Multimodal feedback for the acquisition of small targets

ANDY COCKBURN and STEPHEN BREWSTER[†]

Human-Computer Interaction Lab, Department of Computer Science, University of Canterbury, Christchurch, New Zealand
tel:+64 3 364 2987
fax: +64 3 364 2569
email: andy@cosc.canterbury.ac.nz

[†]*Glasgow Interactive Systems Group, Department of Computing Science, University of Glasgow, Glasgow, Scotland*
tel: +44 141 330 4966
fax: +44 141 330 4913
email: stephen@dcs.gla.ac.uk

This paper examines how multimodal feedback assists small-target acquisition in graphical user interfaces. All combinations of three feedback modes are analysed: non-speech audio, tactile, and pseudo-haptic 'sticky' feedback. The tactile conditions used stimulation through vibration (rather than force-feedback), and the sticky conditions were implemented by dynamically reconfiguring mouse control-display gain as the cursor entered the target. Results show that for small, discretely located targets all feedback modes reduce targeting times, with stickiness providing substantial improvements. Furthermore, stickiness and tactile appear to combine well. However, the results of a more ecologically-oriented menu-selection task show the need for caution, revealing that excessive feedback can damage interaction though 'noise' that interferes with the acquisition of neighbouring targets.

KEYWORDS: Target acquisition, Fitts' Law, multimodal feedback, tactile, non-speech audio, sticky targets.

Multimodal feedback for the acquisition of small targets

ANDY COCKBURN and STEPHEN BREWSTER[†]
Human-Computer Interaction Lab, Department of Computer Science, University of Canterbury, Christchurch, New Zealand
email: andy@cosc.canterbury.ac.nz

[†]*Glasgow Interactive Systems Group, Department of Computing Science, University of Glasgow, Glasgow, Scotland*
email: stephen@dcs.gla.ac.uk

This paper examines how multimodal feedback assists small-target acquisition in graphical user interfaces. All combinations of three feedback modes are analysed: non-speech audio, tactile, and pseudo-haptic 'sticky' feedback. The tactile conditions used stimulation through vibration (rather than force-feedback), and the sticky conditions were implemented by dynamically reconfiguring mouse control-display gain as the cursor entered the target. Results show that for small, discretely located targets all feedback modes reduce targeting times, with stickiness providing substantial improvements. Furthermore, stickiness and tactile appear to combine well. However, the results of a more ecologically-oriented menu-selection task show the need for caution, revealing that excessive feedback can damage interaction though 'noise' that interferes with the acquisition of neighbouring targets.

KEYWORDS: Target acquisition, Fitts' Law, multimodal feedback, tactile, non-speech audio, sticky targets.

1 Introduction

Most graphical user interfaces are heavily dependent on mouse-driven input. Johnson, Hewes, Dropkin and Rempel (1993), for example, showed that up to 65% of computer operator time is spent moving the mouse. As the resolution of computer displays increases, and as the number of components within interfaces increases, the size of many interface items is decreasing. Window borders, margin markers and split pane handles, and many more controls, are all under ten-pixels on one or both dimensions, and when rendered on a modern, high-resolution display they are less than two millimetres wide or high. Acquiring these tiny targets demands a high level precision and dexterity from users in favourable conditions, and can be extremely frustrating in the presence of adverse conditions such as vibration on a train or plane, when bright light reduces screen clarity, or when using a laptop trackpoint or touchpad. Similarly, anyone with visual or motor impairments will find targeting small items difficult.

Often, the only feedback that the user receives in target selection is the visual correspondence between the location of the cursor and the target. Sometimes additional 'over-target' feedback is presented by changing the cursor's representation (for example, an arrow-cursor is often used over the window border). This visual feedback stimulates only the (already heavily loaded) human visual system, leaving the powerful human senses of touch and hearing redundant. Multimodal feedback, stimulating several senses, has been suggested as a way to improve interaction and reduce the load on any one sense (Bolt 1980; Oviatt, DeAngeli and Kuhn 1997; Oviatt 2002).

We are investigating how feedback to different sensory modalities combines to aid target acquisition. In particular, we are examining how audio, tactile and pseudo-haptic 'sticky' components combine in both abstract Fitts' Law target acquisition tasks and in ecologically-oriented menu selection tasks. Our aim is to be able to advise designers on how they should best use the different types of feedback that are available.

The following section describes related work on Fitts' Law models of target acquisition, previous work on auditory and haptic targeting, and on multimodal target acquisition. We then describe our two experiments, and present and discuss the results.

2 Related Work

When interacting in the real world we receive information on multiple sensory channels simultaneously—we see, hear and feel the objects we work with (the senses of smell and taste are currently harder to work with and are not considered here). Multimodal interfaces employ a range of perceptual and expressive channels in facilitating the communication between humans and computers. They have been actively researched since Bolt demonstrated his "put that there" system, which used parallel verbal and pointing controls for object manipulation (Bolt 1980).

Multimodal research can be broadly categorised into input and output. Input researchers tend to examine how humans naturally use multiple expressive channels, and how sensing devices and interfaces can interpret and exploit them. Examples of multimodal input research include Oviatt's studies of synchronised speech and gesture control (Oviatt et al. 1997; Vitense, Jacko and Emery

2003) and the resultant distillation of guidelines for multimodal input (Oviatt 1999); further multimodal design guidelines were recently published by (Reeves, Lai, Larson, Oviatt, Balaji, Buisine, Collings, Cohen, Kraal, Martin, McTear, Raman, Stanney, Su and Wang 2004).

Multimodal output, the subject of this paper, investigates how presenting information to different sensory modalities can be used to enhance interaction with computers. Typical modalities include visual, audible, and tactile stimuli. Research contributions range from analysis of human perceptual limits (such as Hale and Stanney's study of the physiological, psychophysical, and neurological foundations of haptic rendering (Hale and Stanney 2004)) through to empirical observation of the impact of different modalities on human performance.

The remainder of this section describes Fitts' Law (the standard theoretical tool for analysis of human target acquisition) and presents prior work on multimodal feedback in support of target acquisition.

2.1 FITTS' LAW

Fitts' Law (1954) accurately models human psycho-motor performance in rapid aimed pointing tasks. In Human-Computer Interaction research, it is commonly used to compare the effectiveness of cursor-movement using the mouse and other input devices. Several variants of the original Fitts' Law model have been proposed, with all showing that movement time *MT* increases linearly with the Index of Difficulty (ID), see Equation 1. The index of difficulty relies on the logarithm of the distance moved (the amplitude *A*), over the width of the target *W*, and *W* is normally measured using the smallest value of the width and height dimensions (MacKenzie and Buxton 1992). ID is measured in 'bits' to reflect the information content necessary to move the cursor (or limb) to the target. The two constants, *a* and *b*, are determined experimentally and depend on cognition and motor preparation time, and on hand-eye coordination respectively.

$$MT = a + b \times ID$$ **Equation 1.**

The "Shannon formulation" of the index of difficulty is the de-facto standard model used in HCI research (MacKenzie 1992)—see Equation 2. Other forms of ID include Fitts' original formulation $ID=\log_2(2A/W)$ and Card, English and Burr's (1978) $ID= \log_2(A/W+0.5)$. In a recent retrospective of 27 years of Fitts' Law research in HCI, Soukoreff and Mackenzie (2004) state that the Shannon formulation is preferred because it provides a better fit with observations, it cannot produce negative ID values for large close targets, and the resultant regression models are less prone to yielding negative intercepts (implying negative time to select low ID targets).

$$ID = \log_2\left(\frac{A}{W}+1\right).$$ **Equation 2.**

The theoretical foundations and precise formulations of Fitts' Law models are still hotly debated, with several researchers maintaining that Fitts' Law should not be applied to targeting tasks that are completed within approximately 200msec (or when the ID is less than 3 bits). The argument is that these short 'open-loop' or 'ballistic' targeting tasks are not dependent on visual feedback, while higher ID tasks are dependent on 'closed-loop' dynamic interpretation of visual feedback as the cursor is moved to the meet the target (Gann and Hoffmann 1998). Empirical observations almost invariably show a 'flattening' of the linear relationship between movement time and ID at low ID values, but HCI researchers contend that this flattening is best accommodated through an 'adjustment for accuracy' which accounts for the typically higher error rates at higher IDs (Crossman 1957). Crossman's adjustment involves re-calculating the effective target width based on the actual distribution of clicks around the real target: 4.133×σ, where σ is the standard deviation of the distance clicks occur from the target centre. Through this technique low ID targets slightly increase in effective ID value because they have a relatively tight spread around the target, while high ID targets are adjusted to lower values due to their wider distribution. Full discussion of these issues is beyond the scope of this paper, and we refer interested readers to (Soukoreff and MacKenzie 2004).

Fitts' Law also provides a measure of human processing of movement tasks, called the 'Index of Performance' (IP) or 'bandwidth' (measured in 'bits/second'). Bandwidth is useful for comparing the effectiveness of different input devices. It can be measured in two ways (MacKenzie 1992). The more common method is to calculate IP from the reciprocal of the slope constant b, which is determined from Fitts' Law regression analysis (Equation 3). It can also be calculated on a per-task basis by dividing movement time by the index of difficulty (Equation 4). The two calculation techniques produce different values, with the discrepancy increasing with the absolute value of the linear intercept a. Both calculation techniques are used in this paper.

$$IP = \frac{1}{b}$$ **Equation 3.**

$$IP = \frac{MT}{ID}$$ **Equation 4.**

2.2 NON-SPEECH AUDIO INTERFACES AND TARGET ACQUISITION

Many researchers have investigated audio enhancement of graphical user interfaces. Many different interface widgets, such as menus, buttons, scroll-bars, progress bars and tool palettes, have been augmented with non-speech audio earcons reducing error rates, task completion times and subjective workload (Beaudouin-Lafon and Conversy 1996; Brewster 1998a). Brewster and Crease (1997) looked at the use of sound to reduce the incidence of errors in drop-down menu selections, in which users 'slip off' the desired menu item, accidentally selecting the one above or below. This happens partly because the action of releasing the mouse button can move the mouse a little and also because users often start to move the mouse to the location of the next action before the mouse button is released on the current action. Results showed that non-speech audio allowed errors to be corrected significantly faster and with a reduction in subjective workload. In a study on sonically-enhanced drag and drop for discretely positioned targets Brewster (1998b) found that sound significantly reduced subjective workload (using NASA-TLX workload measurements (Hart and Staveland 1988)), error rates and task times. In particular, the target highlight time (the time the cursor was over the target before the drop was made) was reduced by 18% in both simple (one target) and complex (multi-target) interfaces.

Jacko and colleagues have conducted a series of experiments scrutinising the effectiveness of multimedia (including audio) feedback in aiding visually impaired (and non-impaired) users (Jacko, Scott, Sainfort, Barnard, Edwards, Emery, Kongnakorn, Moloney and Zorich 2003; Vitense et al. 2003; Jacko, Barnard, Kongnakorn, Moloney, Edwards, Emery and Sainfort 2004). In drag-and-drop tasks audio improved the performance of both visually impaired and sighted users.

The study most closely related to ours is that of Akamatsu, MacKenzie and Hasbrouc (1995) who conducted a Fitts' Law analysis of abstract target acquisition in the presence of several feedback cues, including audio. In the audio condition, a simple 2KHz tone was played while the cursor was over the target. They found that audio feedback made no significant difference to the overall targeting time, but that it did reduce the time spent over the target (the time between the cursor entering the target and the mouse button being pressed), similar to Brewster's scrollbars above.

2.3 HAPTIC INTERFACES AND TARGET ACQUISITION

The use of haptic or touch-based interfaces has become an important area of research over recent years (Wall, Riedel, Crossan and McGee 2002). The human sense of touch can be roughly split in to two parts: kinaesthetic and cutaneous. "Kinaesthetic" is often used as catch-all term to describe the information arising from forces and positions sensed by the muscles and joints. Force-feedback haptic devices (such as the PHANToM from SensAble, www.sensable.com) are used to present information to the kinaesthetic sense. Cutaneous perception refers to the mechanoreceptors contained within the skin, and includes the sensations of vibration, temperature, pain and indentation. Tactile devices are used to present feedback to the cutaneous sense.

There has been some previous work into the use of tactile displays at the desktop user interface, and how tactile cues such as Tactons (Brewster and Brown 2004) might be designed. Akamatsu, MacKenzie and Hasbrouc (1995) investigated the impact of passive tactile feedback (a vibrating mouse) on single target acquisition. Like audio feedback, they found that tactile feedback did not significantly influence the overall targeting time, but that it did reduce the time over target. Tactile feedback reduced over-target time by more than audio or visual feedback. However, the authors did not consider multiple targets and the distractions that tactile feedback might cause in more complex, real-world displays.

In a subsequent study, Akamatsu and MacKenzie (1996) further analysed target acquisition with tactile feedback (a solenoid-driven pin that stimulated the user’s index finger) and with force feedback (electromagnetic induced drag in the mouse). Unlike their previous study, they found that the feedback modes reliably reduced overall targeting time and error rates. Surprisingly, although force-feedback slightly increased overall acquisition time (compared to only visual feedback), the combination of tactile and force produced the lowest mean acquisition times.

Using a PHANToM force-feedback device, Oakley, McGee, Brewster & Gray (2000) investigated both abstract targeting tasks (using a range of different force-feedback effects such as gravity well, texture, friction and recess overlaid on graphical on-screen buttons) and more ecologically oriented tasks involving the use of a haptically enhanced scrollbar. They found that the force-feedback device did not reliably reduce task completion time, but that it did significantly reduce the number of errors made and subjective workload experienced. They found gravity wells and recess effects much more effective than friction or texture due to the latter two causing perturbations to the users’ desired movements which made small targets harder to hit.

Langdon, Keates, Clarkson and Robinson (2000) found that mouse-generated force-feedback (using the Logitech Wingman force-feedback mouse, www.logitech.com) dramatically improved the targeting performance of motion-impaired users, and that the benefits increased with the severity of the user’s impairment. Dennerlein, Martin and Hasser (2000) also studied mouse-generated force-feedback, but for non-impaired users, and found that performance in certain mouse movement tasks was dramatically improved by force-feedback. Their ‘steering’ tasks involved moving the cursor through a constrained region, simulating the types of movement necessary to, for example, select a cascading menu item where the cursor must move within a constrained vertical region. They also examined a combined steering and targeting task, and although the force-feedback condition again provided reliable performance improvements, the effect was less dramatic. In another study focusing on targeting rather than steering, Dennerlein and Yang (2001) found that attractive force-fields that surround targets improved targeting by approximately 25%.

Although promising, these studies often focused on abstract, single target tasks that ignored issues such as the distraction of forces created by neighbouring targets. Oakley et al. (2000) looked at the addition of force-feedback effects to menus where targets are densely stacked vertically. Their results showed that the naïve application of force could significantly reduce user performance. They compared standard visual menus to those with both fixed and dynamically adjusted forces (based on gravity wells). The dynamic condition lowered force with speed of movement and direction (if users are moving rapidly over an item they are unlikely to be targeting it and so do not need force-feedback). The fixed forces caused much slower performance as users were dragged on to all of the menu items as they moved through a menu. Dynamically adjusted forces significantly reduced task times and subjective workload, as forces were only applied where appropriate. Similar results were found for tool palettes and desktops (Oakley, Adams, Brewster and Gray 2002).

Hwang, Keates, Langdon and Clarkson (2003) examined the usability of force-feedback gravity wells in the presence of distracter targets that were placed at various positions with respect to the actual target. Their comprehensive study also examined performance by motion-impaired and able-bodied users. They found that force-feedback aided targeting, even in cases when the cursor entered a distracter target. The positive effect was particularly strong for motion-impaired users.

There has been little research into the use of tactile feedback in more complex and realistic, multiple target displays. There may be similar problems to those reported in the force-feedback research above, as vibrations produced when moving over distracters may cause problems. This needs to be investigated further so that guidance can be offered to interface designers on how to use effective tactile displays in their systems.

2.4 STICKINESS AND TARGET ACQUISITION

Sticky widgets attempt to aid target acquisition by using a pseudo-haptic metaphor based on gravity, magnetism, or stickiness. Worden et al. (1997) implemented ‘Sticky Icons’ by decreasing the mouse control-display gain (the mapping between the physical mouse movement and the resultant cursor movement) when the cursor enters the icon. In this way, the user must move the mouse further to escape the boundary of the icon, effectively making the icon larger in motor-space without using extra screen space. Worden et al.'s evaluation showed sticky icons to be efficient for selecting small targets, particularly for older users.

Langdon et al. (2000) performed an evaluation of a similar 'force feedback' concept, which warped the user's pointer towards targets. Although this condition was 30% to 50% faster than the normal condition, the technique is of limited utility because of the undesirable impact on selecting near-neighbour items. A scrollbar, for example, would be difficult to use if the pointer continually warped toward the window border.

Keyson {, 1997 #1209} examined target acquisition across four feedback conditions: visual, visual+tactual (force feedback 'trackball'), visual+control-display gain 'stickyness', and visual+tactual+sticky. Although their study did not isolate the impact of tactile feedback, results indicate that it had a much larger positive impact on targeting than control-display gain.

In our previous work, we compared three schemes that aim to aid target acquisition, including sticky widgets (Cockburn and Firth 2003). The two other techniques were bubble targets, which expand when the cursor is close, and goal-crossing targets, which are acquired by sweeping the cursor over the item while holding down a modifier key or mouse button. Results showed that sticky components were both popular and efficient, allowing targets to be selected 28% faster (on average) than normal. More recently, Blanch, Guiard and Beaudouin-Lafon (2004) completed a formal analysis of control-display gain adaptation for target acquisition, including a rigorous Fitts' Law analysis. Their results confirmed that the technique can improve target acquisition—up to 17%, on average, using targets 32 pixels wide or larger.

2.5 TARGETING WITH MODALITIES IN COMBINATION

Several research projects have investigated how feedback modalities combine. Vitense, Jacko and Emery (2003) investigated how all combinations of visual, audio and haptic feedback influenced user performance in drag-and-drop tasks. The feedback modes were activated when the dragged object was over the target. Haptic feedback was provided by vibrating a Logitech Wingman force feedback mouse, and audio feedback consisted of a musical tone. Dependent measures included the total trial time (drag and drop), the highlight time (between entering the target and dropping the object), and subjective responses using NASA-TLX workload measurements. There were several interesting results, including the fact that haptic feedback increased the total trial time but reduced the target highlight time, and that audio increased the highlight time, the opposite result to Brewster's audio drag-and-drop (1998b). Overall their results showed that some combinations of feedback were very effective, but other combinations were not and actually reduced performance. The combination of all three feedback types was not successful, again reducing performance.

As mentioned above, the two studies most closely related to ours are those of Akamatsu, MacKenzie and Hasbrouc (1995) and Akamatsu and MacKenzie (1996). The first study examined whether targeting was aided when sound, tactile-vibration, and colour feedback were used to indicate the mouse-over condition. They compared five conditions: ‘normal’, the three feedback modes in isolation, and a ‘combined’ condition in which all feedback modes were used simultaneously (they did not investigate the complete pair-wise use of the modalities to fully explore the design space). Although they found no significant difference in overall selection time, their results showed that target highlight time was reduced when feedback was present. Tactile feedback appeared to have a greater effect in reducing highlight time than either sound or colour feedback.

In the second study, Akamatsu and MacKenzie (1996) examined the contribution of tactile and force feedback in targeting. They found a significant difference between normal, tactile, force, and force+tactile feedback conditions. Tactile and force+tactile feedback reduced targeting times by 5.6% and 7.6% respectively. Force feedback alone resulted in slightly higher targeting times.

Neither study investigated all possible combinations of feedback modality to understand the interactions between them. The aim of our research is to use the best design of stimuli from work in each of the individual modalities, combine them and then look at how they perform in combination for abstract single target (Experiment One) and more realistic multiple target (Experiment Two) acquisition tasks.

3 Experiment One—Abstract Targets

Our two studies investigate how audio, tactile, and pseudo-haptic (stickiness) feedback combine to aid or hinder both abstract and more ecologically oriented targeting tasks. We introduced pseudo-haptic stickiness as a condition because of its promising results in previous work (Cockburn and Firth 2003), and because actual force-feedback pointing devices (such as the PHANToM) are not in widespread use. In our preliminary experiments, we found that the combination of stickiness and tactile feedback provided a sensation somewhat similar to that of a force-feedback gravity well (Section 2.3). We wished to see if this perception would yield measurable performance differences.

The first experiment investigates how the various feedback modalities combine when used to acquire targets in a simple, one-dimensional movement task. As well as calculating how Fitts’

Law models the tasks we also analyse several metrics that characterise the participants' response to the modalities. The second experiment investigates the use of our feedback modalities in a complex task to understand how they might perform in more realistic situations. The same participants took part in Experiments One and Two, with the participants proceeding to Experiment Two immediately on completing all of the tasks in Experiment One.

3.1 METHOD

All of the interfaces were visually indistinguishable from one another (see Figure 1). The tasks involved clicking on two thin vertical bars in sequence. As soon as one bar was clicked the other bar would move to a new location and be highlighted in green. The participant would then acquire it (clicking on it to complete the task) as quickly as possible. The target bar was a constant width of 8 pixels for all tasks (a constant width was used because of our interest in methods for enhancing the acquisition of small targets).

[Figure 1 around here.]

3.1.1 PARTICIPANTS

The twenty participants (15 male, 5 female) were all graduate and undergraduate Computer Science students from the University of Canterbury and between 20 and 47 years of age (mean 26). All used the mouse in their right hand (by choice). Experiments One and Two lasted approximately thirty minutes in total. Participation was rewarded with a $5 shopping voucher.

3.1.2 APPARATUS

The experiment was run on a Toshiba Tecra 8000 laptop, running Windows 2000, with a 15" display running at 1024×768 resolution. All mouse input was provided by an off-the-shelf Logitech iFeel tactile Mouse (see Figure 2) that was placed on a rubberised mouse mat (the surface on which the iFeel Mouse rests can influence the tactile sensation, and Logitech recommend a soft mouse surface). Mouse acceleration was disabled throughout the experiment, and the mouse speed was set to a control-display gain of 1:1.

[Figure 2 around here.]

The experimental conditions were controlled by a Python program that generated the interface, cued various questionnaire dialogue boxes, and logged all user actions, with each task time being the period between the cursor leaving one bar and clicking on the next (green) target one. In addition, Windows Media Player was used to play a looping audio stream (waves crashing on a beach) to the participants through headphones. This low-level, broadband noise was used to block out sound generated by the motor in the iFeel mouse. The audio feedback was mixed in with this noise at a higher volume and the headphones were worn in all conditions.

In the normal (control) condition the user received no feedback other than the visual correspondence between the location of the cursor and the underlying target. On successfully selecting the green target bar, it was coloured grey, and the other bar would move to a new location and be highlighted green. Missing the target (clicking off the highlighted line) caused no visible change in the interface, and the participants continued as normal until a bar was correctly selected.

In the audio condition simple earcons were used for the feedback, based on previous work on audio widgets (Brewster 1998b). The first earcon was played when the participant moved the mouse over the target. This was a quiet, continuous reed organ sound at pitch C4 (130Hz) (the sound stopped if the user moved off the target, or selected it within 500 ms.). The second earcon was used to indicate that the user had successfully selected the target. The sound was at pitch C4 and was played for 300msec. with a bell timbre. There was no audio feedback if the user missed the target.

A simple Tacton was used for the tactile feedback. This was produced by vibrating the mouse (full force at 200 Hz) while the cursor was over the target. This frequency is in the range where the skin is most sensitive. Feedback continued until the user moved off the target or made a selection, when it stopped.

In the sticky condition the mouse control-display gain was reduced to one-twentieth of its original value (1:0.05) while the cursor was over the target. This was implemented by addressing sub-pixel cursor coordinates, then warping the cursor to the rounded pixel locations. Other than configuring the effective mouse control-display gain, stickiness caused no visual changes to the display.

3.1.3 EXPERIMENTAL DESIGN

For the Fitts' Law analysis, linear regressions between the measured movement time and the index of difficulty (controlled and adjusted for accuracy) were calculated for each combination of feedback of modalities.

The data from several dependent measures were also analysed in a 7×8 repeated measures analysis of variance (ANOVA) for factors 'distance' and 'feedback modality'. The seven levels of distance were 8, 16, 32, 64, 128, 256 and 490 pixels from the target centre. The eight levels of modality consisted of all combinations of stickiness, tactile and audio (from normal, in which all modalities were off, through to all three modalities being on). Although the main dependent measure is the time to select targets (this is the value that most concerns users), several other measures are also analysed to scrutinize user performance with the feedback modalities (reported below).

3.1.4 PROCEDURE

Having answered a preliminary questionnaire on background demographics—handedness, age, gender, and so on—the participants were shown the experimental interface. They were told to click on the green line as quickly as possible, in blocks of 40 selections. They were encouraged to rest between blocks. The inter-block gap was identified by the software which presented a dialogue box asking participants to respond to the question “These settings *(stickiness, tactile, audio* as appropriate*)* helped me to rapidly select targets” using a five-point Likert scale (1-disagree, 5-agree). Prior to presenting this question the participants had not been told which combination of modalities they were using.

Data from the first five selections in each batch were discarded as training tasks, leaving 35 logged selections (five at each of the seven distances) for each modality combination. The order of exposure to each of the eight combinations of feedback modalities, and to each of the distances, was randomised for each participant.

3.2 RESULTS

3.2.1 FITTS' LAW MODELS

As expected, Fitts’ Law accurately modelled the participants’ performance with all combinations of feedback. Figure 3 shows the relationship between mean movement time and Index of Difficulty (with ID calculated using Equation 2). The movement-time flattening previously reported for low ID tasks is visible in the figure. To accommodate for this flattening, Table 1 shows two forms of Fitts’ Law analyses: the left main-column shows the lines of best fit, R^2 values, and the IP measurements for regression analysis based solely on ID > 3; and the right main-column shows the same measures for all IDs with target widths adjusted through Crossman’s (1957) ‘adjustment for accuracy’ of $4.133\times\sigma$ (see Section 2.1). IP values are calculated using Equation 3.

The IP values in both models show a marked increase in throughput when stickiness is present, and that audio feedback has a less dramatic impact. The effectiveness of tactile feedback appears questionable as it produced lower IP values than the normal condition, and also reduced throughput when combined with stickiness. These observations are further discussed below.

[Figure 3 around here.]

[Table 1 around here.]

The R^2 values of the regression analyses are all above 0.95 in the models based on ID>3.0, meaning that more than 95% of the variance in performance is explained by the Fitts' Law models.

3.2.2 TIME AND BANDWIDTH

The Fitts' Law results suggest that some combinations of modalities aided participants in target acquisition. The analysis of variance allows us to scrutinise further how the modalities influenced targeting across different distances and across different portions of the acquisition process. In this section, IP ('bandwidth') measures are calculated using Equation 4.

Across both factors (modality and distance), the mean target acquisition time was 636 ms (s.d. 251), with a mean per-task bandwidth of 5.7 bits/sec (s.d. 2.0). There was a significant main effect for feedback condition ($F_{7,133} = 32.9$, $p< 0.001$), ranging from the fastest performance in the sticky+tactile condition (mean 521ms, s.d. 203, bandwidth 6.9 bits/sec) to the slowest performance in the normal condition (mean 743ms, s.d. 296, bandwidth 4.8 bits/sec).

As Figure 4 shows, there was a relatively dramatic difference between acquisition time with and without stickiness. This is unsurprising as stickiness effectively increases the target size in motor-space. Having predicted this effect, we conducted an analysis of variance across the non-sticky conditions. This showed no significant difference between the non-sticky feedback conditions ($F_{3,57} < 1$, p=0.58).

[Figure 4 around here.]

Naturally, there was a significant main effect for distance ($F_{6,114} = 345$, $p< 0.001$). More interestingly, however, there was a reliable interaction between feedback modality and distance ($F_{42, 798} = 1.99$, $p< 0.001$), meaning that target acquisition deteriorates differently across the feedback conditions as the distance increases. This effect is also caused by stickiness, as the interaction is not reliable on removing stickiness from the analysis ($F_{18, 342} < 1$, p=0.78).

3.2.3 OVER-TARGET TIME

The feedback modalities used in the study have no impact on the user until the cursor enters the target—all of the conditions behave identically during movement towards the target. Although the user is most concerned about acquiring targets rapidly, the measure that provides the best research insights into user response to the feedback is the time spent over the target prior to selection.

To scrutinise the effects of feedback on user's performance we conducted an analysis of variance of the time gap between entering the target and pressing the mouse-button, repeating the 7×8 design used previously.

Across all conditions, the mean over-target time was 211ms, s.d. 79. There was a significant difference between the different modalities ($F_{7,133} = 2.4$, $p < 0.05$) with stickiness slowest (236 ms, s.d. 103), followed by normal (231 ms, s.d. 89), then a jump to audio+sticky (214 ms, s.d. 72), audio (211 ms, s.d. 64), audio+sticky+tactile (210 ms, s.d. 80), tactile (196 ms, s.d. 66), audio+tactile (196 ms, s.d. 69) and finally sticky+tactile (192 ms, s.d. 70). It appears that tactile and audio feedback increased the participants' confidence that they were over the target, allowing quicker selection of the target once movement was complete. In contrast, we suspect that stickiness effectively 'surprised' participants—they would move rapidly to the target, the cursor would 'snap' into it, and they would then have to visually confirm that the cursor was inside the target prior to pressing the button. When tactile or audio feedback was combined with stickiness, the participants could rely on these modes to aid confirmation that the target was successfully acquired. Stickiness and tactile seemed to provide a particularly powerful combination, providing (in our opinion) a sensation that approaches that of a 'gravity well' in force-feedback devices.

There was also a reliable main effect for distance ($F_{6,114} = 9.1$, $p < 0.001$), with short distances resulting in longer over-target times than long ones. For eight pixel movements the mean over-target time was 233ms (s.d. 83), with a rapid drop off to 201ms (s.d. 75) for movements of 64, 128 and 256 pixels. As for stickiness, this effect seems best explained by the time taken to visually confirm the over-target state. In short movements, the user barely has time to begin moving prior to entering the target, and there is therefore less opportunity to anticipate the precise timing of target entry. For longer movements, the user can prepare for the cursor's entry into the object by observing the cursor's rate of movement towards and into the target.

The relatively long time needed to confirm target acquisition in the sticky condition, particularly for short distances, caused a significant interaction between feedback condition and distance ($F_{42, 798} = 2.0$, $p < 0.001$).

Our results agree with Akamatsu and MacKenzie's (1996) study of audio and tactile feedback, in which they found that mouse-over times reduced by a 20% with tactile and by 12% with audio, compared with 15% and 9% in our study.

3.2.4 MISSES

Another important measure of performance in target acquisition is the error-rate—the proportion of clicks that occur outside the target. In Fitts' Law studies users are normally encouraged to adjust their performance to an error rate of approximately 4%, but our targets are intentionally hard to hit because of their small size. We were interested to see how the modalities impacted on the users' error rate without prompting users to adjust their performance.

The 7×8 design was reused, but with the number of off-target clicks per trial as the dependent measure. The analysis revealed interesting and surprising effects of the different feedback modalities.

Across all conditions the mean miss rate was 18% (approximately one miss in five trials). This high error rate is explained by the small target size (8 pixels); substantial dexterity is necessary when selecting such a small item, and our participants were attempting to select them rapidly. There was a significant difference between the miss rates with different feedback modalities ($F_{7,133} = 12.14$, $p < 0.001$). Although we had expected additional feedback to reduce the error rate, we were surprised to find that miss rates were higher when tactile feedback was present. Mean miss rates for the normal, tactile, and audio+tactile conditions were 22%, 30% and 32% respectively. As expected, stickiness dramatically reduced the miss rate (sticky 11%, sticky+tactile 10%, sticky+audio 11%), but stickiness was not solely responsible for the reliable difference between conditions (removing sticky conditions from the analysis gives $F_{3,57} = 6.3$, $p < 0.01$). Audio also appeared to reduce the number of misses (the mean miss rate for audio alone was 5% better than the normal condition at 17%).

There was a significant main effect for distance ($F_{6,114} = 5.7$, $p < 0.001$), with miss rates increasing with distance. There was no feedback × distance interaction.

The observation that tactile feedback increases miss rates is supported by previous work: Akamatsu and MacKenzie (1996) noticed that with their larger targets, error rates rose from 6.6% in the normal condition to 11% with tactile feedback. Their explanation was that the tactile sensation triggers a small reflexive muscle response that is sometimes sufficient to displace the mouse outside the target. This explanation would predict the larger increase in error-rate that we observed, because with our small targets the cursor must be close to the item edge. An alternative explanation could come from Oakley et al. (2000) who found that force-feedback generated textures caused users problems with targeting due to the vibrations perturbing users' motions and throwing them off target. The perturbations caused by the iFeel mouse are much less than with the PHANToM device when rendering textures, but it could be that with small targets the vibrations still cause problems with very precise targeting. One way to overcome this might be to present the tactile feedback to another part of the body (say the forearm or the other hand) to avoid moving the hand or fingers. Further experimentation is needed to find out more about this effect.

3.2.5 OVERSHOOTS

A final dependent measure that characterises the participants' performance with the various modalities is the number of times they overshot the target. Unsurprisingly, stickiness dramatically reduced overshooting—from approximately half of the non-sticky trials to below 10% of the sticky trials. Audio and tactile made no significant difference to overshooting.

3.2.6 SUBJECTIVE MEASURES

After each block of trials with a particular combination of feedback modalities, participants responded to the question "These settings (*list of the modalities used*) helped me to rapidly select targets" using a five-point Likert scale (1-disagree, 5-agree). There was a reliable difference between the participants' ratings for the various modalities (Friedman $\chi^2 = 86$, $p<0.001$), with mean ratings ranging from normal feedback (2.35, s.d. 1.0) though tactile (3.15, s.d. 1.2), audio+tactile (3.3, s.d. 1.4), audio (3.5, s.d. 1.4), to a marked increase to the sticky conditions, all of which exceeded 4.5.

The participants' post-experiment comments reinforced these subjective measures, with many stating that additional feedback dramatically aided targeting. Comments were particularly strongly in favour of stickiness.

3.2.7 DISCUSSION

The results show that using feedback in different modalities can significantly improve targeting in a simple user interface. By increasing the motor size of small targets, stickiness dramatically reduces the time taken to acquire such targets. Stickiness reduced the normal targeting time by 25%. The other feedback modalities also reduced targeting time, but by less dramatic amounts: audio and tactile feedback reduced the mean target acquisition time by 4.2% and 3.5%. The results for tactile and audio are supported by Akamatsu and MacKenzie (1996) who found that tactile reduced targeting time by 5.6%, and by Akamatsu et al. (1995) who showed that tactile and audio feedback reduced acquisition times by 11% and 1% respectively. Much of the other work on sonic enhancement of interfaces has concentrated on error reduction rather than selection time improvements (see Brewster et al. above) so direct comparison with that literature is not possible.

The experiment provided some interesting insights into which feedback modalities combine positively. Although we had expected that all feedback modalities would combine in a positive way, the results suggest that some do, while others do not. For example, although audio and tactile individually improved targeting times by 4.2% and 3.5%, the combination of audio+tactile reduced normal targeting times by only 1.7%. Similarly, although stickiness reduced targeting times by 25%, audio+sticky provided little further benefit. However, the results suggest that sticky+tactile combine positively, with targeting times reducing by an additional 5% beyond stickiness alone.

4 Experiment Two—Ecological Menu Targets

The first experiment examined user performance with combinations of sticky, audio and tactile feedback when acquiring a single small target. Importantly, the target was isolated from any surrounding 'distracter' targets. This allowed us to scrutinise idealised user performance, and reflects a small set of practical uses such as targeting window borders in uncluttered desktops. The focus of the study, however, was largely independent of ecological validity. Our second experiment addresses this and examines how the modalities combine in a more strongly ecological task, specifically menu selection.

4.1 METHOD

The participants proceeded to Experiment Two immediately after completing Experiment One, using the same apparatus, with the same Python program controlling the experimental conditions and logging data.

The interface used in Experiment Two, shown in Figure 5, was visually unaffected by the different feedback modalities. Menus were chosen for the experiment as they had already been enhanced with audio and tactile feedback in earlier studies. In many ways they are a good test platform as the widgets are densely packed, emphasising the issues caused by distracter targets. The left-hand side of the split window in Figure 5 showed the target menu item, while the right hand side showed only a menu button. The participants' tasks involved navigating through cascading menus to select the target item as quickly as possible. On selecting the correct menu item a new target menu item was displayed in the left-hand split-pane. If an incorrect menu item was selected the background of the right-hand split pane was coloured red, and the target item displayed in the left hand pane remained unchanged, requiring the user to re-navigate through the menu. Like normal menus, the menu-item under the cursor was visually highlighted, and users could select items either by dragging or by using clicks to post menus and cascades. Each menu item was twenty pixels high with a seven-pixel desensitised gap between items (where no audio/tactile/sticky feedback was presented). The desensitised area was used to allow users to perceive a gap between menu items. Selecting a gap had the same effect as clicking outside the menu (unposting the menu).

[Figure 5 around here.]

The audio and tactile feedback were unchanged from Experiment One. The level of stickiness was reduced to accommodate the larger targets. Within sticky menu items the control-display gain was attenuated to 40% of its normal value (rather than to the 5% used in Experiment One).

4.1.1 EXPERIMENTAL DESIGN

The experimental used a 3×8 repeated measures design with factors 'menu depth' and 'feedback modality'. The menu depths were either one, two or three depending on whether the target menu item was in the top-level menu or in a second or third level cascade (Figure 5 shows a third level target). The levels of factor 'feedback modality' were the same as Experiment One. The dependent measures were: selection time, selection errors, over-target time, and the same subjective questions as before.

4.1.2 PROCEDURE

Each participant made 144 menu selections in eight blocks of eighteen selections (one block per feedback modality). The first six selections in each block were treated as training tasks (two at each of the three levels of menu depth), and the data were discarded. The remaining twelve selections comprised four selections at each of the three menu depths. The same set of twelve selections was used for all participants, with randomised orders for menu-selection trial and for exposure to the eight feedback modalities.

4.2 RESULTS

4.2.1 SELECTION TIME

The mean time to select menu items across all conditions was 2.15 seconds (s.d. 1.1). Although there was a significant main effect for feedback modality ($F_{7,133} = 10.5$, $p < 0.001$) the effect was due to the poor performance of the sticky conditions—see Figure 6. Analysing the data in a 3×4 ANOVA for the non-sticky conditions shows no significant difference between feedback modalities ($F_{3,57} < 1$, p=0.57). Mean menu selection times were fastest in the tactile condition (1.9secs, s.d. 0.9), closely followed by the other non-sticky conditions. There was then a fairly marked performance drop to the sticky conditions, with mean performance times more than 15% worse than the normal condition.

[Figure 6 around here.]

As expected, there was a strongly significant main effect for menu depth ($F_{2,38} = 294$, $p < 0.001$). There was no depth×feedback interaction ($F_{14,266} = 1.2$, $p=0.3$), meaning that performance with all modalities deteriorated similarly across increasing menu depth.

The poor performance of stickiness was clear to the participants, many of whom made strong statements criticising it. In general, the lower control-display gain of the adjacent menu items failed to produce a ‘sticky’ sensation. Instead stickiness “felt like the mouse was going annoyingly slow” when moving through the menus.

4.2.2 OVER-TARGET TIME

The over-target analysis shows a significant main effect for feedback modality ($F_{7,133} = 4.7$, $p < 0.001$), with participants spending least time over the target in the normal condition (345ms, s.d. 92), slightly more with tactile feedback (361ms, s.d. 134), followed by audio+tactile (372ms, s.d. 106) and audio only (377ms, s.d. 112). Participants hesitated longest over sticky targets, with all sticky means exceeding 390ms. Audio+tactile+sticky produced the highest mean over-target time at 409ms (s.d. 134).

4.2.3 SUBJECTIVE MEASURES

Responses to the five-point Likert-scale question “These settings (*list of the modalities used*) helped me to rapidly select targets” showed the inverse preferences to Experiment One. There was a reliable difference between ratings for the modalities (Friedman $\chi^2 = 78$, $p<0.001$) with mean ratings ranging from 1.6 (s.d. 1.0) in the sticky condition, through 2.9 for the audio, tactile and audio+tactile conditions, to the normal condition at 3.7 (s.d. 1.1).

4.2.4 DISCUSSION

In Experiment One additional feedback reduced selection times and allowed users to quickly select items once the cursor was over the target. Experiment Two shows the opposite, with additional feedback slowing targeting performance and making users more hesitant once the cursor was over the target.

The reason for the discrepancy seems clear. In Experiment One, the feedback is a discrete ‘burst’ of information that is provided only when a desired state is acquired (the cursor being over the target). In Experiment Two, because several candidate targets are adjacent to one another, the user is effectively saturated in feedback from multiple items. The level of feedback becomes noise that distracts the user from the task.

One way to avoid this would be to use dynamically controlled feedback (as suggested by Oakley et al. (2001) above for force-feedback displays). They suggest a reduction in feedback along each axis of movement individually (to zero) in proportion to speed along the opposite axis. This has the effect of providing little feedback as users move rapidly over targets and more when they slow down to begin targeting, plus supporting movement to keep users on targets (e.g. to avoid slipping off menus when dragging down).

With fixed force-feedback cues Oakley found that performance slowed to below that of standard menu usage, but when dynamically adjusted feedback was present it boosted performance back to the normal level with the important effect of significantly reducing wrong target selections.

Brewster and Crease (1997) showed that menu selection errors (either slipping off a menu item on to an adjacent one when making a selection, or slipping off a menu entirely when dragging through it) could be reduced by the addition of non-speech sound, but overall time to make selections was unaffected (similar to the results for our experiment here). This suggests that selection time improvements from the addition of multimodal feedback might not be the only benefit.

Vitense *et al.* found that, in their drag-and-drop tasks, tactile feedback reduced highlight time. Here we found that it was increased. The reason for this is the type of interaction; Vitense's drag-and-drop task did not have the problem of distracter targets giving unnecessary feedback (we too found that tactile feedback reduced highlight time in our first experiment). This suggests that care should be taken to make sure that feedback suits the particular task. It may be that feedback good in one situation is poor in another due to excessive distraction.

5 General Discussion

The results of the two experiments clearly illustrate the need for careful consideration in the way designers use different feedback modalities to aid targeting. Experiment One showed that additional feedback can aid the acquisition of discretely placed targets, while Experiment Two showed that it can harm interaction by overloading the user with superfluous and distracting information. Although the finding of Experiment Two may seem obvious, there are many examples of interfaces that provide excessive feedback for targeting—for example, roll-over audio on neighbouring items often features within Flash websites, and the standard TouchWare software provided with the Logitech iFeel mouse provides independent tactile feedback for each file as the cursor moves over a filename in Windows Explorer or icons on the Windows desktop, yielding an incomprehensible vibration.

For discrete targets, our results support those of prior work, showing that audio and tactile feedback both reduce targeting time by around 4%. We also confirmed the observation that tactile feedback tends to increase the incidence of errors with small targets. Sticky targets reduced the mean selection time for the 8 pixel target by 25%, and it combined well with tactile feedback, giving a sensation approximating that of a force-feedback gravity well, without the cost of a force-feedback device. Our guideline is that designers looking to improve the selection accuracy of discrete targets should use the multimodal combination of tactile and sticky feedback. For more complex tasks we suggest that care should be taken if selection time improvements are required. The guideline here is that inappropriate use of modalities can increase selection times.

Although stickiness and tactile combined well in Experiment One, the results indicate that the three-way combination of stickiness, tactile and audio provided excessive feedback. The total acquisition time and the over-target time increased 9% and 6% over the sticky+tactile condition in the three way condition, suggesting that users were distracted by excessive feedback ‘noise’.

The feedback we used was presented redundantly; the same selection event was presented in the different modalities and perhaps this overloaded the users. The selection events we were indicating were simple and presenting these in several modalities at the same time may have been overkill. It could be more productive to use the different modalities differently. For example, one could be used to aid targeting and others could be used to reduce errors or indicate other types of information, increasing the whole bandwidth of communication.

We also used the same feedback cue each time (for simplicity and consistency between modalities). Brewster and Crease (1997) used two sounds for menu items so that moving from one item to the next caused a change in feedback. This allowed users to recognise that they had slipped off one menu item on to an adjacent one. Manipulating the cues we used in this way would allow us to communicate more, perhaps providing more useful feedback and less ‘noise’.

There are two issues that we wish to pursue in further work. First, our feedback was continuously provided while the cursor was over the target. Although continuous feedback better supports the user’s sensation of the over-target state, it does so in a relatively forceful manner, particularly in the presence of multiple candidate menu item targets in Experiment Two. Instead, discrete feedback could be used to denote attaining and leaving the over-target state, with the audio and tactile cues distinguishing between enter and leave (for example, high frequency for enter and low frequency for leave). We therefore wish to examine how discrete versus continuous multimodal feedback compare for various targeting tasks.

Another issue we wish to investigate is how far objects need to be separated for multimodal feedback to succeed. In Experiment One, the target was entirely separate from all others, and in Experiment Two there was a seven-pixel separation. Somewhere between these two extremes there must be a cross-over point at which the benefits of multimodal feedback balance with the costs of distraction. We wish to investigate where these boundaries lie for different types of targets and different modalities, and to experiment with various metrics for ‘distance’ (pixel distance and temporal ‘distance’ in which there is a delay between the user’s selections).

6 Conclusions

Mouse controlled selection and manipulation of graphical user interface components consumes a large portion of the time spent working with graphical user interfaces. Any improvement to targeting has the potential to yield substantial usability benefits.

This paper investigated how three specific methods of multimodal feedback could combine to assist targeting small interface components. The modalities were non-speech audio, tactile and pseudo-haptic ‘stickiness’. All three modalities are readily available for standard desktop computers, with tactile only requiring a tactile mouse (rather than relatively expensive force-feedback devices). Stickiness is readily implemented by tailoring the mouse control-display gain when the cursor enters a target.

Results showed that, as expected, Fitts’ Law accurately models targeting with all combinations of modalities. Furthermore, when selecting small targets that are physically remote from other targets, stickiness can yield dramatic performance improvements, with the combination of stickiness and tactile appearing to be particularly efficient and appealing. However, the results of an ecological experiment in which the modalities were combined within menu-selection tasks clearly showed that poorly designed feedback can damage interaction by distracting users from tasks.

Future work will investigate the boundary conditions between successful multimodal feedback for discrete targets and the distraction of feedback in selecting from neighbouring targets.

7 Acknowledgements

This work was conducted while Brewster was on sabbatical at the University of Canterbury in Christchurch. His work was funded by an Erskine Visiting Fellowship and EPSRC grant GR/S53244/01.

8 REFERENCES

Akamatsu, M. and MacKenzie, I. S. (1996). "Movement Characteristics Using a Mouse with Tactile and Force Feedback." International Journal of Human-Computer Studies **45**: 483-493.

Akamatsu, M., MacKenzie, I. S. and Hasbrouc, T. (1995). "A Comparison of Tactile, Auditory, and Visual Feedback in a Pointing Task Using a Mouse-Type Device." Ergonomics **38**: 816-827.

Beaudouin-Lafon, M. and Conversy, S. (1996). Auditory Illusions for Audio Feedback. Proceedings of CHI'96 Conference on Human Factors in Computing Systems, Vancouver, Canada. 299-300.

Blanch, R., Guiard, Y. and Beaudouin-Lafon, M. (2004). Semantic pointing: improving target acquisition with control-display ratio adaptation. Proceedings of CHI2004: Conference on Human factors in computing systems, Vienna, Austria, ACM Press. 519-526.

Bolt, R. A. (1980). Put-that-there: Voice and gesture at the graphics interface. Proceedings of the 7th annual conference on Computer graphics and interactive techniques, Seattle, Washington, United States. 262-270.

Brewster, S. (1998a). "The Design of Sonically-Enhanced Widgets." Interacting with Computers: The Interdisciplinary Journal Human-Computer Interaction **11**(2): 211-235.

Brewster, S. (1998b). Sonically-Enhanced Drag and Drop. Proceedings of ICAD'98, Glasgow, UK, British Computer Society

Brewster, S. and Brown, L. (2004). Tactons: Structured Tactile Messages for Non-Visual Information Display. Proceedings of the Fifth Australasian User Interface Conference (AUIC'04), Dunedin, New Zealand. In press.

Brewster, S. and Crease, M. (1997). Making Menus Musical. Proceedings of INTERACT'97: the sixth IFIP conference on Human Computer Interaction, Chapman & Hall. 389-396.

Card, S., English, W. and Burr, B. (1978). "Evaluation of Mouse, Rate-Controlled Isometric Joystick, Step Keys, and Text Keys for Text Selection on a CRT." Ergonomics **21**(8): 601-613.

Cockburn, A. and Firth, A. (2003). Improving the Acquisition of Small Targets. People and Computers XVII (Proceedings of the 2003 British Computer Society Conference on Human-Computer Interaction.), Bath, England. 181-196.

Crossman, E. (1957). The speed and accuracy of hand movements. The nature and acquisition of industrial skill: Report to the MRC and DSIR Joint Committee on Individual Efficiency in Industry.

Dennerlein, J., Martin, D. and Hasser, C. (2000). Force-Feedback Improves Performance for Steering and Combined Steering-Targeting Tasks. Proceedings of CHI'2000 Conference on Human Factors in Computing Systems, The Hague, The Netherlands. 423--430.

Dennerlein, J. and Yang, M. (2001). "Haptic Force-Feedback Devices for the Office Computer: Performance and Musculoskeletal Loading Issues." Human Factors **43**(2): 278-286.

Fitts, P. (1954). "The Information Capacity of the Human Motor System in Controlling the Amplitude of Movement." Journal of Experimental Psychology **47**: 381-391.

Gann, K. and Hoffmann, E. (1998). "Geometrical conditions for ballistic and visually controlled movements." Ergonomics **31**(5): 829-839.

Hale, K. and Stanney, K. M. (2004). "Deriving Haptics Guidelines from Human Physiological, Psychophysical, and Neurological Foundations." IEEE Computer Graphics and Applications: 33-39.

Hart, S. and Staveland, L. (1988). Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. Human Mental Workload. P. a. M. Hancock, N**:** 139--183.

Hwang, F., Keates, S., Langdon, P. and Clarkson, P. (2003). Multiple Haptic Targets for Motion-Impaired Users. Proceedings of CHI'2003 Conference on Human Factors in Computing Systems, Fort Lauderdale, Florida. 41-48.

Jacko, J. A., Barnard, L., Kongnakorn, T., Moloney, K. P., Edwards, P. J., Emery, V. K. and Sainfort, F. (2004). Isolating the effects of visual impairment: exploring the effect of AMD on the utility of multimodal feedback. Proceedings CHI2004: Conference on Human factors in computing systems, Vienna, Austria, ACM Press. 311-318.

Jacko, J. A., Scott, I. U., Sainfort, F., Barnard, L., Edwards, P. J., Emery, V. K., Kongnakorn, T., Moloney, K. P. and Zorich, B. S. (2003). Older adults and visual impairment: what do exposure times and accuracy tell us about performance gains associated with multimodal feedback? Proceedings of CHI2003: Conference on Human factors in computing systems, Ft. Lauderdale, Florida, USA, ACM Press. 33-40.

Johnson, P., Hewes, J., Dropkin, J. and Rempel, D. (1993). Office Ergonomics: Motion Analysis of Computer Mouse Usage. Proceedings of the American Industrial Hygiene Association, Fairfax, VA. 12-13.

Keyson, D. (1997). "Dynamic cursor gain and tactual feedback in the capture of cursor movements." Ergonomics **40**(12): 1287-1298.

Langdon, P., Keates, S., Clarkson, P. and Robinson, P. (2000). Using Haptic Feedback to Enhance Computer Interaction for Motion-Impaired Users. 3rd International Conference on Disability, Virtual Reality and Associated Technologies (ICDVRAT 2000). Alghero, Sardinia, Italy. 25--32.

MacKenzie, I. (1992). Movement Time Prediction in Human-Computer Interfaces. Proceedings of Graphics Interface '92, Toronto, Canadian Information Processing Society. 140-150.

MacKenzie, I. and Buxton, W. (1992). Extending Fitts' Law to Two-Dimensional Tasks. Proceedings of CHI'92 Conference on Human Factors in Computing Systems, Monterey, CA. 219--226.

Oakley, I., Adams, A., Brewster, S. and Gray, P. (2002). Guidelines for the Design of Haptic Widgets. People and Computers XVI (Proceedings of the 2002 British Computer Society Conference on Human-Computer Interaction), London, UK, Springer. 195-212.

Oakley, I., Brewster, S. and Gray, P. (2001). Solving Multi-Target Haptic Problems in Menu Interaction. Extended Abstracts of CHI'2001: ACM Conference on Human Factors in Computing Systems, Seattle, WA, ACM Press. 357-358.

Oakley, I., McGee, M., Brewster, S. and Gray, P. (2000). Putting the Feel in 'Look and Feel'. Proceedings of CHI'2000 Conference on Human Factors in Computing Systems, The Hague, The Netherlands. 415--422.

Oviatt, S. (1999). "Ten myths of multimodal interaction." Commun. ACM **42**(11): 74-81.

Oviatt, S. (2002). Multimodal Interfaces. Handbook of Human-Computer Interaction. J. Jacko and A. Sears. New Jersey, Lawrence Erlbaum.

Oviatt, S., DeAngeli, A. and Kuhn, K. (1997). Integration and synchronization of input modes during multimodal human-computer interaction. Proceedings of CHI'97: Conference on Human Factors in Computing Systems, Atlanta, Georgia, ACM Press. 415-422.

Reeves, L. M., Lai, J., Larson, J. A., Oviatt, S., Balaji, T. S., Buisine, S. p., Collings, P., Cohen, P., Kraal, B., Martin, J.-C., McTear, M., Raman, T., Stanney, K. M., Su, H. and Wang, Q. Y. (2004). "Guidelines for multimodal user interface design." Commun. ACM **47**(1): 57-59.

Soukoreff, R. and MacKenzie, I. S. (2004). "Towards a standard for pointing device evaluation, perspectives on 27 years of Fitts' law research in HCI." International Journal of Human Computer Studies **61**(6): 751-789.

Vitense, H., Jacko, J. and Emery, V. (2003). "Multimodal feedback: An assessment of performance and mental workload." Ergonomics **46**(1-3): 68-87.

Wall, S., Riedel, B., Crossan, A. and McGee, M., Eds. (2002). Proceedings of Eurohaptics. Edinburgh, University of Edinburgh.

Worden, A., Walker, N., Bharat, K. and Hudson, S. (1997). Making Computers Easier for Older Adults to Use: Area Cursors and Sticky Icons. Proceedings of CHI'97 Conference on Human Factors in Computing Systems, Atlanta, Georgia. 266-271.

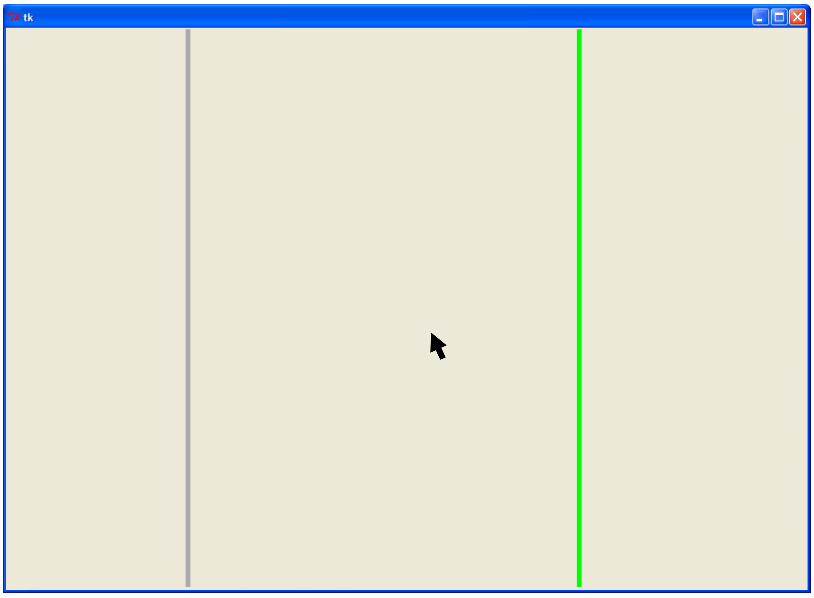

FIGURE 1: The interface used in Experiment One. The user moves from the grey line (on the left) to acquire the green one (on the right) as quickly as possible. On selecting the green line, the grey line moves, the highlighting is toggled, and the user targets the new green line.

FIGURE 2: The Logitech iFeel mouse (www.logitech.com). The mouse looks and operates like an ordinary desktop mouse but contains a motor with an eccentric weight which generates a range of different types of simple vibrotactile feedback to the hand.

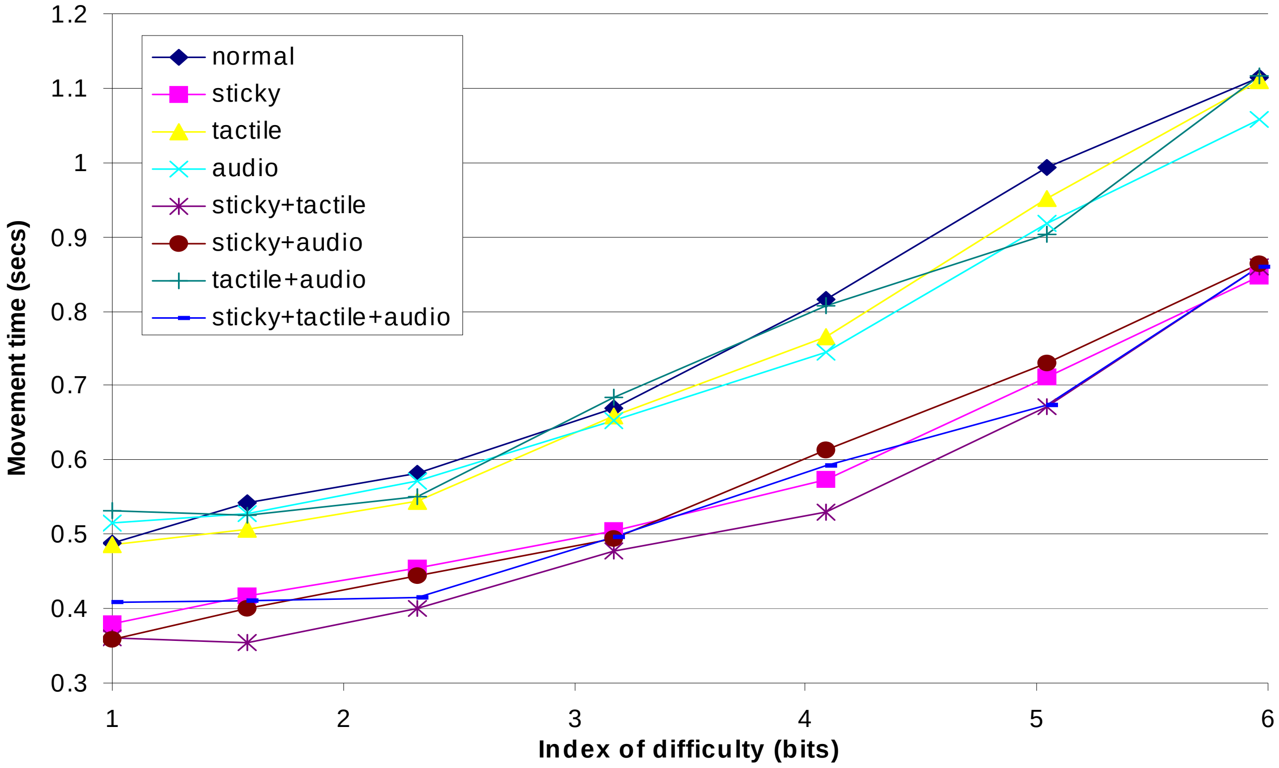

FIGURE 3: Mean movement times plotted against Index of Difficulty for the eight combinations of modality in Experiment One.

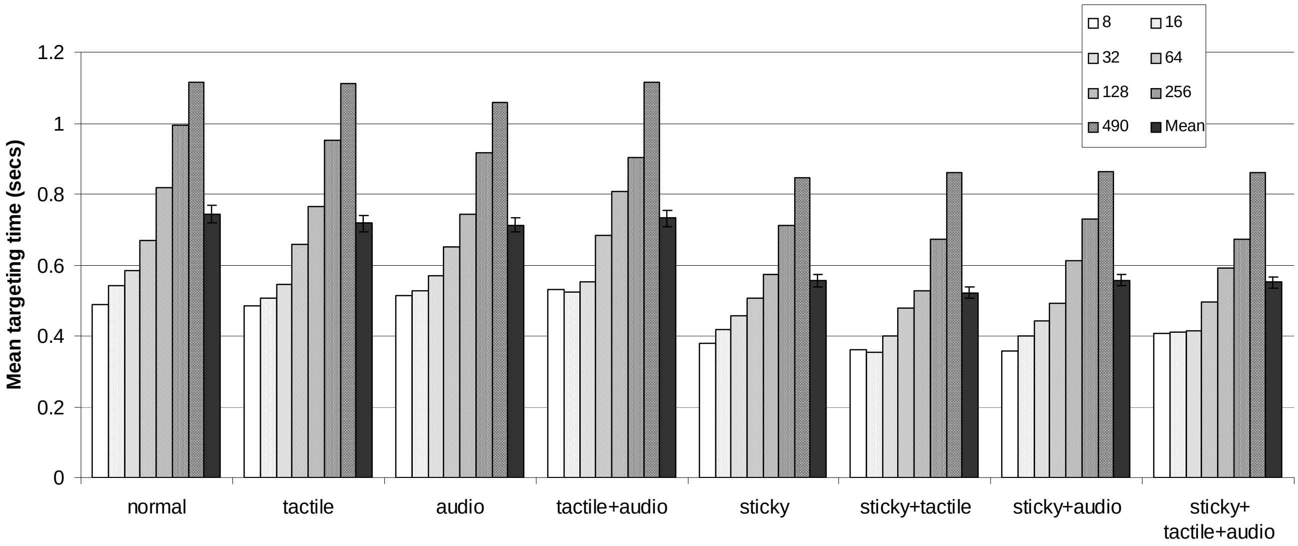

FIGURE 4: Mean target acquisition times for the eight combinations of modality across the seven distances in Experiment One. Error bars on the overall means show one standard error above and below the mean.

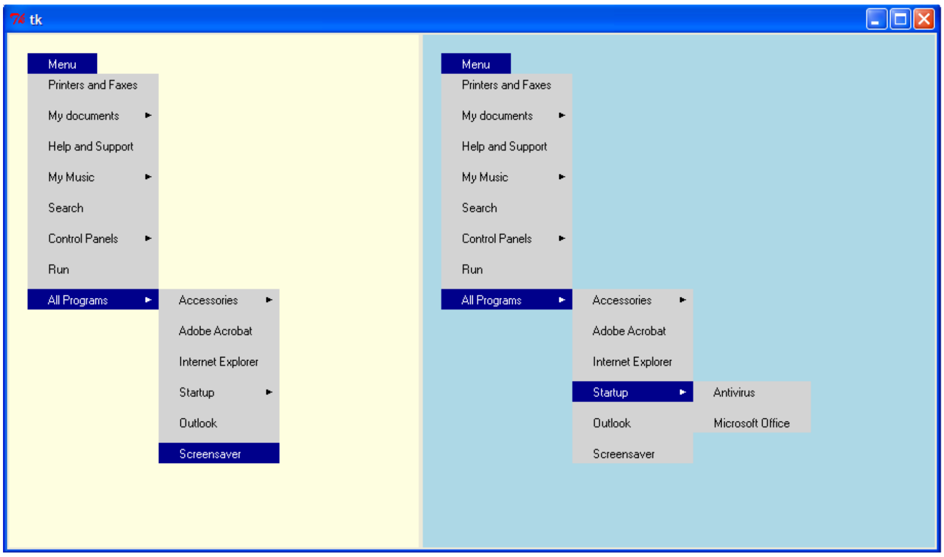

FIGURE 5: Interface used in Experiment Two. The left-hand pane shows the target menu item the user must choose. The right-hand pane contains the menu the user interacts with to select the item.

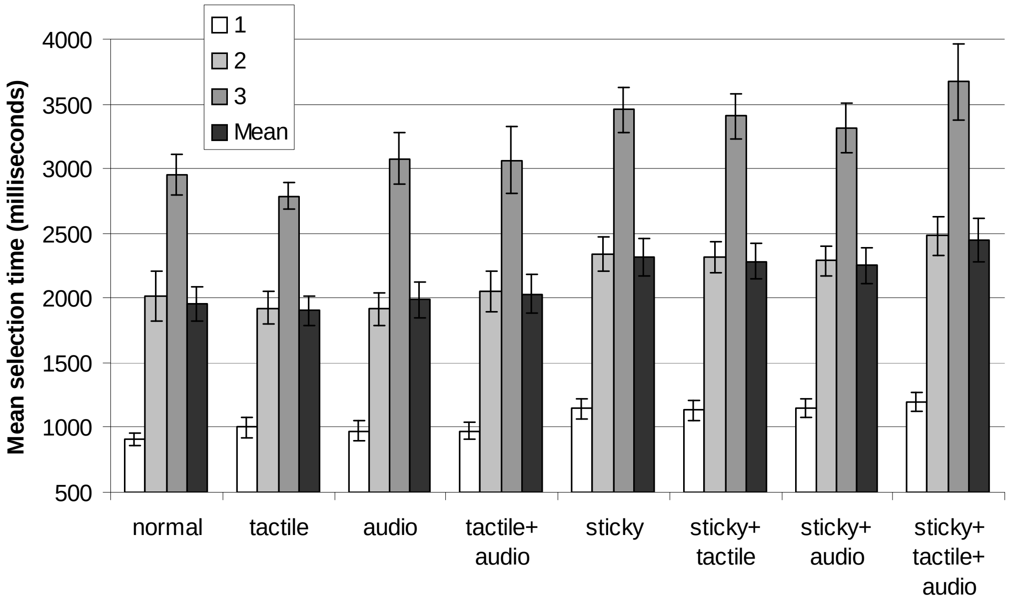

FIGURE 6: Mean menu selection times for the different combinations of feedback modalities across three levels of menu depth in Experiment Two. Error bars show one standard error above and below the mean.

TABLE 1: Linear regression equations, R^2 values, and Indices of Performance for the feedback modes in Experiment One. The left-hand main column shows results based on only those tasks with ID > 3.0. The right-hand main column shows results for all tasks, but using an 'adjustment for accuracy' calculation for target width.

	Models based on ID > 3.0			**Models based on 4.133×σ 'adjustment for accuracy'**		
Method	**Line of best fit**	**R^2**	**IP (bits/sec)**	**Line of best fit**	**R^2**	**IP (bits/sec)**
Normal	MT=158 + 162×ID	0.99	6.17	MT= 251 + 125×ID	0.96	8.03
Sticky	MT=89 + 125×ID	0.98	8.01	MT= 210 + 85×ID	0.95	11.79
Tactile	MT=115 + 166×ID	0.99	6.04	MT= 216 + 130×ID	0.94	7.72
Audio	MT=161 + 149×ID	0.99	6.7	MT= 280 + 108×ID	0.91	9.25
Sticky+Tactile	MT=5 + 138×ID	0.95	7.25	MT= 106 + 102×ID	0.90	9.83
Sticky+Audio	MT=75 + 131×ID	0.97	7.61	MT= 189 + 90×ID	0.89	11.13
Tactile+Audio	MT=197 + 149×ID	0.97	6.7	MT= 317 + 107×ID	0.94	9.38
Sticky+Tactile+Audio	MT=83 + 125×ID	0.96	7.99	MT= 182 + 89×ID	0.92	11.25