Variability in Wrist-Tilt Accelerometer Based Gesture Interfaces

Andrew Crossan,[1] Roderick Murray-Smith[1,2]

[1] Hamilton Institute, National University of Ireland, Maynooth, Co. Kildare Ireland.
`andrew.crossan@may.ie`
`http://www.hamilton.ie/andy`
[2] Department of Computing Science, University of Glasgow, Glasgow, Scotland.
`rod@dcs.gla.ac.uk`
`http://www.dcs.gla.ac.uk/~rod`

Abstract. In this paper we describe a study that examines human performance in a tilt control targeting task on a PDA. A three-degree of freedom accelerometer attached to the base of the PDA allows users to navigate to the targets by tilting their wrist in different directions. Post hoc analysis of performance data has been used to classify the ease of targeting and variability of movement in the different directions. The results show that there is an increase in variability of motions upwards from the centre, compared to downwards motions. Also the variability in the x axis component of the motion was greater than that in the y axis. This information can be used to guide designers as to the ease of various relative motions, and can be used to reshape the dynamics of the interaction to make each direction equally easy to achieve.

1 Introduction

Mobile devices are now widely used for a variety of everyday tasks. However, due to the requirement for a small screen, interacting with these devices often proves to be difficult. On-screen buttons are generally closely grouped together making interactions slow and error prone. This is particularly the case in a mobile context where the user's visual attention may be required elsewhere.

Generally, interaction with these devices has taken the form of discrete messages passed between the user and device. The user will click a button or select a menu item, and the device will supply feedback. This method can be slow and frustrating particularly in situations requiring many button clicks such as typing with an on-screen keyboard.

The development of new interaction techniques and sensors provide more opportunity for a more continuous form of interaction, allowing closed loop interaction between device and the user's motions. In this instance, all of the user's movements affect the interpretation of the interaction and the device can continually change the feedback supplied to the user accordingly. Gesture input is one form of continuous interaction that has been underused in interaction with current systems. Text entry is the one major exception to this where gesturing with a stylus is often used for input-

ting text to a PDA. In this case, it is used to provide a quick, more natural alternative to a screen-based keyboard where the keys may be required to be small and are tightly packed together leading to high error rates. Pirhonen, Brewster and Holguin [6] demonstrate an example of gesturing as an input technique for controlling a PDA based MP3 player. These interactions are designed to be intuitive for the task performed. Pirhonen, Brewster & Holguin were able to demonstrate significant usability benefits with the gesture interface over the standard interface, with users indicating that the gesture system required a lower workload to perform the task.

Recent studies have examined the possibility of using accelerometers attached to a mobile device to provide input. Advantages over most stylus based gesture systems are that they offer the possibility of one handed, screen free gesture control. They are often suggested as useful for continually monitoring background acceleration and providing context information for the current task. The components required for inertial input are also cheap to manufacture. (ca. $2 a device for mass production).

Accelerometers allow a user to input data and commands by tilting the device. Hinkley *et al.* [2] present a study that demonstrates a tilt-based gesture system for scrolling and automatic screen orientation of a PDA. Through user testing, they were able to provide a system that performed screen orientation and scrolling in a manner that was useful and predictable to the user. This study demonstrates the potential for tilt-based gestures to provide a fast, natural method for interaction.

Rekimoto [7] explores the possibility of using tilt input to navigate menus and scroll large documents and maps. The prototype system described allowed users to select items in pie menus although no formal evaluation was carried out.

Williamson and Murray-Smith [9] have developed the Hex system for inputting text on a PDA with accelerometer. This system allows the user to select letters by tilting the PDA to navigate a cursor through a series of tiled hexagons. Through use of a language model, they were able to adjust the feedback given to the user such that probable sequences of characters were easier to perform than non-probable sequences. TiltType presented by Partridge *et al.* [5] is similarly a tilt based text entry method where characters are selected by a combination of button clicks and the orientation of the device. The inertial control allows TiltType to be used on devices with extremely small screens such as a watch.

2 Targeting Tasks

There is a large body of literature studying targeting tasks using many different input devices. Most common are Fitts' Law based studies where users are required to continuously move between two targets (an overview can be found in [3]). Timing and error rates can be gathered for different target widths and separations allowing the experimenter to calculate the comparative difficulty of the task. Most studies work with univariate targets by setting narrow target widths while allowing effectively infinite target heights. Accot and Zhai [1] describe a study that extends Fitts Law to take account of two-dimensional targets. Their experiment was used to select a model that provides the Fitts' Law index of difficulty for two-dimensional targeting.

MacKenzie *et al.* [4] describe methods that are based on the variability in movement rather than error rates. They suggest task metrics suitable for measuring movement variability including slip off errors, mean distance from the task axis, movement variability perpendicular to the task axis, and orthogonal direction changes.

This paper is concerned with gesturing using wrist tilt motions. With all gesturing systems, there will be a degree of variability in the gesture, and therefore uncertainty about the gesture performed. This study examines the variability in movement for short gestures in eight directions. The gestures require users to move a cursor between a series of pairs of points by tilting their wrists. The study hoped to determine areas of difficulties at the limits of comfortable movement in different tilt directions. Both error rate and variability metrics are considered. Speed and accuracy of targeting in different directions is also examined.

3 Experimental Method

3.1 Equipment

The experiment was carried out with an HP 5450 PDA with the Xsens P3C 3 degree of freedom linear acceleration sensor attached to the serial port (shown in Figure 1). Its effect on the balance of the device is negligible (its weight is 10.35g). The accelerometer was used to detect tilt magnitude around the x and y axis of the mobile device, sampling at a rate of 35 samples per second.


Figure 1. PDA with XSens accelerometer attached at the base. The user would move the cursor by tilting the device in the directions shown.

3.2 Task

The experimental environment used is shown in Figure 2. Nine circular targets of radius 15 pixels were placed throughout the environment. One target was placed at the centre of the screen, and eight were spaced at 45-degree angles around the circumference of a circle centred on the initial target such that the radius of this circle was 100 pixels. The gain on the cursor movement was set such that this distance corresponded to a tilt of approximately 48 degrees in the x direction and approximately 36 degrees in the y direction. The difference in these values correspond to a scaling due to screen size such that the same tilt magnitude is required to move to each of the edges of the screen (for a screen of width 240 pixels and height 320 pixels). Due to the different x-y cursor gains, the results section considers comparisons made between targets in opposite directions only.

These values provided a wide range of tilts while still allowing the user to easily view his or her interaction on the screen. A pilot study suggested that screen contrast became an issue with larger tilts in the y direction. The cursor gain was deliberately set to a low value such that large tilts would be required to complete the task and the limits of the movement would therefore be explored.

The task given to participants was to select the highlighted target (in Figure 2 the top centre target is shown to be highlighted). The cursor was controlled by a linear position control mechanism, mapping rotation of the device to movement of the cursor. The device accelerometer was calibrated such that the starting position of the device corresponded to the centre position on the screen. This calibration occurred at the start of each trial. To move the cursor in the x direction, the device was titled left or right, and to move the cursor in the y direction, the device was tilted up or down (shown in Figure 1). Distance of the cursor from the centre position was directly mapped to angle of rotation from the rest position. Therefore, double the rotation angle of the device would lead to the cursor being twice as far from the central position. Since a position control mechanism was employed, if the user held the devices still at any orientation, then the cursor would remain still on the screen.

Users held the device in their dominant hand and were instructed to sit in a comfortable position with the device held such that they could easily see the screen. In practice, all participants sat with the device slightly tilted towards them and leaning forwards slightly over the device.

Selection required the user to hover the cursor over the target for 1.5 seconds. If the cursor slipped off the target before the selection was complete, the target timer was reset and the user was again required to move onto and hover over the target for the full one and a half seconds. Once successful selection of a target was complete, a different target was then highlighted. The sequence of targets was chosen such that highlighted targets alternated between any of the outside target and the centre target. This ensured that all movement was either from the central target to an outer target, or from an outer target to the central target. This sequence was chosen to ensure that the path distance to the next target was always kept constant and that the angle to the next target was restricted to the eight equally spaced angles chosen.

Two competing factors affected the chosen target size. As the trajectory rather than the targeting was the main measurement for the task, the targets needed to be big

enough to allow easy targeting. However, to maintain similar path length between starting position and target position, the targets could not be made to be too large. A diameter of 15 pixels was eventually chosen empirically. A bar at the top of the screen (shown in Figure 2) indicates the time the user has spent over the target. When the bar reaches the right of the screen, target selection has been completed.
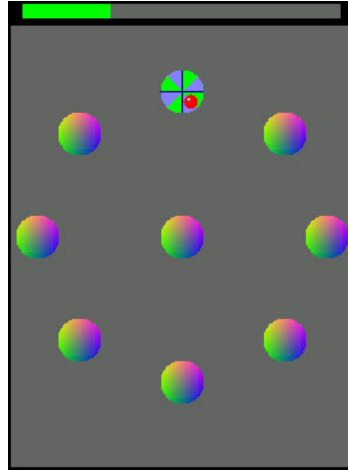

Figure 2. The experimental environment used for the study. The top centre target is the highlighted target. The cursor is the smaller circle within this target.

All participants took part in three experimental sessions with an hour break between each for recovery. The first session was used to train users in the task. The second and third sessions were eventually used when analysing the movement characteristics of different participants. The sessions were designed to be short to minimise user fatigue. No session lasted over five minutes.

3.3 Participants

Twelve participants took part in the training then the two experimental sessions. Their ages ranged from 23 to just under 40 and eleven were male. Two had previous experience with accelerometers and mobile devices, but none had experience with the cursor control mechanism described above. Ten participants were right handed and two were left handed, and all used their dominant hand for this study. The effect of this factor is considered in the next section.

3.4 Hand Used To Tilt The Device

The hand used by the participant to tilt the device is an important factor when it comes to analysing the results. It is not uniformly easy to tilt the wrist in all direc-

tions, and the degree of tilt possible from a given starting position will be different in different directions. For right-handed users, to move the cursor to the right of the screen will require the wrist to be tilted such that the palm of the hand moves towards the wrist. For a left-handed participant moving the palm of the hand towards the arm will move the cursor to the left of the screen. This reversal is only true in the one axis of the wrist. Since this study is examining the restrictions placed by the body on wrist tilting interfaces, when analysing the results we must take into account the hand used by the participant during the study. The correction made for left handed participants is to switch the results obtained for targets on the left with the corresponding target on the right such that the top-left target switches with the top-right target, the rightmost target switches the leftmost target, and the bottom-right target switches with the bottom-left target.

3.5 Measured Factors

Slip Off Errors

A slip off error occurs whenever the user moves off the current target before selection. By measuring slip off errors, we can determine how difficult the targeting task was in the different directions and make comparisons. A slip off error and recovery is demonstrated in Figure 3.

Figure 3. Cursor trace showing a slip off error and recovery.

Trajectory Analysis

It is important to consider the ease of movement in different directions when creating gestures. This is particularly the case for applications involving rotation of the wrist where some directions may be more difficult to tilt in than others. Data was separated into different directions of movement (to the different targets) and analysis was carried out to look for paths that resulted in a high degree of variability from the ideal (direct) path to the different targets. The measure of variability used was the distance travelled when moving between targets. Moving from the central target to the edge of any of the outer targets in a straight line was 85 pixels in length. Excess path length was therefore classified as the distance travelled above this minimum.

Time to Target

This factor will measure the time taken for the user in moving onto the target. It does not include the time required to hover over the target to perform the selection.

Unintentional Movements by the User
This factor measured the noise generated by a user when holding the device still at different angles. The user hovered over a target for one and a half seconds. By analysing the middle second of this data, it is possible to estimate this noise value. This was then compared to noise generated by the sensor on a fix surface.

4. Results and Discussion

4.1 Slip Off Errors

The mean numbers of slip off errors for all users in all directions are shown in Figure 4. These data are shown as the mean number of times that the user slipped off each target during the experiment. Each target had to be selected 12 times by each user.

The mean number of slip offs for each is relatively small when compared to the variability in the data. These data suggest that users found moving to the lower targets easier than moving to the upper targets. For the top centre target, one in four of attempts to select a target resulted in slipping off the target. This is reduced to approximately one in six attempts when targeting the bottom centre target. It must be noted that there is a high level of variability in the data.

In total, there were 309 slip off errors out of 1152 targeting attempts. This number is high compared to targeting studies with other devices. This could indicate the difficulty of the task, but could also be due to the fact that users were required to hover over a target rather than click on it.
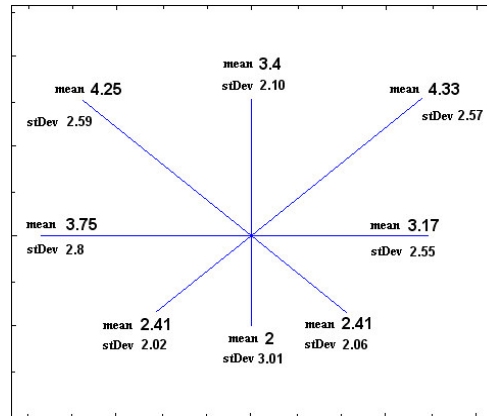

Figure 4. Mean slip off errors for each user for all targets. Each user had 12 attempts to acquire each target.

With the low cursor gain used in this study, a lower number of slip off errors may be expected since a comparatively large tilt is allowed before slipping of the target.

However, in this study, users are being asked to make large movements that required them to rotate their wrist to the limit of movement. Future studies should investigate a higher gain that allows targeting with more comfortable ranges of movement.

4.2 Trajectory Analysis

From analysis of the cursor traces, and post hoc discussion with participants, it became clear that the mapping from wrist orientation to cursor position was confusing in a small number of cases. One user in particular had expectation of the opposite mapping. The trajectory data was initially analysed to detect cases where this occurred. These cases were defined as cases where the user initially moved at least one cursor radius (15 pixels) away from the target from the start position. Three examples of such trajectories are shown in Figure 5. There were 30 out of 1152 such targeting attempts, which were spread over both experimental sessions. One user who expressed a strong opinion for the opposite mapping was responsible for 16 of these trajectories. Although, targeting was achieved without this confusion in the vast majority of the cases, these results suggest that the natural mapping is not as strong as in a similar position-control device, such as a mouse. Unlike when using the mouse, users must map a rotation to a cursor translation. More than one sensible mapping exists and different users may have different preconceptions of this mapping, making it more difficult to learn the opposite mapping. In this study, the cursor could be thought of as a marble attached to a piece of elastic. If you tilt one side of the device downwards from the start position, the cursor will move towards that side. One alternative model would be to think of the cursor as a bubble in liquid beneath the screen. This would correspond to the opposite mapping where tilting one area of the screen upwards would cause the cursor to move towards that area of the screen. These results show, however, that most users were comfortable with the mapping described.


Figure 5. Three examples of the user initially moving in the wrong direction from a centre target to the highlighted outside target.

Directional errors in movements caused by the user mistakenly moving the control device in the wrong direction have been noted by Sheridan [8]. For the errors discovered, the user consistently moved in the opposite direction from the new target. This strongly suggests confusion with the mapping rather than false anticipation of the next target. As these trajectories are most likely an artefact of confusion with the

mapping rather than difficultly in the task, they are excluded from the final analysis of the trajectory lengths.

Figure 6 shows the mean excess difference travelled by all users when travelling to the different targets. It can be seen from this figure that some directions are easier to travel in than others. Generally, the data indicates that users found selecting targets in the bottom half of the screen easier than in the top half. These results similarly suggest that lower targets are easier to select than the higher targets. Although slip off errors will have an effect, this can be considered to be minimal as the user's movements will be comparatively small when close to the target attempting to remain in the target area. Again, the high level of variability in the data should be noted, particularly when comparing variability for the targets in the upper area of the screen with those in the lower half of the screen.
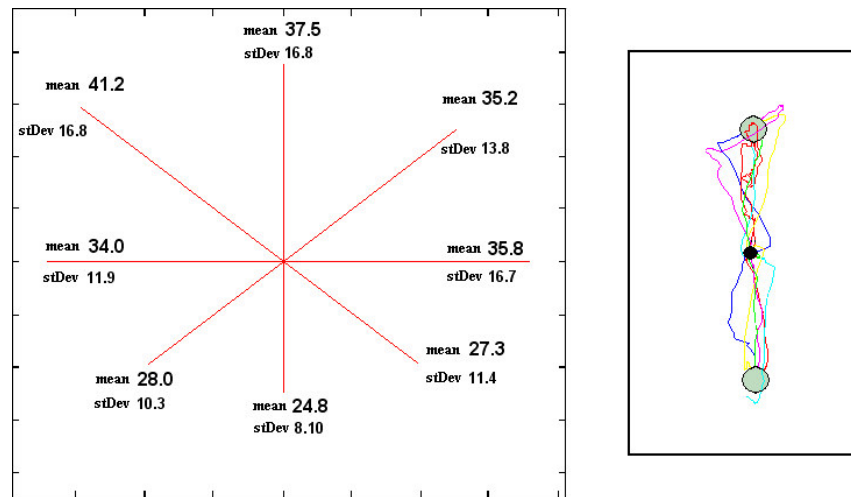

Figure 6. (Left) Mean excess distance travelled to each target in pixels. The distance travelled during target selection is not included in this measurement. (Right) Cursor trace of one user moving to the top and bottom targets six times each during one session demonstrating variability during an individual trial.

The right of Figure 6 displays six trajectories for a typical user targeting the top centre and bottom centre targets in the same experimental session. The variability displayed can be used to explain the longer path length noted when moving to the upper targets. This can be explained by the dynamics of the arm. For a posture where the user holds the device and looks at the screen, it is difficult and uncomfortable to rotate the hand such that the palm faces upwards and the screen is still at the appropriate rotation. There is a far greater range of movement when rotating the palm downwards.

4.3 Timing Data

The mean time to target data for all users is displayed in Figure 7. The differences in time between the upper and lower targets are small in this instance and may be explained by the larger number of slip offs in the upward direction. This suggests that time to target is approximately uniform in all directions for wrist tilt applications.


Figure 7. Targeting time in second for each of the outer targets

4.4 Unintentional Movement by the User

Unintentional tilts generated by the user while hovering over a target were measured in the x and y directions. A measure of the variability was given by taking the standard deviation of the mean change in tilt value during one sample point for each individual target. Only the middle second of data during the target selection was considered to allow for the user moving onto the target and moving in anticipation of the next target. For illustrative purposes, sensor readings have been converted to approximate angle in degrees.

These values are shown to be consistent for all targets in x and in y. Although, differences are small, the variability in the y direction seems to be consistently smaller than the equivalent in the x direction. This could be due to the targets being smaller in the y direction due to the higher y gain. However, since the target radius in each direction would allow for a rotation of approximately 7.2 degrees in the x direction and 5.4 degrees in the y direction which is significantly higher than the variability values recorded. One other possibility to be considered is the positioning of the accelerometer. As the accelerometer is placed at the centre of the base of the device, it is at the centre of rotation in the x direction but offset in the y direction. This means that for the same tilt in x and y, the extra leverage due to the displacement of the

accelerometer in y will lead to higher accelerations in that direction. If this were the cause, the opposite effect would have been expected since smaller tilts in the y direction would have moved the accelerometer a larger distance.


Figure 8. Approximate variability in degrees when hovering over the target in the direction indicated. (Left) is X variability. (Right) is Y variability.

When the device was flat and at rest on a solid surface, the device generated the equivalent of 0.26 degrees tilt in x and 0.24 degrees tilt in y. It can be seen from Figure 8 that these values are far smaller than the measured values for the device when held by a user at different angles.

By constantly monitoring the variability of the sensor readings, it is therefore possible to detect when the user is holding the device in a controlled fashion, and when it is resting on a surface. This provides similar functionality to that proposed by Hinkley *et al.* [1] but using accelerometer data rather than an extra touch sensor. This context information would provide programs running with information about the state of use of the device that can be used to modify its behaviour.

5 Conclusions and Future Work

This paper has examined the variability in movement in different directions for short wrist-based target acquisition with visual feedback. The results demonstrate that the direction of cursor movement affects the performance of the user in a tilting task. With the marble control metaphor described, users displayed more variability and lower performance when moving to targets in the upper half of the screen compared to targets in the lower half of the screen. No time difference was detected when moving to the upper or lower targets. The results suggest a high level of variability in the movements. It should be noted, however, that the system described in this study was not designed to produce optimal targeting results but explore variability in motion. Performance would be expected to improve with a higher cursor gain and different selection mechanism. This information can guide interface designers, as to the relative difficulty of different tilt-motions.

The ease of use of the mouse has demonstrated how a non-linear control display gain can provide a natural mechanism for interaction. Our future work will look at inverting our model for wrist-based tilting to enable us to achieve uniformly easy

tilting behaviour in all directions. There is the potential in tilt-based interfaces to compensate for different levels of variance in different directions by adapting the dynamics of the cursor depending on the state and velocity vector - the handling qualities would be more damped in regions of higher variability. The trajectories will be further analysed to examine the possibility of using the individual user variations and movement characteristics to identify that user.

Future studies will initially examine wrist tilt cursor control with higher gain levels and eventually lead to developing interactive systems that provide changing dynamics to aid the user's movements, and reduce variability. These methods will also be applied to coping with disturbance, particularly for interaction in a mobile context.

Acknowledgements

References

[1]   J. Accot and S. Zhai, "Refining Fitts' law models for bivariate pointing," in Proceedings of ACM CHI, Fort Lauderdale, Florida, 2003, pp. 193-200.

[2]   K. Hinkley, J. Pierce, and E. Horvitz, "Sensing Techniques for Mobile Interaction," in Proceedings of ACM UIST, 2000, pp. 91-100.

[3]   R. J. Jagacinski and J. M. Flach, *Control Theory for Humans*: Lawrence Erlbaum Associates, 2003.

[4]   I. S. MacKenzie, T. Kauppinen, and M. Silfverberg, "Accuracy Measures for Evaluating Computer Pointing Devices," in Proceedings of the ACM CHI, 2001, pp. 9-16.

[5]   K. Partridge, S. Chatterjee, V. Sazawal, G. Borriello, and R. Want, "Tilt-Type: Accelerometer-Supported Text Entry for Very Small Devices," in Proceedings of UIST, 2002.

[6]   A. Pirhonen, S. A. Brewster, and C. Holguin, "Gestural and Audio Metaphors as a Means of Control for Mobile Devices," in Proceedings of ACM CHI, Minneapolis, 2002, pp. 291-198.

[7]   J. Rekimoto, "Tilting Operations for Small Screen Interfaces (Tech Note)," in Proceedings of UIST, 1996.

[8]   T. B. Sheridan and W. R. Ferrell, *Man-Machine Systems: Information, Control, and Decision Models of Human Performance*: MIT Press, 1974.

[9]   J. Williamson and R. Murray-Smith, "Dynamics and probabilistic text entry," Department of Computing Science, Glasgow University DCS Technical Report TR-2003-147, June 2003.