Improving Menu Placement Strategies for Pen Input

by

Mark S. Hancock

B.Sc., Simon Fraser University, 2002

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF

THE REQUIREMENTS FOR THE DEGREE OF

Master of Science

in

THE FACULTY OF GRADUATE STUDIES

(Department of Computer Science)

We accept this thesis as conforming
to the required standard

The University of British Columbia

April 2004

© Mark S. Hancock, 2004

Abstract

Pen-based interaction is becoming a commonplace two degree-of-freedom alternative to the mouse. The use of pen input allows users to acquire targets directly on a computer display. This style of interaction introduces a unique form factor and a new set of considerations in the design of applications for such devices. This thesis presents a series of experiments designed to evaluate the use of pen-input devices on a variety of display setups. In particular, user performance is investigated in terms of menu selections in circular and rectangular pop-up menus using stylus-driven direct input on horizontal and vertical display surfaces. These studies help to clarify effects of hand posture and hand preference. The results of these studies 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. This pattern is different for both vertical and horizontal displays due to a change in hand posture. Implementation details are provided for an automatic menu placement strategy for a tabletop display. Details are presented on how to detect which hand is being used to hold the device and on how to apply the results of the study to display rectangular pop-up menus in a co-located collaborative environment.

Contents

List of Tables

List of Figures

Acknowledgements

I would like to thank my supervisory committee Dr. Kellogg Booth, Dr. Kori Inkpen, and Dr. Joanna McGrenere for their guidance. I also thank Colin Swindells and Regan Mandryk for their help both in designing and running the experiments, Bill Buxton for his helpful comments on the design of my final experiment, and Dr. Brian Fisher for his advice on experimental analysis. I also thank the members of the Imager Lab and the Interaction Design Reading Group (IDRG). Financial and infrastructure support for this research was provided by the Natural Sciences and Engineering Research Council of Canada (NSERC), the BC Advanced Systems Institute (ASI), the Canada Foundation for Innovation (CFI), and the New Media Innovation Centre in Vancouver, British Columbia (NewMIC).

Mark S. Hancock

The University of British Columbia
April 2004

Chapter 1

Introduction

1.1 Motivation

With the emergence of computing devices such as the Tablet PC, large-screen tabletop displays, wall displays, and personal digital assistants, direct pen input is becoming more prevalent in day-to-day computing. With this shift comes the 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, thus introducing an effect of occlusion by the hand holding the device that is not present with indirect mouse-based input. Furthermore, the posture of the hand when using a pen is significantly different than the posture when using a mouse.

This thesis focuses on the use of pen-input devices both directly and indirectly onto horizontal and vertical display surfaces. Specifically, this thesis addresses limitations of pen input due to occlusion and the effect of handedness in terms of context menu selection.

1.2 Terminology

The following terms are used throughout the thesis. Although they have multiple meanings in the literature, we intend their meanings to be as follows unless explicitly stated otherwise.

context menu - Menus associated with objects on the display that appear near the object when the object is selected are a well known interaction technique. Such menus are also referred to as pop-up menus. The term "context menu" is a reminder that the particular menu that pops up usually depends on the object that is selected. A default background object provides the context when no object is selected.

direct vs. indirect interaction - We distinguish between interaction techniques in which the control space where the hand is moving and the display space where visual feedback is provided are superimposed (the direct case) or separated (the indirect case). This distinction is finer than that captured by the higher-level notion of direct manipulation, which can include both direct and indirection interaction.

point of activation - The location on the display where a context menu is initiated is the *point of activation*. For direct interaction this is the physical location of the input device (pen or finger). For indirect interaction this is the physical location of the cursor or other visual feedback.

occlusion - We refer to the obscuration (for rear-projected displays) or distortion (for front-projected displays) of the salient portion of the display caused by the presence of the user's hand or the input device itself as *occlusion*.

handedness - The literature contains references to concepts such as the preferred and non-preferred hand, the dominant and non-dominant hand, and left- and right-handedness. We use each of the terms when discussing the relevant literature, but we will operationally define *handedness* to mean simply the hand with which the user is holding the input device.

hand posture - We use this term in a relatively informal manner to capture a variety of different parameters related to the position, orientation, and relative tension and flexion in the muscles of the hand during an interaction.

target acquisition - We define *target acquisition* to be the movement from a starting position to a final (target) position on a display in either a direct or indirect interaction. This is often one component of a larger interaction task.

menu selection - *Menu selection* is a compound task in which the first target acquisition initiates a context menu and a second target acquisition selects an item in the menu.

1.3 Menu Selection

In the use of pen-input technology, there are a variety of interactions that may be affected by a change in hand posture and by occlusion of the hand. In this thesis, the target acquisition task involved in selection from a pop-up menu was chosen as a representative task for a variety of reasons. Target acquisition is pervasive in computer applications. Studying one form of target acquisition can help to predict user performance in a wide variety of tasks. Context menus, in particular, are one of the most common uses of two-degree-of-freedom input. In most operating systems, including Microsoft® Windows, Mac OS X™ and X Window System™, on-screen objects are enhanced with a context menu that appears at the point of the cursor by clicking with the mouse on the object. A similar interaction technique is used on the Tablet PC™ for pen input by pressing the pen into the tablet. The findings of this thesis can be used to inform both the design of pop-up menus and the design of other interaction techniques involving target acquisition, such as movement toward a button, or selection of objects on the screen.

By studying user performance in this specific context, it is possible to determine effects that are particular to pen use. Effects that may occur due to the difficulty of the task can be eliminated. We performed a series of three experiments that show effects of cognition due to task difficulty can pose a significant problem in isolating performance differences.

1.4 Collaboration

A typical use of pen input is for large-screen displays, such as a tabletop display or an electronic whiteboard. Because of their large size, these displays can more easily be used by groups of people working in collaboration with each other. Collaboration adds another element to the interaction that can complicate the design of applications for these displays. For example, users of these displays typically share one or more of the available input devices. The sharing of devices may lead to a change in hand posture that would be difficult for the system to predict. If one of the users has a different handedness than another, this change in hand posture could be drastic.

In an application designed for a single user, the interface can allow customization to account for these individual differences, but customization may not be feasible in a co-located collaborative environment with a shared display because of the overhead required to set and reset customized parameters.

This thesis addresses the issue of pen sharing and describes how to use knowledge of hand posture, even in dynamically changing collaborative environments to improve user performance in menu selection and other tasks. Many of the results also apply to single-user applications.

1.5 Thesis Outline

Chapter 2 surveys previously published literature on target acquisition, menu design and horizontal display surfaces. Chapter 3 describes three user studies used to determine effects of handedness and occlusion on horizontal tabletop displays. Additionally, the description of each study includes a description of the experimental procedure, analysis, a discussion of the results, the lessons learned and a summary of the knowledge gained from the set of three studies. Chapter 4 explains how to apply the knowledge gained from the experiments in a realistic tabletop display application. It describes an appropriate menu placement strategy suitable for both left-handed and right-handed users and describes the details of its implementation. Concluding remarks and suggestions for future work are discussed in Chapter 5.

Chapter 2

Related Literature

2.1 Target Acquisition

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 [6] 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 [9], relating 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 [31, 33, 32]:

$$MT = a + b\log_2\left(\frac{A}{W} + 1\right) \tag{2.1}$$

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

Welford introduces another model to predict movement time that splits the pointing task into two distinct movements: an initial ballistic impulse movement towards the target followed by a controlled deceleration when homing in on the target [45]. This separation leads to a decomposition of Fitts law described by the following formula:

$$MT = a + b_1 \log_2 A - b_2 \log_2 W \tag{2.2}$$

Graham and MacKenzie [18] compared a display-space analysis to a control-space analysis for a virtual pointing task with the finger using the second formulation. They found the analysis of hand movement (control-space) to be much cleaner and showed separable effects of amplitude and width. This analysis revealed that target width had less of an effect than amplitude on movement time. They hypothesized that this difference is due to the need for a cognitive strategy in the use of vision in virtual pointing, compared to the use of the visuomotor mechanism in physical pointing tasks. In testing the effects of gain for indirect mouse input, Graham [17] also found the movement plan in the initial phase to be relative to the distance moved by the hand, not by the visually perceived distance. Graham and MacKenzie [19] also compared physical and virtual pointing with the finger using a similar analysis. Their results showed no difference in the initial phase of movement, but did show that movement time is more sensitive in virtual pointing to the accuracy constraint of small targets in the second phase of movement.

2.1.1 Differences in Left and Right Hand Movement

It has been suggested [40, 47] that the preferred hand is superior to the non-preferred hand at controlled movements. Flowers [12] applied this theory to both a "ballistic" rhythmic tapping task and a "controlled" variation of Fitts's reciprocal tapping task. His results showed no difference in performance for the ballistic task, but did show that the preferred hand performed better in the controlled task.

An alternate explanation for differences in preferred and non-preferred hands is that the variability in motor output is greater in the non-preferred hand [1, 38]. In pointing tasks, this "noisy" output could result in either a larger difference in movement time or a larger deviation in error for the non-preferred hand as the amplitude of the target gets larger.

A third explanation for these differences is the lateral asymmetry in the left and right brain. Todor and Doane [43] consider the decomposition of movement time that was suggested by Welford [45] in Equation 2.2 and combine it with the theory that the right-hand (controlled by the left-hemisphere of the brain) is superior at sequential processing and the left-hand (controlled by the right-hemisphere) is superior at non-adaptive parallel processing. They hypothesize that the left hand should perform better in the initial distance-covering phase whereas the right hand should perform better in the controlled homing phase and that scaling the amplitude and width, while maintaining a constant ID, should increase the movement time for the right hand and decrease it for the left.

Kabbash et al. [25] corroborate the theory suggested by Todor and Doane [43] in an experiment measuring differences in preferred and non-preferred hands for three indirect pointing devices. They show that for 24 right-handed participants, the right hand was advantaged by larger target widths and the left hand was advantaged by smaller amplitudes, regardless of the index of difficulty.

Boritz, Booth and Cowan [4] compared mouse-based menu selections by left- and right-handed users and found that the 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 [28] performed an experiment to analyze their marking menus. Participants were tested with both a stylus and a mouse while 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 [15, 22] and Fitts law studies [8] 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 or away from 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.

Balakrishnan and I.S. MacKenzie [2] have shown apparent inconsistencies with the assumption that hand movements will be faster than arm movements. They demonstrate that certain combinations of movement can in fact be slower with the hand, such as unsupported index finger movement. They offer several explanations for differences in movement for the various limb segments. In particular, they suggest that the type and range of movement will determine which limb segments will have the best performance. They add that the type of task (serial versus discrete) and the control-display gain can also dictate the performance of each limb.

Kurtenbach et al. [30] 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 the dominant hand and a puck in the non-dominant hand. They utilize the relative device positions to determine the user's handedness. They report that without handedness adaptation, because users would typically take turns, they would only use the pen device and not the puck so as to avoid collisions of the two hands (or the pen and the puck). They use this handedness information to intelligently place pop-up palettes.

2.1.2 Contribution of the Research

The findings reported in this thesis contribute to the 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. An adaptive technique for the placement of menus that demonstrates the usefulness of these findings for pen-based input is described. The system can detect and adapt to handedness using one-handed pen-input to a tabletop display. Our findings serve to corroborate the use of such an adaptation, employed previously by Kurtenbach et al. [30], and in our work for tabletop displays.

2.2 Context Menus

Many different pop-up menuing systems have been proposed for pen-based input techniques. Each technique promotes a method to improve user performance or fluidity of movement so as to enhance pen-based interaction. This section includes a discussion of the various strategies and a description of how our work can provide further insight in the design and evaluation of these and future menuing systems.

2.2.1 Circular Menu Designs to Improve User Performance

Circular menus allow for faster performance than linear menus, as predicted by Fitts law [9]. In circular menus, each item is placed so as to minimize target distance and maximize target width. Callahan et al. [5] performed an experiment comparing pie menu selection to linear menu selection using a mouse and revealed slower seek times for menu items that are further down in linear menus. They claim that seek time in pie menus is "fairly constant". However, in their graph of target location vs. seek time for pie menus with eight items, it appears as though two opposite octants are fastest and that seek time may decrease as the target menu item grows farther in angular distance from these two octants. Callahan et al. give no clear indication of which target location corresponds to which menu item location, and their analysis does not include position as a factor, so the significance of this observation cannot be determined from their discussion. They were successful in proving their hypothesis that pie menus do in fact provide faster performance, but they leave unexplored the importance of menu position, hand posture, and handedness.

The most common form factor for context menus is rectangular, with a linear list of items that perhaps includes items that lead to further sub-menus, which are again linear lists of items. Some early computer-aided design (CAD) and animation systems employed two-dimensional rectangular menus. An article published in 1991 in Dr. Dobbs Journal makes the claim that pie menus were invented in 1986 at the University of Maryland [24]. The Dr. Dobbs Journal article and the CHI '86 conference paper by Callahan et al. [5] that analyzed that implementation is often cited as the first reference to circular menus in the literature. It seems likely that pop-up circular menus had been used prior to that even though they were not described in the literature. For example, Forsey implemented a system that included pie menus in the early 1980s, although the publications describing his system focussed on the multi-threaded system architecture and did not explicitly mention pie menus [13].

Marking Menus [29, 28], Control Menus [36], and FlowMenus [21, 20] also capitalize on the improved performance of a circular layout as predicted by Fitts law. These techniques each improve upon regular pie menu selection. These three techniques are designed and tested for use with pen input. Hierarchical Marking Menus are intended to optimize expert performance by making the menu visible only after a short delay (0.3 seconds). Users can perform the stroke without the menu appearing, if they know the direction in which to move. Control Menus are similar to Marking Menus except that an action begins after moving the "activation distance" from the centre of the menu, without lifting the pen. This technique allows the operation to be both selected and controlled in the same continuous interaction. In FlowMenus, each selection is completed by returning to the centre portion of the menu. This technique allows hierarchical selections to occur within the screen space of a single menu and remain continuous. In the evaluation of both Marking Menus and FlowMenus, effects of menu item position were observed. Neither analysis includes the position of the target menu item as a factor in the analysis.

In the discussion of FlowMenus [21], Grumbretière and Winograd raise the issue of occlusion. They recommend avoiding the use of the occluded octant (bottom-right for right-handers, bottom-left for left-handers) or placing items in this location that are complementary to the opposite octant. They do not suggest how the user should indicate their handedness nor do they discuss any method to determine handedness automatically.

T-Cube [44] is an alternative text-entry method developed by Venolia and

Neiberg that uses pie-shaped menus. The user initiates the text entry by pressing the pen in one of eight octants plus the centre of the T-Cube widget and "flicking" toward one of eight directions. As with Marking Menus, a pie menu will appear after 1/3 of a second with the possible letters as the menu items, however, the menus are displayed offset from the point of activation in order to avoid occlusion by the hand. The menu appears offset to the left for right-handers and offset to the right for left-handers. The user must then perform an indirect interaction to select the menu item by moving in the direction indicated by the visually separated menu, not toward the visible menu item. Venolia and Neiberg do not discuss how to determine the handedness of the user, only that the interaction will be different for each group.

Bullseye menus [14] and Tracking Menus [10] both place each menu item in a separate concentric ring of its own. This design does not capitalize on Fitts law in the same way as a pie menu, but it does eliminate angular positional differences. Each menu item can be acquired by movement in any direction, so these menus can be used with similar or mirrored movements for both left-handed and right-handed users. That is, the menu selection movement requires no angular movement, only radial movement.

2.2.2 Contribution of the Research

In all of the mentioned context menu designs, only a few consider the importance of the position of the menu items. In some cases, the effect of occlusion by the hand is discussed for direct pen-input devices, but the actual movement required to make a selection is not considered. Our experiments demonstrate that the consideration of hand and arm movement will help to improve the design of these types of menus.

Furthermore, our analysis in terms of physical movement reveals a clearer picture of handedness asymmetries.

The results of our experiments provide insights that can be used to inform appropriate placement and arrangement of the menus that have been suggested for pen interfaces. The use of our experimental results is demonstrated in a sample map application for a tabletop display using pen input. This application demonstrates an appropriate placement strategy for rectangular pop-up menus.

2.3 Horizontal Display Surfaces

The application of the findings in this thesis are demonstrated through the design of a tabletop display application. Large, horizontal display surfaces can allow many users to more easily work together on a shared display. In order to support such collaborative work, application designers must consider certain aspects that are unique to tabletop displays. Tabletop displays differ from typical desktop applications in a variety of ways. Their use makes interaction techniques such as the mouse less practical and thus are typically controlled with alternate methods such as direct pen input. Tabletop displays allow users to view the screen from all sides of the table and thus introduce the issue of orientation of objects on the screen, including widgets such as context menus. The collaborative nature of tabletop displays also illustrates the need to support a variety of users simultaneously, who frequently exchange control. This environment poses a uniquely challenging medium in which to apply handedness considerations.

This section includes a brief overview of some of the literature on horizontal displays and their associated interaction techniques.

2.3.1 Direct Input on Horizontal Displays

The horizontal nature of a tabletop display lends itself to having people seated around it. This setup allows the display to be in physical reach of all of the collaborators as well as providing for superimposition of display space on input space. Display proximity and direct input can both contribute to increased explicit and consequential communication between collaborators [34]. Collaborators are much more likely to see their partner's action when using stylus input, or a touch sensitive display as opposed to small mouse cursors that are being controlled indirectly [34]. However, a drawback of direct input is that users obscure information when interacting with a display, both for others and for themselves. This is exacerbated with horizontal displays because users tend to rest their arms and hands on the table.

In addition, the arms and hands of users of top-projected displays may cast shadows onto the tabletop when information projects onto their limbs. Given contextual information about the users, some components could be altered to appear in an accessible location. In Wellner's DigitalDesk [46] system, the translation application assumed that users were right-handed and presented digital information to the right of them for ease of interaction. This inconvenienced left-handed users in a number of ways. Wellner suggested using an overhead camera to detect the dominant hand and then present information accordingly.

2.3.2 Orientation

Tabletop displays are a relatively new form of computing environment with unique challenges and advantages. The horizontal orientation allows users to interact in a variety of configurations including being seated beside each other or across from each other. When collaborating at a table, people use orientation for comprehension (e.g. ease of reading, alternate perspective), coordination (e.g. establishment of personal/groups spaces), and communication [27]. Although support for face-to-face interaction can improve communication between collaborators [34], there is the disadvantage that users view the interface from different directions. As a result, digital information may be oriented correctly for me, but not for my collaborators. Designers of tabletop applications have addressed this issue in different ways.

The Personal Digital Historian (PDH) [39] uses a circular tabletop display that allows users to rotate either individual objects or the whole display, like the rotating "lazy Susan" often seen on tables in Chinese restaurants. When individual objects are moved, they automatically orient radially towards the outside of the display. Individuals can also "magnetize" the entire display so that all objects in the display orient along a predetermined direction. This approach relies on the users to rotate on-screen objects. Instead, a computer system could rotate digital objects automatically using contextual information to determine appropriate orientations for on-screen objects. When multiple ConnecTables [42] are joined to form a larger, integrated display, the orientation of displayed objects depends on the individual ConnecTable surface on which they were invoked. A similar orientation technique is used by the InfoTable [37] to orient display objects toward the table edge closest to a user's laptop. Both systems assume that a user's position will be based on a static "personal" display space, potentially leading to inappropriately oriented objects if users move around the table.

The concept of a Rotating User Interface (RUI) was proposed by Fitzmaurice et al. [11] to support artwork orientation. They defined a RUI as a system where the display could rotate freely, but the user interface components stayed fixed relative to the user. Although the proposed RUI differs from tabletop displays, the underlying issues are the same. Some UI widgets are rotation-sensitive, such as menus and other textual components. Other widgets, such as a color palette, are rotation insensitive [11] and thus would be equally accessible to all users around the table without re-orientation. Pop-up widgets are best suited to re-orientation because they can simply appear correctly for the user who invokes them [11], whereas static widgets need to be rotated by the user, or undergo a dynamic transformation when accessed by a user.

Collaborators around tabletop displays do not simply orient objects for ease of viewing, but use orientation to communicate information to others around the table [26]. Orientation of objects is used to establish public and personal space; personal objects are kept close to the user and oriented appropriately for them while group objects are oriented according to what the group has decided [26]. The use of orientation as a communicative gesture must be considered when determining at what level the orientation should occur. At the lowest level, rotation of widgets could be supported if one collaborator is accessing only one individual widget at any given time. However, rotating an application window including all the contained components might best support collaboration in other situations. Finally, the entire interface including all open applications could be reoriented [39].

2.3.3 Contribution of the Research

Although menus can be used in a collaborative setting, they are not expected to be the focus of the collaborative interaction. They should seamlessly support the collaboration without adding to the complexity of the task at hand. A pop-up menu is invoked by a particular user and each item in the menu is typically selected by that same user. However, the environment should support invocations of menus by multiple users either simultaneously or in tandem and with either shared or separate devices. In order to achieve this style of interaction, pop-up menus at a tabletop display should automatically appear oriented toward the invoker and appear in a suitable location according to their handedness. The work reported in this thesis demonstrates how a user's handedness can automatically be determined in this environment, despite possible frequent sharing of the input devices.

Chapter 3

User Experiments

A series of laboratory experiments were conducted to better understand how to design interfaces that utilize direct pen-input devices on horizontal displays. The following hypotheses were tested.

Hypothesis #1 Occlusion of menu items slows down target acquisition.

Hypothesis #2 Target acquisition times vary according to the position of the target menu items, relative to the point of activation.

Hypothesis #3 Target acquisition times vary according to the handedness of the user.

Hypothesis #4 Effects of occlusion and target position depend in a mirrored way on the handedness of the user.

Hypothesis #5 Positional differences in selection times for pen-input are different from observed patterns for mouse-input reported in the literature.

Hypothesis #6 Direct pen-input is faster than indirect pen-input.

Hypothesis #7 Effects of occlusion, target position and handedness vary according to the type of display.

These expected alternative hypotheses led to the following testable null hypotheses.

H1 Target acquisition time is not affected by occlusion.

H2 Target acquisition time is not affected by the position of the target, relative to the point of activation.

H3 Target acquisition time is not affected by handedness.

H4 Differences in target acquisition times for various target positions and levels of occlusion do not depend on the handedness of the user.

H5 Target acquisition time for pen-input is the same as for mouse-input.

H6 Target acquisition time is the same for indirect and direct pen-input.

H7 Differences in target acquisition times for various target positions and levels of occlusion do not depend on the orientation of the display.

Three experiments were run to test the above hypotheses. The experiments are progressive in nature. The first two experiments, although inconclusive, informed the design of the final experiment, which gives insight into precise patterns of target acquisition. The final experiment also sheds light on some deficiencies in the original experiments. Each experiment is described in detail with the lessons learned in each and how these issues were applied to the design of the final experiment.

Target Acquisition in Menu Selection

The three experiments were designed to identify effects on target acquisition times. We were interested in the target acquisition involved in selecting items from a pop-up menu. Selection from a pop-up menu can involve other subtasks, including a visual search for the menu item. In the first two experiments, the time required for this visual search is included in the dependent measure.

Selection from a pop-up menu also involves the movement to the *point of activation* of the menu. The point of activation is defined to be the location of the pen tip (or cursor for indirect input) when the menu is activated. Pop-up menus are typically placed relative to this point. In the following analyses, the movement from the point of activation to the menu item is considered, and not the initial movement to the point of activation.

3.1 Experiment 1: Rectangular Menu Placement

In our first experiment, we were interested primarily in exploring the issue of occlusion of rectangular menus and how occlusion relates to handedness. We designed our experiment to test hypotheses H1, H2, and H3. The experiment involved the use of a horizontal display with tethered direct pen-input.

3.1.1 Method

Participants

Eight users participated in the experiment of which four were left-handed. Two participants in each handedness group were female and two were male. Participants ranged in age from 21 years to 39 years. None of the users had previous experience with tabletop displays although three had used stylus input on personal digital assistants (PDA).

Apparatus

The tabletop system was top-projected and consisted of a 150 cm by 80 cm white laminate surface onto which output from a Pentium IV 2.0 GHz computer was projected. The projected display was 90 cm by 67 cm with a resolution of 1024 by 768 (see Section 4.3).

To provide styli input, a Polhemus Fastrak was used. A Fastrak is a six-degree-of-freedom (6DOF) magnetic tracking device that can detect the position and orientation of up to four sensors. Each sensor is a small receiver that is either embedded in a stylus device or placed on an arbitrary object, as is used for tracking head positions in head-mounted displays. The position and orientation of the receiver is sensed via a transmitter that emits electromagnetic pulses.

We developed Java-based software to process the Fastrak information sent to the serial port. To reduce magnetic interference, a wooden table was used and we avoided placing metallic or magnetic objects of any kind on or around the table while the Fastrak was in use. The transmitter was mounted underneath the centre of the table using industrial strength Velcro™. In order to improve the accuracy of the measurements, the transmitter was placed underneath the centre of the table so that the average distance to each possible tabletop display coordinate was minimized. Participants were seated at the midpoint of the long edge of the table. To determine acquisition times, the system created a "table touched" event when the stylus came within 0.01 cm of the table's surface and the pen stopped moving for 500 ms. These threshold values were found through pilot testing to minimize unintentional touches, while maintaining some causal association between a user's action and the computer's response.

Procedure

Participants began the experimental session by filling out a background questionnaire to gather information related to their experience with computers and their experience with input devices (see Appendix A.1). The questionnaire was a web-based form presented on the tabletop display, which helped participants to experience the tabletop display and stylus input. After the experimental task was completed, an experimenter engaged the participants in an interview. The interview was designed to elicit opinions on the tabletop display, stylus input, and menu placement.

Each trial in the experiment involved two parts (see Figure 3.1). The first part represented the movement to the point of activation of the menu. Participants were asked to rest the stylus tip in a starting position marked with a white circle. Participants were then asked to acquire a circular target as quickly and accurately as possible once it appeared on the screen. Targets were red, blue, or green and were presented at one of three distances (10.5 cm, 21.0 cm, 42.0 cm), three widths (1.3 cm, 2.6 cm, 5.2 cm), and three angles measured from the mid-line ($0°$, $-45°$, $45°$). The second part represented the movement from the point of activation to the item within the menu. Participants were asked to select from a rectangular pop-up menu the name of the colour of the target that had just been acquired. The pop-up menu was activated by the target acquisition from the first part of the trial. The pop-up menu was 7 cm wide by 8.5 cm high, and was placed in one of three locations relative to the point of activation: the bottom-right of the stylus, the bottom-left of the stylus, or directly above the stylus.

For the menus projected below the stylus, the closest corner of the menu was 1.2 cm from the stylus tip. For the menu placed above, the centre point of the bottom edge of the menu was 12.2 cm above the stylus tip. This distance was chosen to ensure that the menu was not occluded by the participant's hand or arm. The menu placement, position of name within menu, target width, target distance, and angle from start position were crossed to create 243 trials in the experiment. For

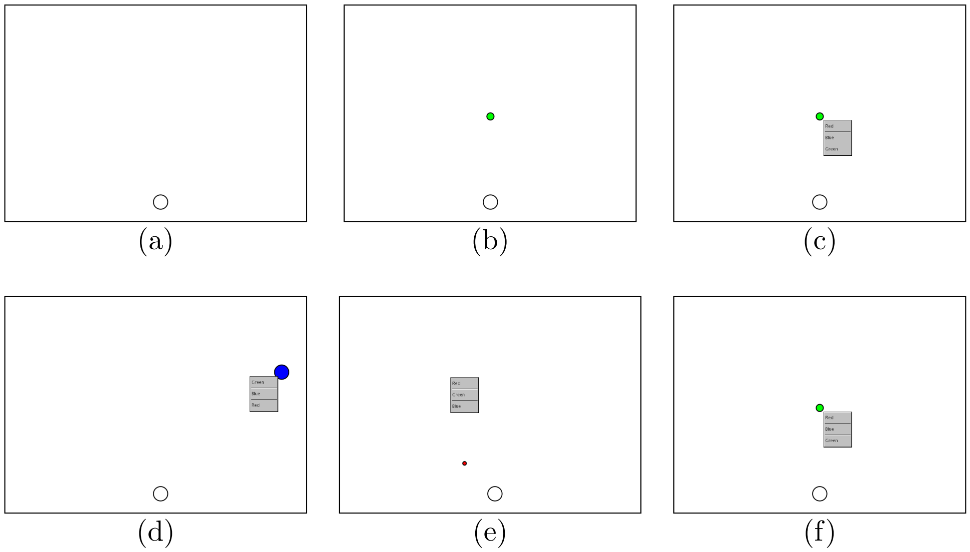

Figure 3.1: (a) The trial begins when the participant touches the start circle with the pen tip. (b) The first part of the trial involves the acquisition of a red, blue or green target (shown here in green). (c) A menu appears upon acquisition of the first target to begin the second part of the trial. The user then selects the name of the color of the target that was just acquired from the menu to complete the trial (back to (a)). The dependent measure is taken as the difference in time from the acquisition of the target in the first part of the trial (b) to the acquisition of the menu item in the second part (c). The menu appears either to the left (d), above (e), or to the right (f) of the point of activation (i.e. the first target).

each trial, a random menu configuration was chosen from the six permutations of the three colours. This configuration, together with the menu item position for that trial determined the colour of the initial target. After a few practice trials, these trials were presented randomly in one block, although participants were encouraged to rest at their convenience between trials.

3.1.2 Results

Analysis Technique

Participant background data were gathered from questionnaires, while preferences were determined from video recordings of participant interviews that followed the experiment. Computer logging was used to determine menu selection times and errors as well as to detect mistrials.

Because both left- and right-handed participants were tested on both sides of the mid-line, the data was mirrored for left-handed users to reflect activities relative to the dominant and non-dominant hands. Specifically, relative menu locations and angles were used. Rather than analyzing right and left menu locations, dominant and non-dominant menu locations were analyzed. For movement towards a target along an angle, the 45° angle towards the participant's dominant hand was called the dominant angle, whereas the 45° angle away from the dominant hand was called the non-dominant angle. Finally, for position within the menu, the positions were named near, middle and far. Note that for the dominant and non-dominant menu placements, the near menu item is at the top of the menu, but that for the menu placed above the hand, the near menu item is at the bottom.

We performed a 2 (handedness) x 3 (distance) x 3 (angle) x 3 (position within menu) x 3 (menu location) mixed design factorial analysis of variance on the selection times. The score for each distance/angle/position/location combination was averaged over the three target widths. Across all subject data, 166 (8.5%) incomplete trials and 71 (3.7%) outliers were removed from a total of 1944 trials. Trials beyond two standard deviations from the mean (across all subjects and conditions) were considered outliers. Three data points had either incomplete trials or outliers

Factor	Level	Mean (s)	F	p
Distance	Near	1.900	4.087	.044
	Middle	1.910		
	Far	1.977		
Position	Near	1.827	31.850	$< .001$
	Middle	1.895		
	Far	2.066		
Menu Location	Above	1.712	26.083	$< .001$
	Non-dominant	1.973		
	Dominant	2.102		
Handedness	Right	1.820	8.039	.030
	Left	2.038		

Table 3.1: A summary of the main effects.

for all three target widths and were filled with the average of all other users for that condition.

Menu Selection Times

A summary of the main effects for target acquisition times in the menu selection task is shown in Table 3.1.

There was a main effect of distance ($F(2,12) = 4.087$, $p = .044$). Post-hoc analysis revealed that at 42.0 cm (far), participants were significantly slower at acquiring menu items than at 10.5 cm (near) ($p = .048$).

There was a main effect of position within the menu ($F(2,12) = 31.850$, $p < .001$). Average selection times for near menu items were faster than selection times for far away items, relative to the point of activation of the menu. There was also a main effect of menu location ($F(2,12) = 26.083$, $p < .001$). For menus placed

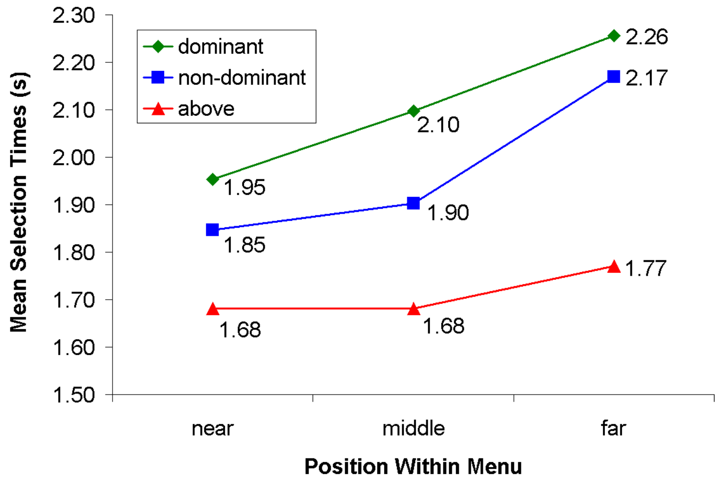

Figure 3.2: Interaction of Menu Location and Position Within Menu. For the above menu location, there was no significant difference between the menu positions.

above the hand and far away, selection times were fastest. The next fastest times were for menus appearing on the non-dominant side, and the slowest times were for menus that appeared on the dominant side. These differences were all statistically significant at $p < .05$.

As shown in Figure 3.2, there was an interaction between menu location and position within the menu ($F(4, 24) = 5.029$, $p = .004$). For all three menu placements, the order of mean selection times for the three item positions was near, middle, then far, in ascending order. These means were all significantly different for the dominant menu ($p < .05$). For the non-dominant menu, the participants selected items in the far position significantly slower than in the near ($p < .001$) and middle positions ($p = .001$), but the near and middle item positions were not significantly different ($p = .262$). For the menu placed above the point of activation, the selection times for the three item positions were not significantly different ($p > .15$).

The analysis revealed a main effect of handedness ($F(1, 6) = 8.039$, $p = .030$)

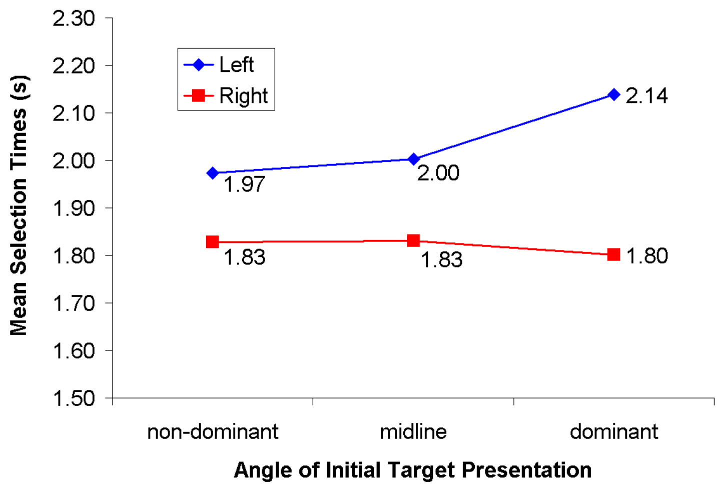

Figure 3.3: Interaction of Angle with Handedness. Left-handed users were only slower along the dominant angle.

with mean selection times for left-handed participants slower than for right-handed participants. In addition, there was an interaction between angle and handedness ($F(2, 12) = 4.217$, $p = .041$). For movement along the non-dominant angle and along the mid-line, left- and right-handed participants showed no significant differences in selection times ($p = .170$ and $p = .088$, respectively). However, along the dominant angle, left-handed participants had slower selection times than did right-handed participants ($p = .007$) (see Figure 3.3).

Errors

Only seven incorrect menu selections were recorded across all participant trials. No statistical analysis was run on these errors due to the fact that they occurred in less than one percent of the trials.

User Preferences

In the post-experiment interview, participants were asked which menu location they found easiest when selecting items. Following their response and explanation, they were then asked whether this was also the menu location they liked the best. Of the eight participants, six considered the non-dominant menu location to be the easiest and most preferred. Explanations were that it was closer than the menu above their hand, and that their hand did not occlude the menu items. The other two subjects preferred the above menu because it was easiest to see. All ten participants reported liking the dominant menu location least because their hand was in the way of menu items. All ten subjects also reported using adaptive strategies in the way they held the pen to minimize interference by their hand.

3.1.3 Discussion

The position within the menu correlates directly with the amplitude of the target and thus the main effect of position within the menu is predicted by Fitts Law. Menu positions that are further from the point of activation are slower to acquire. The main effect of distance may be due to the need for reaching. When the arm is fully extended, the movement required to acquire the menu item target may be more difficult than the movement required when the arm is closer to the body. Thus, the coefficients in Equations 2.1 and 2.2 would be higher for larger distances. Alternatively, the perceived width of the target menu item may be smaller for more distant menu selections, due to a change in visual angle. A smaller target width would correspond to a larger index of difficulty and thus slower acquisition times.

Based on the main effect of menu placement, one might choose to reject H1. Menus placed on the dominant side were reported to be occluded by the participants' hand. Both the dominant and non-dominant menus were placed at equal distances from the point of activation, with menu items of equal size and number. Even though the indices of difficulty were the same for both menus, the dominant menus were selected more slowly. User preference is consistent with the quantitative result; all participants found the dominant-side menu to be the least favourable. However, hand movement may be different for the target acquisition required in selecting the items in the two different menus. The effect of occlusion cannot be isolated as the cause of slower acquisition times for the dominant menus. Thus, this hypothesis cannot be rejected.

H2 cannot be rejected based on the results of this experiment. It is interesting that the menus placed above the hand were selected more quickly than either of the other two. The indices of difficultly for these menu items were *largest* so it was predicted that performance for this location would be the slowest. Several participants stated that they preferred this placement, although most did not. It is clear from this qualitative data that the users lack a consistent intuition about which of the two unoccluded menus is better.

The results show an effect of handedness that seems to depend on the angle of approach of the user. One possible explanation for this effect and interaction is that the device used was tethered, with the cord from the magnetic tracker placed to the left of the user for all participants. The interaction seems consistent with this bias being the explanation for the differences in means. Thus, H3 cannot be rejected based on this experiment alone. This issue is further discussed in the following section.

3.1.4 Lessons Learned

The method used to determine the time of the "table-touched" event introduced a measurement error. An alternative to this method would have been for the user to indicate actions via the button on the stylus sensor of the Polhemus Fastrak. This alternative was decided against because the introduction of such a button press would drastically alter the interaction technique. The technique would involve both a pointing task and a button press task, thus reducing construct validity. In order to account for both measurement error and construct validity, it was necessary to devise a more accurate method of detecting a "table-touched" event from the stylus input. An alternative method of detecting this event was introduced in future experiments.

The experiment showed the menu placement above the user's hand to be superior to both other placements. This result was contrary to our hypothesis, since it was placed sufficiently far from the user so that the index of difficulty was greater than in the other two placements, for all items within the menu. The conditions of the experiment were designed to demonstrate a difference between the other two placements, so no conclusion should be drawn from this result alone. It was decided that in future experiments, a larger set of target placements relative to the point of activation would be necessary.

As mentioned above, one aspect of this experiment that may have impacted left-handed users was the fact that the stylus device was tethered with a cable which was attached to the left-hand side of the table. To minimize interference of the cable, right-handed users often held the cable in their left hands, while left-handed users could not because that is where they held the stylus. The interaction with angle showed that significant differences between right- and left-handed users occurred only along the dominant angle. For left-handed users, this was the side of the table where the cable interfered most. To eliminate this bias, a wireless alternative could be used in future experiments.

The task in this experiment comprised two distinct parts:

1. Movement from the white starting position to the coloured target

2. Movement from the coloured target to the menu item

The first part included three factors (amplitude, target width, angle) and the second included two more (menu location, position of item), resulting in a five-factor design. A design of this size is difficult to analyze and requires many repeated trials for each participant. A reduction in complexity was decided upon for future experiments.

The experimental task performed by each participant was not typical of tabletop display applications, or even to computer applications in general. The lack of realism in the experiment made it difficult to directly apply the results of the experiment to the design of tabletop display applications.

Nevertheless, it is clear from the experiment that a direct application of Fitts law is not sufficient to predict user performance for direct pen-input onto horizontal display surfaces. The placement of menus relative to the hand and the distance to the menu were shown to effect the target acquisition required in selection from a menu, as well as the amplitude and width of the target. These additional factors should therefore be considered in the design of tabletop display applications.

3.2 Experiment 2: Rectangular Menus in Cartography

In the second experiment, we attempted to address some of the deficiencies of the first experiment, but still were testing only H1, H2, and H3. This experiment again involved only rectangular menus and was performed only on the horizontal tabletop display with tethered direct pen-input.

3.2.1 Method

Participants

Two left-handed and two right-handed students and faculty (3 male, 1 female) between the ages of 23 and 56 ($M = 32$, $SD = 15$) from a local university participated in the experiment.

Apparatus

The tabletop display used in this experiment was identical to the one used in the first experiment. The direct pen-input again was achieved using the Polhemus Fastrak 6DOF tracker, however, the "table touched" events were determined differently. The Fastrak pen was modified so that touching the table closed a circuit at the stylus tip, which sent a "mouse-pressed" event to the computer. This method of contact sensing proved to be more accurate than proximity calculations using the z-coordinate. The x- and y-coordinates of the tracker were still used to obtain the position of contact on the tabletop.

During the experiment, participants were asked to use a stylized cartography application. The application displayed a map of the world on the tabletop display. When the user touched a particular country with the pen-input device, a rectangular pop-up menu appeared with the following six options (see Figure 3.4):

Capital - display a label showing the capital of the selected country

Population - display a label showing the population of the selected country

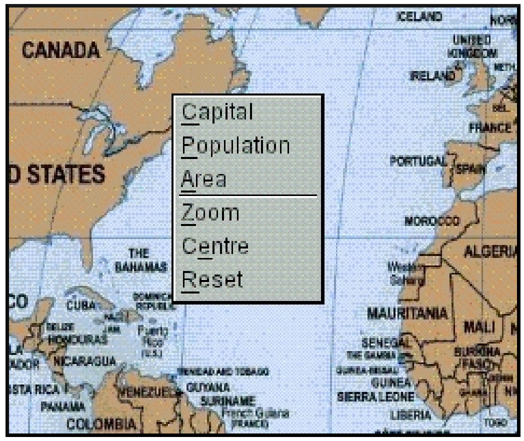

Figure 3.4: The menu used in the experiment. The menu contains six items separated by a single line.

Area - display a label showing the area in square kilometres of the selected country

Separator

Zoom - magnify the display by a constant factor and centre on the selected country

Center - centre the display on the selected country

Reset - reset the zoom to 100% and re-centre the display

Procedure

Each participant began the experiment by filling out a questionnaire intended to elicit background information such as experience with tabletop displays and pen-input devices. Following the questionnaire, participants were asked to sit at the tabletop display opposite the experimenter. The participants were then asked to complete a short training session to familiarize themselves with the tabletop display, the pen device and the application. For each trial, a participant was asked to perform a particular action on a particular country (e.g. "Display the population of Madagascar"). The experimenter provided the participant with each instruction verbally and pointed to the country while giving the instruction.

Participant were asked to perform four blocks of trials, one for each menu position relative to the point of activation (top-left, top-right, bottom-left, and bottom-right). The order of blocks was counter-balanced between subjects. Each block of trials consisted of 30 instructions, twice in each of 15 regions of the display, utilizing each menu item six times. Each block of 30 instructions was composed of six groups each containing 4 to 6 instructions followed by an instruction to reset the display. Four complete scripts were created and the order of presentation for these scripts was counter-balanced. The order of the six groups within each script was fully randomized for each script and each participant. Because the "Reset" menu item was not associated with a particular country, the six reset instructions were not included in the count of 30 trials nor were they included in the analysis.

Following the completion of the fourth block of trials, participants were asked to complete a short questionnaire. For the first part of the questionnaire, participants were shown all four pop-up menu placements simultaneously in each of six regions of the screen and asked to rank them in order of preference by selecting the menus with the stylus in decreasing order of preference. For the second part of the questionnaire, participants were asked to provide feedback about their experience.

3.2.2 Results

To analyze the data gathered in the experiment, a 5 (menu item position) x 4 (menu location) x 2 (handedness) full-factorial ANOVA with acquisition time as the dependent variable was used. Because this experiment was intended only as a pilot, data were collected from only four participants, so the statistical power of the

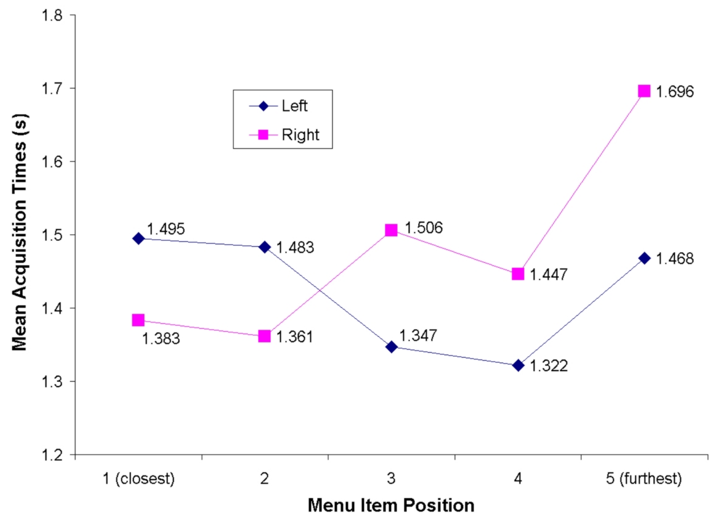

Figure 3.5: An interaction between menu item position and handedness shows that right-handed participants are slower at acquiring the furthest menu item position than are left-handed participants.

experiment was low. The intent of the experiment was to identify effects or trends to be tested in a subsequent experiment.

The analysis resulted in only a marginally significant main effect of menu item position ($F(4,8) = 3.7$, $p = .054$) and a significant interaction between menu item position and handedness ($F(4,8) = 4.2$, $p = .04$). Post-hoc analysis of the interaction revealed that for right-handed participants, the furthest position was significantly slower than both the closest position ($p = .042$) and the second closest position ($p = .026$) and the second furthest position was significantly slower than the third closest ($p = .019$). Figure 3.5 shows this interaction graphically.

User Preference

To analyze user preference data, a Kendall's W test was performed. Responses from each of the six regions of the display were considered as separate data points.

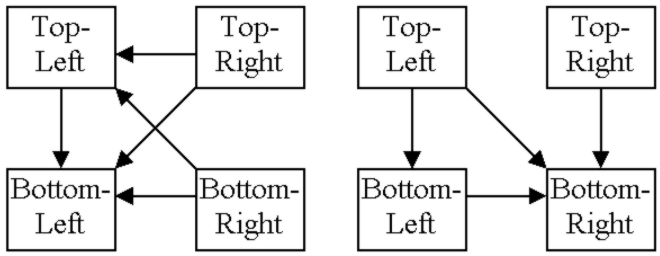

Figure 3.6: Qualitative results show that users prefer some menu placements to others. Arrows indicate an "is preferred to" relationship.

The data was split based on the handedness of the participant. Analysis revealed consistent results for both left-handed ($\chi^2(3, 12) = 31.3$, $p < .001$) and right-handed ($\chi^2(3, 12) = 15.5$, $p = .001$) participants. Post-hoc Wilcoxon signed ranks tests revealed the pattern of preferences shown in Figure 3.6. This pattern suggests that for left-handed users the bottom-left menu position is less preferable than the top-left position ($p = .004$), the top-right position ($p = .002$), or the bottom-right position ($p = .002$). Similarly, for right-handed users, the bottom-right position is less preferable than the top-right position ($p = .002$), the top-left position ($p = .010$), or the bottom-left position ($p = .012$).

3.2.3 Discussion

Due to the lack of significant results, none of the null hypotheses can be rejected. The marginally significant effect of menu item position can be easily explained by Fitts Law. Items that appeared closer to the point of activation were acquired faster than those that appeared further away. Since all targets had the same width, the index of difficulty was greater for the menu items that were selected more slowly.

The only hypothesis that could be considered for rejection is H4, based on the interaction between handedness and menu item position. This interaction shows an interesting pattern of acquisition times for the two handedness groups. The interaction suggests no difference in the five menu item positions for left-handed users. It may be that, because of the way left-handed users hold the device, the closer menu items are obscured by the hand. This would slow down acquisition of closer menu items and balance the effect of target distance for the more distant menu items. The effect also suggests that, for right-handed users, the effect of target distance outweighs the effect of occlusion. This may be because the way that right-handed users hold the pen causes less occlusion in the near menu item positions than for left-handed users. Although this effect appears to distinguish left- and right-handed users, it is insufficient evidence to fully reject the possibility that H4 is true.

Although this experiment failed to demonstrate the expected mirrored pattern of target acquisition for left- and right-handed users, it is clear from the qualitative results that users have a clear dislike for occluded menu placement. Although this data could not be corroborated by our quantitative analysis, user frustration may nonetheless result from poor menu placement. Over time, this frustration may also degrade the performance of the user.

3.2.4 Lessons Learned

The tabletop display used in this experiment was the same as that used in the first experiment. The pen input device was again tethered, which may be the cause of the handedness interaction that was found. To eliminate this possibility, an untethered device should be used in future experiments.

The task in this experiment was chosen primarily for its realism. The application is intended to represent a real use of tabletop display technology where menus may be activated by the stylus input. To provide increased realism, it was necessary to sacrifice some precision in the experimental task. It appears, however, that the size of the effects caused by occlusion and hand movement may be small enough that they are outweighed by effects of cognition in realistic tasks. Because the participants were required to decide which menu item to select, the effects on automatic hand movements were not noticeable in the analysis. To isolate the desired effects from these effects of cognition, a more precise task should be used in future experiments.

Another disadvantage caused by opting for realism was a lack of generalizeability of the results. The task in this experiment is specific to direct pen-input on tabletop displays. Future experiments would benefit from the inclusion of other media, such as vertical large-screen displays, tablet PC displays, as well as other input techniques including indirect pen or mouse input.

3.3 Experiment 3: Target Position in Circular Menus

The two preliminary experiments suggested to us that the movement time varies according to the position of the target relative to the point of activation. In this experiment we directly tested hypotheses H1, H2, H3, H4, H6, and H7.

We believed 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.

The target acquisition task used in the experiment most closely resembles selection from a circular pop-up menu. The results should 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. In order to achieve a higher level of precision, realistic rectangular menus were not directly tested.

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

3.3.1 Method

Participants

Six left-handed and six 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 the experiment.

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 1:6.7 in the indirect condition (since the measurement of the target width is constant in control space, this part of the ratio is assigned a value of 1, alternatively the ratio could be written 0.15:1 if the display space was considered constant). In all three conditions, the participants 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 obtain information about their 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 [35] 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, a participant was asked to point and select one of four regions of the display that was indicated by the border of that region changing to a

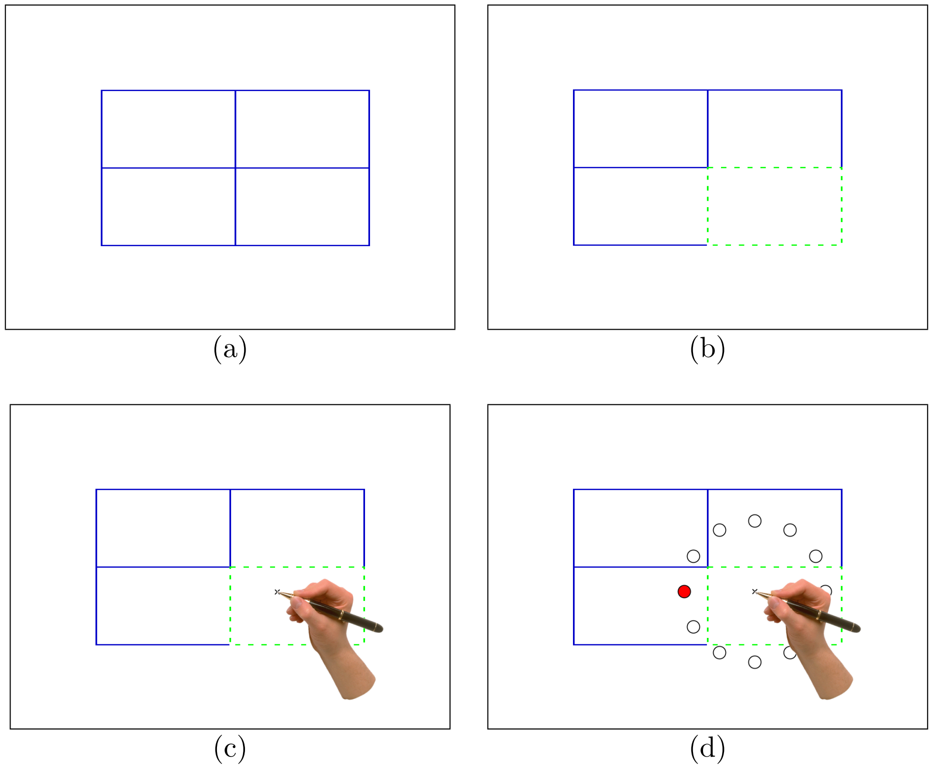

Figure 3.7: (a) To begin each trial, each participant was presented with four regions outlined in blue. (b) To start the trial, the outline of one region changed to a different colour (dashed in this figure, but not in the experiment). (c) The participant then would point and select that region of the display. (d) This action activated a ring of twelve circles, with the target circle in red.

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 so that its center was 350 mm from the point of contact of the stylus (see Figure 3.7). 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. 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. As suggested by Graham [16], 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 Equation 2.1) 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.

Because of the difference in the size and resolution of the displays, the width of the target spans the same number of pixels in both the indirect and horizontal-direct conditions (30 pixels), but a much smaller number of pixels in the vertical-direct condition (5 pixels). This difference has two negative side effects. In drawing the circles, a single pixel accounts for a 3% change in the indirect and horizontal-direct conditions, but a 20% change in the vertical-direct condition. Thus, effects whose size is smaller than 20% may be due to this effect. Also, circles that are drawn with a smaller number of pixels look different than those drawn at a higher resolution. This visual difference may have an effect on the cognitive task required

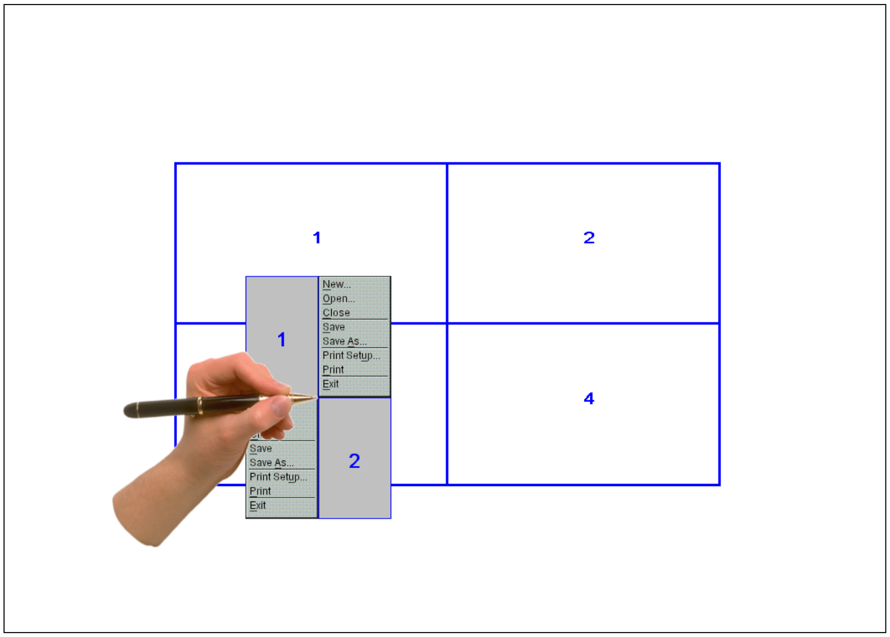

Figure 3.8: Users were asked to rate the four possible rectangular menu placements in order of preference in all four regions of the display. Menus were shown relative to the point of activation (the center of the four menus). The four rectangular menus were replaced by the integers 1-4 as they were clicked in decreasing order of preference.

by the participants.

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 3.8). Second, for each individual menu placement, participants were asked to state on a five-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.

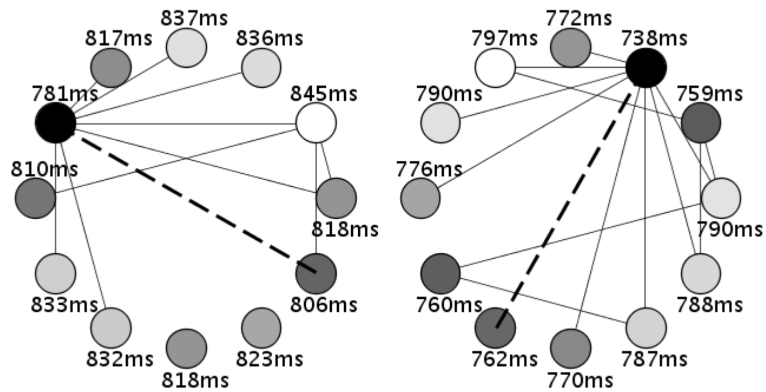

Figure 3.9: An interaction between position and handedness shows that left-handed participants (left) select targets more slowly in the bottom-left and top- right of the targets and faster in the top-left and bottom-right. The mirror effect occurs for right-handed participants (right). Labels represent 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.

3.3.2 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 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 from the actual location of the stylus tip in both of the direct conditions. In the vertical-direct condition, the precision of the SmartBoard touch screen is not below one pixel and with a target width of five pixels, sub-pixel resolution would be required to determine the true distance from the target. Because the targets on the Tablet PC have more pixels, the error can be more accurately measured in the horizontal-direct condition, but the Tablet PC is also sufficiently inaccurate to measure the true distance from the target. This situation also made it impossible to use effective throughput as the dependent measure as suggested by the ISO 9241 standard, Part 9 [7].

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 3.7). 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 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 3.9. 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 condition. Post-hoc analysis revealed more significant differences in the horizontal-direct condition and fewer significant differences using indirect input (see Figure 3.10).

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

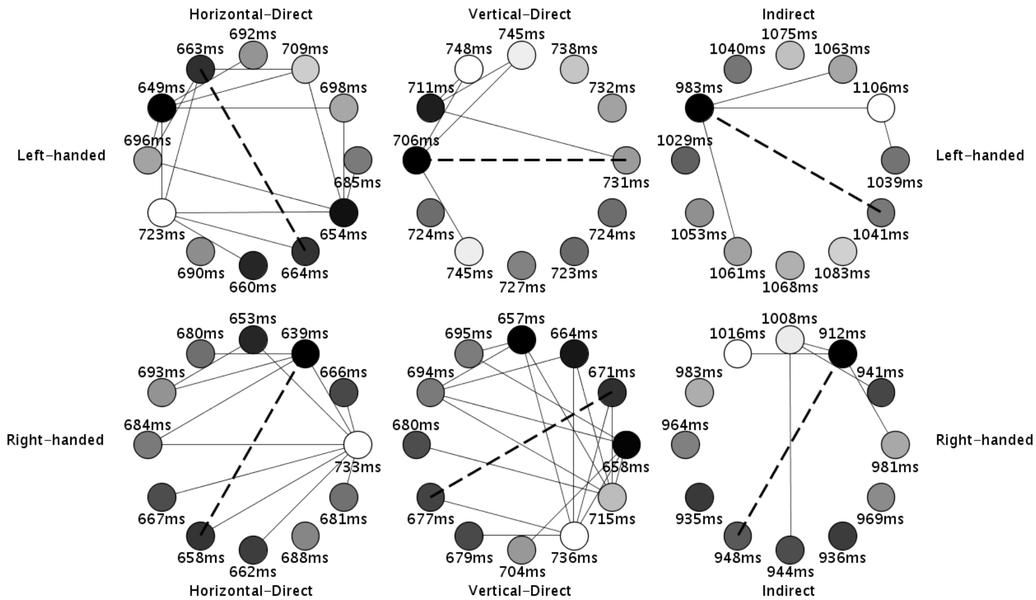

Figure 3.10: A three-way interaction between display condition, handedness and target position shows a different pattern of fastest selection times for all three display conditions. In the horizontal-direct condition, movement along the top-left to bottom-right axis was fastest for left-handed participants (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 3.9.

ANOVA resulted in the same main effects and interactions (see Table 3.2).

A Note on the Analysis

An alternative analysis, sometimes used in the literature, can be performed such that each trial is considered as a separate data point. We did this analysis and found additional effects, including an unexpected main effect of handedness. This analysis increases the degrees of freedom and can artificially reveal effects that would

	Acquisition Time Analysis		Throughput Analysis	
	F	p	F	p
Display	76.4	$< .001$	443.1	$< .001$
Target Position x Handedness	4.1	$< .001$	4.7	$< .001$
Target Position x Handedness x Display	1.8	.017	2.3	.001

Table 3.2: The data were normalized by using throughput as the dependent measure. The analysis revealed the same main effects and interactions.

not be otherwise significant. We believe that the more common method of averaging over trials before applying the ANOVA is a more stringent analysis. The alternative analysis is presented in Appendix B.

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 ($\chi^2(3,6) > 7.0$, $W > .50$, $p < .05$). Table 3.3 summarizes these results. 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

To analyze suitability of menu placement, a series of Kruskal-Wallis tests was used. Results showed significant 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 ($\chi^2(3, 6) > 9.0$, $p < .05$). Table 3.4 summarizes these results. 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 3.11). There were three additional significant pairwise differences

Left-handed		TL	TR	BL	BR
H-D	W	0.467	0.433	0.589	0.456
	$\chi^2(3,6)$	8.4	7.8	10.6	8.2
	p	**0.038**	**0.050**	**0.014**	**0.042**
V-D	W	0.589	0.533	0.611	0.611
	$\chi^2(3,6)$	10.6	9.6	11.0	11.0
	p	**0.014**	**0.022**	**0.012**	**0.012**
I	W	0.167	0.044	0.256	0.344
	$\chi^2(3,6)$	3.0	0.8	4.6	6.2
	p	0.392	0.849	0.204	0.102

Right-handed		TL	TR	BL	BR
H-D	W	0.078	0.256	0.211	0.144
	$\chi^2(3,6)$	1.4	4.6	3.8	2.6
	p	0.706	0.204	0.284	0.457
V-D	W	0.411	0.644	0.233	0.344
	$\chi^2(3,6)$	7.4	11.6	4.2	6.2
	p	0.060	**0.009**	0.241	0.102
I	W	0.300	0.133	0.300	0.167
	$\chi^2(3,6)$	5.4	2.4	5.4	3.0
	p	0.145	0.494	0.145	0.392

Table 3.3: Results of Kendall's W test for user preference in each combination of handedness, display condition, and region of display. Significant results appear in boldface.

Left-handed		TL	TR	BL	BR
H-D	$\chi^2(3,6)$	12.7	13.9	14.6	18.0
	p	**0.005**	**0.003**	**0.002**	**$<$ .001**
V-D	$\chi^2(3,6)$	14.3	16.0	14.8	12.7
	p	**0.002**	**0.001**	**0.002**	**0.005**
I	$\chi^2(3,6)$	6.7	5.0	5.2	4.9
	p	0.082	0.169	0.160	0.179

Right-handed		TL	TR	BL	BR
H-D	$\chi^2(3,6)$	13.3	14.0	11.1	9.2
	p	**0.004**	**0.003**	**0.011**	**0.027**
V-D	$\chi^2(3,6)$	13.9	12.2	15.3	16.5
	p	**0.003**	**0.007**	**0.002**	**0.001**
I	$\chi^2(3,6)$	1.7	9.8	1.3	0.8
	p	0.630	**0.021**	0.740	0.860

Table 3.4: Results of Kruskal-Wallis tests for suitability of menu placement. Significant results appear in boldface text.

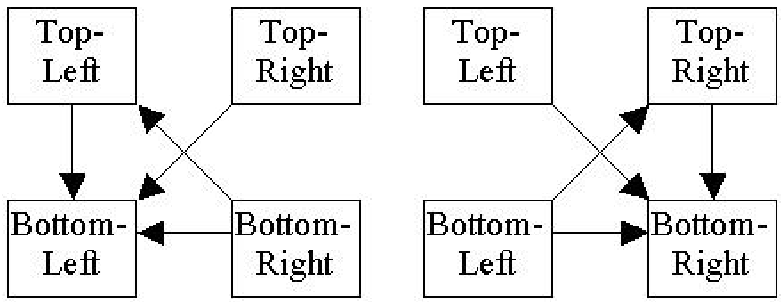

Figure 3.11: Results of suitability ratings suggest a pattern of ratings for left-handed participants (left) that is mirrored for right- handed participants (right). Arrows indicate the "is more suitable than" relation.

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 ($p = .027$).

3.3.3 Discussion

The differences in display conditions suggest that hypotheses H6 and H7 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. The difference in selection times between the horizontal and vertical displays is likely due to an increased fatigue effect in the vertical display. 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. Thus, the conclusion that must be drawn from the results of this experiment are specific to target acquisition and do not generalize sufficiently to inform the choice of input technique when designing applications. However, one can conclude that when the choice to use direct input has already been made, fatigue effects will likely be greater if a vertical display is chosen instead of a horizontal one.

The interactions involving handedness suggest that H4 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.

Hypotheses H1 and H2 can also be rejected. 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 (see Chapter 2, Section 2.1.1). 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 explains 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.

3.3.4 Lessons Learned

This experiment benefited greatly from the lessons learned in the previous two experiments. A wireless pen-input device, used for input to a Tablet PC, was used in all conditions to eliminate the effect of the tether. Both other experiments used the tethered pen-input and suffered from handedness effects that seem to be reduced in this experiment. The experiment included conditions that accounted for various target positions relative to the point of activation in order to accurately assess positional differences. The first experiment clearly did not account for these differences and the second experiment did so insufficiently. The task in this experiment was a simple target acquisition task and did not involve multiple parts, thus simplifying the analysis.

The experiment still suffered, however, from a lack of realism. The task performed was designed to precisely isolate particular effects of pen use on horizontal and vertical display surfaces. To obtain this precision, realism had to be sacrificed. The need for this sacrifice was apparent in the lack of results in the second experiment.

Measurement error was not fully eliminated in this experiment. Because the position of contact was measured differently on different surfaces, the error in performance could not be compared between display conditions. The precision of the wireless device and touch-sensitive surface was also insufficient for error analysis. With the wireless device and vertical touch-screen, it was possible to accurately determine the time at which the user touched the display, but the position of this contact could not be recorded with sufficient accuracy. To obtain the necessary information, the measurement must be made in a uniform way in all conditions with a more accurate device.

3.4 Experimental Conclusions

In this section, the results of the three separate experiments are summarized and are related to similar research experiments. From the three experiments, conclusions can be drawn about all of the hypotheses laid out at the beginning of this chapter. These conclusions describe the effects of occlusion, target position, handedness, input technique, direction (direct vs. indirect input) and display orientation (horizontal vs. vertical).

In all three experiments, the apparatuses involved only the use of pen-input devices. The conclusions drawn here apply only to this input technique, but are contrasted with similar results using mouse-based input. The target acquisition in all three experiments is similar to selection in both circular and rectangular pop-up menus. These conclusions are specific to this method of interaction, but may in some cases generalize to other types of target acquisition.

3.4.1 Effect of Occlusion (H1)

Based on the results of all three experiments, it is not clear that occlusion alone slows down target acquisition in menus for pen-input devices. Occluded menus are doubly disadvantaged because acquisition of targets that are occluded by the hand requires the movement of the arm for selection. Arm movement has been shown in this experiment as well as others to be slower than hand/wrist movement. It is clear that selection is slower in menu placements that are occluded than in some placements that are not, but other unoccluded placements appear to be equally slow to acquire. The effect of occlusion seems to be outweighed by the effect of differing arm/hand motions. This coupling of arm movement with occlusion prevents drawing any conclusion about the effect of occlusion on target acquisition.

Although occlusion cannot be quantitatively assessed as a cause of slower target acquisition, users demonstrate a clear preference for menus that appear unoccluded and consistently rate occluded menus as less suitable. Over time, the frustration of occluded menus may cause a degradation of performance that was not noticeable in these three experiments.

3.4.2 Effect of Target Position (H2)

It is clear from the three experiments that the position of the target relative to the point of activation has an effect on acquisition time. Other studies have shown that target acquisition requiring arm movement is slower than target acquisition requiring only hand/wrist movement. These three experiments are consistent with that result and provide insight into which target locations require arm movement and which do not. However, this result does not appear as a main effect in any of the three experiments. Because hand movement is different for left- and right-handed users and because this movement varies in different display conditions, this effect is only revealed as an interaction with these other factors.

3.4.3 Effects of Handedness (H3 and H4)

Bias in Design

In all three experiments, left-handed participants had slower mean selection times than right-handed participants. This effect was only significant in the first experiment and in the alternate analysis of the last, but the result was consistent for all three. In the first two experiments, the effect of handedness could be explained by the tether on the pen. However, in the final experiment, this tether was not present, yet some participants still commented that the pen was perhaps not designed for left-handed users. This suspicion of bias is likely due to an expectation for this design flaw based on experience with other devices such as the mouse, but may indicate a lack of exact symmetry in the left-handed versus right-handed use of an apparently symmetric device.

In all three experiments, some of the left-handed participants reported using the mouse with their right hand, but no right-handed participants reported using the mouse with their left. Nonetheless, all left-handed participants used their left hand with the pen input in the three experiments. Another explanation for the right-handed bias is a transfer effect from mouse usage. It is possible that left-handed users are slower, simply because they must retrain their left hand to do tasks familiar to their right.

The third experiment also provides evidence consistent with this claim. The alternate analysis of the experiment suggests a potential bias for right-handed users that is less prevalent on horizontal displays with direct input. An interaction between handedness and display shows left-handed users to be most disadvantaged in the indirect condition and least disadvantaged in the horizontal-direct condition. Since pen usage directly on a horizontal display is most similar to writing, both handedness groups have prior training in this condition. Since indirect pen input to a vertical display is most like mouse input, it may be that left-handed users are more disadvantaged by this being a less familiar task.

Another potential explanation for the handedness-display interaction is lateral asymmetry. Since the indirect condition limits the second homing phase of the targeting task to the visual system and disallows the use of the visuomotor mechanism, the indirect condition may be biased toward the left hemisphere of the brain. This bias may lead to slower target acquisition with the left hand for indirect input devices (see Chapter 2, Section 2.1.1).

A Mirrored Pattern for Handedness Groups

As mentioned in Section 3.4.2, the effect of target position is not revealed as a main effect in the three experiments. Instead, the effect appears as an interaction with handedness. This interaction shows a mirrored pattern for optimal target placement. Figure 3.12 shows this effect pictorially.

According to the results of these experiments, a static placement of targets would disadvantage either left-handed or right-handed users. There is no single placement that is optimal for both groups. It is therefore clear that, in the design of applications that use pen input, the system must know the handedness of the user in order to determine an optimal placement strategy.

Although this handedness interaction shows a clear difference between the two handedness groups, it does not provide a complete picture with which a de-

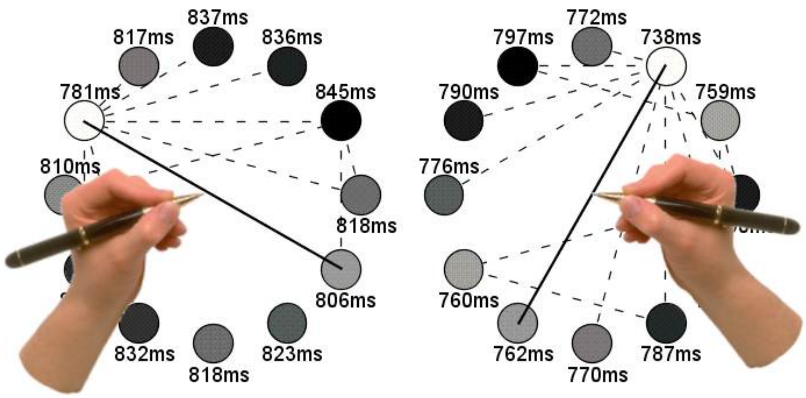

Figure 3.12: The optimal placement of targets is different for left-handed and right-handed users. This picture is a modification of Figure 3.9. Note that the hands are not necessarily drawn to scale.

signer can create an application. Without specific domain knowledge, such as the orientation of the display and the type of pen input (direct or indirect), a target placement based on handedness information alone can still be sub-optimal.

3.4.4 Pen Input vs. Mouse Input (H5)

The three experiments do not consider mouse input directly, and so no conclusion can be drawn from these experiments alone. However, other mouse-based studies have shown similar patterns with which to compare pen-based input. The position of the hand when using the mouse is drastically different than when using a pen (see Figure 3.13). Thus, as with pen-input in various modalities, mouse-based input has a different pattern for optimal target placement. For mice, hand/wrist movement occurs in the left to right direction, and arm movement occurs in the front to back direction, for both left-handed and right-handed usage. Thus, optimal target locations would be to the left and to the right of the point of activation.

This explanation is consistent with other mouse-based results. Boritz et al.

(a) (b) (c)

Figure 3.13: The positions of the arm and hand are quite different for mouse-based input (a) than for pen-based input either horizontally (b) or vertically (c).

[4] performed a study that measured performance of target acquisition with a mouse at various angles from the point of activation. They found that target placements below the point of activation were significantly slower for right-handed participants. Kurtenbach and Buxton [28] measured performance of mouse-based menu selection in marking menus and found that participants selected items faster for on-axis items (sometimes involving only left to right movement) than for off-axis items (always involving some arm movement). These studies, together with the three experiments performed for pen-based input, show that optimal placement for targets is different for pen-based input than for mouse-based input.

When designing applications, it is important to consider the input technique when deciding where to place on-screen targets. This collection of studies provides insight into the optimal placement, given the input technique used.

3.4.5 Indirect vs. Direct Pen Input (H6)

Of the three experiments, only the third measured differences in direct versus indirect input. This experiment suggests that direct pen input is indeed faster than indirect pen input, in general. However, direct pen input on a vertical display was slower than direct pen input on a horizontal display, suggesting a possible effect of fatigue. Although direct pen input to the vertical display may be faster than indirect input, the effect of fatigue may cause performance degradation in the long-term in favor of the indirect method.

In the third experiment, hand placement for the horizontal-direct condition was similar to hand placement for the indirect condition. Although the hand positions are the same, the variance in mean selection times for the different target positions was far greater for the indirect condition. This difference is possibly because of training effects due to the lack of familiarity of the indirect task. It may also be that occlusion effects are greater on the horizontal-direct condition that are not present in the indirect condition, but the difference in variance makes it impossible to distinguish these two possibilities.

3.4.6 Horizontal vs. Vertical Display Surfaces (H7)

Of the three experiments, only the third explores differences between horizontal and vertical displays. This experiment shows a different pattern of target acquisition means for each type of display. Users tend to hold their hands differently when using a pen directly on the horizontal surface than when using a pen on the vertical surface. Thus, the optimal placement of targets may be slightly different for horizontal and vertical displays. Figure 3.14 shows this difference pictorially. Optimal target placement for horizontal displays is to the left and right of the hand. This placement allows for the use of only wrist movement to acquire targets. For vertical displays, the optimal placement is clearly different.

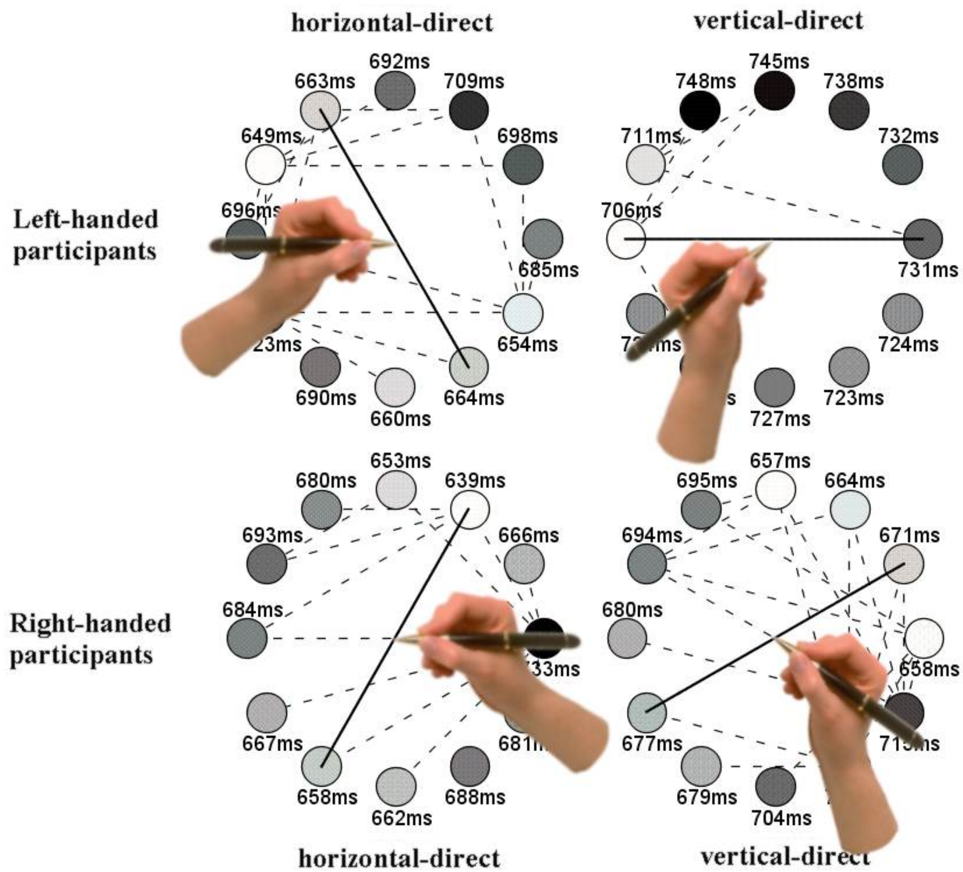

Figure 3.14: The position of the user's hand is slightly different for horizontal and vertical displays. This picture is a modification of Figure 3.10. Note that the hands are not necessarily drawn to scale.

Chapter 4

Menu Placement Strategies

The results of the experiments show that if a pen-input interface does not account for the handedness of users, either left- or right-handers will be disadvantaged. Since acquisition times for targets in the experiment showed a mirrored pattern, menus presented in a static position, relative to the point of activation, cannot possibly maximize performance for both groups. In order to eliminate this bias, it is necessary to create an interface that is capable of presenting menus differently to each user group.

There are two potential methods to provide support for both left-handed users and right-handed users. One method is to provide an adaptable display that allows the user to choose the appropriate placement according to preference. The experiments show 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 the user. This method requires a model of the user that includes handedness.

4.1 Adapting to a User Model

We begin with a few definitions:

user action - an action performed by the user in order to complete a task.

behaviour - the response of the system to a specific user action.

context of use - the information present in the system about the user at the time of a user action.

The behaviour of a system can be placed into any of three categories:

static behaviour - behaviour that is determined completely by user actions

adaptable behaviour - behaviour that can be modified explicitly by the user

adaptive behaviour - behaviour that is determined by the context of use

Interfaces that make use of primarily one of these three types of behaviour can be referred to as static, adaptable, or adaptive interfaces, respectively. Interfaces that make use of a mixture of these three can be named accordingly. Adaptable and adaptive behaviours can also be grouped into a single category, namely **dynamic behaviour**.

Arguments have been made both in favour and against all three types of behaviour. In this particular case, the *user action* in question is the activation of a menu and the *context of use* is the handedness of the user. However, the following discussion of the preferred behaviour applies in general to any system where the context of use involves an unchanging feature of the user.

As mentioned above, it has been shown that a static interface would disadvantage at least some subset of users. Choosing to use a static interface in this case is analogous to providing an Arts & Crafts classroom with only right-handed scissors. If the class size is large enough, some of the students are guaranteed some level of frustration.

A disadvantage to using an adaptable interface is that the user is required to manually perform the adaptation. At least two types of overhead necessarily result from this required step. A temporal overhead occurs when a user must take the time to perform the adaptation, and thus increase the overall time to complete their task. A cognitive overhead occurs when a user is required to learn a process by which they can alter the placement of pop-up menus. The user may be hindered by additional visual cues on the display, by the need to remember a gestural command, or by whatever mental step required for the adaptation. Both the temporal and cognitive overhead decrease the overall efficiency of an adaptable interface.

Adaptive systems often suffer from the effect of a dynamically changing interface that can sometimes lead to confusion in the user. When the user is not provided with control of the adaptation, they may not understand the reason for a change in the behaviour of the system. Furthermore, a dynamically changing interface can hinder the visuomotor memory of the user, thus decreasing the learnability and long-term efficiency of the interface. However, the adaptive system that is suggested here does not have this disadvantage. As long as the user model can accurately predict the user's handedness, each user will have a consistent interface.

Some software for the Tablet PC already include an adaptable interface in the form of 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 an option is a minimum requirement for such applications. In collaborative environments, however, this minimum requirement is no longer sufficient. Kurtenbach et al. [30] 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, a model of each user is created 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. Three different methods of obtaining a user model were tested.

4.2 Adaptive Map Application

A sample tabletop display application (see Figure 4.1) is used to demonstrate the use of the results of the experiments in application design. This application displays a map of the world containing information about individual countries. To display

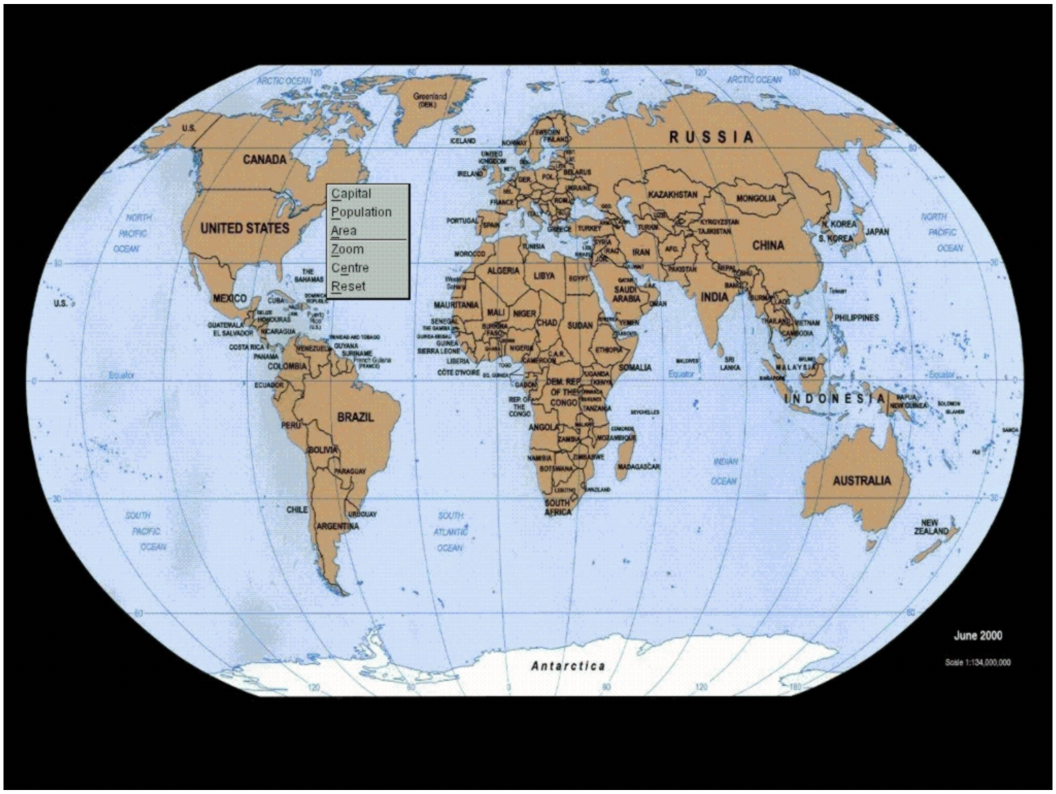

Figure 4.1: The sample map application used to demonstrate the handedness detection software.

this information, the user must tap on a country using a pen-input device and choose from one of six items in a 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 can be invoked from any of the four sides of the table and it will be automatically oriented so that it is readable from the location it was invoked.

The six menu items are as follows:

Capital - display a label showing the capital of the selected country

Population - display a label showing the population of the selected country

Area - display a label showing the area in square kilometres of the selected country

Separator

Zoom - magnify the display by a constant factor and centre on the selected country

Center - centre the display on the selected country

Reset - reset the zoom to 100% and re-centre the display

This application demonstrates the benefit of the experimental results by providing a pop-up menu that is placed so as to optimize both user performance and to increase user satisfaction. The menus could alternatively be placed up and to the right for right-handed users and up and to the left for left-handed users. This placement would also optimize performance, however qualitative results suggest this location is less preferable.

4.3 Adaptive Tabletop System

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 centre 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

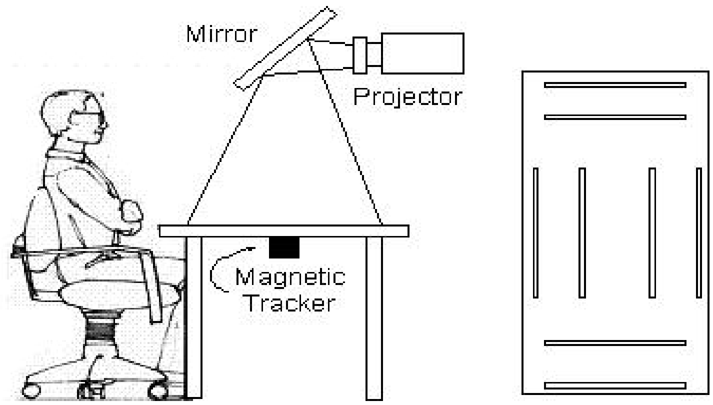

Figure 4.2: The tabletop display is projected onto the surface from above with the Fastrak cube placed underneath the table (left). The force sensing resistors are distributed evenly across the table's surface as seen from the top (right).

Figure 4.2).

In all three implementations, the handedness of the user is determined for a single point in time. That is, the position and orientation of the Fastrak and the side of the table of the user at an instant in time is used to predict the handedness of the user holding the device at that instant. In the sample map application, the instant that is used as input to the system is the time at which the menu was activated.

The side of the table is determined by the FSR strip that had the maximum force applied within the five seconds prior to and including the instant used. The threshold value of five seconds introduces a short delay before the system can detect a change in side when the user moves or passes the device. This threshold may vary between applications and should depend on the context of use. That is, applications that involve frequent sharing of devices or movement around the table should have lower threshold values, whereas applications that involve little transfer of input device or movement should have a higher threshold value. Alternatively, this threshold could be chosen dynamically.

4.3.1 Simple Heuristics

In order to demonstrate the need for a slightly more complex system, the first model that we explore 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. Thus, the formula used to determine handedness is:

$$handedness = \begin{cases} right, & 0 < azi <= \pi \\ left, & \pi < azi <= 2\pi \end{cases} \tag{4.1}$$

4.3.2 Neural Network

The second model utilizes a feed forward neural network, to determine the handedness of the user. The input layer has a node for each of the 6DOF, 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 (see Figure 4.3).

Training the Network

Before the network can be used to determine the side of the table with any degree of accuracy, it must be trained. Training is performed using the backpropagation algorithm. That is, for each instance in the training corpus, the input is used as activation for the input layer and is propagated to the output layer. The received output is then compared to the desired output and an error value is calculated for each node in the output layer. The weights on edges going into the output layer are adjusted by a small amount relative to the error value. This error is

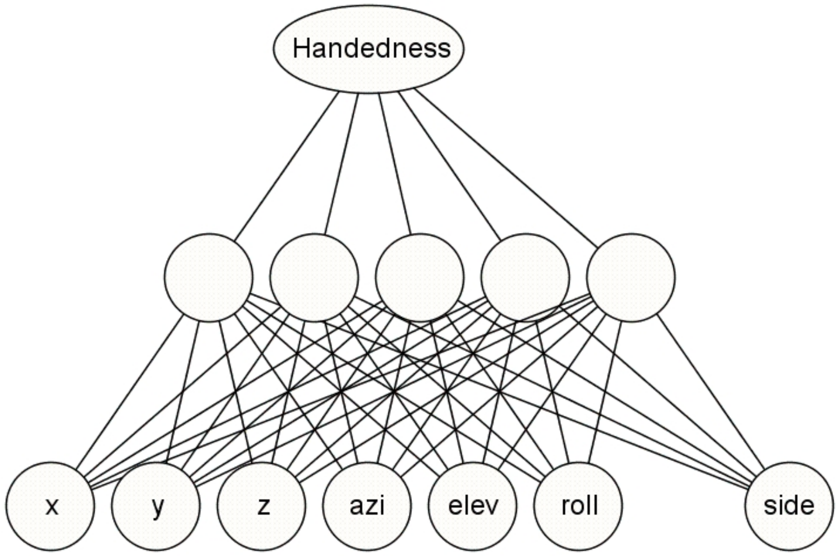

Figure 4.3: Neural Network Model of the Interactive Tabletop Display System

propagated backwards through the network to correct edge weights at all levels. For this particular network, a learning rate of 0.5 is used. The training corpus is passed through the network 100 times.

Utilizing the Network

The network functions by propagating the activation from the input layer towards the output layer. First, the activation of each hidden node is calculated as a weighted sum of the activation of its adjacent nodes in the input layer. That is, the activation of the j^{th} node in the hidden layer is determined by the following equation:

$$a_j = \sum_{i=1}^{7} w_{ij} a_i \tag{4.2}$$

where a_1,...,a_7 are the activation levels of the seven input nodes. Similarly, the activation of the node in the output layer is calculated by a weighted sum of activation of adjacent nodes in the hidden layer.

Complexity

To utilize the network, the summation in Equation 4.2 must be performed for each node in the hidden layer and each node in the output layer. Thus, the complexity of the algorithm to propagate through the network once is $O(e)$, where e is the number of edges in the network. In this case the number of edges in the network is held constant at $e = 40$, and so the complexity is essentially $O(1)$ to determine a particular user's handedness, given the input from the Fastrak and the location of the user. To train the network, propagation must be performed for each training instance, so the complexity is $O(n)$, where n is the size of the training corpus.

Advantages and Disadvantages

The neural network utilizes stochastic training methods to determine the handedness of the user. Because the neural network determines the output based on data from actual use, the network is more likely to be accurate than a heuristic "guess".

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.

Another disadvantage of the neural network is the need for training. Although testing shows that the network can be trained for a large population with a small sample (see Section 4.4), it is still possible for the network to incorrectly identify handedness. Because the network uses statistical data to determine its output, the user cannot easily determine under what conditions to expect failure.

4.3.3 Bayesian Network

The third model used to determine handedness is a Bayesian Network. This network utilizes a model of causation to probabilistically determine the user's handedness. Like the neural network, the Bayesian network is trained using actual data, however, the causal relationships between nodes in the network are determined based on a priori knowledge of the tabletop display environment.

The model used for this tabletop display system contains ten variables, four discrete and six continuous (See Figure 4.4). Using this model, the system can observe the measured coordinates from the Polhemus Fastrak device and eliminate variables to determine any of the following information about the person using the given input device:

- The handedness of the user
- The side of the table at which the user is sitting
- The true position and orientation of the user's device

Discrete Variables

The side, handedness, actual position, and the actual orientation are discretized in the Bayesian network model. Probability tables for the actual position and orientation variables were obtained by training using the Naive Bayesian rule. Each side and each handedness value were given equal prior probabilities. Although this may not reflect a true tabletop display environment, these probabilities were used to test the ability of the network to detect the true value of these variables *without* prior knowledge. If used in practice, these prior probabilities should be adjusted accordingly.

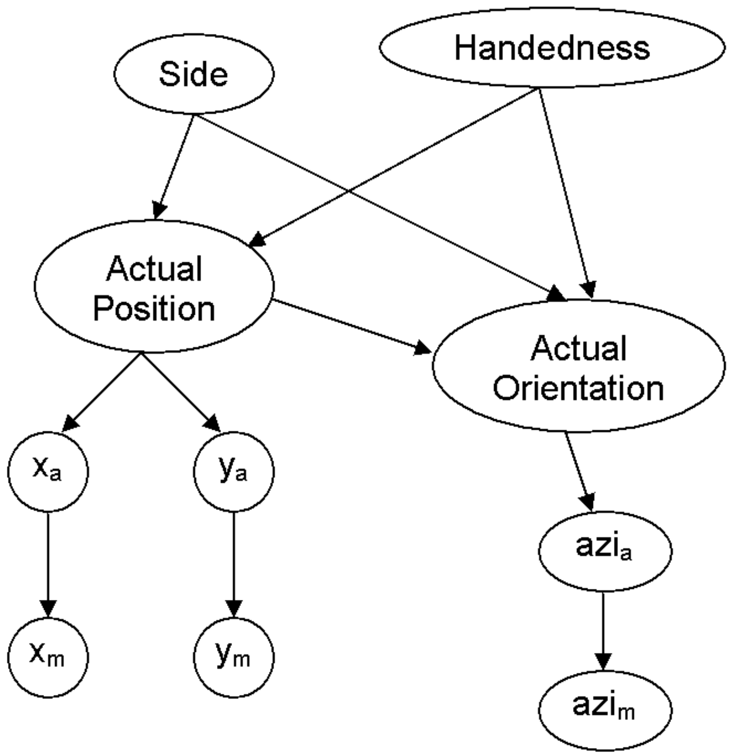

Figure 4.4: Bayesian Network Model of the Interactive Tabletop Display System

Continuous Variables

Each of the measured input coordinates (x_m,y_m,azimuth$_m$) are essentially continuous and must be discretized in some way. To achieve this discretization, each measured coordinate is modeled with with a continuous Gaussian distribution representing the likelihood of error in the device. Each actual position (x_a,y_a,azimuth$_a$) coordinate was then given a uniform distribution across a 3 by 3 grid of the horizontal display surface. The actual azimuth angle was uniformly distributed over fifteen discrete ranges of angle. The size of the grid and the number of discrete angle ranges were varied, but these values were chosen to be sufficient to obtain accurate results.

For both the x- and y-coordinates, the same Gaussian probability distribution was used. The measured coordinates are assumed to be normally distributed with a mean at the *actual* x-coordinate with a standard deviation of 10 pixels. For the azimuth angle, the measured angle is similarly normal with a mean of the actual angle, but a standard deviation relative to the measured elevation of the stylus. The standard deviation is varied because azimuth measurements are most accurate when the elevation is 0° and least accurate when the elevation is 90°.

Thus, the probabilities are as follows:

$$P(x_m|x_a) = \frac{1}{\sqrt{2\pi}\sigma} e^{\frac{-(x_a - x_m)^2}{2\sigma}} \tag{4.3}$$

(y is similar)

$$P(azi_m|azi_a) = \frac{1}{\sqrt{2\pi} f(elev_m)} e^{\frac{-(azi_a - azi_m)^2}{2f(elev_m)}}, \tag{4.4}$$

where

$$f(elev_m) = \frac{6}{\pi} elev_m \tag{4.5}$$

Elimination of Variables

The continuous variables are eliminated from the network first. Observing each measured coordinate results in the functions seen above, so no work need be done. To then eliminate the remaining continuous variables (the actual coordinates), the integral of each function is taken in the appropriate interval for each quadrant of position and each class of orientation angle. The Gaussian integral is approximated using the continued fraction:

$$\int_0^a e^{-t^2}\, dt = \frac{\sqrt{\pi}}{2} - \frac{\frac{1}{2} e^{-a^2}}{a + \frac{1}{2a + \frac{2}{a + \frac{3}{\cdots}}}} \tag{4.6}$$

After eliminating the continuous variables, the side can be observed. The variables are then eliminated in the following order: orientation, position, followed by handedness. If the side of the table is not known, the network can still be used to predict both location and handedness by eliminating variables in the following order(s): orientation, position, side, then handedness (to get handedness) and orientation, position, handedness, followed by side (to get side). The first two eliminations need only be executed once. Furthermore, the location can be determined more precisely if the handedness of the user is known by first observing the handedness before eliminating the remaining three variables.

Training the Network

The probability tables for the discrete variables are obtained using the Naive Bayesian method. The probabilities that need to be determined are:
$\forall x \in \{N, E, S, W\}, \forall y \in \{left, right\}, \forall i \in \{1..9\}, \forall j \in \{1..15\}$

- $P(side = x)$
- $P(hand = y)$
- $P(pos = i | side = x, hand = y)$
- $P(ori = j | side = x, hand = y)$

A priori probabilities for $P(side = x)$ and $P(hand = x)$ are artificially set to be equal for all values of x. The remaining tables are obtained stochastically with simple Naive equations. Thus, the formulas used to determine the discrete probabilities are:
$\forall x \in \{N, E, S, W\}, \forall y \in \{left, right\}, \forall i \in \{1..9\}, \forall j \in \{1..15\}$

$$P(side = x) = \frac{1}{4} \tag{4.7}$$

$$P(hand = y) = \frac{1}{2} \tag{4.8}$$

$$P(pos = i|side = x, hand = y) = \frac{\#I(pos = i \wedge side = x \wedge hand = y)}{\#I(side = x \wedge hand = y)} \tag{4.9}$$

$$P(ori = j|side = x, hand = y) = \frac{\#I(ori = j \wedge side = x \wedge hand = y)}{\#I(side = x \wedge hand = y)} \tag{4.10}$$

where

$\#I(p)$ = number of instances in the training corpus for which p is true.

Simplifications

Because user profile information is extracted as the user interacts with the table, the system must still respond in real-time as the information is collected. Thus, the variable elimination must be done quickly so that the user does not notice any delay in the performance of the pen-input device. For this reason, a simpler model is preferred.

Firstly, the measured z-coordinate and roll angle of the stylus is ignored. These degrees-of-freedom are not likely to be useful in determining either handedness or location. Furthermore, the elevation angle of the stylus is not directly modeled by the network, but rather included by varying the standard deviation of the error in azimuth angle in relation to the measured elevation.

Complexity

The complexity of the Bayesian network depends on both the number of edges and on the number of discretizations used. More precisely, the complexity of inference is exponential in the number of discrete variables inside each clique in the network and linear in the number of cliques. Since the size of the network is constant in this case, the complexity for a single inference is $O(1)$, as is the neural network. However, the size of the constant is much larger than for the neural network (2 cliques $\times$ 2 levels of handedness $\times$ 4 sides $\times$ 9 levels of position $\times$ 15 levels of orientation $=$ 2160). Similarly, the training has complexity $O(n)$, where n is the number of instances in the training corpus.

Advantages and Disadvantages

Similar to the neural network, the stochastic method of training is advantageous due to increased accuracy. The Bayesian network also uses a priori knowledge to determine causal relationships. The use of this knowledge makes the Bayesian network still more accurate than the neural network.

Another advantage of the structure of the network is its extensibility to include other variables, such as information from a video camera. This model can also be extended to include contextual information, such as the locations of interface components. To add such information involves the addition of another node and its associated probability table. Note, however, that adding more nodes increases the complexity of the network and thus slows down its use.

This model suffers from a larger complexity than the other two methods. Both training and use of the network are significantly slower than for the neural network. However, testing shows that the overhead involved in using this method is sufficiently small to be usable in an application.

The Bayesian network suffers the same disadvantage as the neural network over the heuristic rules. It is difficult for the user to easily identify when the network *will* fail to correctly identify handedness. It is possible that the use of a reasonable model of causation will allow a user to predict these errors better than for neural networks, but certainly not better than for the heuristic model.

4.4 Model Comparison

To test the accuracy of the three models, data were collected from 2 computer science graduate students and faculty (1 left-handed and 1 right-handed) using the tabletop display. Users were asked to use the adaptive map application (see Section4.2) on all four sides of the table. The 10-fold cross-validation technique [41] 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. The average times to make a prediction of handedness were recorded for each network.

4.4.1 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%). The Bayesian network took the longest to predict results (M = 18.04 ms, SD = 3.50 ms), followed by the neural network (M = 0.05 ms, SD = 0.01 ms).

4.4.2 Discussion

Because the Bayesian network is both more accurate and more extensible than both of the other methods, we believe it to be the best method for handedness determination. The Bayesian network has also been shown to be sufficiently fast to be usable in an application, despite its larger complexity.

The neural network appears to be an adequate alternative, and in situations that involve the need for incredibly high speed response, it may be superior. However, it is less accurate and less extensible, so for most applications, the Bayesian network is preferred.

The heuristic model appears to also have a high accuracy rate (approximately 3 errors in 125), but with regular usage of this model, the user is guaranteed to be exposed to a large number of incorrect guesses. It's primary advantage is that users can easily predict the behaviour of the system. Although the accuracy measure given here suggests a lower performance, another interpretation of the results is that users almost always adhere to this obvious pattern.

Chapter 5

Future Work and Conclusions

5.1 Future Work

We have described some of the disadvantages of an adaptable interface solution for determining handedness for pen input by having users specify this directly. We believe that an adaptive solution is best, where the system discovers users' handedness dynamically, although hybrid solutions should be explored. Future work will attempt to improve upon existing methods of handedness customization and will compare existing methods to both the improved methods and the automatic method described in this thesis.

The experiments in this thesis are specific to pen input on a variety of displays. Another common input technique is finger-based input. 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.

5.2 Conclusions

As computer technology advances from the typical single-user desktop computer with a mouse and vertical display to more sophisticated interaction techniques, possibly involving multiple users and large, shared displays, the need to support alternate input devices increases. Pen input offers an alternative to mouse input that can be used on a variety of vertical or horizontal, small or large-screen displays. Although the pen allows for the same two degree-of-freedom input as the mouse, the form factor and style of interaction can vary quite drastically. The hand posture and form factor of the pen itself can also vary between the different display setups.

The results of our experiments show specific differences in hand postures for three different display conditions. They also show that hand posture for pen input differs from that used with the mouse. These results can be used to inform the design of applications for a variety of display setups.

The analysis of the data in our experiments benefited from consideration of hand movement that is often ignored in analyses. Designers do pay attention to the effects of occlusion caused by the hand in direct pen-input applications, but may fail to recognize the degraded performance due to arm movements that could be more easily accomplished with wrist or hand movements. Future designs will benefit from attention paid to the hand and arm movements used in the target acquisitions required by the application.

The results of our experiments also show that handedness issues are of paramount importance for applications that utilize pen input. The results provide clear suggestions for how to best place targets for a variety of different display conditions so as to improve user performance for both left-handed and right-handed users. The results demonstrate that no single placement (e.g. "always above the hand") will work for both left- and right-handed users, so a system must have knowledge of which hand the user is employing for each input task.

Qualitative data also suggest that users tend not to modify default settings for handedness that already exist in desktop applications. Thus, the addition of an option with which the users can specify their preferred hand may not be sufficient. With no such option, the only way for the system to improve performance for both groups is to automatically detect the handedness of the user in some way. It may be that existing methods for modifying this default setting are somehow deficient. An improved method may encourage users to specify handedness and thus allow the application to optimize for performance. When automatic methods fail, a hybrid solution may also be a viable alternative. Such solutions have not been explored, but may prove superior to an unaided adaptation.

Although a handedness setting may be sufficient in some media, co-located, collaborative applications typically allow multiple users to share the same input device and display. In such environments, the overhead required to customize the display for one's self may be too great if input devices are frequently passed between users. The performance gained with the knowledge of handedness may be overridden by the time required to change settings. In such environments, automatic detection becomes essential. The sample map application discussed here demonstrates the feasibility of automatic detection. In particular, it is possible to automatically detect the handedness of the user at a tabletop display using our method. Other methods, such as a vision-based technique, could be used to achieve the same goal in this and other collaborative environments. Our implementation shows that simple methods may not suffice, but that the use of slightly more complex techniques may be enough to automatically predict handedness in a sufficiently accurate way.

The sample map application also demonstrates how to use handedness information to place targets on the screen. In particular, the application includes a rectangular pop-up menu system that is suitable for both left-handers and right-handers.

In the design of pop-up context menus and other interaction techniques involving target acquisition, the consideration of hand preference and posture can help to improve both user performance and user satisfaction. In this design process, one must consider the type and size of display, the input device (e.g. mouse, pen, finger) and style of interaction (direct vs. indirect input) as well as the context of use (e.g. collaborative or single-user). Only with complete knowledge of these parameters can hand posture be predicted reliably and a suitable target placement strategy be decided upon.

Bibliography

[1] J. Annett, M. Annett, P. T. W. Hudson, and A. Turner. The control of movement in the preferred and non-preferred hands. *Quarterly Journal of Experimental Psychology*, 31:641–652, 1979.

[2] R. Balakrishnan and I. S. MacKenzie. Performance differences in the fingers, wrist, and forearm in computer input control. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 303–310. ACM Press, 1997.

[3] K. Booth, W. B. Cowan, and D. R. Forsey. Multitasking support in a graphics workstation. In *First IEEE International Conference on Computer Workstations*, pages 82–89. IEEE Computer Society Press, June 1985.

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

[5] J. Callahan, D. Hopkins, M. Weiser, and B. Shneiderman. An empirical comparison of pie vs. linear menus. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 95–100. ACM Press, 1988.

[6] S. K. Card, T. Moran, and A. Newell. The keystroke-level model for user performance time with interactive systems. *Communications of the ACM*, 23(7):396–410, 1980.

[7] 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.

[8] 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.

[9] P. M. Fitts and J. R. Peterson. Information capacity of discrete motor responses. *Journal of Experimental Psychology*, 67(2):103–112, 1964.

[10] G. Fitzmaurice, A. Khan, R. Pieké, B. Buxton, and G. Kurtenbach. Tracking menus. In *Proceedings of the 16th annual ACM symposium on User interface software and technology*, pages 71–79. ACM Press, 2003.

[11] G. W. Fitzmaurice, R. Balakrishan, G. Kurtenbach, and W. Buxton. An exploration into supporting artwork orientation in the user interface. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 167–174. ACM Press, 1999.

[12] K. Flowers. Handedness and controlled movement. *British Journal of Psychology*, 66:39–52, 1975.

[13] D. R. Forsey. Transposing Harmony. Unpublished master's thesis, University of Waterloo, Department of Computer Science, 1985.

[14] N. Friedlander, K. Schlueter, and M. Mantei. Bullseye! When Fitts' law doesn't fit. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 257–264. ACM Press/Addison-Wesley Publishing Co., 1998.

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

[16] E. D. Graham. *Pointing on a Computer Display.* PhD thesis, Simon Fraser University, Department of Computer Science, 1996.

[17] E. D. Graham. Virtual pointing on a computer display: Non-linear control-display mappings. In *Proceedings of Graphics Interface*, pages 39–46, 1996.

[18] E. D. Graham and C. L. MacKenzie. Pointing on a computer display. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 314–315. ACM Press, 1995.

[19] E. D. Graham and C. L. MacKenzie. Physical versus virtual pointing. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 292–299. ACM Press, 1996.

[20] F. Guimbretière, A. Martin, and T. Winograd. Measuring FlowMenu performance. CS Technical Report CS-TR-2001-02, Stanford, September 2001.

[21] F. Guimbretière and T. Winograd. FlowMenu: Combining command, text, and data entry. In *Proceedings of the 13th annual ACM symposium on User interface software and technology*, pages 213–216. ACM Press, 2000.

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

[23] M. S. Hancock and K. S. Booth. Improving menu placement strategies for pen input. In *Proceedings of Graphics Interface*. ACM Press, 2004. In Press.

[24] D. Hopkins. The design and implementation of pie menus. `http://www.art.net/~hopkins/Don/piemenus/ddj/piemenus.html`, December 1991.

[25] P. Kabbash, I. S. MacKenzie, and W. Buxton. Human performance using computer input devices in the preferred and non-preferred hands. In *Proceedings of the conference on Human factors in computing systems*, pages 474–481. Addison-Wesley Longman Publishing Co., Inc., 1993.

[26] R. Kruger, S. Carpendale, and S. Greenberg. Collaborating over physical and electronic tables. In *Conference Supplement to CSCW 2002*, pages 139–140, 2002.

[27] R. Kruger, S. Carpendale, S. D. Scott, and S. Greenberg. How people use orientation on tables: Comprehension, coordination and communication. In *Proceedings of the 2003 international ACM SIGGROUP conference on Supporting group work*, pages 369–378. ACM Press, 2003.

[28] G. Kurtenbach and W. Buxton. The limits of expert performance using hierarchic marking menus. In *Proceedings of the conference on Human factors in computing systems*, pages 482–487. Addison-Wesley Longman Publishing Co., Inc., 1993.

[29] G. Kurtenbach and W. Buxton. User learning and performance with marking menus. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 258–264. ACM Press, 1994.

[30] 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 the SIGCHI conference on Human factors in computing systems*, pages 35–42. ACM Press, 1997.

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

[32] I. S. MacKenzie. Fitts' law as a research and design tool in human-computer interaction. *Human Computer Interaction*, 7:91–139, 1992.

[33] 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*. Morgan Kaufmann, second edition, 1995.

[34] R. L. Mandryk, S. D. Scott, and K. M. Inkpen. Display factors influencing co-located collaboration. In *Conference Supplement of CSCW 2002*, pages 137–138, 2002.

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

[36] S. Pook, E. Lecolinet, G. Vaysseix, and E. Barillot. Control menus: Execution and control in a single interactor. In *CHI '00 extended abstracts on Human factors in computing systems*, pages 263–264. ACM Press, 2000.

[37] J. Rekimoto and M. Saitoh. Augmented surfaces: A spatially continuous workspace for hybrid computing environments. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 378–385. ACM Press, 1999.

[38] R. A. Schmidt, H. Zelaznik, B. Hawkins, J. S. Frank, and J. T. Quinn. Motor-output variability: A theory for the accuracy of rapid motor acts. *Psychological Review*, 86:415–451, 1979.

[39] C. Shen, N. Lesh, B. Moghaddam, P. Beardsley, and R. S. Bardsley. Personal digital historian: user interface design. In *CHI '01 extended abstracts on Human factors in computing systems*, pages 29–30. ACM Press, 2001.

[40] M. R. Sheridan. Effects of S-R compatibility and task difficulty on unimanual movement time. *Journal of Motor Behavior*, 5:199–205, 1973.

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

[42] P. Tandler, T. Prante, C. Mller-Tomfelde, N. Streitz, and R. Steinmetz. ConnecTables: Dynamic coupling of displays for the flexible creation of shared workspaces. In *Proceedings of UIST*, pages 11–20. ACM Press, 2001.

[43] I. I. Todor and T. Doane. Handedness and hemispheric asymmetry in the control of movements. *Journal of Motor Behavior*, 10:295–300, 1978.

[44] D. Venolia and F. Neiberg. T-Cube: A fast, self-disclosing pen-based alphabet. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 265–270. ACM Press, 1994.

[45] A. T. Welford. *Fundamentals of Skill.* Methuen, 1968.

[46] P. Wellner. Interacting with paper on the digital desk. *Communications of the ACM*, 36(7):87–96, 1993.

[47] R. S. Woodworth. The accuracy of voluntary movements. *Psychological Review [monograph supplement]*, 3(2, Whole No. 13), 1899.

[48] S. Zhai, P. Milgram, and W. Buxton. The influence of muscle groups on performance of multiple degree-of-freedom input. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 308–315. ACM Press, 1996.

Appendix A

Questionnaires

A.1 Study 1 Background Questionnaire

Background Questionnaire

1. Gender:

- ○ Male
- ○ Female

2. Which is your dominant writing hand?

- ○ Left
- ○ Right

3. How often do you use:

	Never	Rarely	Sometimes	Often	Always
Computers (PC, handheld, laptop...)?	○	○	○	○	○
A mouse when working on a computer?	○	○	○	○	○
A stylus when working on a computer?	○	○	○	○	○

4. Do you use computers for (please select as many as appropriate):

- ☐ Fun
- ☐ At Work
- ☐ School
- ☐ Work at Home

6. List any other input devices you commonly use when working on a computer:

7. Have you ever interacted with a tabletop display before?

- ○ Yes
- ○ No

7b. If yes, describe the system and where you used it:

8. What is your age?

A.2 Study 2 Background Questionnaire

<u>Background Questionnaire</u>

Age: ______
Sex (circle one): Male Female

1. Do you normally write with your left or right hand (circle one)? Left Right

2. Do you normally use a mouse with your left or right hand (circle one)? Left Right

3.
   a. Do you have any experience using large screen displays, such as electronic whiteboards or tabletop displays?

      Yes No

   b. If yes, please briefly describe the system that you used:
   ____
   ____
   ____
   ____

4. How often do you use:

	Never	Rarely	Sometimes	Often	Always
Computers (PC, handheld, laptop...)?	1	2	3	4	5
A mouse when working on a computer?	1	2	3	4	5
A stylus when working on a computer?	1	2	3	4	5
Pop-up menus?	1	2	3	4	5

5. Name the 3 most common applications or games in which you use pop-up menus?

   1) ___________________________
   2) ___________________________
   3) ___________________________

A.3 Study 2 Post-experimental Questionnaire

Post-Experimental Questionnaire

1.
   a. Describe what you liked about the tabletop display:

   b. Describe what you disliked about the tabletop display:

2.
   a. Describe what you liked about the menu system in the Map application:

   b. Describe what you disliked about the menu system in the Map application:

3.

a. (Please proceed to the tabletop display once more to complete this section of the questionnaire)

   In each of the 6 boxes displayed on the screen, please touch the area with the stylus. 4 menus should appear surrounding the location that you touched. Touch each menu in order of preference (1=most preferred, 4=least preferred). A number will appear in place of the menu to indicate your preference. You may redo a selection by touching the menu again. For each position, once you have completed filling in your preference, please wait for the experimenter to clear the screen.

   Using the following scale, for each menu location in each position, please indicate whether or not you agree or disagree with the following statement:

   The menu placement is suitable for use in a tabletop display application.

   Scale:

Strongly Agree	**Agree**	**Neutral**	**Disagree**	**Strongly Disagree**
1	2	3	4	5

   Write the appropriate number in the corresponding box below:

Position 1:		Position 2:		Position 3:	
☐	☐	☐	☐	☐	☐
☐	☐	☐	☐	☐	☐

Position 4:		Position 5:		Position 6:	
☐	☐	☐	☐	☐	☐
☐	☐	☐	☐	☐	☐

b. Please comment on your choice of menu placement:

   ____
   ____
   ____

A.4 Study 3 Background Questionnaire

Background Questionnaire

Age: ______
Sex (circle one): Male Female

1. Please state your preferred hand for the following activities:

	Left Only	Left Mostly	Right Mostly	Right Only	No Experience
Writing	1	2	3	4	5
Drawing	1	2	3	4	5
Throwing	1	2	3	4	5
Scissors	1	2	3	4	5
Toothbrush	1	2	3	4	5
Knife (without a fork)	1	2	3	4	5
Spoon	1	2	3	4	5
Broom (upper hand)	1	2	3	4	5
Striking a match	1	2	3	4	5
Opening a lid	1	2	3	4	5
Mouse	1	2	3	4	5

2. a. Do you have any experience using large screen displays, such as electronic whiteboards or tabletop displays?

   Yes No

   b. If yes, please briefly describe the system that you used:

   ____
   ____
   ____
   ____

3. How often do you use:

	Never	Rarely	Sometimes	Often	Always
Computers (PC, handheld, laptop...)?	1	2	3	4	5
A Tablet PC?	1	2	3	4	5
A mouse when working on a computer?	1	2	3	4	5
A stylus when working on a computer?	1	2	3	4	5
Pop-up menus?	1	2	3	4	5

A.5 Study 3 Post-experimental Questionnaire

Preference and Suitability Questionnaire

For this portion of the experiment, the experimenter will direct you on how to complete the following questions using the input technique for the trial set that you have just completed.

1. *User Preference*
For each menu location in each region, please indicate the order you would prefer to have the menus appear (1=most preferred, 4=least preferred).

2. *Suitability of Menu Placement*
Using the following scale, for each menu location in each region, please indicate whether or not you agree or disagree with the following statement:

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

Strongly Agree	Agree	Neutral	Disagree	Strongly Disagree
1	2	3	4	5

Appendix B

Alternate Analysis for Study 3

B.1 Results

There was a main effect of display ($F(2, 372) = 788.6$, $p < .001$). Post-hoc analysis showed that a horizontal display with direct input was significantly faster ($M =$ 680ms, $SD = 4$ ms) than both the vertical display with direct input ($p < .001$) and the vertical display with indirect input ($p < .001$) and that the vertical-direct condition ($M = 708$ ms, $SD = 6$ ms) had significantly faster ($p < .001$) selection times than the indirect condition ($M = 1009$ ms, $SD = 13$ ms). A main effect of handedness ($F(1, 186) = 11.9$, $p = .001$) showed that left-handed participants selected targets significantly slower ($M = 822$ms, $SD = 9$ ms) than right-handed participants ($M = 776$ ms, $SD = 9$ ms). There was no significant main effect of target position.

A two-way interaction between display and handedness ($F(2, 372) = 10.2$, $p < .001$) suggested that the effect of handedness depends on the display condition. Post-hoc analysis revealed that left-handed participants were significantly slower than right-handed participants in both the indirect condition ($p = .001$) and the vertical-direct condition ($p < .001$), but that this difference was not significant for the horizontal-direct condition ($p = .442$).

There was a two-way interaction between display and target position ($F(22, 4092) = 2.1$, $p = .002$). This interaction suggested that the effect of display condition depends on the position of the target.

A two-way interaction effect between target position and handedness ($F(11, 2046) = 6.3$, $p < .001$) suggested that the effect of position depends on the handedness of the participant.

A three-way interaction ($F(22, 4092) = 2.6$, $p < .001$) suggested that the two-way interaction between target position and handedness depends on the particular display condition. Post-hoc analysis revealed more significant differences in the horizontal-direct condition and the least significant differences using indirect input.

B.2 Adjustment for Perceived Target Width

In order 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. The only notable difference appears in the interaction effect between handedness and display condition. The post-hoc analysis of this interaction effect revealed left-handed participants were significantly slower than right-handed participants in both the indirect condition ($p < .001$) and the vertical-direct condition ($p < .001$), but that this difference was not significant for the horizontal-direct condition ($p = .191$), as before. However, the mean difference in throughput for the indirect condition was smaller (0.17 bits/s) than for the vertical-direct condition (0.28 bits/s).