

Tilt to Scroll: Evaluating a Motion Based Vibrotactile Mobile Interface

Ian Oakley Sile O'Modhrain

Palpable Machines Research Group, Media Lab Europe, Bellevue, Dublin 8, Ireland

E-mail: ian@medialabeurope.org, sile@media.mit.edu

Abstract

The increasing complexity and pervasiveness of handheld devices is demanding the development of interaction techniques explicitly tailored for mobile scenarios. This paper examines one such technique: the use of device orientation to navigate through one dimensional lists or menus. We describe a novel input mapping, directly relating list position to device orientation, and two studies quantifying user performance with this system. We believe vibrotactile feedback will be significant in movement-based mobile interfaces, and also examine its influence. Our results indicate that both our input mapping and the inclusion of vibrotactile feedback positively affect user performance and experience. We conclude by discussing the design implications of this work.

1. Introduction

Handheld devices are now a part of everyday life, and increasingly feature more and more complex functionality. New devices offer services such as web browsing, complex document creation and simple image and video capture and editing. Despite these substantial developments, there has been little corresponding improvement in the sophistication of mobile interfaces. Arguably, many handheld interfaces are simply adaptations of desktop GUIs, despite the fact that the ergonomics of mobile devices (in terms of the richness of the input and output mechanisms they support), and the context in which they are typically used, differs greatly from that found in a desktop environment.

Driven by this discrepancy, there is a growing body of research exploring alternative interaction paradigms designed expressly for use in mobile scenarios [3, 11]. Within this context, movement seems a promising input modality and has attracted considerable attention. Rekimoto [10] highlights a number of reasons for this. Firstly, as a user naturally holds a mobile device, no additional step is required to use movement as input.

Secondly, single-handed interaction is automatically supported. Finally, movement forms a rich, 6 DOF input channel capable of supporting complex interactions.

Extending this reasoning, Oakley *et al.* [7] point out that there is a direct trade off between using movement as input and graphical feedback as output. Essentially, they highlight the fact that moving a handheld device to control an interface reduces a user's ability to visually observe the state of that interface. Furthermore, as the movements being considered become richer and more dynamic, graphical feedback becomes less and less effective. One potential solution to this problem is through the inclusion of non-visual feedback.

To this end, a number of researchers have examined mobile interactions supported by audio [8, 11] and vibrotactile cues [9]. A compelling case can be made that vibrotactile feedback in particular is an appropriate output modality for movement-based interfaces. Perhaps the most important factor is that it can be fully integrated into a handheld device, and does not require users to don additional equipment such as headphones. Due to the naturally established contact with the hand, it is also highly immune to environmental noise. Furthermore, it is discreet; feedback can be delivered to the user holding the device without this being observed by others who might be nearby. Finally, due to the inclusion of a vibration alert in the vast majority of mobile phones, many users are familiar and comfortable with basic vibrotactile display.

The most common application area for movement-based interfaces has been the use of device orientation to control position in some navigable virtual space. Examples of such tilt interfaces include scrolling on a two-dimensional canvas such as a map [10] and navigating through one-dimensional menus or lists [9]. This second example strongly relates to current mobile interfaces: mobile phones usage is heavily based on menu selection operations [12].




Fig. 1. Tilt input to control position in a list.

One similarity among existing discussions of tilt-driven interfaces, is that they map device orientation to rate of scrolling, creating a control that operates much like a vehicle's accelerator pedal – the further it is pushed, the faster the vehicle moves. However, authors also report that one of the significant problems with tilt interfaces is that users find it difficult to stop at their desired destinations, often overshooting, then having to reposition themselves. This behavior is reported as a source of increased task completion times, error rates and frustration [9].

We propose an alternative input mapping to address this problem. Instead of using device orientation to control the rate of list traversal, we are concerned with the use of orientation to directly control list position, with linking specific orientations to specific list items. This is illustrated in Figure 1. We believe this simpler mapping, featuring first rather than second order control, will confer a number of benefits. Primarily, as it directly translates orientations to list positions, we suggest that users will be able to take more advantage of their kinesthetic sense – their awareness of the positions and orientations of their limbs – to reach list items than they would when using a rate-control metaphor. We anticipate this will lead to faster, more accurate and more satisfying interactions. However, as the range of movements required for this kind of position-based interaction is likely to obscure the screen of a handheld device, we also believe that vibrotactile feedback will be important in designing an effective and useable system.

With these factors in mind, we conducted two empirical studies to explore the potential of this idea. The first attempted to characterize user performance using a tilt driven list navigation task and an input metaphor mapping orientations directly to list position. We were interested in examining how list size and the inclusion of vibrotactile feedback affected user performance and satisfaction. Armed and informed by the results of this experiment, we then conducted a follow up study directly comparing the use of



Fig. 2. iPaq and MESH.

orientation to control list position against the use of orientation to control rate of list traversal. These studies are described below, followed by a discussion of the design implications of their results.

2. Experimental Overview

2.1. Hardware, Software and Stimuli

The experiments were conducted using a Compaq iPaq 5550 handheld computer running PocketPC 2003. All software was written in C++, using GapiDraw [1] to render all graphical elements. We used an iPaq expansion pack known as MESH [6] to sense orientations and produce vibrotactile cues. It features 3-axial acceleration sensing and a transducer based vibrotactile display. An iPaq and MESH are illustrated in Figure 2. The accelerometers in this system natively have a 10-bit accuracy with a bandwidth of 30Hz over a 5G range (an accuracy of approximately 0.005G). For use in these studies, this is filtered in software to 12-bit accuracy at 8Hz (approximately 0.00125G). This in turn can accurately measure the orientation of gravity (and therefore of the device) at the level of a ninth of a degree. This sensor system is capable of detecting orientations considerably more accurately than the human body is capable of reliably producing [13]. It is, however, worth noting that the accelerometers measure more than simply the vector of gravity: lateral movements of device also contribute to the accelerations reported.

Two vibrotactile samples were used in the studies. The first was a simple 250 Hz sine wave, which was displayed at various different intensities and for various different purposes in the two studies, and is discussed in more detail in those descriptions. The second stimulus featured unvaryingly in both studies and was a pop or click-like vibrotactile sample of 225ms in duration and composed of a 250 Hz sine wave with a gradually descending intensity.

2.2 Interface, Procedure and Measures

Both experiments were simple target acquisition studies, and featured a similar GUI, procedure and set of measures. Prior to each trial, a message was displayed asking users to orient the PDA at a comfortable angle before continuing. A 500ms minimum time between trials was enforced.

Figure 3 shows the graphical interface to the studies. At the top of the display was the list from which subjects were required to select items, towards the center of the screen lay the list item they must currently select. The visual presentation of the list always consisted of five items, regardless of the actual size of the list. A dark bar highlighted the currently selected item, and a scroll bar to the right of the textual labels indicated overall list position. The presentation of the list was such that the highlighted item remained in the center of the display until the extremes of list (the first, second, penultimate and ultimate members) were reached. At this stage, to ensure 5 items were consistently on screen, the highlight moved to the top or bottom of the list.

Trials were begun and ended by pressing a button on the PDA. To try to control the potentially confounding influence of posture and freedom of movement on the results, all participants were required to stand, to hold the PDA in their left hands, and to use a button situated under their thumbs to start and stop trials. The decision to use the left hand comes from the ergonomics of the PDA itself: it is designed for use when a stylus is employed in the dominant, typically right, hand.

The experimental measures were identical for both studies. For each trial we captured the duration,



Fig. 3. Graphical interface to studies.

whether it was a success or failure, and the number of times the correct target was entered before it was selected. The last measure reflects an observation made in the previous literature [9] that users often overshoot their desired destinations when using a tilt driven interface, and allows us to explore the validity of this suggestion. At the end of each condition participants completed an electronically administered NASA TLX questionnaire [5], an established measure of subjective workload.

3. Experiment 1

3.1. Experimental Design, Participants and Hypothesis

This study investigated user performance navigating lists when using position-based input. The effects of two variables were studied: list length and the vibrotactile feedback used to support the navigation. Lists consisting of 6, 9 and 15 items were examined. The range of PDA orientations mapped to list position was 90 degrees, from a horizontal orientation with the PDA's screen facing upwards to a vertical orientation with the screen facing towards the user. Therefore, as the size of the list increased, the angular size of each target decreased from 15 degrees, to 10 degrees, to 6 degrees respectively. To reduce rapid oscillations when moving between adjacent items we dynamically enlarged the currently highlighted item by 6 degrees, creating in effect a detent. Psychophysical literature suggests that the JND for orientation of the hand is in the order of 2.5 degrees [13]; all items sizes fall considerably outside this practical limit.

We studied three feedback conditions: Control, Transition, and Position. The Control condition featured no vibrotactile feedback. The Transition condition employed the clicking or popping stimulus described in the experimental overview whenever a user moved to a new item. The Position condition involved the continuous display of the 250Hz waveform, with its intensity determined by the currently highlighted list position. If this lay at the start of the list, a very low intensity waveform was displayed; at the end of the list, a waveform of the maximum magnitude was produced. The Transition condition was intended to provide users with information relating to a change of state of the device. The Position condition was designed to do not only this, but to also simultaneously provide information about the new state. As a user moves between items in the Position condition, the corresponding change in the intensity of the displayed waveform indicates the list

Condition abbreviation	Details
6C	6 Items, Control
9C	9 Items, Control
15C	15 Items, Control
6T	6 Items, Transition
9T	9 Items, Transition
15T	15 Items, Transition
6P	6 Items, Position
9P	9 Items, Position
15P	15 Items, Position

Table 1. Abbreviations for conditions.

Sub No	Condition order
1	6C, 9C, 15C, 6T, 9T, 15T, 6P, 9P, 15P
2	6C, 9C, 15C, 6P, 9P, 15P, 6T, 9T, 15T
3	6T, 9T, 15T, 6C, 9C, 15C, 6P, 9P, 15P
4	6T, 9T, 15T, 6P, 9P, 15P, 6C, 9C, 15C
5	6P, 9P, 15P, 6C, 9C, 15C, 6T, 9T, 15T
6	6P, 9P, 15P, 6T, 9T, 15T, 6C, 9C, 15C
7	6C, 6T, 6P, 9C, 9T, 9P, 15C, 15T, 15P
8	6C, 6T, 6P, 15C, 15T, 15P, 9C, 9T, 9P
9	9C, 9T, 9P, 6C, 6T, 6P, 15C, 15T, 15P
10	9C, 9T, 9P, 15C, 15T, 15P, 6C, 6T, 6P
11	15C, 15T, 15P, 6C, 6T, 6P, 9C, 9T, 9P
12	15C, 15T, 15P, 9C, 9T, 9P, 6C, 6T, 6P

Table 2. Condition order for each subject.

position has altered, and the new intensity actually represents the new position. While it may be debatable whether or not a user can reliably and consistently distinguish among a large set of vibrotactile intensities [14], simply the direction of the magnitude change itself indicates the scrolling direction and may be useful.

This experiment featured twelve participants, eight male and four female, all recruited from within our research lab. Their mean age was 28. Three were left handed, nine right handed. Although none had more than passing experience with tilt driven interfaces, all were familiar and comfortable with computers and mobile devices.

The experiment adopted a fully repeated measures design. Each subject performed each list-length condition with each feedback condition, leading to a total of 9 conditions. Each condition consisted of 30 trials; there were 270 trials in the entire study. Within each condition trials were presented in a random order to each participant, but were always drawn from the

same set of target items. In the 6 and 15 item length conditions, each target had to be selected 5 and 2 times respectively, while in the condition featuring 9 items, each target had to be selected a minimum of 3 times, with the first, fifth and ninth items gaining an additional fourth selection.

To attempt to control for order effects, each of the subjects experienced the conditions in a unique order, as shown in Tables 1 and 2. Immediately prior to the experiment, subjects completed a practice session of exactly half the length of the experimental session. Condition order was always the same in the practice and experimental sessions.

This goal of this study was to gain a general understanding of how users perform when using position-based tilt input. Although fundamentally exploratory in nature, some specific hypotheses were made. Firstly, we expected user performance to degrade with longer lists. Secondly, we predicted that the vibrotactile feedback would improve user performance and experience. Finally, we were interested in whether specific list items, and therefore specific device orientations, were easier or faster to reach than others.

3.2. Results

Initial analysis was performed using a 3x3 repeated measures ANOVA, followed by *post hoc* t-tests using Bonferroni confidence interval adjustments. List size exerted an effect on task completion time ($F=78.624$, $p<0.001$), as did vibrotactile feedback applied ($F=8.15$, $p<0.002$). There were no interactions ($F=1.424$, $p=0.242$). Follow up t-tests showed significant differences between all list length conditions (all at $p<0.003$ or better), indicating that users perform more rapidly with shorter lists. The Transition condition also led to faster results than the Control and Position conditions (both $p<0.05$). These data are shown in Figure 4.

List size also significantly effected error rate ($F=19.91$, $p<0.001$) and the number of times a target was entered before it was selected ($F=20.426$, $P<0.001$). However, feedback applied had no effect on either of these measures (respectively, $F=0.1$, $p=0.905$ and $F=1.647$, $p=0.216$), nor were any interactions uncovered ($F=0.127$, $p=0.972$ and $F=0.755$, $p=0.56$). A comparison of means showed the three list size conditions differed significantly from one another (at $p<0.01$ or better), on both error rate and number of times over target. Higher error rates and increased numbers of movements over targets occurred with longer list sizes. These data are shown in Figs 5 and 6.

The results of the TLX subjective measures questionnaires are shown in Figures 7 and 8, adjusted so that higher values consistently indicate greater workload. Broadly speaking the main trends to come from this data are that increased list size had no significant effects on subjective workload and that the Position condition led to increased workload (when compared to both the Control and Transition conditions) in Overall Workload and in all individual categories except Time Pressure (all at $p < 0.05$ or better).

After this comparison among the conditions, we also performed a preliminary examination of the effect of the position of list items on performance. We feel a full analysis of this data would be moving beyond both what the data captured in this study can support, and the scope of this paper, but the error data reveals an interesting curve, shown in Figure 9. In each of the three list sizes, it can be seen observed that excellent performance can be achieved at the peripheries of the lists, while considerably poorer performance is attained towards the center of the list. ANOVA revealed these trends to be significant for all list lengths: 6 ($F=25.036, p < 0.001$), 9 ($F=7.079, p < 0.001$) and 15 ($F=3.388, p < 0.001$) items.

3.3. Discussion

List length exerted a strong effect on the objective measures; this experimental hypothesis was confirmed. Task completion times, errors rates and instability when over the desired target all increased dramatically with longer lists. Somewhat curiously, there was no corresponding increase in reported workload, although informal conversations with participants after the study indicated that they tended to find the longer lists more difficult to navigate. This data suggests that 15 item lists would be a practical limit for this input technique, and that considerably improved performance is attainable when using shorter lists.

Vibrotactile feedback in the form of the Transition condition yielded a 10 percent improvement in task completion time compared to the control condition, a concrete performance boost, supporting the experimental hypothesis regarding the usefulness of non-visual feedback in this task. However, there was no improvement in error rate, stability over target, or subjective workload. Informal user comments were more in line with the temporal data and suggested that the presence of the Transition vibrotactile feedback was strongly appreciated. One user remarked it was reassuringly similar to the physical feedback found in traditional mechanical controls.

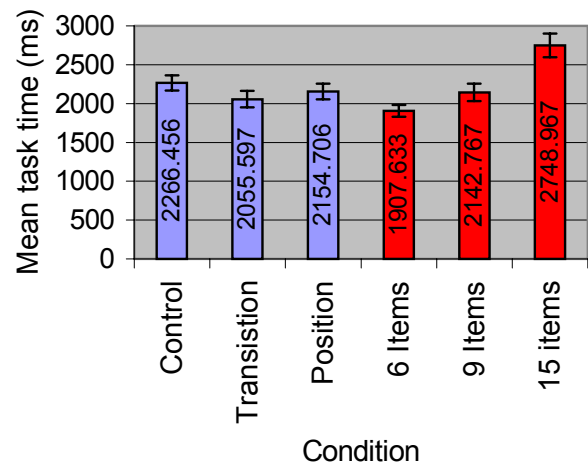


Fig. 4. Task completion times in expt 1.




Fig. 5. Mean number of errors in expt 1.




Fig. 6. Mean times over target in expt 1.

The Position condition led to no improvements in the objective measures, and slight but significant increases in subjective workload. Once again, user comments cast some light on this result; it was generally negative, with the exception of one user who reported they found the directional cues it provided informative and a useful augmentation to the graphical interface. Other comments included one user reporting an inability to distinguish between the different intensities, a suggestion that the sensation was evocative of the PDA being broken or malfunctioning, and one user even stating that the vibration made her anxious (she reported needing to go for a walk after the study in order to calm down). These comments are interesting, as they highlight the diversity of how this very simple feedback can be received and interpreted, both simply functionally (in terms of whether or not it is meaningful), and also, perhaps more importantly, emotionally. These subjective results highlight the evocative nature of haptic cues, and the importance of considering more than simply perceptual factors in their design.

The cursory examination of the influence of list item position on user performance suggests that it is an important factor to consider with regard to this input technique. Users can reliably reach the ends of a list, but find it more difficult to accurately target the center portion; a good interface design should exploit these properties.

Several informative comments can also be made from observations of the general usage of the system. Firstly, it was clear that most users employed a movement strategy that involved relatively large-scale motions, but that minimized viewing problems due to extreme screen angles. Basically, they adopted coordinated movements of their entire arm in conjunction with their neck to ensure the PDA's screen remained perpendicular to the angle of their gaze throughout the 90-degree range of rotation. Secondly, the majority of users reported a trade off between pressing the button to perform item selection and successfully hovering over the desired target. Essentially, they suggested the finger motion required to press the button could unintentionally disturb the orientation of the PDA, on occasion causing the highlight to move to an adjacent item just prior to the actual selection operation. This was a source of errors and frustration. It could be addressed in a number of ways, perhaps most simply through the addition of greater latency into the system by either the inclusion of an explicit delay or by lowering the cut-off frequency of the filter function that smoothes the input data.

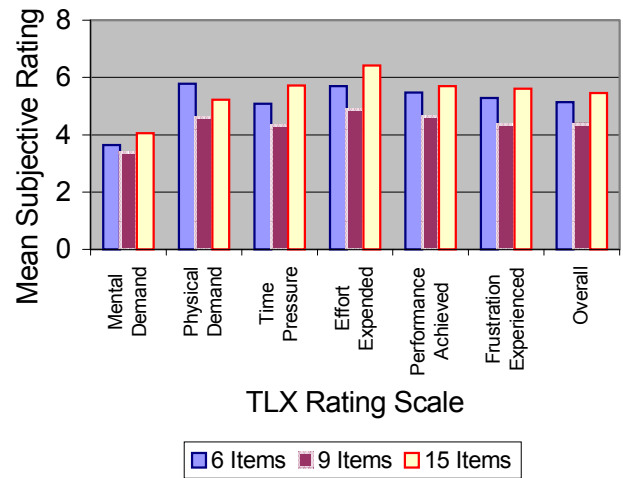


Fig. 7. TLX data by list size in expt 1.

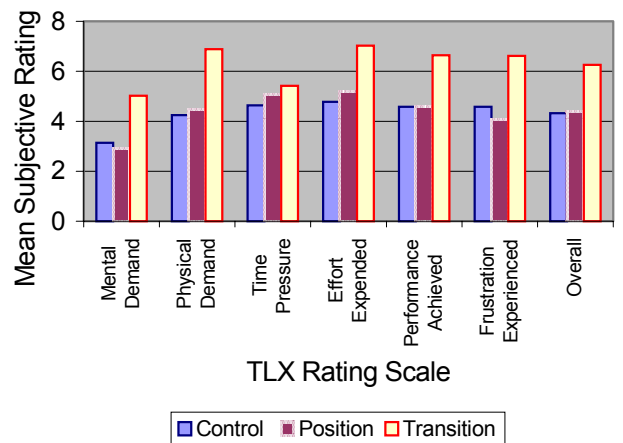


Fig. 8. TLX data by feedback applied in expt 1.

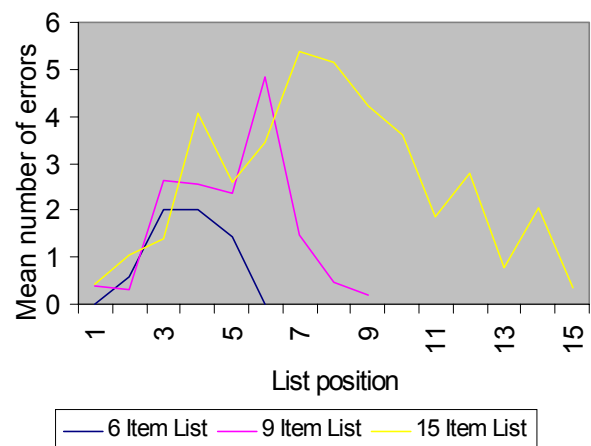


Fig. 9. Mean errors by item position in expt 1.

To situate the data from this experiment within a wider context it is important to compare it to list and menu selection data gathered by other researchers in other contexts. GUI literature varies dramatically in its reports of menu selection times on desktop computers from 2.24 seconds for an 8-item menu [2], to 0.947 seconds for a 9-item one [15]. It is beyond the scope of this work to unify this diverse data, but it is promising to note, than even without the benefit of extended practice, our results (of 2.142 seconds on a 9 item menu) fit within this range.

More recently, research has examined completion times on typical button based cell-phone menu selection tasks. Principally, St. Amant *et al.* [12] evaluate user-interface models (such as GOMS [4] and ACT-R) against data gathered from an empirical study. They conclude GOMS is a good predictor of task performance, and decompose menu selection tasks into two discreet actions: scrolling to an adjacent item and item selection. Respectively, these have a time cost of 0.505 seconds and 0.616 seconds. Fitting these figures to our experimental scenario, it suggests a specific item could be selected from a 9-item menu in a minimum of 0.616 seconds, a maximum of 5.161 seconds, and an approximate mean of 2.888 seconds. This suggests our interface may be considerably quicker than conventional menu selection techniques. One caveat to this conclusion is that our list is composed of ordered items, whereas Anant's model is derived from the actual interface to a mobile phone; the cognitive load required to process the terms used in this interface may contribute to the task times generated from the model.

Finally, it is worth comparing our results to empirical data from previous studies of list navigation using tilting interfaces. Pouprey *et al.* [9] report such a study. They use a rate control interface to measure user performance with and without the support of vibrotactile feedback (described in more detail in the section on experiment 2). In their study, task completion times ranged from 1.9 seconds to select an adjacent list item when supported by vibrotactile cues, through 3.1 seconds to select a target 6 items distant, to 3.7 seconds to select an item that was initially 12 items away. Task completion times without vibrotactile feedback were all significantly greater than these figures. Although this work is not directly comparable to the study reported in this paper, it does cast it in a favorable light. The position-based input metaphor results in a 1.9 second target selection time when distinguishing between 6 items, and 2.7 seconds when selecting from a range of 15 available targets. However, given the prevalence of the rate-based metaphor in tilt interfaces, we felt this topic warranted direct examination.

4. Experiment 2

Our initial experiment indicated that user performance using the position-based input was generally good, and allowed us to select optimum parameters for position-based list navigation. However, most work on tilt driven input has focused on rate-based navigation. In order to directly compare between these two techniques, we ran a short follow-up study.

3.1. Experimental Design, Participants and Hypothesis

This study had a straightforward design including just two conditions, one incorporating rate-based navigation, the other position-based navigation. All trials featured lists of 9 items in length. Feedback and input in the position-based condition was identical to that in the Transition condition in the first study. The rate-based condition was heavily based on that described and evaluated by Pouprey *et al.* [9]; this study directly links into that work.

In this condition, device orientation controls the rate at which navigation in the list takes place. The orientation at the start of each trial is taken as a neutral orientation; no list traversal occurs in a 10-degree band centered on this point. When outside of this band movement in the list occurs at two distinct speeds. If a user is within a 20-degree band from the center point for 500ms, the highlight moves to the list adjacent item. If they are outside of this band movement between items occurs every 250ms. In this condition two vibrotactile stimuli are employed. Firstly, the click used in the previous study marks the movement to a new item. Secondly, when users are in the neutral band, a low intensity 250Hz sine wave is continuously played to indicate their position is currently stable.

6 users participated in this study. 4 were male, 2 female, 4 were right-handed, 2 left-handed. The mean age was 26. One again, all subjects were employees at our lab and were comfortable with mobile devices, but had no more than trivial experience with tilt driven interfaces. None had participated in the previous study.

This study had a repeated measures experimental design. There were two order conditions; half of the participants were allocated to each. Each condition included 54 trials, arranged such that users had to select each list item 6 times. Trials were delivered in a random order. A practice session of exactly half the length of the experimental session took place immediately before the experiment.

This study contrasts the use of tilt in position-based and rate-based control metaphors. Based on our own observations and the results from the first study, the experimental hypotheses were that that position-based control would yield better performance, and decreased subjective workload when compared to the rate-based technique.

3.2. Results

All analyses were conducted using paired sample t-tests and assuming a one-way hypothesis. In the objective data, the position-based condition yielded significantly lower task completion times ($p < 0.05$) than the rate-control condition. The difference in error rates approached significance ($p = 0.054$), and no difference was observed in the number of times users entered the correct target prior to selection ($p = 0.259$). These data are shown in Figure 10, 11, and 12. The TLX data (shown in Figure 13) also supports the position-based condition. It produces not only significantly lower overall workload ($p < 0.05$) but also reductions in the individual categories of mental demand, effort expended and frustration experienced (all at $p < 0.05$).

3.3. Discussion

Despite the statistical limitations imposed by its small size, this study provides compelling results. The position-based input technique leads to significantly improved task completion times, and a substantial decrease in error rates when compared against the more established rate-based metaphor. Subjective measures also indicate users strongly prefer position-based input.

The data from this study also reinforces that generated in first study: tasks completion times and error rates (when adjusted for the differences in the number of trials forming each condition) are similar to those found in the 9-item Transition condition. The data from the position-based condition also fits in well with that reported by Poupreve *et al.* [9]. In each trial in this study, users moved between zero and four items prior to target selection (a mean of 2.222 items), and the task completion time of 2.665 seconds falls neatly between Poupreve's data of 1.9 seconds for adjacent items and 3.1 seconds for the selection of a target initially 6 items distant.

User comments were not as straightforward as the workload data might suggest. Although there were no negative comments about the position-based condition, several users were positive about the rate-based

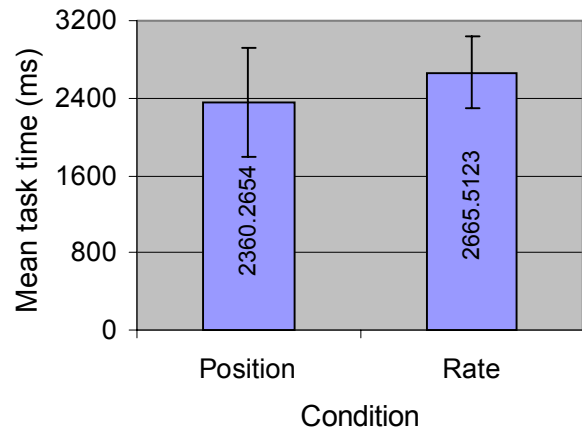


Fig. 10. Mean task time in expt 2.




Fig. 12. Mean times over target in expt 2.




Fig. 11. Mean errors in expt 2.

condition, citing the fact that they found it easier to see the screen as the main reason for this. Other users however, were strongly in favor of the position-based condition, relating their difficulties in targeting specific items as the main cause of their disfavor. These users reported a number of coping strategies for this problem including trying to develop a “flicking” motion that could be used to move to adjacent list items.

From observing users interacting with both techniques it seems that one of the chief benefits of the position-based control is its lack of reliance on any hidden virtual model. The rate-based condition operates through the definition of an abstract neutral orientation, and changes in state are linked to a user’s relationship to that object. In contrast, the position-based model simply relies on the directly observable orientation of the device. It exhibits no more complex behavior: if the device is rotated in one direction, movement in the list occurs proportionally to that. Unlike the rate-based metaphor, which includes concepts such as slowing down and speeding up, there is no conceptual distance between a user’s manipulation of the device and of the list. It seems that most users appreciated this simplicity.

4. Design Implications

The results of these studies support the development movement based mobile interfaces; user performance compares favorably to that achieved using existing button-based interaction techniques. However, it is clear that optimal performance will only be gained with systems that are carefully designed. Consequently, here we attempt to make some general observations that could serve to guide such design processes.

The position control mapping proposed in this paper yields good performance, but it seems likely it will only be useful in specific circumstances. The practical limit on list length suggests that it might be an appropriate technique for navigating relatively short menus, but not for tasks featuring longer lists, or lists with a variable number of items. A common example of such a complex list would be an address book.

This suggestion is reinforced by the fact that the error rate is dependent on list item position. A static menu can be designed such that it is easy to reach commonly used items – indeed it would seem an imperative concern when considering this technique – but an address book cannot. An address book is not only of variable length, but additions to the list can appear at any point, causing the position of all subsequent items to change. A good menu design for

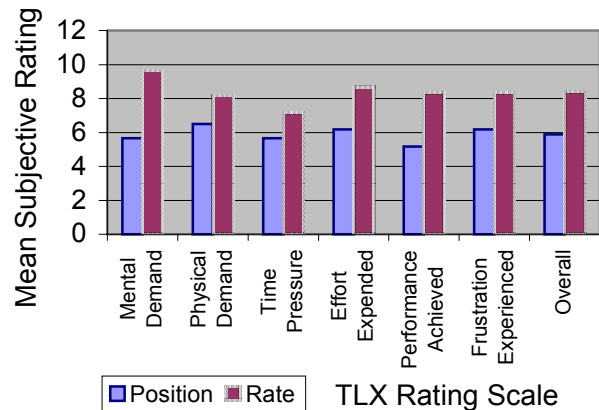


Fig. 13. TLX data from experiment 2.

the position-based input should include the most commonly selected items on the extremities of the list, where they can be reached most reliably. Based on the informal observation that most users hold a mobile device in a relatively static comfortable orientation (of approximately 30 degrees from horizontal), it may also be helpful to include a default menu item in that region.

A final important property of the position-based input mapping is its learnability; the fact that a specific orientation corresponds to a specific item. An expert user could learn the orientations required to execute specific commands, potentially beginning to use motion in open loop interactions. This has specific implications for mobile devices, as the proliferation of this technology is leading to a similar proliferation of expert users. They are a demanding user group and one that is becoming an increasingly important segment to consider in the design of handheld interfaces.

The rate-control input mapping suits different scenarios. Specially, although it affords reduced user performance when compared to the position control mapping, it is more general purpose. With it there is no inherent limit on list length; once scrolling is initiated, it can be continued indefinitely. The only directly learnable behavior it features is a default initially highlighted item. Consequently, it seems unlikely that error rate will vary with list position, or that the addition of extra list items will exert such a dramatic effect as it would with the position-based input.

Vibrotactile feedback is also an important component of motion-based interfaces. While including it can lead to objective improvements in performance, the most compelling results stem from user commentary. Appropriately designed feedback (such as the click used in the studies described here to mark the transition between list items) yields almost universally positive remarks, while other, inappropriately designed feedback (such as the

continual vibration used on the Position condition in the first study) yields consistently negative feedback. This observation highlights the importance of quality, and not only efficiency, in interaction design. Correspondingly, we suggest that designers of systems to include vibrotactile display pay close attention to user opinions of the stimuli they employ. Based on our experiences, it appears that cues that resemble those that might be encountered when manipulating a mechanical input device are likely to be well received, and can significantly add to the perceived quality of a virtual interaction.

5. Conclusions and Future Work

We believe the position-based input mapping proposed in this paper has considerable potential. One key reason for this is that users can learn to reach specific list items, which in turn, may lead to an open loop interaction, where an item can be selected with confidence without explicitly requiring feedback from the system. It would be valuable to empirically examine the validity of this idea.

It would also be interesting to look at the implications of using this input technique with nested menu systems, where multiple selections are made to reach a single goal. Specifically, to examine whether users can learn to string together multiple selection operations into more fluid gestural motions, much as they do with marking menu systems [16].

The role of vibrotactile feedback in motion-based interfaces could also be developed further. Significant questions include whether well-designed vibrotactile cues can support a user learning an interface, or provide meaningful contextual information.

In conclusion, we believe that interfaces based around motion input and vibrotactile output have an important role to play in the next generation of mobile interaction techniques. We have demonstrated that appropriately designed systems can yield high levels of performance, and provide users with a satisfying experience. In short, we believe that haptic interaction has a role to play beyond the desktop, and that this work takes steps towards that aim.

References

- [1] "Develant.com - Cross Platform Tools for Mobile Game Development", www.gapidraw.com, 2004.
- [2] J. Callahan, D. Hopkins, M. Weiser, and B. Shneiderman, "An Empirical Comparison of Pie vs Linear Menus," presented at ACM CHI'88, 1988.
- [3] K. Ehrlich and A. Henderson, "Design: (Inter)facing the millennium: where are we (going)?," *interactions*, vol. 7, pp. 19-30, 2000.
- [4] W. D. Gray, B. E. John, and M. E. Atwood, "Project Ernestine: A validation of GOMS for prediction and explanation of real-world task performance," *Human-Computer Interaction*, vol. 8, pp. 209-237, 1993.
- [5] S. Hart and L. Staveland, "Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research," in *Human mental workload*, P. Hancock and N. Meshkati, Eds. Amsterdam, 1988, pp. 139-183.
- [6] S. Hughes, I. Oakley, and S. O'Modhrain, "MESH: Supporting Mobile Multi-modal Interfaces," presented at ACM UIST'04, Santa Fe, NM, 2004.
- [7] I. Oakley, J. Ängeslevä, S. Hughes, and S. O'Modhrain, "Tilt and Feel: Scrolling with Vibrotactile Display," presented at EuroHaptics 2004, Munich, Germany, 2004.
- [8] A. Pirhonen, S. A. Brewster, and C. Holguin, "Gestural and Audio Metaphors as a Means of Control for Mobile Devices," presented at ACM CHI'02, Minneapolis, MN, 2002.
- [9] I. Poupyrev, S. Maruyama, and J. Rekimoto, "Ambient touch: designing tactile interfaces for handheld devices," presented at ACM UIST'02, Paris, France, 2002.
- [10] J. Rekimoto, "Tilting Operations for Small Screen Interfaces," *UIST*, 1996.
- [11] M. Rinott, "Sonified Interactions with Mobile Devices," presented at International Workshop on Interactive Sonification, Bielefeld, 2004.
- [12] R. St. Amant, T. E. Horton, and F. E. Ritter, "Model-based Evaluation of Cell Phone Menu Interaction," presented at ACM CHI'04, Vienna, Austria, 2004.
- [13] H. Z. Tan, M. A. Srinivasan, B. Eberman, and B. Cheng, "Human factors for the design of force-reflecting haptic interfaces," presented at ASME Dynamic Systems and Control Division, Chicago, IL, 1994.
- [14] J. van Erp, "Guidelines for the Use of Vibrotactile Displays in Human Computer Interactions," presented at EuroHaptics'02, Edinburgh, UK, 2002.
- [15] N. Walker and J. B. Smelter, "A Comparison of Selection Times from Walking and Pull-Down Menus," presented at ACM CHI'90, 1990.
- [16] S. Zhao and R. Balakrishnan, "Simple vs. Compound Mark Hierarchical Marking Menus," presented at ACM UIST'04, Santa Fe NM, 2004.