Improving Menu Placement Strategies for Pen Input

Mark S. Hancock Kellogg S. Booth

Department of Computer Science
The University of British Columbia
2366 Main Mall, Vancouver, BC, Canada V6T 1Z4
{mhancock, ksbooth}@cs.ubc.ca

Abstract

We investigate menu selection in circular and rectangular pop-up menus using stylus-driven direct input on horizontal and vertical display surfaces. An experiment measured performance in a target acquisition task in three different conditions: direct input on a horizontal display surface, direct input on a vertical display and indirect input to a vertical display. The third condition allows comparison of direct and indirect techniques commonly used for vertical displays. The results of the study show that both left-handed and right-handed users demonstrate a consistent, but mirrored pattern of selection times that is corroborated by qualitative measures of user preference. We describe a menu placement strategy for a tabletop display that detects the handedness of the user and displays rectangular pop-up menus. This placement is based on the results of our study.

Key words: Pen-input devices, horizontal display, vertical display, direct input, indirect input, handedness, tabletop display.

1 Introduction

The emergence of computing devices such as the Tablet PC™, large-screen tabletop displays, wall displays, and personal digital assistants has increased the prevalence of direct pen input. This shift has created a need for suitable interaction styles.

With some exceptions, most applications have chosen to utilize pen-input devices simply as a replacement for mouse input. Although mice and pens (or styli) both provide two degrees of freedom for input, the form factor of each device is unique and should be considered when designing applications. An important distinction is that pen-input devices can be (and typically are) used as direct input devices (the control or hand space is the same as the display space), thus introducing an effect of occlusion by the hand holding the device. Occlusion is not present with indirect input techniques.

We describe an experimental study comparing direct pen input on horizontal and vertical surfaces, with indirect pen input on a vertical surface. These conditions model Tablet PC, tabletop, and whiteboard display techniques. The results suggest effective techniques for menu placement when user handedness and position are known. We have used this empirical data to inform the design of an adaptive menu system for pen input to interactive tabletop applications that displays pop-up menus to each user in an appropriate location according to the user's handedness and position around the table.

2 Background

Several models have been proposed to predict the movement time for target acquisition tasks that resemble menu selection. The well-known keystroke model of Card, Moran and Newell [2] suggests that target selection time is the sum of four subtasks: mental preparation, acquiring the mouse, pressing the button(s), and moving the mouse with the hand to the target. Further decomposition suggests that movement time (MT) can be predicted using Fitts Law [5], which is a function of target width W and target distance (or amplitude) A. The literature argues that such hand movement most closely follows the Shannon formulation of Fitts Law [11, 12]:

$$MT = a + b\log_2(A/W + 1)$$

where a and b are constants determined by linear regression. The logarithmic term is referred to as the index of difficulty (ID).

Boritz, Booth and Cowan [1] compared mouse-based menu selections by left- and right-handed users and found that angle of approach affected selection times. Mouse movement towards the user was slowest. For right-hand-dominant participants, movement to the right was fastest with the right hand and movement to the left was fastest with the left hand. They found no effect for left-handers.

Kurtenbach and Buxton [9] performed an experiment to analyze their marking menus. Participants were tested with both a stylus and a mouse making selections in circular marking menus. They found slower performance for selection of "off axis" menu items than for "on axis" menu items. An interaction effect suggested less performance degradation "off-axis" with the stylus.

Both neurophysiological studies [6, 7] and Fitts Law studies [5] suggest that the finer motor control achieved with the hand results in better performance than motor control with the arm. The findings of Boritz et al. are consistent with this literature, because mouse movement to the left or right is made with the hand, whereas mouse movement towards one's own body is made with the arm. This is also consistent with the findings of Kurtenbach and Buxton for on- and off-axis mouse movement. It might even explain the interaction effect they found between the axis of movement and input device, because left and right movement with a stylus does not necessarily align with the axes of the display as well as does mouse-based input.

Kurtenbach et al. [10] demonstrate the need and use for automatic handedness detection in a 2D drawing application. Their system requires that the user hold a stylus device in their dominant hand and a puck in their non-dominant hand. With their application, users tended to share the device and frequently passed control between one another. They report that, when the system did not adapt to the handedness of the users, the users would only use the pen device and not the puck. To determine handedness, they utilize the relative device positions and use this information to intelligently place pop-up palettes.

Our findings contribute to this literature by providing specific evidence that selection times in pop-up menus depend on the position and orientation of menu items relative to the point of activation of the menu, and that this effect depends in a consistently mirrored way on the handedness of the user. We describe an adaptive technique for the placement of menus that demonstrates the application of these findings for pen-based input. This system can detect and adapt to handedness using only one-handed pen input to a tabletop display. The results of the study corroborate the usefulness of the technique used both in our system as well as the two-handed version of the technique implemented by Kurtenbach et al. [10].

3 Experiment

We were primarily interested in pop-up menu selection tasks with a stylus input device on a horizontal display surface. Three different pilot studies suggested to us that movement time varies according to the position of the target relative to the point of activation of the menu. We hypothesized that occlusion of the target by the user's hand and by the stylus input device increases mental preparation time. We expected that the positional effects on acquisition times would vary with the handedness of the user and that this variation would have a mirrored pattern.

We expected the effect of occlusion to be most prevalent in conditions involving horizontal display surfaces that utilize direct input. To isolate this effect we included three combinations of input technique and display orientation: direct input onto a horizontal display surface, direct input onto a vertical display surface, and indirect horizontal input to a vertical display. We hypothesized that the positional differences in menu selection times would be greater on horizontal displays with direct input, such as tabletop displays and Tablet PCs, and virtually non-existent on vertical displays with indirect input.

Because we were interested in analyzing pop-up menu selection, we designed an experiment to measure performance relative to the "point of activation". This point is defined to be the location that the user selects to initiate the pop-up menu. The following null hypotheses were directly tested in our experiment:

- H-1 Users can acquire targets at the same speed for all target positions relative to the point of activation.
- H-2 Users can acquire targets at the same speed on both vertical and horizontal displays.
- H-3 Users can acquire targets at the same speed using both direct and indirect pen input.
- H-4 Left- and right-handed users have the same pattern of target acquisition speeds, relative to the point of activation.

As with any single experiment, one must forfeit some level of precision, realism or generalizability. The target acquisition task used in the experiment most closely resembles selection from a circular pop-up menu. We expect our results do generalize to placement of rectangular menus, because target positions in the experiment correspond both to items in a circular menu and to typical placement of rectangular pop-up menus relative to the point of activation. To achieve a higher level of precision, realistic rectangular menus were not directly tested.

We did, however, separately collect qualitative measurements of user preference in relation to placement of rectangular pop-up menus. We expected that user's would prefer menu placements that allow faster menu selections.

The experiment tested the potential utility of an adaptive interface that models the user's handedness. Adaptation provides the ability to display menus in an appropriate location depending on the handedness of each user. We expected to determine an appropriate adaptation based on the results of the experiment if our hypotheses were borne out.

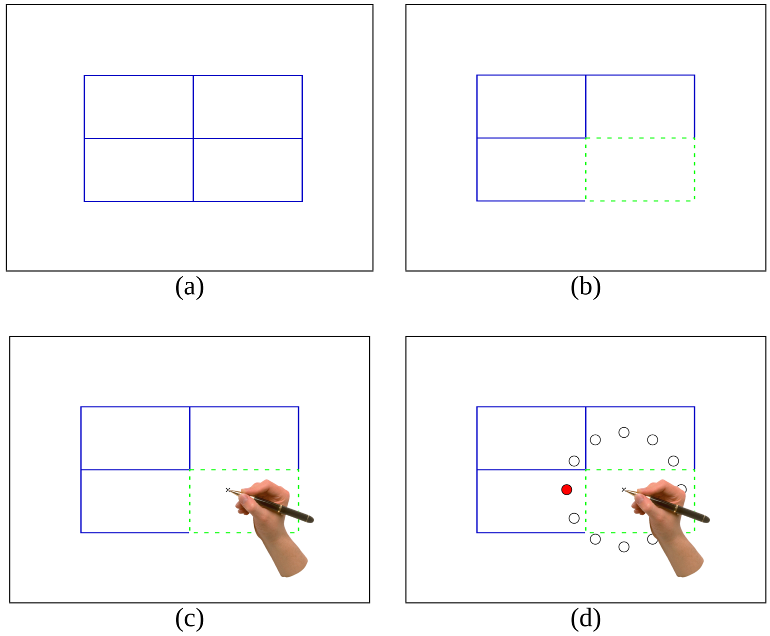

Figure 1: (a) To begin each trial, each participant was presented with four regions outlined in blue. (b) To start the trial, the outline of one region changed to a different colour (dashed in this figure, but not in the experiment). (c) The participant then would point and select that region of the display. (d) This action activated a ring of twelve circles, with the target circle in red.

3.1 Method

Participants

6 left-handed and 6 right-handed students (7 male, 5 female) between the ages of 19 and 35 ($M = 25$, $SD = 4.4$) from a local university participated in our study.

Apparatus

Participants were asked to select targets in one of three combinations of input technique and display surface. In the horizontal-direct condition, target selections were made directly on a Tablet PC with a 21 cm by 16 cm display, mounted horizontally on a table's surface. Participants were instructed to adjust the height of the seat to suit their comfort. In the vertical-direct condition, participants were asked to select targets on a touch-sensitive SmartBoard with a 141 cm by 102 cm display. Participants were told to stand directly in front of the SmartBoard at an arm's length distance, so they could comfortably reach the display. For the indirect condition, participants were asked to select targets using the blank screen of the Tablet PC as the input device, with the output only shown on the SmartBoard display. Participants were seated exactly as they were for the first condition. The SmartBoard monitor was located 173 cm from the participant. The control space to display space ratio (c:d) was thus 1:1 (by definition) for both direct input conditions and was measured to be approximately 0.15:1 in the indirect condition (targets were the same size on all displays). In all three conditions, the participant used the pen-input device provided with the Tablet PC. The Tablet PC had a 1 GHz Transmeta Crusoe 5800 processor and the SmartBoard had an IBM compatible computer with a 2.66 GHz Pentium IV processor. Both displays had a resolution of 1024 by 768 pixels. The software for the experiment was written in Java.

Procedure

To begin the experiment, participants were asked to complete a background questionnaire in order to establish experience with pen input, large-screen displays, and pop-up menus. To ensure an appropriate assignment to each handedness condition, we utilized the Edinburgh Inventory [13] to separate participants into left- and right-handed groups. In addition to this inventory, participants were also asked with which hand they used the mouse.

To begin each trial, the participant was asked to point and select one of four regions of the display, indicated by the border changing to a different colour. Upon selection, a ring of twelve circles would pop up surrounding the point of contact. One of the twelve circles appeared in red and the others in white. The participant was then asked to point and select the red circle as quickly and accurately as possible.

We used a 12 (circle position) x 4 (starting position) x 3 (display) x 2 (handedness) mixed factorial design. Each of the twelve circles was 61 mm wide and was displayed

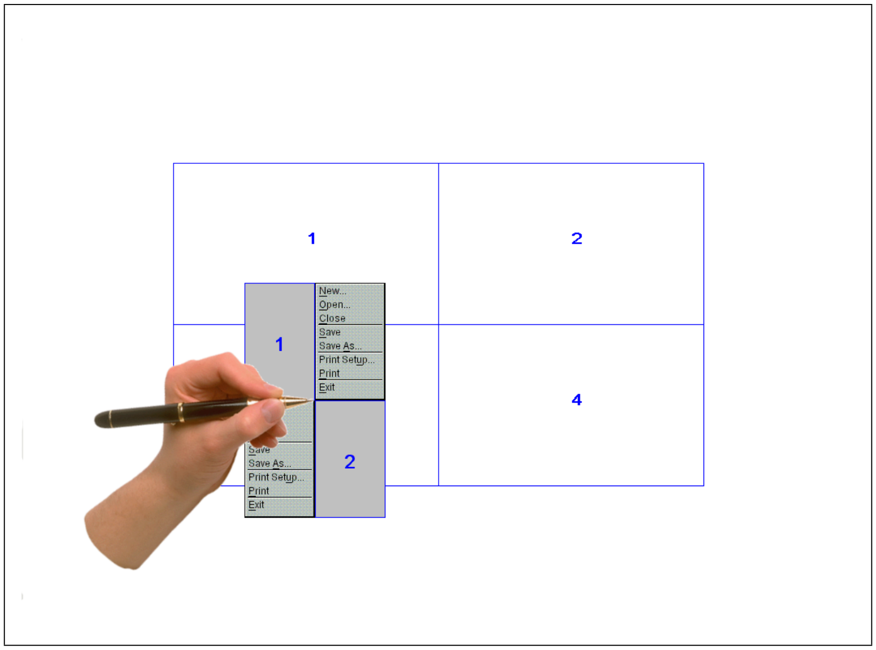

Figure 2: Users were asked to rate the four possible rectangular menu placements in order of preference in all four regions of the display. Menus were shown relative to the point of activation (the center of the four menus). The four rectangular menus were replaced by the integers 1-4 as they were clicked in decreasing order of preference.

so that its center was 350 mm from the point of contact of the stylus (see Figure 1). Participants performed the experiment in each of three different display conditions: a horizontal display with direct input (horizontal-direct), a vertical display with direct input (vertical-direct), and a vertical display with indirect horizontal input (indirect). Participants performed selections in each circle position and each starting position four times in fully randomized order in all three display conditions for a total of 576 trials per participant. The displays were presented in counter-balanced order for both left- and right-handed participants. Activation times and target acquisition times were recorded as well as the positions of each action.

To maintain a consistent index of difficulty (ID), measurements were taken in the control space of the input device. Thus the resolution was smaller for the vertical-direct condition. If the amplitude and widths of the targets are measured in the same space (display or control space), the IDs in each display condition are the same. However, if the perceived width is measured in display space, adjusting for the difference in visual angle due to distance from the screen, the ID in the direct conditions is 1.4 times larger (using the Shannon formulation) than in the indirect condition. To account for this variation, the analysis was done first with no adjustments and then with data normalized using this ratio.

After each display condition was completed, participants were asked to answer two questions about the placement of menus with the current combination of display and input technique. First, participants were shown four pop-up menu placements in each region of the screen and asked to rank them in order of preference by selecting the menus with the stylus (see Figure 2). Second, for each individual menu placement, participants were asked to state on a 5-point scale, whether or not they agreed or disagreed with the following statement:

> The menu placement is suitable for use on this display in an application.

3.2 Results

Three left-handed and one right-handed participant reported having experience with large-screen displays. All participants reported that they had never used a Tablet PC. One left-handed participant reported sometimes using a stylus input device, three left-handed participants reported rarely using a stylus and the remaining participants reported never using a stylus. All participants reported that they sometimes, often or always used pop-up menus.

Of the six left-handed participants, three reported using the mouse mostly with their left hand, one reported using the mouse mostly with his right, and two reported using the mouse only with their right hand. Of the six right-handed participants, one reported using the mouse mostly with his right hand and the remaining five reported using the mouse only with their right hand. All twelve

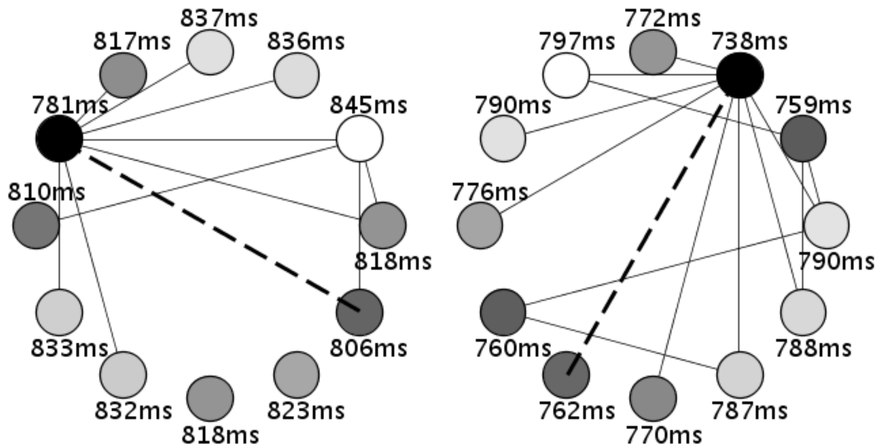

Figure 3: An interaction between position and handedness shows that left-handed participants (left) select targets more slowly in the bottom-left and top- right of the targets and faster in the top-left and bottom-right. The mirror effect occurs for right-handed participants (right). Labels represent positional means and circles are shaded linearly between the fastest (black) and slowest (white) in each condition. Significant pairwise differences ($p < .05$) are shown with a thin solid line. The thick dashed line represents the axis of expected best performance and is calculated as the line with the maximum sum of positional mean selection times, weighted by distance from the axis.

participants used the stylus input device with their dominant hand.

Two trials were removed from the data due to system error in recording selection times and two trials were removed because the end position was recorded to be more than 160 pixels from the target location. We believe that the latter two trials were system error due to the participants accidentally touching the SmartBoard with something other than the stylus device. No errors were detected in the remaining trials.

It was not possible to reliably analyze accuracy in our experiment. The recorded location of the cursor upon target acquisition is different than the actual location of the stylus tip in both of the direct conditions. This error also made it impossible to use effective throughput as the dependent measure as suggested by the ISO 9241 standard, Part 9 [4].

Before running the experiment, we decided to not include starting position as a factor in the analysis of selection times. We felt that the small distance between starting positions was not likely to produce any significant effects in selection times. Furthermore, the experimenter noted that many participants recognized that they could minimize hand movement between trials by activating the circular targets as close to the center of the screen as possible, further reducing the actual distance between starting positions (see Figure 1). Target selection times were analyzed using a full factorial Analysis of Variance (ANOVA) on the remaining three factors.

There was a main effect of display ($F(2, 20) = 76.4$, $p < .001$). Post-hoc analysis showed that a horizontal display with direct input ($M = 679$ ms, $SD = 17$ ms) was marginally faster than the vertical display with direct input ($p = .053$) and significantly faster than the vertical display with indirect input ($p < .001$) and that the vertical-direct condition ($M = 707$ ms, $SD = 22$ ms) had significantly faster ($p < .001$) selection times than the indirect condition ($M = 1007$ ms, $SD = 45$ ms). There was no significant main effect of handedness nor of target position.

A two-way interaction effect between target position and handedness ($F(11, 110) = 4.1$, $p < .001$) suggested that the effect of position depends on the handedness of the participant. Post-hoc analysis revealed pairwise differences that are shown pictorially in Figure 3. There was no significant interaction between display and handedness, nor between display and target position.

A three-way interaction ($F(22, 220) = 1.8$, $p = .017$) suggested that the two-way interaction between target position and handedness depends on the particular display

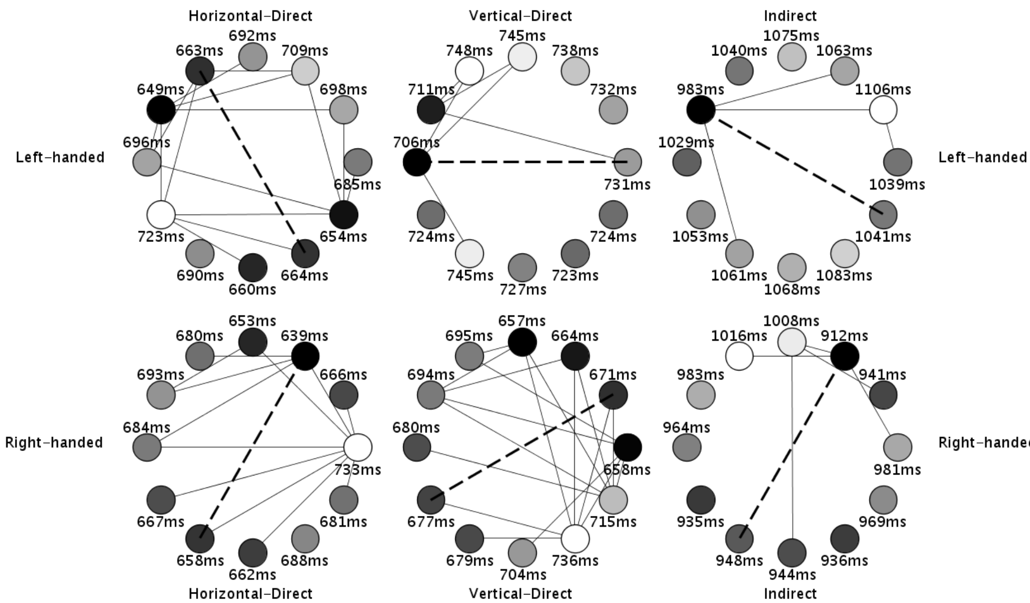

Figure 4: A three-way interaction between display condition, handedness and target position shows a different pattern of fastest selection times for all three display conditions. In the horizontal-direct condition, movement along the top-left to bottom-right axis was fastest for left-handed participants (top-left). For right-handed participants in the same condition, movement along the top-right to bottom-left axis was fastest (bottom-left). For the vertical condition, the axis of best performance is along a more horizontal axis for both left-handed participants (top-middle) and right-handed participants (bottom-middle). There are fewer significant differences in the positional means for the indirect condition. The lines and shading are as in Figure 3.

condition. Post-hoc analysis revealed more significant differences in the horizontal-direct condition and fewer significant differences using indirect input (see Figure 4).

Adjustment for Perceived Target Width

To account for the discrepancy in perceived target width between display conditions, the dependent measure of throughput was used with normalized indices of difficulty. The ANOVA was then rerun on this normalized data. The results of this factorial ANOVA resulted in the same main effects and interactions.

User Preference

To analyze order of preference for menu placement, we performed a Kendall's W test for each combination of handedness, display condition and region of the display. Left-handed participants had consistent preference ratings in all four regions of the display in both the horizontal-direct condition and the vertical-direct condition ($\chi^2(3,6) > 7.0$, $W > .50$, $p < .05$). For right-handed participants, the only consistent preference was found in the horizontal-direct condition in the top-right of the display ($\chi^2(3,6) > 11.6$, $W = .644$, $p = .009$). Preferences for the indirect condition were not significantly consistent. A Wilcoxon Signed Ranks post-hoc test was performed to analyze the pairwise preference orderings.

In both the horizontal-direct and vertical-direct conditions in all four regions of the displays, left-handed participants rated the bottom-left menu placement as significantly less preferred than the bottom-right and top-right menu placements ($p < .03$), with two exceptions. In the top-left of the Tablet screen, the bottom-left and bottom-right placements were not statistically different ($p = .057$), and in the top-right of the Tablet screen, the bottom-left and top-right placements were not significantly different ($p = .056$). In the bottom half of the vertical-direct display, these participants also rated the bottom-left menu placement as significantly less preferable than the top-left menu placement ($p = .026$ and $p = .024$).

Right-handed participants consistently preferred the bottom-left menu placement to the top-left placement ($p = .026$), the top-right placement ($p = .024$) and the bottom-right placement ($p = .023$) only in the top-right region of the horizontal display.

Suitability of Menu Placement

To analyze suitability of menu placement, a series of Kruskal-Wallis tests was used. Results showed signif-

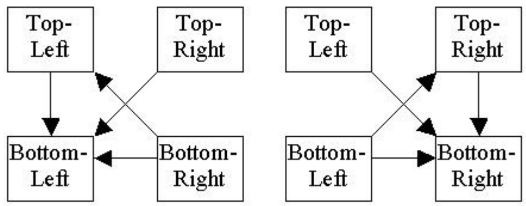

Figure 5: Results of suitability ratings suggest a pattern of ratings for left-handed participants (left) that is mirrored for right- handed participants (right). Arrows indicate the "is more suitable than" relation.

icant differences for both left- and right-handed participants in the ranking of menu placements in both the horizontal-direct condition and the vertical-direct condition in all regions of the display ($\chi^2(3,6) > 9.0$, $p < .05$). In the indirect condition, the only significant difference in rankings was found for right-handed users in the top-right of the display ($\chi^2(3,6) = 9.779$, $p = .021$).

A post-hoc Mann-Whitney test was used to determine pairwise differences in suitability ratings. In both direct input display conditions and all regions, left-handed participants tended to follow a consistent pattern of significant pairwise differences in suitability ratings ($p < .05$) that is mirrored for right-handed participants (see Figure 5). There were three additional significant pairwise differences in suitability ratings for right-handed participants. In the horizontal-direct condition, they rated the bottom-left menu placement higher than the top-left menu placement in the top-right of the display ($p = .021$) and the top-left menu placement higher than the top-right menu placement in the bottom-right of the display ($p = .036$). In the vertical-direct condition, right-handed participants rated the top-left menu placement higher than the bottom-left placement in the bottom-right region of the display ($p = .027$).

3.3 Discussion

The differences in display conditions suggest that null hypotheses H-1, H-2, and H-3 are false. Slower selection times for the indirect condition suggest that users have less difficulty selecting targets with a stylus when they interact directly with the display. On the horizontal display, participants were observed resting their hand on the display during the trials. This resting position likely reduced fatigue and increased ability to acquire targets utilizing wrist movement instead of arm movement. This beneficial hand position may explain the smaller selection times for horizontal displays than for vertical ones. Although direct input appears to be a faster method for selection, direct input has the disadvantage of occlusion that does not exist with indirect input. This occlusion affects many aspects of the user interface besides menu placement and should be considered carefully when choosing between these two methods. The conclusion that can be drawn from the results of this experiment are specific to target acquisition and may not generalize sufficiently to inform the choice of input device when designing applications. However, when the choice to use direct pen input has already been made, fatigue effects and inhibitory arm movement will likely occur more frequently on a vertical display than on a horizontal one.

The interactions involving handedness suggest that H-4 is also false. These interaction effects were predicted by our hypotheses and give clear suggestions for optimal placement of menus and menu items relative to the handedness of the user. These suggestions are also consistent with user preference and suitability ratings, with a few minor exceptions.

The results of this experiment show that there is a discrepancy between left- and right-handed users about the fastest target location relative to the point of activation. Left-handed users clearly are faster in the upper-left and lower-right quadrants and right-handed users clearly are faster in the upper-right and lower-left quadrants. Despite subtle differences between display conditions, this effect is consistent for selection times in all display conditions, and for user preference and suitability ratings. This effect is to be expected, since the faster quadrants require only left to right movement of the hand which utilizes a faster muscle group than do forward and backward arm motion in the respectively opposite two quadrants. For direct input devices, there is also an additional effect of occlusion. Targets appear occluded when underneath the user's hand, and so the time to acquire the targets in these positions is increased. The hand and stylus occlude the display the most on the horizontal display with direct input and least on the vertical display with indirect input, which may explain the predominance of the positional differences in the former and their absence in the latter.

Despite the decreased effect in both vertical display conditions, the optimal menu placement strategy suggested by the results of the experiment provides fast menu selection times in all three conditions. By providing this same strategy on all three displays, designers can account for handedness effects and still provide a consistent interface for all display devices that utilize pen input.

4 Appropriate Menu Placement

There are two potential methods to compensate for the discrepancy between stylus-driven menu selection performances for left- and right-handed users in the design of applications. One method is to provide an adaptable display that allows the user to choose the appropriate placement according to their preference. Our experiment shows that user preference is consistent with better performance, which demonstrates the viability of such an adaptable interface. The second method is to automatically adapt the display to respond to the handedness of each user. This method requires a model of the user that includes handedness.

Some software for the Tablet PC already include an option to specify the handedness of the user. There are several disadvantages to this approach. In our experiment, all six of the left-handed participants reported using the mouse only with the left hand, three of which reported using the mouse with the right hand *more* frequently than with the left. This result is evidence that users have a tendency to not alter this particular default setting. Furthermore, in co-located, collaborative applications, control of the input device is frequently passed between several users, some of whom may differ in handedness. In this environment, the need to specify one's handedness explicitly becomes too great of an overhead for the user to benefit from any advantage the system might provide.

It may be possible to improve the method of explicitly specifying handedness for pen-input devices that are typically used by only one person (or very few people), such as the Tablet PC or a Wacom digitizing tablet. The results of the experiment demonstrate that such and option is a minimum requirement for such applications. In collaborative environments, however, this minimum requirement is no longer sufficient. Kurtenbach et al. [10] demonstrate a method of automatically determining handedness for a particular collaborative application that utilizes two-handed input where one hand is used for stylus input. We add to this work by demonstrating a technique for automatically determining handedness for one-handed pen input to a collaborative application on a large-screen tabletop display.

In order to determine the user's handedness, we create a model of each user that includes the position and orientation of the user's stylus input device, the side of the table at which the user is sitting and the handedness of the user. We tested three different methods of obtaining this user model.

Our system is intended to demonstrate the feasibility of an adaptive solution to this problem and does not purport to be the only method of obtaining handedness and orientation information. Other solutions, such as the DiamondTouch [3] or computer vision could also be used to obtain the necessary information about users at a tabletop display. Our intention is only to demonstrate how the results of the experiment can inform the design of tabletop display systems that utilize pen input, not to promote a particular sensing technique.

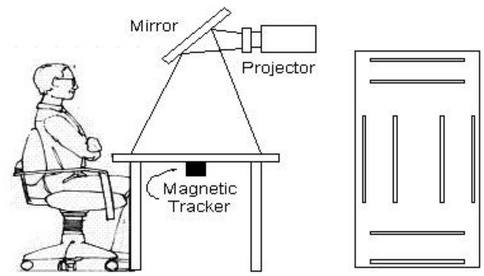

Figure 6: The tabletop display is projected onto the surface from above with the Fastrak cube placed underneath the table (left). Force sensing resistors (seen from the top) are distributed evenly across the table's surface (right).

4.1 Sample Map Application

To demonstrate the use of our technology, we created a sample tabletop display application. This application displays a map of the world containing information about individual countries. To display this information, the user must tap on a country and choose one of six items in a rectangular pop-up menu. This menu appears down and to the left for right-handed users and down and to the right for left-handed users. The menu is oriented to face the user, regardless of the side of the table from which it has been activated.

4.2 System Description

The position and orientation of the stylus are obtained from a Polhemus Fastrak. The table's surface has been modified with Force Sensing Resistors (FSR) to determine the side of the table of the user. This combined information provides the input to the user model from which handedness can be determined. We describe three potential methods of obtaining this model and compare the accuracy of each.

The computer display is projected from above onto a 150 cm by 80 cm white laminate surface at a resolution of 1024 by 768 pixels. The magnetic tracker cube is placed underneath and at the center of the table in order to minimize the distance from the stylus to the tracker and thus maximize accuracy. Eight 61.0 cm x 1.5 cm x 0.5 cm FSR strips are placed on the surface of the table. The table is then covered with white poster board so as not to interfere with the projected image (see Figure 6).

Model 1: Simple Heuristics

To demonstrate the need for a slightly more complex system, the first model that we explored is simplistic in nature. This model determines the handedness of the user based solely on the azimuth angle of the stylus input device. Given the side of the table at which the user is sitting, if the azimuth angle is between $0°$ and $180°$ relative to this location, the system predicts that the user is right-handed, and if the angle is between $180°$ and $360°$, the system predicts that the user is left-handed.

Model 2: Neural Network

The second model utilizes a feed forward neural network, trained using back propagation to determine the handedness of the user. The input layer has a node for each of the six degrees of freedom, and a node for the side of the table of the user. The output layer has a single node to represent the user's handedness. The hidden layer has five nodes. For our particular network, a learning rate of 0.5 is used. The training corpus is passed through the network 100 times.

One disadvantage of the neural network is that it does not utilize a priori knowledge about the environment. Thus, the network relies on training to infer the appropriate relationship between input device and the user model. It is therefore more difficult to extend the model to include contextual information or other input devices.

Model 3: Bayesian Network

The third approach models the tabletop display environment using a Bayesian network trained using a naive algorithm. This network consists of 10 variables, six continuous and four discrete. In contrast to the neural network, only the x- and y-coordinates and the azimuth angle are modeled directly. The elevation angle is used indirectly to vary the probability distribution for the actual azimuth angle. The network models both the measured and actual value for each of these degrees of freedom. Two discrete variables are used to model handedness and location of the user. Variable elimination is used to determine the most probable value for each of these nodes.

The Bayesian network has several advantages over the neural network. First, the Bayesian network only requires four of the six degrees of freedom as input. Second, the Bayesian network can more easily be extended to include input from other sources of information such as a video camera. Finally, contextual information, such as the locations of interface components, can more easily be added to the Bayesian network.

4.3 Model Comparison

To test the accuracy of the three models, data were collected from 2 computer science graduate students (1 left-handed and 1 right-handed) using the tabletop display. These users include one of the authors. Users were asked to use the sample map application on all four sides of the table. The 10-fold cross-validation technique [14] was used to separate the training corpus from the test corpus for the neural and Bayesian networks. Accuracy measures for each model are given as an average percentage of correctly classified results.

Accuracy Results

The Bayesian network correctly predicted the handedness of the user with the highest accuracy ($M = 100.0\%$, $SD = 0.0\%$), followed by the neural network ($M = 99.9\%$, $SD = 0.2\%$), and the simple heuristics had the lowest accuracy ($M = 97.6\%$).

5 Conclusion

The results of our experiment clearly show that handedness issues are of paramount importance for applications that utilize pen input. Our findings illustrate that a single static interface for such devices will neccessarily disadvantage either left-handed or right-handed users. We demonstrate a dynamic technique to improve pop-up menu selection for one-handed pen input in a collaborative application, without the need for explicitly specifying each user's handedness.

6 Future Work

We describe some of the disadvantages of an adaptable interface solution to the issue of handedness for pen input. Future work will involve an attempt to improve the method of adaptation and to compare existing methods to both this improved adaptable method and the automatic method described in this paper.

At the suggestion of one of the anonymous reviewers, we are planning to investigate finger-based target selection. The hand posture when using a pen to point and select is different than when using the finger. This change in posture may result in varied selection times for targets relative to the point of contact of the finger. Further experimentation is required to study these differences.

7 Acknowledgments

We would like to thank Kori Inkpen, Regan Mandryk, and Colin Swindells for help on pilot experimentation leading up to this work, Bill Buxton for his useful suggestions on drafts of the paper and on the experimental design, and the anonymous reviewers for their helpful comments. Financial and infrastructure support for this research was provided by the Natural Sciences and Engineering Research Council of Canada (NSERC), the Canada Foundation for Innovation (CFI), and the New Media Innovation Centre in Vancouver, British Columbia (NewMIC).

References

[1] J. Boritz, K. S. Booth, and W. B. Cowan. Fitts's law studies of directional mouse movement. In *Graphics Interface '91*, pages 216–223, 1991.

[2] S. K. Card, T.P. Moran, and A. Newell. The keystroke-level model for user performance time

with interactive systems. *Communications of the ACM*, 23(7):396–410, 1980.

[3] P. Deitz and D. Leigh. DiamondTouch: a multi-user touch technology. In *Proceedings of UIST*, pages 219–226. ACM Press, 2001.

[4] S. A. Douglas, A. E. Kirkpatrick, and I. S. MacKenzie. Testing pointing device performance and user assessment with the ISO 9241, Part 9 standard. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 215–222. ACM Press, 1999.

[5] P. M. Fitts. The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology*, 47(6):381–391, 1954.

[6] C.B. Gibbs. Controller design: interactions of controlling limbs, time-lags and gains in positional and velocity systems. *Ergonomics*, 5(2):385–402, 1962.

[7] M. Hammerton and A. H. Tickner. An investigation into the comparative suitability of forearm, hand and thumb controls in acquisition tasks. *Ergonomics*, 9(2):125–130, 1966.

[8] M. S. Hancock. Improving menu placement strategies for pen input. Master's thesis, The University of British Columbia, Vancouver, BC, Canada, May 2004.

[9] G. Kurtenbach and W. Buxton. The limits of expert performance using hierarchic marking menus. In *Proceedings of INTERCHI*, pages 482–487. ACM Press, 1993.

[10] G. Kurtenbach, G. Fitzmaurice, T. Baudel, and W. Buxton. The design of a GUI paradigm based on tablets, two-hands, and transparency. In *Proceedings of CHI*, pages 35–42. ACM Press, 1997.

[11] I. S. MacKenzie. A note on the information-theoretic basis for Fitts' Law. *Journal of Motor Behavior*, 21:323–330, 1989.

[12] I. S. Mackenzie. Movement time prediction in human-computer interfaces. In R. M. Baecker, W. A. S. Buxton, J. Grudin, and S. Greenberg, editors, *Readings in Human-Computer Interaction*. Kaufmann, second edition, 1995.

[13] R. C. Oldfield. The assessment and analysis of handedness: the Edinburgh Inventory. *Neuropsychologia*, 9:97–113, 1971.

[14] M. Stone. Cross-validatory choice and assessment of statistical predictors. *Journal of the Royal Statistical Society*, B36:111–147, 1974.