

Evaluating Tactile Feedback and Direct vs. Indirect Stylus Input in Pointing and Crossing Selection Tasks

Clifton Forlines^{1,2}

¹Mitsubishi Electric Research Labs
Cambridge, MA, USA
www.merl.com
forlines@merl.com

Ravin Balakrishnan²

²Department of Computer Science
University of Toronto
www.dgp.toronto.edu
ravin@dgp.toronto.edu

ABSTRACT

We present a pair of experiments that explore the effects of tactile-feedback and direct vs. indirect pen input on pointing and crossing selection tasks. While previous work has demonstrated the validity of crossing as a useful selection mechanism for pen-based computing, those experiments were conducted using an indirect input device – one in which the pen-input and display were separated. We investigate users' performance with pointing and crossing interfaces controlled via not only an indirect input device, but also a direct input device – one in which the pen-input and display are co-located. Results show that direct input significantly outperforms indirect input for crossing selection, but the two modalities are essentially equivalent in pointing selection. A small amount of tactile feedback is shown to be beneficial for both pointing and crossing selection, most noticeably in crossing tasks when using direct input where visual feedback is often occluded by a hand or stylus.

ACM Classification: H.5.2. [User Interfaces]: Input devices and strategies.

General terms: Experimentation, Human Factors

Keywords: Haptic & tactile feedback, crossing interfaces, pen/stylus input, direct & indirect input

INTRODUCTION

Researchers have recently been investigating crossing interfaces [1, 4, 5] as an alternative to pointing-and-clicking interfaces. In a crossing interface, an action occurs when the user moves a pointer through a boundary instead of when the user taps within a target. It has been argued [4] that crossing may be preferable to pointing for pen-based computing because it supports fluid transitions from one action to another. In a study aimed at understanding the properties of crossing performance, Accot and Zhai [1] compared four crossing selection techniques and two pointing selection techniques. They found that crossing obeyed a Fitts' law [8] model and was as fast, or faster, than pointing for a variety of selection tasks. However, their experimental apparatus consisted of a Wacom tablet and separate

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2008, April 5–10, 2008, Florence, Italy.

Copyright 2008 ACM 978-1-60558-011-1/08/04...\$5.00.

CRT monitor, resulting in an *indirect* input to display mapping (Figure 1, left) which is quite different from the *direct* input to display mapping found in current tablet computers (Figure 1, right). In contrast, Ren and Moriya [16] measured user performance with *direct* stylus input (Figure 1, right) and found no difference in selection time between sliding (i.e., crossing) into a target and clicking directly on a target, although sliding had significantly fewer errors than clicking. The literature is thus inconclusive as to the impact of direct vs. indirect stylus input on crossing and pointing tasks, and a systematic study of this issue is crucial in order to guide the designs of future pen-based interfaces.

Another important issue to consider is that while direct input provides strong affordances for interaction [18], partial occlusion of the target and screen by a hand or stylus can make visual feedback difficult to perceive. This is particularly problematic when the appearance of a target being honed in upon, and thus visually occluded, is meant to convey state information to the user. It is thus worth exploring if tactile feedback is a viable second channel for conveying state information to the user in a private fashion.

In this paper, we present a pair of experiments designed to answer two questions. First, how does using direct vs. indirect stylus input impact user performance in pointing and crossing selection tasks? Second, does tactile feedback provided via a stylus equipped with a solenoid [10] improve user performance in such tasks, particularly when the visual occlusion problem is present with direct stylus input? Note that our focus is on the impact of tactile feedback and direct vs. indirect input on both pointing and crossing selection tasks, and not on pointing vs. crossing selection per se. We conclude by discussing how our experimental results can inform the design of future pen-based interfaces.




Figure 1. (left) *Indirect* input: stylus input and display are separated. (right) *Direct* input: stylus input and display are co-incident. The same tactile feedback stylus is used in both cases.

TACTILE FEEDBACK

Several terms are used interchangeably by researchers investigating tactile and force-feedback for enhancing interfaces [2, 3, 12-15], so we must define our terminology. We use the term *tactile feedback* to mean that the system is indicating state to the user through the sense of touch, and *force feedback* when the system is providing physical resistance to the user in response to the user's movements.

Akamatsu et al. [3] modified a computer mouse to provide tactile feedback via a pin projecting through a hole in a mouse button. In a target selection study, they found tactile feedback improved selection times by 10.6%. In another study, Akamatsu et al. [2] reported that tactile feedback improved selection times by 5.6% at the expense of a 65.2% increase in error rate. Note that these experiments [2, 3] examined selecting a series of single isolated targets. As Oakley et al. [12] point out, many studies that demonstrate the benefits of force feedback typically ignore the common situation where numerous "haptic enhanced" targets can act as distracters when they have to be passed over en-route to the target of interest. Indeed, Oakley et al. [12] found that widgets providing force-feedback using a Phantom perform worse or no better than those without such feedback when used in a multi-target environment.

Stylus with Tactile Feedback

We modified a TabletPC stylus to provide tactile feedback in a manner similar to the Haptic Pen of Lee et al. [10]. Our stylus (Figure 1) has a 20V solenoid attached to the eraser end. We modified the design of Lee et al. [10] by attaching a spring to the end in the solenoid so that it does not rely only on weight and gravity to return to its resting state, thus working correctly when the pen is used at an angle (as is the case when operated on a handheld TabletPC) or even completely inverted. Note that this design provides tactile feedback, rather than force feedback found in, for example, the Phantom which requires a fixed point of reference and would be inappropriate for a mobile device.

While Lee et al. [10] explored a wide range of mappings between pen pressure and types of feedback, we used a single type of feedback in our studies. Our stylus indicated a successful selection with a simple click, which gave the tactile impression of running the stylus tip over a small bump in the screen. This provides tactile and aural feedback similar to what a user experiences when using a mouse, where the feedback is inherent in the mechanical properties of a mouse button's micro-switch. In contrast, the typical electronic stylus provides no such mechanical click feedback when the pen tip crosses a virtual boundary.

Placing the tactile feedback actuator inside the stylus is in contrast to approaches that place actuators behind the display glass [9, 14]. While this alternative eliminates the need for an active stylus, it has three major drawbacks – it may not scale to large wall-sized displays, it cannot provide tactile feedback when the stylus is hovering above the display, and, if used in a multi-user system, all of the users would be forced to share a single channel of tactile feedback.

Poupyrev et al. [15] described a study in which participants used an actuated-display-glass system [14] to perform pointing selection in a direct input configuration. They hypothesized that tactile feedback would improve standard Fitt's-style tapping tasks; however, their results showed no significant difference in selection times between trials with and without tactile feedback. Two important differences exist between the experiment described by Poupyrev et al. [15] and those described in this paper. First, it is important to note that Poupyrev et al.'s experiment was performed on a *direct* input device, and so it is not clear if their results would transfer to an *indirect* input device. Second, because their experiment used an actuated display screen, tactile feedback was only possible when the pen tip was in contact with the display. In a tapping task, the duration of this contact is very short, which results in a short window in which to provide feedback. Indeed, many of their participants did not notice the tactile feedback at all. In contrast, the haptic pen used in our study is capable of providing tactile feedback even when the pen is lifted from the display, which could lead to differences in performance when tactile feedback is used in pointing tasks.

DIRECT vs. INDIRECT and RELATIVE vs. ABSOLUTE

As discussed in the introduction and Figure 1, we use the terms *direct* and *indirect* to refer to the co-locality of input device and output display. The relationship between the input device and the virtual cursor, however, is another important issue. We use the term *absolute* input when the device's position in its operating region (e.g., stylus on a tablet) is reported relative to an origin that remains invariant throughout use, and is mapped to the cursor's position on screen in a direct one-to-one mapping with constant gain. In contrast, *relative* input occurs when the reference origin can be dynamically altered by a clutching mechanism, and where the device to cursor mapping can have a non one-to-one mapping with variable gain. The mouse is the canonical example of a relative device. A stylus on a tablet is typically used in absolute mode, but could operate in a relative manner if a clutching mechanism is used. In the present work, we use the stylus in the usual *absolute* mode in both *direct* and *indirect* device-display settings.

EXPERIMENT ONE – POINTING SELECTION

Participants

24 participants, 10 women and 14 men, between 18 and 37 years old, recruited mainly from local universities, volunteered for the experiment. All were regular computer users but did not have significant experience using a computer stylus or a TabletPC. They were paid \$20 each,

Apparatus

The experiment was conducted on a 1.7 GHz Toshiba Portege M200 TabletPC running WindowsXP TabletPC Edition. The 12.1" screen had a resolution of 1400x1050 pixels. The computer was setup in the tablet configuration (with its screen folded to cover the keyboard and track pad) and positioned horizontally on a desk in front of the user. Participants sat during the experiment and operated the

tablet with our custom tactile feedback stylus in their dominant hand. In the *direct* input configuration (Figure 1), the tablet input was co-incident with the display screen beneath. In the *indirect* input configuration (Figure 1), participants used the same tablet and stylus for input, but the underlying display was turned off and a vertical desktop display of equal size and resolution was used instead.

Task and Stimuli

In order to build upon the literature, our tasks were designed to match those in Accot and Zhai’s [1] experiment, and for consistency we also use their terminology. Although replication of experiments is uncommon in the field of human-computer interaction, it is standard practice in other domains and is crucial in ensuring the validity of research results and continued investigation into important topics. By following Accot and Zhai’s design, but with a different range of target widths and distances and additional input and feedback conditions, we not only shed light on the impact of tactile feedback and direct vs. indirect input but also strengthen the validity and reliability of the prior research.

In Experiment One, we studied two pointing tasks (Figure 2). These tasks differed in target constraint – constraining targets either collinearly or orthogonally to the main direction of stylus movement.

CP – Pointing with Collinear Constraint (Figure 2 left). This matches the traditional Fitts reciprocal tapping task [8]. Participants tap reciprocally between two rectangles of width W , an effectively infinite height, and distance D apart measured from the target centers.

OP – Pointing with Orthogonal Constraint (Figure 2 right). Targets are positioned so that they have a height of W and an effectively infinite width. The inside edges of the two targets are separated by a distance of D . The constraint is named orthogonal because the vertical constraint is orthogonal to the horizontal direction of stylus movement.

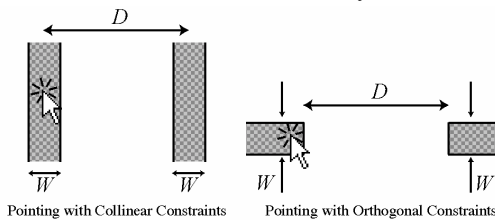


Figure 2. Pointing techniques in Expt. 1. Images and terminology modeled after Accot and Zhai [1].

Participants received two feedback types – *visual* or *visual-plus-tactile*. In both conditions, the active target was displayed in green, and the inactive target in grey. A correctly selected target would flash orange before turning grey. Upon the successful selection of a target, the opposite target would turn green. In addition, in the *visual-plus-tactile* condition the tactile feedback stylus would produce a soft click. This feels much like a mouse-button click, and is clearly distinguishable from the feeling of the stylus tip contacting the display glass. When a target was missed, the participant received no tactile feedback.

Design

The independent variables for our two pointing tasks were *input type* (direct and indirect), *target constraint* (Orthogonal and Collinear), *target distance* D (256 and 1024 pixels), *target width* W (4, 8, 16, and 32 pixels), and *feedback type* (visual and visual-plus-tactile). A repeated-measures within-participant design was used with input type as a between-participants variable, with participants split into 2 groups of 12. We deliberately made input type a between-participants variable in order to reduce the possibility of asymmetric skill transfer, and to reduce the complexity of the experiment. The order of presentation of the two tasks was counterbalanced within each group.

The smallest width of 4 pixels was known to be very difficult, but was chosen for external validity reasons because it is the smallest size target that WindowsXP Tablet Edition requires users to be able to select.

For each task, participants performed three blocks of trials for each of the two feedback conditions. The order of presentation of the feedback condition was counterbalanced between participants. The first block was practice and the last two were data collection sessions. Within each block, participants performed seven repetitions for each of the 8 D - W combinations, which were presented in random order. The first of these seven repetitions was discarded, because of the uncontrolled starting position of the stylus at the start of each set. Participants could take breaks between feedback conditions and were forced to take breaks between the two tasks. In summary, the design was as follows:

12 participants x
 2 input types (direct and indirect, between-subjects) x
 2 pointing tasks (CP, OP) x
 2 feedback type conditions (visual and visual-plus-tactile) x
 2 blocks x
 8 D - W combinations x
 6 repetitions per D - W combination =
 9216 selections in total.

For each trial, participants had to successfully select the target before moving on to the next trial, even if it required multiple taps with the stylus. Trials where participants did not successfully select the target on the first attempt were marked as errors. This design effectively removes the possibility that participants may try to “race” through the experiment by selecting wildly anywhere.

Hypotheses

- H1. Input type will affect selection time in pointing tasks
- H2. Input type will affect error rates in pointing tasks
- H3. Feedback type will affect selection time in pointing tasks
- H4. Feedback type will not affect error rates in pointing tasks
- H5. Tactile feedback will have a greater impact on performance with direct than with indirect input, due to the hand occlusion present in direct input situations.

Experiment One Results

Selection Time Analysis

The results in this section exclude trials in which there were errors as well as 200 trials (2.1% of our data) with selection times more than three standard deviations from the mean. The first block of trials was counted as practice and not used in the data analysis. A repeated-measures ANOVA on the remaining two blocks shows no effect of block in terms of selection time ($F_{1,22} = 0.39$, $p = 0.54$).

We found no significant effect for *input type* on selection time ($F_{1,22} = 0.64$, $p = 0.43$), thus hypothesis H1 is not confirmed. There was a significant interaction between *target width* and *input type* ($F_{1,22} = 6.39$, $p = 0.02$) (Figure 3), with direct input resulting in faster performance relative to indirect input as target widths increased (i.e., became easier to acquire); however, there was also a significant interaction between *target distance* and *input type* ($F_{1,22} = 9.37$, $p = 0.006$), with direct input resulting in faster performance relative to indirect input as distances increased (Figure 4).

Feedback type had a significant main effect on selection time ($F_{1,22} = 26.08$, $p < 0.001$), thus confirming H3. Visual-plus-tactile feedback was 6.9% faster than visual feedback alone, with means of 1.04 and 1.12 seconds for visual-plus-tactile and visual feedback respectively. There was also a significant interaction between *feedback type* and *input type* ($F_{1,22} = 13.15$, $p = 0.001$). In contrast to H5, tactile feedback improved selection time in the *indirect input* condition (by 11.9%) but had no effect with direct input (Figure 5).

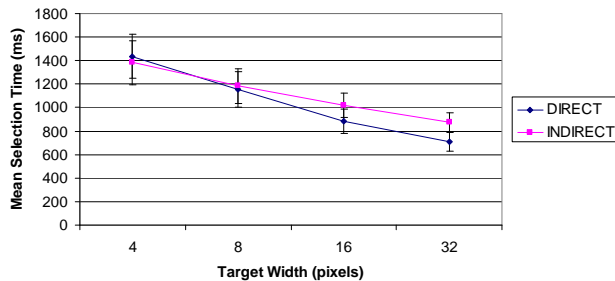


Figure 3. Mean selection times for the two pointing tasks, by *target width* and *input type*.

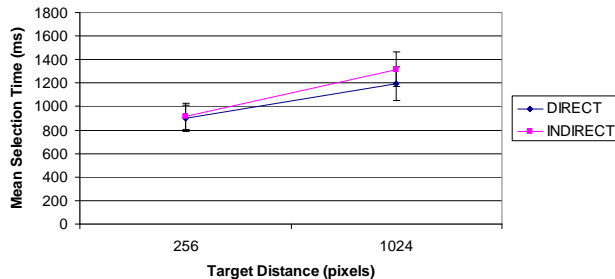


Figure 4. Mean selection times for the two pointing tasks, by *target distance* and *input type*.

Error Rate Analysis

Again, the first of the three blocks of trials was counted as practice and removed from the data analysis. A repeated-measures ANOVA on the remaining two blocks shows no effect of block on error rate ($F_{1,22} = 1.94$, $p = 0.18$).

There was a significant effect for *input type* on error rate ($F_{1,22} = 4.60$, $p = 0.043$), with mean error rates of 19.6% and 34.2% for direct and indirect input respectively, thus confirming H2. Note that reasons for these relatively high error rates are discussed in the “Overall Discussion” section at the end of the paper. There was a significant interaction between *target constraint* and *input type* ($F_{1,22} = 5.08$, $p = 0.034$). There were fewer errors in the direct input condition for targets with collinear constraints, and in the indirect input condition for targets with orthogonal constraints. Finally, as one would expect with feedback that occurs after selection, *feedback type* had no significant effect on error rate ($F_{1,22} = 0.013$, $p = 0.91$), thus H4 is confirmed. No other significant effects with respect to selection time or error rate for the two pointing tasks were observed.

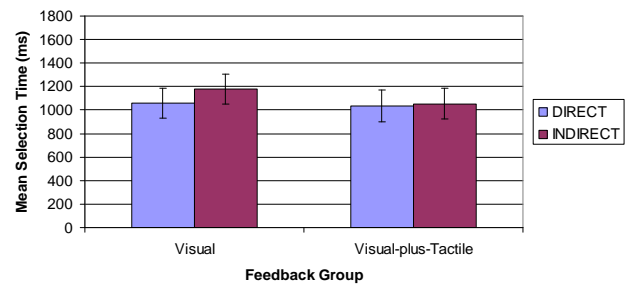


Figure 5. Mean selection times for the two pointing tasks, by *feedback type* and *input type*.

Fitts' Law Analysis

Fitts' law is widely used to model pointing tasks, and Accot and Zhai [1] showed that it also applied to crossing tasks. Fitts' law states that movement time $MT = a + b \log(D/W + 1)$, where a and b are determined empirically, and the log term is referred to as the task's index of difficulty (ID). The ID s in our tasks ranged from 3.2 to 8: a larger range than used in typical Fitts' studies. Fitts' models for each pointing task per input and feedback type resulted in good fits to the data ($r^2 > 0.92$ in all cases) (Table 1).

Input, Feedback Type	Model	r^2
Indirect, Visual+Tactile	$Time = 135.0 + 167.7 ID$	0.92
Indirect, Visual	$Time = 108.3 + 190.7 ID$	0.92
Direct, Visual+Tactile	$Time = -41.1 + 190.1 ID$	0.94
Direct, Visual	$Time = -98.6 + 206.7 ID$	0.93

Table 1. Fitts' law models for the two pointing tasks, by *input type* and *feedback type*.

In the Fitts' law literature and the ISO 9241-9 standard on input performance, there has been an attempt to use a single measure called index of performance ($IP = ID / MT$) to capture the performance using Fitts' law. However, Zhai [19] correctly argues that this single measure of performance is fraught with problems and that it is preferable to report results as the complete Fitts' model, including both a and b as separate numbers. An interesting observation with our data is that the mean selection times (Figure 5) show that performance with direct input is generally faster than indirect input. A traditional calculation of $IP = 1/b$, however, would erroneously indicate that direct input results in

lower performance than indirect input. This is because the calculation of IP , as Zhai [19] argues, ignores the differences in a between different conditions. In our data, a (the y-intercept) is different for the indirect and direct conditions, showing essentially a constant overhead cost incurred by the indirect technique. It is difficult to ascertain exactly what causes this overhead cost, but it is likely due to the separation between input and output spaces.

Experiment One Discussion

One explanation for the interaction between *target width* and *input type* is based on the two targeting methods possible in the direct input condition. When selecting with direct input, users may track the focus of interaction with either the system pointer or the physical tip of the stylus. As shown in Figure 3, there was little difference between direct and indirect input for small targets. Because of the parallax between the stylus tip and system pointer due to the thickness of the glass surface of the display, small targets could only accurately be selected in the direct input condition by tracking with the system pointer. Larger targets could be easily selected by tracking with the stylus tip. In the indirect input condition, participants had to track with the system pointer for all target sizes – so they were not able to benefit from tracking the stylus tip, even for large targets.

The interaction between *target distance* and *input type* could be explained as follows: When the stylus is lifted away from the tablet surface, it eventually moves out of the tablet’s sensing range and the system pointer’s position on screen is no longer updated to reflect the current position of the stylus. When the stylus moves back into the sensing range of the tablet, the pointer instantly warps to the current position. In the direct input condition, the physical tip of the stylus provides excellent tracking information to the user even when the pen moves out of range. This occurs because the user’s visual focus is on the display that is coincident with the tablet. In the indirect input condition, the user’s visual focus is on the vertical screen and the physical stylus tip is thus not seen. Further, longer distances are more likely to lead to a lifting of the stylus to an out-of-range position and then returning within range when homing in on the target. Thus, selection times in the indirect input condition are penalized by a greater degree when distances are large.

Although counter to our hypothesis H5, our data corresponds with Poupyrev et al.’s [15] in showing that tactile feedback does not aid pointing times when used in a direct input configuration (Figure 5). Further, our data adds to the literature by showing that tactile feedback does aid pointing selection when used in an indirect input configuration. This suggests that input type (direct or indirect) is an important factor to be considered when deciding whether or not to use tactile feedback for pointing tasks.

EXPERIMENT TWO – CROSSING SELECTION

The same 24 participants in Experiment One took part in Experiment Two, using the same TabletPC and vertical desktop display as they used in Experiment One.

Task and Stimuli

Our four crossing tasks (Figure 6) match those in Accot and Zhai’s [1] experiment, differing in the target constraint and in the continuity of contact between the stylus and the display. For consistency we again use their terminology.

D/CC – Discrete Collinear Crossing (Figure 6, upper left). Participants reciprocally cross two horizontal targets of width W , height one, and separated by a distance D . Participants were asked to cross these targets from the top through to the bottom. This condition was named *discrete* because participants were asked to lift their stylus from the tablet surface in-between target selections. A thick vertical line split the display in half, and if participants stroked over this line with their stylus, the system played a beep sound to remind them to lift their stylus in-between targets.

D/OC – Discrete Orthogonal Crossing (Figure 6, upper right). Participants discretely crossed two vertical targets of width W separated by a distance D . The same thick vertical line as in the *D/CC* condition reminded participants to lift their stylus from the tablet in-between target selections.

C/CC – Continuous Collinear Crossing (Figure 6, lower left). Participants reciprocally selected two targets of width W separated by a distance D by crossing through them with the stylus. Participants were asked to select these targets by stroking down from above the target through to below the target. This condition was named *continuous* because participants were asked to keep their stylus in continuous contact with the screen. If the stylus was lifted from the tablet surface, the system repeatedly played a beeping sound to remind users to keep the stylus in contact with the surface.

C/OC – Continuous Orthogonal Crossing (Figure 6, lower right). Participants reciprocally selected two targets of height W , width one, and separated by a distance D . Participants were asked to stroke the stylus from the left side to the right side of the target. As in the *C/CC* condition, the system would play a beeping sound whenever the participant lifted the stylus from the tablet to remind them to keep the stylus tip in continuous contact with the display.

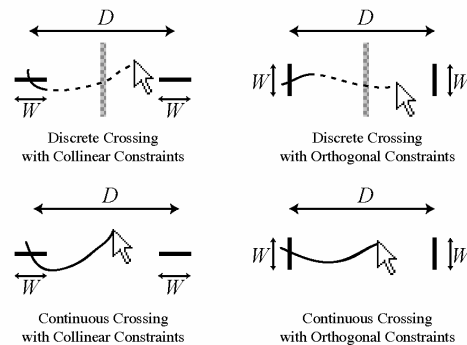


Figure 6. Four crossing selection techniques used in Experiment Two. Images and terminology modeled after Accot and Zhai [1].

Participants again received two types of feedback – *visual* or *visual-plus-tactile*. When stroking through a crossing-based target, the tactile feedback feels like dragging a pencil over a thin wire.

Design

The independent variables for the four crossing tasks were *input type* (direct and indirect), *target constraint* (Orthogonal and Collinear), *target distance D* (256 and 1024 pixels), *target width W* (4, 8, 16, and 32 pixels), *feedback type* (visual and visual-plus-tactile), and *continuity of contact* (Continuous and Discrete).

A repeated-measures within-participant design was used with input type as a between-participants variable. The order of presentation of the four tasks was counterbalanced within each group using a Latin-squares design. Participants could take breaks between feedback conditions, and were forced to take breaks between each of the six tasks. In summary, the design was as follows (excluding warm-ups):

12 participants x
2 input types (direct and indirect, between-subjects) x
4 crossing tasks (D/CC, D/OC, C/CC, and C/OC) x
2 feedback type conditions (visual and visual-plus-tactile) x
2 blocks x
8 D-W combinations x
6 repetitions per D-W combination =
18432 selections in total.

Hypotheses

- H6. Input type will affect selection time in crossing tasks
- H7. Input type will affect error rates in crossing tasks
- H8. Feedback type will affect selection time in crossing tasks
- H9. Feedback type will not affect error rates in crossing tasks
- H10. Tactile feedback will have a greater impact on direct input than indirect input performance, due to the hand occlusion present in direct input situations.

Experiment Two Results

Selection Time Analysis

Selection time was defined as the amount of time taken between a target turning green and the successful crossing selection of that target. The results in this section include only trials in which there were no errors (i.e., those where the target was successfully selected on the first attempt). Also, 537 trials (2.9% of our data) with selection times that were more than three standard deviations from the mean were identified as outliers and removed from the analysis.

The first of the three blocks of trials was counted as practice and not used in the data analysis. A repeated-measures ANOVA on the remaining two blocks showed no effect of block on selection time ($F_{1,22} = 0.46$, $p = 0.51$).

A repeated-measures ANOVA showed a significant main effect for *input type* on selection time ($F_{1,22} = 5.93$, $p = 0.023$), thus confirming hypothesis H6. Averaged across all four crossing tasks, direct input was 16% faster than indirect input, with a mean selection time of 1.05 seconds and 1.22 seconds for direct and indirect input respectively.

As shown in Figure 7, there was a significant interaction between *target width* and *input type* ($F_{3,66} = 3.07$, $p = 0.03$), with direct input resulting in faster performance relative to

indirect input as target widths increased (i.e., became easier to acquire). However, there was also a significant interaction between *target distance* and *input type* ($F_{1,22} = 7.74$, $p = 0.02$), with direct input resulting in faster performance relative to indirect input as distances increased (Figure 8).

Feedback type had a significant main effect on selection time ($F_{1,22} = 41.86$, $p < 0.001$), thus confirming hypothesis H8. Averaged across all four crossing tasks, visual-plus-tactile feedback was 6.3% faster than visual feedback alone, with mean times of 1.10 and 1.17 seconds for visual-plus-tactile and visual feedback respectively. There was also a significant interaction between *feedback type* and *input type* ($F_{1,22} = 29.52$, $p < 0.001$) with visual-plus-tactile feedback providing almost no benefit in the indirect input condition yet leading to about 11% faster selection times in the direct input condition, thus confirming hypothesis H10. Figure 9 illustrates these effects. There was also an interaction between *feedback type* and *target width* on selection time ($F_{1,22} = 17.35$, $p < 0.001$), with smaller targets benefitting more from tactile feedback (Figure 10).

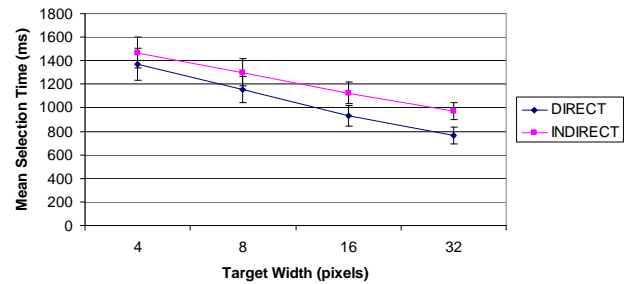


Figure 7. Mean selection times for the four crossing tasks, by *target width* and *input type*.

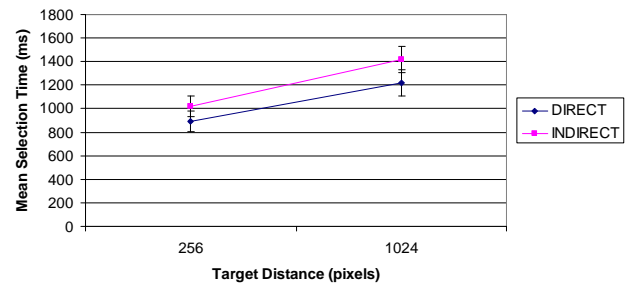


Figure 8. Mean selection times for the four crossing tasks, by *target distance* and *input type*.

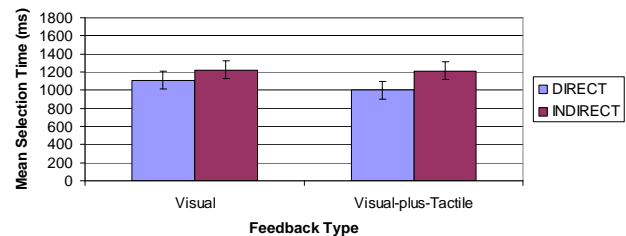


Figure 9. Mean selection times for the four crossing tasks, by *feedback type* and *input type*.

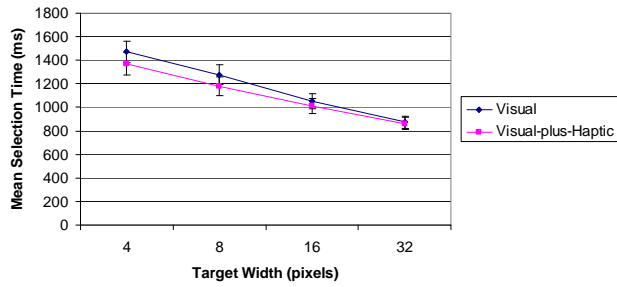


Figure 10. Mean selection times for the four crossing tasks, by *target width* and *feedback type*.

There was a significant interaction between *continuity of contact* and *input type* ($F_{1,22} = 23.93$, $p < 0.001$). For direct input, the discrete tasks D/CC and D/OC which required lifting the stylus in-between selections resulted in 11% lower selection times than the continuous tasks C/CC and C/OC (1.00 vs. 1.11 seconds respectively). In contrast, for indirect input, the discrete tasks resulted in 14% *higher* selection times than the continuous tasks (1.30 vs. 1.14 seconds respectively).

Error Rate Analysis

The first block of trials was counted as practice and removed from the data analysis. A repeated-measures ANOVA on the remaining two blocks shows no effect of block on error rate ($F_{1,22} = 0.23$, $p = 0.63$).

Input type did not have a significant effect on error rate ($F_{1,22} = 0.074$, $p = 0.79$), with mean error rates of 19.3% and 20.3% for direct and indirect input respectively; thus, H7 was not confirmed. As shown in Figure 11, there was a significant interaction between *continuity of contact* and *input type* ($F_{1,22} = 5.15$, $p = 0.03$), with continuous input resulting in similar errors for both direct and indirect input and discrete input resulting in a lower error rate for direct input relative to indirect input. Unlike selection time, there was no significant interaction between *target width* and *input type* ($F_{1,22} = 0.059$, $p = 0.81$). However, there was a significant interaction between *target distance* and *input type* ($F_{1,22} = 7.55$, $p = 0.012$), with indirect input resulting in higher error rates relative to direct input as target distances increased (Figure 12). As expected with feedback that occurs after the selection, *feedback type* had no significant effect on error rate ($F_{1,22} = 0.57$, $p = 0.46$), with mean error rates of 19.6% and 20.1% for visual and visual-plus-tactile feedback respectively. No other significant effects with respect to feedback type and/or input type on error rates were observed for the four crossing tasks.

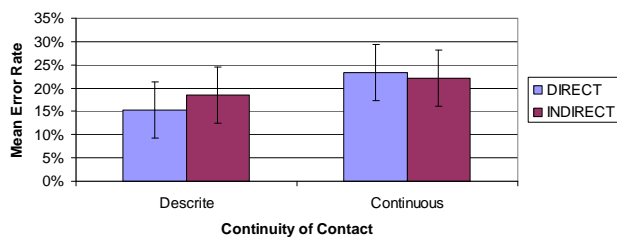


Figure 11. Mean error rates for the four crossing tasks, by *continuity of contact* and *input type*.

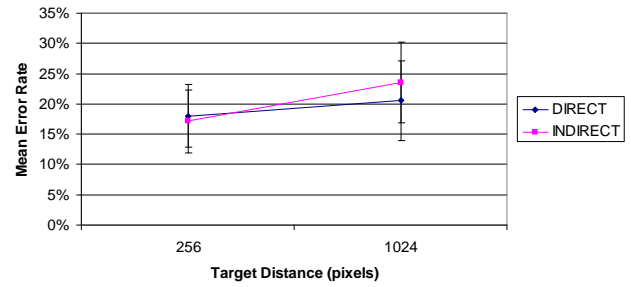


Figure 12. Mean error rates for the four crossing tasks, by *target distance* and *input type*.

Fitts' Law Analysis

Individual Fitts' models constructed for each of our four crossing tasks per input type and feedback type resulted in very good fits to the data ($r^2 > 0.87$ in all cases) (Table 2). We see that the difference in y-intercepts between direct and indirect input are even more pronounced than they were for pointing tasks. Again, through comparing not only IP, but also intercept, we see the relative performance between the two input types (Figure 9).

Input, Feedback Type	Model	r^2
Indirect, Visual+Tactile	$Time = 275.9 + 169.5 ID$	0.87
Indirect, Visual	$Time = 158.4 + 194.1 ID$	0.90
Direct, Visual+Tactile	$Time = 14.9 + 180.5 ID$	0.94
Direct, Visual	$Time = -18.4 + 202.1 ID$	0.91

Table 2. Fitts' law models for the four crossing tasks, by *input type* and *feedback type*.

Experiment Two Discussion

Our results provide some interesting insights into the differences between direct and indirect input on crossing selection, as well as the impact of tactile feedback. At the highest level, our results clearly show that crossing tasks can be performed more efficiently with direct input, and that tactile feedback further enhances direct input. These results have important implications to the design of direct stylus input interfaces, such as for TabletPCs. In particular, the value of augmenting the stylus with simple tactile feedback mechanism is clearly demonstrated.

The significant interaction between *feedback type* and *input type* in terms of selection time may be best explained by target occlusion in the direct input condition. When using direct input, a participant's hand and stylus were placed directly on the co-incident tablet and display, and the stylus was typically directly on the target – the very item that is meant to provide visual feedback of a successful selection. This occlusion was not a problem in the indirect input condition, which would explain why tactile feedback had little impact in the indirect setting (Figure 9). The linking of the benefits of tactile feedback and target occlusion are further supported by the interaction between *feedback type* and *target width* (Figure 10). Participants benefited more from tactile feedback with the smaller, harder to see targets than with the larger targets that provided better and less occluded visual feedback.

The interaction between *continuity of contact* and *input type* for selection time and error rate in these crossing tasks may be due to the differences between direct and indirect input in terms of how the user tracks the cursor and the stylus tip, as discussed previously in the context of experiment one. In the case of indirect input, continuous contact with the tablet means that the stylus never moves out of range of the sensing hardware, and thus, the system is able to provide continuous cursor tracking information to the user. Further, in the direct input condition, target occlusion is less likely for discrete input where the user lifts their hand from the display as part of the task, than for continuous input where the user’s hand always rests on the display. By lifting the hand, even momentarily, one has a better chance to clearly see the next target. Target occlusion in the continuous contact condition also provides an explanation for the effects of *continuity of contact* on selection error rate.

The interaction between *target distance* and *input type* gives further evidence to the hypothesis that tracking is a major difference between direct and indirect input. For targets that are farther apart, users were more likely to lift their stylus out of range as they moved between targets. For close targets, the stylus more often remained within sensing range, giving users accurate and timely tracking information in both the direct and indirect input types.

OVERALL DISCUSSION: EXPERIMENTS ONE AND TWO

As noted in the introduction, the focus of our work was on direct vs. indirect stylus input, with and without tactile feedback, on crossing and pointing selection tasks rather than on the differences between crossing and pointing tasks per se. However, it is still worthwhile discussing the two experiments as a whole, including making some observations as to how user performance differed between the two pointing and four crossing tasks we studied and the relative impact of input type on these two task genres.

and Figure 14 show the overall mean selection times and error rates for each of the six selection techniques for both the direct and indirect input conditions. In terms of selection time, pointing and crossing selection with direct input were almost identical, with means of 1.04 vs. 1.05 seconds for pointing and crossing respectively; however, for indirect input, pointing interaction averaged about 9% lower selection times than crossing, with mean selection times of 1.12 vs. 1.22 seconds for pointing and crossing.

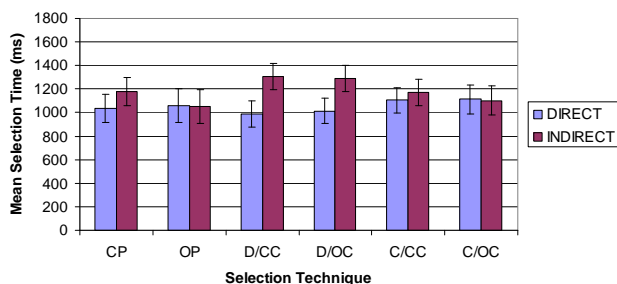


Figure 13. Mean selection times for both pointing and crossing tasks, by *technique* and *input type*.

For indirect input, the speed advantages of pointing were balanced by higher error rates of 34.2% and 20.3% for pointing and crossing respectively. For direct input, pointing and crossing were equivalent, with mean error rates of 19.6% and 19.2% for pointing and crossing respectively.

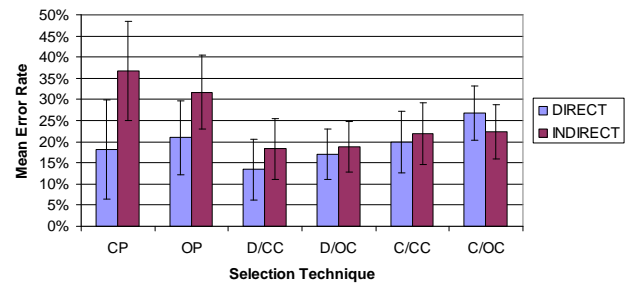


Figure 14. Mean error rates for both pointing and crossing tasks, by *technique* and *input type*.

These results are similar to those presented by Ren and Moriya [15], who found “slide touching” (i.e. crossing) to have similar selection times to pointing with a stylus in a direct setting. Unlike our study, they found that subjects committed many more errors while pointing than while “slide touching”, although differences in the method of counting errors between the studies makes this comparison difficult.

Accot and Zhai [1] reported similar error rates for indirect pointing and indirect crossing as a whole, with pointing error rates falling between discrete and continuous crossing error rates. They too found similar selection times between pointing and crossing, and suggest that other factors, such as accessibility and device form factor, might be better metrics to choose selection methods with.

It is important to recognize and account for the relatively high error rates in our experiment, particularly in the pointing tasks. Note that our experiment design required that participants successfully select each target before proceeding to the next trial. This design was chosen to ensure that participants had no incentive to “race through the experiment” by simply clicking anywhere, and as such one might expect error rates to hover around 4%, representing the normal distribution tradeoff between speed and accuracy. An examination of our data indicated two reasons for the higher than expected error rates.

First, in the pointing task, the trigger for trial completion was a “pen-down” action. The currently predominant explanation for the motor processes involved when pointing is the iterative corrections model [17], where one makes multiple corrective sub-movements towards to the target before finally hitting it. In making these sub-movements, users often inadvertently triggered a pen-down event even though they knew that they had not yet completed the trial and were simply honing in on it. Thus, many of the trials that we marked as errors could legitimately have been considered as successful trials since the additional pen-down events are simply a regular part of the target acquisition process. Note that we could have used a separate button to

indicate trial completion rather than “pen-down”, but this would have introduced a somewhat unnatural technique.

As a robustness check, we repeated all selection time analysis with the error trial selection times included, but found no important or qualitative differences from the conclusions reported. Another method for incorporating error rates into the overall analysis is to compute Fitts’ law models with the index of difficulty adjusted for error rate (i.e., resulting in an *effective ID*) [10]; however, there is significant disagreement in the literature as to the validity of such adjustments [19, 20] and we therefore chose not to use it.

The second reason for the higher error rates is the relatively small sizes of some of our targets. Accot and Zhai [1] chose target widths of 8, 16, 32, 64 and 128 pixels while ours were 4, 8, 16, and 32 pixels. The true physical dimensions (in mm) of their targets were also slightly larger than ours (Figure 15). Thus, 40% of their targets were significantly easier to select than ours. While we do not have a breakdown of their error data by target width, it is worth comparing their mean error rate to the error rates for our 32 pixel targets (i.e., their median target width). This comparison is only for indirect input, since Accot and Zhai did not study direct input. Our error rates are similar to theirs in the crossing tasks, but much higher in the pointing tasks. We hypothesize that this difference is due to users losing cursor tracking and having to reacquire the cursor in the pointing tasks where they tend to lift the pen beyond the small hover zone (< 0.5 inch in our TabletPC). Accot and Zhai used a regular Wacom tablet that likely had a larger ~1 inch hover zone which might have mitigated this issue. Also, they used a gain of 1.6 which means their users made smaller physical movements than ours, despite the visual space distances being identical. Thus, it is less likely that their users left the hover zone since smaller motor movements do not seem to incur the upward hand movement.

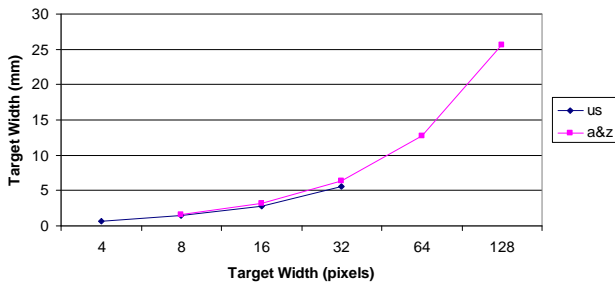


Figure 15. Pixel and actual physical target sizes for our experiment vs. Accot and Zhai [1].

CONCLUSIONS and IMPLICATIONS for DESIGN

We have presented our findings of the impact of direct and indirect input on pointing and crossing based selection. We found that the method of input, either direct or indirect, had strong effects on both pointing and crossing in terms of selection time and error rate. These effects varied by selection technique, target size, and target distance.

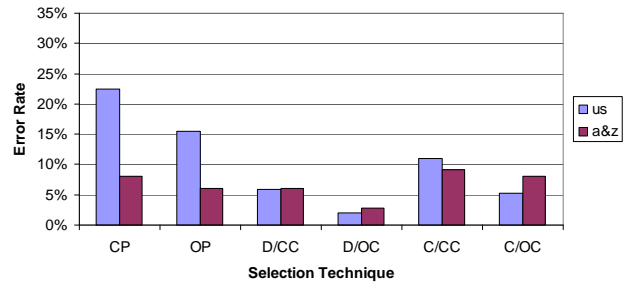


Figure 16. Error rates for our experiment and Accot and Zhai [1], for indirect input, 32 pixel targets.

We also studied the effects of tactile feedback on performance in pointing and crossing selection tasks. Our results suggest that such feedback improves selection times, most noticeably for crossing tasks when input and output are directly co-incident. One might ask how a change in the type of feedback could effect selection times since feedback is, by definition, after-the-fact. Our hypothesis is that tactile confirmation of a successful selection allows users to more quickly move onto the next target selection in the reciprocal task we studied. In other words, the benefit of tactile feedback is not that it improves selection time per se, but rather that it provides a confirmation of the selection *without the need for visual attention*. Knowing that the command was successful, users quickly move to the next target. It is important to caution that this benefit may not be as valuable for non-reciprocal selection tasks. A follow-up study should investigate this issue further.

The alternative of displaying more noticeable visual feedback, or feedback in an area of the screen that is less likely to be occluded by a hand or stylus are both possible, but drawbacks to these solutions include occlusion of other elements on the screen and the potential distraction of other users in a multi-user environment. For example, in shared-display groupware, the screen is a shared resource, so there are limits to the visual feedback that a single user’s actions can trigger. Aural feedback is an appealing alternative, but it has drawbacks including the lack of privacy and, again, the potential for the distraction of other people in the area. Sound is a shared feedback channel, and sounds meant for one user are often misinterpreted by other users in a multi-user environment [7]. Headphones would solve both of these problems, but the wearing of headphones may isolate users from one another, which is detrimental to collaboration [6]. In contrast, tactile feedback provides a largely private channel for each user in a multi-user environment. Signals meant for one user should not reach other users. Finally, tactile feedback has the added benefit that it does not interfere with other people working in the same area of the display.

Much of the previous work with haptics focuses on preventing errors, often through force feedback or through tactile feedback when the system pointer is hovering over a target. It is not immediately clear what it means to hover over a goal, so this type of haptic feedback may not be pos-

sible for crossing selection. Our tactile feedback occurred after the target was successfully selected; thus, error prevention was not a realistic goal and we did not expect, nor did we see, differences in error rate between the feedback conditions.

The idea presented by Aplitz and Guimbretiere [4] that crossing interfaces are well suited for commands in which there is a fluid transition from one action to another, such as selection from a hierarchical menu, is supported by the interaction between *target distance* and *continuity of contact*. For targets that are close together, crossing continuously through them led to lower selection times than for targets that were farther apart.

Our results indicate that indirect input relatively outperforms direct for more difficult targets. This has an important implication for the design of pen-based interfaces since we now know to avoid small targets when using direct pen input although these targets may be reasonable for indirect pen input. Thus, a pen-based interface for direct-input tabletPC's might need to be retooled if used with a separate tablet and display.

In summary, our findings suggest that the relationship between control and display space must be taken into consideration when choosing between pointing and crossing selection for pen based interfaces. Further, we have shown that tactile feedback can be beneficial in some situations, and if such feedback can be provided at reasonable cost it would be worthwhile considering its inclusion in pen based interface design.

REFERENCES

1. Accot, J. and Zhai, S. (2002). More than dotting the i's - foundations for crossing-based interfaces. *ACM CHI Conference on Human Factors in Computing Systems*. p. 73-80.
2. Akamatsu, M. and MacKenzie, S. (1996). Movement characteristics using a mouse with tactile and force feedback. *International Journal Human-Computer Studies*, 45(4). p. 483-493.
3. Akamatsu, M., MacKenzie, S., and Hasbrouc, T. (1995). A comparison of tactile, auditory, and visual feedback in a pointing task using a mouse-type device. *Ergonomics*, 38(815-827).
4. Aplitz, G. and Guimbretière, F. (2004). CrossY: a crossing-based drawing application. *ACM UIST Symposium on User Interface Software and Technology*. p. 3-12.
5. Baudisch, P. (1998). Don't click, paint! Using toggle maps to manipulate sets of toggle switches. *ACM UIST Symposium on User Interface Software and Technology*. p. 65-66.
6. Blaine, T. and Perkis, T. (2000). Jam-O-Drum: A study in interaction design. *ACM DIS Conference on Designing Interactive Systems*. p. 165-173.
7. Everitt, K., Forlines, C., Ryall, K., and Shen, C. (2005). Observations of a shared tabletop user study. Interactive poster in *ACM CSCW Conference on Computer-Supported Cooperative Work*.
8. Fitts, P.M. (1954). The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology*, 47. p. 381-391.
9. Fukumoto, M. and Sugimura, T. (2001). Active click tactile feedback for touch panels. *Extended Abstracts of the ACM CHI Conference on Human Factors in Computing Systems*. p. 121-122.
10. Lee, J., Dietz, P., Leigh, D., Yerazunis, W., and Hudson, S. (2004). Haptic pen: a tactile feedback stylus for touch screens. *ACM UIST Symposium on User Interface Software and Technology*. p. 291-294.
11. MacKenzie, S. (1992). Fitts' law as a research and design tool in human-computer interaction. *Human-Computer Interaction*, 7. p. 91-139.
12. Oakley, I., Adams, A., Brewster, S., and Gray, P. (2002). Guidelines for the design of haptic widgets. *Proceedings of the British HCI Conference*. p. 195-212.
13. Oakley, I., Brewster, S., and Gray, P. (2001). Solving multi-target haptic problems in menu interaction. *Extended Abstracts of ACM CHI Conference on Human Factors in Computing Systems*. p. 357-358.
14. Poupyrev, I. and Maruyama, S. (2003). Tactile interfaces for small touch screens. *ACM UIST Symposium on User Interface Software and Technology*. p. 217-220.
15. Poupyrev, I., Okabe, M., and Maruyama, S. (2004). Haptic feedback for pen computing: directions and strategies. *Extended Abstracts of ACM CHI Conference on Human Factors in Computing Systems*. p. 1309-1312.
16. Ren, X. and Moriya, S. (2000). Improving selection performance on pen-based systems: a study of pen-based interaction for selection tasks. *ACM Transactions on Computer-Human Interaction*, 7(3). p. 384-416.
17. Rosenbaum, D. (1991). *Human motor control*. San Diego, CA: Academic Press.
18. Schneiderman, B. (1991). Touch screens now offer compelling uses, *IEEE Software*, 8(2). p. 93-94.
19. Zhai, S. (2004). Characterizing computer input with Fitts' law parameters — The information and non-information aspects of pointing. *International Journal of Human-Computer Studies*, 61(6). p. 791-809.
20. Zhai, S., Kong, J., and Ren, X. (2004). Speed-accuracy trade-off in Fitts' law tasks — On the equivalency of actual and nominal pointing precision. *International Journal of Human-Computer Studies*, 61(6). p. 823-856.