simple 12-button keypad, which is an inherently poor tool for generating phrases for a 26-letter alphabet. Using traditional text-entry techniques, such as *MultiTap*, an average text message of 7 words requires roughly 70 key presses. Given estimates (www.gsmworld.com) that in 2003 nearly 500 billion text messages will be sent worldwide from mobile phones, entry using current techniques will require approximately 35 trillion key presses worldwide this year. While much research effort has gone into devising a variety of more efficient text input techniques [9, 15] which have all shown various improvements to the status-quo, none has yet emerged as a new standard. As such, there remains considerable opportunity for researchers to influence this area by developing new techniques.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. UIST '03 Vancouver, BC, Canada

© 2003 ACM 1-58113-636-6/03/0010 \$5.00

TiltText

We have developed a new text input technique, called *TiltText*, which uses the standard 12-button mobile phone keypad augmented with a low-cost tilt sensor. Similar to *TiltType* described by Partridge et al. [10], *TiltText* uses a combination of a button press and tilting of the device to determine the desired letter. Our technique differs from *TiltType* in the keypad used and in the sensing algorithms, which we discuss in detail later in this paper. The standard phone keypad mapping assigns three or four alphabetic characters, and one number, to each key. For example, the 2 key also has the characters a, b, and c assigned to it. *TiltText* assigns an additional mapping by specifying a tilt direction for each of the characters on a key, removing any ambiguity from the button press. The user presses a key while simultaneously tilting the phone in one of four directions (left, forward, right, back) to input the desired character (Figure 1). For example, pressing the 2 key and tilting to the left inputs the character a, while tilting to the right inputs the character c. By requiring only a single keypress and slight tilt to input alphanumeric characters, the overall speed of text entry can be increased. Further, unlike some techniques [15] that improve on the status quo, *TiltText* is not language dependent, and thus can be used by experts without visually attending to the display screen.

In this paper we first review related work, then discuss implementation issues, then present an experiment comparing the performance of *TiltText* to the most common existing technique, *MultiTap*. We conclude by discussing the characteristics of *TiltText* compared to other techniques, including implications for evaluation metrics.



Figure 1. *TiltText*. The center picture shows the untilted phone where pressing a key enters its numeric value. Left picture: left tilt enters first character on key. Top picture: forward tilt enters second character. Right picture: right tilt enters third character. Bottom picture: tilting towards the user enters fourth character if one exists for that key.

RELATED WORK

There are two areas of research that are relevant to our work: text input techniques for mobile phones, and the use of tilt transducers in mobile devices.

Text Input Techniques for Mobile Phones

A small number of mobile phones today utilize QWERTY style keypads that enable text entry with techniques similar to typing on a regular keyboard, albeit at a much smaller physical scale (e.g., Nokia 5510 www.nokia.com). More recently, hybrid devices that combine phones with PDAs, such as the Handspring Treo (www.handspring.com) and PocketPC Phone (www.microsoft.com), utilize pen-based text input techniques common to PDA's such as Graffiti. While these devices are making small inroads into the mobile phone market, the vast majority of mobile phones are equipped with the standard keypad (Figure 2) which has 12 keys: 0-9, *, and #.



Figure 2. Standard 12-key mobile phone keypad

Entering text from a 26 character alphabet using this keypad forces a mapping of more than one character per button of the keypad. A typical mapping has keys 2-9 representing either three or four characters, with space and punctuation mapped to the other buttons. All text input techniques that use this standard keypad have to somehow resolve the ambiguity that arises from this multiplexed mapping. There are three main techniques for overcoming this ambiguity: *MultiTap*, two-key, and linguistic disambiguation. We now review these techniques briefly, and refer the reader to Soukoreff and MacKenzie [15] for a more comprehensive review that is beyond the scope of the present paper.

MultiTap

MultiTap works by requiring the user to make multiple presses of each key to indicate which letter on that key is desired. For example, the letters pqr s traditionally appear on the 7 key. Pressing that key once yields p, twice q, etc. A problem arises when the user attempts to enter two consecutive letters on the same button. For example, tapping the 2 key three times could result in either c or ab. To overcome this, *MultiTap* employs a time-out on the button presses, typically 1-2 seconds, so that not pressing a button for the length of the timeout indicates that you are done entering that letter. Entering ab under this scheme has the user press the 2 key once for a, wait for the timeout, then press 2 twice more to enter b. To overcome the time overhead this incurs, many implementations add a "timeout kill" button that allows the user to skip the timeout. If we assume that 0 is the timeout kill button, this makes the sequence of button presses to enter ab: 2, 0, 2, 2.

MultiTap eliminates any ambiguity, but can be quite slow, with a keystrokes per character (KSPC) rate of approximately 2.03 [8].

Two-key Disambiguation

The two-key technique requires the user to press two keys in quick succession to enter a character. The first keypress selects the appropriate group of characters, while the second identifies the position of the desired character within that group. For example, to enter the character e, the user presses the 3 key to select the group def, followed by the 2 key since e is in the second position within the group. This technique, while quite simple, has failed to gain popularity for Roman alphabets. It has an obvious KSPC rate of 2.

Linguistic Disambiguation

There are a number of linguistic disambiguation schemes that utilize knowledge of the language to aid the text entry process. One example is T9 (www.tegic.com) that renders all possible permutations of a sequence of button presses and looks them up in a dictionary. For example, the key sequence 5, 3, 8 could indicate any of 27 possible renderings (3x3x3 letters on each of those keys). Most of these renderings have no meaning, and so are rejected. Looking each of them up in a dictionary tells the system that only jet is an English word, and so it is the one rendered. Ambiguity can, however, arise if there is more than one valid rendering in the language, in which case the most common is presented. For example, the sequence 6, 6 could indicate either on or no. If the system renders the wrong word, a "next" key allows the user to cycle through the other valid permutations. An analysis of this technique for entering text from an English corpus found a KSPC close to 1 [8]. Newer linguistic disambiguation techniques such as *LetterWise* [9] and *WordWise* (www.eatoni.com) also perform similarly well, with subtle advantages over earlier techniques. While these all have excellent KSPC rates, the success of linguistic-based systems depends on the assumption that users tend to enter "English like" words when sending text messages. As Mackenzie et al. [9] note, users often use abbreviations, and not complete English when text messaging. Further, users of text messaging often communicate in acronyms or combinations of letters and numbers (e.g., b4 for before). Another problem with these linguistic techniques is that users have to visually monitor the screen in order to resolve potential ambiguities, whereas the *MultiTap* and two-key techniques can be operated "eyes-free" by skilled users.

As a result of these limitations of current keypad text input techniques, the quest for a widely applicable, low KSPC, text input technique continues.

Using Tilt Sensors in Mobile Devices

Several researchers have recently proposed interesting interaction techniques that are enabled by incorporating a low-cost tilt sensor within mobile devices [3-7, 10, 12, 13]. While some of this prior art (e.g., [3-7, 12]) do not concern

text entry techniques per se, they do add to the set of possible interactions that could take advantage of tilt sensors embedded in mobile devices, thus providing further justification for the incremental cost of the sensor.

Of particular relevance to our work are two techniques for text entry that use tilt information. Both of these techniques focus on very small devices lacking a large number of buttons, and were not optimized or evaluated for speed of entry. *Unigesture* [13] used tilt as an alternative to button pressing, eliminating the need for buttons for text entry. Rather than having the user make one of 8 ambiguous button presses (as is the present case with mobile phones), *Unigesture* has the user tilt the device in one of 7 directions to specify the group, or “zone”, of the character that is desired. The ambiguity of the tilt is then resolved by using dictionary-based disambiguation.

TiltType [10] refines *Unigesture* by adding the combination of button pressing and tilt for entering unambiguous text. *TiltType* was designed to enter text into a small, watch-like device with 4 buttons. Pressing a button triggered an on-screen display of the characters that could be entered by tilting the device in one of eight directions, the appropriate tilt was then made, and the button released. *TiltType* has the same root concept as our *TiltText* technique, in that tilt is used to disambiguate button presses.

Our present work builds upon *TiltType* in several significant ways. First, neither *TiltType* nor *Unigesture* were designed for use with mobile phone keypads, as we are proposing with our *TiltText* technique. We believe that using the standard mobile phone keypad will significantly increase the viability of tilting text input as a real, usable, technique. Second, while *TiltType* uses eight tilt directions, we only use a maximum of four tilt directions, reducing the accuracy demands on the user when tilting. Third, the algorithm used for detecting tilt in the *TiltType* technique is one which we dub *key tilt*, which, as is discussed later in our paper, is not the most optimal tilt detection mechanism for speedy text entry. We develop two alternative tilt detection mechanisms that improve upon *key tilt*. Finally, we present the results of a controlled experiment that provides the first set of usability data with regards to using tilt for text input.

DESIGN ISSUES

TiltText uses the orientation of the phone along two axes to disambiguate the meaning of button presses. Tilting the phone to the left selects the first letter of the key, away from the body the second, to the right the third, and, if present, towards the body the fourth (see figure 1). Pressing a key without tilting results in entering the numeric value of the key. Space and backspace operations are carried out by pressing unambiguous single-function buttons (as in *MultiTap*).

Supporting both lowercase and uppercase characters would require a further disambiguation step since a total of seven characters per key would need to be mapped for keys 2-6

and 8, and nine characters each for the 7 and 9 keys (Figure 2). Adding case sensitivity could be done by either requiring the pressing of a “sticky” shift-key, or considering the magnitude of the tilt as a disambiguator where greater magnitude tilts result in upper case letters, as Figure 3 illustrates. The latter technique, however, would likely make eyes-free entry more difficult.

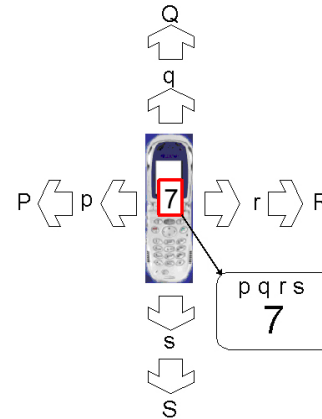


Figure 3. Uppercase text entry with *TiltText*. Tilting beyond a threshold makes the character uppercase.

Techniques for Calculating Tilt

The tilt of the phone is taken as whichever direction has the greatest tilt relative to an initial “origin” value. After exploring various options during our development process, we have found that there are three main ways to determine the tilt value: *key tilt*, *absolute tilt*, and *relative tilt*.

Key Tilt

With this technique, first seen in the *TiltType* work, the amount of tilt is calculated as the difference in the value of the tilt sensors at key down and key up. This requires the user to carry out three distinct movements once the button has been located: push the button, tilt the phone, release the button. We conducted a pilot experiment comparing a *TiltText* implementation that used *key tilt*, and found that user performance with this implementation was much slower than the traditional *MultiTap* technique. For this reason, *key tilt* was not used in our final experiment.

Absolute Tilt

This technique compares the tilt sensor’s value at any given time to a “fixed” absolute origin. Only two distinct movements are required to enter a character: tilt the phone, then press the key. In contrast, *key tilt* requires that the key be pressed first, phone tilted, then key released. However, this approach is also not ideal, since in practice users do not maintain a constant arm posture. In order for the tilt value to be meaningful, the fixed origin will have to be reset every time the user’s gross arm posture changes. Further, when using *TiltText* to enter two characters requiring tilt in opposite directions, more movement is required using this absolute approach, since the first tilt must be undone, then the new tilt applied. For example, entering the letters ac using the 2 key requires an initial tilt of some angle Δ to

the left to enter the a. Then, the user has to tilt the same angle Δ in the reverse direction to return to the origin, before tilting another angle θ to the right to enter the letter c. The total amount of movement is $2\Delta + \theta$, instead of the smaller $\Delta + \theta$ that one may expect. However, one advantage of this method over *key tilt* is that if successive characters with the same tilt direction are to be entered, then the user can keep the phone tilted at that direction for the successive keypresses.

Relative Tilt

This approach calculates the tilt relative to a floating origin that is set when a tilt gesture begins. The beginning of a gesture is determined by continuously watching for a change in orientation or a change in the direction of a tilting gesture. This approach solves both problems of the absolute tilt method. Since all tilts are relative to the beginning of the gesture, there is no absolute origin that need be reset when changing arm position. Further, opposite direction tilts do not require double tilting, since the second tilt's origin is the end of the first tilt's gesture. So, entering the letters ac requires a tilt of some angle Δ to the left to enter a, then another tilt of angle θ to the right to enter the c, for a total movement of $\Delta + \theta$. Note that, just like with *absolute tilt*, when entering only letters, we can enter successive characters with the same tilt direction without re-tilting the phone, by looking at the last significant tilt.

EVALUATION

Goals

We sought to compare the performance of *TiltText* to other techniques for text entry into mobile phones. For this first experiment, we chose as a comparison the most commonly implemented technique, *MultiTap*, because it is the common baseline in almost every other evaluation of text entry techniques reported to date [9, 14, 15]. As such, while the present study only directly compares *TiltText* to *MultiTap*, we can indirectly make comparisons of *TiltText*'s performance relative to other techniques via the results reported in the literature.

Apparatus

Hardware

We used a Motorola i95cl phone. The phone was equipped with an Analog Devices ADXL202EB-232 2-axis accelerometer board to enable tilt sensing, connected to the phone via a serial cable (with the addition of an external power line).

An implementation of a relative tilt system would require regular sampling from the tilt sensor. Unfortunately, our hardware allowed only a reliable rate of ~ 10 Hz. Pilot studies, using a relative tilt implementation of *TiltText*, showed that participants' text entry speed increased to a point such that our mechanism for determining origins (repeated sampling of tilt at 10Hz) proved insufficient for determining tilt accurately. The resulting device-caused errors in recognition made entering text at higher speeds

frustrating. From these pilot studies, we believe that rates of 20-50Hz would be required for a robust relative tilt implementation. As a result, despite the limitations discussed earlier, we had to implement an absolute tilt approach, allowing the user to reset the origin at any time by holding the phone at the desired orientation and pressing "0". The additional movement required by this approach, however, is acceptable for our evaluation purposes because if we can demonstrate that *TiltText* performs well despite this additional movement, then any more robust implementation using a relative tilt approach can only do better. In other words, our evaluation is biased against our new technique.

Because the ADXL board can detect a tilt of only fractions of a degree, only a very small tilt of the phone is necessary to disambiguate a button press. The maximum of the tilt in either axis was taken to be the intended tilt, with a 10% bias towards forward/back. This bias is included based on our pilot studies which revealed a tendency to pitch to the dominant side when tilting forward with the wrist.

Software

The software to read tilts and render text, as well as conduct the experiment, was written in Java 2 Micro-Edition (source at www.dgp.toronto.edu/research/tilttext) using classes from both the Mobile Devices Information Profile (MIDP 1.0) and proprietary i95cl specific classes.

The experiment was conducted entirely on the mobile phone rather than simulating a mobile phone keypad on some other device. All software, including those implementing the text entry techniques, and data presentation and collection software used in the experiment ran on the phone. No connection to an external computing device beyond the tilt sensor was required. Of the major techniques for text entry into mobile phones evaluated in the literature over the last several years, this seems to be the first to do so on an actual mobile phone keypad / display. We believe that this helps to more closely represent real use, and so enhances the external validity of our experiment.

Our *MultiTap* implementation used the i95cl's built-in *MultiTap* engine, with a 2 second timeout and timeout kill. We only considered lowercase text entry in this evaluation. As such, the *MultiTap* engine was modified slightly to remove characters from the key mapping that were not on the face of the button, so that the options available were only the lower case letters and numeral on the key. This matches the traditional *MultiTap* implementation in past experiments, such as *LetterWise* [9].

Participants

Ten participants volunteered for the experiment. They were recruited from within the university community. There were 5 men and 5 women of whom 3 were left-handed and 7 right-handed. Participants were pre-screened so that no one with any experience composing text using either technique was included. Participants did not receive any tangible compensation for their participation.

Procedure

Participants entered short phrases of text selected from among those in MacKenzie's English phrase dictionary (www.yorku.ca/mack/phrases2.txt). These phrases were chosen because they have been used in previous text entry studies involving *MultiTap* [9], allowing us to leverage this previous work. This corpus' high correlation of frequencies of letters to English is a benefit, although it does not take into account abbreviations commonly used in texting.

The desired text phrases were shown to participants on the screen on the phone. For consistency with past *MultiTap* experiments, participants were instructed to enter text only with the thumb of the hand with which they held the phone, and not to change hands during the experiment.

Timing began when participants entered the first character of the phrase, and ended when the phrase was entered completely and correctly. If an erroneous character was entered, the phone alerted the user by vibrating, and the user was required to correct their error. With this procedure, the end result is error-free in the sense that the correct phrase is captured. Also, the phrase completion time incorporates the time taken to correct for errors.

Before beginning each treatment, participants were told to read and understand the displayed phrase before entering it, and were given instructions for that treatment as follows:

MultiTap instructions: to enter a character using the *MultiTap* technique, first find the key that is labeled with that character. Press that key repeatedly until the desired character is reached. Press once for the first character, twice for the second, three times for the third, and, if present, four times for the fourth. Once you have found the correct letter, and are ready for the next one, you simply repeat the process. If the letter you wish to enter next is on the same key, you must first either press the "right" arrow on the phone or wait two seconds for the cursor to advance.

TiltText instructions: The technique works by tilting the phone in the direction of the letter you wish to enter, then pressing the key on which it is inscribed. For the first letter, tilt left. For the second letter, tilt forward. For the third letter, tilt to the right. For the fourth letter, tilt towards you. The direction of tilt is measured relative to the "centre" or "origin" position of the phone. You can reset the origin at any time by pressing the 0 key.

The experimenter then demonstrated the relevant technique. To ensure that participants understood how the technique worked, they were asked to enter a single phrase that would require tilting in all four directions for *TiltText*, or two successive letters on the same key for *MultiTap*.

Additional instructions were given for both techniques to describe space and delete keys, as well as to enter an extra space at the end of the phrase to indicate completion. The process for error correction was also explained to them. Participants were also directed to rest as they liked between phrases, but to continue as quickly as possible once they had started entering a phrase.

Design

A within-subjects design was used. Participants were randomly assigned to two groups of 5 participants each. The first group performed the experiment with the *MultiTap* technique first, followed by *TiltText*, while the second group did it in the reverse order.

For each technique, participants were asked to complete two sessions of 8 blocks of trials each. Each block required the entry of 2 identical practice phrases, followed by 20 different phrases randomly selected from the corpus. The selection of phrases for each of the 16 blocks were done before the experiment, and presented in the same order to each participant. Phrases were selected such that all blocks had similar average phrase length. The same set of phrases and blocks were used for both techniques. In other words, all participants were required to enter identical phrases in the same order, the only difference being which technique they used first. Participants were asked to rest for at least 5 minutes between each block, and each session of 8 blocks each was conducted on separate days. At least 24 hours passed between sessions for the different techniques to limit interference. In summary, the design was as follows:

10 participants x
2 techniques (*MultiTap* and *TiltText*) x
2 sessions per technique x
8 blocks per session x
20 phrases per block (excluding practice phrases)
= 6400 phrases entered in total.

Results

Data Summary

The data collected from 10 participants took an average of 10.3 minutes per block. A total of 145360 correct characters of input were entered for the 6400 phrases.

Physical Comfort

Some participants reported that their thumb became sore while using both techniques. When this was reported, the participants were encouraged to rest until they felt comfortable to proceed. No participant reported pain or discomfort in their wrist or arms.

Text Entry Speed

We use the standard wpm (words-per-minute) measure to describe text entry speed. This is traditionally calculated as characters per second * 60 / 5. Because timing in our experiment started only after entering the first character, that character should not be included in calculations of entry speed. Thus, for the purposes of these computations, the length of a phrase is n-1 characters. Also, to signify completion, users had to enter an extra space at the end of each phrase. However, our timing considers the entry of the last real character of the phrase to be the end time.

The average text entry speed for all blocks were 11.76 wpm and 10.11 wpm for *TiltText* and *MultiTap* respectively. Overall, *TiltText* was 16.3% faster than *MultiTap*.

The means for the first block of trials were 7.42 wpm and 7.53 wpm, for *TiltText* and *MultiTap* respectively.

Performance in both techniques increased steadily, with the means for the last (16th) block of trials of 13.57 wpm for *TiltText* and 11.04 wpm for *MultiTap*. While subjects performed marginally better with *MultiTap* initially, they improved considerably faster with *TiltText*, with the spread between the techniques reaching 22.9% in favour of *TiltText* by the end of the experiment.

Analysis of variance indicated significant main effects for technique ($F_{1,8} = 615.8, p < .0001$), and block ($F_{15,120} = 145.2, p < .0001$). There was also a significant technique x block interaction ($F_{15,120} = 20.5, p < .0001$), indicating that participants improved at different rates for the different techniques. Figure 4 illustrates these effects.

The use of a within-subjects design in our experiment raises the possibility of transfer effects. We attempted to balance these transfer effects by assigning subjects to two different groups, one group using *TiltText* first followed by *MultiTap*, and the other group doing it in the reverse order. However, as Poulton [11] points out, for this counterbalancing to work, any skill transfer that occurs must be symmetric. Analysis of variance showed a significant order x technique interaction ($F_{15,120} = 481.5, p < .0001$), indicating that our participants exhibited *asymmetric* transfer effects. This essentially means that participants' experience with the first technique they used affected their performance with the second technique by different amounts depending on whether the first technique was *TiltText* or *MultiTap*. As such, we did additional analyses by looking only at the data for the first technique participants performed (i.e., half the experimental data). This in effect results in a between-subjects design. Analysis of this data indicated effects similar to that of the full within-subjects data set: significant main effects for technique ($F_{1,8} = 120.9, p < .0001$), block ($F_{15,120} = 80.6, p < .0001$), and a significant technique x block interaction ($F_{15,120} = 14.5, p < .0001$). Figure 5 illustrates these effects.

From our analysis and Figure 5, we see that without prior experience with either technique, *TiltText* started out performing worse than *MultiTap*, only crossing over at block 4. This is likely because *TiltText* required participants to master two distinctly different motor skills: pressing the key, and tilting the phone. Whereas *MultiTap* required only a single type of motor action: multiple presses of the key.

Figure 6 shows data after participants switched techniques (i.e., the second half of the experiment). We see here that *TiltText* starts off faster than *MultiTap*, indicating that participants' were able to take advantage of and transfer their previous experience with *MultiTap* in the first half of the experiment. This is a positive indication since it means that real users with lots of experience with *MultiTap* can transfer at least some of that skill if they switch to *TiltText*. Note, however, that there is quite a bit more variability in the performance for *TiltText*, as indicated by the poorer fit of the power curve as compared to Figures 4 and 5. This indicates that participants experienced some level of interference due to previous experience with *MultiTap*.

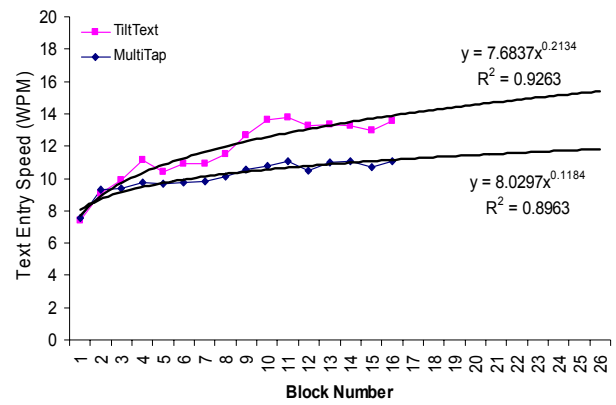


Figure 4. Entry speed (wpm) by technique and block for entire experiment. Best-fit power law of learning curve shows projected progress beyond the measured data in the first 16 blocks.

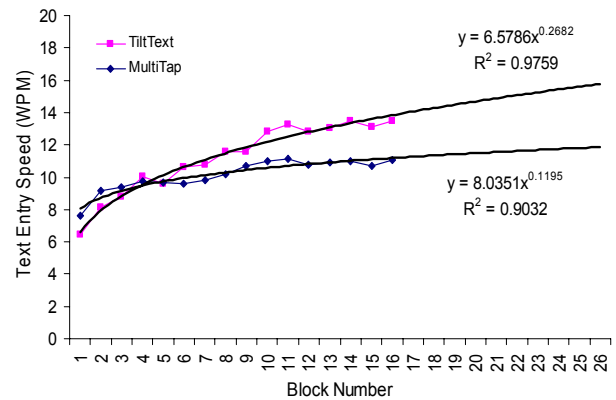


Figure 5. Entry speed (wpm) by technique and block, for the first half of the experiment, *before* participants switched techniques. Best-fit power law of learning curve shows projected progress beyond the measured data in the first 16 blocks.

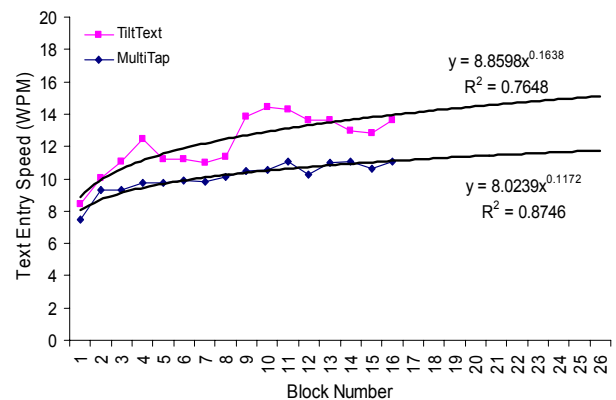


Figure 6. Entry speed (wpm) by technique and block, for the second half of the experiment, *after* participants switched techniques. Best-fit power law of learning curve shows projected progress beyond the measured data in the first 16 blocks.

Error Rates

Given that our experimental procedure required participants to make corrections as they proceeded, with an end result of a completely correctly entered phrase, the entry speed results discussed previously incorporate the cost of error correction. However, it is still useful to look at a more explicit error rate. We calculate percentage error rate as the number of characters entered that did not match the expected character, divided by the length of the phrase. In this case, we used the actual length of the phrase, and not (n-1) as in the wpm rate.

Overall, error rates were much higher for *TiltText* (11%) than for *MultiTap* (3%). This effect was statistically significant ($F_{1,8} = 1378.8, p < .0001$). There was also a significant effect for blocks ($F_{15,120} = 21.1, p < .0001$). A significant technique \times block interaction ($F_{15,120} = 23.3, p < .0001$) and Figure 7 indicate that while the error rates for *MultiTap* remain quite constant throughout the experiment, the error rates for *TiltText* drop rapidly over the first 8 blocks, and begin to asymptote from block 9 onwards.

As with the entry time analysis, a significant order \times technique interaction ($F_{15,120} = 168.9, p < .0001$) indicates that our participants exhibited *asymmetric* transfer effects. An analysis of the first half of the data (i.e., before participants switched techniques) indicates main effects similar to that of the entire dataset: technique ($F_{1,8} = 632.4, p < .0001$), blocks ($F_{15,120} = 7.3, p < .0001$), and technique \times block interaction ($F_{15,120} = 10.4, p < .0001$). Figure 8 illustrates. Interestingly, the mean *TiltText* error rate (8.6%) was lower than for the entire data set, indicating that the lack of interference from *MultiTap* was beneficial.

Figure 9 illustrates the data from trials in the second half of the experiment (i.e., after participants switched techniques). Comparing this to Figure 8, we see that the mean *TiltText* error rate of 13.5% is much higher than the mean 8.6% rate in the first half of the experiment. Further, the first 8 blocks of trials did not exhibit a constant trend for *TiltText*. Clearly, participants' previous experience with the *MultiTap* technique was having a detrimental effect on their ability to use *TiltText* right after the switch in technique occurs. This is consistent with the effect observed in the text entry speed data illustrated earlier in Figure 6. However, this effect wears off roughly after block 8. This is interesting, since this is the last block completed before participants broke for the day, resuming 24 hours later with block 9. The interference effect drops after this break.

To examine the cause of the higher *TiltText* error rate, it is useful to group errors into two categories: *tilt errors* and *button errors*. *Tilt errors* are those where the participant entered a letter that appears on the same button as the correct letter, indicating that an erroneous tilt was made. *Button errors* are those where the participant entered a letter that appeared on a different button.

We had anticipated that *button errors* would have similar rates for *TiltText* and *MultiTap*. However, the results showed a significant difference ($F_{1,8} = 320.67, p < .0001$),

where 3% of characters entered in the *MultiTap* trials were *button errors*, but only 1.5% of characters entered in *TiltText* showed this type of error.

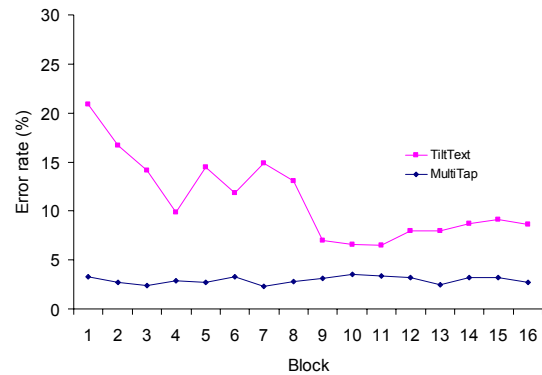


Figure 7. Error rates (%) by technique and block for entire experiment

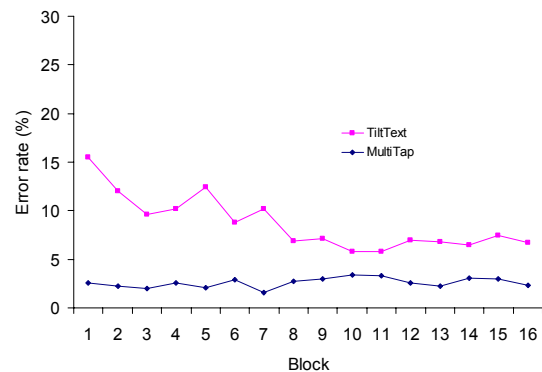


Figure 8. Error rates (%) by technique and block for the first half of the experiment, before participants switched techniques

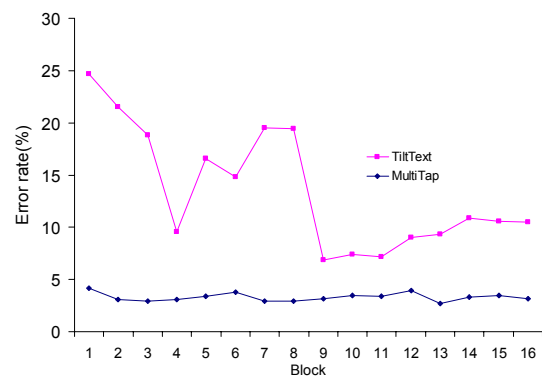


Figure 9. Error rates (%) by technique and block for the second half of the experiment, after participants switched techniques

Reconciling this low *button error* rate with the high overall error rate for *TiltText*, it is clear that most of the errors committed while using *TiltText* were *tilt errors*. Breaking down the *tilt error* rate by letter shows that participants committed significantly more errors for some letters than others ($F_{25,200} = 2.47, p < .001$), as Figure 10 illustrates.

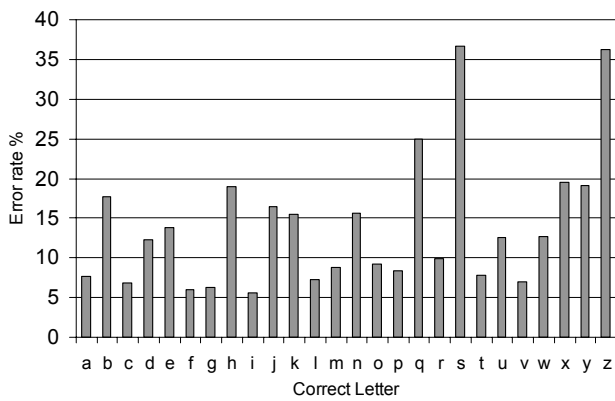


Figure 10. *Tilt* error rates (%) for each letter for the entire experiment.

Analysis of variance showed a significant main effect for tilt direction on *tilt error* rate ($F_{3,24} = 37.6, p < .0001$). Pairwise means comparisons showed a significantly higher *tilt error* rate for those letters requiring forward or backward tilting than those requiring right or left tilting. In particular, backwards tilt results in significantly higher errors than all the other tilt directions. Figure 11 illustrates this trend.

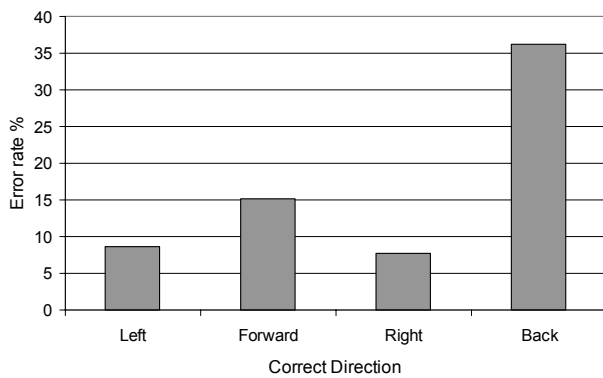


Figure 11. Error rates (%) by direction of tilt for correct letter for the entire experiment.

As was discussed previously, the overall *TiltText* error rate decreases with practice. As such, it is possible that the high *tilt error* rate for backward tilt (letters “s” and “z”) is due to the limited amount of practice users had with entering the letter “z”, which was only entered 3 times in the experiment by each participant. However, the other letter that required backward tilting, “s”, also showed a similarly high *tilt error* rate, despite being entered very frequently during the experiment. In other words, additional practice did not seem to decrease the backward *tilt error* rate significantly, indicating that users had an inherent difficulty with backward tilting actions.

When participants committed a *tilt error* for a letter requiring a left or right gesture, 82% of the time they ended up entering the forward-tilt letter on that button. This indicates that the 10% bias we introduced in our algorithm seems to overcompensate. Reducing this compensation factor may lower *tilt error* rate for left/right tilts.

DISCUSSION

As discussed in [9], learning of text entry techniques can be roughly divided into three stages: an initial discovery stage where users’ performance is significantly affected by prior experience, a second stage where users are developing their motor skills with the new techniques, and a third stage where expert users perform all actions reflexively.

In our experiment, it is quite clear that learning is still continuing, even at the end of the experiment (Figures 4-9), and thus the third stage is not reached. Most interestingly, we found that when participants did the experiment with *MultiTap* first, followed by *TiltText*, their experience with *MultiTap* very significantly affected their ability to use *TiltText* as evidenced by the higher errors rates. However, our analysis indicated that this prior experience did *not* affect the participants’ entry speed when they switched to *TiltText*. Indeed, even though they made more errors, they were able to compensate for them by higher speeds of action. It appears from our analysis that participants move to the second stage at around block 9, or about 180 phrases (~4000 keystrokes).

One of the reasons why we chose to compare *TiltText* to *MultiTap* was because *MultiTap* is used as a baseline technique in most studies of text entry techniques. Comparing our results to that of [9], we see approximately that the number of phrases entered per technique by block 16 in our experiment is roughly equivalent to the 6th or 7th session in MacKenzie et al.’s experiment. At this point, their data for *MultiTap* is roughly in the 11 wpm range, which is very close to ours. At the same point in time, entry rates for *LetterWise* are about 14 wpm, which is in the same range as our experiment’s rate for *TiltText* of 13.57 wpm. While the different experimental designs, number of phrases per block, and other factors necessarily imply that these cross-experiment comparisons are not precise, this rough analysis does give us a ballpark sense of how *TiltText* compares to techniques other than *MultiTap*.

The increased error rate for forward/backward movements is possibly explained by limitations of the absolute tilt system used in our experiment. Participants tended to set the origin, then have the phone slowly “creep” forward as they made far more forward than back tilting gestures. As a result, the phone was always tilted somewhat forward. This meant that an exaggerated back gesture was required to enter “s” or “z”, which users often failed to accomplish on the first attempt. The tendency to hold the phone in a forward position also explains why most *tilt errors* resulted in entering the forward tilt letter on the same button. Due to hardware constraints, we used an absolute tilt implementation of *TiltText* in our experiment, instead of the more efficient and accurate relative tilt method. We posit that error rates could be improved if we use relative tilt, since users would not have to tilt a specific amount past the origin between characters as required in the absolute tilt method, and they also would not have to exaggerate the back tilts to overcome the forward posture. These extra

requirements are a likely cause of errors, particularly if users attempt to perform the technique quickly without watching the screen. Relative tilt would be more amenable to fast, “eyes-free”, use.

Previous work [8, 9, 15] present very detailed analyses with regards to the various components of text entry time using various existing techniques. We refer readers to those papers for the details, however, it is worthwhile discussing how *TiltText* differs. First, all other techniques have only a single motor component: that of selecting the correct key and pressing it a certain number of times. In contrast, *TiltText* has two motor components: selecting and pressing the correct key as with other techniques, and tilting the phone in the appropriate direction. One could probably argue that there is also an additional cognitive component required in determining which direction to tilt, at least during the initial discovery and skill development stages of learning. This raises an interesting question with regards to the applicability of the standard metric for evaluating text entry techniques put forth by [8]: KeyStrokes Per Character (KSPC). *TiltText* by definition has a KSPC of 1, not including error correction. However, KSPC does not fully capture the second, tilting, aspect of *TiltText*. While the tilting movements required by *TiltText* are small, particularly if a relative tilt implementation is used, Fitts’ law tells us that there is nevertheless a cost associated with this movement. It is difficult to determine exactly the parameters of the tilting movement in order to calculate a precise estimate of its cost, however, we can say that the tilt angle required ranges from a little more than 0 degrees to an approximate maximum of 90 degrees. From our observations of participants in our experiment, it appears that the average tilt angle is probably around 30 degrees. With a more definitive determination of this parameter or at least a smaller bound on its range, it would be possible to develop a model that more accurately describes *TiltText* than KSPC. Since there is no fixed target that the tilt must accomplish, a simple application of Fitts’ law is not appropriate. However, the work of Accot and Zhai on the “steering law” [1] and “crossing based interfaces” [2] provide a foundation on which an analysis could be performed. We leave the development of such a model to future work, once a more detailed determination of the tilting parameter is performed.

CONCLUSIONS and FUTURE WORK

We have described *TiltText*: a new technique for text input to mobile phones that combines a low cost tilt sensor with the standard numeric keypad. An initial experiment indicates that *TiltText* significantly outperforms the commonly used *MultiTap* technique. Comparisons to previous work indicate that *TiltText* approaches the performance of linguistically-based techniques, without the limitation of requiring words to be from fixed language set or requiring the constant monitoring of the visual display. Given the huge installed base of the standard numeric keypad, and other potential uses of tilt sensors in mobile

devices [5-7], we believe that *TiltText* could potentially have a good chance of adoption.

We plan in future to develop a working relative-tilt implementation of *TiltText* in an effort to demonstrate even better speed results. We believe a relative tilt system will also reduce errors, but other work on reducing the error rate via possibly better sensing hardware and algorithms is clearly needed.

We will also examine whether different tilting directions are more natural for the wrist, and perhaps replace the simplistic “left/right, back/forward” technique examined here. We will also test a *TiltText* implementation that groups letters to minimize the need to re-tilt between character entries.

ACKNOWLEDGEMENTS

We thank Microsoft Research for supporting this work, Michael McGuffin for help with the statistics, Richard Watson for background help with software development, all who participated in our experiment, and members of the Dynamic Graphics Project lab (www.dgp.toronto.edu) at the University of Toronto for valuable discussions.

REFERENCES

1. Accot, J., & Zhai, S. (1997). Beyond Fitts' Law: Models for trajectory-based HCI tasks. *ACM CHI Conference on Human Factors in Computing Systems*. p. 295-302.
2. Accot, J., & Zhai, S. (2002). More than dotting the i's - foundations for crossing-based interfaces. *ACM CHI Conference on Human Factors in Computing Systems*. p. 73-80.
3. Bartlett, J. (2000). Rock 'n' Scroll is here to stay. *IEEE Computer Graphics and Applications, May/June 2000*. p. 40-45.
4. Fishkin, K., Gujar, A., Harrison, B., Moran, T., & Want, R. (2000). Embodied user interfaces for really direct manipulation. *Communications of the ACM*, 43(9). p. 75-80.
5. Harrison, B., Fishkin, K., Gujar, A., Mochon, C., & Want, R. (1998). Squeeze me, hold me, tilt me! An exploration of manipulative user interfaces. *ACM CHI Conference on Human Factors in Computing Systems*. p. 17-24.
6. Hinckley, K., & Horvitz, E. (2001). Toward more sensitive mobile phones. *ACM UIST Symposium on User Interface Software and Technology*. p. 191-192.
7. Hinckley, K., Pierce, J., Sinclair, M., & Horvitz, E. (2000). Sensing techniques for mobile interaction. *ACM UIST Symposium on User Interface Software and Technology*. p. 91-100.
8. MacKenzie, I.S. (2002). KSPC (keystrokes per character) as a characteristic of text entry techniques.

Fourth International Symposium on Human-Computer Interaction with Mobile Devices. p. 195-210.

9. MacKenzie, I.S., Kober, H., Smith, D., Jones, T., & Skepner, E. (2001). LetterWise: Prefix-based disambiguation for mobile text input. *ACM UIST Symposium on User Interface Software and Technology*. p. 111-120.
10. Partridge, K., Chatterjee, S., Sazawal, V., Borriello, G., & Want, R. (2002). TiltType: accelerometer-supported text entry for very small devices. *ACM UIST Symposium on User Interface Software and Technology*. p. 201-204.
11. Poulton, E.C., Bias in quantifying judgements. 1989, Hillsdale, NJ: Lawrence Erlbaum.
12. Rekimoto, J. (1996). Tilting operations for small screen interfaces. *ACM UIST Symposium on User Interface Software and Technology*. p. 167-168.
13. Sazawal, V., Want, R., & Borriello, G. (2002). The Unigesture Approach. One-Handed Text Entry for Small Devices. *Mobile HCI*. p. 256-270.
14. Silfverberg, M., I. Scott MacKenzie, & Korhonen, P. (2000). Predicting text entry speeds on mobile phones. *ACM CHI Conference on Human Factors in Computing Systems*. p. 9-16.
15. Soukoreff, W., & MacKenzie, I.S. (2002). Text entry for mobile computing: Models and methods, theory and practice. *Human-Computer Interaction*, 17. p. 147-198.