Pressure-Based Menu Selection for Mobile Devices

Graham Wilson, Craig Stewart and Stephen A. Brewster
Glasgow Interactive Systems Group, Department of Computing Science
University of Glasgow, Glasgow, G12 8QQ, UK
+44 141 330 3451/4966

{gawilson, craig, stephen}@dcs.gla.ac.uk www.dcs.gla.ac.uk/~stephen

ABSTRACT
Despite many successes in desktop applications, little work has looked at the use of pressure input on mobile devices and the different issues associated with mobile interactions e.g. non-visual feedback. This study examined pressure input on a mobile device using a single Force Sensing Resistor (FSR) with linearised output as a means of target selection within a menu, where target menu items varied in size and location along the z-axis. Comparing visual and audio feedback, results showed that, overall, eyes-free pressure interaction reached a mean level of 74% accuracy. With visual feedback mean accuracy reached 85%. Participants could accurately distinguish up to 10 pressure levels when given adequate feedback indicating a high level of control.

Categories and Subject Descriptors
H.5.2 [**User Interfaces**]: Haptic IO

General Terms
Performance, Design, Human Factors.

Keywords
Pressure input, non-visual feedback, mobile interaction.

1. INTRODUCTION
Isometric input, more commonly referred to as force or pressure input, has seen a rise in both popularity and coverage recently, not just in academic usability research (e.g. [1-3]) but also in prototype products such as the Microsoft Pressure Sensitive Keyboard [4] and Synaptics' Fuse™ concept phone[1]. Human control [5] and perception [6, 7] of force output are highly accurate as well as efficient [8] and this fine control is becoming better realised in HCI research. Pressure input has been found to provide accurate and fine-grained control over single axis manipulations within desktop environments, including linear targeting [9-11], shape translation [3], zooming [12, 13] and password security[2]. A small number of studies have also found effective use of pressure input on mobile devices [2, 14, 15]. However, despite these successes, there is far more work on desktop implementations than on mobile devices.

The proliferation of touchscreen phones, many with multitouch capabilities, has improved the degree of control that users have over mobile applications, with accelerometer or gyroscope sensors adding further interaction possibilities. However, despite this move towards a richer touch and manual experience, the potential benefits of pressure input on mobile devices have yet to be properly explored. Pressure can add real-time and fine-grained control over a single axis of manipulation such as menu traversal, zooming and scrolling. These are frequently implemented on mobile touchscreen devices as "flick" or "pinch" gestures. Although an application can react to the speed of the flick or the rate of pinch, navigating down a long menu/document or zooming any distance requires multiple flicks or pinches, which can then occlude part of the display. Having a pressure sensor on the device would allow for continuous control of these, and other, manipulations leaving the screen fully visible. The amount of zoom, or the rate of traversal/scrolling, could be continuously based on the amount of pressure applied.

A further problem with touchscreen interaction is controlling an application while mobile. Mobile devices with physical buttons can still be operated to an extent "in pocket" as the user can identify buttons through touch in order to turn down volume or skip a track while listening to music. Due to smooth, featureless touchscreens, nearly all functions require that the user look at the display in order to press the right virtual button, which in turn requires the user to divert visual attention away from his/her surrounding environment. Touchscreen devices could benefit from pressure interaction coupled with audio feedback in such situations. Pressing on the device could operate traversal of a simple linear menu non-visually, with the interaction being aided by audio feedback.

The purpose of the study presented in this paper is to further investigate the usability of pressure input on a mobile device. Specifically, it will examine control of pressure input using non-visual feedback, something that would benefit interactions in a mobile environment. It will also emphasise the importance of using a properly calibrated pressure sensor, using one developed by Stewart *et al.*. [16] as well as how it can drastically improve control of pressure for user interactions in comparison to raw sensor output and different sensor value-to-pressure mappings.

2. BACKGROUND

2.1 Control of Pressure

2.1.1 The Right Number of Levels
The majority of investigations into the controllability of pressure in HCI have been based on how many distinct divisions, or levels, of pressure a user can accurately differentiate and apply, thus indicating the resolution (and bandwidth) of a pressure-based interaction. In this paper the term "pressure space" is used to refer to the amount of navigable pressure involved in an interaction, from 0 pressure up to the maximum that is used in the interaction (be that the maximum a sensor can detect or a limit enforced by the experimenters). Pressure input can generally be divided into 2 forms: *position-based* and *rate-based*. *Position-based* pressure input maps the absolute amount of pressure to position within an

[1] http://www.synaptics.com
[2] http://jdadesign.net/?page_id=37

MobileHCI'10 September 7 - 10, 2010, Lisbon, Portugal.
ACM 978-1-60558-835-3/10/09.

interaction (e.g. cursor position). *Rate-based* input maps the amount of pressure to the speed of interaction (e.g. rate of list traversal).

Pressure input has been studied most in desktop environments, with initial work using graphics styli housing a pressure-sensitive tip. In a simple set-up, the width or density of lines and shapes drawn using the stylus can be increased by applying pressure to the tip. Ramos *et al.*. [9] used such styli to investigate the feasibility of using pressure for more general interactions, including the manipulation of graphical widgets. They recommended that any (positional) pressure-based widget should employ a maximum of 6 distinct levels, as above this number, accuracy and control declines considerably. An important finding of this study, which was to influence many subsequent studies, was subjective reports of control difficulty at low levels of pressure (as much as the first 30% of the pressure space) corroborated by measures of control at these levels. The authors suggest this may be due to poorer human control at these levels, however, as will be discussed later, it is much more likely that this is an artefact of sensor design.

Cechanowicz *et al.*. [10] introduced pressure input into more traditional and common desktop control scheme by attaching dedicated Force-Sensing Resistors (FSR) to the body of a mouse to be operated by the thumb. Noting the lack of control at low levels in [9], Cechanowicz *et al.*. applied a quadratic discretization function to the sensor output so as to improve accuracy, making low pressure targets larger and high pressure targets smaller. They found that users could control up to 8 or 10 discrete levels. Lamenting the generally high error rates being exhibited in pressure interaction, particularly the difficulty in mapping pressure input to sensor output, Shi *et al.*. [11] applied a fish-eye mapping to the sensor output so as to artificially increase the size and visibility of the currently highlighted level. By doing this they found users could control up to 10 discrete levels at approx 75-80% accuracy or up to 16 discrete levels with approx 60% accuracy.

There has been far less work on pressure input on mobile devices. Mizobuchi *et al.*. [15] looked at using a stylus to press on a pressure-sensitive screen as a means of target selection similar to those above. They calibrated their sensors so as to divide all the targets into equal sizes (in Newtons) and found that users could select 10 levels with >= 85% accuracy. Moving away from target selection, Brewster & Hughes [2] looked at using pressure input to aid text entry on a mobile device with a pressure-sensitive screen (on a Nokia N810 Internet Tablet). Using just 2 levels of pressure, the lower level generated lower case letters while the higher level generated uppercase letters. They found that the pressure version was quicker than the normal method for inputting uppercase letters, which necessitates pressing the shift key. An important contribution of this paper was that users also carried out pressure-augmented text entry while walking, with results showing that pressure control was robust even while in motion. McCallum *et al.*. [14] also found pressure input to benefit text entry, using a pressure-sensitive variant of a traditional phone keypad. As single keys input multiple letters (e.g. 2 = a, b or c), the authors substituted multiple presses for increasing pressure levels where a single soft press outputs "a" and a single hard press outputs "c". Clarkson *et al.*. [17] suggest further uses for pressure-augmented keypads such as preview zooming, 3D navigation or "affective input" where emotional state is derived from the degree of force.

Even from the small number of studies conducted on mobile devices, it can be seen that pressure can be a usable and beneficial augmentation to mobile interactions, with comparable performance to desktop variations. However, all of the above results involved providing the user with continuous visual feedback and, as will be seen, the amount and type of feedback provided can severely influence the accurate application of pressure. Due to the limitations of mobile devices in terms of input and output, and the need to focus visual attention on the environment when mobile, being able to operate functions non-visually is an important interaction issue. Therefore it is important to understand how feedback affects the accuracy of pressure application so as to best design non-visual feedback to mitigate these effects. Only a small number of studies have looked at non-visual feedback in isolation in relation to control of pressure [16][18, 19] so there are currently no guidelines on best practice.

2.1.2 The Role of Feedback

Providing a user with external feedback regarding the amount of pressure being applied (i.e. not from proprioceptive or afferent signals within the body) has been found to aid accurate production of target force levels [5, 18, 19]. This suggests that we cannot simply rely on how much our muscles or sense of touch is telling us that we are pressing.

Ramos *et al.*. [9] found that reducing the amount of visual feedback that users were given, even after having had one hour of experience using the device with full visual feedback, adversely affected both accuracy and control of pressure input. They go on to suggest that full and continuous feedback is necessary, particularly in the early stages of learning an interaction. This is in line with research on the interaction between feedback frequency and perceptual-motor learning where frequent feedback aids in the initial acquisition of tasks [20]. Mizobuchi *et al.* [15] compared continuous visual feedback with discrete visual and no feedback and found that discrete, and particularly no feedback, produced worse performance than continuous visual feedback. Ramos & Balakrishnan [13] reported that audio feedback was appreciated by users when used in conjunction with visual feedback, as it gave information that otherwise required monitoring a small on-screen widget. The design of the feedback is not described, other than referring to it as "brief" and it was used to indicate the "clutching" (holding) of a pressure level.

Stewart *et al.* [16] looked at user ability to acquire one of 3 targets with only visual, audio or vibrotactile feedback (or a combination of both audio and tactile feedback). Although visual feedback produced near perfect performance, audio-only produced a mean of 69% accuracy, vibrotactile a mean of 82% and the combination a mean of 71%, suggesting pressure can be controlled with varying success, with non-visual forms of feedback. This disparity in performance between visual and non-visual feedback came predominantly from errors acquiring the middle target of the three. It is not clear why audio-only feedback performed worse than vibrotactile-only however, it has been suggested that vibrotactile feedback has a more direct mapping to pressure input and that interactions benefit from having the same input and output channel, for example touch output and vibrotactile input [18].

From these results it appears that providing full and continuous feedback concerning the level of pressure being applied is essential for accurate control, suggesting non-visual feedback should do the same. It should also be noted, however, that although increased feedback benefits acquisition of a task (improving initial performance compared to less feedback) the review by Schmidt & Bjork [20] suggests that less frequent feedback is beneficial, or even necessary, for longer-term retention of a skill, suggesting that lessening feedback over time improves retention without adversely affecting acquisition. This may mean that the effects of reducing or changing feedback may be mitigated over time and that expert performers will have different feedback needs from novice users.

30% of participants in [15] were quoted as preferring discrete feedback to continuous feedback. In the experiment continuous feedback came in the form of a coloured bar which filled dynamically relative to the amount of pressure; the discrete feedback was simply a number displayed on screen indicating the current pressure level that the user was in. Those expressing preference for the number said so partially because the continuous feedback was "distracting" as it was constantly changing. Those preferring the number did not benefit from the continuous feedback. This may suggest that, regardless of experience level, there are individual differences in feedback requirements.

For the purposes of our study we chose audio feedback as the non-visual channel. A number of participants reported appreciating audio feedback in [13] and [19] found that audio feedback provided quicker pressure-responses than vibrotactile feedback (at the cost of accuracy). They asked users to perform pressure "chords" with three fingers where the amount they were to press (low, medium or high) with each finger was indicated through either one of three colours (Green, Blue or Red), simple audio tones (of 100, 600 or 1100Hz) or through frequency of tactile burst (1.5, 3 or 6Hz). Finally, the amount of information that can be provided through audio remains greater than that which can be provided through vibrotactile feedback. Although Tactons [21-23] have greatly expanded the communication bandwidth of vibrotactile feedback, the vibrations remain comparatively basic when contrasted with the rich tapestry of real and abstract sounds, as well as speech, available through audio feedback. This is predominantly due to the basic output capabilities of current commercial vibrotactile actuators, rather than limitations in human perception, however.

3. EVALUATION

3.1 Pilot Study

A pilot study was run as an initial investigation into non-visual pressure-based interaction. Target selection along a single axis (adapted from [9]) has been a common and effective way of demonstrating control of pressure in many other studies and so is used again here for comparison. This task involves dividing the pressure-space into a set number of levels, or divisions, of equal width (in Newtons), placed along the given axis and having the user move a cursor along the axis by applying pressure using position-based input (i.e. the amount of pressure dictates the location of the cursor; see Figure 1**)**. When users have applied enough force to place the cursor in the target level, they activate a selection mechanism to confirm selection. A continuously moving cursor was chosen over discrete feedback (for example highlighting levels relative to applied pressure) as previous work has suggested that continuous feedback is beneficial for perceptual-motor control, especially when first learning such actions [9, 20]. By dividing the same pressure space up into larger numbers of levels, the levels become smaller, giving an indication of the limits of control. Varying the position of the active target throughout the pressure-space gives an indication of control of pressure at different levels. Previous work suggests that user accuracy remains high up to 10 levels when using positional pressure input and so this was chosen as the maximum number of divisions in this study as along with 4, 6 and 8 levels.

3.1.1 Selection Techniques

The selection techniques compared in this study were *Quick Release* and *Dwell*. *Quick Release* involves lifting the finger/thumb from the sensor when the cursor is in the target level, whereas *Dwell* requires the user to remain in the given target level for a set length of time to confirm selection. In this study a *Dwell* duration of 1 second was chosen. In initial testing 500ms was chosen so as to increase the speed of interaction and this was also the length of time used successfully in [2], however after a high number of erroneous selections this was increased to 1 second. This length has been found to be a suitable length of time in other similar interactions [10]. The *Quick Release* mechanism is generally more error prone than *Dwell* but is usually much faster [2, 9, 10].

3.1.2 Feedback Design

To give the task more relevance to real-life mobile use, the interaction was designed to resemble traversing a flat linear menu and selecting menu options, with each pressure level being given a unique label that one might find in a typical application. The labels chosen are common menu items found in various applications:

File, Edit, View, Format, Bookmarks, Text, Tools, Window, Help, Exit.

The order of the items never changed, only the number that were placed on screen, starting with File, so pressure menus with 4 items ended at "Format", 6 items went up to "Text" and so on. The visual feedback displayed the pressure levels as equal sized grey rectangles aligned vertically in the middle of the screen (see Figure 1). A small cursor moved vertically just outside the menu, indicating the amount of pressure being applied in a continuous form. The active target for any given trial was displayed briefly in bright green at the start of the trial. This design provided the continuous feedback necessary for successful target acquisition in a perceptual-motor task as identified by [20] and [9]. Additionally, we gave the pressure levels common labels to aid familiarisation with the interaction.

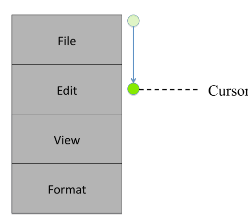

Figure 1. Visual layout of pressure menu.

For the audio feedback conditions the screen was left blank. The movement of the cursor in relation to pressure input, and the positions and layout of the menu items were all the same as in the visual condition, only the participant could not see them. To inform users of which menu item they were in, the items label was spoken in synthetic speech once as the cursor entered the item from either side (i.e. entering it by increasing pressure from 'below' or by decreasing pressure from 'above'). If the cursor moved so fast as to enter another item before such speech had finished playing, the initial audio was stopped and the newly entered items label was spoken. To help users identify when they were on the verge of crossing over into the next menu item, a warning tone (chord of 2 notes: 350Hz & 460Hz) was played when the cursor enters the last 25% of any menu item. This would help avoid accidentally moving into a menu item unintentionally, or would help participants know how much further to go to deliberately move into the next item.

3.1.3 Participants

Fourteen participants (7 Male, 7 Female) aged between 20 and 32 (mean = 22.2 years) took part in the evaluation, all of whom were studying or working in the University of Glasgow. All were right-handed and all were paid £10 for participation, which lasted no more than 90 minutes.

3.1.4 Experimental Design & Procedure

The experimental software was run on a Nokia N810 (see device in Figure 3) using the pressure-sensitive resistive touchscreen as the input sensor. The sensor output a value between 0 and 1 relative to the amount of pressure being applied and this was divided into 1024 pressure levels to allow comparison with previous work [9-11]. Unfortunately, the sensitivity of the sensor was not uniform around the screen resulting in uneven output depending on the point of contact. To minimise this effect a specific location was chosen as the contact point for all participants. A black square outline was placed on screen to indicate where participants were to press. Due to the uneven behaviour of the sensor it was not possible to accurately calibrate the sensitivity of the sensor (and so the size of the pressure space) in Newtons.

Participants held the device in both hands, using their right thumb to press on the screen. It is important to note, however, that the pressing action more closely resembled a pinch between thumb and first or second finger, due to the way the device was held. The thumb pressed against the device, which was then resisted by the fingers behind the device mimicking a thumb-finger squeeze. Audio feedback was delivered through headphones from the audio jack of the N810.

3.1.4.1 Variables

Ramos *et al.* [9] designed their experimental task so as to be able to measure conformity of pressure-based target selection to Fitts' Law [28]. Because dividing the pressure space into differing numbers of levels produced different sizes of targets, this meant that they were different distances away from the start point. The authors therefore chose four targets from each number of divisions which have within them a common distance. This meant they could compare acquisition of smaller targets at similar distances. These same four distances were used in the pilot study and equate to 205, 410, 615 and 820 sensor values.

A selection was considered an error (ER) if it occurred outside the target pressure level, by either lifting the thumb outside that level in *Quick Release* or by remaining in any non-target menu item for 1 second in *Dwell*. If an error occurred an error tone was played to the user. There was no direct feedback for correct selections and after either an error or correct selection the next trial started after a pause of 2 seconds (which was accompanied by a blank screen in the visual condition). It was decided to make it impossible for an individual to "overshoot" the last item in a menu, as it was assumed that this would be the case in a real-life implementation of this type of task in an application. If the cursor entered a target level and subsequently exited it again, this was counted as a crossing, and was used as a measure of control (number of crossings or NC), where a lower number of crossings was equated with a higher degree of control. Movement time (MT) measured the time from the first non-zero reading from the sensor up until selection, be that an error or a correct selection.

The experiment used a within-subjects repeated-measures design with 4 independent variables:

- Number of menu items: 4, 6, 8 and 10
- Target distance: 205, 410, 615 and 820.
- Selection method: Quick Release and Dwell
- Feedback modality: Visual and Audio

With the dependent variables:

- Number of errors (ER)
- Movement time (MT)
- Number of target crossings (NC)

3.1.4.2 Procedure

The whole task was split into 2 halves: one using only the *Quick Release* selection technique and one using only the *Dwell* technique. Within these conditions were one visual-only and one audio-only feedback condition leaving a total of 4 conditions referred to here as Quick-Visual; Quick-Audio; Dwell-Visual and Dwell-Audio. In order to remove potential ordering issues, half of the participants took part in the Quick Release conditions first and the other half took part in the Dwell conditions first.

All participants took part in all conditions, with the ordering of conditions counterbalanced except for the first 2 feedback conditions. To facilitate familiarisation with the interaction as a whole, all users first engaged in a visual feedback condition under their first selection technique, followed by an audio condition. The order of conditions under the second selection technique was then counter-balanced to reduce possible bias towards audio feedback.

Under each selection technique x feedback pairing (e.g. Quick-Visual) there were 3 blocks each including 8 target acquisitions of all 4 target sizes (each of 4 target distances was presented at random twice). This gave a total trial count of: 14 participants x 2 selection techniques x 2 feedback techniques x 3 blocks x 4 target sizes x 4 target distances x 2 repetitions = 5376 trials.

Each trial involved the user acquiring one target pressure level. For the visual conditions the whole menu would be displayed in grey apart from the target item, which was shown for 1 second in bright green before returning to the same grey colour. The participant would then use the sensor to move the cursor to the desired level before either lifting their finger when using *Quick Release* or remaining in the item for 1 second when using *Dwell* to confirm their selection. For the audio conditions the screen would be blank. The participant would first be informed of the number of menu items through the spoken phrase "[number] items" where [number] is either 4, 6, 8 or 10. The name of the target item to acquire would then be provided by the phrase "Get [item]" where [item] is one of the labels given above. The user then presses on the sensor to move the cursor to the desired pressure level before again confirming selection through *Quick Release* or *Dwell*.

3.2 Results

All analyses involved 2 x 2 x 4 repeated measures ANOVA.

<u>*Error Rate*</u>

Analysis revealed a significant main effect of selection technique ($F_{1,153} = 72.463$, $p<.001$): *Quick Release* had significantly fewer errors (32%) than *Dwell* (50%). There was also a significant effect of feedback type ($F_{1,153} = 313.672$, $p<.001$): Audio feedback produced significantly more errors (56%) than Visual Feedback

(26%). Error rate also increased as the number of menu items increased with mean error rates of 31%, 36%, 45% and 50% for 4, 6, 8 and 10 items respectively.

Movement Time

Both selection technique ($F_{1,1234} = 22.752$, $p<.001$) and feedback type ($F_{1,1234} = 59.115$, $p<.001$) had a significant effect on movement time: *Dwell* produced a significantly higher average movement time (2.7s) compared to *Quick Release* (2.3s). Audio feedback had a significantly higher average MT (2.8s) than Visual Feedback (2.2s).

Number of Crossings

Similarly, both selection technique ($F_{1,1234} = 74.289$, $p<.001$) and feedback type ($F_{1,1234} = 44.434$, $p<.001$) had a significant effect on control. *Dwell* technique lead to a significantly higher average number of crossings (2.3) per target compared to *Quick Release* (1.4). Visual feedback produced a significantly higher average number of crossings (2.2) per target compared to Audio feedback (1.5).

3.3 Discussion

The findings of the pilot were somewhat disappointing with high error rates for the Dwell technique and for Audio feedback. Also the error rates found for all number of menu items were well below the results of [9] and also worse than [10] achieved with higher numbers of levels (12 or even 16). From the data and from subjective reports by users, two primary contributing factors were identified for the poor results: the pressure-sensitive screen used and the audio feedback.

Sensor Deficiencies

Although different sensors use different analogue-to-digital converters, there is a common problem in that they are often disproportionately more sensitive to light touches compared to moderate or high pressure. This was found to be the case by [9] as well as [10, 11, 16]. This was also the case with the N810 screen as users complained that the low levels were much less controllable and error prone than farther levels. This lack of a uniform, or linear, relationship between pressure and cursor behaviour confused users and made holding the cursor at a given level (particularly low levels) much more difficult. Being less able to accurately hold the cursor at a desired level had a greater negative effect on the *Dwell* condition, which requires precise control over time. For the *Quick* condition users stated that they simply lifted their thumb "as soon as they heard" the target label, requiring little 'fine-tuning' of cursor position. Given the common problems across digital pressure sensors this suggests a fundamental problem with their use in HCI particularly with such a fine and accurate input channel as human application of pressure.

As can be seen from the research in the background section above, there have been a number of different sensors used (e.g. styli, single FSRs, touchscreens, and screens with FSRs underneath) in a number of different studies using a number of different ways to translate pressure input into an appropriate sensor output (e.g. simple sensor-to-pixel mapping, quadratic discretization, fish-eye function etc). All of these mappings are attempts to make an interaction more usable by making the correlation between pressure and sensor behaviour more linear. Mizobuchi *et al.* [15] calibrated their device so that the targets were equal in size in Newtons, regardless of the sensor output, and they found no such errors or difficulty in acquiring targets at low-pressure levels.

By attempting to make a non-linear output appear more linear through changing how a device treats the output signal, some authors are perhaps treating the symptoms and not the cause of non-linear sensors. By using different sensors with different treatments of the signal, no two studies are truly directly comparable. However, having a linear output from the sensor itself there would be closer correlation between pressure and output and a much more standardised foundation on which to build an interaction, giving a common sensory baseline for studies using such a device.

Such a sensor was recently developed and tested by Stewart *et al.* [16]. A full description is not within the scope of this paper and interested readers are directed to the original paper for a more detailed explanation of design and implementation. They attached an opamp-based current to voltage converter to FSRs (which was then attached to an Arduino[3] interface for A-to-D conversion and output) to produce a good fit to a linear function ($p = 0.0008x + 0.0339$; $R^2 = 0.97$) between pressure applied and the output signal. The authors compared the linear signal to a quadratic mapping (similar to that used by [10]) and found that the linear sensor allowed for a greater degree of control with users performing significantly better than with a non-linear output for a 3 pressure level task. Therefore outside of the benefit of having similar sensors, using properly linear sensors also benefits control of pressure in interactions.

Audio Feedback

The poor results for audio feedback suggest that the design choices were not as useful for orienting around the menu as initially hoped. Participants were encouraged to familiarise themselves with the order and layout of the menu items during their first visual condition which would then aid them in navigating the audio feedback. This proved highly troublesome, however, as users were unable to familiarise themselves well enough and often became "lost" within the menu, not knowing where they were or where the target item was relative to their current position. The general difficulty traversing the audio menu was exacerbated by the poor control offered by the N810 screen. These problems influenced a redesign of the feedback for the main evaluation.

3.4 Main Evaluation

3.4.1 Pressure Sensor Used

Due to the promising results of [16] and the poor accuracy of the N810 pressure-sensitive screen, we chose to conduct the main evaluation using the dedicated FSR-based sensor developed by Stewart *et al.*.

3.4.2 Task

The experimental task was identical to that of the pilot study with one exception. In the pilot we followed the design of [9] of comparing only those targets that lie at similar distances (we chose this design so as to provide comparison of results with Ramos rather than to compare our results to Fitts' Law). Other studies have also followed this experimental design [10, 11] however, in doing so, the results can only ever examine selection at 4 distances, not the full number of levels stipulated in the interaction (i.e. 6, 8, 10, 16, 64). Sacrificing comparisons with Fitts Law, we looked at selecting targets at all distances from 0. It is hoped this will give a clearer picture of pressure control across the entire

[3] http://www.arduino.cc

interaction space (a future experiment will then look at the fit with Fitts' Law). The selection mechanisms used were exactly the same as in the pilot: *Quick Release* and 1-second *Dwell*.

3.4.3 Feedback Design

The visual feedback for this study was identical to the visual feedback in the pilot. However, the audio feedback design was changed significantly. The main problem from the pilot seemed to stem from a lack of positioning. In the pilot study, participants complained of being "lost" in the menu, i.e. not knowing where they were or where other items were in relation to their position. Given the spatial nature of visual feedback it is easy to see where display elements are relative to others. The audio was simply presented monaurally in the pilot study providing no such spatialisation as in the visual feedback. We attempted to rectify this by using panned audio around the head for the main evaluation. The audio was now laid out as if the menu ran horizontally across the front of the user so that the first menu item was always on their far right and the last item always on their far left (see Figure 2). Panning was achieved by simply altering the volume output (0 to 100) to the left and right ears so that, for example, a volume of 0 (left) and 100 (right) indicated positioning at the far right and 70 (left) and 30 (right) indicated position left of centre. Although this is a different orientation compared to the visual menu, several studies have found spatialised audio around the head to be suitable for mobile interactions [24-26].

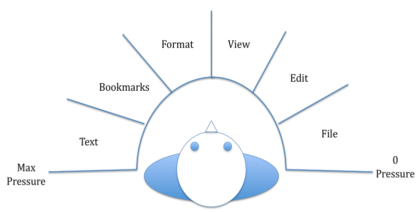

Figure 2. Panned audio design for main evaluation.

The label of each menu item was still spoken by synthetic speech whenever the cursor entered that item from either side, identical to the pilot. Each item was also given a unique musical note that played for the duration of the time that the cursor was in that item. Both the label and the note were played in panned audio in the position of where the item lay on the menu e.g. "File" was always heard on the far right. Mizobuchi *et al.* [15] reported that users instinctively aimed for the centre of the targets in their study and some participants in our pilot study reported gaining no benefit from the warning tone at an items edge so, rather than a warning tone, a second note, one octave above the given item's unique note, was played when the cursor was in the central third of the item.

Therefore as the cursor moves up (or down) the menu, the names and notes of each menu item play in 1-dimensional space around the user from right-to-left (or left-to-right on the way down). From hearing the location of the label or note in its position relative to the start (right) and finish (left), the user gets a spatial clue as to it's location in the whole menu. For example, hearing "Bookmarks" slightly to the left of the head tells the user it is quite far up the menu. One final addition changed the way the user was informed of any given active target. In the visual feedback, the user can see what the active target is (it is briefly coloured bright green) and, automatically, can see how far down the menu it is. In the pilot study it was simply spoken to the user in the form of e.g. "Get Bookmarks". But unless the user is familiar with the layout and ordering of the menu items, this does not indicate where *Bookmarks* is in the menu, unlike the visual feedback. Therefore, for this evaluation, during the phrase "Get [label]", the name and note of the target item was played in its relative panned position before each trial, indicating where in the menu that item was.

3.4.4 Participants

Seventeen male participants aged between 19 and 35 (mean = 21.5 years) took part in the evaluation, all of which were studying or working in the University of Glasgow. The gender bias was not intentional; a request for volunteers was issued and acceptance was only received from these male participants. Sixteen were right-handed and all were paid £10 for participation, which lasted no more than 90 minutes. None had taken part in the pilot study.

3.4.5 Experimental Design & Procedure

The experiment used a within-subjects repeated-measures design with the same 4 independent variables as the pilot with the exception of:

- Target Distance: 28 distances within the 4, 6, 8 and 10 item menus.

Figure 3. Experimental set up for main evaluation. Circular FSR is under white adhesive tape.

3.4.5.1 Apparatus

The apparatus was set up as seen in Figure 3. The Nokia N810 was used to run the experimental software and provide the visual and audio feedback. The FSR was attached to a piece of firm Perspex (under white adhesive tape) so as to allow for a squeezing/pinching action with the thumb contacting the sensor and the forefinger providing resistance, similar to the action from the pilot study. Initially the sensor was attached to the body of the N810 to the right of the screen so that the device could be held, as it would be in normal use. However, this positioning led to flexing of the strip connecting the sensor pad with the Arduino interface, which in turn affected the behaviour of the sensor, and it was necessary to ensure the strip remained stationary. The resulting interaction mechanics, where the sensor is manipulated in a pinch grip between thumb and forefinger, is very similar to the way it would be manipulated if the sensor were attached to the device. The sensor is powered and driven by the N810 showing that the interaction is feasible on more limited hardware. Audio feedback was delivered through headphones from the audio jack of the N810.

Due to the linear output of the sensor it was possible to accurately calibrate the output and so measure the pressure space in Newtons. Mizobuchi *et al.* [15] reported that users suffered fatigue when applying 4N, which in turn led to tremor at these levels and so decreased accuracy, suggesting a smaller pressure space would be less error prone, and more comfortable at the high end of the space. However, by reducing the pressure space, and so reducing the size of the interaction space, it also means that any number of levels must be squeezed into a smaller space making them harder to acquire. The sensor used in this study can detect a maximum of 12N, but the limit was set at 3.5N (0-3.5N) for the interaction so as to provide a large enough interaction space while reducing the potential fatigue at higher levels. The full 3.5N was split into 4, 6, 8 or 10 equal-width 'bins' of approx 0.87N, 0.58N, 0.44N and 0.35N each respectively, representing the menu items.

3.4.5.2 Procedure

The procedure for the main evaluation was identical to that of the pilot other than two details. Because of the choice to require that users acquire all target distances in a given number of menu items, there are an uneven number of selections for each number condition (4 selections for the 4 item menu, 6 for 6 items, etc.). As this would make any comparison between numbers of items uneven, only those targets that were identified by [9] as having similar distances were used to compare performance across numbers of items. For all other analyses concerning performance all target distances would be considered. Also, as each target distance was to be selected, requiring participants to acquire 2 of each distance (as was the case in the pilot) would have increased the total number of trials and total task time beyond what was considered reasonable considering participant fatigue. Therefore each distance was only acquired once giving a total of: 17 participants x 2 selection techniques x 2 feedback techniques x 3 blocks x 28 target distances = 5712 trials. Participants completed NASA TLX workload estimation forms after each condition.

3.4.6 Hypotheses

H1: There will be fewer errors in the Dwell conditions than in the Quick Release conditions.

H2: Errors will increase as the number of menu items increases.

H3: That movement time will be lower in the Quick Release conditions than in the Dwell conditions.

H4: That movement time will be lower in the visual conditions compared to the audio conditions.

H5: That the number of crossings will increase as the number of menu items increases.

H6: There will be more crossings in the audio conditions compared to the visual conditions.

3.5 Results & Initial Discussion

We removed outliers from the data set. A trial was considered an outlier if the pressure value of the selection was more than 2 standard deviations outside of the mean selection location for that target distance. A total of 291 trials were removed constituting 5% of all trials.

3.5.1 Error Rate (ER)

<u>*Learning Effects*</u>

A 2 x 2 x 3 (selection technique x feedback x block) repeated-measures ANOVA showed no significant effect of block on error rate ($F_{2,950}$ = .237, p>.05). This suggests there were no learning effects and performance did not change significantly over time.

<u>*Selection Technique & Feedback Type*</u>

Figure 4 shows mean error rate for all conditions. The mean overall error rate across all conditions was 20.5%. A 2 x 2 (selection technique x feedback) repeated-measures ANOVA showed a significant main effect of selection technique on the mean number errors ($F_{1,1168}$ = 280.908, p<.001), a significant main effect of feedback type on mean number of errors ($F_{1,153}$ = 107.070, p<.001) and an interaction between selection technique and feedback type ($F_{1,153}$ = 47.798, p<.001). Dwell had a lower error rate (11%) than Quick Release (30%). Visual feedback had a lower error rate (15%) than Audio feedback (26%). The difference in ER between the Dwell-Visual and Dwell-Audio conditions was much larger than Quick-Visual compared to Quick-Audio, leading to the interaction effect.

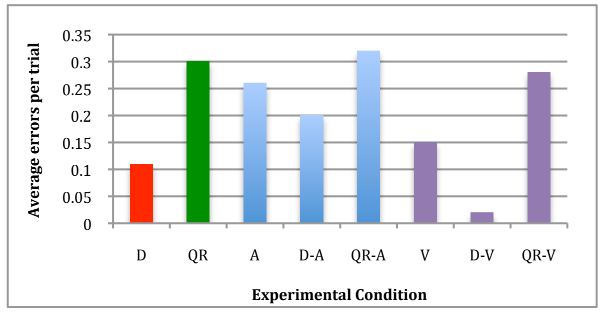

Figure 4. Average number of errors per trial for all conditions (D: Dwell; QR: Quick Release; A: Audio; V: Visual).

<u>*Number of Menu Items*</u>

Figure 5 shows the average number of errors per trial for each menu size. For this comparison only target selections from the 4 common-distance items from each number of divisions was considered (as suggested by [9]). A 2 x 2 x 4 (selection x feedback x number of items) repeated-measures ANOVA showed a significant main effect of selection technique ($F_{1,109}$ = 63.105, p<.001) and feedback type ($F_{1,109}$ = 40.100, p<.001) on mean error rate in the same directions as those above. It showed a significant main effect of the number of menu items on mean error rate ($F_{3,327}$ = 22.405, p<.001) as well as a significant interaction between selection technique and feedback type ($F_{1,109}$ = 12.082, p<.01) as above, a significant interaction between selection technique and number of items ($F_{3,327}$ = 2.665, p<.05), a significant interaction between feedback and number of items ($F_{3,327}$ = 6.115, p<.01) and a 3-way interaction between selection technique, feedback and number of items ($F_{3,327}$ = 2.758, p<.05). Mean error rate increased as the number of menu items increased with mean error rates of 9%, 17%, 25% and 26% for 4, 6, 8 and 10 items respectively. Pairwise comparisons revealed that the number of errors differed significantly for all pairs of menu sizes at significance p<.001 except for 8 x 10 items which was non-significant (p>.05).

The interaction between selection technique and number of items may exist because, in the *Dwell* conditions, 10 menu items produced fewer errors than 8 items, whereas, in the *Quick Release* conditions, 10 items produced more errors than 8 items. Error rates for both feedback conditions increase from 4 to 8 items. Upon simple inspection the feedback x number of items interaction may come from a similar uneven change in error rate from 8 items to 10 items, as it drops from 8 to 10 items under audio feedback but increases from 8 to 10 items under visual feedback. As for the 3-way interaction, error increases with increased number of items under all selection-feedback pairs except for Dwell-Audio which increases to 8 items before dropping in error rate from 8 to 10 items (see square points in Figure 5).

The results here for error rate support the acceptance of hypotheses H1 and H2, and are much more encouraging than in the pilot study and suggest that near-perfect performance is possible in pressure interaction, even with as many as 10 distinct pressure levels (in this case using the *Dwell* selection technique and visual feedback, triangles/lowermost line in Figure 5). It also suggests that non-visual interaction is also highly usable if the number of pressure levels is kept at or below 8 (again using the *Dwell* technique, square line in Figure 5). Poor performance using the *Quick Release* technique, however, was quite surprising with this being more evident in the Audio feedback condition.

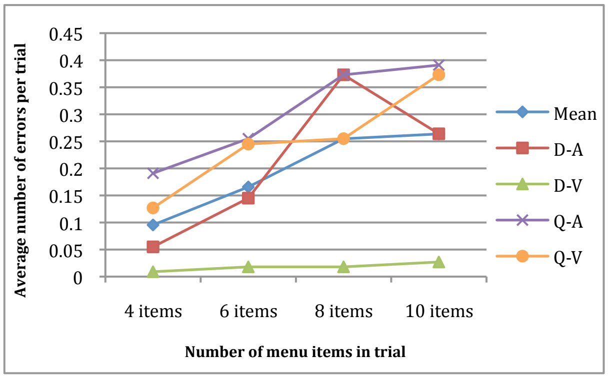

Figure 5. Average number of errors per trial for all numbers of menu items. Lines correspond to selection technique-feedback pairs.

3.5.2 Movement Time (MT)

Selection Technique & Feedback Type

A 2 x 2 repeated measures ANOVA showed a significant main effect of selection technique on movement time ($F_{1,1427}$ = 136.529, p<.001) and a significant main effect of feedback type on movement time ($F_{1,1427}$ = 565.253, p<.001). *Dwell* had a higher average movement time (3.4 seconds) compared to *Quick Release* (2.7 seconds) and Audio had a higher average movement time (3.8 seconds) than Visual (2.2 seconds; see Figure 6).

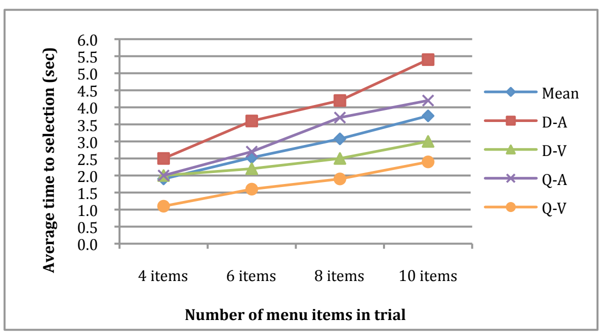

Figure 6. Average movement time (MT) per trial in seconds. Lines correspond to selection technique-feedback pairs.

Number of Menu Items

Average movement time increased as the number of items increased, with means of 1.9s, 2.5s, 3.1s and 3.8s for 4, 6, 8 and 10 menu items respectively. Average movement time also increased as target distance increased for all number of items under all conditions. In a similar trend to error rates, the last item frequently had lower MT.

MT results support rejection of the null hypothesis in favour of adopting hypotheses H3 and H4 as *Quick Release* trials were on average faster than *Dwell* trials and Visual feedback allowed quicker average selection times than Audio feedback. Audio feedback increases selection time by almost 75%.

3.5.3 Number of Crossings (NC)

Selection Technique & Feedback Type

A 2 x 2 x 4 repeated measures ANOVA showed a significant effect of selection technique on number of crossings ($F_{1,203}$ = 72.174, p<.001) as well as a significant effect of feedback ($F_{1,203}$ = 59.676, p<.001). *Dwell* had a higher average number of crossings (7.2) compared to *Quick Release* (4.7), while Audio feedback produced more crossings (7.1) than Visual feedback (4.8).

Number of Menu Items

A further 2 x 2 x 4 repeated measures ANOVA showed a significant main effect of number of menu items on the number of crossings. Mauchly's test indicated a violation in the assumption of sphericity of variance for number of items (chi-square = 105.804, p<.001), therefore degrees of freedom were corrected using Greenhouse-Geisser estimates of sphericity (epsilon = 0.74). Under this correction the significance stood at $F_{3,609}$ = 156.458, p<.001.The number of crossings increased as the number of menu items increased with means of 2.2, 4.5, 6.4 and 10.6 crossings for 4, 6, 8 and 10 items respectively. Pairwise comparisons revealed that the number of crossings differed significantly for all pairs of number of menu items at significance p<.001.

Again, the results for NC support acceptance of alternative hypotheses H5 and H6. In a very similar trend to MT, NC also increases as the number of items increases, which suggests that users take more time oscillating back and forth over targets as they become smaller.

4. DISCUSSION AND CONCLUIONS

The results from the main evaluation suggest a much better audio design was employed, compared to the pilot study, and they show that both visual and non-visual pressure-based interaction with a mobile device are both useable and highly accurate. Several accuracy rates shown here are above those found in previous studies using non-linear sensors. Shi *et al.* [11] found 78% accuracy with visual feedback using the *Dwell* selection technique, whereas we found 83% accuracy with visual feedback and *Dwell*. In the current study participants managed 10 levels at 73% accuracy using only audio feedback, almost equaling that of [11]. It would appear from the results that a pressure space of 3.5N allows for good control at up to 10 levels, particularly when using the *Dwell* selection technique with visual feedback and a linearised sensor. Contrary to the findings of [15] we found no extreme fatigue at acquiring targets at the farthest end of the pressure space. Although errors did increase as the distance increased, as they did in that study, subjective reports (NASA TLX) of thumb fatigue peaked at 8.5 out of 21, with an average report of 7.2. A particularly encouraging set of results is the near-perfect accuracy rates for all numbers of menu items under visual feedback using the *Dwell* technique. The worst performance was still only at 3% errors for 10 menu items with perfect 0% errors for 4 items.

However, in comparison to both [11] and [9], our MT and NC results were worse, suggesting that improved accuracy in the current study came at the cost of speed of interaction. Both the average movement time and the average number of crossings increased as the number of items increased, but they also tended to increase as the distance to target increased, although this was not apparent across all conditions. Given the very similar increase in both MT and NC this suggests that, rather than deliberately taking more time to carefully orient towards targets, participants are more likely unintentionally moving the cursor back-and-forth over a target in an attempt to pinpoint the small target size. Because only 25% of the sensor's range was used in the study, it is possible that there would have been more noise in the output than if the whole range had been used. The sensor can be calibrated so as to be linear across any pressure range so perhaps reducing the sensor range to 3.5N would improve control, and consequently MT and NC, even further.

There is a clear difference between the accuracy of selections for Dwell-Visual compared to all other conditions (see Figures 4 & 5). Examining the selection points (pressure value where selection occurred), whereas almost all selection points fall within the target boundaries for this condition, the majority of all misses in the other conditions occur within a relatively short 'distance' of the lower target boundaries. We went back to the pressure behaviour to try to determine why this was, why there were so few errors where users have 'overshot' the target (other than the last menu item). The factors influencing the *Quick Release* selection will be discussed shortly, but for the *Dwell* trials it seems as though participants simply did not press enough to get to the target. They would take too long to press enough and accidentally remain in a non-target item for the 1-second timer. This hesitancy or lack of speed could come from lack of familiarity of the menu, as they are not sure where to go. However, the errors are spread across all blocks, so they would be expected to have a firmer knowledge of item positioning. It may be, however, that, as in the pilot study, there remained insufficient information in the audio feedback to properly facilitate accurate positioning. If this was the case, however, we might expect more errors above the items as well. Another factor could be explained by the finding by [8] that, when gripping objects, we apply a 'just enough' grip-force strategy where we apply only enough force (or a "small safety margin") so as to avoid slippage (due to object weight and friction), therefore not risking damaging the object or unnecessarily over-exerting ourselves. This 'least necessary effort' could account for the low levels of pressure, as it seems we may have a natural tendency to err on the side of less force.

4.1.1 Quick-Release Performance

Performance using the *Quick Release* mechanism was surprising and disappointing. Although *Quick Release* has been found to be generally more error-prone [2, 9, 10], performance here, particularly when using audio feedback, was noticeably worse (see Figures 5 and 7). The same possible factors outlined in relation to the *Dwell* technique are also relevant to the *Quick Release* trials. However, looking at the pressure behaviour, one of the primary contributing factors appears to be the *Quick* selection mechanism itself.

Designing an accurate *Quick Release* mechanism is troublesome because it is difficult to identify a common and clear pattern of sensor behaviour from which user intent can be unambiguously retrieved. Because the sensor sampled at 52Hz it was almost unavoidable that samples would be taken between lift-off and a 0 reading. The selection method used in the main evaluation used a simple algorithm comparing where and when samples were taken to decide on the lift-off point. Looking at the pressure behaviour of participants it became clear, however, that the algorithm might not always look far enough back along the sensor value timeline. Occasionally, therefore, even if a participant lifted within the target, the algorithm would take a sample outside of that target (on the way back to 0 pressure) instead and that would be taken as the selected item. To evaluate the effect of this problem we went back to the experimental data and improved the way the algorithm looked for the selection value. The pressure behaviour therefore remained the same, we simply used a different method to identify lift-off. Plotting the "corrected" *Quick Release* selection values (Figure 7) highlights what would have been a marked change in accuracy, should this variation of the mechanism have been used for the main evaluation.

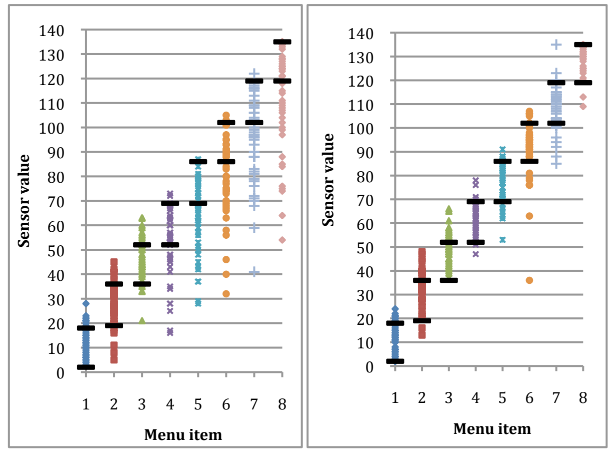

Figure 7. Right: example "corrected" selection distribution for Quick-Release-Audio condition compared to original selection distribution (left).

The "corrected" selections were much more accurate, decreasing ER rates by up to 50% (of original ER). Given this apparent improvement in the robustness of the *Quick Release* mechanism, we will carry out a further study to test its true validity. The drawback of this mechanism has always been that it is more error prone than *Dwell*, while retaining the benefit of speed. If refinement of the *Quick Release* mechanism can bring down the error rate as hinted at here, it could become the ideal mechanism.

4.1.2 The Importance of the Sensor

As was mentioned in the evaluation section, the choice of sensor for pressure interaction is highly important. The pilot study here suffered due to non-linear output from the pressure-sensitive screen of the Nokia N810, as other studies have suffered from non-linear sensors [9-11]. Stewart *et al.* [16] proposed the benefits of a properly linearised sensor, and the results here demonstrate the accuracy of this claim. Accuracy rates using the linear sensor were better than in our pilot study, and were better than results from other work using non-linear sensors. The previous section outlined the potential benefits that a better *Quick Release* mechanism would bring to this interaction, and with the majority of all errors in this evaluation came from the *Quick Release* trials, these accuracy rates may improve even further.

The task time for this interaction is still above that found in many other similar target selection tasks, and is something that requires investigating in order to improve. As mentioned before, altering the size of the sensor pressure space, so as to make full use of the resolution of the sensor, may help this.

4.1.3 Comparing Pressure-Spaces

A final note worth discussing concerns the amounts of pressure used within HCI interactions. Much of the work in psychology and psychophysics research considers performance in relation to the number of Newtons involved in the task, either in absolute terms or in relation to the individuals Maximum Voluntary Contraction (MVC; the maximum amount of force the individual can apply). Accuracy at orienting to and maintaining levels of force are affected by the degree of the force itself, both in isolation [5], as well as in relation to several other factors: the individual's characteristics [8,19]; the number of digits used in the interaction [27] and the feedback provided [5, 18, 19]. Not all studies in HCI have reported the amounts of force involved in their interaction or even the maximum amount of pressure involved in the task. Many have simply focused on the sensor values and so related performance to numbers of sensor values or abstract "low"/"high" pressure. Several studies have mentioned the maximum detectable pressure from the chosen sensor but they have then applied mappings, functions or treatments of the sensor signal leaving it unclear if this maximum remains and how the interaction is distributed over the pressure-space. Those studies that have reported maximum Newton levels range from approx 0.9N to 4N. As mentioned above, given the influence the amount of pressure involved can have on control, it is difficult to know how the results from one study may relate to others. Without referring back to the levels of pressure involved in their own study, it becomes difficult to formulate a map of exactly how different levels of pressure affect performance in an HCI setting and so how to best design interactions. This study used 3.5N of pressure space (from 0 to 3.5N) from a linearised pressure sensor, with the 3.5N divided into 4, 6, 8 or 10 equal sized divisions (of approx 0.87N, 0.58N, 0.44N and 0.35N respectively).

In conclusion, this paper presented two studies investigating the use of pressure interaction on mobile devices while stationary using both visual and audio feedback. We asked users to accurately acquire pressure levels of varying size and position through a menu-based target selection task. The results showed that using a *Quick Release* selection mechanism provided faster interaction times, but at the cost of more errors compared to a *Dwell* selection technique. Visual feedback provided better accuracy than audio feedback, although accuracy under audio feedback was still encouraging. All performance was best when using a linearised pressure sensor compared to non-linear sensor mappings.

5. REFERENCES

[1] Scott, J., Brown, L. M. and Molloy, M. Mobile Device Interaction with Force Sensing. *IEEE Pervasive '09*, 1-18.

[2] Brewster, S. and Hughes, M. Pressure-Based Text Entry for Mobile Devices. *MobileHCI 2009*.

[3] Shi, K., Subramanian, S. and Irani, P. PressureMove: Pressure Input with Mouse Movement. *IFIP INTERACT* 2009.

[4] Dietz, P., Eidelson, B., Westhues, J. and Bathiche, S. A practical pressure sensitive computer keyboard. *Proceedings of ACM UIST 2009*.

[5] Srinivasan, M. and Chen, J. Human performance in controlling normal forces of contact with rigid objects. *Winter Annual Meeting of ASME 1993*.

[6] Pang, X., Tan, H. and Durlach, N. Manual discrimination of force using active finger motion. *Perception & Psychophysics, 49(6), 531-540*.

[7] Tan, H., Pang, X. and Durlach, N. Manual resolution of length, force, and compliance. *Advances in Robotics*, 42.

[8] Johansson, R. and Westling, G. Roles of glabrous skin receptors and sensorimotor memory in automatic control of precision grip when lifting rougher or more slippery objects. *Experimental Brain Research*, 56, 550-564.

[9] Ramos, G., Boulos, M. and Balakrishnan, R. Pressure widgets. *Proceedings of CHI 2004*, Vienna, Austria.

[10] Cechanowicz, J., Irani, P. and Subramanian, S. Augmenting the mouse with pressure sensitive input. *Proceedings of CHI 2007,* San Jose, California.

[11] Shi, K., Irani, P., Gustafson, S. and Subramanian, S. PressureFish: a method to improve control of discrete pressure-based input. *Proceedings of CHI 2008,*Florence,Italy.

[12] Rekimoto, J. and Schwesig, C. PreSenseII: bi-directional touch and pressure sensing interactions with tactile feedback. *Proceedings of CHI 2006*, Montreal, Quebec.

[13] Ramos, G. and Balakrishnan, R. Zliding: fluid zooming and sliding for high precision parameter manipulation. *Proceedings of ACM UIST 2005*.

[14] McCallum, D., Mak, E., Irani, P. and Subramanian, S. PressureText: pressure input for mobile phone text entry. *Proceedings of CHI 2009*, Boston, MA.

[15] Mizobuchi, S., Terasaki, S., Keski-Jaskari, T., Nousiainen, J., Ryynanen, M. and Silfverberg, M. Making an impression: force-controlled pen input for handheld devices. *Proceedings of CHI 2005*, Portland, Oregon.

[16] Stewart, C., Rohs, M., Kratz, S. and Essl, G. Characteristics of Pressure-Based Input for Mobile Devices. *Proceedings of CHI 2010,* Atlanta, Georgia.

[17] Clarkson, E. C., Patel, S. N., Jeffrey, S. P. and Abowd, G. D. Exploring Continuous Pressure Input for Mobile Phones. *Proceedings of ACM UIST 2005*.

[18] Tang, H., Beebe, D. and Kramer, A. Comparison of tactile and visual feedback for a multi-state inputmechanism. *Engineering in Medicine and Biology Society 1997*.

[19] Tang, H., Beebe, D. J. and Kramer, A. F. A multilevel input system with force-sensitive elements. *International Journal of Human-Computer Studies*, 54, 2001, 495-507.

[20] Schmidt, R. and Bjork, R. New conceptualizations of practice: Common principles in three paradigms suggest new concepts for training. *Psychological Science*, 3, (4), 207-217.

[21] Brewster, S. and Brown, L. Tactons - Structured Tactile Messages for Non-Visual Information Display *Australasian User Interface Conference,* 2004, 15-23.

[22] Brown, L., Brewster, S. A. and Purchase, H. C. A First Investigation into the Effectiveness of Tactons. *IEEE WorldHaptics*, 2005.

[23] Brown, L., Brewster, S. and Purchase, H. Multidimensional Tactons for Non-Visual Information Presentation in Mobile Devices. *MobileHCI 2006*.

[24] Marentakis, G. and Brewster, S. Effects of Feedback, Mobility and Index of Difficulty on Deictic Spatial Audio Target Acquisition in the Horizontal Plane. *Proceedings of CHI 2006,* Montreal, Quebec.

[25] Brewster, S., Lumsden, J., Bell, M., Hall, M. and Tasker, S. Multimodal 'Eyes-Free' Interaction Techniques for Wearable Devices. *Proceedings of CHI 2003*, Ft. Lauderdale, Florida.

[26] Sawhney, N. and Schmandt, C. Nomadic radio: speech and audio interaction for contextual messaging in nomadic environments. *ACM transactions on computer-human interaction*, 7, 3, 353-383.

[27] Newell, K. and McDonald, P. V. Information, coordination modes and control in a prehensile force task. *Human Movement Science*, 13, 3-4, 375-391.

[28] Fitts, P.M. The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology*, 47 (6), 381-391.