

IntelliTilt: An Enhanced Tilt Interaction Technique for Mobile Map-Based Applications

Bradley van Tonder and Janet Wesson

Department of Computing Sciences, Nelson Mandela Metropolitan University
PO Box 77000, Port Elizabeth, South Africa, 6031
{Bradley.vanTonder, Janet.Wesson}@nmmu.ac.za

Abstract. Current interaction techniques for mobile map-based applications suffer from several usability problems. Tilt interaction provides an alternative form of interaction which combines the benefits of one handed interaction with intuitive physical gestures. Research has shown that tilt interaction suffers from a lack of controllability, high mental demand and practical concerns. In this paper, the design and evaluation of a new tilt interaction technique, called IntelliTilt, is described. IntelliTilt incorporates several intelligent techniques to address the shortcomings of tilt interaction. IntelliTilt was compared to a basic tilt interaction technique using a prototype mobile map-based application in an experiment. The results of this experiment showed that IntelliTilt was preferred by the participants and that it offered significant advantages in terms of mental demand, perceived efficiency and controllability.

Keywords: Tilt interaction, mobile map applications, sensor-based interaction.

1 Introduction

Map-based applications are rapidly becoming standard features of mobile phones. Touch-screen and keypad interaction are currently the preferred interaction techniques for interacting with mobile map-based applications. These techniques, however, suffer from several shortcomings. Touch-screen interaction can result in the display being obscured by the user's fingers during interaction [1]. Touch-screen interaction also requires both hands. Previous research has shown that users prefer to only use one hand when interacting with maps on mobile phones [2]. Keypad interaction limits the user's freedom of movement and provides little control over panning speed [3, 4].

Tilt interaction offers an intuitive alternative, with the user being able to interact with mobile map-based applications using natural gestures [3]. Tilt interaction typically allows users to interact with the display using rotations about the x and y axes. Tilt interaction does, however, suffer from several shortcomings. Previous studies have also shown no significant benefit in terms of effectiveness or efficiency [5].

Basic tilt interaction using Speed Dependant Automatic Zooming (SDAZ) was previously shown to have several shortcomings in comparison to keypad interaction for mobile map-based applications [6]. This paper proposes an enhanced tilt interaction

technique, called IntelliTilt, which incorporates several design features to address the shortcomings of basic tilt interaction. These features include making use of visual and vibrotactile feedback to improve the ease with which users can control panning and selection operations. Attractor mechanisms are included to make it easier to settle the cursor on a particular icon. Sensitivity adaptation is included to compensate for variability in accelerometer data while walking. A gesture zooming technique was developed to provide an intuitive method of zooming. These features were designed to address the shortcomings of basic tilt interaction identified earlier.

A user study was conducted to compare IntelliTilt to basic tilt interaction incorporating SDAZ. The purpose of the user study was to determine whether the design features of IntelliTilt helped to address the previously identified shortcomings of basic tilt interaction. Participants were required to perform typical mobile map-based tasks while seated and while walking. Touch-screen interaction was not considered in this research as our focus was on investigating one-handed interaction techniques for mobile devices.

This paper begins with a discussion of related work. The design of the IntelliTilt interaction technique is then described. A prototype mobile map-based application, called MapExplorer, which was used for evaluation purposes, is described. Experimental results are then presented, followed by a discussion of the implications of these results. Finally, conclusions and ideas for future work are presented.

2 Related Work

2.1 Mobile Map-Based Applications

Mobile map-based applications typically support users in performing the following tasks [7, 8]:

- *Locating*: Identifying the position of something (e.g. where am I?);
- *Searching*: Identifying facilities matching certain criteria (e.g. where is the nearest hotel?);
- *Navigating*: Finding a route between two points or navigating along a route;
- *Checking*: Determining the condition of a person or place (e.g. operating hours of a business); and
- *Identifying*: Identifying and recognizing people, places or objects (e.g. the name of a business).

These five high-level tasks are accomplished through several low-level operations. Irrespective of the interaction technique used, three low-level operations are commonly used, namely panning, zooming and selection [9].

2.2 Tilt Interaction

Tilt interaction was first proposed almost 15 years ago as an experimental form of interaction for mobile devices [10]. Research into the use of tilt interaction has since been extended to a range of domains including menu navigation, text entry, mobile museum guides, photo browsing and interacting with mobile map-based applications [3, 4, 11]. Most existing research has focused on the use of tilt as a means of performing panning, where the tilt angles along the x and y axes are mapped onto

panning speeds in the horizontal and vertical directions. Most of this research relied on accelerometer data to determine tilt angles. Recently, gyroscopes and digital compasses have also started to be integrated into mobile phones, allowing for more accurate measurement of pitch and roll angles [12].

Tilt interaction techniques allow for one-handed interaction with mobile devices. This is desirable in a mobile context of use and allows the user an unobstructed view of the display. Furthermore, one-handed interaction requires less visual attention than bimanual control [13]. Existing studies have shown that users like tilt interaction because of its natural, expressive and intuitive nature [3]. Unlike keypad interaction, tilt interaction allows for fine-grained control over panning speed.

Different implementations of tilt interaction differ in terms of the mapping between tilt input and the corresponding effect. Some implementations identify discrete inputs, where tilting past a certain threshold is interpreted similarly to a button press [14]. Other implementations allow for continuous input, where tilt input is used to continuously perform panning [15]. The mapping between tilt input and panning effect can be done using position control, rate control or inertial control [14].

Rate control maps the rate at which the display pans onto the angle at which the device is tilted relative to the neutral position. For example, the further the device is tilted to the left, the faster the display pans left [16]. The function mapping tilt input onto panning speed does not need to be linear. It has been suggested that a linear mapping for small tilt angles allows for fine-grained control, while the amplification of larger tilt-angles makes it easier to pan faster over long distances [15]. This approach has previously been used in a mobile map-based application [17].

Discretization can be used to improve the level of control users have over tilt interaction. Different tilt angle intervals are mapped onto different tilt speeds to minimize jitter and allow for finer-grained control. Linear, quadratic and sigmoidal discretization are all possible. Existing research has provided evidence that quadratic discretization allows for the highest level of user control [18].

Some mobile map-based applications allow panning and zooming to be performed simultaneously. Most of these applications implement Speed Dependent Automatic Zooming (SDAZ). SDAZ was proposed to address some of the problems and shortcomings of traditional zooming and panning techniques [19]. These problems include blurring (and disorientation) when panning at high speeds and the need for the user to switch focus between the display and user interface controls [19, 20]. SDAZ adjusts the zoom-level automatically in response to the speed at which the user pans the display. SDAZ has since been implemented for touch-screen mobile devices [21].

SDAZ has been implemented using tilt interaction to perform combined zooming and panning in a mobile document browser [22]. Kratz and Rohs[1] combined an element of manual control over the zoom level with a tilt-controlled SDAZ implementation. Their technique, called semi-automatic zooming (SAZ), allows users to control the zoom level using a touch-screen slider control. SAZ was shown to offer significant efficiency, workload and user satisfaction advantages over SDAZ.

2.3 Problems and Shortcomings

An experiment was previously conducted by the authors to compare tilt interaction (incorporating SDAZ) and keypad interaction [6]. Participants were required to

perform typical mobile map-based tasks (locating, navigating and checking) using a prototype mobile map-based application, called MapExplorer. This experiment took the form of a lab-based user study and involved 32 participants.

The results of this experiment allowed several shortcomings of basic tilt interaction to be identified:

- *Controllability*: Tilt interaction was shown to be difficult to control precisely, particularly for performing selections. This problem is not restricted to our implementation and is a well-documented shortcoming of tilt interaction [2]. Tilt interaction used for panning also tends to suffer from overshooting problems, where users pan past their intended target and have to reverse direction [23].
- *Zooming*: The use of SDAZ in MapExplorer yielded mixed results. While SDAZ did allow users to pan longer distances more rapidly, it was criticized because it was too easy to trigger accidentally and took control away from the user.
- *Mental Demand*: Tilt interaction was reported to be more mentally demanding than keypad interaction, particularly for locating and checking tasks.
- *Sensitivity*: Several participants felt that the sensitivity of the tilt interaction did not match their expectations (either too sensitive or not sensitive enough).
- *Practicality*: Some participants commented that they felt that tilt interaction was unlikely to be feasible in other contexts of use, such as while walking. This is due to the fact that tilt interaction requires users to exert control over their physical movements while interacting with the mobile device.

Several design modifications were made to the basic tilt interaction technique used in the above experiment in order to address the shortcomings identified. The design of this modified tilt interaction technique, called IntelliTilt, is described below.

3 IntelliTilt: Design and Implementation

IntelliTilt was designed to address the shortcomings of basic tilt interaction discussed above. Table 1 summarizes the relationships between the problems identified in basic tilt interaction and the features of IntelliTilt designed to address these shortcomings.

Table 1. Shortcomings of tilt interaction and IntelliTilt features designed to address these shortcomings

<i>Shortcoming</i>	<i>Feature(s)</i>
Controllability	Visual and vibrotactile feedback, attractor mechanisms
Zooming	Gesture zooming
Mental Demand	Visual and vibrotactile feedback
Sensitivity	Sensitivity adaptation
Practicality	Sensitivity adaptation, dwell-time selection

The design of each of the above features of IntelliTilt is now described in more detail.

3.1 Visual and Vibrotactile Feedback

The results of the previous experiment showed that participants struggled to control tilt interaction. Qualitative comments revealed that some participants found the exact effects of tilt gestures (in terms of panning speed and direction) difficult to predict. In order to accurately control tilt interaction, users need to be aware of the effect the current orientation of the mobile phone has on the direction and speed of panning.

Various modalities have previously been employed to aid in controlling tilt interaction, including visual, vibrotactile and audio feedback [1, 24, 25]. Audio feedback was considered to be impractical because it is likely to be annoying to other people (and the use of headphones is not always feasible). Previous research has only made limited use of visual feedback to indicate cursor position and to control automatic zooming [1].

IntelliTilt makes use of visual feedback in the form of arrows attached to the cursor, showing the vertical and horizontal panning speeds (Figure 1). The length of these arrows is used to denote speed (the longer the arrow, the faster the panning speed). A simple linear mapping between panning speed and arrow length is used.

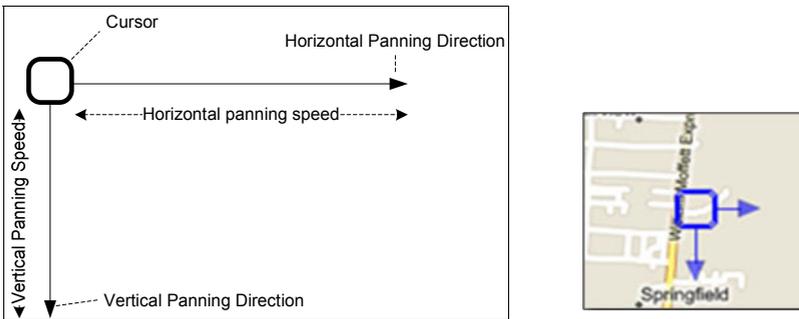


Fig. 1. Design of visual feedback in IntelliTilt (left) and an example of its use (right)

IntelliTilt also makes use of vibrotactile feedback to reinforce the visual feedback channel. This form of feedback is particularly useful in situations where the user may be partially distracted (e.g. while walking). Previous research regarding the use of vibrotactile feedback in conjunction with tilt interaction for menu navigation showed that users liked the use of short vibration pulses to indicate movement from one menu item to another [24]. IntelliTilt uses short vibration pulses (250 ms) to indicate when the cursor moves over point of interest (POI) icons and route markers. This duration was selected as it has previously been employed for similar purposes and experimentation revealed it to be long enough to be noticeable, but short enough to not irritate users [24]. The available map area is divided into grid cells based on latitude and longitude, with POI icons indexed into these cells using a hash table. This enables IntelliTilt to efficiently determine whether the cursor is close enough to a POI or route marker icon to trigger vibrotactile feedback.

3.2 Attractor Mechanisms

One of the most significant problems encountered during the use of tilt interaction in mobile map-based applications is that users struggle to settle the cursor on a particular target icon. In a previous system using tilt interaction to browse photo collections on a mobile phone, attractors were employed to make it easier to settle on a particular photo [3].

This idea was extended to the two-dimensional domain of map browsing. The attractors are designed to work in conjunction with a discretization approach, which splits tilt input into discrete speed levels. If both the current horizontal and vertical panning speeds are slow (in which case it is likely that the user is trying to select an icon), the algorithm then determines whether any selectable icons are within a specified distance. Nearby icons are identified using the same indexing method used for vibrotactile feedback. If a nearby icon exists, the position of the cursor is adjusted vertically and horizontally to draw the cursor towards the centre of the nearest icon. Figure 2 illustrates the use of attractors to draw the cursor (shown in blue) towards the nearest POI within range of the cursor.

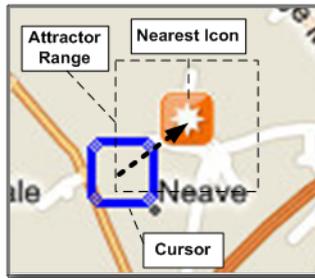


Fig. 2. Example of the use of attractors to aid selection

The resulting effect is one of drawing the cursor slowly towards selectable icons. The effects of the attractors are easily overcome by increasing panning speed, so as to avoid being drawn onto icons when this is not desired. The use of attractors eliminates the need to exactly position the cursor, as the user is able to approximately position the cursor and is assisted by the attractor to perform exact selection.

3.3 Sensitivity Adaptation

One of the problems identified in the previous experiment was that users felt that tilt interaction would be impractical to use in a mobile context due to the sensitivity of this form of interaction [6]. If tilt interaction is not sensitive enough, it is inefficient, while if it is too sensitive it is difficult to use while walking. Sensitivity adaptation was implemented in IntelliTilt to address this problem.

Streams of accelerometer data along the x, y and z axes are constantly recorded and monitored. Accelerometer data is sampled at approximately 20 Hz and a sliding window of data from the last second (last 20 samples) is used to identify whether the user is currently stationary or mobile. Such a sampling rate has previously been

successfully employed in a similar mobile application [1]. Several data samples were recorded of users using MapExplorer while stationary and while walking in order to measure the baseline acceleration variances for the three axes. While performing tilt operations, spikes in acceleration along one or more of the axes are recorded, so a single spike in acceleration is not sufficient to identify when the user is walking. Variances were calculated for acceleration along the three axes, as these provide a more accurate measurement of variability in the acceleration data than the raw values themselves. Analysis of the recorded data showed that while walking, all three acceleration variance values were typically greater than 0.125 (where acceleration is measured in g-forces). This is as a result of the linear acceleration detected by the accelerometer while the user is walking. While seated, acceleration is largely as a result of deliberate tilting gestures, which typically involve acceleration along only one or two axes. As a result, all three values rarely exceeded this threshold at the same time while the user was seated. This observation allows a computationally inexpensive method of identifying when the user is walking that can be performed quickly enough not to negatively affect system performance. Figure 3 illustrates the acceleration variances using actual recorded data while seated and walking. In order to improve the accuracy of the algorithm, all observations within a second need to indicate a change in order for the algorithm to transition between walking and stationary states. Preliminary testing revealed this method to be extremely accurate at identifying when users were walking (> 98% accuracy in informal trials).

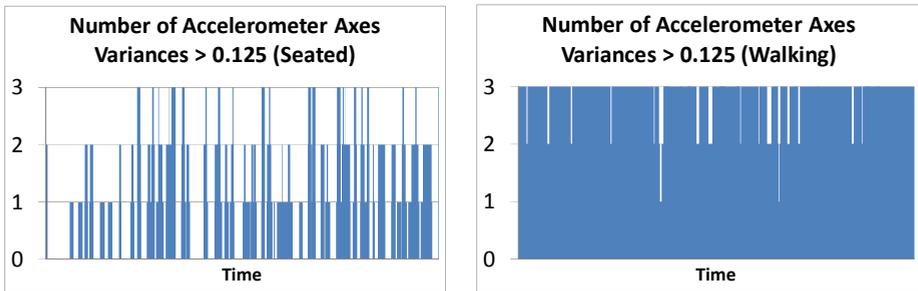


Fig. 3. Comparison of the number of axes recording acceleration variances of greater than 0.125 while seated (left) and while walking (right). Acceleration measured in g-forces.

In order to compensate for increased acceleration variability when walking, IntelliTilt decreases the sensitivity of tilt interaction (Figure 7). Tilt interaction techniques typically employ a “dead zone”, where small tilt angles are ignored to allow the user to maintain a stable cursor position. IntelliTilt also increases the size of the dead zone when the user is walking. IntelliTilt constantly observes the variation in the acceleration data in order to identify when to implement sensitivity adaptation.

3.4 Gesture Zooming

Previous implementations of tilt interaction have largely relied on SDAZ, where the zoom level is linked to the panning speed [1, 6, 26]. This form of zooming, however,

has been shown to be frustrating for users who feel that it takes control away from them. In an attempt to find a form of zooming that would offer better usability and user satisfaction, a gesture-based zooming technique was implemented in IntelliTilt.



Fig. 4. Zooming out using gesture zooming in IntelliTilt

Gesture zooming in IntelliTilt relies on physical movement of the mobile phone. Zooming out was implemented using a backward movement of the device (Figure 4), while zooming in was implemented using a forward movement of the device, relative to the z-axis (which is perpendicular to the phone’s screen). This mapping was designed to be conceptually the same as moving a camera away or towards the map. Such an approach has previously been used in a camera-controlled mobile map-based application [27]. These gestures are identified by constantly examining the z-axis acceleration data received from the accelerometer to identify spikes in the data.

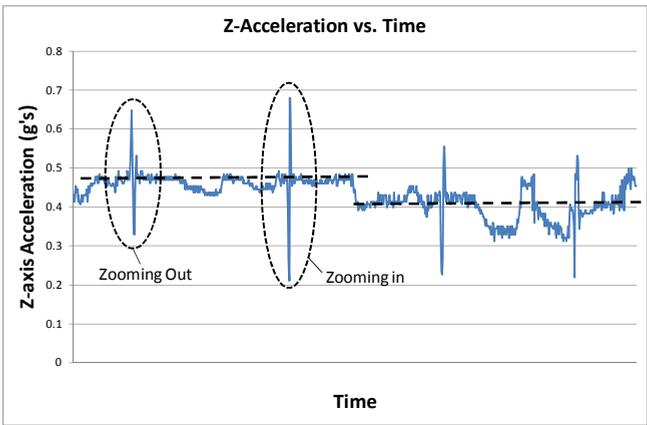


Fig. 5. Z-axis acceleration data showing zoom gestures. The dotted line shows the baseline acceleration which changes depending on the orientation of the phone.

Acceleration data was captured and stored of users performing the zooming gestures using a Nokia N97. Figure 5 shows an example of the z-axis acceleration data captured. The spikes in the data show where zooming gestures were performed. Zooming in and zooming out gestures can be differentiated from each other based on

the order in which the lower and upper spikes occur (this can be used to infer the direction the phone is moving). Experimental trials identified changes of 0.15g either side of the baseline z-acceleration (when the phone is at rest) as suitable to identify zooming in and zooming out gestures. While walking, this interval is increased to allow for the greater variability in z-axis acceleration data. The acceleration values reported during zooming gestures must take the starting orientation into account. As shown in Figure 5 (dotted line), the baseline z-acceleration value will differ depending on the orientation of the phone. As a further measure to avoid accidental zooming, zooming gestures have to be completed within 300ms (normal movement is generally more gradual). This threshold was determined after experimentation to identify the most suitable interval.

3.5 Dwell-Time Selection

During the previous experiment, it was observed that users would often accidentally pan away from a target icon while touching the screen to perform a selection operation [6]. Furthermore, the original design required both hands to perform selection operations. In order to combat this, dwell-time selection was implemented in IntelliTilt. If the cursor is over an icon for longer than 2.75 seconds, the icon is selected without the user having to do anything. Vibration pulses of increasing durations (25ms, 50ms and 75ms) are used to indicate to the user that a selection operation is about to be performed. The duration of 2.75 seconds was selected after experimentation and was set relatively long as the timing starts when the user is within 25 pixels of any icon, rather than directly over an icon. Shorter durations resulted in unintentional selections taking place. Dwell-time selection can be deactivated when not desired and manual touch-screen selection is always possible.

4 Prototype: MapExplorer

A prototype mobile map-based application supporting both basic tilt and IntelliTilt was developed for experimental purposes. The prototype, called MapExplorer, was implemented in Java ME and Python and tested using a Nokia N97. Microsoft's Bing Maps service was used as the source of the maps in the application, with map tiles being saved to the phone's memory to improve performance. MapExplorer allows users to browse maps at ten different zoom levels (from 1 pixel = 4891.97m to 1 pixel = 9.56m). A caching system is used, whereby the most recently loaded map tiles are stored in memory, allowing these tiles to be efficiently loaded the next time they are needed. When the maximum cache size is reached (40 tiles), the least recently used tiles are removed from memory to make space for the new tiles. The basic functions of MapExplorer are first described, along with the implementation of tilt interaction. The implementation of SDAZ, used in the basic tilt technique, is also described.

4.1 Functionality

MapExplorer supports four of the five typical mobile map-based tasks (locating, navigating, identifying and checking). Searching was not implemented in order to ensure that users interacted with the map-based display.

Three categories of POIs are included in the system, namely restaurants, hotels and tourist attractions, each denoted by a different icon (Figure 6). Selecting a POI icon displays detailed information such as a description, contact details and opening hours. MapExplorer also allows users to plan and follow routes. A route planning web service was used to calculate routes and retrieve routing instructions. By moving the cursor over route information markers, users are able to read routing instructions (Figure 6).

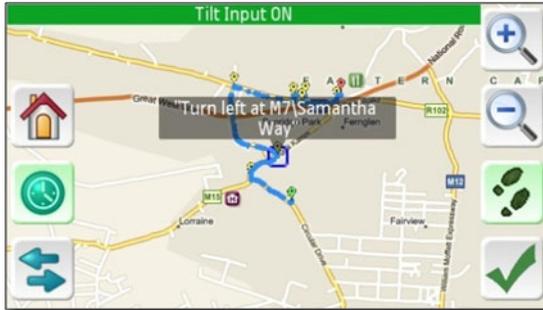


Fig. 6. Navigating using MapExplorer

Tilt interaction in MapExplorer is implemented using a rate control approach. The further the device is tilted in a certain direction, the faster the display is panned in that direction. The user is able to finely control the speed and direction of panning by adjusting the tilt angle in any direction.

The accelerometer in the Nokia N97 is able to detect acceleration between $-2g$ and $2g$ relative to the x, y and z axes. The Python sensor module was used to access the acceleration data. Accelerometer values along the x, y and z axes are used to calculate pitch and roll angles [28].

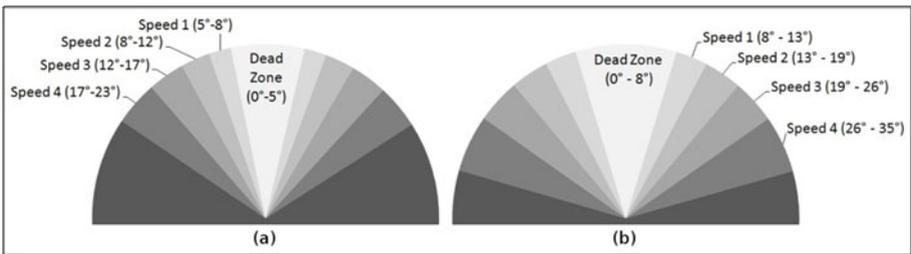


Fig. 7. Discretisation where the user is (a) stationary and (b) mobile

Discretization is used to map pitch and roll angles onto panning speeds (Figure 7). This process consists of splitting the range of motion into intervals, with tilt angles falling in the same interval resulting in the same effect. Discretization intervals are increased for larger angles, as previous research has shown that this approach allows for the highest level of user control. Panning speed increases linearly for the first two

discretisation intervals and quadratically for the last three. Previous research has shown that this allows for fine control with small tilt angles and faster panning with larger angles. A dead zone is included to avoid unintentional panning as a result of minor orientation adjustments and hand tremor.

Tilt control is activated and deactivated by tapping anywhere on the display. This allows the user to lock the display and prevents accidental panning from taking place. The neutral orientation, relative to which tilt angles are calculated, is the orientation of the phone when tilt interaction is activated.

4.2 SDAZ Implementation

The zoom level is linked to panning speed in order to prevent the problem of extreme visual flow, where the display can become blurred when panning at high speeds [19]. In order to combat this problem, the display is automatically zoomed out in order to prevent blurring when panning speed exceeds a certain threshold. The zoom level is a function of the panning speed. According to the original implementation of SDAZ, the zoom level is calculated as the normalized displacement between the current panning speed and the minimum panning speed which triggers zooming [19]. That is:

$$scale = s_o \frac{(d_y - d_0)}{(d_1 - d_0)} \quad (1)$$

Where s_o is the minimum scale, d_y is the current panning speed, d_0 is the minimum panning speed which triggers zooming out and d_1 is the maximum panning speed.

When automatic zooming takes place, a red rectangle is displayed to denote the area that will fill the screen when the panning speed is decreased again (Figure 8). A zoom level indicator is also temporarily displayed while the user is zooming. Automatic zooming was limited to only zooming out a single discrete zoom level to minimize user disorientation. Manual zooming can still be performed using the on-screen controls.



Fig. 8. Panning at a high speed using basic tilt interaction and SDAZ

5 User Study

In order to determine whether IntelliTilt helped to address the shortcomings identified in the previous experiment, a user study was conducted to compare the tilt interaction

technique developed in the original version of MapExplorer (referred to hereafter as basic tilt) with the modified interaction technique, called IntelliTilt.

5.1 Method

Participants. Sixteen participants (eleven male, five female) took part in the experiment. Nine participants had prior experience with tilt interaction, but in most cases this was only limited use of mobile games or the Nintendo Wii. All participants were computer science students between the ages of 20 and 29. Four participants were left handed and most indicated occasional use of mobile map-based applications.

Experimental Design and Tasks. A within-subjects approach was used with all participants using both interaction techniques. Two similar task sets were used to offset any order effects and the order in which the task sets were used was counterbalanced across participants. The order in which participants used the two interaction techniques was counterbalanced to offset any learning effect. The two techniques were simply named “version 1” and “version 2” (depending on the order in which they were used) in all questionnaires and documentation to avoid any bias. Participants were required to complete three types of tasks:

- *Locating tasks:* participants were required to find a particular POI icon (e.g. Locate the Ritz Hotel). Locating tasks therefore incorporate identify tasks.
- *Navigating tasks:* participants were required to plan and follow a route from a particular start point to a particular end point (e.g. Plan a route from the beach to the airport. Follow the route from start to finish).
- *Checking tasks:* participants were required to check specific information for a particular POI (e.g. check the opening hours of Starlight Restaurant).

Participants had to complete two tasks of each task type while seated followed by two tasks of each type while walking. Walking tasks were included in order to evaluate the two interaction techniques in situations in which there is likely to be more variability in the accelerometer data as a result of movement. Walking tasks were conducted indoors in a laboratory environment. All tasks started at the scale 1 pixel=38.22m. Targets were located in a major metropolitan area.

Metrics. The following three sets of metrics were collected:

- *Perceived workload:* Participants were required to complete a post-task questionnaire after using each interaction technique. This questionnaire included the six perceived workload measures from the NASA-TLX questionnaire (mental, physical and temporal demand, performance, effort and frustration) [29].
- *User satisfaction:* Six questions were included in the post-task questionnaire to measure perceived efficiency, effectiveness, ease of use, controllability, ease of performing selections and ease of use while walking. This section was based on the standard After-Scenario Questionnaire (ASQ) [30].
- *Performance:* Tasks were built into the system, allowing for easy measurement of task times, which were recorded in a log file stored on the phone.

The post-test questionnaire required participants to rate their preferred interaction technique while seated and while walking and to identify their preferred zooming technique. Participants were also asked to rate the usefulness of the visual and vibrotactile feedback and to identify positive and negative aspects of each interaction technique. Seven point semantic differential scales were used throughout.

5.2 Perceived Workload and User Satisfaction Results

Figure 9 shows the mean perceived workload ratings for the two interaction techniques. Perceived workload was lower (better) in all categories for IntelliTilt. These differences were significant for mental demand (Wilcoxon $Z = 2.78$, $p = 0.005$), temporal demand (Wilcoxon $Z = 2.20$, $p = 0.03$) and effort (Wilcoxon $Z = 2.55$, $p = 0.01$).

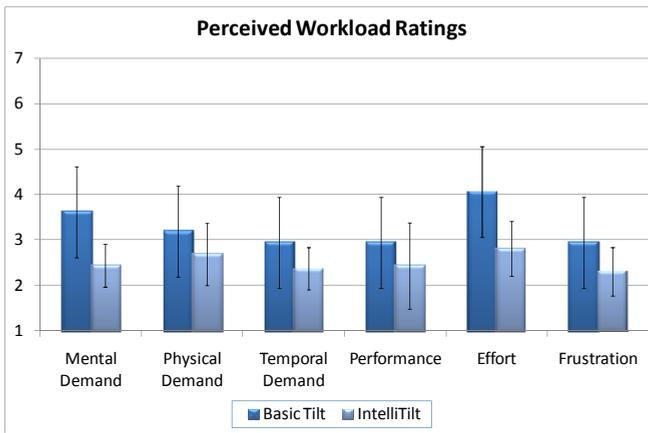


Fig. 9. Mean perceived workload ratings (95% confidence intervals shown) ($n = 16$). Lower performance values indicate better performance

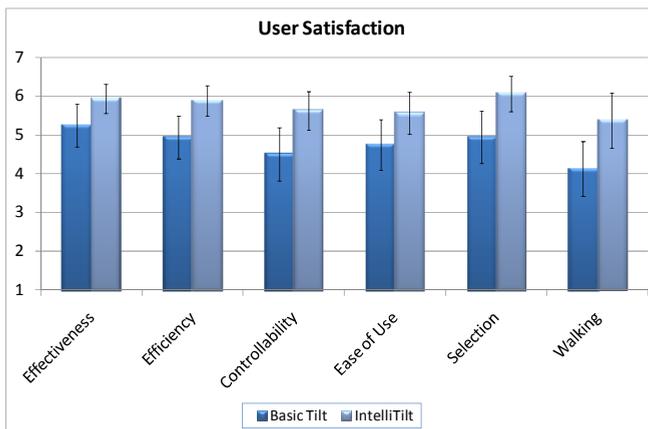


Fig. 10. Mean user satisfaction ratings (95% confidence intervals shown) ($n = 16$)

Figure 10 shows the mean user satisfaction ratings for the two interaction techniques. Participants were asked to rate the two techniques on a seven point semantic differential scale in terms of effectiveness, efficiency, controllability, ease of use, ease of selection and ease of use while walking. The results show that the IntelliTilt received higher user satisfaction ratings for all six satisfaction questions. These differences were statistically significant for effectiveness (Wilcoxon $Z = 2.19$, $p = 0.03$), efficiency (Wilcoxon $Z = 2.71$, $p = 0.01$), controllability (Wilcoxon $Z = 2.17$, $p = 0.03$), ease of selection (Wilcoxon $Z = 2.82$, $p = 0.01$) and ease of use while walking (Wilcoxon $Z = 2.13$, $p = 0.03$).

Table 2 shows the participant preference ratings from the post-test questionnaire. In Table 2, a value of 1 indicates a preference for basic tilt and 7 indicates a preference for IntelliTilt. Participants preferred IntelliTilt both while seated and while walking. The preference for IntelliTilt was stronger when participants were walking (mean of 5.88 while walking as opposed to 4.94 while seated). The difference between the mean values and the neutral value was statistically significant in both cases. Gesture zooming was preferred over SDAZ, with a mean value of 4.44. The difference between this mean value and the neutral value was not, however, statistically significant.

Table 2. Participant preferences (1=Basic tilt, 7 = IntelliTilt). Standard deviation values are shown in parenthesis (n = 16)

	<i>Mean</i>	<i>T-stat</i>	<i>P-value</i>
Preferred Technique (seated)	4.94 (1.69)	2.22	0.03
Preferred Technique (walking)	5.88 (1.45)	5.16	< 0.01
Preferred Zooming Technique	4.44 (2.00)	0.88	0.39

Participants were also asked to indicate whether they felt the vibration feedback and visual feedback were useful. Mean values of 6.25 and 6.00 were recorded for vibration and visual feedback, indicating that participants found these forms of feedback to be very useful. The difference between the mean and the neutral value of four was statistically significant for vibration feedback ($t_{15} = 7.32$, $p < 0.01$) and for visual feedback ($t_{15} = 6.32$, $p < 0.01$).

5.3 Performance Results

Table 3 shows the mean task times while seated and walking for the two interaction techniques. IntelliTilt proved slightly more efficient for seated tasks, while basic tilt proved slightly more efficient for walking tasks. The mean total time for both walking and seated tasks was almost identical. None of the differences between the two techniques were statistically significant using Wilcoxon ranked paired tests and alpha values of 0.05.

Mean task times for different task types are shown in Table 4. IntelliTilt achieved marginally better task times for locating and checking tasks, while basic tilt was slightly more efficient for navigating tasks. None of these differences were statistically significant using Wilcoxon ranked paired tests and alpha values of 0.05.

Table 3. Mean task times (in seconds) for seated and walking tasks (standard deviation shown in brackets) (n = 16)

	<i>Seated</i>	<i>Walking</i>	<i>Total</i>
Basic tilt	275.73 (60.49)	267.74 (66.38)	543.47 (113.13)
IntelliTilt	259.81 (82.54)	284.90 (91.66)	544.71 (156.98)

Table 4. Mean task times (in seconds) for different task types (standard deviation shown in brackets) (n = 16)

	<i>Locating</i>	<i>Navigating</i>	<i>Checking</i>
Basic tilt	124.38 (20.15)	314.16 (27.20)	104.93 (19.85)
IntelliTilt	116.47 (16.92)	328.56 (44.90)	99.68 (19.26)

5.4 Qualitative Feedback

Positive feedback about IntelliTilt focused on improved controllability. Several participants found the attractor mechanisms to be useful when performing selections. Participants also liked the visual and vibrotactile feedback. Positive feedback regarding basic tilt was received from some participants who felt that the use of SDAZ allowed them to be more efficient when locating POIs, particularly while seated. Some participants also liked the relative simplicity of this technique.

Negative feedback regarding IntelliTilt mainly centered on issues with gesture zooming, with some participants struggling to get used to this form of interaction. Negative feedback about basic tilt focused on the controllability of this technique, particularly while walking. Participants complained that SDAZ was not required when following a route and took place too easily while walking.

Participants who preferred gesture zooming found this form of zooming to be more intuitive than SDAZ. Several participants stated that they found SDAZ to be difficult and disorienting to use while walking. Participants who preferred the SDAZ zooming technique stated that they found this form of zooming to be more efficient and some found gesture zooming difficult to learn to use.

6 Discussion

IntelliTilt provided statistically significant improvements in three aspects of perceived workload, namely mental demand, temporal demand and effort. The improvement in mental demand is particularly encouraging, as this is one of the areas in which tilt interaction was previously identified as being inferior to keypad interaction [6]. Given the strongly positive comments regarding the use of attractors and visual and vibrotactile feedback, it is reasonable to assume that these features contributed to the improved mental demand.

The user satisfaction results were all higher for IntelliTilt. Of particular interest is the fact that users felt that it was easier to perform selection operations accurately

using IntelliTilt. Given that the additional feedback and attractor mechanisms were designed to improve selection, this result is encouraging. It is also encouraging to note that participants felt that IntelliTilt was significantly easier to use while walking, since sensitivity adaptation was designed for this purpose. Participants were noticeably more relaxed when using IntelliTilt while walking and were generally able to maintain their normal walking pace. Using basic tilt, participants frequently had to slow down to maintain control. Participants made frequent use of dwell-time selection, even when both hands were available.

There was very little difference in terms of efficiency for the two interaction techniques and overall task times were almost identical. No significant trends were evident in the task times for the different types of tasks. Despite the fact that there was little difference between the task times of the two interaction techniques, the participants perceived IntelliTilt to be more efficient. Given that controllability has previously been identified as a problem with tilt interaction, it is encouraging to note that participants found IntelliTilt easier to control.

User feedback regarding the use of visual and vibrotactile feedback was very positive. Participants rated both forms of feedback as very useful. Qualitative comments regarding these forms of feedback was also very positive, with only one participant remarking that vibrotactile feedback might become annoying after prolonged use.

IntelliTilt was preferred both while seated and while walking. This preference was stronger while walking. Several possible reasons for this exist. Firstly, IntelliTilt included sensitivity adaptation which makes allowance for the greater variability in accelerometer data while walking. Secondly, it is likely that the attractor mechanisms and feedback played a greater role in improving controllability while the user was mobile. Finally, SDAZ was perceived to be difficult to use while walking.

Opinions regarding the preferred zooming technique were divided. While some users quickly mastered the gesture zooming technique, others found the learnability of this technique to be problematic. Gestures were sometimes not detected due to participants moving the phone vertically, rather than backwards or forwards along the z-axis (which is dependent on the phone's orientation). SDAZ also resulted in divided opinions. SDAZ was identified as being useful for locating POIs, but difficult to use when following a route. Participants also found SDAZ to be difficult to use when walking.

7 Design Implications

Analysis of the above results provided insight into the design of tilt interaction for mobile map-based applications. The following list contains some implications for the design of tilt interaction for mobile devices:

1. Mechanisms should be provided to compensate for the inexact nature of tilt interaction and facilitate selection by providing automated assistance to the user.
2. Multimodal feedback can be used to improve controllability and reduce mental demand. Vibrotactile feedback allows for non-visual reinforcement that is useful in a mobile context.
3. Compensation should be made for variability in accelerometer data while walking by decreasing the sensitivity of tilt interaction.

4. One-handed selection can be facilitated using tilt interaction. Dwell-time selection can provide one possible implementation.

8 Conclusions and Future Work

This paper proposed an enhanced tilt interaction technique called IntelliTilt, which was designed to address the current shortcomings of basic tilt interaction. IntelliTilt was compared to basic tilt interaction using a prototype mobile map-based application, which allowed participants to browse maps in a wide geographic area at a wide range of zoom levels. The inclusion of walking tasks allowed IntelliTilt to be evaluated in a realistic context of use which is essential for mobile interaction research.

IntelliTilt was designed to address specific shortcomings of basic tilt interaction that were previously identified. Statistically significant improvements were achieved in several important areas, including perceived mental demand, controllability and ease of selection. These results provide empirical evidence that the use of visual and vibrotactile feedback and attractor mechanisms can aid users in controlling tilt interaction and performing selections. Participants also indicated that they preferred IntelliTilt, particularly when walking, suggesting that sensitivity adaptation was beneficial in this regard. Results regarding gesture zooming were less conclusive, however, with participants divided between SDAZ and the new gesture zooming technique.

Future work will involve investigating improvements to gesture zooming. The use of gyroscope and magnetometer sensors in conjunction with accelerometers will be investigated in order to improve gesture recognition and responsiveness.

References

1. Kratz, S., Brodien, I., Rohs, M.: Semi-Automatic Zooming for Mobile Map Navigation. In: Proc. MobileHCI 2010, pp. 63–72. ACM Press, New York (2010)
2. Karlson, A., Bederson, B., Contreras-Vidal, J.: Understanding Single-Handed Mobile Device Interaction. Technical Report HCIL-2006-02 (2006)
3. Cho, S.-J., Sung, Y., Murray-Smith, R., Lee, K., Choi, C., King, Y.-B.: Dynamics of Tilt-Based Browsing on Mobile Devices. In: Proc. CHI 2007, pp. 1947–1952. ACM Press, New York (2007)
4. Kratz, S., Rohs, M.: Navigating Dynamically-Generated High Quality Maps on Tilt-Sensing Mobile Devices. In: Proc. of Workshop, MEIS 2008 (2008)
5. Dong, L., Watters, C., Duffy, J.: Comparing Two One-handed Access Methods on a PDA. In: Proc. Mobile HCI 2005, pp. 235–238. ACM Press, New York (2005)
6. van Tonder, B., Wesson, J.: Is Tilt Interaction Better Than Keypad Interaction for Mobile Map-Based Applications? In: Proc. SAICSIT 2010, pp. 322–331. ACM Press, New York (2010)
7. Reichenbacher, T.: Mobile Cartography - Adaptive Visualisation of Geographic Information on Mobile Devices. Institut für Photogrammetrie und Kartographie., PhD Thesis, Technischen Universität (2004)

8. Von Hunolstein, S., Zipf, A.: Towards Task-Oriented Map-based Mobile Guides. In: Proc. International Workshop "HCI in Mobile Guides", Mobile HCI 2003 (2003)
9. Looije, R., te Brake, G., Neerinx, M.: Usability Engineering for Mobile Maps. In: Proc. Mobility 2007, pp. 532–539 (2007)
10. Rekimoto, J.: Tilting Operations for Small Screen Interfaces (Tech Note). In: Proc. UIST 1996, pp. 167–168. ACM Press, New York (1996)
11. Mantyjarvi, J., Paternò, F., Salvador, Z., Santoro, C.: Scan and Tilt –Towards Natural Interaction for Mobile Museum Guides. In: Proc. MobileHCI 2006, pp. 191–194. ACM Press, New York (2006)
12. Apple. iPhone 4, <http://www.apple.com/iphone>
13. Eslambolchilar, P., Murray-Smith, R.: Interact, Excite, Feel. In: Proc. the Second International Conference on Tangible and Embedded Interaction (TEI 2008), pp. 131–138. ACM Press, New York (2008)
14. Oakley, I., Ängeslevä, J., Hughes, S., O'Modhrain, S.: Tilt and Feel: Scrolling with Vibrotactile Display. In: Proc. EuroHaptics 2004, pp. 316–323 (2004)
15. Dachsel, R., Buchholz, R.: Throw and Tilt – Seamless Interaction across Devices Using Mobile Phone Gestures. In: Proc. MEIS 2008, pp. 272–278 (2008)
16. Hinckley, K., Pierce, J., Horvitz, E., Sinclair, M.: Foreground and Background Interaction with Sensor-Enhanced Mobile Devices. *ACM Transactions on Computer-Human Interaction* 12(1) (2005)
17. Mock, M., Rohs, M.: A GPS Tracking Application with a Tilt- and Motion-Sensing Interface. In: Proc. Workshop Mobile and Interactive Systems, MEIS 2008 (2008)
18. Rahman, M., Gustafson, S., Irani, P., Subramanian, S.: Tilt Techniques: Investigating the Dexterity of Wrist-based Input. In: Proc. CHI 2009, pp. 1943–1952. ACM Press, New York (2009)
19. Igarashi, T., Hinckley, K.: Speed-dependent automatic zooming for browsing large documents. In: Proc. 13th Annual ACM Symposium on User Interface Software and Technology, pp. 139–148. ACM Press, New York (2000)
20. Cockburn, A., Savage, J.: Comparing speed-dependent automatic zooming with traditional scroll, pan, and zoom methods. In: Proc. People and Computers XVII: British Computer Society Conference on Human Computer Interaction, pp. 87–102 (2003)
21. Jones, S., Jones, M., Marsden, G., Patel, D., Cockburn, A.: An evaluation of integrated zooming and scrolling on small screens. *International Journal of Human-Computer Studies* (63), 271–303 (2005)
22. Eslambolchilar, P., Murray-Smith, R.: Tilt-Based Automatic Zooming and Scaling in Mobile Devices - a state-space implementation. In: Proc. Mobile HCI 2004, pp. 120–131 (2004)
23. Cho, S.-J., Murray-Smith, R., Kim, Y.-B.: Multi-Context Photo Browsing on Mobile Devices Based on Tilt Dynamics. In: Proc. Mobile HCI 2007, pp. 190–197. ACM, New York (2007)
24. Oakley, I., O'Modhrain, S.: Tilt to Scroll: Evaluating a Motion Based Vibrotactile Mobile Interface. In: Proc. of EuroHaptics 2005, pp. 40–49 (2005)
25. Eslambolchilar, P., Williamson, J., Murray-Smith, R.: Multimodal Feedback for Tilt Controlled Speed Dependent Automatic Zooming. In: Proc. UIST 2004. ACM Press, New York (2004)
26. Eslambolchilar, P., Murray-Smith, R.: Control centric approach in designing scrolling and zooming user interfaces. *International Journal of Human-Computer Studies* 66(12), 838–856 (2008)

27. Mooser, J., You, S., Neumann, U.: Large Document, Small Screen: A Camera Driven Scroll and Zoom Control for Mobile Devices. In: Proc. I3D, pp. 27–34. ACM Press, New York (2008)
28. Tuck, K.: Tilt Sensing Using Linear Accelerometers, Freescale Semiconductor (2007), http://www.freescale.com/files/sensors/doc/app_note/AN3461.pdf
29. Hart, S.G., Staveland, L.E.: Development of NASA-TLX (task load index): results of empirical and theoretical research. *Human Mental Workload*, 139–183 (1988)
30. Lewis, J.R.: Psychometric evaluation of an after-scenario questionnaire for computer usability studies: The ASQ. *SIGCHI Bulletin* 23(1), 78–81 (1991)