Understanding and Improving Distal Pointing Interaction

Régis Augusto Poli Kopper

A dissertation submitted to the faculty of
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
In
Computer Science and Applications

Doug A. Bowman, Chair
Ravin Balakrishnan
D. Scott McCrickard
Chris North
Deborah Tatar

June 9, 2011
Blacksburg, Virginia

Understanding and Improving Distal Pointing Interaction

Régis Augusto Poli Kopper

Abstract

Distal pointing is the interaction style defined by directly pointing at targets from a distance. It follows a laser pointer metaphor and the position of the cursor is determined by the intersection of a vector extending the pointing device with the display surface. Distal pointing as a basic interaction style poses several challenges for the user, mainly because of the lack of precision humans have when using it.

The focus of this thesis is to understand and improve distal pointing, making it a viable interaction metaphor to be used in a wide variety of applications. We achieve this by proposing and validating a predictive model of distal pointing that is inspired by Fitts' law, but which contains some unique features. The difficulty of a distal pointing task is best described by the angular size of the target and the angular distance that the cursor needs to go across to reach the target from the input device perspective. The practical impact of this is that the user's relative position to the target should be taken into account. Based on the model we derived, we proposed a set of design guidelines for high-precision distal pointing techniques. The main guideline from the model is that increasing the target size is much more important than reducing the distance to the target.

In order to improve distal pointing, we followed the model guidelines and designed interaction techniques that aim at improving the precision of distal pointing tasks. Absolute and Relative Mapping (ARM) distal pointing increases precision by offering the user a toggle which changes the control/display (CD) ratio such that a large movement of the input device is mapped to a small movement of the cursor. Dynamic Control Display Ratio (DyCoDiR) automatically increases distal pointing precision, as the user needs it. DyCoDiR takes into account the user distance to the interaction area and the speed at which the user moves the input device to dynamically calculate an increased CD ratio, making the action more precise the steadier the user tries to be. We performed an evaluation of ARM and DyCoDiR comparing them to basic distal pointing in a realistic context. In this experiment, we also provided variations of the techniques which increased the visual perception of targets through zooming in the area around the cursor when precision was needed. Results from the study show that ARM and DyCoDiR are significantly faster and more accurate than basic distal pointing with tasks that require very high precision. We analyzed user navigation strategies and found that the high precision techniques afford users to remain stationary while performing interactions. However, we also found that individual differences have a strong impact on the decision to walk or not, and that, sometimes, is more important than the technique affordance. We provided a validation for the distal pointing model through the analysis of expected difficulty of distal pointing tasks in light of each technique tested.

We propose selection by progressive refinement, a new design concept for distal pointing selection techniques, whose goal is to offer the ability to achieve near perfect accuracy in selection at very cluttered environments. The idea of selection by progressive refinement is to gradually eliminate possible targets from the set of selectable objects until only one object is available for selection. We implemented SQUAD, a selection by progressive refinement distal pointing technique, and performed a controlled experiment comparing it to basic distal pointing. We found that there is a clear tradeoff between immediate selections that require high precision and selections by progressive refinement which always require low precision. We validated the model by fitting the distal pointing data and proposed a new model, which has a linear growth in time, for SQUAD selection.

Keywords: HCI models of human motor behavior, Fitts' law, distal pointing, 3D interaction, 3D selection, progressive refinement.

Acknowledgments

Coming to the completion of a doctoral research is an achievement that I don't claim to be mine alone. Many people helped in different ways to make this journey possible.

I would like to thank my beloved wife, Cristiane, for her unconditional support throughout the five years of this research. She was not only there for me to celebrate the good times of accomplishments, but also and most importantly in the difficult times of struggle. To my parents, Lauro and Maria de Lourdes, thank you for giving me my most basic education and always encouraging me to be inquisitive and to fight for what I want. I would also like to thank my older brother Roberto and sister Patricia for serving as my role models.

I would like to express my special gratitude for Dr. Doug Bowman, my research advisor. His guidance was fundamental for allowing me to keep in the right track when I wasn't sure where to go. Without his encouragement, understanding and patience, this work would have been so much more difficult, if even possible, to achieve. I would also like to thank the members of my dissertation committee, Dr. Ravin Balakrishnan, Dr. Scott McCrickard, Dr. Chris North and Dr. Deborah Tatar for their always helpful comments, guidance and discussions.

I also want to thank the members of the 3DI group at Virginia Tech. Thanks to Ryan McMahan, Mara G. Silva and Felipe Bacim for the collaborations we had in different parts of the work presented in this thesis. Thanks to Eric Ragan for all the helpful discussions and for some of the most hilarious alternative presentation slides one could ever hope. I would like to also thank the members of the Center for Human-Computer Interaction at Virginia Tech for the great environment I had to work and play.

Table of Contents

List of Figures

List of Tables

Definition of Key Terms

Before we begin the main part of this thesis, we introduce key terms that are used throughout this document. These definitions are intended to help the reader to become familiar with and understand the terms in the context in which they are used.

Distal Pointing: An interaction metaphor in which the user interacts by pointing directly at the display from a distance, using a laser pointer metaphor. The metaphor allows users to in teract while standing up and walking in the area in front of the display.

Predictive Model: A formula used to allow the analysis of tasks based on parameters that determine the time it takes for a given task to be performed. The outcome of the model is the predicted time it takes for a task with parameters input into the formula to be completed.

Motor Behavior: A description of the actions performed by one or more of a user's muscles required to perform a physical task.

Fitts' law: A predictive human motor behavior model used in human-computer interaction (HCI) to predict task performance with two-dimensional input devices. It predicts the time it takes to point at a target of as a function of target width (W) and distance to target (A). Its general formulation is $MT = a + bID$, where MT is the movement time to complete an aimed pointing task.

Index of Difficulty (ID): The term that determines the difficulty of a Fitts' Law task. In HCI it is most often described as $\log_2\left(\frac{A}{W}+1\right)$.

Index of Difficulty of Distal Pointing Tasks (ID_{DP}): The difficulty term of the distal pointing model we derived. It is based on the relationship between the angular distance to the target (α) and the angular target size (ω) and is defined as $\left(\log_2\left(\frac{\alpha}{\omega^k}+1\right)\right)^2$, where k is an empirically determined constant.

Physical Navigation: Navigation by physical body movements, such as walking. It is distinguished from virtual navigation, in which the control is done indirectly, through the use of, for example, a joystick.

Progressive Refinement: An interaction style characterized by the gradual reduction of the set of possible objects until the target is the only one left. Progressive refinement interaction techniques often do not require the user to act with precision at any point during the task.

Control/Display (CD) Ratio: The mapping of the input device movement to cursor movement. With distal pointing, the CD ratio is often 1, as the position of the cursor is determined by the intersection of the vector extending the input device with the display. However, the CD ratio can be increased to allow for a larger movement of the input device to cause a smaller movement of the cursor, providing more precision.

Visual Acuity: The ability a person has to visibly distinguish objects that are in sight. It is related to the spatial resolution of the visual processing system. In this thesis, we define visual acuity as the ability to correctly perceive the cursor and targets on a display.

Large High-Resolution Display: A wall-sized display that provides large visualizations and also offers high levels of detail. Usually, large high-resolution displays are achieved by combining a large number of desktop monitors.

Chapter 1

Introduction

1.1 Problem Statement

Recently, new types of interactive systems and applications allow and afford users to stand while interacting. Examples of such systems include very large high-resolution displays (e.g., Vogel and Balakrishnan, 2005), and video games (e.g., Nintendo Wii, 2011; Playstation Move, 2011)). Users of large high-resolution displays can especially benefit from standing and walking while using applications, as very detailed visualizations allow users to physically navigate the data to understand it better (Ball et al., 2007). Users can also stand and walk while directly interacting with the environment (Simon et al., 2008; Wilson and Shafer, 2003).

Interfaces that use traditional input devices, such as mice, trackballs, and track pads, however, discourage this physical navigation, since these devices need to be on a flat surface to work properly, thus hindering the possibilities for the user to move around freely. Particularly with very high-resolution displays, traditional interaction techniques pose additional disadvantages. One of the characteristics of these displays is that they can accommodate massive amounts of data, and interacting using traditional input devices can be challenging. Figure 1.1 shows the NASA Hyper Wall 2, reportedly the world's highest resolution visualization system (Mewhinney and Fay, 2008). We can see the problems of using a mouse to interact with such a large display. At a very high resolution, it may be very easy to lose track of the cursor position, and either too much clutching or high levels of cursor acceleration may be necessary to move the cursor between two distant areas of the screen. Too much clutching leads to user fatigue and slow interaction, while high acceleration levels result in a loss of precision.

One of the main uses of very large high-resolution displays is to analyze and make sense of data at multiple levels of scale by physically walking in order to perceive details that otherwise couldn't be seen from a distance (Yost et al., 2007). However, while visualizing and analyzing the data is a primary task with high-resolution visualizations, other interactions are often necessary. It has been suggested that large high-resolution displays require novel types of interaction techniques (Ni et al., 2006).

Figure 1.1 An example of a very large high-resolution display.

1.1.1 Distal Pointing

In many cases, these novel interaction techniques are based on advanced input devices such as six degrees-of-freedom (DOF) position trackers, which allow the user to stand up, walk around, and interact by pointing directly at the display from a distance. We call this interaction *distal pointing*. Distal pointing follows a laser pointing metaphor and is based on ray-casting (Bowman et al., 2004), in which the intersection of a vector coming from the input device with the display determines the cursor position. Distal pointing is an alternative pointing metaphor that enables users to move around freely while still being able to interact with the information. Distal pointing is very effective for interacting with large spaces, as large areas can be covered with a quick rotation of the wrist, and there is no delay related to rolling and clutching an indirect input device, such as a mouse.

One important point to make is that distal pointing is beneficial not only to very large high-resolution display interaction. Other domains such as 3D virtual environments (Mine, 1995) and direct interaction with the environment, as envisioned by the xWand concept (Wilson and Shafer, 2003), can also benefit from distal pointing. In fact, distal pointing can be used in any application scenario that allows the user to be standing and moving while performing interaction tasks.

1.1.2 Usability Issues

Distal pointing is a viable alternative to traditional techniques that do not provide the same mobility, but we acknowledge that there are issues that limit its usability in many situations. The basic form of distal pointing has several usability issues related to different factors, which we present below.

- *Hand tremor* – The hand has a physiological tremor around 8-12 Hz (Riviere et al., 1997). This tremor causes jitter in the input device, which is transferred to the cursor position. If the input device is very near to the display, causing the interaction to be touch-screen-like, hand tremor does not cause a large problem. However, as the interaction becomes more distant from the display, the linear error in the cursor position caused by hand tremor grows exponentially, as it is related to the intersection of the ray extending from the input device with the screen. Since the control of the cursor with distal pointing is done by wrist rotations, even a tiny unintentional rotation of the wrist caused by natural hand tremor can result in a large movement of the cursor on the display.

- *No parkability* – A mouse is considered *parkable*, in that the user can release the device without causing any change in the cursor position. With distal pointing, however, the user needs to hold the input device in place for as long as the controlled cursor needs to be in a particular position. The user needs to actively point at the screen for the cursor to be displayed. This can potentially cause fatigue, as the user's arm needs to be held up for as long as the task requires. If the user drops her arm, the cursor disappears, since the intersection of the ray falls off the display.

- *Lack of a supporting surface* – The mouse allows small adjustments to be made by using the desk surface as a support for finer movements, but this is not available with distal pointing. At best, the user can take advantage of his other hand to secure his grip on the input device. However, the support of the other hand will only increase steadiness so much, and is not available if the interaction technique requires both hands.

- *Difficulty to acquire small targets from a remote position* – With distal pointing, the effective size of a target decreases as the user moves far from the display surface due to the reduction of the angle within which the intersection of the ray determined by the direction of pointing and the screen is inside the target area.

- *Heisenberg effect* – The so-called Heisenberg effect (Bowman et al., 2002) refers to the disturbance caused to the tracked input device when a discrete input (e.g., a button press) is performed by the user. This problem is particularly important for interactions with small visual targets. The cursor may be at the right position when the target is active for selection, but as the user starts to press the button, the input device moves, ultimately causing the target to be missed.

1.1.3 Basic Enhancements to Distal Pointing

There are some basic precision enhancements that can be applied to distal pointing without changing the basic laser pointer metaphor. A common enhancement to reduce jitter caused by hand tremor and tracking error involves filtering the low frequencies caused by the input device. A typical approach is to use low-pass Kalman filters, which filter out any movement of the input device whose frequency is below a threshold, causing the cursor to appear smoother (Welch and Bishop, 1995). However, this method often causes "stickiness" of the cursor, as slow, precise movements can be mistaken for noise and be eliminated. To eliminate this problem, Vogel and Balakrishnan (2005)) proposed the use of dynamic recursive low-pass filters, which use a linear interpolation based on the velocity of the input device to determine the amount of jitter to be filtered out.

Another simple enhancement can be used to reduce the Heisenberg effect. If we can estimate how long it takes for a button to be pressed, we can use the position of the cursor right before the click action started, overcoming the afore mentioned problem of the unintentional displacement of the cursor.

More elaborate enhancements can be added to basic distal pointing in order to further increase precision. These may include area cursors rather than point cursors (Tse et al., 2007), the ability to zoom in on a particular area of the display (Forlines et al., 2005), or the ability to control cursor speed (Vogel and Balakrishnan, 2005). Chapter 2 provides a thorough review of existing distal pointing interaction techniques.

1.1.4 Design Challenges

Distal pointing can be a great platform for interaction in open spaces and with large high-resolution displays. However, interaction by distal pointing is more complex than interaction by traditional, mouse-based techniques, and designing effective distal pointing interaction techniques can be a challenging endeavor.

While traditional techniques use an input device with a limited input space on a flat surface, distal pointing affords free three-dimensional (3D) movements of the input device. The user almost always stays at a fixed position – usually sitting – when interacting with traditional devices. With distal pointing, however, the user is not constrained to any fixed location and can move freely relative to the display and the interaction region while performing tasks. This is a design challenge because the lack of constraints gives users many ways of performing the same task, and that can have a direct impact on the expected success or failure of the task.

As an example, consider the task of moving an object across the display using distal pointing. One possible strategy employed by the user would be to walk towards the object, perform the selection and walk along the display to the target region while dragging the object, and finally dropping the object to finish the task. Another way of performing the same task would be for the user to start at a distance from the display, select the object and drag it while standing still, and finally drop the object at the target region without ever taking a step.

We can see how these two ways of completing the task can have very difficult outcomes. The user who chose to walk to complete the task may have taken longer, but was more likely to be successful than the user who completed the task at a distance, since the objects were visually larger and easier to select. The user who was at a distance could quickly point from one side of the display to the other with a mere wrist rotation, but the refinements necessary to complete the task from a distance may have increased the chances of error.

It is also important to note that the strategy involved in completing a distal pointing task may not only be related to user preferences, but also to the application or to the task. For example, a task may involve reading small labels on objects to decide what to select. That would encourage physical navigation. On the other hand, a task involving a search of visual patterns in a visualization to decide what task to perform encourages the user to remain at a distance, keeping an overview of the data.

Apart from the varied strategies that can be employed with distal pointing, this interaction style also affords other input modalities that are usually not encountered with traditional techniques. For example, distal pointing techniques may involve bi-manual interaction and non-linear mappings in order to provide more precision to the user.

All these possibilities make designing an effective distal pointing technique quite challenging, and designers would benefit from understanding the important factors affecting performance with distal pointing interaction and the tradeoffs involved in various design choices. Prior to the work presented in this thesis, there has not been any systematic approach to understanding distal pointing that provides a solution proven to be usable for high-precision interaction.

1.1.5 Evaluation Challenges

Another aspect that makes distal pointing a challenging interaction metaphor is the difficulty involved in evaluating distal pointing techniques. Because of the range of strategies that can be used to complete any given distal pointing task, it is hard to establish controlled environments for the comparative evaluation of distal pointing-based techniques.

This spread of strategies is not limited to physical navigation, but also to technique-specific features. For example, with certain techniques the user may need to decide when and where to zoom in, or when and how to change the cursor speed.

Contrary to traditional, mouse-based interfaces in which a pointing task depends mostly on task-specific features, such as target sizes and movement distances, with distal pointing, the strategy the user takes can have a direct impact on the difficulty and performance of a pointing task. Performance will depend on whether the user chose the best strategy, and individual differences may be large.

1.2 Research Questions

The problem statement above opens up a range of opportunities towards the goal of understanding and improving distal pointing interaction. Based on the analysis of the problem, we defined five broad research questions, which we addressed through the research presented in this thesis:

1. **How can we effectively understand human motor behavior with distal pointing interaction?**

   This is the fundamental question of this work.

   Observing and analyzing users' motor behavior with distal pointing provided us with knowledge about the primary factors that determine the difficulty and can predict the performance of distal pointing tasks. This led to our model of distal pointing and offered us the base knowledge necessary to answer the remaining research questions.

   The goal of this question was to provide objective knowledge, through a predictive model of performance, about the tradeoffs involved with distal pointing interaction. Rather than looking at the physiological issues involved with distal pointing motor behavior, we took a more pragmatic approach, aiming at understanding how different parameters of a distal pointing task (e.g., target sizes, user physical position) affected user performance.

2. **What guidelines for designing high-precision distal pointing interaction techniques can we pose based on the model we developed?**

   This question aims to provide designers and researchers with take-home messages from the model.

   Practically, what are the implications of the model? How can a designer use the knowledge contained in the model to improve distal pointing interaction? These questions can be answered by providing design guidelines, i.e., recommendations for the design of best practices distal pointing interaction techniques.

3. **Can we design effective high-precision distal pointing interaction techniques based on the guidelines created in response to question 2?**

   This question aims to validate the design guidelines and provide the community with a set of effective interaction techniques for precise distal pointing.

   Following the design guidelines, are we able to create interaction techniques that, in fact, improve distal pointing precision? By addressing this question, we provided evidence that the design guidelines succeed in providing direction for the design of high-precision distal pointing techniques.

4. **Will the analytic evaluation of distal pointing motor behaviors, using the model developed in response to question 1, reflect actual user performance when completing tasks with interaction techniques based on distal pointing?**

This question aims to provide evidence of the validity of the model in realistic contexts.

Given the predictive model of performance developed in response to research question 1, we are able to assess different techniques without the need to run empirical studies. We can do so analytically, based on the knowledge provided by the model. To validate the model, however, we need to confirm that empirical performance data matches the predictions of the model.

5. **How do the high-precision techniques designed in response to question 3 affect user strategy and performance when used in a realistic task setting?**

   The answer to this provided knowledge about individual differences related to distal pointing interaction as well as a validation of the affordances provided by the high-precision features of the interaction techniques designed in response to question 5.

   What do we find when we evaluate the high-precision techniques from question 3? With this question, we aimed at understanding the relationship between the affordances provided by different techniques and individual differences, particularly in terms of physical navigation.

1.3 Approach

The research presented in this thesis was carried out through three main instruments: analysis, evaluation and design. We analyzed distal pointing characteristics to determine candidate models for distal pointing. A series of user studies were conducted and statistical analysis was performed in order to gather evidence deemed sufficient to address the research questions. Based on the guidelines developed from the model, three high-precision distal pointing interaction techniques were designed and evaluated through user studies.

Research question 1 was addressed by the definition of a predictive model of performance for distal pointing tasks. The model was inspired by Fitts' law (Fitts, 1954), and contains some fundamental differences from model formulations intended for two-dimensional inputs. First, the user's relative position to the target on the display needs to be taken into account. Thus, angular measurements are preferred over the linear ones used for modeling traditional interaction. The other main difference from other Fitts' law-like models is that the target size is weighed much more heavily than the movement amplitude. The analysis of the tradeoffs presented in the (ID_{DP}) of the distal pointing model led to the development of the design guidelines in response to research question 2. The research conducted in this phase is detailed in chapter 3.

To address research question 3 we developed a set of high-precision interaction techniques that were designed based on the model guidelines. We also discovered a new concept for distal pointing selection, which we called Selection by Progressive Refinement (section 4.7). Of the interaction techniques we designed, three are aimed at increasing the performance of distal pointing by varying the CD ratio – ARM (section 4.5) and DyCoDiR (section 4.6) – and/or increasing visual size to address visual acuity limitations – ZELDA (section 4.4.2) and zoom variations of ARM and

DyCoDiR. Another interaction technique we designed to increase precision of distal pointing tasks is SQUAD (section 4.7.2), a selection technique aimed at maximizing the accuracy and increasing performance of selection of very small objects in cluttered environments. SQUAD “beats” the distal pointing model in that target size is not a factor in task difficulty.

We demonstrated the effectiveness of the techniques we designed by performing comparisons of the high-precision techniques to basic distal pointing. Chapter 5 presents the study of the immediate techniques, while chapter 6 describes the progressive refinement selection experiment.

Research question 4 was tackled by analyzing the results of both the immediate (section 5.9) and the progressive refinement (section 6.4) techniques. While we were able to accurately validate the model in terms of time predicted to complete a task in the progressive refinement study, we did not achieve full validation based on the results of the immediate techniques experiment. This was likely due to the extreme precision required by the tasks in this study, and we discuss this in section 5.9.

Finally, we addressed research question 5 by performing a detailed analysis of the strategies employed by users of the high-precision immediate distal pointing techniques as compared to basic distal pointing. Activation of high precision mode was analyzed as a component of user strategy (section 5.8.2) for the ARM techniques. We also analyzed the physical navigation strategies (section 5.8.4) employed per technique and per participant in order to determine what factors affected the choice to walk or not.

1.4 Contributions

We can frame the contributions of this work into three main categories, related to the problem statement presented in section 1.1.

- Understanding of distal pointing

    1. We provide a predictive model of human motor behavior for distal pointing tasks, which was derived through experimentation (Kopper et al., 2010).
    2. We propose the concept of selection by progressive refinement, which provides a design space for accurate selection of targets that are otherwise very difficult to select with immediate distal pointing techniques (Kopper et al., 2011).
    3. We offer a set of design guidelines based on knowledge gathered from the model to aid in the design of distal pointing-based interaction techniques.
    4. We demonstrate that increasing visual size to address visual acuity limitations has an impact on distal pointing performance, in addition to increasing motor precision.

- Design of distal pointing techniques

1. We provide a set of high-precision distal pointing interaction techniques that have been empirically verified to be effective.
2. We offer the design of SQUAD, a selection by progressive refinement distal pointing technique whose performance is not affected by target size and successfully maximizes accuracy while keeping selection fast and scalable.
3. We demonstrate that it is possible to design interaction techniques using progressive refinement that allow lazy yet accurate interaction in very complex environments.
4. We show that high-precision distal pointing techniques afford interaction at a distance, but also that individual differences play a large role in the physical navigation strategy, independent of the interaction technique.

- Evaluation of distal pointing techniques

  1. We show that the distal pointing model can be successfully used to analytically evaluate distal pointing interaction techniques.
  2. We provide validation for the predictive model of performance in a selection by progressive refinement distal pointing technique.

Chapter 2

Literature Review

This chapter covers the literature related to the research presented in this thesis.

2.1 Models of Human Motor Behavior

It has been argued that the HCI problems and challenges should be reduced to well formed and rigid models in order to have a hard science of interface design (Newell and Card, 1985). However much discussion has been carried in the literature on whether the above concept is valid (Carroll and Campbell, 1986; Newell and Card, 1985, 1986), in reality, very few models are robust enough to make interface design an exact science. Models of human motor behavior have shown to be one exception.

Most human motor behavior models used in HCI provide guidelines for the design of desktop interfaces, and predict time and difficulty of tasks using traditional interaction techniques. When it comes to novel interfaces that require other types of input, there is little development of human motor behavior models. Exceptions are a model of interaction for volumetric displays (Grossman and Balakrishnan, 2004) and a model of ray-casting in an immersive 3D virtual environment (VE) (Wingrave and Bowman, 2005). Wingrave and Bowman (2005) studied the factors influencing ray-casting tasks and, although it presents a step towards a model for distal pointing tasks, it's not generalizable outside immersive ray-casting interaction and was not validated. As much as we can tell from literature review, to this date, no general models of motor behavior for interactive tasks that involve distal pointing have been studied. There is even an urge in the community for the definition of such models (Grossman, 2008).

Understanding how humans behave in distal pointing tasks can also guide the design of interfaces in several domain areas. For example, new gaming systems, such as the Wii, employ distal pointing techniques and ray-casting selection in 3D VEs and, although with different perceptual cues, they use the same motor aspects of distal pointing. Further, large high-resolution displays provide applications for areas such as geospatial visualization and visual analytics. Nonetheless, large displays are not intended only for visualizations, but also for interactive applications in which data can be manipulated across the entire display.

Most of the models that have been proposed for pointing and movement tasks are flavors or extensions of Fitts' law (Accot and Zhai, 1997, 2003; MacKenzie, 1992; Mackenzie and Buxton, 1992). The goal of these models is to predict the time to acquire a target as a function of different factors, such as the size of the target, the width of the path and the amplitude of the movement. Grossman and Balakrishnan (2005b) proposed a different approach, using a probabilistic model for the prediction of 2D target acquisition time. Although a probabilistic model of distal pointing can be useful as future work, we believe that a univariate model is fundamental for understanding the factors that determine the difficulty of distal pointing tasks. Most prior research has focused on models in which the user either touches a target directly, or translates an input device, to cause a proportional translation of the cursor. In distal pointing, however, different types of movement are used (especially involving wrist rotation), and both the position and orientation of the input device determine the position of the cursor on the display.

In this section, we provide an overview of human motor behavior models that have been proposed over the years and how they evolved to become more robust and expressive.

2.1.1 An Overview of Motor Behavior Modeling

Modeling human motor behavior has been the subject of research for over fifty years. Fitts (1954) was the first to study the information capacity of the human motor system. He hypothesized that the information theory (Shannon and Weaver, 1949), with its roots in communications, applies to people when performing tasks that only involved their motor system. Fitts defines the information capacity of the motor system as "its ability to produce consistently one class of movement from among several alternative movement classes". Since these movements are dependent only on measurable values, such as the direction and amplitude of the motor response, the information capacity can be inferred from the variability of the measurement of successive responses to different types of movements. What Fitts suggested is that it is possible to determine a constant variation of the human motor response time as the movement amplitude and width of the target movement point vary.

Fitts was motivated by the need of unifying the concept of motor capacity, which was then not well understood as the interrelation of the movement duration with its amplitude and target width was not consensual. Card et al. (1978) were the first to bring Fitts' law to HCI, as they used the model to compare different input devices. Providing means for comparison of input devices is one of the benefits of Fitts' law in HCI. The development of guidelines for the design of user interfaces is the other. Other models based on and inspired by Fitts' law focusing in HCI have been proposed. In this section, we detail Fitts' law and review the literature on models of human motor behavior for HCI.

Fitts' Law

Fitts' law is a predictive model of human motor behavior for pointing tasks. It is defined as a function of the amplitude of the movement and the size of a target to be acquired. In the work that introduced the original Fitts' law (Fitts, 1954), Paul Fitts performed three experiments that measured the time and error rate of subjects when performing motor tasks. In all experiments, the subject was instructed to perform several repeated movements as fast as possible maintaining accuracy. The tasks in each experiment were, respectively, reciprocal tapping two plates of width W separated by amplitude A; disc transfer between two pins, with disc hole of diameter W and distance between pins A; and pin transfer between two holes of diameter W separated by distance A. A regression analysis of the trials provided the following speed-accuracy tradeoff for the mean time to complete the task:

$$MT = a + b\log_2\left(\frac{2A}{W}\right), \tag{2.1}$$

where MT is the mean movement time and a and b are empirically determined constants. The logarithmic factor is referred to as the Index of Difficulty (ID) of a task. After the original proposition of the Fitts' law, numerous studies were carried out to find a better data-to-model fit. Currently, the most widely adopted form of Fitts' law is the one proposed by MacKenzie (1992), which is a direct analogy to Shannon's Theorem 17 of information theory (Shannon and Weaver, 1949). In the new formulation, ID is defined as

$$MT = a + b\log_2\left(\frac{A}{W} + 1\right), \tag{2.2}$$

and solves several issues existent in the ID of the original formulation (Equation 2.1). With the new formulation (Equation 2.2), it is impossible to get a negative ID, thus solving the problem of the model theoretically allowing a negative value of MT.

Fitts' law has been applied to the HCI domain with success (Card et al., 1978) for tasks in which only the initial cursor position and a one-dimensional target are taken in consideration. Other refined models were proposed in order to describe more complex tasks.

Refined Models of Fitts' Law

While Fitts' law successfully models human motor behavior of reciprocal selection tasks with a mouse (Card et al., 1978), it only takes one dimension of the target in consideration.

The first model to consider two dimensions of targets was proposed by Crossman (1956). Hoffmann and Sheikh (1994) reported that Crossman (1956) found that the target height had an effect similar to the width, and proposed the following model for two dimensional tasks:

$$MT = a + b\log_2\left(\frac{2A}{W}\right) + c\log_2\left(\frac{2A}{H}\right), \tag{2.3}$$

with a, b and c being constants. Hoffmann and Sheikh (1994) disagrees with such model and propose that width and height should interact, in the sense that as one gets higher, it reduces or eliminates the effect of the other.

Mackenzie and Buxton (1992) performed a similar study, in which they varied the amplitude of the movement, the target width and height, and the angle of approach. They found that using the minimum of the target width and height provided the best correlation with the experimental data, and should be used instead of the traditional Fitts' *ID*. Hoffmann and Sheikh (1994) independently found the same model. The *ID* for two-dimensional targets proposed by Mackenzie and Buxton (1992) is

$$MT_{min(W,H)} = a + b\log_2\left(\frac{A}{min(W,H)} + 1\right). \tag{2.4}$$

Although the $ID_{min(W,H)}$ provided a high correlation of over 95%, Accot and Zhai (2003) argue that such model isn't expressive enough, especially when the ratio of W and H is between 1 and 2. They state that a good bivariate pointing model should be scale independent, conform to limit tasks, provide a dominance effect to the smaller of W and H, have H and W in the same model, of similar nature and continuously represented.

Accot and Zhai (2003) tested the hypothesis that the target width and height interact in the *ID* with different weights. They used weighted ℓp-norms to define the "appropriate 'distance' in a two-dimensional space", and found that the best model that describes bivariate pointing is the Euclidean model with one free weight:

$$MT = a + b\log_2\left(\sqrt{\left(\frac{A}{W}\right)^2 + \eta\left(\frac{A}{H}\right)^2} + 1\right), \tag{2.5}$$

with η being a free weight empirically determined. This model provided a high correlation (95%) while keeping the number of parameters low, thus reducing the problem of overfitting the data.

The models described in this section are refinements of the original Fitts' law, all modeling *time-minimizing* tasks, or *rapid aimed movements* (MacKenzie, 2003). According to MacKenzie (2003), there is another type of tasks that can be modeled similarly to Fitts' law – *the space-minimizing tasks*, or tasks in which the cursor needs to follow a predefined path. In the following section, we examine Accot and Zhai (1997) model of a *space-minimizing* task.

Models of Human Motor Behavior for Space-Minimizing Tasks

While the purpose of Fitts' law is to model human motor behavior for binary tasks in which only the initial and final position of the target is taken in consideration, other models that extended the original law to accommodate for more complex tasks have been proposed. Accot and Zhai (1997) performed a study that extrapolated Fitts' law and modeled "tunnel traveling" tasks, in which not only the initial and final positions of the target mattered, but the trajectory of movement. This became known as *steering law*.

To model such task, first, the authors demonstrated that a task in which the cursor needs to cross two consecutive goals follows the same model as the original Fitts' tapping task. Then, they defined a recursion, by dividing the amplitudes of the movement and multiplying the number of divisions by the *ID* successively until the limit condition of *A* tending to zero, as shown in Figure 2.1. Such recursion leads to the following model of motor behavior for tunnel crossing:

$$MT = a + b\frac{A}{W}. \tag{2.6}$$

The authors performed other two experiments describing different trajectories, one for a narrowing tunnel and another for a spiral tunnel and specified the task generally, for any kind of trajectory:

$$MT = a + b\int_C \frac{ds}{W(s)}, \tag{2.7}$$

where C is the generic curved path. Figure 2.2 illustrates the integration.

In another application area, the *steeringlaw* has been demonstrated to apply not only for tasks that involve hand movements, but also for locomotion tasks, such as driving a car, in virtual environments (Zhai and Woltjer, 2003).

Probabilistic Approach to Model Pointing

Differently from approaches that aim at using information theory as an analogy to human pointing (MacKenzie, 1992), and from approaches whose models are empirically based (Accot and Zhai, 2003), Grossman and Balakrishnan (2005b) proposed a probabilistic approach to model two-dimensional pointing. The idea presented by the authors was to use a bivariate normal density function to predict the time to perform a pointing task based on a scatter of regions where single-phase ballistic pointing tasks were performed.

Having a spread of hit marks originally targeted at a point, but with only the ballistic phase of the movement being used to create the scatter, allowed for the determination of the probability with which the rapid, ballistic movement would fall inside the target, not requiring any refinement. The idea, then, was to map that probability to an index of difficulty, in a way that the less likely the cursor is to fall in the target only with a ballistic movement, the more refinements would be

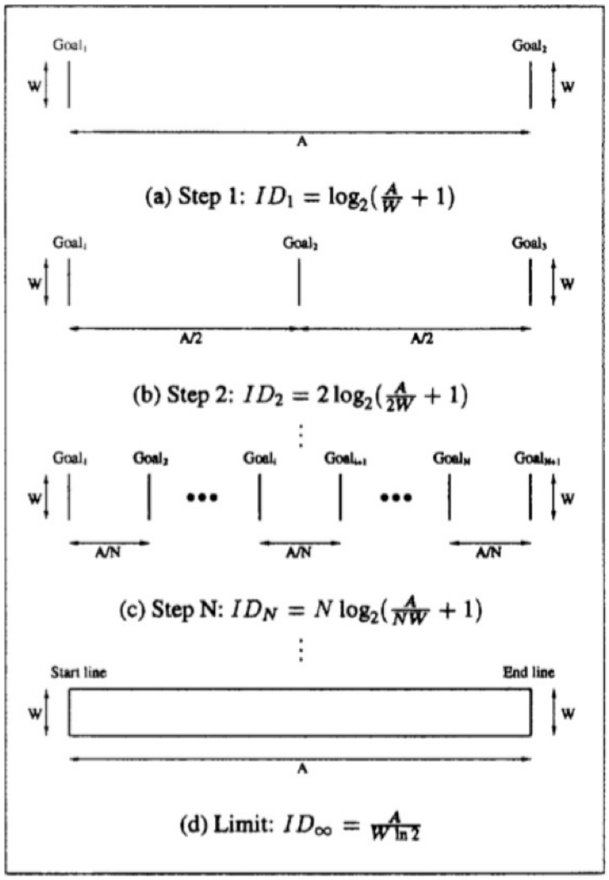

Figure 2.1 Defining a recursion with goal passing tasks. (Extracted from Accot and Zhai (1997).)

necessary to complete the task, thus increasing its difficulty, and, ultimately, the expected time to perform a pointing task.

2.1.2 Advanced Motor Behavior Models

There has been some effort in the research community to establish models for pointing behaviors other than those related to traditional mouse-based interaction techniques.

Kondraske (1994) proposed a model of direct target acquisition that used angular measures in the index of difficulty, motivated by the use of joint angles to determine end-effector position in biomechanical modeling. Although our model also uses angular measurements, we model a different task – distal pointing.

Murata and Iwase (2001) proposed a model for pointing at large information spaces. Their work, however, focused on up-close pointing in which the users touch the target with their fingers,

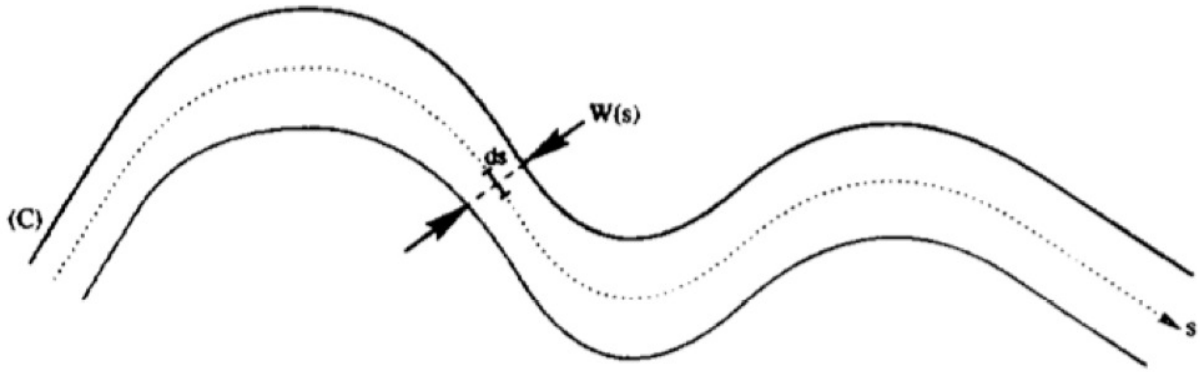

Figure 2.2 Integrating along a curve. (Extracted from Accot and Zhai (1997).)

not on distal pointing.

Grossman and Balakrishnan (2004) extended the Accot and Zhai (1997) bivariate pointing model for three-dimensional environments and proposed a model of trivariate pointing using volumetric displays. The modeled task is that of selecting 3D targets in a volumetric display as a factor of the width, height and depth of the target, as well as the amplitude of the movement and the angle of selection. As in our work, the input device they used was a 6-DOF tracker. However, their technique used a one-to-one mapping of the tracker position to the 3D display area. Our techniques, on the other hand, create a ray from the 6-DOF input device which intersect with a flat display to determine a cursor position.

The authors were able to successfully predict performance for pointing at trivariate targets in a volumetric display using the Euclidean model with the parameters being a function of the movement angle :

$$MT = a + b\log_2\left(\sqrt{f_W(\theta)\left(\frac{A}{W}\right)^2 + f_H(\theta)\left(\frac{A}{H}\right)^2 + f_D(\theta)\left(\frac{A}{D}\right)^2} + 1\right), \tag{2.8}$$

where W, H and D are the target width, height and depth, respectively; A is the movement amplitude and $f_W(\theta)$, $f_H(\theta)$, and $f_D(\theta)$ are empirically determined parameters based on the movement angle.

Perhaps the most similar work to our own is that by Stefels et al. (2007), who evaluated different pointing devices to be used by a surgeon in an operating room. They used Fitts' law to model performance with a regular mouse, a gyroscopic relative input device and the UI Wand, which is a input device design for distal pointing. In their experiment, performance with the UI Wand could be modeled accurately using the standard form of Fitts' law. Although they evaluated distal pointing, they did not vary the user distance to the display, which was fixed at $1.5m$, with the user remaining seated. Our work, however, focuses on conditions in which the user is standing and may interact from different distances to the display, and we eventually found that Fitts' law was not sufficient to model these more general types of distal pointing tasks.

2.1.3 Summary

Most laws presented in the previous sections describe the motor behavior model of moving the hand and lower arm on a supporting surface through an amplitude of no more than a desktop screen size. A few exceptions are the trivariate model (Grossman and Balakrishnan, 2004), where there is no supporting surface, thus leading to more muscle strain; the model evaluated by Stefels et al. (2007), in which the distance to the display was constant and rather close; and the steering law for locomotion tasks, in which the limiting aspects of the tasks are more cognitive than physical. Most models, then, express motor behaviors that are ideal for mouse, pen or touch-based interaction. Such models provide a good background for the research presented here, but they cannot be taken as applicable directly to distal pointing tasks, since they involve other types of motor behavior.

Although not directly applicable for distal pointing tasks, existing motor behavior models such as Fitts' law provide a good starting pointing for the development of a distal pointing model. We believe that we can use the knowledge from existing laws and extend from there to the development of the model we present in this research. In many ways, we want to parallel the work that has been done for modeling motor behavior for traditional interaction techniques in a distal pointing context.

2.2 Effects of Visual Acuity on Motor Precision

Strictly speaking, the motor precision is all that matters to determine the difficulty, and, by extension, the time it takes to perform a distal pointing task. However, even if the precision is provided to the user, this might not be enough if the user cannot perceive the target well enough. We, as humans, rely on what we see to drive our actions (Connolly and Goodale, 1999).

Several factors related to visual acuity can influence the performance of distal pointing tasks. From the literature, we didn't find, to the best of our abilities, any previous research that focuses on understanding the quality of the feedback control loop in relation to limited visual acuity. We can, however, speculate on possible effects by looking at other studies relating to the visual feedback loop and to visual acuity.

Vetter et al. (2010) demonstrated that performance of target detection tasks on a desktop computer is reduced in aging users with reduced visual acuity. This is evidence that there is a relationship between the time it takes to detect a target, or, in a pointing task, to detect that the cursor is on a target, and the ability the user has to perceive the target well.

Previous research has shown that closed loop movements – those in which there is continuous visual feedback of the body part in motion – are faster and more accurate than open loop movements – those in which there is no visual, but only proprioceptive feedback (Miall, 1996). The human visual feedback loop has a delay of 100-200ms (Carl and Gellman, 1987), which means that the motor system takes some time to adjust and correct for sudden changes of behavior. This is especially true in placement tasks on very tight spaces. Further, Ware and Balakrishnan (1994) have shown that lag causes small targets to be more difficult to select. Carlton (1981) suggests that reduced visual acuity causes delayed motor response.

All these pieces of results provide us hints that distal pointing tasks that have very small targets can compare to poor feedback loop, as, since the elements of the task cannot be clearly noticed, the feedback loop may not by completely closed. More is taken on the cognitive part from the user, which needs to squint and take a longer time to process whether the distal pointing task is ready to be completed.

It has been shown that people adapt to varied delays of feedback loop (Foulkes and Miall, 2000), and different offsets between the hand tracked position and the displayed position are easily corrected by humans (Vercher and Gauthier, 1992). These are indications that changing the mapping of the input device to the cursor movement, as we did in the design of ARM and DyCoDiR may be easily adjusted and accounted for by users.

2.3 Increasing Pointing Precision

Previous research focused on enhancing the precision of pointing, especially in large, high-resolution displays. This section covers the related literature on pointing enhancing techniques, both in terms of improving traditional techniques, as well as increasing the performance of distal-pointing-based tasks. Prior research has looked into using less orthodox input technologies for interaction with large displays. Examples are eye-gaze-based input (Zhai et al., 1999) and handheld projector input (Beardsley et al., 2005). In this section, however, we focus in methods to increase pointing precision based on the mouse and distal pointing techniques.

2.3.1 Mouse-Enhancing Techinques

As technologies become cheap and available, large display configurations are becoming popular by the day. Large displays can't benefit completely from traditional desktop computer input devices, interaction techniques and metaphors. Two recent literature surveys on large displays (Badillo et al., 2006; Ni et al., 2006) expressed these issues by listing interaction challenges that need to be overcome so that large display applications become usable and effective. Both works state that target acquisition and cursor tracking are among the most important issues that need to be solved. Target acquisition becomes a problem when the displays have a large physical size and resolution. Using the mouse to traverse long distances requires high cursor acceleration, which in turn leads to its loss and difficult tracking. It is a consensus in the research community (Badillo et al., 2006; Baudisch, 2006; Dasiyici, 2008; Ni et al., 2006) that novel interaction techniques for large displays are needed.

While the mouse as it works in desktop computers is not effective for large format displays, several mouse-based techniques that improve performance on large displays have been proposed.

Some techniques have been developed in order to overcome such problem. Baudisch et al. (2003b) proposed the high-density cursor (Figure 2.3), a technique in which cursor images were rendered to fill-in large gaps between cursor images when large cursor movements are made, thus keeping the cursor easy to follow. This technique was shown to be unobtrusive, with most subjects not noticing the presence of the cursor images, but yet provided significant improvement on the time to complete tasks.

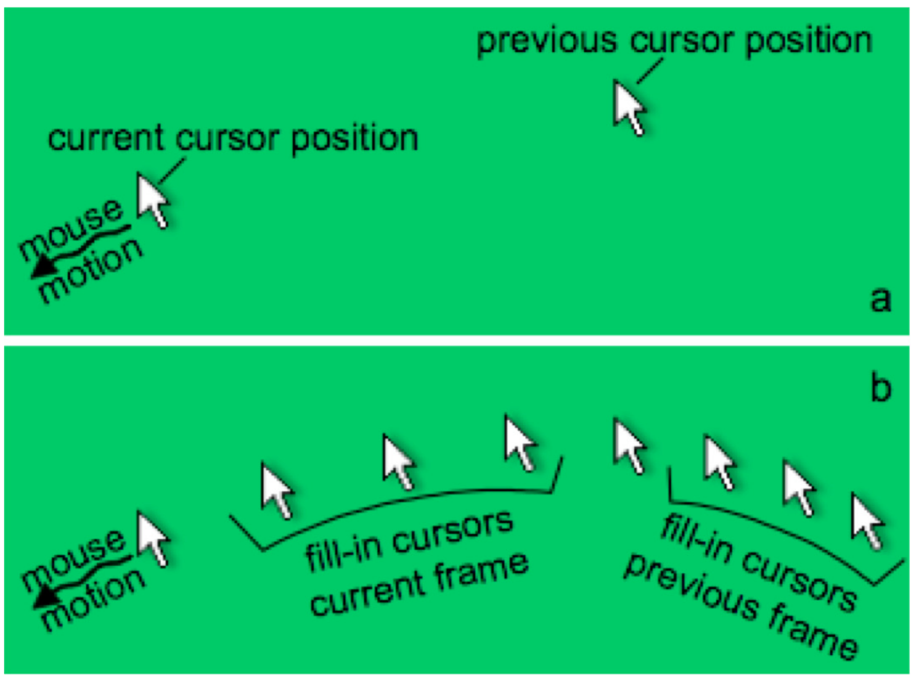

Figure 2.3 High-density cursor. (Extracted from (Baudisch et al., 2003b).)

Dasiyici (2008) tackled the cursor tracking and remote target acquisition by developing the Multi-Scale cursor (Figure 2.4) which provided a dynamic control-display (C/D) ratio of the cursor size and speed. With this technique, subjects were able to follow a very large cursor when it was moving fast and, for precise and refine movements, the cursor returned to its original size.

Figure 2.4 The MultiScale cursor. (Extracted from (Dasiyici, 2008).)

Another way of improving target acquisition in large displays is to snap the cursor to possible targets in the application. Object pointing (Guiard et al., 2004) is one example in which the cursor travel through objects on the application, creating "void-phobic" interaction, or interaction in which the cursor is never in an empty space when a selection is to be made. Another technique that snaps the cursor to targets is called bubble cursor (Grossman and Balakrishnan, 2005a). In this case, an area cursor changes its size dynamically to encompass the closest target, or it changes its shape so that the closest target is completely inside the area cursor.

These examples of mouse enhanced interaction techniques give evidence that it is possible to adapt the mouse for large high-resolution displays. However, mouse-based interfaces do not afford the user to stand and walk around area in front of the display, which has been found to be a common behavior among users who improved the performance of search tasks compared to virtual navigation (Ball et al., 2007). In such cases, different techniques should be sought.

One option for interaction with large displays that afford the users to walk across the area in front of the display surface are up close interaction techniques, such as touch screens or pen-based interfaces. A technique of this type is HybridPoint (Forlines et al., 2006), which is a pen-based interaction technique that allows remote target acquisition by switching between absolute and relative pointing modes. Another example is Drag-and-Pop and Drag-and-Pick (Baudisch et al., 2003a), in which proxies of possible targets are temporarily displayed near the interaction point of the application. A similar technique that addresses some usability issues of Drag-and-Pop is called Vaccum (Bezerianos and Balakrishnan, 2005).

2.3.2 Distal Pointing

Pointing is one of the fundamental classes of 3D interaction techniques (Bowman et al., 2004). Distal pointing is the term we use to describe a style of interaction in which the intersection of an invisible ray extending from the user's hand, or hand-operated input device, with the display surface determines the interaction point, usually the cursor, of the application.

Distal pointing with large displays provides users the freedom to move around in front of the display and supports moving the cursor rapidly to any point on the display. It is also very simple and direct. Unfortunately, its lack of precision makes it an impractical technique to use with large, high-resolution displays (Vogel and Balakrishnan, 2005). This is evident when trying to select and manipulate small targets. On the other hand, an advantage of techniques based on distal pointing is that they do not suffer too much with cursor tracking issues, because the cursor position is determined by the direction where the user points with the input device.

Previous research used the angular motion of input devices to determine the linear cursor position on a display. Gibbs (1962) compared several levels of angular gain (both positional and velocity) in the performance of target selection on a display, while Buck (1980) studied how angular gains on different joysticks interacted with different gains related to target widths on the display.

Several interaction techniques based on distal pointing have been proposed. Some of these techniques have used laser pointers as an input device (Dan R. Olsen and Nielsen, 2001; Matveyev and Göbel, 2003; Myers et al., 2002a). Vogel and Balakrishnan (2005) used distant freehand pointing to interact with large displays. Jiang et al. (2006) created Direct Pointer, an interaction technique that uses a camera on a handheld device as input. Using computer vision to track the input device called VisionWand, Cao and Balakrishnan (2003) developed a set of interaction techniques based on distal pointing and wand gestures. We also note the similarity of distal pointing techniques to ray-casting techniques in 3D virtual environments (Bowman et al., 2004).

Techniques for up close interaction with large displays afford the users to remain always right in front of the display, while interacting with it. In many cases, the user needs to be at a distance from the display to perform an interaction, for example, with the overview of a dataset.

Using freehand gestures for the interaction with large displays is an alternative that allows the users to stand at any distance from the display. Vogel and Balakrishnan (2005) presented a set of techniques based on freehand gestures. They tackled the precision pointing problem by switching between absolute and relative modes. Attempting to improve interaction from a distance at large displays, Malik et al. (2005) created a multi-finger gestural interface where the user sits at a table in front, but far from the display and perform operations on a tablet. While this work does not involve distal pointing, it shares some of the same goals of allowing precise distal interaction.

Ray-Casting

Ray-casting (Mine, 1995) is the most basic distal pointing interaction technique. Traditionally, ray-casting has been used as a basic interaction technique for immersive VEs (Bowman et al., 2004). With ray-casting, the user points with a virtual ray extending from the hand or input device to specify an object in the scene. Although it is very simple, this technique in its pure form suffers from a number of issues, mostly because of natural hand tremor, which varies between 8-12 Hz (Riviere et al., 1997) and tracker jitter, which makes it difficult for the user to control the origin and orientation of the ray. This is a bigger issue with ray-casting than with other techniques because the small hand movements are amplified at the end of long rays, causing ray-casting to be less precise as target objects get farther from the user. These issues make ray-casting difficult to use when the objects have a small visual size (Steed and Parker, 2004), as selecting such objects by pointing requires high levels of precision.

In VEs, several improvements to ray-casting have been proposed. Bowman and Hodges (1997) proposed the HOMER (Hand-centered Object Manipulation Extending Ray-casting) technique, that facilitates the manipulation of distant objects by attaching the user's virtual hand to the remote object. After manipulation is completed, the hand returns to its original position. Although it improves the precision of object manipulation, HOMER does not solve the precise selection problems. To address that, other enhancements were proposed, such as snapping to objects (Wingrave et al., 2002), in which the ray doesn't move continuously, but "jumps" from object to object in the VE.

Unfortunately, some of the VEs ray-casting-based techniques are not suited for large format display interaction. Since there is no virtual hand that can be placed anywhere in the environment, the enhancements need to be in the display surface. Another issue with large display interaction is that the “ray of light” present in VEs cannot be displayed in the real environment of large format displays, which is replaced by an invisible ray.

In order to address these issues, a number of improvements have been proposed. Even though such techniques improve selection performance in general, they can have a negative effect in very cluttered environments. Cone-casting (Liang and Green, 1993), for example, extends ray-casting by adding a cone-shaped volume to the ray to make it easier to select objects that are distant. In cluttered environments many objects will fall inside the cone, so that the user still has to point precisely to select the desired object. IntenSelect is a snapping technique presented by Haan *et al.* (Haan et al., 2005) which uses a selection volume to calculate and accumulate scores over time for each object. This way, it can estimate which object the user wants. The bubble-cursor (Grossman and Balakrishnan, 2005a) is a 2D technique that dynamically resizes a circular cursor so that it only contains one object. A 3D extension of the bubble-cursor, which uses a sphere instead of a circle, was presented by Vanacken et al. (2007). Both of these techniques may actually perform worse in cluttered environments, since even small movements will cause the ray or the cursor to constantly snap or resize to select new targets.

Laser Pointers

Laser pointers have been used as input devices for large display interaction (Dan R. Olsen and Nielsen, 2001). Myers et al. (2002b) performed a comparative study with different laser pointers and with a laser pointer and three other input devices. They showed that laser pointer interaction has poor performance, and time to select targets was high. Most studies with laser pointers used low-resolution projection screens, and even in such conditions pointing was problematic due to hand jitter, camera tracking and target size. For interaction on high-resolution displays, Davis and Chen (2002) presented a technique for multi user input using laser pointers for tiled displays.

Although laser pointers can be used to interact with large displays, as the resolution increases, and the target sizes get small, precision problems start to occur.

Adaptive CD Ratio

A number of approaches have been proposed to enhance distal pointing precision by dynamically changing the CD ratio based on the input device velocity. PRISM (Frees et al., 2007) focus on scaling down movements of a ray cast onto an immersive virtual environment as the user tries to be precise. PRISM is intended to perform small movements at large arm rotations, providing a lot of precision, but at a low interaction space. Wilkes and Bowman (2008) have extended the HOMER technique (Bowman et al., 1999) to provide high-precision selection and manipulation of objects using the original technique.

Specifically for distal pointing at large high-resolution displays, König et al. (2009) introduces the adaptive pointing technique, which uses several criteria to determine the CD ratio to be applied to the pointing position. While this technique successfully improves precision, reducing pointing time and error, it does not take into account the relative distance of the user to the display, and its implementation is rather complex, as compared to DyCoDiR (section 4.6).

2.3.3 Distal Pointing by Progressive Refinement

Bowman et al. (2004) divided selection techniques into four main categories: selection by pointing (e.g., ray-casting (Mine, 1995)), selection by touching (e.g., virtual-hand (Mine et al., 1997)), selection by occlusion (e.g., image-plane techniques (Pierce et al., 1997)) and indirect selection (e.g., selection by attributes (Bowman et al., 2004)). Most of these techniques can be classified as immediate selection, since they only require a single high-precision selection without refinement.

Other techniques improve ray-casting accuracy by changing the control-display ratio, either automatically (e.g., PRISM (Frees et al., 2007), when moving slowly) or manually (e.g., ARM (Kopper et al., 2010), by pressing a button), or by providing the ability to zoom (e.g., zoom-and-pick (Forlines et al., 2005)). While these techniques can achieve very high levels of precision, all of them have limitations. PRISM and ARM cause a significant mismatch of the physical pointing direction to the perceived pointing position, and the mapping is nonlinear. Zoom-based techniques suffer from potential loss of detail. Finally, these techniques require the user to interact very carefully and with full attention. Our proposed approach of selection by progressive refinement aims to allow "lazy" interaction with high accuracy.

There are existing progressive refinement techniques in the literature. For example, the shadow cone-casting technique (Steed and Parker, 2005) uses continuous movement with cone-casting to disambiguate selection. Steed (2006) presented a general model for selection using 3D gestures and proposed a range of techniques that can use the same concepts. The depth-ray technique (Grossman and Balakrishnan, 2006), which adds depth control to the classic ray-casting technique to select occluded objects, requires two actions to specify the target. PORT (Lucas, 2005) allows the selection of multiple objects and uses a series of movement and resizing actions to define the set of targets.

The flower ray technique (Grossman and Balakrishnan, 2006), in which occluding targets are concurrently selected by ray-casting and disambiguated in a second phase by a marking menu is perhaps the closest existing technique to our approach. Although this technique is suited for highly cluttered environments, it requires high precision for the ray selection and does not scale well, since the marking menu object specification is done in a single phase.

To the best of our knowledge, there has been no prior generalization of the progressive refinement concept, and no comparison of progressive refinement techniques to immediate techniques.

2.3.4 Summary

Table 2.1 shows a summary of existing types of pointing enhancements that have been proposed. Most of the previous work, however, does not focus on interaction techniques that afford the user to walk freely in front of the display while interacting with the application, which is an important affordance (Ball et al., 2007). Further, there is little theoretical background in the previous works. We believe that we can investigate distal pointing in order to provide theoretically valid guidelines for the design of effective and usable user interfaces.

Table 2.1 Types of pointing precision enhancements and example techniques.

Type of Input	Type of Enhancement	Interaction Techniques
Mouse	Gap filling	High-Density Cursor (Baudisch et al., 2003b)
	Dynamic C/D ratio	Multiscale Cursor (Dasiyici, 2008)
	Snapping	Object Pointing (Guiard et al., 2004)
		Bubble Cursor (Grossman et al., 2005a)
Touch	Absolute/relative mode switching	HybridPoint (Forlines et al., 2006)
	Target proxies	Drag-and-Pop (Baudisch et al., 2003a)
		Drag-and-Pick (Baudisch et al., 2003a)
		Vaccum (Bezerianos et al., 2005)
Distal Pointing	Absolute/relative mode switching	Hybrid RayToRelative (Vogel et al., 2005)
		ARM (Kopper et al., 2010)
	Snapping	IntenSelect (Haan et al., 2005)
		3D Bubble Cursor (Vanacken et al., 2007)
	Adaptive C/D Ratio	PRISM (Frees et al., 2007)
		Scaled HOMER (Wilkes and Bowman, 2008)
		DyCoDiR (section 4.6)
	Progressive refinement	Shadow Cone-Casting (Steed et al., 2005)
		Depth Ray (Grossman et al. 2006)
		Flower Ray (Grossman et al. 2006)
		PORT (Lucas, 2005)
		SQUAD (section 4.7.2)

2.4 Physical Navigation

One of the goals of this study is to understand users choice to physically navigate while interacting with distal pointing. The motivation for this research question comes from previous studies, which have shown physical navigation to improve the performance of tasks on large high-resolution displays. In contrast, we know from experience that many users do not want to physically navigate while performing a distal pointing task.

This section covers some prior research on the effects of physical navigation on interactive tasks.

A solid body of prior work in immersive virtual environments (VEs) show that physical navigation can be beneficial in the completion of tasks in virtual reality (VR) applications. For example, Pausch et al. (1997) demonstrated that visual search and wayfinding improved when head tracking was present. Further, research has shown that spatial orientation is kept better by users if they can physically, as opposed to virtually, navigate in immersive VEs (e.g., Bakker et al., 1998; Chance et al., 1998).

The tradeoff between physical and virtual navigation has been the topic of previous research. Ball and North (2005) studied visualization and navigation tasks on different sized multi-monitor systems and concluded that participants favored physical navigation over virtual panning and zooming. These results were further verified in a larger study involving a bigger display with higher resolution (Ball et al., 2007). The tasks in this study involved participants walking to a target, finding targets, pattern finding, and insight gathering on a map that contained abstract information overlaid. The results show that not only participants perform better with physical navigation, but physical navigation was preferred over virtual panning and zooming.

Yost et al. (2007) has also shown that visualizations can improve from much details that go beyond visual acuity. In those cases, participants need to physically navigate in order to gather detail from the data, but they were successful in doing so.

It is important to note that, for the visual search tasks involved in the afore mentioned studies, physical navigation was compared to virtual navigation. For interactive tasks that use distal pointing, virtual navigation is not possible, but we analyze the tradeoff between performing a task aided by high-precision enhancements while stationary and walking closer to the targets to get precision.

These results from previous studies contrast with our experience that participants tend to not physically navigate a lot, and especially, not walk towards targets (Kopper et al., 2008). We have partially verified this initial idea in the study presented in chapter 5. Some participants indeed never chose to physically navigate. We believe that this seemingly contradiction between our results from the results from prior work is so because the tasks that were performed by all studies described above involved solely visual search and data navigation. Interaction with application elements was never required, and we hypothesize that the cause for some participants to not walk is because they are performing interactions.

Chapter 3

Distal Pointing Model

With a predictive, analytical model, we can address the three distal pointing challenges we identified in this research. The understanding of distal pointing interaction is gathered from the model itself and the knowledge it captures. The design of distal-pointing-based interaction techniques is facilitated by the provision of guidelines derived from characteristics of the model. Finally, the evaluation of distal pointing interaction can be done through the analytic assessment of different techniques using the model.

This chapter presents the process to propose and evaluate a human motor behavior model for distal pointing tasks.

3.1 Overview

It has been argued that human-computer interaction (HCI) problems and challenges should be reduced to well-formed and rigid models in order to have a hard science of interface design (Newell and Card, 1985). In reality, very few models are robust enough that they can be used in such a way. Models of human motor behavior, however, have proven to be one exception.

Most motor human behavior models used in HCI, such as Fitts' law (Fitts, 1954) and the Law of Steering (Accot and Zhai, 1997), provide guidelines for the design of desktop interfaces, and predict time and difficulty of mouse-related tasks. When it comes to novel interfaces that require other types of input, there is little development of human motor behavior models – a model for direct interaction with volumetric displays (Grossman and Balakrishnan, 2004) and a model of raycasting in a virtual environment (Wingrave and Bowman, 2005) are notable exceptions.

Recently, many interactive systems allow the user to interact while standing up and walking around, by pointing his hand or input device directly at the display from a distance. We call this type of interaction *distal pointing* (Bangerter and Oppenheimer, 2006; Pfeiffer et al., 2008; Po, 2005), and distinguish it from other methods of indicating a location on a display, such as direct touching using touch screens or indirect cursor control using a mouse. With a typical distal pointing technique (i.e., raycasting), the cursor position is determined by the intersection of a ray emanating from the input device with the display.

Several existing systems and research projects use some form of distal pointing to control the cursor. Examples include home entertainment system control (Patel and Abowd, 2003) and laser pointer interaction (Dan R. Olsen and Nielsen, 2001). In fact, distal pointing interaction has made its way into consumer products with the Nintendo Wii (2011), which uses distal pointing to control an on-screen cursor. Interaction with large high-resolution displays can also benefit from distal pointing (Forlines et al., 2005; Malik et al., 2005; Matveyev and Göbel, 2003; Vogel and Balakrishnan, 2005).

One important aspect of distal pointing is that it has inherent precision issues, due to hand jitter, lack of a supporting surface, and the difficulty of acquiring small targets from a remote position. To accommodate for such lack of precision, distal pointing techniques often include enhancements to allow higher precision. Enhancements may include input filtering (Vogel and Balakrishnan, 2005), area cursors rather than point cursors (Tse et al., 2007), the ability to zoom in on a particular area of the display (Forlines et al., 2005), or the ability to control cursor speed (Vogel and Balakrishnan, 2005).

Because of the complexity added by these enhancements, a user may develop several different strategies to achieve the same goal. The user must decide when and where to walk, when and where to zoom in, or when and how to change the cursor speed. This dependence on strategy makes it difficult to empirically compare distal pointing techniques with traditional techniques or with each other. Performance will depend on whether the user chose the best strategy, and individual differences may be large.

A predictive, analytical model for distal pointing could be used to evaluate distal pointing interaction techniques, which would allow us to overcome the limitations of empirical studies. With such a model, we would be able to predict user performance in the presence of various strategies. An analytical model could also offer guidelines for the design of techniques that afford the use of good strategies, because we would know the most critical factors that affect performance on distal pointing tasks, and would design techniques accordingly.

In this chapter, we present an experiment that led to a model of human behavior, inspired by Fitts' law (Fitts, 1954), for distal pointing tasks. While similar to Fitts' law, the distinctive idea that differentiates our model from most Fitts' law approaches (Accot and Zhai, 1997, 2003; Grossman and Balakrishnan, 2004; MacKenzie, 1992) is that angular target sizes and movement amplitudes, rather than linear measures, are used as input parameters for the model. Although we are not the first to use angular measures in pointing tasks (Buck, 1980; Gibbs, 1962; Kondraske, 1994), no previous research has looked into using angular measurements as factors for a general model of distal pointing. Figure 3.1 illustrates how arm and wrist rotations are important in a distal pointing task, which we have verified by observing users pointing directly at targets from a distance. In addition, our model shows that angular target size is more significant than angular movement amplitude for distal pointing tasks, and that the difficulty of distal pointing tasks increases non-linearly.

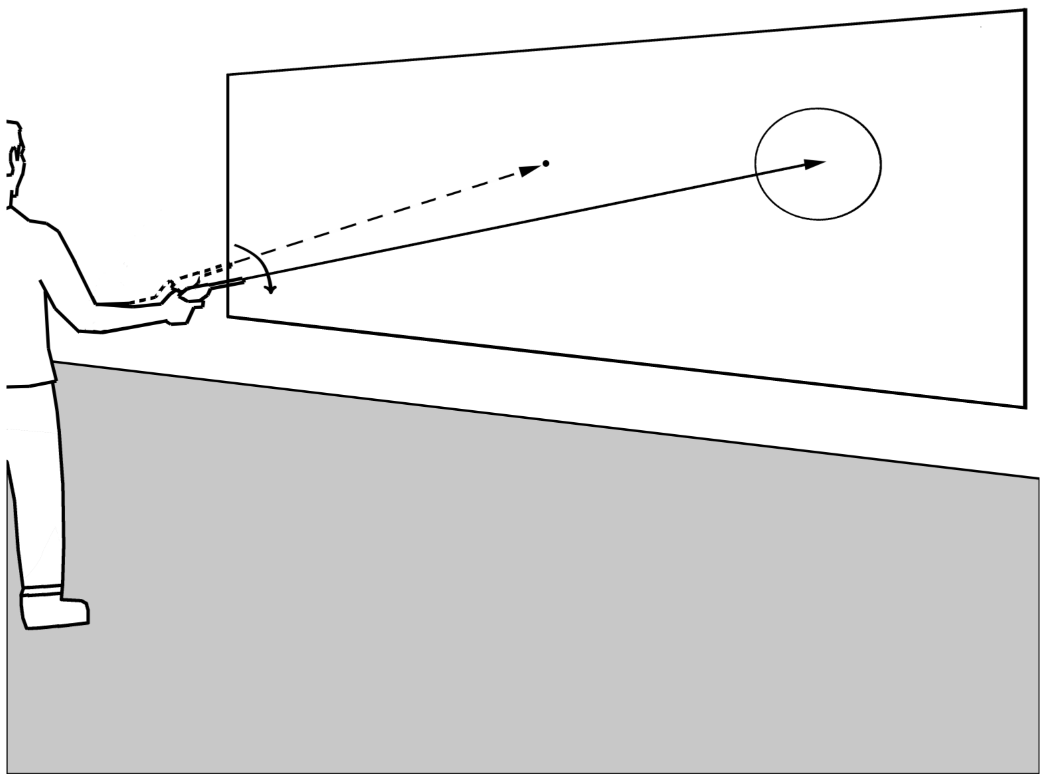

Figure 3.1 Distal pointing movements require arm and wrist rotations, which are best expressed as angles in the model.

3.2 Modeling Distal Pointing

The best-known human motor behavior model for pointing tasks is Fitts' law (Fitts, 1954). Its simplicity and robustness is perhaps the reason for its heavy use in the design of graphical user interfaces (GUIs). It can be used to show that the time to acquire a target is dependent on its size and on the amplitude of movement. Generally, it can be expressed as

$$MT = a + b \cdot ID, \tag{3.1}$$

where MT is the movement time to complete the task, a and b are empirically determined constants, and ID is the index of difficulty of the task, which is a function of amplitude and target size.

We hypothesize that the difficulty of tasks based on distal pointing can be modeled linearly, with the slope and intercept of the regression line being determined empirically. We further hypothesize that this model contains an ID that expresses the relationship among the parameters of the task, which, for distal pointing, include target size, path length and user distance to the display. Finally, we hypothesize that, as in Fitts' law, ID a logarithmic factor, following the transmission of information theory from the Shannon's Theorem (Shannon and Weaver, 1949).

In this section, we discuss candidate formulations of *ID*, and discuss potential benefits and disadvantages for each of the proposed formulations.

3.2.1 Original Fitts' *ID*

Due to the uniqueness of distal pointing, as compared to other uses of Fitts' law in HCI, we believe that the traditional *ID* is not adequate to model distal pointing tasks. In its most widely accepted form, the original Fitts' *ID* (Accot and Zhai, 1999; MacKenzie, 1992; Mackenzie and Buxton, 1992) is expressed as

$$ID = \log_2\left(\frac{A}{W} + 1\right), \tag{3.2}$$

where A is the movement amplitude and W is the width of the target. We believe that the user distance to the display is an important factor that needs to be absorbed in the model.

Stefels et al. (2007) found a good fit of their distal pointing data to the original Fitts' law model from a fixed distance of $1.5m$ to the display. Knowing that the distance to the display is an important element in task performance, we propose to determine whether the original Fitts' model can be used for different distances of the user to the display and affect only the slope and intercept of the Fitts' regression line (coefficients a and b from Equation 3.1) without loss in the accuracy of the model.

However good the fit from the original Fitts' *ID* is for a given distance of the user to the display, we still believe that this factor should be incorporated in the index of difficulty of the task. We wish to use our model of distal pointing to analytically evaluate distal pointing interaction techniques and strategies. Indeed, an empirical study of such techniques (Kopper et al., 2008) showed that users physically navigate relative to the display when performing realistic tasks. It is important, then, that we be able to include the user distance to the display as a parameter when doing performance predictions.

3.2.2 Integrating *D* Into Fitts' *ID*

For distal pointing tasks, movement is not constrained to a fixed plane, such as a table, and the physical position of the user plays an important role in the difficulty of a task. To incorporate this into the Fitts' *ID*, we can use the raw parameters of the task, namely the amplitude of movement (A), the width of the target (W) and the distance of the user to the display surface (D), leading to an index of difficulty expressed as

$$ID_{RAW} = \log_2\left(\frac{A \cdot D}{W^2} + 1\right). \tag{3.3}$$

The reason for the square of the target width in this ID is that we hypothesize that the decrease in performance as W gets smaller is approximately proportional to the decrease in performance as A gets larger, or to the decrease in performance as D gets larger. When both A and D are placed in the numerator, therefore, W^2 is required in the denominator.

ID_{RAW} is more expressive than ID (Equation 3.2) since it accounts for the user distance to the display. It is also straightforward, since it requires only the linear parameters of the task to predict distal pointing performance. On the other hand, ID_{RAW} may not be very generalizable. Realistically, users may stand in any position in front of the display and point in any direction to perform a task. As we can see in Figure 3.2, it becomes unclear which value should be used for D in situations in which distance to the initial pointing location is different from the distance to the final pointing location.

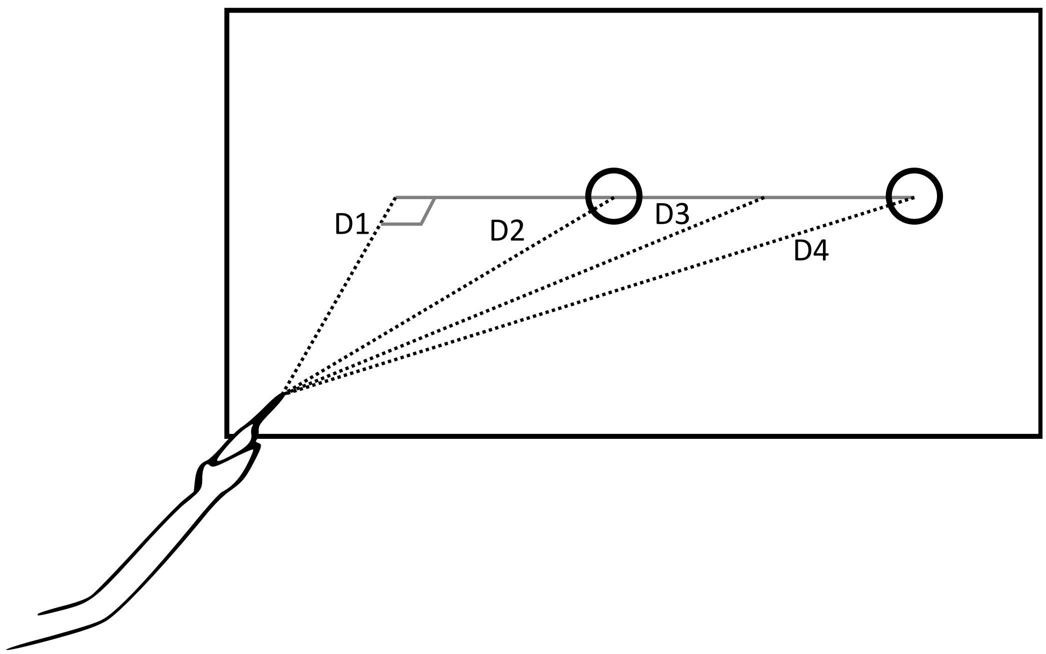

Figure 3.2 Ambiguity in the user distance to the display. It is not clear which distance should be used in ID_{RAW}.

We could resolve this ambiguity by using angular measurements of target size and movement amplitude in the index of difficulty for distal pointing tasks. This leads to the next proposed model.

3.2.3 Using Angular Measurements for *ID*

In Fitts' law terms, the amplitude of user movement in a distal pointing task decreases as the user moves away from the display because the arm or wrist rotation is smaller. For the same reason, the target width, in terms of required user movement, is smaller the farther the user is from the display.

In other words, the angular amplitude of the movement (α) and the angular size of the target (ω) may be more appropriate parameters for the distal pointing model than the linear amplitude of the cursor movement or the linear width of the target object on the display. In our experimental task (see section 3.3), in which the user is always in the center of movement, α is formulated by

$$\alpha = 2\arctan\left(\frac{0.5A}{D}\right), \tag{3.4}$$

and ω is defined as

$$\omega = \arctan\left(\frac{0.5\,(A+W)}{D}\right) - \arctan\left(\frac{0.5\,(A-W)}{D}\right), \tag{3.5}$$

where A is the amplitude of movement, W is the width of the target, and D is the perpendicular distance from the user to the display surface. Figure 3.3 illustrates the analogy between the linear and the angular values.

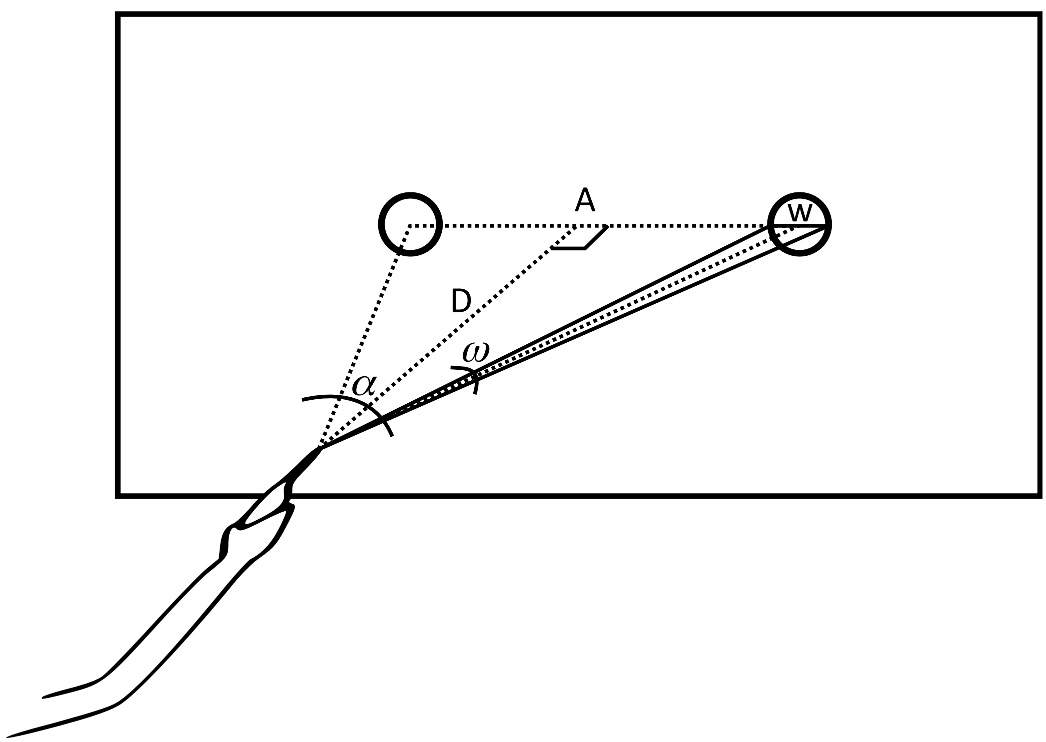

Figure 3.3 Relationship between α and A and between ω and W.

We propose to incorporate angular measurements to the distal pointing model as a direct analogy to the original Fitts' law. This is represented as

$$ID_{ANGULAR} = \log_2\left(\frac{\alpha}{\omega^k} + 1\right). \tag{3.6}$$

with α being the angular amplitude of movement, ω the angular width of the target and k is a constant power factor determining the relative weights of ω and α.

The reason for the constant k as a power of ω is that there is not always a linear relationship between α and ω. When using distal pointing, the user's movement consists of at least two phases: ballistic and correction (Grossman and Balakrishnan, 2005b; Liu et al., 2009; Woodworth, 1899). In the ballistic phase, the pointer moves very rapidly from one point to another using wrist rotation. These rapid movements place the pointer in the target region, and then, in the correction phase, fine-grained adjustments to acquire the target occur. In our experience with distal pointing techniques, this second phase of movement takes the majority of the time needed to complete the task (suggesting a value of k greater than 1). Further, natural hand tremor and the movement the cursor makes when a button is pressed to acquire a target – the so-called *Heisenberg effect* (Bowman et al., 2002) – add to the precision issues that make a remote target difficult and slow to acquire. We believe that different experimental settings, such as the type of input device and tracking jitter will affect the value of k.

3.3 Experiment

We designed and conducted an empirical study, which was similar to classical Fitts' law studies. The task we were modeling was a reciprocal selection task: the user points to and acquires, by clicking, two consecutive graphical objects.

3.3.1 Apparatus

We used a flat tiled display consisting of 50 NEC MultiSync LCD2080UXi monitors in a 10×5 configuration (Figure 3.4). Each monitor's resolution was 1600×1200 pixels, resulting in a total resolution of 16000×6000 pixels. Although the display we used was a grid of monitors that contain bezels, previous research has shown that they do not affect performance of target selection tasks (Bi et al., 2010). Further, we took measures to minimize the problem of fluidity of perception and interaction by positioning targets such that they did not cross any bezels.

A wireless Iogear Phaser Mouse was used as the input device (Figure 3.5). A trigger-like button was used to send mouse click events, so that movement of the device due to button clicks was minimized. To enable 6-DOF input, we attached reflective markers to the wireless mouse, which were tracked by a VICON MX system with eight cameras. The cameras were distributed evenly in a rough circle of 4.5m diameter, in such a way that tracking accuracy was consistent throughout the whole region in front of the display. The tracking system accuracy was on the order of one millimeter, with the input device's angular granularity being considerably smaller than human natural hand jitter. A dynamic recursive low pass filter (Vogel and Balakrishnan, 2005) was applied to the raw position data read from the tracker, which visibly reduced jitter without compromising the response time. We also determined the 3D position of the display surface,

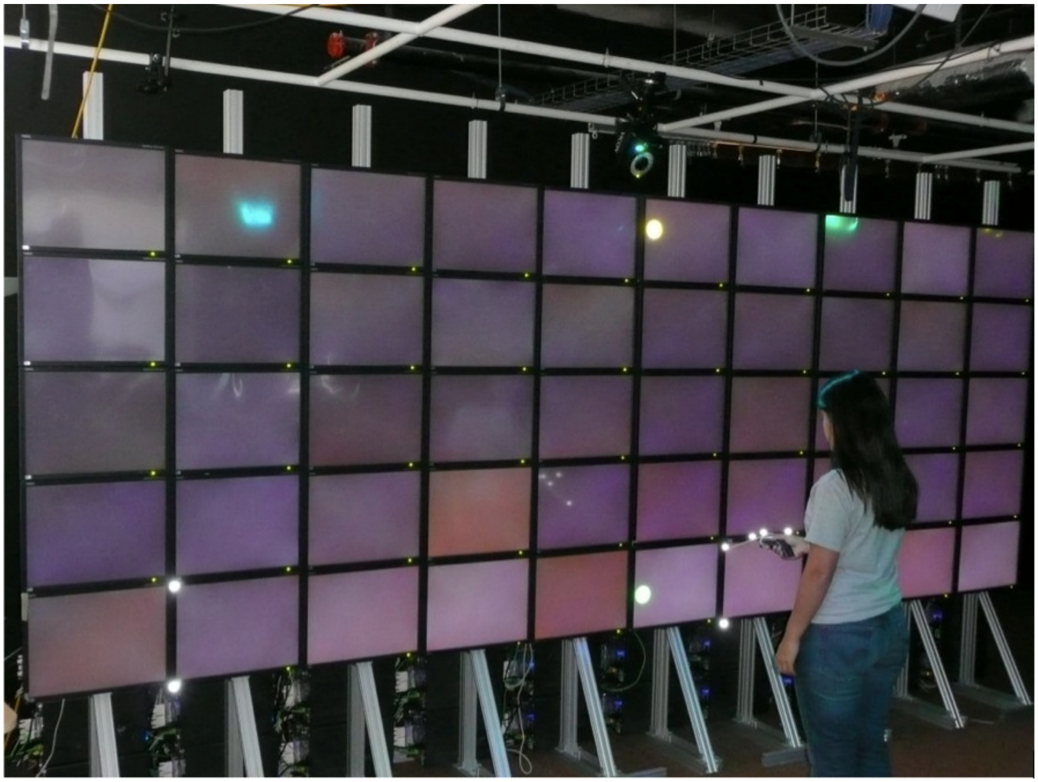

Figure 3.4 Large display used in our experiment.

enabling us to intersect a virtual ray emanating from the front of the input device with the plane of the display, and thus determining the location of the cursor. We implemented a simple application using the OpenScenegraph (www.openscenegraph.org) library. The user was presented with a black screen containing colored circular objects. Two input clients were created to handle button press events and tracker data from the input device to the application, and the data was transmitted over an Ethernet connection through sockets.

3.3.2 Participants

Twenty-one participants (three females) were recruited from the campus community to participate in the experiment. Two were left-handed, and all subjects were instructed to hold the input device in their dominant hand. Their ages ranged from 20 to 40 years old.

3.3.3 Procedure

Subjects first filled in a background questionnaire and signed an informed consent form, after which they were read general instructions about the experiment.

Figure 3.5 Input device used in the experiment.

The participants were instructed to complete tasks that consisted in selecting, via a click on the input device, two consecutive circular targets. Only two targets were shown on the screen at a time, the first was green and the second was yellow. Once the first target was selected, the targets switched colors to indicate that the subject should select the second target.

In order to get used to the interface, each subject practiced on randomly selected task combinations for five minutes. Following the practice section, the users went on to the experimental session, which is detailed below.

3.3.4 Design

We used a factorial within-subjects design with repeated measures. There were four independent variables: the movement amplitude *A* (1000*px* or 0.2758*m*, 3000*px* or 0.8274*m*, 5000*px* or 1.379*m*), the target width *W* (64*px* or 17.65*mm*, 128*px* or 35.3*mm*, 256*px* or 70.6*mm*, the distance of the user to the display *D* (1, 2, 3 meters), and the direction of movement *dir* (down, up, toward dominant side, toward non-dominant side). *A* and *W* are the classical Fitts' law parameters. We included *D* based on our experience that distal pointing increases in difficulty farther from the display. We included *dir* to verify the hypothesis that different movement directions do not significantly affect task completion time.

The user always stood on a line extending at a right angle from the center of the display, and the center of the movement was always the center point of the screen. Thus, for horizontal tasks, the movement always went across the user's body, and for vertical tasks the movement was always directly in front of the user.

The dependent variable was task completion time (denoted in the results as *MT*). Each trial began when the user acquired the first target, and finished when the user clicked on the second target; the experimental software measured the trial time in microseconds. Clicks outside of the target were considered errors and were recorded, but did not invalidate the trial.

During the experiment, subjects completed three sets of trials, blocked on the distance *D*. We counterbalanced the order of presentation of the three distances. We made this choice so that movement fatigue and delays for adjusting to the correct position would not occur. Within each block, we randomized the order of presentation of the 36 combinations of *A*, *W* and *dir*. Each subject performed five consecutive trials for each of the 108 combinations, totaling 540 data samples per subject.

3.4 Results

Trials in which a click aimed at the second target was more than half the amplitude of movement away from the center of the target were considered mistrials and removed from the data analysis. This occurred occasionally due to the sensitivity of the trigger button in the input device, and most of the misclicks occurred inside the first target. 2.5% of the trials contained misclicks and were removed.

We considered trials that were at least two standard deviations away from the mean for each condition as outliers and removed them from the data analysis (Grossman and Balakrishnan, 2005b). A total of 4.9% of the trials were removed as outliers. For the repeated measures analysis of variance, data was averaged across conditions per subject, so that one point per condition per subject was included in the analysis. Of the 2268 data points, 16 were missing after the removal of outliers, due to all five trials of that participant in that condition being outliers. To perform the analysis including those subjects with a few missing values, we then replaced the missing values with the mean of that condition plus two standard deviations (Tabachnick and Fidell, 2006).

Due to the randomized order of trials, we did not test for learning or order effects. We did not observe any significant increases in performance during the experiment. Considering the amount of training combined with the simplicity of the task, we believe that participants achieved asymptotic performance before they started the trials.

3.4.1 Error Analysis

We counted any incorrect attempt to acquire the target as an error, but we considered the trial valid even when errors occurred. We decided not to remove trials that contained errors due to the fact that for the most difficult tasks, where D and A are high and W is low, the error rate was quite large. We can explain this due to the fact that, even with the low pass filter, hand tremor was a significant problem, especially with high values of D. The greater the length of the ray, the more sensitive will be the cursor movements. The low pass filter could have been enhanced to remove more jitter, but increasing its effect too much would lead to intentional movements of the input device not being captured by the application. We tested several cut-off frequencies for the low pass filter and found a good compromise, which reduced jitter notably without filtering out any intentional movement of the user.

Main Effects

We found significant main effects on mean error rate for A ($F_{2,40} = 30.79, p < 0.0001$), D ($F_{2,40} = 75.40, p < 0.0001$), W($F_{2,40} = 89.89, p < 0.0001$) and dir ($F_{3,60} = 6.525, p < 0.05$). Pairwise comparisons show that, at $alpha = 5\%$ all the values of D were significantly different, with more errors at the further distances to the display. Similar trends were found for A, with more errors committed at the highest amplitudes and for W, with the lowest width having the greater number of errors. Figure 3.6 shows the least square means plot for the direction of movement.

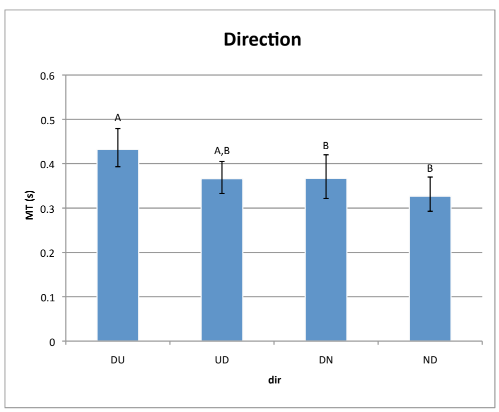

Figure 3.6 Least square means plot of errors for movement direction. Levels not connected by the same letter are significantly different from one another.

Interactions

With respect to the number of errors committed, there were significant interactions between A and W ($F_{4,80} = 7.05, p < 0.0001$), W and D ($F_{4,80} = 65.00, p < 0.0001$) and a four-way interaction among all the factors ($F_{24,480} = 1.56, p < 0.05$). These interactions are reflected in our proposed model (section 3.4.3).

Although error rate was high for the most difficult tasks, all trials were completed, since an error did not invalidate a trial. Thus, the whole set of trials was used in the movement time analysis.

3.4.2 Movement Time Analysis

We performed a full factorial analysis of variance (ANOVA) with repeated measures for task completion time. To reduce the chance of Type I errors, we used a Bonferroni correction in the analysis.

Main Effects

We found significant main effects for the independent variables A ($F_{2,40} = 760.1, p < 0.0001$), D ($F_{2,40} = 148.2, p < 0.0001$), W ($F_{2,40} = 885.5, p < 0.0001$) and dir ($F_{3,60} = 14.73, p < 0.0001$). Figure 3.7 shows the least square means plots for the main effects. With $alpha$ set at 5%, there were significant differences among all levels of A, D and W. The factor dir had differences among most of the levels, as depicted in the figure.

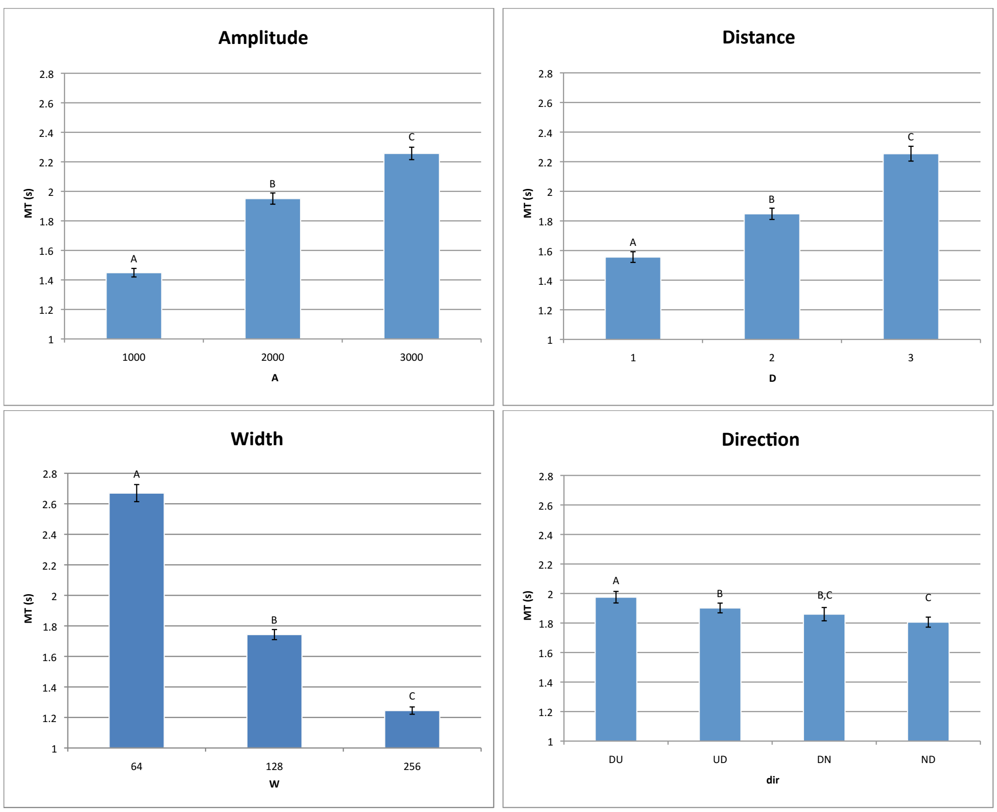

Figure 3.7 Least square means plots of time for significant main effects. For each factor, levels not connected by the same letter are significantly different from one another. Note that Y axes start at $1s$.

Although there is a significant difference among most levels of dir, it is clear that they are much smaller than the differences across the levels of A, D and W. We wish to provide a robust model that does not need to be tied to the direction of movement, so based on these small differences, we chose to evaluate models that are generic in terms of the direction of movement. We show in section 3.4.3 that there are, indeed, no practical differences in prediction time for each of the directions tested.

Interactions

We found significant interactions of *dir* and A ($F_{6,120} = 3.93, p < 0.005$), *dir* and W ($F_{6,120} = 4.70, p < 0.0001$), A and W($F_{4,80} = 26.78, p < 0.0001$), *dir* amd D ($F_{6,120} = 6.67, p < 0.0001$), and W and D ($F_{4,80} = 122.61, p < 0.0001$). There was also a three-way interaction among *dir*, A and D ($F_{12,240} = 2.85, p < 0.005$), and a four-way interaction among all the factors ($F_{24,480} = 1.66, p < 0.05$). Except for the interactions involving *dir*, all the significant interactions are reflected in our proposed model (section 3.4.3).

3.4.3 Regression Analysis

In order to find the best model of human motor behavior for distal pointing, we performed regression analysis using various possibilities for *ID*, as described in section 3.2. The regression was based on a single data point per condition, representing the average performance of all subjects in that condition.

Analysis Based on the Original Fitts' *ID*

As we hypothesized, the most common form of the Fitts' *ID* does not accurately model distal pointing. The regressed model of Equation 3.2 provided fit of $R^2 = 0.686$, which means that over 40% of the data points can't be explained by the model. This is an obvious result, since, as we can see in Figure 3.8, there are 3 distinct points for each *ID* value, each of which corresponds to one of the values of D.

We still need to verify if distal pointing could be modeled by the original Fitts' law if we have different a and b constants for each distance to the display. We regressed our data using the same *ID* from Equation 3.2, but this time, once for each of level of D.

In Figure 3.8 we can see that the variance increases as the distance to the display increases. Table 3.1 provides the coefficient of determination (R^2) for *ID* per distance to the display. For the up-close distance, the fit is very good, but it decreases rapidly as the distance gets larger. This analysis shows that the original Fitts' law is reliable only with distal pointing tasks with the user near the display, and is congruent with the findings of Stefels et al. (2007).

Besides lacking accuracy for higher values of D, omitting it from the index of difficulty of the task is not ideal. A more expressive model would account for different user positions relative to the display. To achieve such expressiveness, we analyzed the data based on ID_{RAW}.

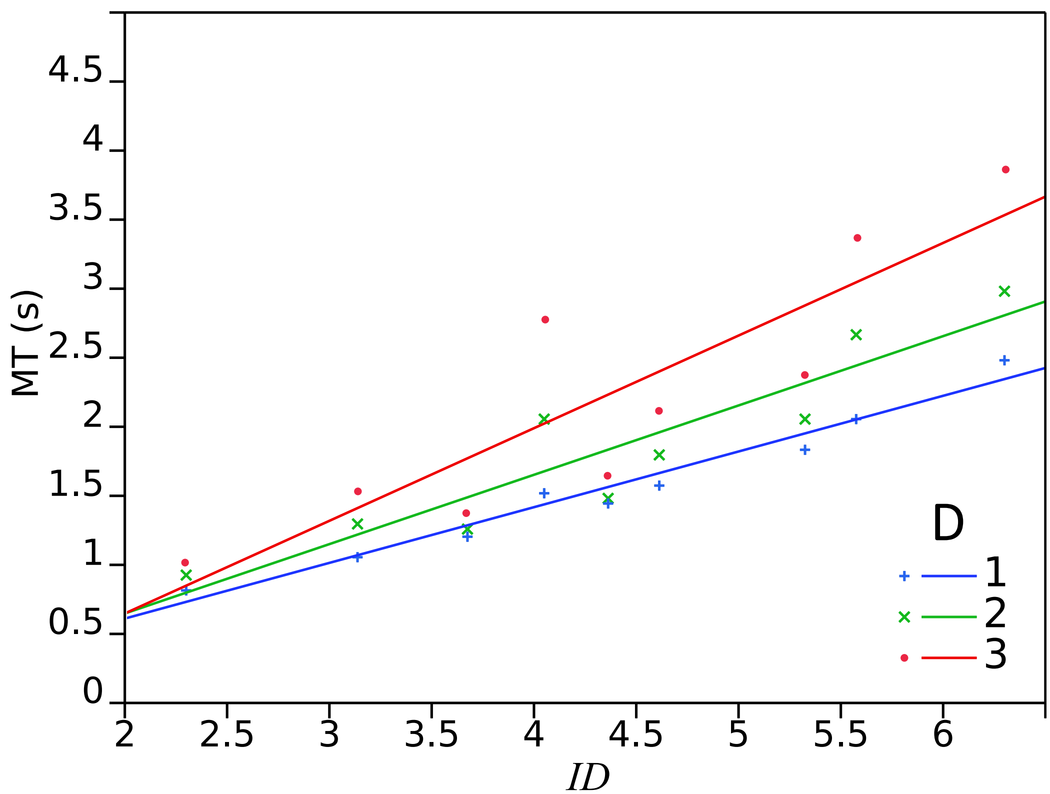

Figure 3.8 Fitts' law regression lines for each distance to the display based on the experimental data.

Table 3.1 Fit of Fitts' law for each distance to the display. *a* and *b* are the coefficients from Fitts' generic model, R^2 is the coefficient of determination and RMS is the Root Mean Square Error

***D*(m)**	***a***	***b***	**RMS**	R^2
1	−0.204	0.402	0.106	0.963
2	−0.362	0.502	0.267	0.864
3	−0.707	0.672	0.484	0.776

Analysis Based on ID_{RAW}

With ID_{RAW} (Equation 3.3), we are able to incorporate the user distance to the display in the difficulty of the task, Figure 3.9 shows the linear regression of the experimental data using ID_{RAW}, with a fit of $R^2 = 0.928$.

Although the model based on ID_{RAW} fits our experimental data quite well, we believe that

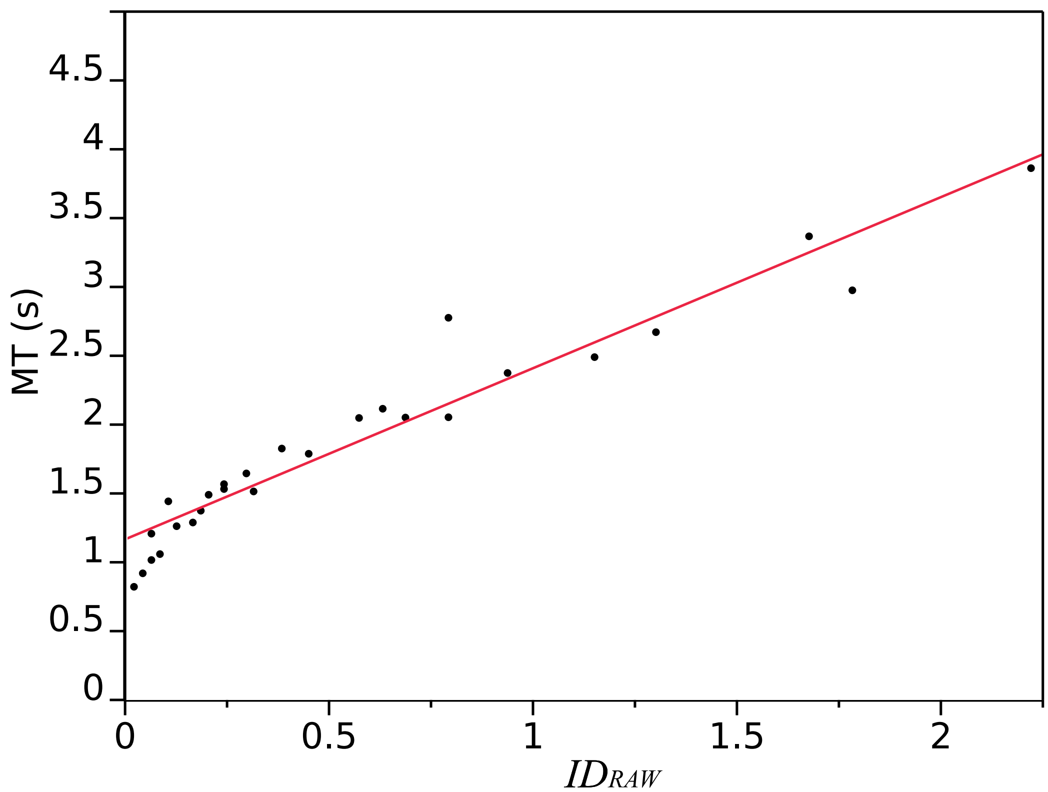

Figure 3.9 Regression line for ID_{RAW}.

we can provide a more generic model if we use angular measurements in the index of difficulty. In our experiment, users always stood in the center of the movement, and we would not be able to guarantee that the same model would apply if the user stood in different positions relative to the targets on the display. By using a model that considers angular measurements, we overcome this limitation, since the angular amplitude and target width will change according to the relative position of the user to the targets in the display.

Analysis Based on Angular Measurements

The model based on $ID_{ANGULAR}$ (Equation 3.6) results in a fit of $R^2 = 0.929$ for our experimental data, using $k = 3$, and Figure 3.10 shows the regression line. The correlation is almost identical to the one found with ID_{RAW}, but, as argued in section 3.2.3, $ID_{ANGULAR}$ is more generic and expressive.

This model has a good fit and overcomes the limitations of the previous models, but it still has problems. We observed the presence of outliers at the two highest values of $ID_{ANGULAR}$, which could be suggesting an exponential trend. Such a trend makes sense, since targets with very small angular widths are very difficult to acquire. We believe that when the angular width gets

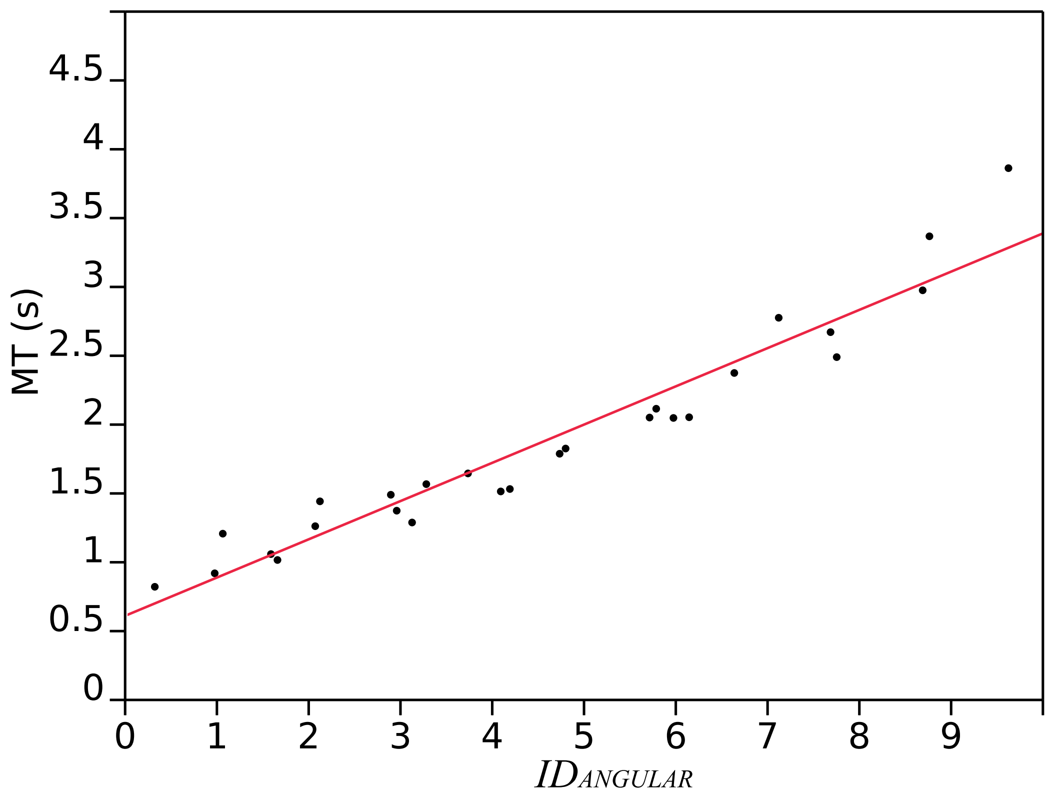

Figure 3.10 Regression line for $ID_{ANGULAR}$.

extremely small, a linear increase in the time it takes to acquire the target is no longer adequate. The movement time increases exponentially, as hand tremor and the Heisenberg Effect (Bowman et al., 2002) make it very difficult for the user to precisely position the cursor over the target and successfully acquire it.

Proposed Model

In order to address the limitations of $ID_{ANGULAR}$ we examined, and ultimately adopted as our preferred model, the following *ID* for distal pointing:

$$ID_{DP} = \left[\log_2\left(\frac{\alpha}{\omega^k} + 1\right)\right]^2. \tag{3.7}$$

where α is the angular amplitude of the movement, and ω is the angular width of the target. To avoid requiring two parameters to denote the size of the target (i.e., width and height), we assume that the largest dimension of the target is parallel to the direction of movement. Using the angular width and amplitude in the *ID*, it is possible to account for the user distance and position relative to the display, since both α and ω will vary accordingly. Using ID_{DP} from Equation 3.7 we were able to fit the data to the model with a coefficient of determination (R^2) as high as 0.961 (when the value of k is 3).

It is notable that ID_{DP} consists of the square of the logarithmic factor. One could argue from intuition that the degradation of accuracy for distal pointing as the angular amplitude and size increase is more than linear. Imagine the use of a laser pointer to light up a fixed spot on a wall. When one is close to the wall, the laser dot hardly moves at all. When one steps back from the wall the laser dot begins to jitter more and more, with the position of the dot limited by a rough circle whose radius increases in inverse proportion to the distance to the wall. The amount of jitter can be expressed by the area of the circle, which increases quadratically as the radius increases linearly. This argument is supported by the results of our experiment, which show that ID_{DP} fits the data better than any of the other candidate models.

We performed regression analysis of models containing different ID_{DP} by varying the value of the constant k in Equation 3.7. From the ANOVA results, we saw that there is more variance in the factors related to ω than to α, so we did not expect to see good fits for k values smaller than 1. The k value that best fits the data is 3.14, but we decided to use the rounded value of 3 for the sake of simplicity to discuss our regression analysis.

The predictive model of performance for distal pointing tasks under our experimental conditions is described as

$$MT_{DP} = 1.091 + 0.028 ID_{DP}, \tag{3.8}$$

where MT_{DP} is the movement time to complete a task and ID_{DP}, expressed in bits is

$$ID_{DP} = \left[\log_2\left(\frac{\alpha}{\omega^3} + 1\right)\right]^2. \tag{3.9}$$

A scatter plot with the regression line of this model is shown in Figure 3.11.

The fit of 96.1% and the scatter plot shown in Figure 3.11 were computed based on the mean for all the subjects per ID_{DP} value. This method has been used before (Accot and Zhai, 1997) and obviously provides the best fit of the data, since variance among subjects is not considered. To show that, even with the high variance from individual differences, we still can get a reasonable fit for the model when one data point per subject is used, we provide Figure 3.12. The fit of the model, for one point per subject per ID_{DP} value is $R^2 = 0.864$. In the figure, one can see how the variance increases with ID_{DP}, showing that the most difficult tasks depend more on individual differences, such as concentration and hand steadiness, than the easier ones.

These findings offer evidence that the model using ID_{DP} (Equation 3.7) accurately models distal pointing tasks. As shown in section 3.4, there was a main effect of the direction of movement. Based on this effect, one could argue that a single model should not be used to predict time independent of the direction of the movement. Table 3.2 shows four regression lines based on each movement direction, along with the fitted model for ID_{DP}. Neither the intercept nor the slope changed by a meaningful amount in these four models as compared to the single generic model. In

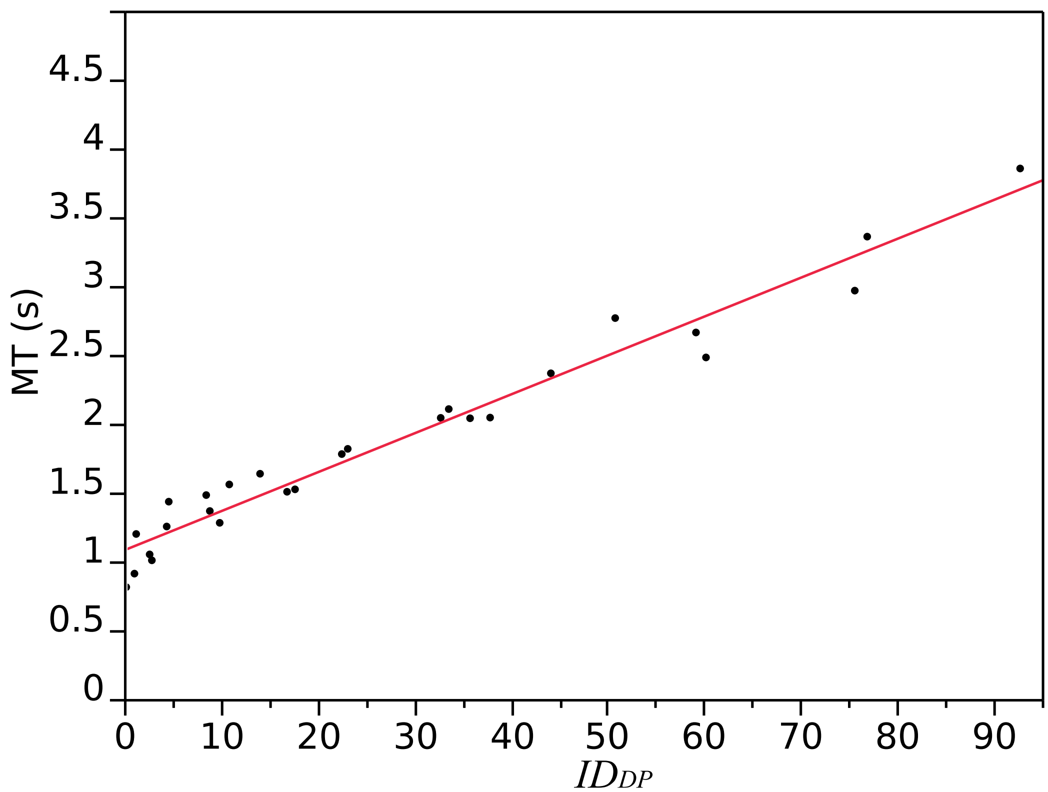

Figure 3.11 The scatter plot with the regression line for the fit of the model shown in Equation 3.8, with $R^2 = 0.961$.

practice, the predicted time for a very hard task ($ID_{DP} = 100$) would differ by a maximum of 0.3s or 7.7%. We conclude that there is no practical advantage for the use of one model per movement direction. The robustness of a single model leads us to recommend one generic model regardless of the movement direction.

Table 3.2 Fit of the model with ID_{DP} for each movement direction. a and b are the intercept and slope of the regression line, respectively, and R^2 is the coefficient of determination

dir	***a***	***b***	R^2
All four directions	*1.091*	*0.028*	*0.961*
Towards dominant	1.072	0.026	0.951
Towards non-dominant	1.087	0.028	0.945
Up	1.091	0.031	0.956
Down	1.123	0.027	0.960

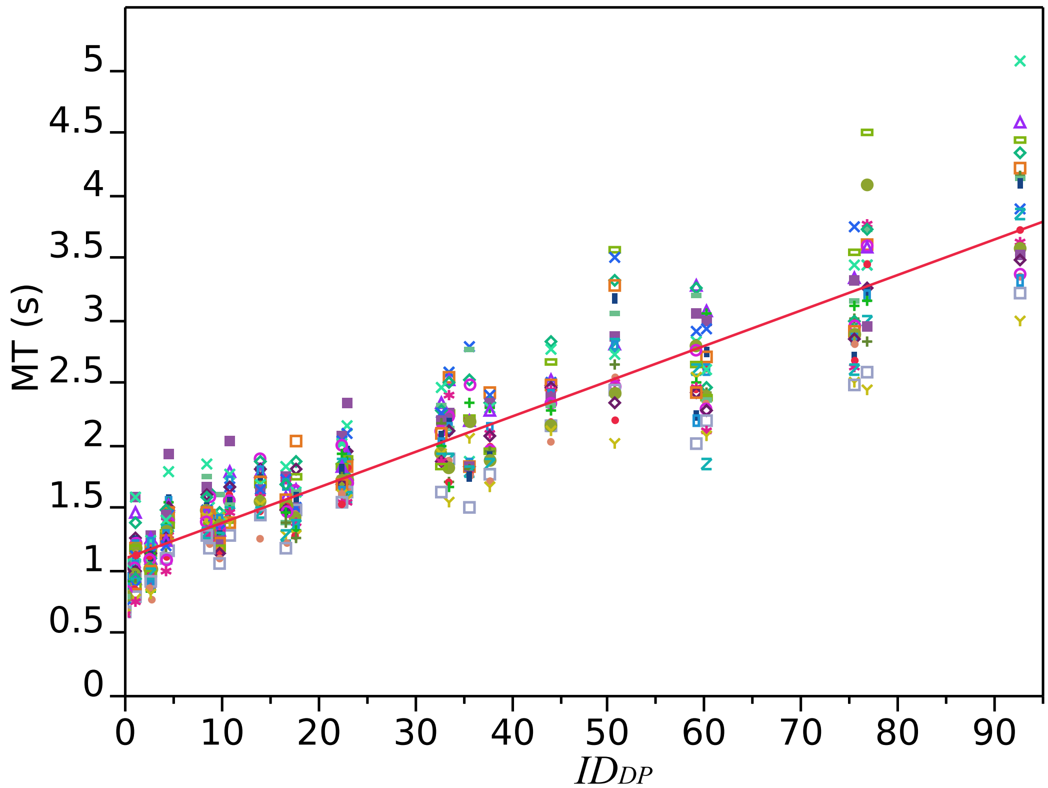

Figure 3.12 Scatter plot of the fit of the model with one point per subject per ID_{DP} value, with $R^2 = 0.864$. Each subject is represented by a different marker/color combination.

3.5 Design Implications

The model of distal pointing presented in section 3.4 is not only of theoretical, but also practical significance. The index of difficulty ID_{DP} has several important implications for the design of UIs that involve distal pointing.

First, and most obviously, our model indicates that the angular measurements of the target size and movement amplitude are the critical factors in task performance, rather than the linear measurements. In other words, the distance of the user from the target is highly significant. A target that may seem large when the user is directly in front of it may actually be quite difficult to acquire from a large distance (e.g., if the user walks backward, or if the target and user are at opposite ends of the display).

Designers could account for this in a number of ways. The entire UI could be designed to be usable from an assumed maximum distance, which would result in uniformly large targets; however, this approach is not likely to be practical for many applications. Alternatively, the size of targets, granularity of interaction, and/or layout of the UI could adapt based on the user's distance from the display (Peck et al., 2009). The distance from the display alone, however, does not determine the difficulty of distal pointing. A target at the opposite end of the display may subtend a very small angle when the user is near to the display, and of course the angular amplitude of movements will increase as the user moves closer to the display. An adaptive UI could therefore take as input both the user's position relative to the display and the area of the display to which the user is pointing. UIs that allow for interaction in the region of the display nearest to the user (e.g., pop-up menus instead of fixed-location pull-down menus) will also tend to reduce the angular amplitude and target size.

Second, our model clearly shows that angular target size has more influence on the difficulty of distal pointing tasks than angular amplitude. While the value of k may not be as high as three for all distal pointing setups, we expect that it will always be higher than one, because of the ease and speed with which users rotate their wrists in the ballistic phase of the movement, as compared to the great difficulty of holding the input device steady and making precise movements in the correction phase. This disproportionate influence of amplitude and size is the most important distinction, in our experience, between pointing movements based on translation across a supporting surface and pointing movements based on rotation in free space. Since there is a tradeoff between angular size and angular amplitude as the user moves closer to or farther from the display, it is very useful to understand that target size should be weighted more heavily.

Thus, increasing the size of targets should be the primary concern of the UI designer. Because of limited screen space, layout concerns, or aesthetic considerations, this will not be possible to achieve directly in most cases. Fortunately, there is a wide variety of techniques that can increase the effective target size without simply increasing the scale of the entire UI. Such techniques include utilizing the edge of the display for targets (as in the Macintosh menu bar), area cursors (Tse et al., 2007), bubble cursors (Grossman and Balakrishnan, 2005a), Object Pointing Guiard et al. (2004), expanding targets (McGuffin and Balakrishnan, 2002), sticky targets (Worden et al., 1997) user-controlled or automatic zooming (Forlines et al., 2005; Ramos et al., 2007; Worden et al., 1997), or automatic/manual control-display ratio adaptation (Forlines et al., 2006; Vogel and Balakrishnan, 2005). Although some of these techniques may be challenging to adapt to distal pointing tasks using 3D input devices, most of them should be applicable at least in concept. In fact, some of these approaches have already been used to design new interaction techniques for distal pointing (e.g., Vogel and Balakrishnan, 2005).

Finally, the quadratic growth of the index of difficulty in our model indicates that distal pointing tasks can get more difficult very quickly as the ratio α/ω^k grows. There comes a point, especially with very small targets from large distances, where such tasks become nearly impossible. Thus, besides designing UI layouts and interaction techniques that minimize angular amplitude and maximize angular target size, designers should also provide alternatives to distal pointing to allow users to continue interacting with some reasonable level of usability regardless of the user's position or the interface elements on the screen. These alternatives could range from discrete input such as button presses to cycle through the available targets, to keyboard shortcuts, to voice interfaces.

A prototype interaction technique developed in our laboratory serves as an illustration of a distal pointing technique for which our model would predict good performance. This technique, which we call Absolute and Relative Mapping (ARM) Raycasting, uses manual control of the C/D ratio to allow users to increase the effective angular width of targets as needed.

By default, ARM simply uses an absolute raycasting technique in which the handheld input device defines a pointing ray, and the cursor appears at the intersection of this ray with the screen. When finer control is needed, the user presses a button that temporarily invokes a “precision mode” with a 10 : 1 C/D ratio, increasing the effective angular width of nearby targets by a factor of ten. Its effect is that a wrist rotation causes the cursor to move by a much smaller offset from the point where the precision mode was activated, giving the impression of a “slow motion” cursor. When the button is released, the cursor jumps back to the position determined by absolute raycasting.

The increased C/D ratio decreases the value of ID_{DP} significantly, but users may still have trouble acquiring very small targets if they cannot perceive whether the cursor is over the target. Therefore, ARM also includes a zoom lens that appears around the cursor when the precision mode is active. The user can control the level of zoom with a scroll wheel on the handheld pointing device.

We have performed informal tests with a prototype ARM technique, and have found that it clearly increases the ease and precision of selection and placement tasks with small targets on a large high-resolution display. In the near future we will improve its design and investigate ways to automatically infer the C/D ratio and zoom level depending on the user’s distance and angle to the cursor position, which represents the area of interest on the display.

3.6 Summary

In this chapter, we have shown the rationale behind the definition of a predictive model of distal pointing performance. We were able to successfully answer research question number 1 (section 1.2), in that we were able to effectively model distal pointing behavior with a very high fit for the data.

Although the model was inspired by Fitts’ law, and characterizes essentially the same cognitive tasks – that of performing an aimed movement – we found some fundamental differences between distal pointing and motor behaviors expressed by Fitts’ law. First of all, the angular, rather than linear, measurements of the task parameters are of importance. By using angular measurements, the model accounts for the user relative position to the target, and absorbs the fact that distand objects are more difficult to point at.

A second peculiarity of the distal pointing model, compared to Fitts’ law, is that the movement amplitude is weighed different between the movement amplitude and target size. The very nature of distal pointing affords very quick wrist rotations to move the cursor through vast areas. The difficulty of precise cursor refinements, however increase very rapidly as targets become smaller.

Chapter 4

Designing High-Precision Distal Pointing

The predictive model of distal pointing interaction introduced a set of design guidelines that, when followed, should provide usable and precise distal pointing techniques. The goal of this chapter is to address such guidelines and propose a set of interaction techniques intended to overcome distal pointing limitations that are evident in the model.

First we discuss inherent precision issues with distal pointing and propose a number of design choices to overcome these problems. Then, we elaborate on four design choices to increase precision of distal pointing interaction, along with the design of the interaction techniques based on each choice.

Three of the design choices are intended for immediate interaction tasks. Immediate tasks are those which are performed by a single, immediate action, such as the click of a button. The last design choice to improve distal pointing presents the new concept of selection by progressive refinement. This style of interaction consists in breaking down a selection task that requires high precision for an immediate action into several coarse, low precision phases.

4.1 Precision Issues

Distal pointing is a practical way of achieving rapid interaction with remote displays while standing and being able to walk. However, it lacks the necessary precision to be a useful interaction metaphor. The human hand has a natural tremor for which we cannot control, which ranges from [8-10Hz] (Riviere et al., 1997). On top of that, jitter from tracking systems also detriment the precision of distal pointing. It is clear that distal pointing in its original form is not precise enough when the visual size of targets are small.

The amount of precision required to point at a target from a distance relates to its angular, rather than to its linear, size. We can illustrate the reason for the relationship of distal pointing precision to the target's angular size by imagining someone trying to point steadily at a distal target. The average linear velocity of the cursor in that situation, taking into account the hand tremor, increases rapidly as the pointing device moves away from the interaction surface. In other words, the area of uncertainty where the cursor position can be increases the farther the input device is from the display.

4.1.1 Basic Distal Pointing

The basic distal pointing technique was implemented by defining a vector coming from the input device towards the display. The position of the cursor was determined by the intersection of the vector with the display surface. This technique served as a basis for all precision-enhanced techniques. To reduce the effects of hand and tracker jitter, we used a dynamic recursive low pass filter, proposed by Vogel and Balakrishnan (2005). All techniques we implemented, including basic distal pointing, uses this filter on top of the raw tracking data, and before any improvements were made. This allowed for a smoother cursor, but was highly insufficient to provide the high levels of precision required to point at small targets from a distance.

For the distal-pointing-based interaction techniques we designed, we assume that we have a 6-degrees-of-freedom (6DOF) tracked input device that provides a direction vector to the display. We acknowledge that this is not a setting that can easily be made widely available, as it requires precise tracking. However, consumer-level devices, such as Sony's Playstation Move (2011), are starting to provide 6DOF tracking with relatively high levels of precision. Even in the absence of absolute pointing devices, approximations to the assumption to 6DOF tracking could be achieved by the use of relative pointing devices, such as gyros and accelerometers that are ubiquitous in today's world in mobile devices. Such approaches would require some sort of calibration, like the indication of pre-determined points on the display, so that the input device had a frame of reference. We acknowledge that current technologies present in mobile devices do not provide very precise readings. Thus, to avoid loss of precision over time, some form of re-calibration should be performed when interaction was required for longs periods of time.

For the intents and purposes of this research, however, we assumed that we have sub-millimiter 6DOF tracking precision available. To achieve that level of precision we used, initially, a VICON MX system with eight cameras, and, later, an Optitrak system with eleven cameras, both of which provide a robust tracking space for our experimental environments.

Usability Improvements

In addition to the base interaction technique used for the high-precision distal pointing techniques we developed, we furnished the testbed application for the interaction techniques with some usability enhancements.

In the application scenario that we used, there was no need to point outside the display, so we implemented a "snap to edge" of the cursor to the display. That caused the cursor to always be inside the display, even if the absolute pointing position would determine the cursor to be outside. The practical advantage of such a feature is that it increases the effective size of objects which are located on the edges of the display, as the cursor stops moving before it leaves the display. This is similar to the behavior of the mouse cursor in desktop computers, which has long been taken advantage of by the Mac OS menu bar, which is reachable even if the movement of the mouse would cause an overshoot of the cursor.

In order to improve the visual perception of the cursor, we increased its size and gave it a color that contrasted highly with the background.

4.2 Alternatives to Increase Precision

In order to overcome the lack of precision inherent to distal pointing, we investigated several approaches to increase the ability to precisely point at a distance. First, a low-pass filter can be used to smooth out natural hand tremor as well as tracking jitter. This filter however, should care not to to offer cursor stickiness, and thus, we employed a dynamic recursive low pass filter, described by Vogel and Balakrishnan (2005).

We note, however, that the design guidelines suggested by the distal pointing model lead us to understand that smoothing out the cursor may not solve the problem of precision entirely. In fact, as the user gets far from the display, the difficulty to point at a target becomes exponentially harder, as the visual size of the targets become smaller. Thus, we need to find other ways to address the lack of precision introduced by distal pointing.

Precision enhancements can be categorized in two actions that deal with the motor precision required to point, and one action related to the visual perception of the target. To increase the motor precision, the possible approaches are to increase the effective target size and to reduce the effective movement amplitude. We also know from the model that it is more important to increase the target size than to reduce the movement amplitude, so enhancements should care more of making targets bigger than bringing them closer together.

We explored several ways of increasing the effective target size. One approach, which also increases the visual perception of the target, is to zoom-in in the whole area around the cursor, for as long as interaction is needed. This approach, however, may introduce a loss of application context, as the peripheral area around the zoomed-in area is lost.

Another alternative is to increase the control/display (CD) ratio to achieve more precision. By increasing the CD ratio, a small physical area is matched to a large movement on the input device, causing the interaction space to be much larger than the visual space. This is an interesting approach, but it needs to be used wisely, as an increased CD ratio for large movements cause the effective movement amplitude to be much larger than it needs to be, and may lead to excessive fatigue and tasks that are harder than they need to be. A good design choice is to provide an increased CD ratio only when precision is needed, namely, to increase the effective size of small visual targets.

Other form of increasing the effective target size is to change the mapping mode of pointing. With absolute pointing, as well as with increased CD ratio, the position of the cursor is determined by a function of the point of intersection of the ray emanating from the input device with the display surface, as a direct transfer with absolute pointing, or by multiplying by a constant factor with increased CD ratio. By changing the mapping mode, the position of the cursor is determined not by the ray intersection, but by a relative mapping of the input device rotation to the displacement of the cursor. This choice solves the problem introduced by constant increase in CD ratio of having a varying linear speed of the cursor based on the angle of the ray to the display. Depending on the angle of the ray with the display, the same rotation of the input device causes a different cursor displacement, with the precision provided by the increased CD ratio reducing drastically as the angle between the ray and the display approaches zero. When a mapping between the pointing device orientation to the cursor displacement is provided, the amount the cursor will move is constant as long as the orientation of the input device changes by the same amount. This precision improvement also needs to be used consciously, as rotating the input device by a large amount causes the cursor to move a small distance, and should only be used during the refinement phase of movement. Again, it is important that the enhancement be offered only when precision is needed, with small visual targets.

A way to solve the problem of deciding when to provide precision enhancements is to use an adaptive CD ratio (Frees et al., 2007; König et al., 2009) based on the user requirement for high precision. As users try to be precise, the rotational speed of their wrists, and, by extension, the input device will be reduced. Taking that into account, the CD ratio can increase as the rotational speed of the wrist decreases, providing smoothly more precision as the user needs it.

Visual Acuity

We believe that not only the motor precision influence the difficulty of a distal pointing task, but also the visual acuity. Visual acuity in a distal pointing task can be characterized by the quality of the visual feedback control loop – the ability the user has to know whether the cursor is inside the target by means of the visual system. Ultimately, this relates to the visual target size. Previous studies (e.g., Carlton, 1981) have hinted that visual acuity may have an impact on task performance with distal pointing .

4.3 Physical Navigation Considerations

Previous work has found that users take advantage of physical navigation to find and understand data, particularly in large high-resolution displays. Ball et al. (2007) showed that larger displays afford more physical navigation, which leads to improvement in performance and user preference. The nature of the tasks conducted by Ball et al. (2007) was of visual search and cognition, and the only interaction needed was to navigate the data, either physically, by walking and moving the head, or virtually, by panning and zooming a small display. In that work, the authors did not test physical navigation effects on interactive tasks – particularly tasks which required selection and manipulation of objects in the visualization application. These types of tasks are the focus of the research presented here.

With distal pointing, which is used specifically to select and manipulate objects in a visualization distant from the user, we believe that physical navigation may have a different impact than to search and cognition tasks. When the task requires the user to directly point at the display to control a cursor, physical navigation can have a negative impact, because the walking may distract the user from the distal pointing task and the user may need to switch modes from walking to pointing. We have also found in previous studies and observations (e.g., Kopper et al., 2008) that many times users prefer not to physically navigate, or to only navigate sideways, but always keeping a distance from the display. We believe that is accounted due to the fact that some users want to be able to perform distal pointing tasks in an "lazy" manner, as long as the pointing technique they are using provides for that. In fact, even in cases in which the pointing technique was the basic one, without any precision enhancements, a large number of participants still chose to not walk, and got very frustrated with the most difficult tasks, causing long pointing times and high error rates.

We can explain this apparent contradiction from our findings to previous studies if we look at the different task modalities in the application scenarios from each study. With the previous study conducted by Ball et al. (2007), no direct interaction with the visualization was necessary. Our task scenarios, however, not only required direct interaction with the visualization, but did that in a pure form, that is, without requiring any cognitive understanding of the data. We don't know what impact the lack of physical navigation on pointing tasks have on the overall understanding of the data, and that is an important piece of work which this research does not investigate, but future work in that direction is encouraged.

The observations we had in the past are one more reason to support the development of high-precision distal pointing techniques, as they may allow users to exercise that desire of lazy interaction, without taking back the choice to walk if they want. While basic distal pointing requires users to walk up close to the display to interact precisely, high-precision techniques offer users the choice of not walking while being precise.

4.4 Zoom as High-Precision Distal Pointing

One of the initial ideas we had to increase precision of distal pointing was to zoom-in into the portion of the visualization in which high-precision interaction was necessary.

4.4.1 Rationale

The rationale for using zoom to increase precision is that when zooming-in, the visual size of the zoomed-in objects increase, and their selectable area increase proportionally with the zoom factor. This can theoretically provide infinite precision, as the zoom factor can always increase as more precision is needed. However, one of the great advantages of large high-resolution displays is the amount of data that can be visualized at each point in time, providing great levels of detail, but also an overview of the data, without having to change the visualization (Ball et al., 2007). This can cause zoom techniques to not take advantage, particularly, of large high-resolution displays.

4.4.2 ZELDA – Zoom for Enhanced Large Display Acuity

As a first attempt at a high-precision distal pointing technique that uses zoom as the means to achieve high precision, we designed the Zoom for Enhanced Large Display Acuity (ZELDA) techniques (Kopper et al., 2008). ZELDA uses bi-manual distal pointing and is very powerful, but also quite complex.

Design

ZELDA uses two distal pointing devices to control a zoom window in addition to the regular cursor. Figure 4.1 illustrates its functionality. The main idea behind ZELDA is to provide a magnified view of objects in the zoom window, which not only improves precision, but also addresses human visual acuity limitations by improving the quality of the visual feedback loop, which may lead to improved performance (Carlton, 1981).

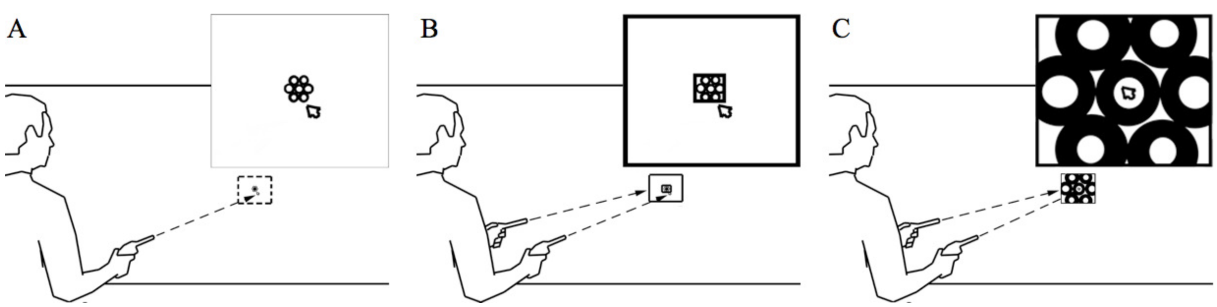

Figure 4.1 ZELDA technique. The insets show a detailed view of the cursor area. (A) The user just points at the display. (B) The user moves the zoom window over the objects so that the region of interest is within the inner rectangle. (C) The user freezes the zoom window and the area covered by the inner rectangle is magnified.

The ZELDA user uses her dominant hand for basic distal pointing and for left-click functionality. The device in the non-dominant hand is used to control the zoom window's position, size and zoom factor.

The zoom window can be either moving or frozen and a button in the non-dominant hand device is used to switch modes. When the zoom window is moving, the user is able to place it with basic distal pointing. So as to not lose contextual information, the zoom window is semi-transparent when moving, and contains an inner rectangle that shows the area that it will zoom to when the window is frozen (figure 4.1b). A scroll wheel in the non-dominant hand device is used to change the zoom factor when the user is pointing towards the display (figure 4.1c).

When the user points to the side, the scroll wheel can be used to resize the zoom window's horizontal dimension; the vertical dimension can be resized when the user points up or down. In either case, dark yellow lines are drawn on two sides of the zoom window to indicate that resizing is possible. The zoom window is always resized around its center. In our implementation, the zoom window starts as a square of 1536px on a side, with the zoom factor set to 2, so that a square area of 768px on a side is doubled in size when the zoom window is frozen.

By offering the zoomed-in view, ZELDA not only increases precision for selection and manipulation tasks, but also addresses visual acuity limitations in the zoomed area. This feature is useful for several reasons. First, for tasks that require very high precision, such as selection of very small objects or placement in very tight places, it may not be enough to increase the precision of movement if the objects are too small to see or the placement area is too tight to verify if a given object is completely within this area. Another advantage of zooming is that it can be used for tasks that involve visual perception as well as interaction. For example, in tasks in which the user wishes to select only icons with a particular label, if the label is too small to read, the zoom window can allow the user both to read the label and select the icon. Finally, ZELDA's zoom window can be used by itself for pure visualization applications that contain very dense and detailed information that cannot be seen at the original zoom level. One example is geospatial visualization, in which very detailed maps are shown on large high-resolution displays. ZELDA can be used to display areas of such maps in more detail without requiring the user to walk up to the display.

Physical Navigation

Since, with ZELDA, high precision is provided on demand by the non-dominant hand, the user has the option to freely walk and perform tasks getting precision from physical navigation, rather than using the high-precision feature from the technique. However, it does afford users to remain at a distance from the target and still be successful. With ZELDA, the gain is both in precision and visual size, and they are coupled, such that the higher the zoom level, the more precision and visual size is provided.

Limitations

Although ZELDA is quite a powerful technique, it also carries a lot of overhead. Controlling the zoom window in the non-dominant hand, and offering the ability to position, resize and change the zoom factor causes a simple selection task to potentially include many subtasks which makes the overall task much more complex then it needs to be.

Initial Findings and Lessons Learned

We performed an comparison of ZELDA with basic distal pointing and an early design of ARM (section 4.5) (Kopper et al., 2008). We found that, indeed, performance with ZELDA was highly dependent in user strategy, and in many cases its performance was poorer than either ARM or basic distal pointing, or both. We conclude that the power of controlling the zoom window and allowing the user to decide how much precision is needed, as well as how big the zoomed area must be increases the complexity of the task more than it needs. However, we acknowledge that using zoom as a visual aid may be an important factor in decreasing the difficulty of a task, especially in situations in which the target zone is very narrow, and the user is at a distance to the target, either due to personal preference or to application requirements[1].

Improvements

Precision can be achieved through other means then zooming, such as increasing the CD ratio or using a mapping of the device rotation to the displacement of the cursor. However, assuming that the user needs or wants to remain at a distance from the display, only zooming can address visual acuity limitations, which may be quite important for the ease of tasks that require very high precision. In such tasks, the margin within which the cursor is on the target can be very narrow and it may be difficult to visually know whether the target is active for selection even if the motor precision allows the user to be successful. For those situations, the use of zooming can be beneficial, and we present variations of the techniques we evaluated in which a zoom window is attached to the cursor when precision is required (sections 4.5.2 and 4.6.2).

4.5 Discrete High-Precision Modes

Rather than having the user controlling all aspects of high precision, we propose a constant gain in precision that is explicitly activated by the user.

4.5.1 Rationale

We learned from ZELDA that it can be more important to make a technique simple, yet not as powerful as it could be, in order to increase its usability. One choice for achieving that is to provide a discrete control to activate high precision when needed. In this case, the application has a pre-defined precision increase factor, which causes the cursor to move slower than it would with basic distal pointing. This approach is very simple, and still gives the user control of whether he needs high precision or not.

[1] We believe that some applications, such as interactive map visualizations, may benefit from interaction while maintaining an overview. With a high-precision distal pointing technique, that can be achieved, and the matter of having users away from the targets becomes an application demand, rather then a personal preference.

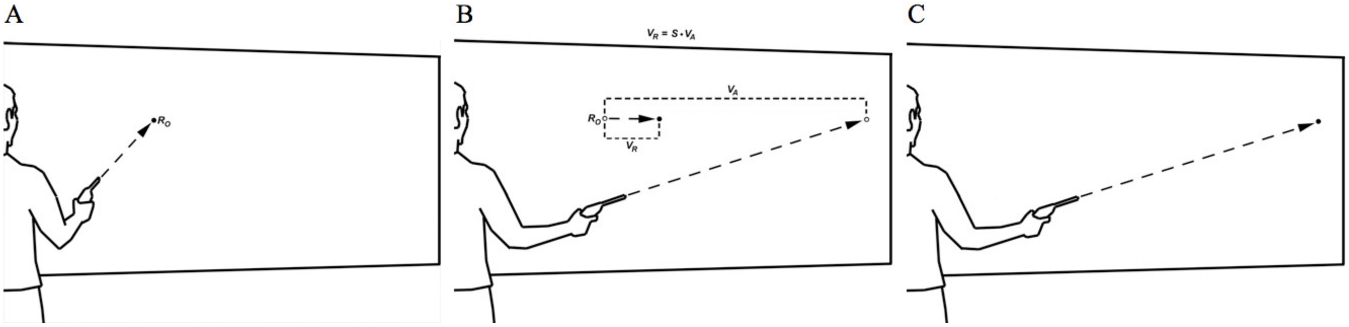

Figure 4.2 ARM Raycasting. (A) The user activates relative mapping at the R_0 position. (B) The user points absolutely to the new position determined by the vector V_A but the cursor appears in the relative position determined by V_R, which is $V_A S$. (C) The user deactivates relative mapping, causing the cursor to return to the absolute mapping position.

4.5.2 ARM - Absolute and Relative Mapping distal pointing

We designed a technique that offers a high-precision toggle when the user needs to increase the precision of distal pointing. There were several design attempts and we made choices to determine the high-precision activation toggle, the type of mapping that would be used, and the amount of precision increase provided.

Initial Design Attempts

Initially, we designed ARM by changing the control/display (CD) ratio. That was done by mapping the movement with the input device to a small movement area of the cursor, as can be seen in figure 4.2. After trying several mappings, we found that a 10:1 CD ratio provided a high enough precision while keeping the cursor responsive.

This design choice, of remapping the cursor movement area, provides high precision but has some drawbacks. Since the mapping is relative to the absolute pointing position, the cursor position is defined basically as a multiplication of the displacement of the absolute pointing position from the moment high-precision mode was activated by the inverse of the CD ratio. In other words, the cursor displacement when arm was active was $1/10$ of the absolute pointing displacement. Since the absolute pointing position is determined by the intersection of a vector coming from the input device with the display plane, the velocity of the cursor depends on the angle of the pointing vector to the display, being higher the more acute the angle is. That means that the absolute cursor velocity will tend to infinity as the angle between the pointing vector and the display plane tends to zero, and the $1/10$ mapping of the ARM cursor position will be less precise as the angle gets smaller. The practical effect of this is that the larger the ARM movement, the less precise the pointing is, which somewhat defeats the purpose of the increase in precision with ARM. It is important to note that, when ARM is not active and the cursor position is determined by intersection of the pointing vector with the display, the same problem with losing precision as the angle between the ray and the screen exists, but to a lesser degree, since, with ARM, the input device active angle – the angle the input device can be rotated with the cursor being visible at the display – is 10 times larger than that of absolute pointing, causing the projected absolute cursor position to be very high, and leading to a rapid cursor velocity even with ARM activated.

Apart from the problem stated above, the original arm design, as a constant increase in the CD ratio, also causes the precision required to point with ARM to vary based on the user distance to the display. Since it is a direct mapping from absolute pointing, as the absolute pointing becomes less precise, so does the initial design of ARM.

High-Precision Activation

An important design choice with ARM is how to activate the high-precision mode. An important choice to make was what type of control to use to designate the high-precision mode. It could either be a toggle or a press-and-hold action to indicate high-precision mode. The advantage of using a toggle is that the user does not have to continuously press in order to keep the high-precision mode activated. However, we expect that high-precision mode would be provided only for a small part of the selection task – when precision was needed. Another problem with a toggle is that some sort of feedback must be given, since there would be no continuous feedback from the user continuous action. Since we expect that high-precision mode would be active only for short periods of time, we chose to using a press-and-hold action for high precision.

Having determined the action to activate high precision, we had to decide on how to present it to the user. Ideally, we would like to provide the high-precision mode activation in the same hand as the one used for pointing, as to keep the other hand unoccupied. However, previous research on bimanual and inter-finger coordination suggests that people differentiate better between limbs than between fingers in the same hand (Fleishman and Hempel Jr, 1956; Guilford, 1958; Seifert and Chollet, 2008). Especially, with a press-and-hold of a button to keep high precision, coupled with a press-and-release action of selecting an object caused a lot of overhead to users in exploratory analyses. We, thus, discarded the use of a secondary button to express the high-precision mode with ARM.

Another alternative we tried to keep the technique single-handed was to take advantage of the assumption that the grip of the input device is firmer when the user requires more precision. We adapted a Phidgets (2011) pressure sensor to the pointing device in order to have the high precision activated when the device had a firm grip. Figure 4.3 shows the pointing device with the pressure sensor attached. Again, we performed exploratory analysis and found that forcing users to a firmer grip caused a lot of trouble. We had to set the threshold somewhat high as to not cause the high-precision to be activated unintentionally, and the firm grip ended up causing a large movement of the input device, and, by extension, the cursor, before high precision was activated. This problem has been referred to as the *Heisenberg Effect* (Bowman et al., 2002). This meant that the cursor would provide high precision in an area far from the target, and a very large movement of the input device was necessary in order for the cursor to go back to the area of interest, and this design alternative proved not to be usable.

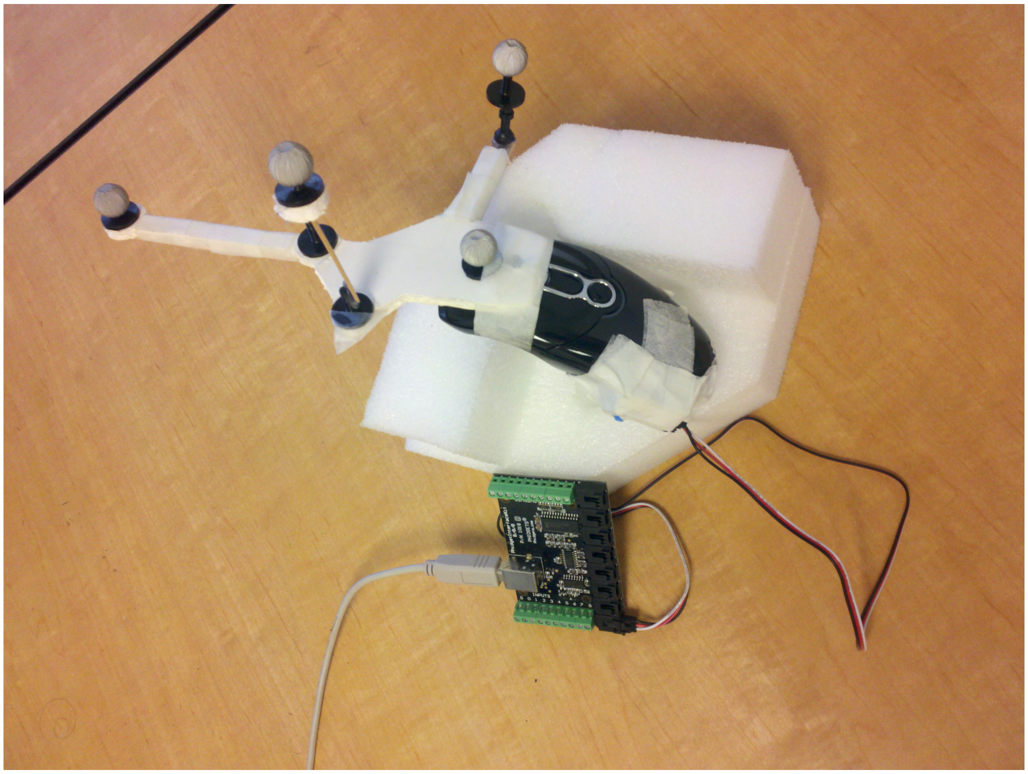

Figure 4.3 Input device modified to activate high-precision mode by pressure of grip.

We finally decided to use a single button press in the non-dominant hand, which proved to be easily separable from the input device, and control was easy for all users.

Final Design

Considering the issues raised from the initial design attempts, we decided to redesign ARM in a way to increase the precision by the same amount independently of the precision required by absolute pointing. We considered using varying CD ratios based on the absolute position velocity, but since ARM has a discreet high-precision toggle, we thought it would be best to change the pointing mode altogether. In order to achieve the goal of keeping the pointing precision independent from the absolute pointing position, we decided to map the device rotation to the cursor displacement when high-precision mode is active. When high precision is activated, the mapping changes from the intersection of the pointing vector with the display to the rotation of the input device. The rotation is mapped to an arc of circle, whose radius was arbitrarily set to $200px$ (based on informal experimentation), and that arc is transferred to the cursor position, causing the cursor movement to depend only on the velocity of the input device rotation which is controlled by the user, as opposed to a function of the absolute pointing position. This design causes the gain in precision to be constant independently of the position of the user relative to the target.

Physical Navigation

As with ZELDA, ARM offers the user the choice to use or not the high-precision mode, giving the option to achieve high precision through physical navigation, rather than the technique feature. However, it affords the user to remain at a distance, as long as the user has enough visual acuity to complete the task. Here, visual acuity relates to the ability the user has to visually assess whether the cursor is at the correct spot before committing a selection, or a dragged object is inside the target area before it is released. In order to accommodate for situations in which there is not enough visual acuity (i.e. the target is too small or the margin for drop is too narrow to see from a distance), we provide a variation of ARM that addresses these limitations by providing a zoomed view.

Zoom Variation

ARM with zoom consists in a zoom window that pops around the cursor every time the high-precision mode is activated. That causes the area around the cursor to zoom-in every time more precision is required. The disadvantage of this approach is that there is a loss of context, since anything behind the zoom window, except the zoomed area is occluded. Since the zoom is provided only when high precision is needed, which is during the refinement phase of a pointing task and is not expected to be long, we believe that this break in context may not affect the understanding of the information as much as the gain in visual acuity will help the task completion, but that remains to be tested.

With the zoom variation, ARM affords users to remain at a distance with any task, as the gain in visual acuity enables users to correctly perceive any task.

Limitations

The design of ARM enable a simple transition between normal and high-precision mode without the overhead of ZELDA. However, it does so by allowing the user to decide when high precision should be used, and, thus, the user should employ a good high-precision activation strategy to maximize performance. If high precision is not activated when a task has a high index of difficulty, the chance of error will increase, and the time to perform a pointing task will be longer. On the other hand, if ARM is activated for an easy task, it may cause more distraction than improve performance. More importantly, the distance of the cursor to the target when ARM is activated will determine how much α – the angular amplitude of the movement (eq. 3.4) – will increase. From the distal pointing model, we know that the greater the value of α, the higher the index of difficulty of a task. If high precision is activated too far from the target, alpha will be too high and the index of difficulty of the task will be much higher. With the ARM circle radius set to $200px$, a full $360\,\mathrm{deg}$ rotation of the input device is required to traverse $200\pi px$, or $630px$.

Another potential limitation of ARM is that the gain in precision is always the same, independently of the required precision by the task.

With the zoom variation, there is a perception incongruence that users have to adapt for. Since the zoom window is shown around the cursor, the zoomed-in area moves with the cursor, causing the magnified view of the objects in it to appear to move against the cursor movement.

4.6 Continuous Gains in Precision

In order to address the limitations of high-precision activation strategy and the constant gain in precision, we propose the idea of continuously increase the precision, as required by the user.

4.6.1 Rationale

We believe that ARM effectively increases precision of distal pointing and makes selection of small targets much easier from a distance. However, we see that ARM has a few characteristics that may not be ideal in all situations. First, it uses an explicit high-precision toggle, which gives the user the choice of activating it when needed. This gives power to the user, but it may be at the expense of causing some users to employ poor strategies when using ARM for high-precision distal pointing. Also, ARM provides a constant mapping of the wrist rotation to the cursor displacement when high precision is active. Of course, the mapping can be changed to make the high-precision mode more or less precise, but it's done at the application level, and not at the task level. This causes a tradeoff between having too much precision offered for not so difficult tasks or too little precision for tasks that are very difficult.

The idea of a continuous gain in precision, is to provide the precision the user requires for a particular task. For that, we use a mapping of the CD ratio of the input device to the cursor displacement based on the rotational velocity of the input device. Aimed pointing is defined in two steps, namely a ballistic and a correction phase (Grossman and Balakrishnan, 2005b; Liu et al., 2009; Woodworth, 1899). During the ballistic phase, a rapid movement is made in such a way that very little visual feedback is used and is intended to throw the cursor in the general target region. Once the cursor gets there, small motor corrections are made based on the user perception of the position of the cursor in relation to the target, and, generally, the closer the cursor is to the target, the slower the user will rotate her wrist, until the cursor is on the target. The techniques we designed – DyCoDiR (section 4.6.2) – uses that information and applies a higher CD ratio the slower the wrist rotates.

This approach has been used before for manipulation and interaction with objects in immersive virtual environments (Frees et al., 2007), and for distal pointing interaction at large high-resolution displays (König et al., 2009). While PRISM Frees et al. (2007) offers a set of interaction techniques that use dynamic CD ratio control to manipulate virtual object in a large scale, König et al. (2009)'s Adaptive Pointing technique shares several characteristics with DyCoDiR. However, it uses a more complex implementation, and, more importantly, Adaptive Pointing does not take into account the user's relative position to the interaction region on the display.

4.6.2 DyCoDiR – Dynamic Control/Display Ratio Distal Pointing

The second high-precision distal pointing technique that we designed, called **Dy**namic **Co**ntrol/ **Di**splay **R**atio (DyCoDiR) distal pointing works by linearly increasing the CD ratio as a function of the wrist rotation speed.

Acknowledging that the precision required to point at a target increases as the user moves far from the target, due to the projection of the input device jitter onto the display causing an increased uncertainty radius of the cursor position, DyCoDiR provides more precision when the user is far away from the interaction region. As the user gets closer, the gain in CD ratio is reduced up to the point in which the input device is touching the cursor, where the CD ratio remains at 1:1 no matter how steady the user is. That is because the ray emanating from the input device has no length, and what is emulated is a touch screen interaction, in which the display itself can be used as a supporting surface, allowing the input device to be steadier. It's important to note that this is a limit situation, and very rarely, the distance between the input device and the cursor will be zero. Most likely, there will always be some precision enhancement, only to a lesser extent as the user gets closer to the cursor.

Equation 4.1 illustrates how the CD ratio for the cursor is calculated, where v is the cursor velocity and d is the distance of the input device to the cursor.

$$CDRatio = \begin{cases} 1 & \text{if } v \geq 5m/s \\ min(d\left(\frac{5}{v}\right), 1) & \text{if } v < 5m/s \text{ and } d < 2m \\ d\left(\frac{5}{v}\right)^{d/2} & \text{if } v < 5m/s \text{ and } d \geq 2m \end{cases} \tag{4.1}$$

Note that, as the user moves away from the cursor, the CD Ratio increases exponentially when the distance is more than two meters. That causes the gain in precision to be virtually unlimited, with the user being able to get sub-pixel precision even at a large distance.

As the user interacts precisely, an offset between the input device absolute pointing position and the cursor position is introduced. This offset is never likely to be large, since closed feedback loop corrections are performed within close range. Even so, as to not accumulate displacement between the input device and the cursor position, once the input device moves above threshold of $5m/s$ again, there is a period of "catching up" of the cursor to the input device. This can happen in one of two situations, both vertically and horizontally. If the input device is moving towards the cursor in one or both dimensions, the cursor does not move at that dimension until the input device reaches the absolute coordinate of the cursor, when the mapping goes back to 1:1. If the input device is moving away from the target, the cursor will move fast in a lower than 1:1 CD ratio as to catch up to the input device in 0.5 seconds. Thus, the catch up phase is time, and not speed, dependent, causing the cursor to move a different speeds depending on the distance to the target and the time elapsed. In practice, from observations with experimental participants and colleagues, we did not notice any annoyance with the catch up phase, either when the input device catches up to the cursor, or vice-versa.

Physical Navigation

DyCoDiR offers the ability for the user to be at a distance, as long as there is enough visual acuity to complete the task. The user is also afforded to freely walk to complete a pointing task, and the technique accounts for the distance of the user in determining the amount of high-precision gain. As the user moves closer to the display, particularly, to the cursor position, less artificial gain in precision is required, since more precision is provided due to the smaller pointing vector intersection with the display. DyCoDiR task care of that by reducing the precision gain as the pointing device gets closer to the center, up to no gain in precision when the pointing device touches the cursor. This provides natural interaction as the user physically navigates, and there is no excessive precision when it is not required, as the user chooses to have more precision by walking closer to the object.

Zoom Variation

As with ARM, DyCoDiR also provides a version that addresses visual acuity limitations. In this case, the zoom window pops around the cursor every time CD ratio is greater than a pre-set threshold. After the threshold is passed a timeout timer of $0.5s$ is reached, the zoom window is displayed with a fade-in animation that takes $0.1s$. The user can then perceive the objects better, and, once the cursor moves faster again and the CD ratio goes below the threshold, the zoom window disappears immediately after the same time as it takes for it to appear. The reason for the timeout timer is that rapid changes in velocity of the cursor would cause the zoom window appear and disappear too many times, resulting in confusion for the user.

With the zoom variation, DyCoDiR also affords users to remain at a distance with any task, as the gain in visual acuity enables users to correctly perceive any task.

4.7 Selection by Progressive Refinement

Selection, which involves the specification of one or more objects by the user, is one of the fundamental tasks in 3D environments (Bowman et al., 2004). Although various metaphors for selection of single objects have been developed, such as virtual-hand (Stoakley et al., 1995) and image-plane techniques (Pierce et al., 1997), basic distal pointing – or ray-casting (Mine, 1995), as it's referred to in virtual environments – is perhaps the most popular selection style in VEs due to its simplicity and generality. Basic distal pointing requires only two degrees of freedom and works at any distance, while virtual hand techniques require at least three degrees of freedom and are often limited to a certain distance from the user. However, even though basic distal pointing provides better performance than virtual hand techniques in many applications, it has limitations. When the visual size of the target is small, due to the object size, occlusion, or distance from the user, basic distal pointing is slow and error-prone (Steed and Parker, 2004), because it does not provide high-precision pointing at a distance.

The first IEEE 3DUI Grand Prize (Figueroa et al., 2010) provides an example of a selection task for which basic distal pointing is unsuitable. The contest environment was a virtual supermarket where, among other tasks, selection of occluded objects in a highly cluttered environment was required. To simply use basic distal pointing for this task would require the user to select partially occluded targets very precisely from a distance, to remove occluding objects to increase the visual size of the targets, or to spend time traveling close to the targets to make them easier to select. Any of these options would have a high error probability and/or take a long time to achieve.

A number of techniques have been proposed to deal with the precision limitations of basic distal pointing. Examples include snapping (Haan et al., 2005), cone-casting (Liang and Green, 1993), 3D bubble cursor (Vanacken et al., 2007) and PRISM-style pointing (Frees et al., 2007). However, some of these techniques do not perform well in highly cluttered environments. They require users to interact very carefully to accomplish a single precise selection, and may actually provide worse performance than basic distal pointing in some situations. There are several tasks that could benefit from techniques that allow accurate selection in cluttered environments without requiring users to be precise. Examples include tasks that involve interaction with very large data sets, such as astrophysical or atomic datasets, in addition to the supermarket task already mentioned.

4.7.1 Concept

To address this challenge, we propose a selection method that uses progressive refinement of the set of selectable objects. The main idea is to gradually reduce the set of selectable objects until the target is the only one left without requiring the user to be precise at any point during the refinement. This is in contrast to traditional distal pointing techniques that use immediate selection, requiring precision. We can see an inherent tradeoff between these two categories of techniques. Progressive refinement requires a process of selection, often using multiple steps, although each step can be very fast and accurate. Immediate techniques, on the other hand, involve a single high-precision spatial selection at the expense of being slow and having a higher error probability. The goal of the work presented in this chapter is to explore this tradeoff to determine whether progressive refinement techniques are useful. In other words, we want to know when it makes sense to sacrifice the simplicity of immediate selection in order to improve speed and accuracy.

As an example of selection by progressive refinement, we designed the Sphere-casting refined by QUAD-menu (SQUAD) selection technique, which uses two distinct refinement phases. In the first phase, the user specifies a volume containing the target object. The user then refines the initial selection progressively by selecting the subset of objects containing the target from a four-item menu displaying all the remaining objects, until the target is finally selected. SQUAD makes it possible to accomplish precise selection without requiring the user to use precise actions at any moment during the selection task.

We hypothesize that SQUAD and other progressive refinement selection techniques have nearly perfect accuracy. This is because these techniques can successfully select any object, no matter how cluttered the scene, or how small the object, by allowing users to make refinements in an imprecise, careless manner. We also believe that selection techniques based on progressive refinement can be faster than immediate techniques in cases where targets are small, as long as the number of required refinements is not too high. To test these hypotheses, we compared SQUAD to basic distal pointing analytically, using the predictive model of distal pointing (chapter 3; Kopper et al., 2010)and a novel progressive refinement selection model, and empirically, through a controlled user study. These analyses can be found in chapter 6.

The concept of selection by progressive refinement is to gradually reduce the set of selectable objects until only the target remains. We identify three dimensions to the design space, which are shown in Figure 4.4.

- Type of progression
  - discrete
  - continuous
- Refinement criteria
  - spatial
  - by object attributes
  - out-of-context
- Display of selectable objects
  - in context
  - out-of-context

Figure 4.4 Design space of selection by progressive refinement.

First, progressive refinement can be done either through several discrete steps, as in SQUAD, or with a continuous process, as in shadow cone-casting (Steed and Parker, 2005).

The method of refinement defines another dimension of the design space. This refers to the criteria that are used to reduce the set of selectable objects. Refinement can be specified spatially within the environment context, for example through the use of a volume or area in the image plane, limiting the region of the environment where the target can be. Refinement can also be by the specification of attributes of the desired object, such as color, size or shape. Refinement can also be done through "out-of-context" subset specification, which involves picking a subset of objects from a list or menu instead of from the environment. The quad-menu refinement in

SQUAD is an example.

The design space is further defined by the method used to display the current set of selectable objects. Subsets of selectable objects can be displayed in context, for example through zooming, visual explosion, highlighting, moving the viewpoint closer to the subset or through the removal or dimming of non-selectable objects. The subset of selectable objects can also be displayed out-of-context, through the use of menus, which may be sorted in some way or arranged randomly.

We can also characterize progressive refinement selection techniques along a continuum based on the gradualness of refinement. At one end of the spectrum we have the immediate techniques, which directly specify the target object. This can be thought of as a "refinement" from the entire set of selectable objects in the environment to one or zero (in case of a failed selection) in a single step. At this end of the continuum, too much precision may be required, as an exact element needs to be specified immediately. At the other end of the continuum we can imagine a technique that has many refinement steps, with an extreme case being a technique where each refinement simply excludes one object from the set of selectable objects. Here precision is also required, and in fact such a technique requires many high-precision selections. In the middle of the continuum are the techniques of interest, where the reduction in the set of selectable objects is rapid and accurate.

4.7.2 SQUAD Selection

We designed a progressive refinement selection technique that falls in the middle of the gradualness continuum. It uses discrete progression, a combination of spatial and out-of-context refinement methods, and a combination of in-context and out-of-context display of the subset of selectable objects. We call this technique Sphere-casting refined by quad-menu (SQUAD).

We designed SQUAD as part of our entry to the 3DUI Grand Prize contest, described in Figueroa et al. (2010). The main challenge proposed by the contest was to design techniques that support interaction in a highly cluttered environment. In a virtual supermarket, users had to select specific objects identified by textures with unique characteristics. To achieve rapid yet precise selection, we designed SQUAD as a progressive refinement technique that divides selection into two discrete steps, the first being spatial and in-context and the second being out-of-context.

The first step uses a modified version of basic distal pointing that casts a sphere onto the nearest intersecting surface to determine which objects will be selectable. We call this subtask sphere-casting. The user simply has to ensure that the desired object is inside or touching the sphere, so that it can be picked from among the other objects in the next phase. Items that will be made selectable are highlighted. In order to improve confidence that the desired object will be available, the sphere's radius increases the farther the user is from the nearest intersecting surface, thus increasing the overall number of objects available in the second phase. Figure 4.5 illustrates this selection phase. (Note that in the study described in chapter 6, however, the sphere size is fixed since all objects are placed at the same distance from the user.) Sphere-casting avoids the precision issues of basic distal pointing, and also allows selection of occluded objects.

Upon completion of the first phase, all objects that were inside or touching the sphere are

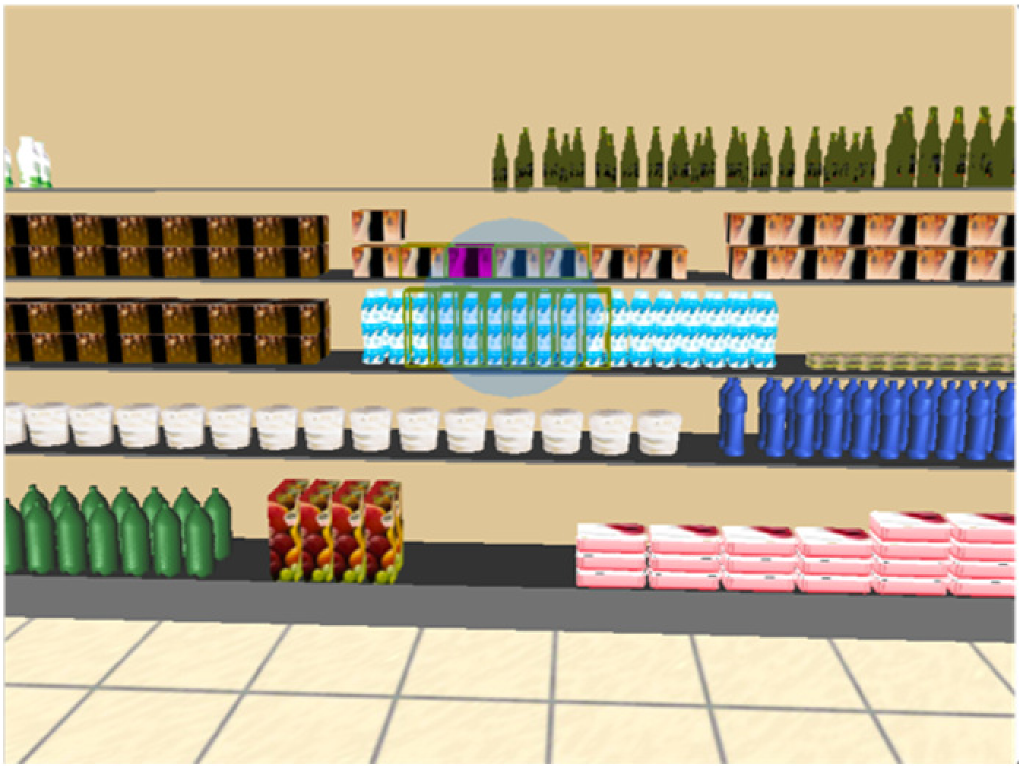

Figure 4.5 Sphere-casting.

evenly distributed among four quadrants on the screen, without regard for the spatial locations of the objects in the 3D environment. We call this the quad-menu, and note its similarity to marking menus. Contrary to zone menus (Zhao et al., 2006), where breadth of selection is achieved by relative position of multiple marking gestures, in the quad menu phase users refine the selection by repeatedly pointing anywhere in the quadrant that contains the item they are looking for, each time reducing the number of objects per quadrant until the desired object is the only one left. This process is illustrated in Figure 4.6. The maximum number of selections necessary in the quad menu is $\lceil \log_4 n \rceil$, where n is the initial number of items. For example, if the sphere has between 17 and 64 objects inside it, our technique would require at most four clicks to select the target (one click for sphere-casting and three clicks for the quad-menu).

SQUAD is an example of a progressive refinement technique that works well in environments where there are many objects that are arranged along a surface, and where the desired object is visually distinct from the rest. For other selection tasks or environments, however, different design choices (e.g., using a cone as the selection volume or distributing items in the menu based on spatial location) might be preferred.

From a distal pointing model (chapter 3) standpoint, each progression of SQUAD is a very effective technique since it removes the target size of the equation, and only rough directional movements are needed in order to complete a selection.

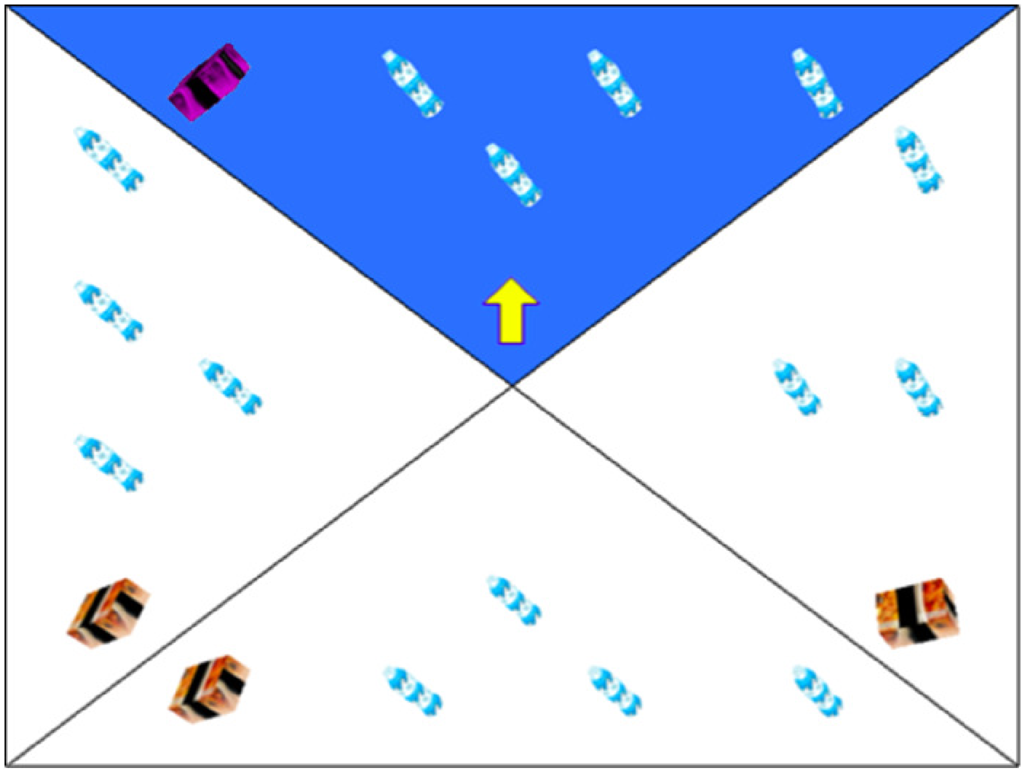

Figure 4.6 Quad-menu. Note that the target object needs to be visually distinct for the selection to be feasible.

4.8 Discussion

ARM and DyCoDiR were designed for distal pointing with 2D data in mind, as opposed to SQUAD, which was designed with 3D virtual environments as the motivation application. We believe that both the immediate techniques, as well as the progressive refinement one can be applied to 2D and 3D environments in mind.

The main challenge with designing an immediate high-precision distal pointing technique for interaction with 3D objects is that there is not a display surface, but the objects are spread in the tridimensional space. However, better feedback can be provided, by utilizing a visible ray coming from the pointing device. ARM can be directly applied to 3D environments, as its implementation is independent of the user's relative position to the interaction region. DyCoDiR, on the other hand, would need some adjustments to reach its full potential for interaction in 3D worlds. The technique takes into account the distance of the user to the cursor to determine the scale factor that impacts the CD ratio. In virtual environments, however, any point across the ray cast from the input device can select an object, and some way to determine at what depth the object is should be sought. In addition, existing ray-based techniques for virtual environments, such as cone-casting (Liang and Green, 1993) and HOMER (Bowman and Hodges, 1997) could benefit from the high-precision improvements provided by ARM and DyCoDiR.

Progressive refinement techniques, and SQUAD in particular, can be adapted for interaction with 2D data with little adjustments. Instead of a sphere cast onto the environment, a circle, or another 2D shape should be used for the first refinement phase, and refinements through the quad-menu are already done in 2D. We believe that there are 2D applications that could make use of selection by progressive refinement. For example, progressive refinement could be used to organize large amounts of data, such as documents, images and multimedia files. If there is too many clutter in the screen, as is the case with, for example, visual analytics applications (Andrews et al., 2010), progressive refinement may be used to both identify occluded documents as well as to select and sort them.

4.9 Summary

The goal of this chapter was to partially address research question 3 (section 1.2), which asks about the design of effective techniques based on the guidelines gathered from the model. Here, we presented the design of high-precision distal pointing techniques that follow the model guidelines, and chapters 5 and 6 show that these techniques are indeed effective.

This chapter focused on putting the design guidelines from the distal pointing model from chapter 3 in practice. We presented the design of three immediate distal pointing interaction techniques – ZELDA, ARM and DyCoDiR. Each technique aimed at addressing distal pointing inherent lack of precision, and all are justifiable to be highly effective.

Early experiments showed that ZELDA, which uses a bimanual style with a zoom window to provide increased precision and to address visual acuity limitations was shown to be too complex to be studied further. The idea of using increased visual size to improve performance, however, was carried on to the other two high-precision immediate distal pointing techniques that we designed.

ARM uses an explicit high-precision toggle and gives the user the choice to be precise when judged necessary. DyCoDiR, on the other hand, attempts to increase precision implicitly, as the user attempts to make refined movements. Both techniques offer a variation in which zooming is provided in order for the user to have a better perception of tasks for which the techniques already offer high precision.

This chapter also presented the design of a new distal pointing interaction paradigm – Selection by Progressive Refinement. After discussing the concept of selection by progressive refinement and presenting a design space which contains progressive refinement techniques, SQUAD is presented as an example of a distal pointing selection technique by progressive refinement. As with the immediate techniques, SQUAD follows the model guidelines but to a new level. SQUAD proposed to eliminate the target size from the task difficulty, at the expense of multiple steps to complete a task.

Chapter 5

Evaluation of High-Precision Immediate Distal Pointing Techniques

We conducted an evaluation to verify the benefits of the high-precision improvements provided by ARM and DyCoDiR. The idea of this experiment was to provide a realistic task setting in which very high precision was required. We wanted to ensure that very high precision was needed to emphasize the benefits of the enhancements provided by the techniques. However, all techniques were feasible with basic distal pointing if the user walked up to the target location at the screen.

A secondary goal of the study was to verify our previous findings that users tend to not want to walk if they don't have to. Indeed, we found that some participants did not walk when it wasn't necessary to increase precision. In fact, some participants did not ever walked, and were very unsuccessful and frustrated in the basic distal pointing trials.

5.1 Goals and Hypotheses

The main goal of this study was to evaluate the tradeoffs presented by the immediate high-precision distal pointing techniques we developed as compared to basic distal pointing. We wanted to understand performance, in terms of task completion time and errors, and user strategies employed with each technique that emphasized a particular tradeoff.

5.1.1 Tradeoffs

Table 5.1 shows the tradeoffs we evaluated with this experiment. The first tradeoff we intended to evaluate was that of high-precision activation mode. ARM uses an explicit high-precision activation, through a button in the non-dominant hand, while DyCoDiR uses an implicit high precision that is dynamically calculated based on the input device rotation speed and the distance to the cursor. This also evaluates the tradeoff between a constant versus a dynamic gain in precision.

Another tradeoff we wanted to look at was that of visual acuity. Varying visual size emphasized the tradeoff of providing zoom to address visual acuity limitations at the expense of a loss of context compared to the original view. The last tradeoff we wanted to evaluate was that of how high-precision could be achieved. It could either be with requiring the user to walk to achieve high precision, or enabling it through the indirect mappings provided by ARM and DyCoDiR.

Table 5.1 Tradeoffs evaluated in the experiment

High-precision activation	Visual Acuity	Achievement of High Precision
Implicit (dynamic)	Zoomed-In View	Physical Navigation Required
Explicit (constant)	Original View	Possible Through Indirect Mappings

5.1.2 Research Questions

With this experiment, we wanted to address the research questions 3, 4 and 5 (section 1.2), reproduced below.

3. **Can we design effective high-precision distal pointing interaction techniques based on the guidelines created in response to question 2?**

   We addressed this question by evaluating the ARM and DyCoDiR techniques that were designed based on the distal pointing model guidelines.

   Based on the tradeoffs we wanted to evaluate with this experiment, we propose two sub questions related to techniques' performance.

   (a) **Will an implicit dynamic high-precision mode lead to better performance of distal pointing tasks that require high levels of precision?**
   With this question, we wanted to test the hypothesis that taking the strategy from the user in terms of high-precision activation leads to increased performance. Specifically, we hypothesized that DyCoDiR would outperform ARM.

   (b) **Will the provision of increased visual size through zooming lead to better performance of distal pointing tasks?**
   We hypothesized that increased visual size would improve performance of the high-precision techniques, but it would hurt performance for basic distal pointing, since the perception of the task would be of a CD ratio lower than 1:1, and the zoom would only cause confusion.

4. **Will the analytic evaluation of distal pointing motor behaviors, using the model developed in response to question 1, reflect actual user performance when completing tasks with interaction techniques based on distal pointing?**

   The study presented in this chapter consists of a realistic task setting with a high level of ecological validity at the expense of control. Without a highly controlled environment, we believe it would be difficult to perform an analytical evaluation of the model in the most objective way, by plugging in values for ID_{DP} and coming out with a predicted time.

   We hypothesize, however, that we can provide validation for the model by performing an analysis based on the relative ID_{DP}'s afforded by each technique. This way, we would still be able to perform analytic evaluation in light of a wide range of strategy possibilities that were present in the study.

5. **How do the high-precision techniques designed in response to question 3 affect user strategy and performance when used in a realistic task setting?**

   We expected that ARM and DyCoDiR would lead to less physical navigation than basic distal pointing, since they allow users to be precise from a distance. Further, based on our previous observations, we expected that some users would not walk towards the display in any of the trials, even with basic distal pointing.

5.2 Techniques

In order to address the research questions posed above, we designed an experiment to evaluate the tradeoffs present in distal pointing interaction. The interaction techniques we used in the study were ARM, DyCoDiR and Basic Distal Pointing, whose designs were detailed in chapter 4. We evaluated each technique in its pure form and by providing increased visual size through zooming. Thus, there were six interaction techniques evaluated in the experiment, as shown in table 5.2.

Table 5.2 Techniques evaluated in the experiment

ARM	DyCoDiR	Basic Distal Pointing
ARM with Zoom	DyCoDiR with zoom	Basic Distal Pointing with Zoom

We expected that zoom would have a positive or null effect with the high-precision techniques, but that it would have a negative effect with Basic Distal Pointing. We believe that increased visual size only aggregates to a task that is feasible in terms of motor precision. For the tasks that we evaluated, which were designed to have very high indexes of difficulty, we hypothesize that basic distal pointing would actually be hurt by zoom. Since precision wasn't enhanced, the perceived CD ratio was actually lower than one, making the cursor movements to appear faster and less precise. The reasons we added the basic distal pointing technique with zoom were mainly to have a complete design and to verify this hypothesis.

5.3 Tasks

The task scenario we used was that of marking regions in a map[1] using predefined objects and targets. This provided a realistic setting in which participants had a good mental model of the task goal. Figure 5.1 shows the map with all objects overlaid. Each task consisted in selecting an object at the edge of the display and dragging it completely inside the target region on the map. The tasks were presented one at a time, in such way that the only interactive elements displayed each time were the target and the object to be placed on top of it (figure 5.3).

Figure 5.1 An overview of all the tasks of the experiment. Objects are spread across the edges of the map, while targets are overlaid in the map. Note that all tasks are displayed in the figure for illustration purposes, but only one task was presented at a time.

The tasks were designed to require very high precision, and we did it so because we wanted to evaluate the techniques we designed when a lot of precision was needed. Tasks consisted in two distinct phases – selection and placement – which can be seen as two individual tasks that can be analyzed separately. In the first phase, the participant pointed at the target, which was initially orange in color, and clicked on it to initiate the selection phase of the task. Then, the user had to, as fast as possible, click the object at the edge of the display to select and drag it. The second phase of the task consisted in dragging the selected object completely inside the target region in the map.

It is important to note that objects and targets had the same shape, with the only difference that targets were a magnified version of the objects. This way we could control for the pixel precision

[1]The map used in the experiment was that of Porto Alegre, Brazil, the author's home town.

required for each task, since the effective width – in other words the margin of placement – allowed for a fixed pixel area within which the object would be completely inside the target.

We were most interested in the placement phase of the task, since that's the one which required high precision. The tasks were, then, loosely controlled with three independent factors: margin of placement, object size and object complexity. Margin of placement was varied with three levels: 4, 8 and 12 pixels, meaning that the object, when centered with the target, would have that number of pixels between each edge and the target's corresponding edge. The object size was loosely varied as small or large and the object complexity also had two levels, with simple objects containing fewer sharp edges and complex ones being more convex. It is important to mention that, apart from the margin of placement, the other two variables were loosely defined and there wasn't a specific control for them. They were chosen to warrant a wide spread of different tasks.

There were a total of 12 tasks in the performance assessment phase of the experiment and five tasks during the training phases. The parameters of the tasks performed during the performance assessment phase are shown in table 5.3.

Table 5.3 Characteristics of all tasks performed in the experiment

Task	Size	Complexity	Margin of Placement	Amplitude	Shape
1	Large	Complex	$4px$	$5557px$	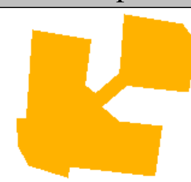
2	Small	Simple	$4px$	$1655px$	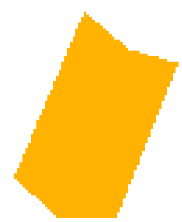
3	Small	Simple	$8px$	$6793px$	
4	Small	Complex	$8px$	$5411px$	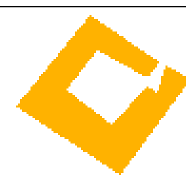

Continued on Next Page…

Table 5.3 – Continued

Task	Size	Complexity	Margin of Placement	Amplitude	Shape
5	Small	Complex	$4px$	$3198px$	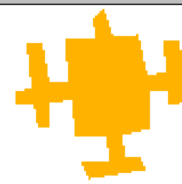
6	Large	Complex	$12px$	$7772px$	
7	Large	Complex	$8px$	$3204px$	
8	Small	Complex	$12px$	$4671px$	
9	Large	Simple	$8px$	$5919px$	
10	Large	Simple	$4px$	$1782px$	
11	Small	Simple	$12px$	$2525px$	

Continued on Next Page...

Table 5.3 – Continued

Task	Size	Complexity	Margin of Placement	Amplitude	Shape
12	Large	Simple	$12px$	$7036px$	

5.4 Apparatus

The apparatus we used was similar to the one we used to perform the distal pointing model experiment (chapter 3).

We used a flat tiled display consisting of 50 NEC MultiSync LCD2080UXi monitors in a 10×5 configuration (Figure 3.4). Each monitor's resolution was 1600×1200 pixels, resulting in a total resolution of 16000×6000 pixels (96 million pixels). Although the display we used was a grid of monitors that contain bezels, previous research has shown that they do not affect performance of target selection tasks (Bi et al., 2010). Further, we took measures to minimize the problem of fluidity of perception and interaction by positioning targets such that they did not cross any bezels.

A wireless Gear Head Wireless Optical Presenter Mouse was used as the input device (Figure 5.2). A trigger-like button was used to send mouse click events, so that movement of the device due to button clicks was minimized. To enable 6-DOF input, we attached reflective markers to the wireless mouse, which were tracked by an Optitrak system with eleven cameras. The cameras were distributed evenly in a rough circle of 4.5m diameter, in such a way that tracking accuracy was consistent throughout the whole region in front of the display. The tracking system accuracy was on the order of one millimeter, with the input device's angular granularity being considerably smaller than human natural hand jitter. A dynamic recursive low pass filter (Vogel and Balakrishnan, 2005) was applied to the raw position data read from the tracker for all techniques and before any high-precision enhancement was provided. That visibly reduced jitter without compromising the response time. We also determined the 3D position of the display surface, enabling us to intersect a virtual ray emanating from the front of the input device with the plane of the display, and thus determining the location of the cursor. We implemented the application using the OpenScenegraph (www.openscenegraph.org) library.

The application started with a blue screen with the task objects but no map or cursor. That was done so that participants didn't have to search for the target which could not be so easily found on top of the map. To display the map and activate the cursor, the participant was instructed to point diagonally down such that there was no intersection point with the screen and press the trigger

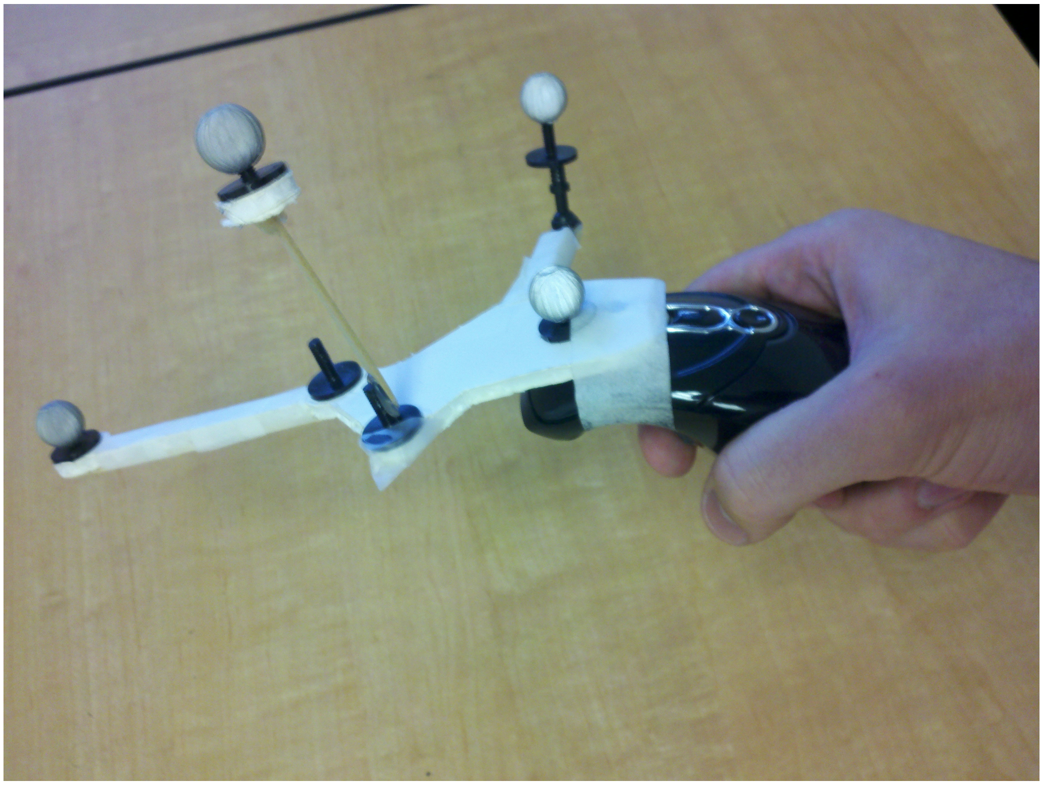

Figure 5.2 Input device used in the experiment.

button. After that, the map was displayed and the object was on the edge of the display (Figure 5.3), and the task was ready to start.

During the trials, users also wore a head band which contained markers, so that they're position could be tracked. The positions of the user' head, the input device, the cursor as well as the cursor speed were recorded for use in the analysis of the results.

5.5 Participants

Twenty-three voluntary and unpaid participants (10 female) performed the experiment. This was one participant short of a full counterbalanced design, due to time constraints. That means that all interaction technique vs. zoom variation permutations were performed by all participants twice, except the order ARM-DyCoDyR-basic distal pointing with the with-zoom variation done first, which was done only once.

All but one participant were undergraduate students with ages ranging from 18 to 22. Participants majored in different disciplines, with computer science and psychology being the most frequent ones. One participant was a 37 year old computer science graduate student.

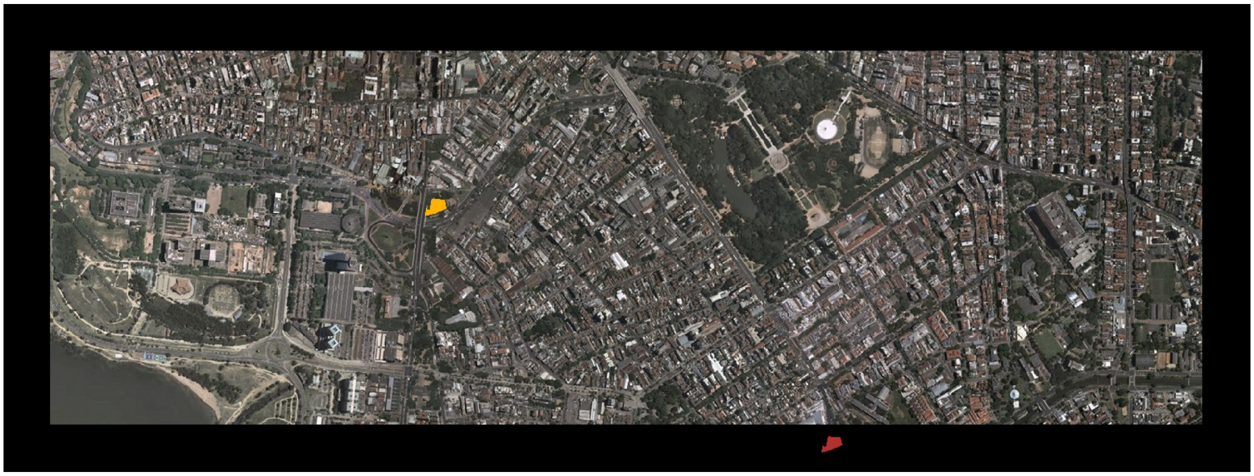

Figure 5.3 The map environment with one target overlaid and one object in its edge representing a task in the experiment.

In order to stimulate the participants to perform their best, a cash prize was offered to the participant with the best performance in each technique, and to the participant with the best overall performance. One participant could not receive two prizes. In order not to bias recruitment, participants were only told about the prize money once they arrived at the laboratory to perform the experiment, and no mention of it was made in the recruiting letters.

5.6 Procedure

Participants arrived at the experiment location and filled out an informed consent form. After that, they underwent a color blindness screening test and followed on to fill a background questionnaire. Then they were given general instructions about the study and were told about the cash prizes. They were introduced to the input device and were asked to wear the head band.

The application started with the task instruction phase. This phase was intended for participants to get familiar with the task. Basic distal pointing without zoom was used and no map was displayed. They were instructed to make all possible mistakes (e.g., double click the target or place the object off the target) until they performed the task correctly. The first task was rather easy, but the second one required more precision and participants were told that they could freely walk in order to be more precise if they judged that would help them.

After the task instruction phase, the tutorial phase for first the technique and zoom variation was done. That involved introducing the participants to the map and explaining the context of the task and also learning how to perform the technique. When zoom was present, it was explained. After learning the technique, participants practiced it with five tasks that were different from the ones which were used during the assessment phase. These tasks were presented all in the same order, from easiest to most difficult in terms of the margin of placement.

After completing the tutorial phase, a message appeared requesting the participant to wait for instructions, and they were told that they would perform two sets of 12 tasks, with a break in between. After completing the first set, a message appeared asking the participant to take a break, after which they continued on to the second set of tasks. The order of the tasks in each set was randomized.

Once the participant completed the trials with a given technique, they took a break (maximum 5 minutes) and were instructed to fill out a preference questionnaire, which asked their opinion about difficulty, enjoyment, fatigue and frustration while performing the tasks. After the second zoom variation of a given technique, participants were also asked to compare the versions of the technique with and without zoom.

After filling out the questionnaire and finishing the break, participants started on the tutorial for the next technique, and this procedure was repeated until all techniques were performed. At the end of the experiment, participants ranked all the techniques in terms of their preference.

5.6.1 Error Handling

When an error occurred, the task was invalidated and it was sent to a queue to be presented after the end of the current attempts of tasks. Since the study was designed to contain very difficult tasks, we expected that some participants would make many mistakes, especially when basic distal pointing was used. In order to decrease the level of frustration and also to keep the experiment in a manageable time – under two hours – the maximum number of attempts at each task was three, after which the task was deemed unsuccessful and not presented again for that set. There was always at most one successful task per set, meaning that participants did not have to repeat a task that was completed successfully at that set.

5.7 Design

The design of the experiment was within subjects with four independent variables: Interaction technique (ARM, DyCoDiR and basic distal pointing), zoom variation (with and without zoom), order of presentation (whether part of the first or the second set) and task. The complete design of the study was, then, $3 \times 2 \times 2 \times 12$, totaling 144 tasks performed by each participant, with a maximum of three attempts per task. We counterbalanced the order of the interaction technique and the order of zoom variation independently, such that for a given participant, the order of the zoom variation was consistent for all techniques, but varied between subjects. That gave a total of 12 permutations of interaction technique and zoom variation.

We measured the time to complete the task, the error rate and recorded the input device and head band positions to analyze physical navigation employed by the users.

5.8 Results

We performed a repeated measures analysis of variance on the data from the experiment. Since there was a maximum of three attempts per trial, many of the tasks didn't have time data. To accommodate for that, we replaced the missing values with the average time for the task in the technique and zoom variation. We made this choice for a few reasons. Repeated measures designs require data to be complete (Field, 2009), and missing data is removed listwise, meaning that all data from a participant with a single value missing is deleted. If we considered missing values to be outliers, we could replace them with the average plus two standard deviations of the condition, but we preferred to use the more conservative average, since our interpretation was that the missing values were high-error trials, but that does not necessarily equivalates to slow trial completion time.

The dependent variable time referred to the time it took for the participant to complete a task (either selection or placement) successfully. Times for unsuccessful trials were not used in the analysis, but, instead, we looked at the error rate for each trial. Error rate was calculated as the ratio of the number of attempts divided by the number of errors made. There was a maximum of three attempts per task, so the error rate of a particular trial was one of the following: 0 if the task was successful in the first attempt, .333 if one error was made, .666 if two errors were made or 1 if three errors were made, and the task wasn't completed successfully at all.

5.8.1 Selection Task

The selection phase of the task was not our main concern, but we performed an analysis of the overall results to see whether the main effects would be consistent with our hypotheses.

Figure 5.4 shows the results of overall selection time. There wasn't a main effect of task completion time for selection in light of the different techniques and zoom variations ($F_{5,18} = 1.983, p = .130$). This does not surprise us, since the selection part of the tasks did not require high precision.

However, figure 5.5 shows that there was a main effect of error rate ($F_{5,110} = 54.292, p < .0001$), even with the low precision required for the selection tasks. We believe that, even with low precision being required, basic distal pointing incurred in more errors because of the cursor being more difficult to control.

Perhaps the most important conclusion we can derive from looking at the selection tasks results is that ARM and DyCoDiR were always at least as good as basic distal pointing, which indicates that performance with the high-precision techniques isn't hurt when precision is not a requirement for the task.

Although we found these results for the selection task, we are much more interested in the placement task, and that is the focus of the rest of this section.

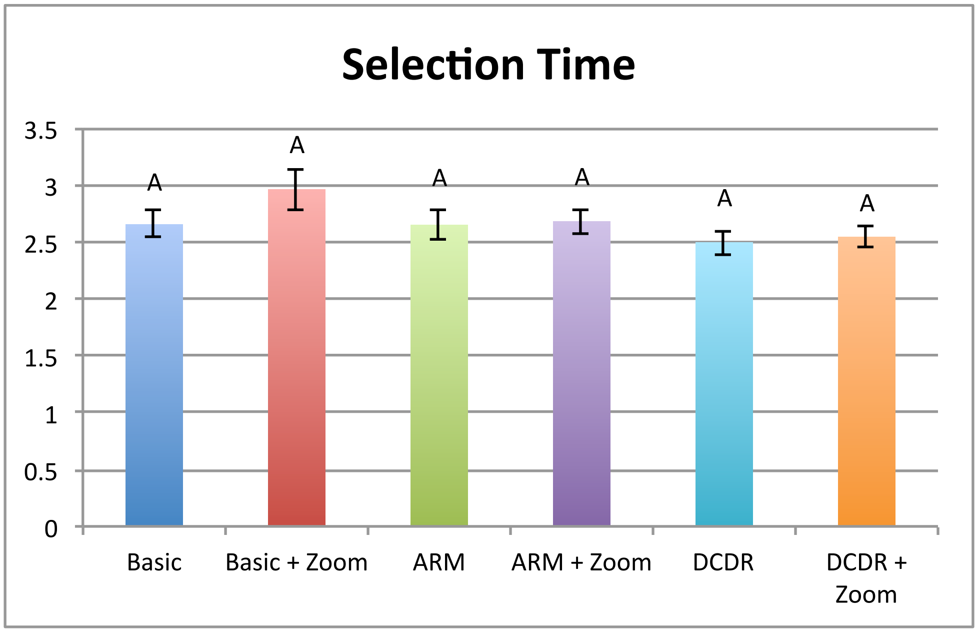

Figure 5.4 Overall result of task selection time.

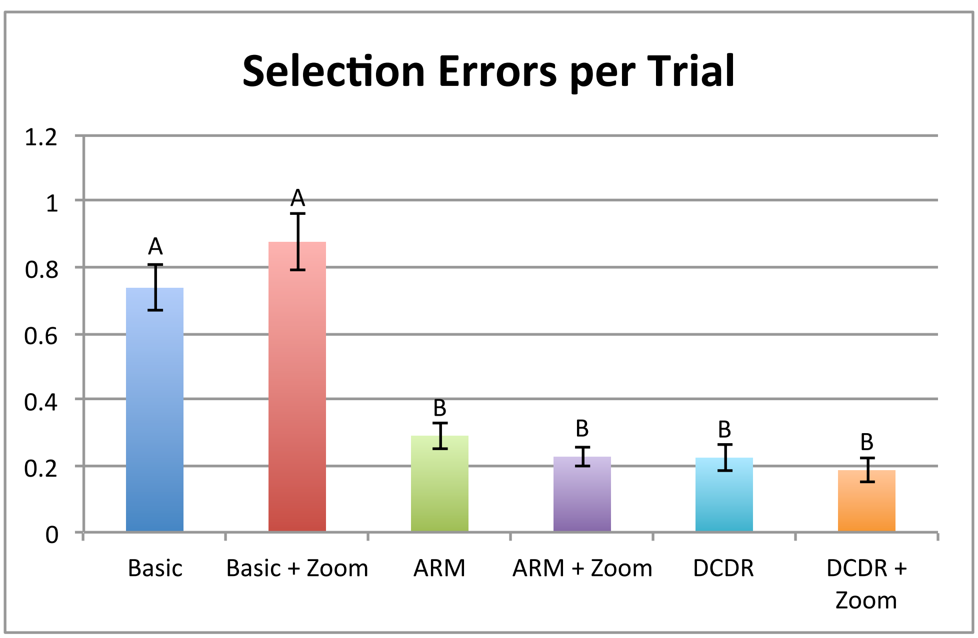

Figure 5.5 Overall result of task selection error rate.

5.8.2 High-Precision Enhancements

The main hypothesis of this research is that the high-precision enhancements provided by ARM and DyCoDiR in fact improve performance of pointing tasks. There was a main effect of technique in the time it took to complete the tasks ($F_{2,44} = 97.434, p < .0001$), as shown in figure 5.6. A post-hoc pairwise comparison using a Bonferroni adjustment for multiple comparisons show that both ARM and DyCoDiR were significantly faster than basic distal pointing ($p < .0001$), but they weren't statistically different from each other ($p = .408$).

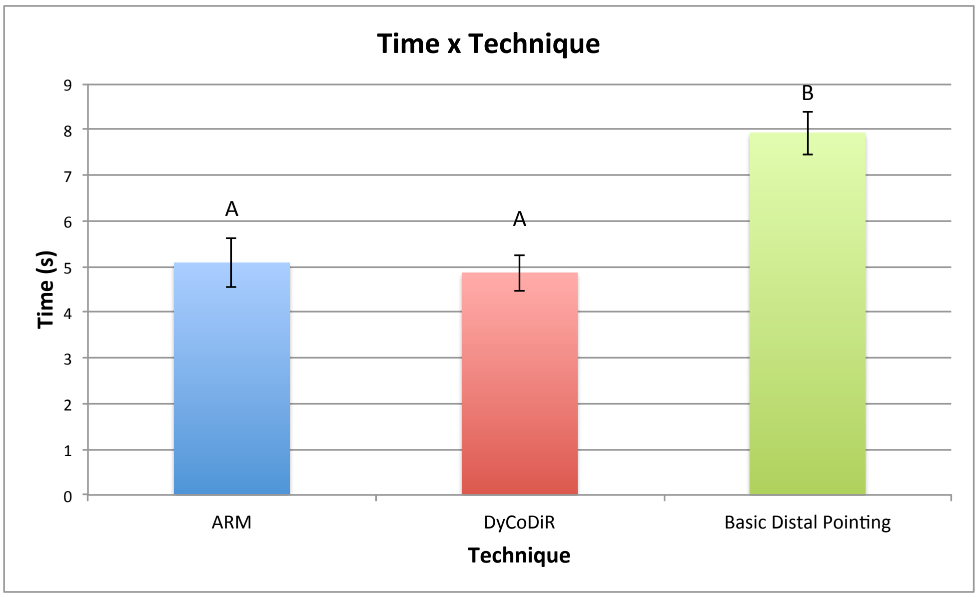

Figure 5.6 Average task completion time with the three distal pointing techniques. Note that this results combine both zoom variations. Error bars represent 95% confidence interval and levels not connected by the same letter are significantly different from each other.

Similarly, we found a main effect of technique in terms of task error rate ($F_{2,44} = 111.149, p < .0001$), which can be seen in figure 5.7. Basic distal pointing was significantly more imprecise than both ARM and DyCoDiR ($p < .0001$). Interestingly, though, DyCoDiR was marginally more precise than ARM ($p = .099$). We believe that this close-to-significant difference is due to the nature of the high-precision enhancement of DyCoDiR as compared to ARM. DyCoDiR increases precision dynamically, and there is no upper limit to the CD ratio that it provides. ARM, however, increases precision by a constant amount. This may have caused the most difficult tasks to be hard even with ARM, but not so much with DyCoDiR. It is also notable how much more imprecise is basic distal pointing compared to the high-precision techniques. The difference is much greater, than that of time, which indicates that, even if a task can be completed in a relatively short time, it will still have a much higher probability of error than ARM and DyCoDiR.

These results give us a clear indication that the high-precision-enhanced techniques were faster

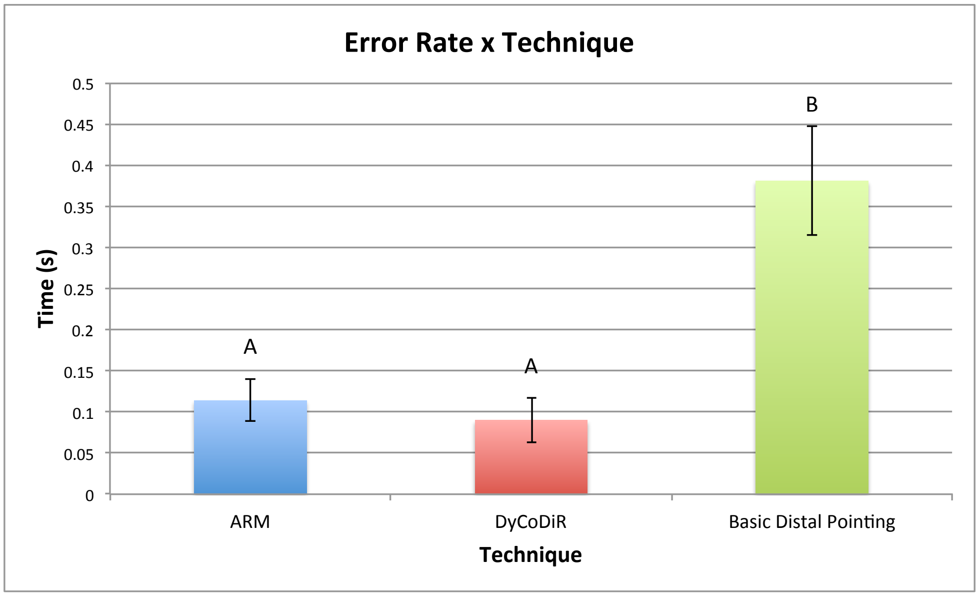

Figure 5.7 Average error rate per task with the three distal pointing techniques. Note that this results combine both zoom variations. Error bars represent 95% confidence interval and levels not connected by the same letter are significantly different from each other.

than basic distal pointing, but that does not tell us how was performance per task. That's interesting to look at, since there could be tasks that benefited more from the enhancements than other ones. For that, we performed the analysis of the interaction between technique and task, which is portrayed in figure 5.8. There was a significant interaction of technique and task with respect to time to complete the task ($F_{22,484} = 24.181, p < .0001$). Looking at the figure, however, we see that basic distal pointing was consistently and significantly slower than the high-precision techniques. Also, for most of the tasks ARM and DyCoDiR were not significantly different, but DyCoDiR was systematically faster on average than ARM. However, the only task in which there was a significant difference between the high-precision distal pointing techniques was task 10, with ARM being significantly faster than DyCoDiR ($p < .01$). We conclude that the interaction is mostly caused by the different amount of improvement of the high-precision techniques as compared to basic distal pointing. Most importantly, there was never a decrease in performance time with the high-precision techniques, which indicates that the techniques don't have a bad impact in the precision of tasks that require high precision.

There was also a significant interaction of task and technique with respect to task error rate ($F_{22,484} = 6.77, p < .0001$), as shown in figure 5.9. Again, we see a consistent trend of basic distal pointing being less precise than ARM and DyCoDiR, with the amplitude of difference varying based on task difficulty. We also note that, although there were no significant differences between ARM and DyCoDiR with respect to error rate, DyCoDiR yielded consistently less errors than ARM on average. This is congruent with the idea that DyCoDiR provides more precision as needed, whereas ARM does not vary the amount of increase in precision.

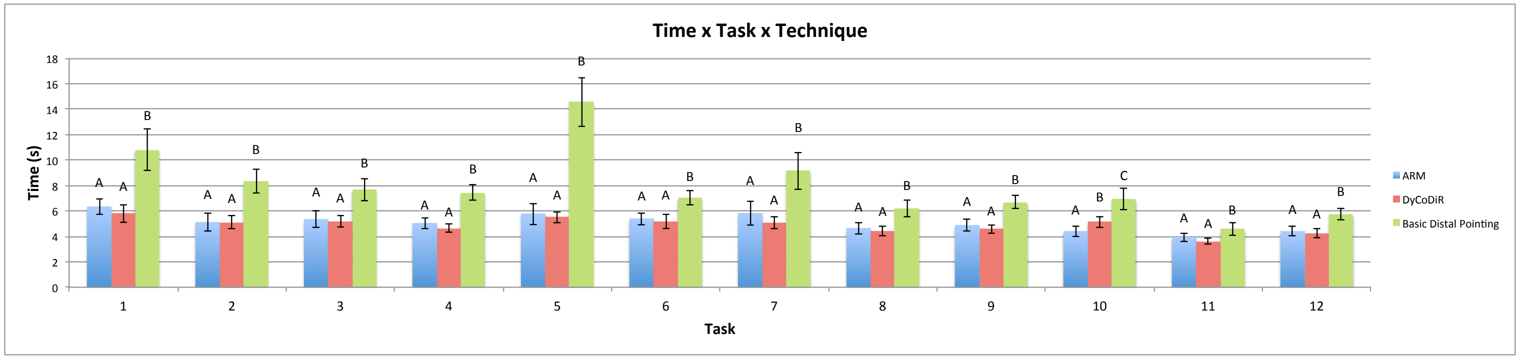

Figure 5.8 Interaction of task and distal pointing technique with respect to time. Error bars represent 95% confidence interval and levels not connected by the same letter within each task are significantly different from each other.

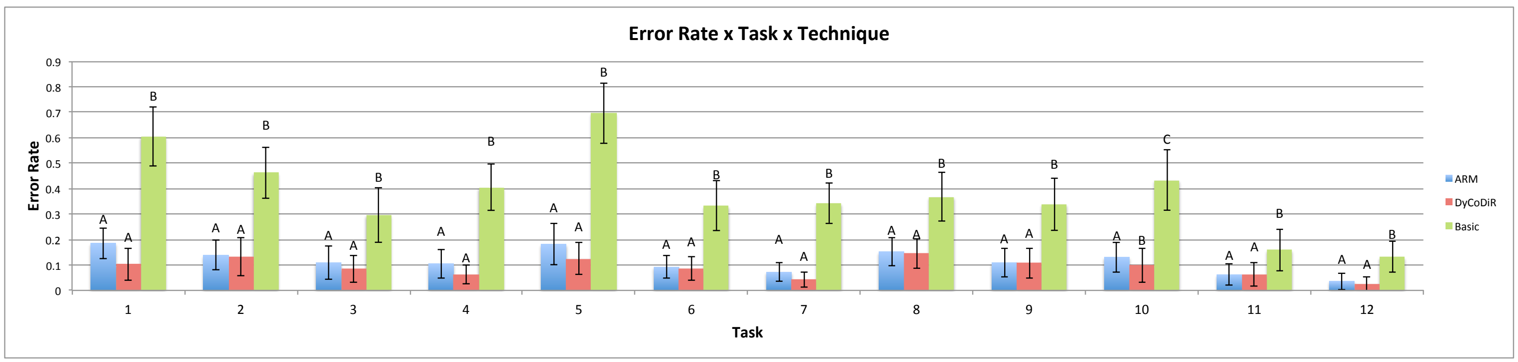

Figure 5.9 Interaction of task and distal pointing technique with respect to task error rate. Error bars represent 95% confidence interval and levels not connected by the same letter within each task are significantly different from each other.

Looking at figures 5.8 and 5.9, we can see that the error bars are often narrower with DyCoDiR than with ARM, which indicates lower variance, consistent with our hypothesis. That indicates that individual differences and strategy played more of a role in the performance of ARM than it did with DyCoDiR. That is what we expected, since DyCoDiR provided high precision implicitly as the user needed, whereas ARM high-precision mode was toggled explicitly by the user.

Zoom Performance

One question that we wanted to answer with the study was whether there was an effect of increased visual size in the performance of distal pointing tasks. We hypothesized that zooming would help the high-precision techniques, since it provides a visual aid when the user can point precisely. With basic distal pointing, however, we hypothesized a negative effect of zooming. Since there is no improvement in the task precision, and only the visual acuity is addressed, the effect it has is that of a lower-than-one CD ratio, as the cursor appears to move faster in the zoomed view in order to cover the same space that it would in the regular view.

We found a strong interaction of zoom and technique with respect to the time it took to complete the task ($F_{2,44} = 35.911, p < .0001$). Post-hoc pairwise analysis with a Bonferroni adjustment for multiple comparisons show that ARM was significantly faster with zoom ($p < .0001$), DyCoDiR did not yield a significant difference with respect to zoom ($p = .145$) and basic distal pointing was significantly slower with zoom ($p < .0001$). These results are partially consistent with our hypothesis and can be seen in figure 5.10. We indeed expected a negative effect of zoom for basic distal pointing, but we did not expect that DyCoDiR with zoom would not be significantly better than the variation without increased visual size. We believe that the lack of a main effect here is due to the fact that DyCoDiR had the zoom activated automatically, after a delay during which the CD ratio was higher then a threshold. We believe that this may have confused participants, since the zoom may have come unexpectedly, even if participants knew to expect it. With ARM, however, zoom was activated at the very moment when the user triggered the high-precision mode, so it was more predictable. A redesign of the zoom activation with DyCoDiR may allow zoom to actually improve performance. With the way that it is currently implemented, we believe that the benefits brought by increased visual size were outweighed by the confusion caused by the zoom suddenly appearing.

High-Precision Activation Strategies

The ARM technique provides an explicit high-precision toggle, which the user decides to activate when she judges it's necessary to increase precision. Ideally, the user should activate ARM only once per trial, after the ballistic movement towards the target, in order to have more precision to perform the refinements necessary to complete the task successfully. If a participant activated ARM more than once, that would cause a breakdown in the task performance, as the cursor would jump back to the absolute pointing position, which typically would be far from the target. Multiple high-precision activations generally occurred because the participant activated the mode when the

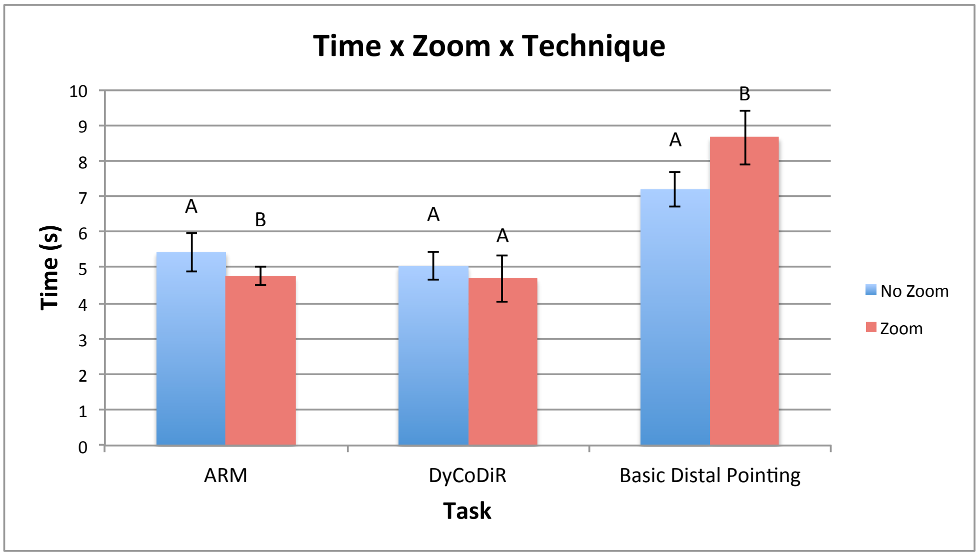

Figure 5.10 Interaction of zoom and distal pointing technique with respect to time. Error bars represent 95% confidence interval and levels not connected by the same letter within each task are significantly different from each other.

cursor was still too far from the target to provide an easy movement of the cursor towards it. Tasks in which participants did not activate the high-precision mode at all were equivalent to the basic distal pointing technique, except that the zoom variation of ARM would never show a zoomed-in view if high-precision wasn't activated.

We performed an analysis of the ARM techniques in light of the strategies participants took when completing the tasks. Table 5.4 summarizes the results.

Table 5.4 High-precision activations with ARM per drag task.

Technique	# Activations	# Successful	# Errors	Error Rate	% of Trials
ARM	Zero	24	14	36.8%	5.9%
	One	481	84	14.9%	87.9%
	> One	35	5	12.5%	6.2%
ARM+Zoom	Zero	5	13	72.2%	2.8%
	One	526	69	11.6%	93.3%
	> One	19	6	24.0%	3.9%
Total	Zero	29	27	48.2%	4.4%
	One	1007	153	13.2%	90.6%
	> One	54	11	16.9%	5.1%

Overall, we can see that ARM was generally well used, with over 90% of the trials being completed with only one activation. The success rate of ARM was also quite high when used properly, given the difficulty of the tasks. As expected, the error rate is quite high for the trials in which high precision wasn't activated, as the task was performed with basic distal pointing, without any precision enhancements.

One participant jumps to attention, in that she did not seem to have understood the ARM techniques. Of the 23 successful trials performed by participant 8, 22 were performed without activating the high-precision mode, while, the 20 failed attempts by participant 8, 14 occurred when high precision was activated once during the trial. This is interesting, as it is an opposite behavior from the other 22 participants, and further investigation should be carried out to find out if the participant is an outlier, or if that hints the existence of users that have difficulty understanding the concept behind ARM. Even more interesting, with ARM in the zoom variation, participant 8 attempted more trials activating high precision, but still had a high error rate for those tasks that high precision was used (6 out of 20 trials were failed). This may indicate that participant 8 was more encouraged to use the high-precision mode of ARM with zoom because of the clear visual feedback provided.

5.8.3 User Ratings

The high-precision techniques were largely preferred over basic distal pointing. Participants were asked to rate the techniques, on a 7-point Likert scale, about how much they liked the technique they had just performed, and what was the maximum level of frustration the technique caused to the participant. Figure 5.11 shows the average ratings of the techniques based on how much they liked using them.

We ran a Wilcoxon Signed Ranks Test on the questionnaire data, and found that basic distal pointing was rated significantly lower than the high-precision techniques ($Z < -3.5, p < 0.0001$). Also, in accordance with the hypothesis that increased visual size would actually hurt basic distal pointing, we found that basic distal pointing with zoom was rated significantly lower than the regular variation of basic distal pointing ($Z = -2.077, < 0.05$). We found no significant differences in the ratings across high-precision techniques with or without zoom.

Participants were also asked about the maximum level of frustration they felt while performing the tasks. The results are shown in figure 5.12. The results from the Wilcoxon Signed Ranks Test shows that the basic distal pointing techniques led to significantly more frustration than the high-precision techniques ($Z < -3.2, p < 0.0001$). Again, basic distal pointing with zoom led to more frustration than the variation of the techniques without zoom ($Z = -2.146, p < 0.05$).

Participants were also asked, at the end of the experiment, to rank the techniques from favorite to least preferred. Figure 5.13 shows the results. As expected, the basic distal pointing techniques we the significantly least preferred ones, and the only significant difference among the high-precision technique is that ARM with the zoom variation was preferred over regular ARM.

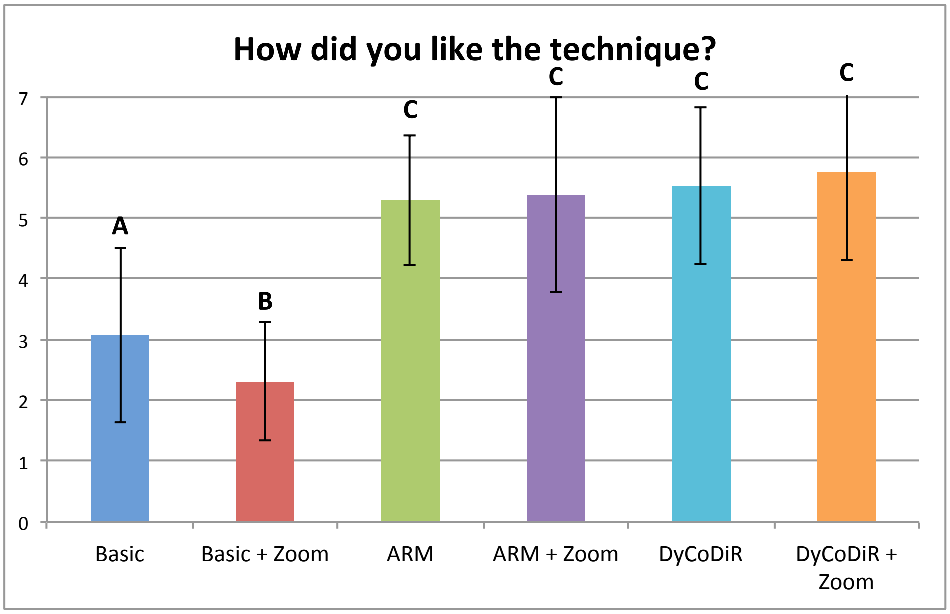

Figure 5.11 Rating of techniques based on participant impression, answering the question "How did you like the technique you just used?". Error bars represent standard deviations and levels not connected by the same letter are significantly different from each other.

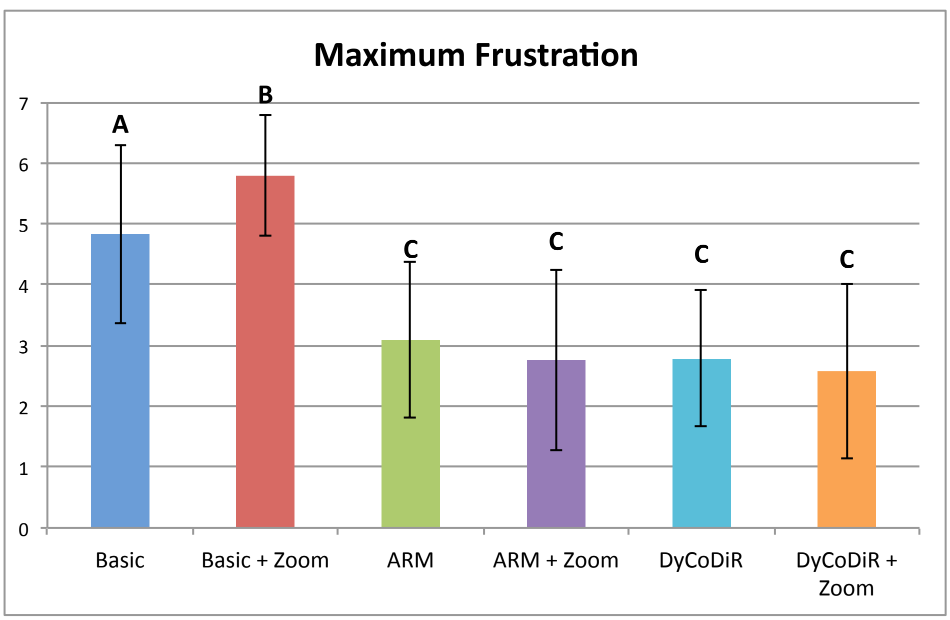

Figure 5.12 Rating of techniques based on participant impression, answering about the maximum frustration they felt while performing the tasks with each given technique. Error bars represent standard deviations and levels not connected by the same letter are significantly different from each other.

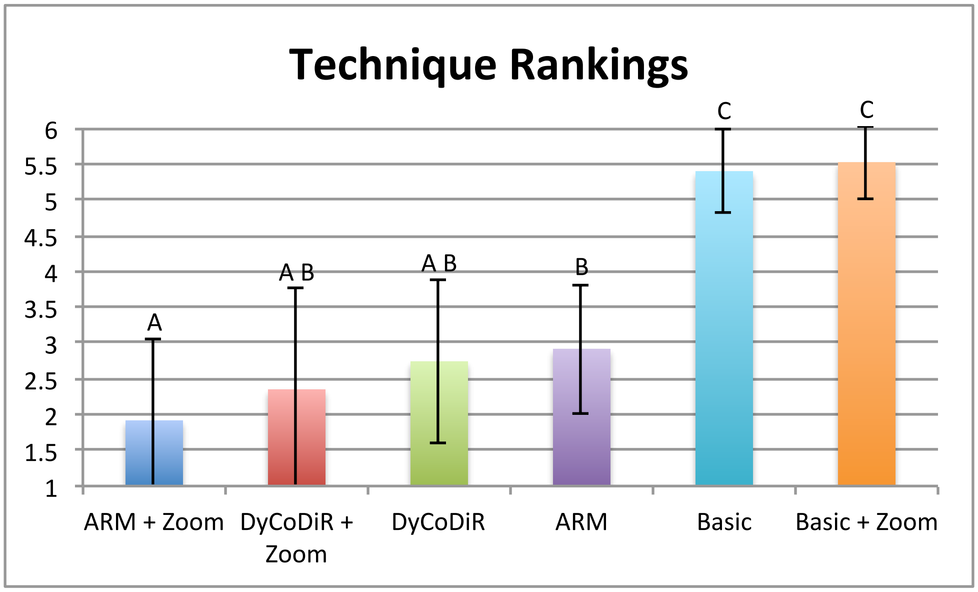

Figure 5.13 Average technique rankings. Error bars represent standard deviations and levels not connected by the same letter are significantly different from each other.

5.8.4 Physical Navigation

Looking at the data, we found that most of the time the offset between the head position and the input device position is close to constant, which means that there wasn't a lot of arm movement, but full body movements were much more common. We, thus, only analyzed physical navigation using the input device data, as head data wasn't as robust since tracking was lost sometimes due to the height of the markers on some participants.

In order to understand participants't choice to physically navigate or not, we looked at the distance they walked from the beginning of the task until its completion. This analysis enabled us to understand participants choices in light of the technique they were using as well as personal differences. We wanted to understand if the interaction technique played the biggest role in determining walking strategies, or if there was a component of personal preference that made some participants be consistent with their walking strategies throughout the study, despite the technique they used.

Encoding of Trials

We categorized each successful trial participants performed with respect to their physical navigation strategies. We observed from plots of participants' positions while performing tasks that in some cases the strategy taken was to stand still, but starting the task already near the target. In other cases, participants stood still and far away from the target. In other situations, participants walked while performing the task, and that was done either ending far away from the target, as would be the case with a lateral movement, and those cases in which participants started far away from the target and finished the task close to it, in a walking towards movement. We summarize the walking strategies in table 5.5.

Table 5.5 Categories of physical navigation for the drag task employed by participants during the experiment.

Amount of walking	Final position	Number of Trials
Stationary (total movement $< 30cm$)	Near target ($< 50cm$)	228 (7.58%)
	Far from target ($> 50cm$)	1285 (42.7%)
Walking (total movement $> 30cm$)	Near target ($< 50cm$)	853 (28.3%)
	Far from target ($> 50cm$)	643 (21.4%)

Figure 5.14 shows examples of each physical navigation strategies. We can see in the table that roughly half of the trials were performed with the participant not moving, and the other half had participants walking. Of the stationary trials, most had participants standing far from the target, while the majority of trials who had participant walking, the task ended with them near the target.

Physical Navigation per Technique

Table 5.5 shows that the physical navigation strategies employed by the participants varied widely, but is does not tell what impact each technique had on the choice to walk or stand still while completing a task.

Figure 5.15 shows a breakdown of the physical navigation strategy per technique. It is interesting to note that the high-precision techniques afforded participants to end the task at a distance from the target. Clear, also, is the fact that basic distal pointing caused participants to walk more than all the high-precision techniques. As expected, the zoomed variation of ARM and DyCoDiR caused people to remain at a distance more than the non-zoom counterparts.

Looking at each technique in detail, we can derive some interesting observations. About half of the trials with ARM were performed with participants standing still and far away from the target. Only a small amount of trials were performed with participants starting and ending the trials with ARM near the target. However, about one in every five participants started the task far away from the target and ended near it, while in the remainder quarter of trials, participants walked towards the target, but never got too close to it. What this analysis of the ARM trials in terms of physical navigation tells us is that the technique afforded participants to remain at a distance, with about 80% of trials being completed that way. However, it also did not hinder participants necessity to walk, either due to preference or to a task that would benefit from physical navigation, as 45% of the trials had participants walking while completing them.

When we contrast the results of ARM with the zoom variation, we see that the main difference is the amount of trials that were completed with participants standing still and far from the target.

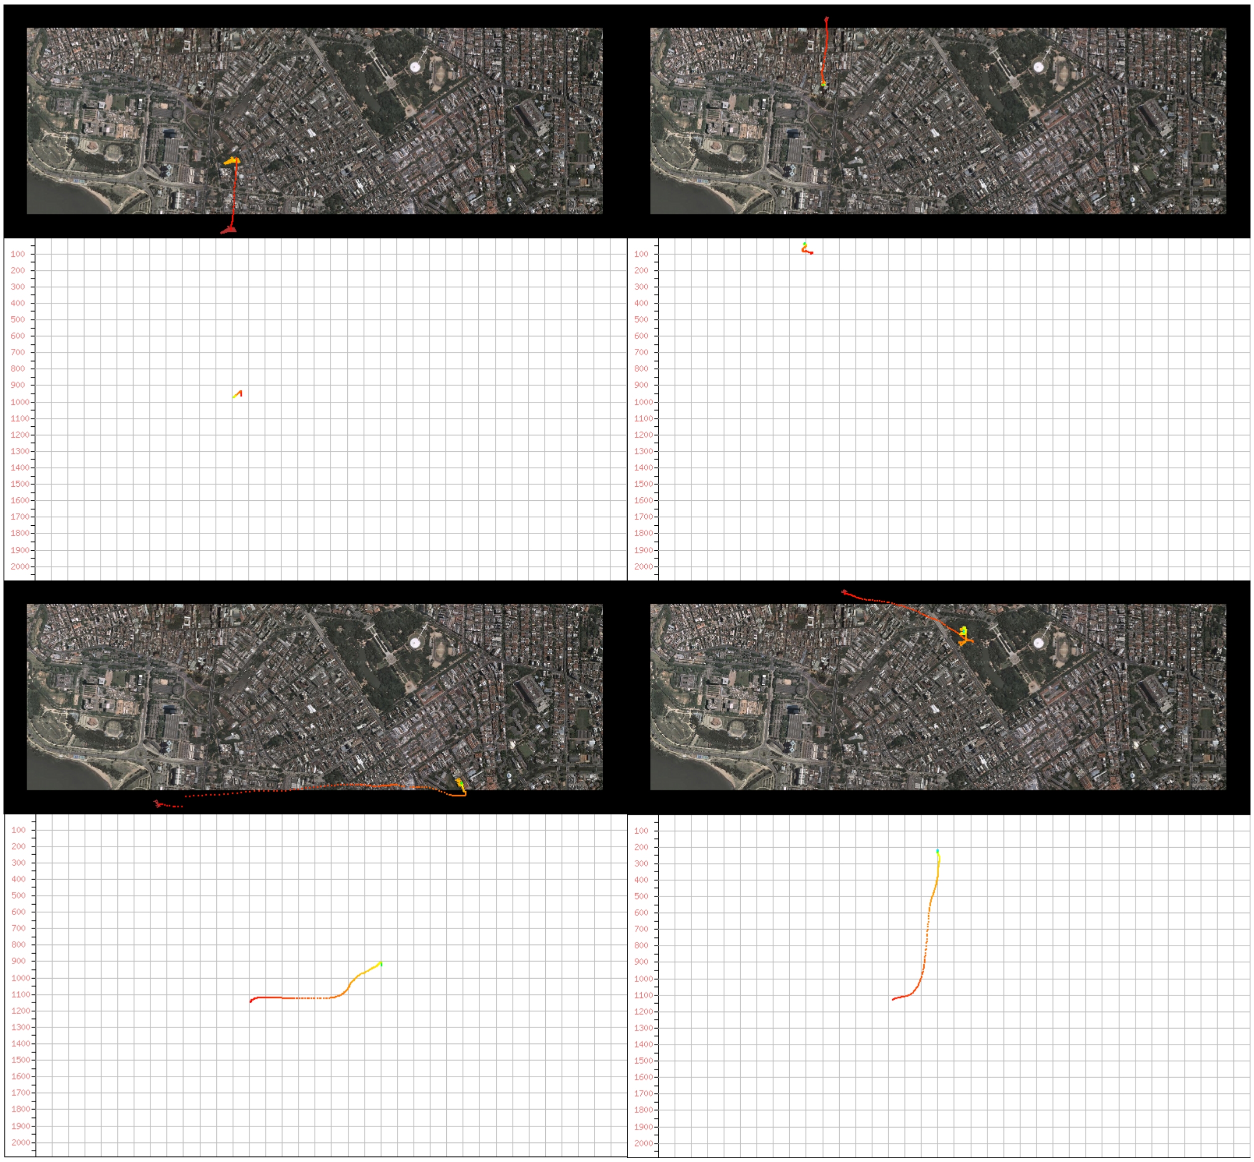

Figure 5.14 Physical navigation strategies. From red to green, the map shows the cursor path for the drag task, and the grid shows the top-down view of the tracked area in front of the display. Each grid square denotes an area of 100×100mm. Top-left: stationary far; top-right: stationary near; bottom-left: towards far; bottom right: towards near.

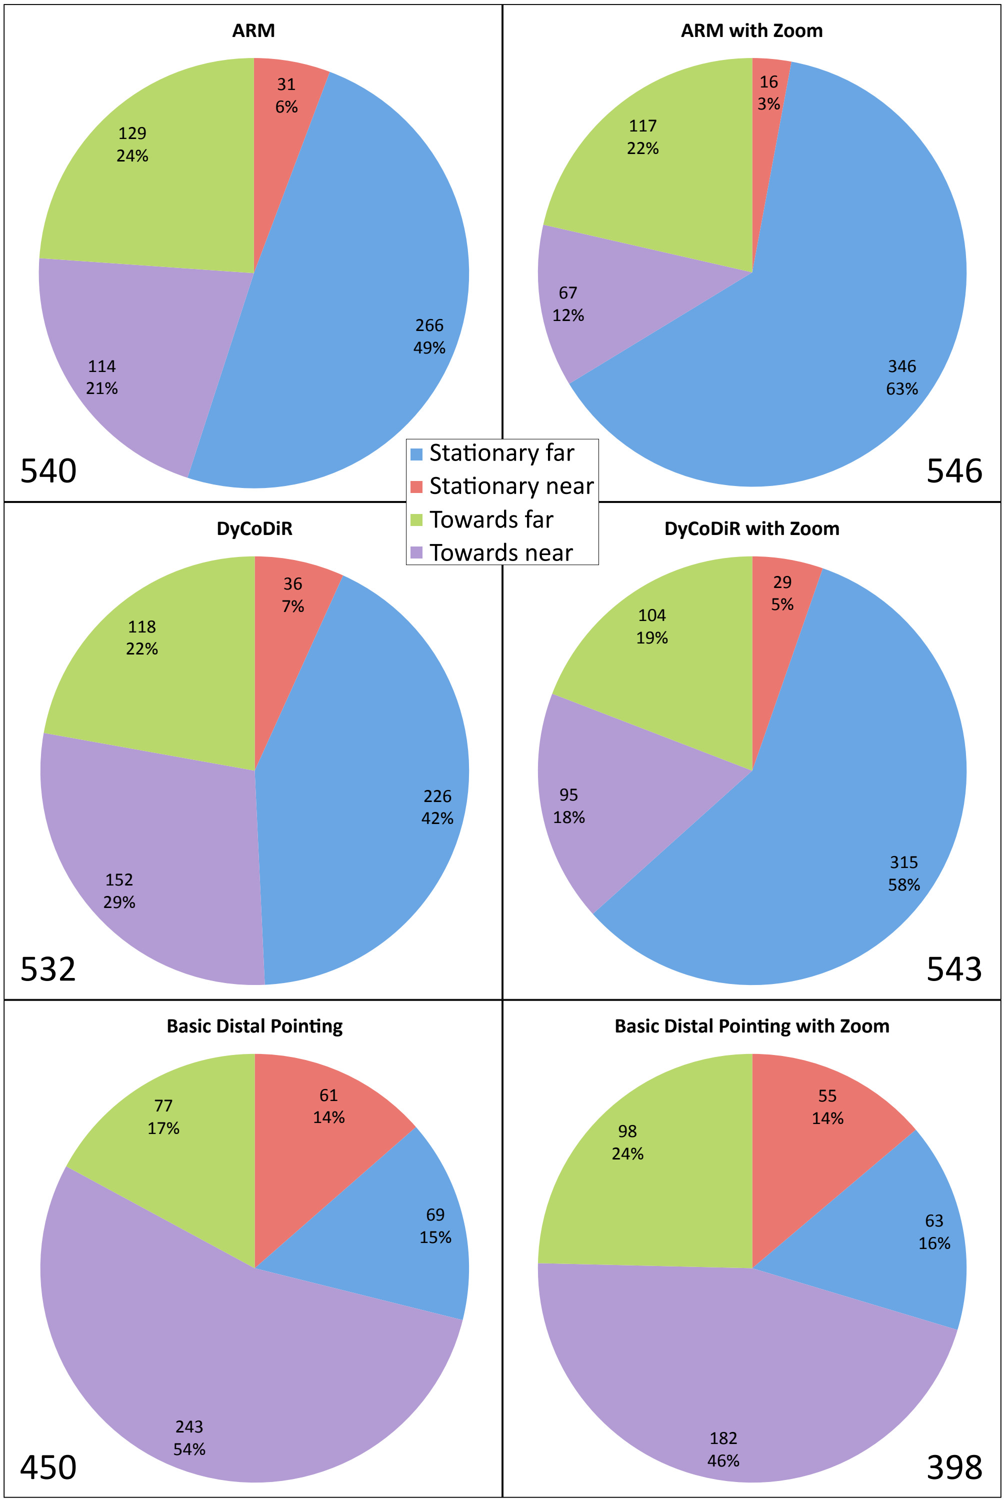

Figure 5.15 Physical navigation strategies per technique. The values outside each pie chart represent the total number of successful trials with each technique.

We believe that the advent of zooming helped participants perceive the targets better, and, thus, did not require them to walk up close to the target to be able to complete a task. In this regard, we note that the conditions that contained fewer trials with ARM when zoom was present were the ones in which participants either walked towards the target, or completed the whole drag task while close to it.

While the DyCoDiR techniques have a somewhat similar behavior in terms of the physical navigation strategies employed as ARM, there are some noteworthy differences to observe. When comparing to ARM, more trials were performed with participants walking towards the target, with fewer stationary and away trials. Since DyCoDiR offers a implicit gain in precision, we speculate that participants were less aware of the high precision they could get while being at a distance to the target, and that caused them to walk more. Again, we see a similar trend of participants walking on more trials in the zoom variation of DyCoDiR as compared to ARM with zoom, but a higher number of trials afforded participants to remain at a distance when zoom was present.

The basic distal pointing techniques present an opposite trend to that of the high-precision ones. With basic distal pointing, participants had to be near the target to succeed in the more difficult tasks, and that is expressed in the chart, in which more than 60% of the trials were performed with participants either standing near the target, or walking towards and ending the task close to it. We do see a trend of participants remaining farther from the target when zoom was present with basic distal pointing, but that also caused participants to be less precise.

Individual Differences

We know from the analysis of the physical navigation strategies per technique that the distal pointing technique directly influences the choice of walking strategy to use. However, we would also like to know whether individual differences play a role that can override the affordance to walk or not provided by the interaction technique. Figure 5.16 shows the walking strategy employed by all participants with all interaction techniques.

There are a few interesting notes to take from observing the figure. We can see that some participants chose to walk towards the display always, no matter what technique they would use. Others, on the other hand, chose to remain at a distance for all techniques that afforded that. In fact, some participants chose to never walk up to the display, and were very unsuccessful with the basic distal pointing trials. It is clear, by observing the figure, that participants varied they're walking strategy, sometimes independently of technique. We can categorize participants with respect to their walking strategies in three groups, as described below.

1. Participants who walked to the target only when they needed in order to be more precise.
2. Participants who tended to always complete the tasks while near the target, independent of technique.
3. Participants who tended to remain at a distance from the target, independent of technique.

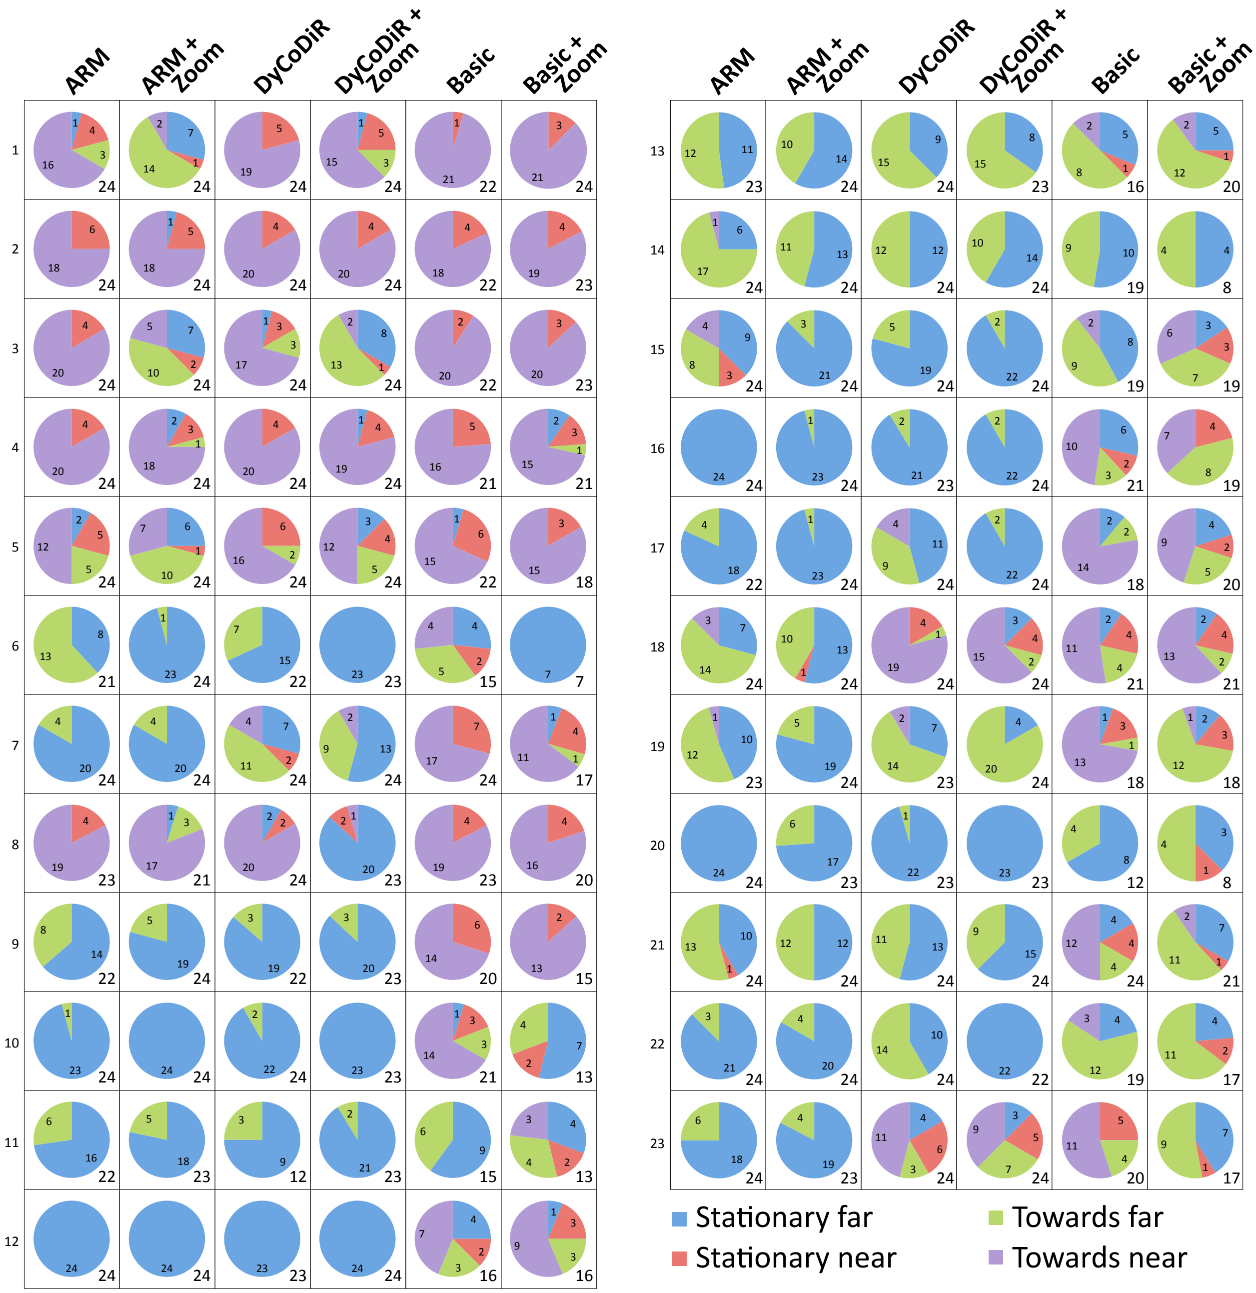

Figure 5.16 Physical navigation strategies per technique for each participant. The labels at the bottom right corner of each pie chart represent the number successful trials with each technique per participant. A maximum of 24 trials were attempted per technique.

The first group describe those users who performed the optimal strategies with each technique. They are participants numbers 7, 9, 10, 12, 16, 17, 18, 19 and 23. They usually remained at a distance if the technique provided high precision without requiring them to move close to the target, but did walk when they needed, with basic distal pointing. A few observations are in place. Some participants (e.g., 7 and 23) tended to walk more with DyCoDiR than with ARM, even though DyCoDiR can provide very high precision. We believe this is so because, since DyCoDiR implicitly improves precision, some participants did not realize that they could be precise while at a distance. With ARM, on the other hand, the high-precision mode is prompted by the user, and suddenly increases the precision, making it much more obvious. Even though most participants were able to use DyCoDiR while not having to walk up close to the display, we believe that some design improvements could be made in order for the high-precision be more explicit to the user. For example, some visual feedback could be provided indicating that high precision is being offered to the user. That said, it is not necessarily negative that participants walk while using DyCoDiR. In fact, the technique encourages absolute pointing interaction as the users moves closer to the interaction point, so that artificial gain in precision is only provided when the user needs it, when she is at a distance from the interaction.

In the second group, we see participants 1 through 5 and participant 8. These participants chose to complete their tasks most of the time by standing near the target, and being able to be precise despite the interaction technique. Interestingly, participant 8 chose to remain at a distance when using DyCoDiR with zoom, probably because of the automatic activation of zoom, which addressed visual acuity limitations without the need to walk close to the target. It is interesting to note, also, that these participants tended to underuse the high-precision mode of ARM, which caused them to walk more during the ARM trials.

The third group contains participants 6, 11, 13, 14, 15, 20 and 22. These participants seldom completed the tasks by going near the target, and that made them quite unsuccessful with the basic distal pointing techniques, as it can be seen by the number of successful trials with these techniques, compared to the 24 trials that were attempted. It is important to note that all participants were aware of their affordance to walk up close to the display and, in fact, were enforced to do so during training. These participants tell us that the high-precision enhancements provided by ARM and DyCoDiR are important not only in application settings that require users to interact while at a distance; it also shows that individual differences affect the choice of users not to walk up to the display to perform an interaction. The high-precision techniques afford these users who do not wish to interact up close to the display to still be successful.

If we tally the participants in each of the afore mentioned groups, we see that, out of the 23 participants, nine chose a strategy that minimized their walking when the technique offered them to perform the task successfully from a distance. Six participants decided to use physical navigation as their principal means to increase precision, while the remainder seven chose not to walk close to the target, even at the expense of having a very poor performance in the basic distal pointing tasks.

We believe that high-precision distal pointing techniques give users the choice to not walk, to interact in a "lazy" manner, without having to be to careful to perform precise interactions. However, the techniques, as we saw in figure 5.16, allow participants to be successful even if they choose to walk. The characteristics of both ARM and DyCoDiR allow users to take advantage from the natural precision they can have if they decided to walk towards the interaction area.

Further, ARM and DyCoDiR allow precise interaction with applications that require users to be at a distance. For example, in geospatial visualization, often analysts need to keep an overview of the data, but they may need to perform some sort of interaction with it. Without high-precision distal pointing techniques, they either need to walk up close to the display, loosing the overview, or use an indirect control to perform the interaction, which also may cause a cognitive disruption. Enabling users to interact while at-a-distance is one of the main contributions from the high-precision distal pointing techniques we designed and evaluated.

5.9 Model Validation

The study presented in this chapter consists of a realistic task setting with a high level of ecological validity at the expense of control. Without a highly controlled environment, we believe it would be difficult to perform an analytical evaluation of the model in the most objective way, by plugging in values for ID_{DP} and coming out with a predicted time.

We hypothesize, however, that we can provide validation for the model by performing an analysis based on the relative ID_{DP}'s afforded by each technique. This way, we would still be able to perform analytic evaluation in light of a wide range of strategy possibilities that were present in the study.

The lack of control from the experiment enabled us to gather important findings about user performance and strategy in natural, realistic settings. However, that led to a high variability of the data, especially in light of tasks that required very high precision, and made the analytical fitting of the experimental data to the distal pointing model described in chapter 3 impractical.

The main contribution of the distal pointing model, however, is not to predict the exact time to complete a distal pointing task based on its parameters and user strategy. The idea for the model is to enable designers and researchers to design techniques that follow the model guidelines. The guidelines, in turn, relate to the index of difficulty of the distal pointing model.

We propose, then, to perform an analytical evaluation of the motor behaviors described by the model through a comparative analysis of the effect that each technique had on ID_{DP}.

First, we shall look at the basic distal pointing technique. This technique can be seen as a baseline for ARM and DyCoDiR to be compared to. Since the experiment gave users freedom to perform a task, the ID_{DP} of a trial varied constantly based on the participant's relative position to the target.

Figure 5.17 shows an example of the progression of ID_{DP} across the performance of a task using basic distal pointing. In the figure, the graph on top shows the values of the calculated ID_{DP} based on the participant's position relative to the task elements, and the grid shows the physical navigation. Each point represents one frame of data captured at $100Hz$ and the grid is divided into $100mm$ squares. The colors show the path from start to finish.

Figures 5.18 and 5.19 show the progression of ID_{DP} for ARM and DyCoDiR, respectively. It is clear that both ARM and DyCoDiR yield lower ID_{DP} towards the end of the task, even with less physical navigation than basic distal pointing. This illustrates the features of both high-precision techniques to reduce ID_{DP} as compared to basic distal pointing.

Comparing the values of the ID_{DP}'s for each technique during the course of a task with the empirical results from figures 5.8 and 5.9, we can infer that the increase in performance provided by ARM and DyCoDiR are due to the reduced indexes of difficulty that they provide for distal pointing tasks.

We can further gather evidence to validate the model when we look specifically at ARM as compared to DyCoDiR. We can see in figure 5.19 that DyCoDiR causes ID_{DP} to be very low, close to zero towards the end of the task, while when ARM is activated, ID_{DP} keeps constant at a higher value. This happens because DyCoDiR provides no limit to the increase in CD ratio, which causes effective target sizes (ω) to be as high as the user needs, reducing dramatically the overall difficulty of the task. Thus, an analytical evaluation of DyCoDiR as compared to ARM would cause us to infer that DyCoDiR outperforms ARM in the most difficult tasks. Figures 5.8 and 5.9 shows that, although ARM and DyCoDiR are not statistically different in most cases, possibly due to high variance caused by lack of control, we see a consistent trend of DyCoDiR outperforming ARM, which provides evidence for the validation of the model.

Ultimately, the best validation of the model that we can get from this experiment is that the model can in fact be used to guide the design of techniques that are prone to be more effective than basic distal pointing.

5.10 Discussion

The experiment we performed gave us many insights into how users perform using distal pointing in light of different tradeoffs.

When we examine the tradeoff between explicit and implicit high-precision activation, we see that implicit activation may be better suited for high-precision distal pointing tasks. DyCoDiR, which uses implicit activation, carries less overhead than the explicitly activated ARM. That caused some users to not follow the best strategy of high-precision activation, perhaps activating too early, or trying to complete the task without even activating the high precision. With DyCoDiR, on the other hand, the strategy is taken out of the equation, as the technique itself provides high precision as the user needs.

We note, however, that there wasn't a main effect of zoom with DyCoDiR. We believe that, since the zoom was displayed automatically when precision was required, it may have confused some participants who weren't expecting the zoomed-in view to appear suddenly. We believe that some other form of zoom activation should be provided with DyCoDiR so that the benefits of addressing visual acuity limitations by increasing the visual size are not outweighed by the confusion caused by how it's design.

In terms of physical navigation, it is clear from the results that participants walked less with the

Figure 5.17 Progression of instant ID_{DP} across the completion of task 4 using basic distal pointing. The only factor that causes ID_{DP} to change is the user's physical navigation.

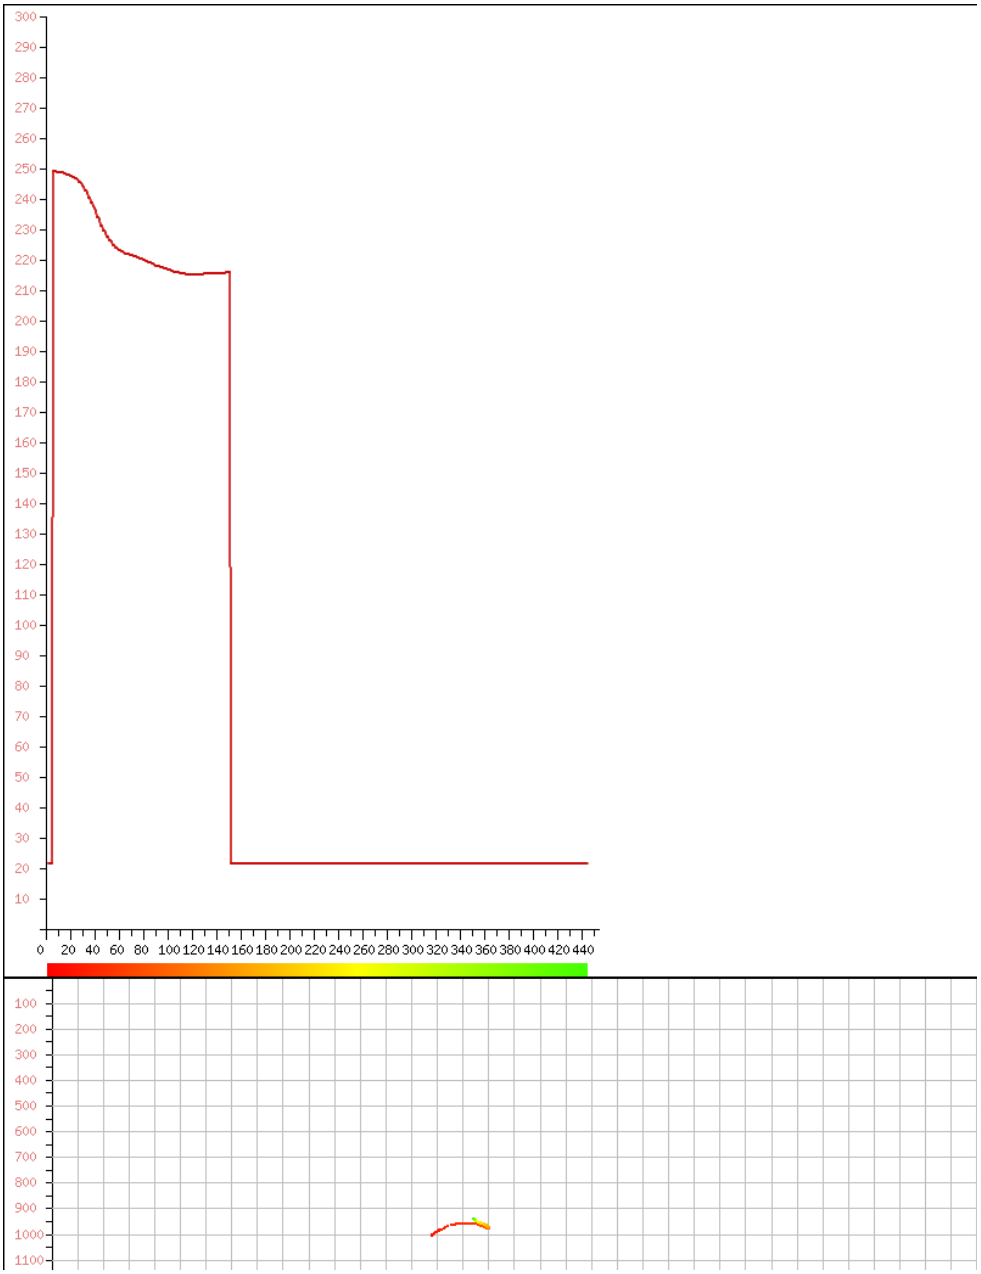

Figure 5.18 Progression of instant ID_{DP} across the completion of task 4 using ARM. ID_{DP} varies according to physical navigation until the high precision is activated. Then, it becomes constant.

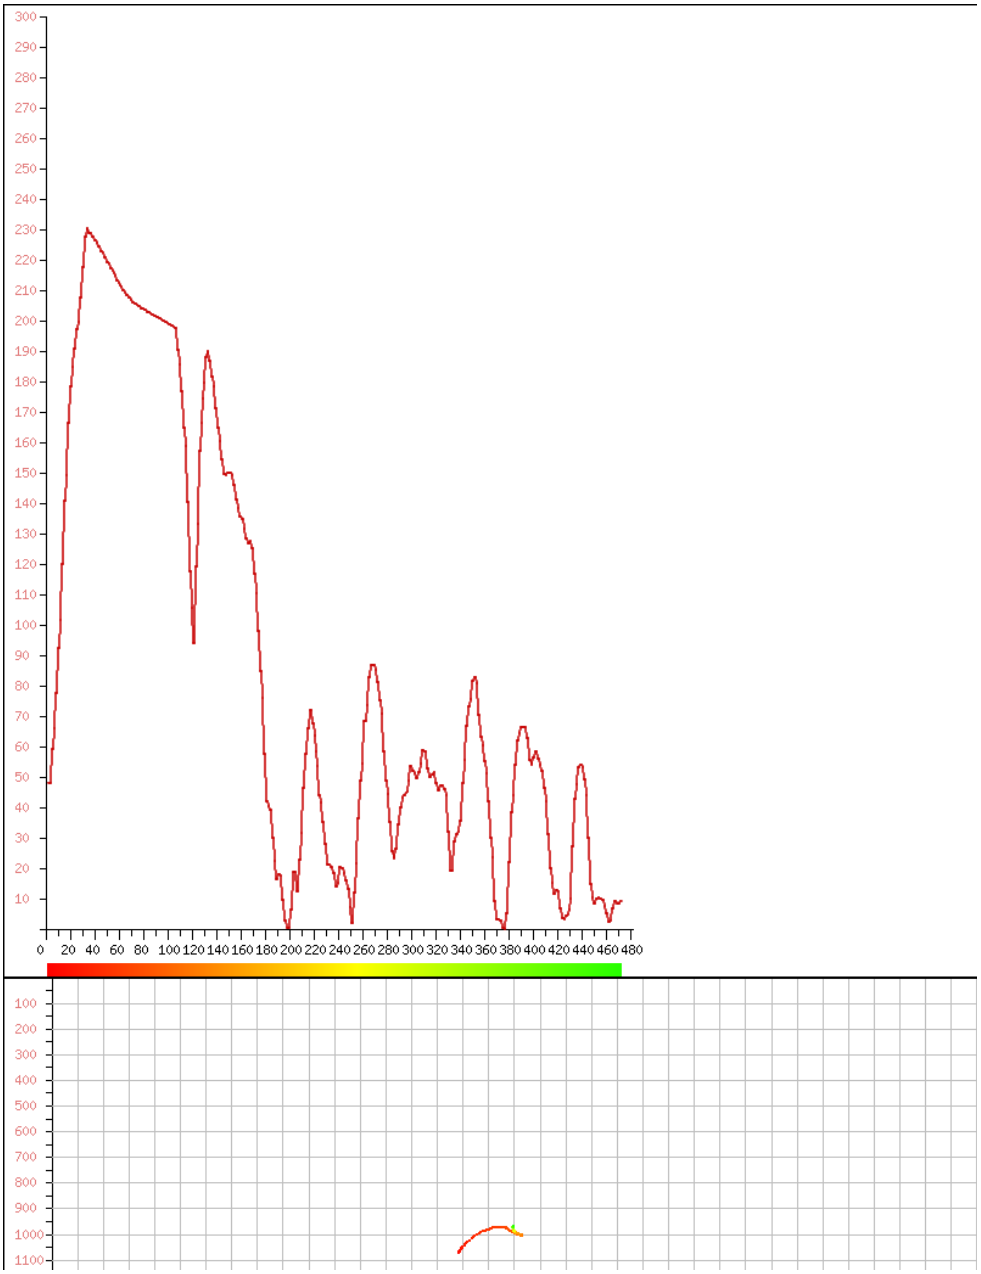

Figure 5.19 Progression of instant ID_{DP} across the completion of task 4 using DyCoDiR. ID_{DP} varies principally according to the input device speed, and we can see how the ID_{DP} reduces as the task nears the end.

high-precision techniques, and even less when visual acuity limitations were addressed. However, we also saw that some participants did not walk close to the targets at all, even when using basic distal pointing. That caused participants to get very frustrated with basic distal pointing, but even then some participants wouldn't walk.

That leads us to conclude that individual preferences play a large role in distal pointing physical navigation strategies. ARM and DyCoDiR do not only allow users to be precise, but they give users the choice to be precise while not having to walk close to the display. This has important implications, both in terms of application scenarios that require users to interact while at a distance, but also for users who prefer to interact at a distance even if the application allows up-close interaction.

5.11 Summary

In this chapter, we addressed research questions 3, 4 and 5 (section 1.2).

Research question 3, which asks if we can design effective high-precision distal pointing techniques based on the model guidelines, was addressed by the results showing that the techniques were both faster and more precise than basic distal pointing. We addressed the specific questions about the effects of visual size and high-precision activation. The results suggest that zooming can increase performance when coupled with high precision, but it should be explicitly activated. We also found a trend suggesting that performance is improved when the amount of precision is implicitly calculated.

We were able to address research question 4, about model validation, by performing an analysis based on the impact that each technique had on the indexes of difficulties of the tasks. Although the uncontrolled nature of the study did not allow for a good fit of the data to the distal pointing model, we were able to show that techniques that are designed with a goal of reducing task difficulty do indeed increase performance of distal pointing tasks.

Finally, we analyzed user behavior when performing distal pointing tasks to address research question 5. We found that participants tended to walk less with the high-precision techniques, using the techniques features rather than physical navigation in order to be precise. We also found that not only the technique impact users' decision to walk, but personal differences play a big role that sometimes takes precedence over the affordances provided by the techniques.

Chapter 6

Evaluation of a High-Precision Progressive Refinement Distal Pointing Technique

We conducted a formal evaluation comparing SQUAD to basic distal pointing.

6.1 Experiment Overview

We evaluated the task of pointing at circular targets that varied in radius, on a screen that was filled with distractor objects varying in number and density.

6.1.1 Goals and Hypotheses

The overall goal of the experiment was to explore the tradeoff between basic distal pointing and SQUAD. While basic distal pointing requires only one click, it requires precision with visually small targets. SQUAD, on the other hand, requires very little precision from the user, at the expense of multiple steps until the desired target is selected.

With this tradeoff in mind, we expected there to be an interaction between technique and target size. We hypothesized that SQUAD would take constant time with respect to target size, while basic distal pointing would be slow with small targets and fast with large targets. We were unsure how the constant SQUAD times would compare to the times for basic distal pointing with the various target sizes, but expected that SQUAD would be faster in at least some target size conditions.

We also hypothesized that the number of distractor objects around the target would have a significant effect on time to select with SQUAD, but that the number of distractors would have no effect on basic distal pointing. We expected that SQUAD would outperform basic distal pointing when the number of distractors was small.

With respect to accuracy, we hypothesized that SQUAD would yield virtually no errors, due to its low required precision, whereas basic distal pointing would have more errors as the target sizes decrease.

Finally, we hypothesized that situations in which the tracking has more jitter would result in more errors and slower time for basic distal pointing, but would not impact SQUAD, as all the steps of the technique require very low pointing precision.

6.1.2 Design

We used a factorial within-subject design with repeated measures. There were four independent variables: *technique* (basic distal pointing, SQUAD), *tracking* (normal, jittery), *target size* (radii $0.01m$ or $0.26°$, $0.015m$ or $0.40°$, $0.04m$ or $1.06°$), and the number of distractors inside the selection sphere (referred to as distractor density) (16, 64, 256). After the guidelines from the distal pointing model (chapter 3), we emphasized in varying the target size, while keeping the movement amplitude within a roughly constant range (see section 6.3.1). The design was, thus, $2 \times 2 \times 3 \times 3$.

The order of presentation of technique and tracking was counterbalanced, blocked by technique, such that each participant performed both tracking conditions within the same technique before moving to the next one. Within the combinations of technique and tracking, each of the nine conditions of target size *vs.* distractor density was repeated eight times and presented in random order.

6.2 Analytic Evaluation

Before running an empirical study (section 6.3), we analytically evaluated performance in our experimental conditions based on predictive models.

The tradeoff between speed and accuracy described in Fitts' law is well known for pointing tasks MacKenzie (1992); Zhai et al. (2004). In chapter 3, we derived a model that applies for distal pointing tasks, in which the input device is remotely located with respect to the display area and the pointing is done in a direct fashion, as opposed to indirectly, for example, through the use of a mouse. SQUAD and basic distal pointing both use the motor behaviors described by the model, making it relevant to our study.

The predictive model of distal pointing states that the time to acquire a distal target through direct pointing depends strongly on the angular width of the target and, to a much lesser degree, on the angular amplitude of the wrist/arm movement required to complete the task. The difficulty of the task is expressed as

$$ID_{DP} = \left[\log_2\left(\frac{\alpha}{\omega^k} + 1\right)\right]^2, \tag{6.1}$$

where ID_{DP} is the index of difficulty, α is the angular amplitude of the movement and ω is the angular width of the target. The constant k is a power factor greater than one that expresses the higher importance of the target width relative to movement amplitude. The value of k was shown to be around three in the experimental setting from which we gathered data to derive the model (section 3.4.3) While this study used a different environment, we believe that it was similar and the value of the constant k should be approximately the same.

The goal of the progressive refinement technique that we propose is to reduce the index of difficulty of an individual pointing action to a minimum at the expense of increasing the number of actions needed to achieve the goal of selecting a single unique object in a highly cluttered environment.

In order to reduce ID_{DP} to a minimum in our study, we set the diameter of the selection sphere to 26.3°. The targets were chosen within a constant range from the starting point, so that the movement amplitude was selected randomly between 10.0° and 17.9°, with an average α of 14.0°. This yields an ID_{DP} of

$$ID_{DP} = \left[\log_2\left(\frac{14.0}{26.3^3}+1\right)\right]^2 \approx 1.23\times 10^{-6}. \tag{6.2}$$

Thus, the index of difficulty of the task of selecting the target region becomes virtually zero, and the expected time to select the target is very small. Similarly, the difficulty of selecting a quadrant in the quad-menu is minimal, as the angular width of each of the quadrants is 45°, yielding an ID_{DP} very near zero. According to Kopper *et al.*'s model, the intercept of the regression line for predicted selection times (when ID_DP tends to zero) is $1.091s$. However, values of ID_{DP} this close to zero have not been tested experimentally. With ω higher than α, we anecdotally observed that selection time is typically under the lower limit of $1s$ set in Kopper *et al.*'s model.

During the quad-menu phase of selection, the user needs to first find the quadrant containing the intended target, then point and click to select it. Although the target stands out and is easily distinguishable from the distractors, in theory the time for visual search will be greater when the number of distractors is larger. We base this assumption on the facts that the visual size of the target is smaller in the quad-menu when there are a lot of distractors, and that the perceived contrast diminishes as the objects become smaller Jr and Fullenkamp (1988). Thus, we hypothesize that the time it takes to select a target using SQUAD selection is

$$MT_{SQUAD} = c + \sum_{i=1}^{N}(c+v_i), \tag{6.3}$$

expressed in seconds, where N is the number of refinement iterations required during the quad-menu phase of the technique, c is an empirically determined constant related to the time it takes to point at a target whose difficulty tends to zero, and v_i is the visual search time to find the target in the quad-menu before movement starts. We expect v_1 to take the longest time, because, first, there is a switch in interaction mode, from sphere-casting to quad-menu selection, and a change in the visual environment. Also, the number of distractors is at its maximum, and it decreases as refinements are made, reducing the target search space and time. Due to the visual distinction of the target in relation to the distractors in our experimental setting, we expect v_i to be low in all phases of refinement and to not affect selection time by a large amount.

Here, we note some interesting characteristics of SQUAD selection as compared to basic distal pointing. First, the target size plays no role in the time it takes to complete a selection. We acknowledge that there may be a longer search time for visual segmentation in highly dense environments with small and occluded targets, but the motor movement time is constant once the target has been found. Second, the time it takes to select a target with SQUAD selection is directly proportional to the amount of clutter – or the number of distractor objects that exist in the region of the desired target. While the time it takes to select a target grows linearly with the increased number of iterations, the growth in the number of iterations is rather slow, on the order of $\lceil log_4(n) \rceil$, where n is the number of objects inside the sphere Figueroa et al. (2010).

In order to compare basic distal pointing with SQUAD, we decided to vary both the target size and the number of distractor objects that fall inside the selection sphere at any given time. We defined target ωs as 0.53°, 0.80° and 2.12°, yielding for basic distal pointing an ID_{DP} of 42.9, 23.4 and 1.67, respectively. The predicted time to complete the basic distal pointing tasks for each of the respective target sizes was 2.29s, 1.74s and 1.14s, respectively. We set the number of distractor objects inside the sphere to be 16, 64 and 256, yielding a total of 3, 4, and 5 clicks to select the target with SQUAD in each distractor density condition. This leads to a theorized $3c + v_{t_0}$, $4c + v_{t_1}$ and $5c + v_{t_2}$ seconds to select a target, where v_{t_i} is the total visual search time across all refinement phases in each condition. The value of c needs to be empirically determined, but we expect it to be less than one. Again, we believe v_{t_i} to be low and not affect movement time by a large amount.

6.3 Empirical Evaluation

In order to empirically validate the results from our analytic evaluation, we performed a comparative study of SQUAD and basic distal pointing.

6.3.1 Apparatus

We used a back-projected VisBox-SX system, with only one projector (monoscopic) to display the experimental environment on a 2.29m x 3.05m screen. The resolution of the graphics was 1400px x 1050px. A wireless Intersense IS-900 Wand was used for controlling the cursor on the screen.

The experimental software was written using the Vizard Virtual Reality Toolkit by WorldViz. It ran under Microsoft Windows XP on a workstation with an Intel Core2 660 CPU at 2.40GHz and 2GB of RAM. The frame rate was fixed at 55 frames per second for all conditions except with the high-density distractor conditions, in which it went down to around 15 frames per second in the sphere-casting phase only, because many collision tests with the selection sphere were necessary.

We were comfortable with the drop in frame rate for that one condition because the sphere-casting selection was very easy to perform.

The environment consisted of circular objects as shown in Figure 6.1. The user stood at the center of an invisible sphere of $2.155m$ radius, at an orthogonal distance of $1.52m$ to the display surface. The red target and the gray distractors were evenly distributed on the surface of the sphere. There was no head tracking or any virtual navigation of the environment and the user remained at a fixed location. The perspective projection of the objects caused them to have the correct visual size from the user's point of view at the center of the sphere. We made the decision to render the circles on the surface of a virtual sphere, as opposed to on the flat screen plane, because the effective angular width of objects displayed far from the center of a flat screen decreases (section 3.2.3). The perspective rendering of the circles near the edges of the display compensated for this effect in our environment, such that all objects had the same angular width from the user's point of view.

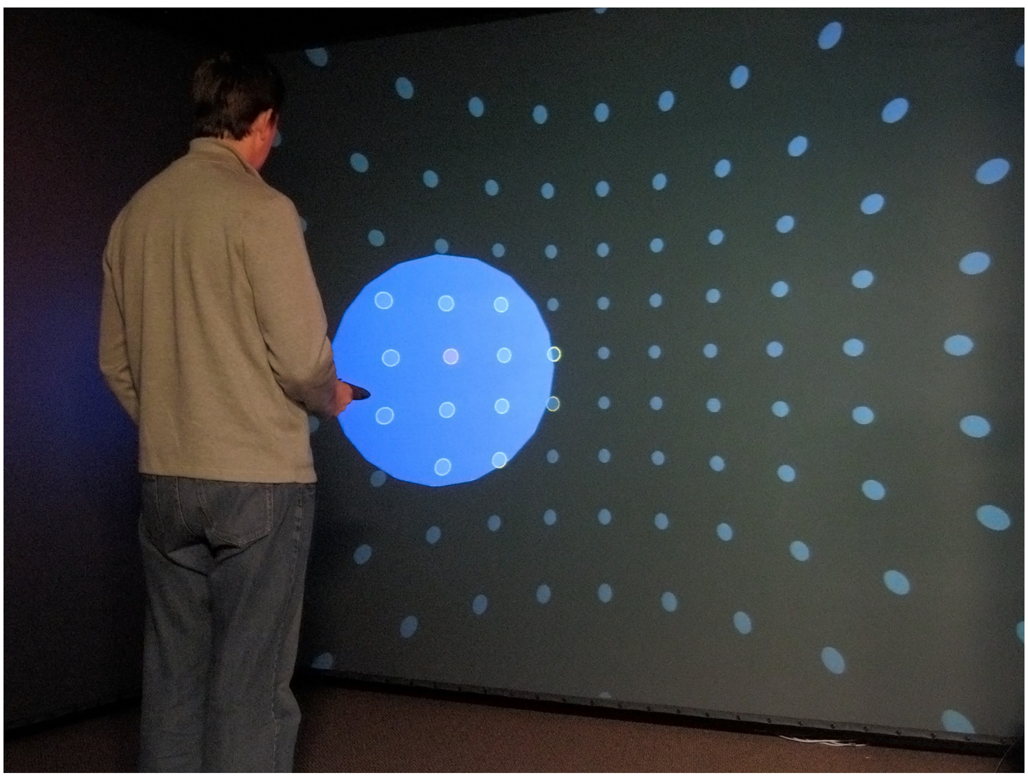

Figure 6.1 Experimental setup with sphere-casting.

The target position was randomly selected from a list of candidate targets that fell inside a torus-shaped section of the display sphere, and was limited by a small radius of $0.52m$ and a large radius of $0.77m$. This ensured that targets were presented in all directions from the center of the projection.

The cursor position was determined by a function of the yaw and pitch of the IS-900 wand and the display's field of view. With the user standing at a fixed position in front of the display, the position of the cursor closely matched the ray extending from the wand. We decided to rely only on the angular readings of the wand, rather then implementing 3D ray-casting based on the combination of position and orientation information, because we wanted to keep the motor difficulty to complete the task constant. Participants were told to keep their hand position within a small range over a mark on the floor that determined the center of the virtual sphere, and not to reach out with their arms. With the cursor position dependent solely on the wand's yaw and pitch, we were able to keep the motor behavior identical to that of basic distal pointing from the sweet spot at the center of the virtual sphere. There was then, of course, a mismatch between the position of the displayed cursor and that of the position of the 3D ray extending from the user hand with the screen. This offset was, however, minimal and no participant seemed to mind, or even notice, the difference.

Each task began with only two objects on the screen: a large yellow circle in the center, and the red target. Once the user clicked the large yellow object, it disappeared and the rest of the screen was filled with distractor objects. We did this for two reasons. First, clicking at a pre-determined spot before the start of a task meant that the angular amplitude of the movement was kept in a controlled range. Second, by not showing the distractors in the beginning, the user could find the target location before starting the task, reducing any cognitive time to segment the target from the distractors to a minimum.

For the basic distal pointing condition, a crosshair represented the cursor and the task was finished when the user clicked the trigger button on the IS-900 wand. When the cursor intersected with an object, either the target or a distractor, the object was highlighted with a yellow border.

In the SQUAD condition, after the user clicked at the yellow object in the center of the display to begin the task, the cursor changed to a sphere (Figure 6.1). All objects that were inside or intersecting with the surface of the sphere were rendered with a highlight, indicating that they were active for selection. The sphere-casting action was committed by a click with the trigger button, and the display changed to the quad-menu (Figure 6.2). In order to maintain experimental control, for each distractor density, the quad-menu contained the same number of elements, even if the sphere did not have exactly that number of objects inside. These numbers were close enough that no participant ever noticed a mismatch between the objects inside the sphere and the objects displayed in the quad-menu. In the quad-menu, we decided to limit the display of the objects to approximately 50° of the view-field, as opposed to the full 90° of the projection screen. We made this decision to minimize the potential visual search time after the menu was displayed, and we found that 50° was enough to display a large number of objects, while still allowing the user to spot the target without any head movement. To refine the quad-menu selection, the user only needed to point anywhere in the quadrant that contained the target and click the trigger button.

We applied a dynamic recursive low-pass filter Vogel and Balakrishnan (2005) to the raw pitch and yaw data from the IS-900 wand. This filter provided a rapid response time while reducing tracking jitter to a minimum (the Kalman filters provided by the IS-900 system had a significant lag when the cursor was moving precisely, causing strange cursor "stickiness" effects).

For the jittery conditions, we applied a random offset between $-0.21°$ and $0.21°$ at each frame

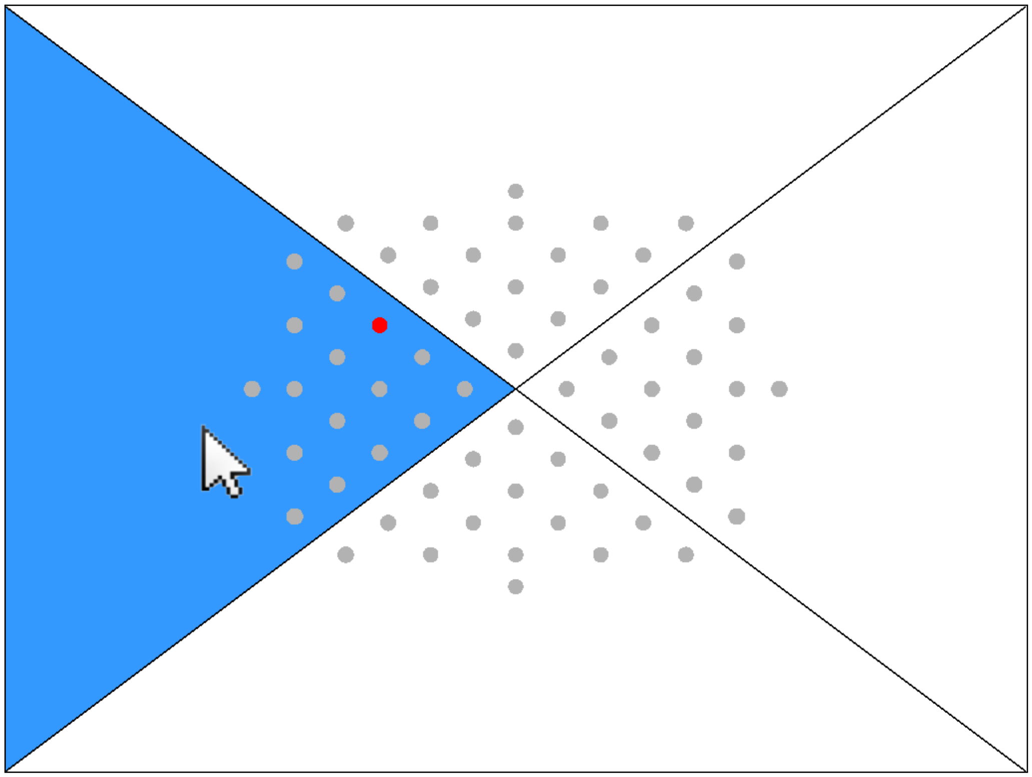

Figure 6.2 Quad-menu stage of SQUAD in the experiment.

to the filtered yaw and pitch readings of the wand. This resulted in a maximum error of 80% of the smallest target width, such that participants had a reasonable chance of successful selection in the hardest basic distal pointing condition.

6.3.2 Participants

We recruited 16 voluntary unpaid participants from the campus community to perform the study. Participants' ages ranged from 20 to 31 years old, with a median age of 22.5. Nine of the participants were female.

6.3.3 Procedure

Upon arrival, participants were greeted by the experimenter and given an informed consent form to read and sign. They were then given a color blindness screening test and proceeded to complete a background questionnaire. After that, they were shown the experimental setting and started learning the first technique *vs.* tracking combination. The learning was done with an easy condition so they could understand the technique without making errors. They were then given a practice session, in which they had to practice all nine target size *vs.* distractor density combinations in the current condition for at least 90*s*.

After practicing, they were reminded that they had to perform the trials as quickly as possible while trying not to make errors, and then performed eight sets of each of the nine combinations. When errors were made, the application displayed a message ("Not quite!") for 0.7*s* and the next task was displayed. The erroneous trial was then put into an array of trials that was presented in a new random order after the end of the set of trials for the current technique *vs.* tracking combination. This process was repeated to a maximum of five attempts per trial. If the user made five errors on a trial, it was deemed failed and was not presented again. The target position was the same for all attempts of a given trial.

After the end of each technique *vs.* tracking session, the participant completed a set of rating-scale questions and rested for up to two minutes. They then moved on to the next condition, following the same protocol, until all four technique *vs.* tracking combinations were completed.

Finally, the participant filled out a post-hoc questionnaire, comparing both techniques overall and in light of the other variables.

6.3.4 Results

We performed a factorial ANOVA with repeated measures on both dependent variables: time to complete a task and mean number of errors per trial.

Time

Overall, basic distal pointing was significantly faster than SQUAD ($F_{1,15} = 4.92$, $p < 0.05$), but only by two-tenths of a second. Interestingly, there was no main effect of tracking ($F_{1,15} = 0.001$, $p = 0.979$).

There were main effects of both distractor density ($F_{1,15} = 398.6$, $p < 0.0001$) and target size ($F_{1,15} = 153.4$, $p < 0.0001$). This indicates that the effects of distractor density on SQUAD and target size on basic distal pointing were so large that the variables were significant overall, but when we examine the interactions, we get a clearer picture of the effects.

The tradeoff of a single precise selection compared to multiple coarse selections can be clearly seen in the interactions of technique with target size and distractor density. Figure 6.3 shows the significant interaction between technique and distractor density ($F_{2,30} = 290.51$, $p < 0.0001$). The 95% confidence interval, at each density level, showed that SQUAD was significantly faster in the low density, that there was no significant difference for the medium density, and that basic distal pointing was significantly faster with high density.

As expected, all densities were significantly different from each other with SQUAD, while there was no statistical evidence for a difference for basic distal pointing between any distractor density pairs. There was, however, a slight increase in the mean task completion time for basic distal pointing as the distractor density increased. We believe that this may have been caused by an increase in visual processing time, as the distractors were highlighted as the cursor intersected with them. However, further studies should be done to verify this effect.

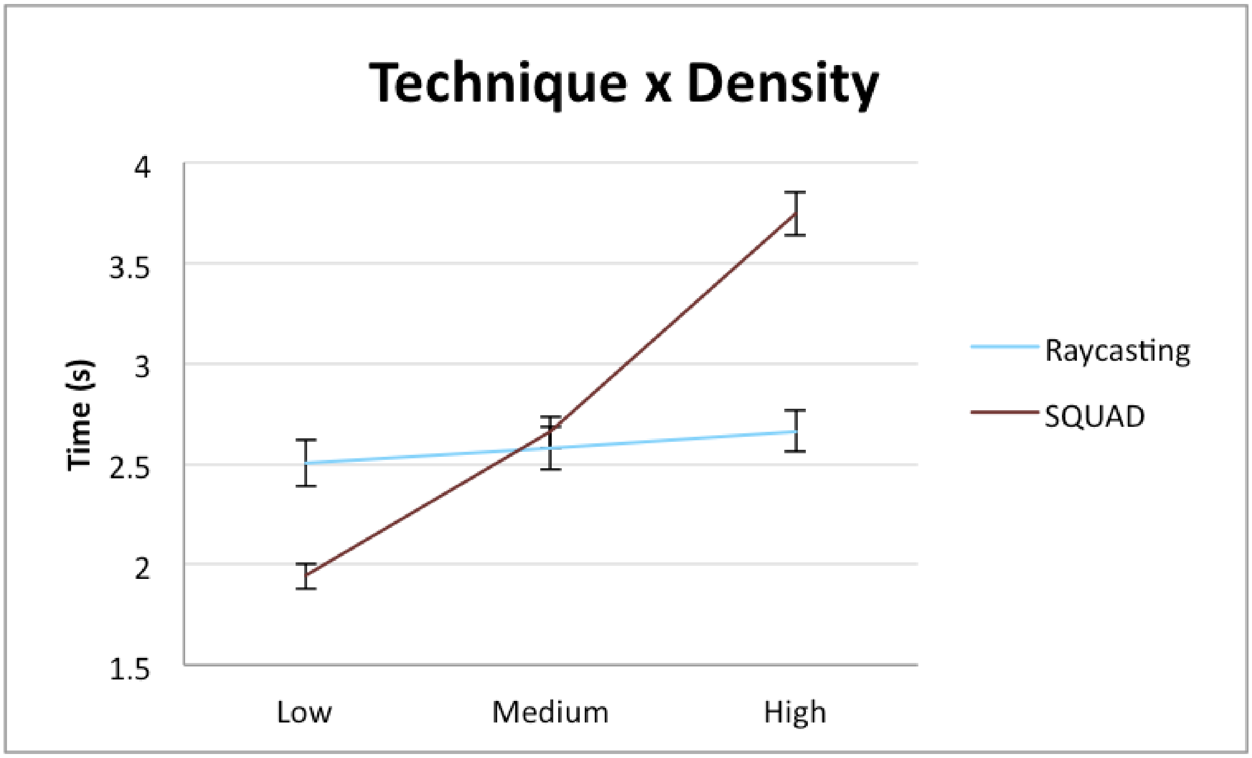

Figure 6.3 Interaction between technique and distractor density. The error bars represent standard error.

The other interaction that evidences the tradeoff is that of technique with target size. There was a highly significant interaction of these factors ($F_{2,30} = 135.17$, $p < 0.0001$), illustrated by Figure 6.4. As expected and predicted by the distal pointing model, pairwise comparisons showed that the smallest targets took significantly longer to select with basic distal pointing, while there were no significant differences among target sizes for SQUAD.

Looking at the interaction between technique and tracking, we expected to see that basic distal pointing would be slower with jittery tracking, while SQUAD would not be affected by tracking jitter. However, this interaction was not significant at a 95% confidence level ($F_{1,15} = 3.93$, $p = 0.066$). Despite near-significance, the mean difference in time with basic distal pointing was only about 0.1s, and more errors were made with bad tracking, which could indicate that participants favored speed over accuracy, even if instructed otherwise.

No other significant interactions were found, which is consistent with our hypotheses.

Figure 6.5 shows the mean results for all technique-density-target size combinations (the three densities for basic distal pointing are averaged in this graph, since density had no effect on basic

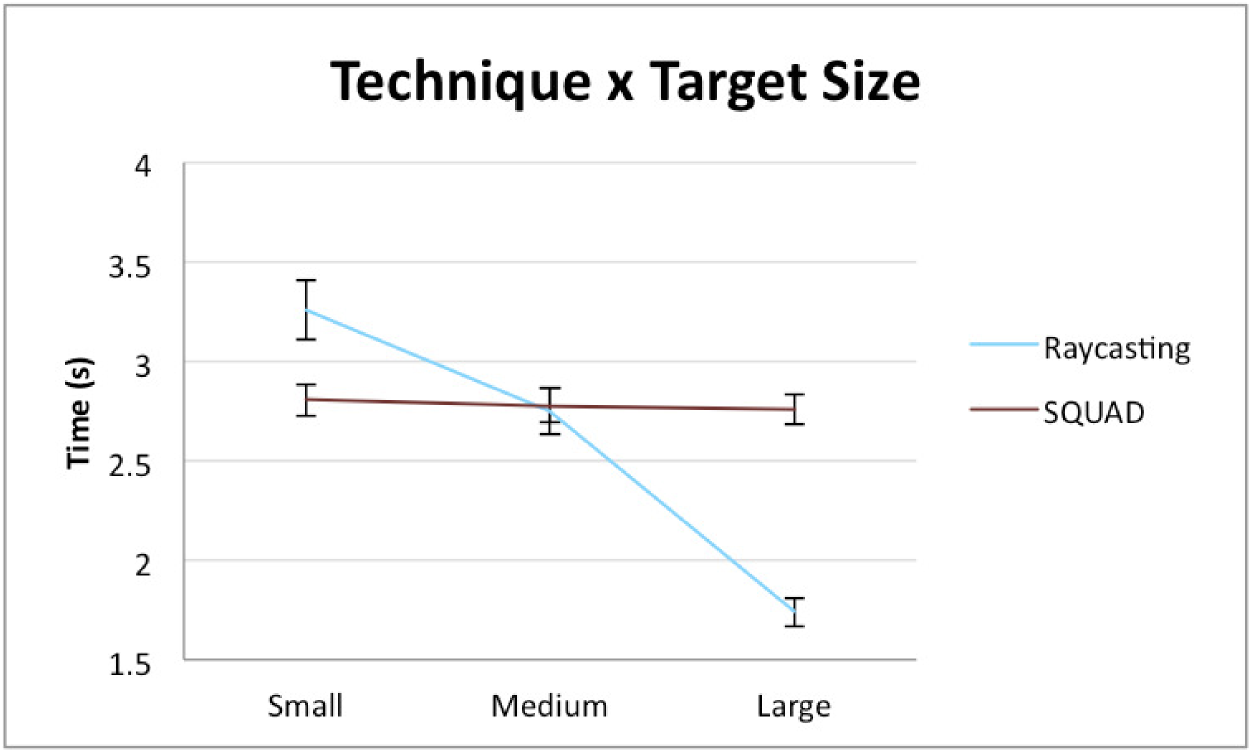

Figure 6.4 Interaction between technique and target size. The error bars represent standard error.

distal pointing performance). It is clear from this graph that SQUAD was significantly faster than basic distal pointing with low density and either small or medium size targets, and with medium density and small targets. In two other conditions, there was no significant difference between the two techniques. Finally, there are four conditions where basic distal pointing is significantly faster than SQUAD.

Errors

As expected, there was a significant main effect of technique with respect to errors ($F_{1,15} = 56.86$, $p < 0.0001$), with more errors being made with basic distal pointing. In fact, virtually no errors were made with SQUAD. The overall error rate with this technique was 0.007 errors per trial.

The lack of errors with SQUAD makes it interesting to look at the effects of tracking, target size and distractor density on basic distal pointing. Thus, we performed a new repeated measures ANOVA removing all the SQUAD conditions.

As expected, there was a significant main effect of target size on the number of errors per trial with basic distal pointing ($F_{2,30} = 46.21$, $p < 0.0001$), with more errors made with the smallest targets (Figure 6.6).

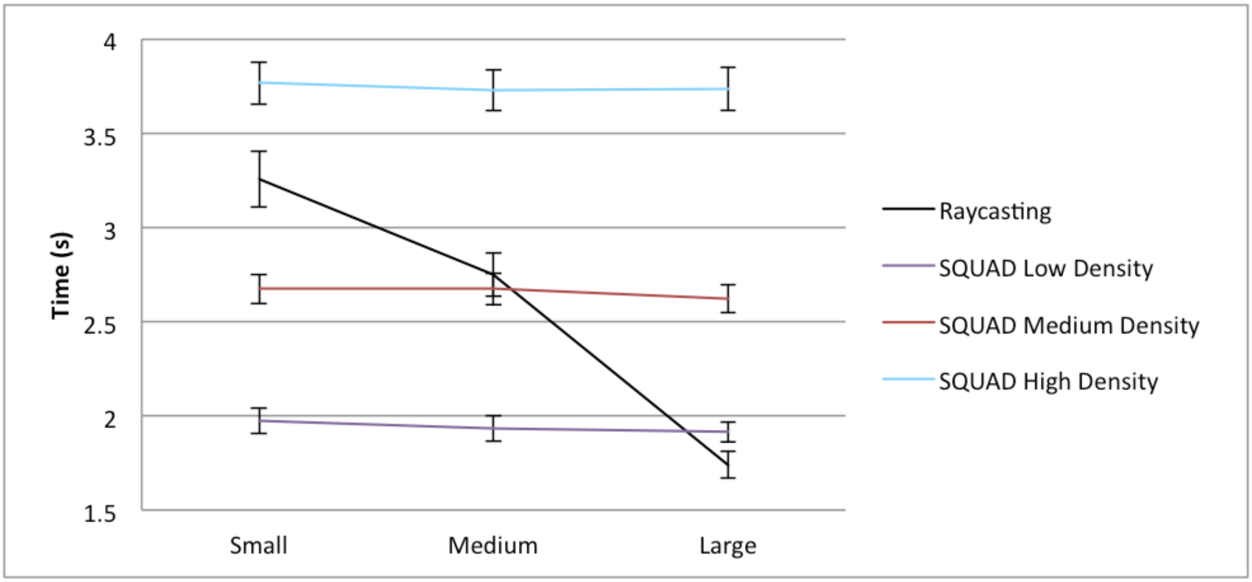

Figure 6.5 Mean results for all technique-density-target size combinations. Note that all basic distal pointing densities are averaged and displayed in a single line, since there was no significant difference among them. The error bars represent standard error.

Although the average number of errors was higher for the jittery conditions, we found no statistical evidence of this difference ($F_{1,15} = 1.12$, $p < 0.31$). We believe that, since the amount of jitter was controlled, users were able to learn and compensate for it, since the low pass filter applied before the jitter allowed most participants to keep the cursor fixed on an exact pixel when needed. That, combined with the fact that the maximum jitter was 80% of the minimum target width and the continuous clear highlighting feedback of cursor intersection may have caused users to adapt and learn to select accurately with basic distal pointing despite the jittery cursor.

User Preference

SQUAD was largely preferred by all participants of the experiment. When asked which technique they favored overall, when the cursor was jittery and when the targets were small, all 16 participants answered SQUAD. When asked about which technique they preferred when there were many distractors in the scene, the majority (nine) still preferred SQUAD, suggesting that the increased number of steps did not outweigh the overall preference of the technique; two participants were undecided, and the remaining five preferred basic distal pointing when many distractors were present.

It is also interesting to look at subjective ratings the participants gave for various aspects of both techniques. Participants were instructed to fill out a survey immediately after completing each of the techniques, and to rate the techniques on a seven-point scale for ease of learning, ease of use, and how hard the techniques were in various conditions (when the cursor contained artificial jitter,

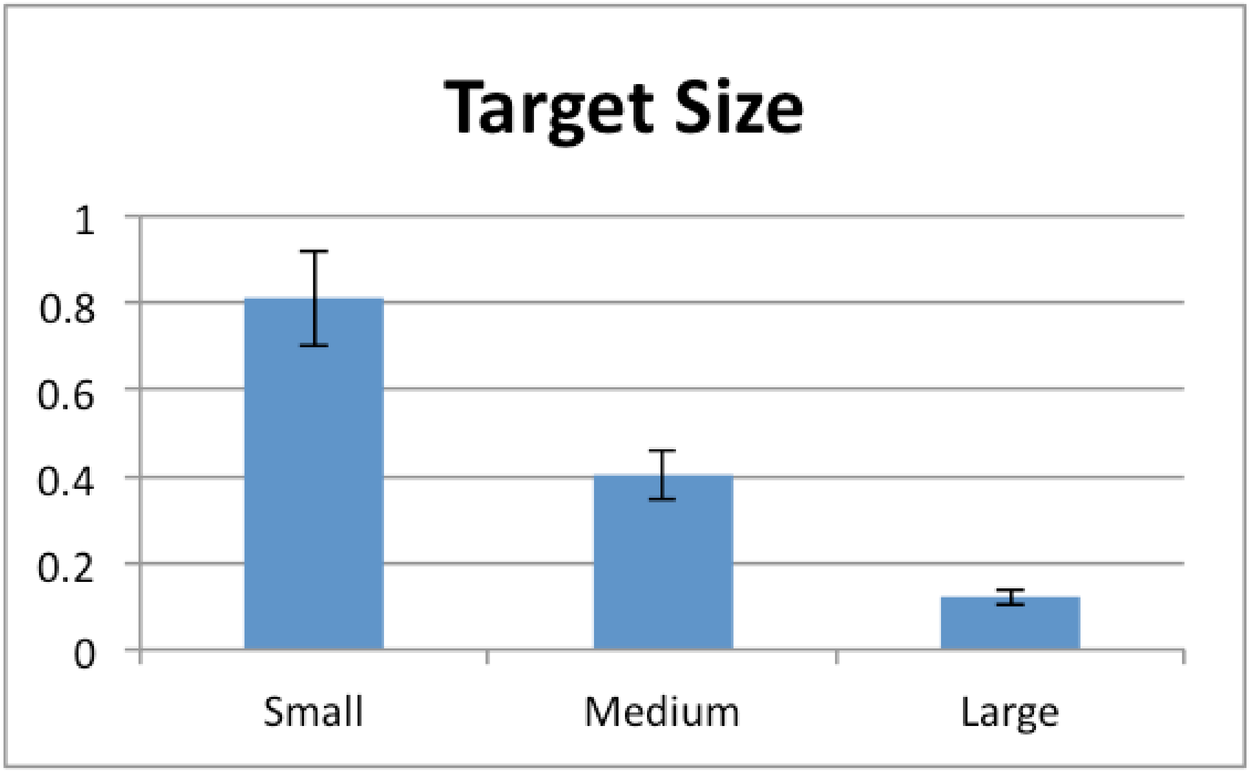

Figure 6.6 Mean number of errors per trial with basic distal pointing. The error bars represent standard error.

the targets were small, and there were many distractors). Also on a seven-point scale, participants were asked to rate their wrist, leg and back fatigue. Before answering the survey after the last technique, participants were instructed to respond independently of the answers to the first one.

We performed Wilcoxon Signed Rank tests on each of the questions. There was no significant difference in the reported ease of learning between the techniques ($n = 7, W = 20, insignificant$). For ease of use, basic distal pointing ($mdn = 4.5$) was ranked significantly more difficult than SQUAD ($mdn = 1$) ($z = 3.16, p < 0.001$). Participants found basic distal pointing significantly more difficult when the cursor was jittery ($z = 3.24, p < 0.001$) and when the targets were small ($z = 3.5, p < 0.001$). There was no significant difference with respect to task difficulty when many distractors were present ($mdn_{distal pointing} = 2, mdn_{SQUAD} = 3.5, z = -0.18, p = 0.19$).

Participants reported significantly more arm fatigue with basic distal pointing ($mdn = 5$) than with SQUAD ($mdn = 3.5$) ($z = 2.65, p < 0.05$). No significant difference was found between the two techniques for leg ($n = 5, W = 5, insignificant$) and back fatigue ($mdn_{basic distal pointing} = 3, mdn_{SQUAD} = 2, z = 1.52, p = 0.064$).

6.4 Model Validation

Based on the analytic evaluation and the empirical results of both techniques, we can validate the predictive models for the basic distal pointing and SQUAD pointing tasks.

Figure 6.7 shows the close linear fit of the ID_{DP}'s of each basic distal pointing conditions based on the model to the actual task performance time. However, we note that the intercept of the regression line is quite a bit higher than that predicted by the original model proposed by Kopper *et al.* Kopper et al. (2010). We believe that this is due to the nature of how errors were considered in the two experiments. In Kopper's experiment, errors did not invalidate a trial, so participants could be more careless when trying to select a target, as they could click multiple times to achieve the selection. In our case, on the other hand, the whole trial had to be attempted again, so we believe participants were more careful and certain that the cursor was inside the target area before they clicked. This resulted in a higher minimum time to complete a trial. The slope of the regression line is quite similar (0.028 in the original model and 0.037 in our experiment), and the correlation coefficient (R^2) is as high as 97.5%, which provides evidence that the distal pointing model was valid in our experimental environment.

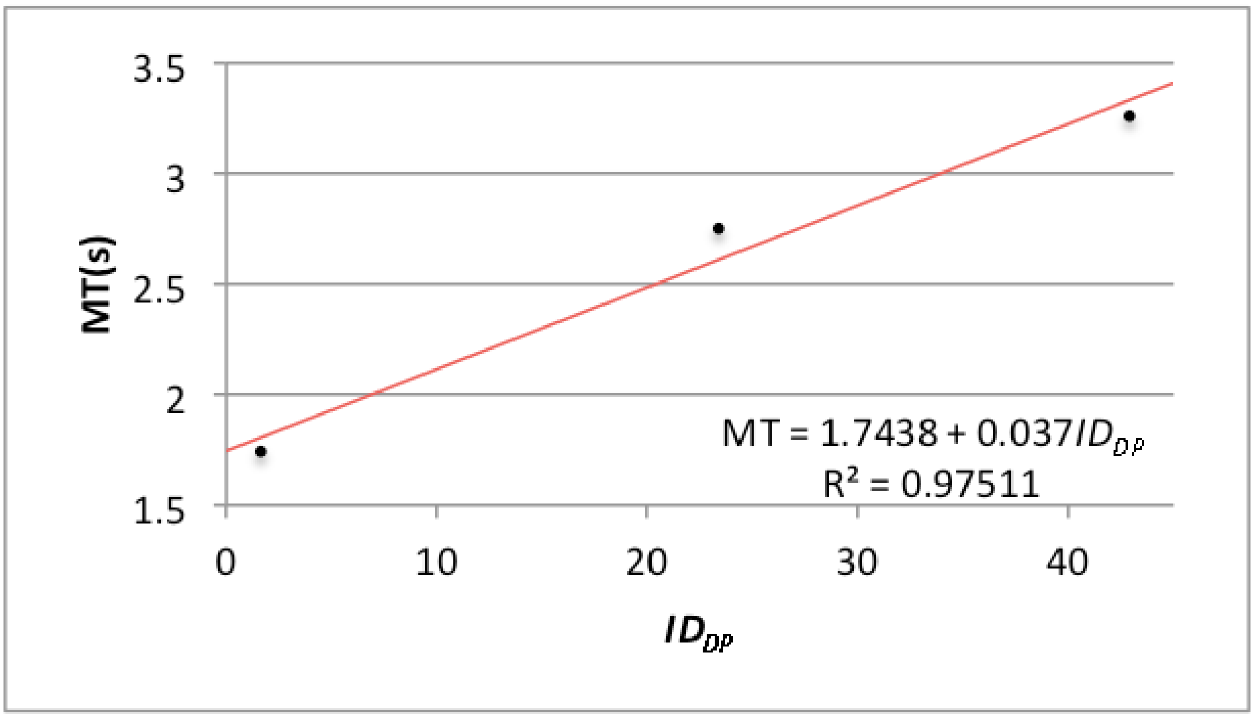

Figure 6.7 Scatter plot and regression line for the basic distal pointing conditions.

We can analyze the SQUAD pointing trials based on the time it took for each of the phases, which consisted of sphere-casting followed by two, three or four refinements. Overall, as we predicted, the selection time has a linear relationship with the number of refinements, as shown in

Figure 6.8. Notice that the growth is linear and the intercept is very close to zero, which emphasizes the constant increase in time as more refinements are needed.

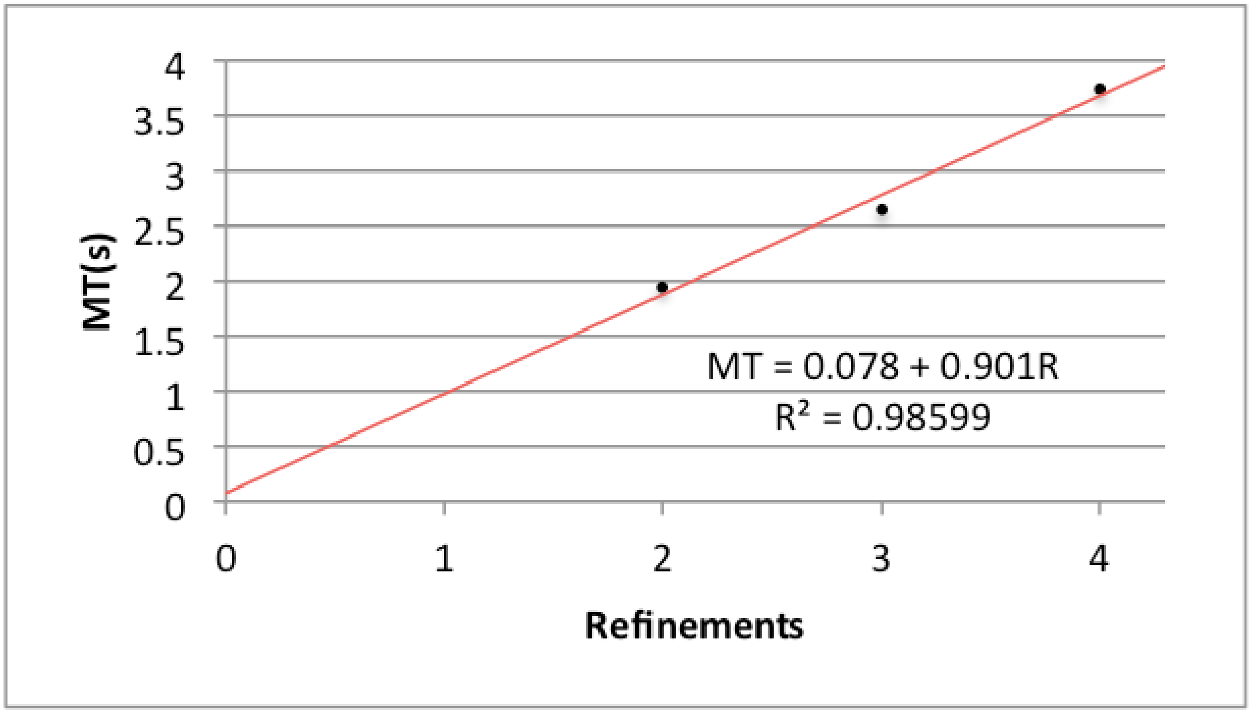

Figure 6.8 Scatter plot and regression line for the SQUAD pointing conditions.

Interestingly, there was a significantly longer time for the quad-menu selection in the first refinement step of the high-density distractors condition, in which there were a total of 256 objects in the quad-menu. The difference was on the order of $0.2s$ longer than in any other refinement phase, which were all within $0.05s$. The conclusion we derive from this is that visual search time was only meaningful when there were a very large number of objects in the quad-menu, while in all other conditions, this time was negligible. However, the time to complete the first refinement phase was significantly higher for all three target densities.

6.5 Discussion

The analytical evaluation of SQUAD selection was backed up by the results of an empirical study comparing it to basic distal pointing. We verified that the there is, indeed, a performance trade-off between immediate techniques that use one precise action to select an object and progressive refinement techniques that require very low precision at the expense of multiple steps. The use of SQUAD, and, by extension, other progressive refinement selection techniques, should be based on a consideration of this tradeoff. We found that SQUAD is significantly faster for selection of small objects and selection in low-density environments. When errors are considered, the case for SQUAD is even stronger, as it achieved near-perfect accuracy. A positive aspect of our approach is that the increase in the number of refinements needed grows very slowly, such that the task is not likely to take many refinement phases. Another interesting aspect of SQUAD is that the time to complete a task grows linearly as the number of refinements increase, whereas with basic distal pointing, the time increase is exponential as targets become smaller.

There is potential for much further research on selection by progressive refinement techniques. SQUAD is only one of a large set of techniques that fall within the design space of progressive refinement techniques. While SQUAD was highly efficient with near-zero error rates and better time than basic distal pointing with small targets, it has some limitations that need to be acknowledged. SQUAD was designed with a particular application setting in mind, and the nature of the task, which involved distant objects roughly arranged on a surface, influenced the design of SQUAD. Its sphere casting component is not well-suited for selecting from among items distributed in depth. However, SQUAD can be adapted to work well in such situations. For example, instead of a selection sphere, a cone or cylinder could be used to specify a deeper region of initial selection for further refinement in the quad-menu phase.

In addition, SQUAD works well for tasks that require the selection of objects that are visually distinct from other possible objects in the vicinity, and that do not depend on the spatial context. Selection tasks that depend on object location rather than visual features, for instance, cannot be achieved by SQUAD, but still may be achieved effectively with selection by progressive refinement. For example, objects can be kept within their spatial context if refinement is accomplished by zooming. In this case, the refinement would consist in the specification of an area in the view that would zoom to fill the display, decreasing the number of selectable objects. This is not necessarily equivalent to navigating close to objects to select them; the zooming could be done discretely. After the selection task, the viewpoint could return to the original position.

We found that the model proposed in chapter 3 accurately predicted performance with basic distal pointing. We also were able to find evidence that the performance of discrete progressive refinement selection techniques can be modeled by a direct linear relationship to the number of refinements necessary for completing a task. The use of analytical models in the evaluation of such techniques can provide a benefit in more realistic settings, in which control can be traded off for ecological validity. In such situations, many aspects, such as user strategy, confound the experimental control, and using reliable analytical models in the performance assessment may be the best choice.

6.6 Summary

With this chapter, we addressed research questions 1, 3 and 4.

We tackled research question 3 by demonstrating that SQUAD, a technique designed based on the distal pointing model guidelines, is effective and optimal for some environments. SQUAD specifically follows the guideline about increasing the effective targets' size. It does so by always providing a very large effective target size, as the actual target sizes are not relevant for SQUAD difficulty.

Research question 1 asks about understanding distal pointing. We addressed it by providing a model of performance of SQUAD selection, which defines a limiting distal pointing condition, in which the movement requires no precision and consists of only a ballistic phase.

Finally, we were able to address research question 4 about analytical evaluation by showing that the analytical evaluation of both SQUAD and basic distal pointing closely match the empirical performance. We were able to get very good fits of the data because this experiment was performed in a controlled environment. However, the empirical demonstration that SQUAD outperforms basic distal pointing in certain conditions is sufficient to provide validation for the model guidelines.

Chapter 7

Conclusions and Future Work

The goal of this thesis was to to understand distal pointing in order to design good distal pointing interfaces and be able to evaluate them. In light of the overall goal, this thesis has made three broad contributions.

- It advances the understanding of distal pointing motor behaviors.
- It provides the community with a set of distal pointing interaction techniques that were shown to provide high precision and accuracy.
- It demonstrates that distal pointing interaction techniques can be analytically evaluated based on a model of performance.

Understanding Distal Pointing

We advanced the understanding of distal pointing by proposing and deriving a model of human performance for distal pointing based on the results of an empirical study. The angular amplitude of movement and angular target width are the main parameters of the model, and we found that target size has a much more important weight in the difficulty of performing distal pointing tasks.

We provided a set of design guidelines for researchers and practitioners to follow when creating interfaces that can benefit from distal pointing. We found that the most important design guideline in distal pointing is make targets large is much more important than reducing the distances. This is an interesting finding that contrasts a well known guideline from Fitts' law (Fitts, 1954), which says that targets should be larger at the same rate as distances should be reduced.

We have discovered a new concept of interaction that is directly applicable to distal pointing which we called selection by progressive refinement. This interaction style offers a new avenue for the design of very accurate distal pointing interaction while keeping it fast and usable.

We found a tradeoff between the number of refinements necessary with progressive refinement and the required pointing accuracy that must be taken into account for the design of distal pointing techniques. When the visual size of objects is very small and the density of the environment is not too high, selection can be achieved more efficiently by progressive refinement.

We have also demonstrated that increasing visual size to address visual acuity limitations has a positive impact in distal pointing performance, independently of motor precision.

Design of Distal Pointing Techiques

We designed a set of interaction techniques based on distal pointing and verified through experimentation that they are effective in increasing the precision of distal pointing tasks.

We designed two high-precision immediate distal pointing techniques – ARM and DyCoDiR – which were shown to be successful in following the model guidelines. Each technique emphasizes different aspects of precise distal pointing interaction, and we showed that they are both successful in increasing precision of distal pointing while being usable. From an early design of ZELDA, a bi-manual distal pointing techniques aimed at increasing precision through zooming, we designed variations of ARM and DyCoDiR to address visual acuity limitations. We found that increasing the visual size of tasks that require a lot of precision lead to better performance. We found, however that the zoom activation should be done explicitly by the user, and not automatically by the application.

We designed SQUAD, a selection by progressive refinement distal pointing technique that was shown to provide virtually perfect accuracy and increase performance of distal pointing tasks especially in highly cluttered environments. Through SQUAD, we've shown that it is possible to design usable distal pointing techniques that allow the user to be lazy yet accurate.

Finally, we found that distal pointing techniques that provide high precision encourage interaction from a distance, but also that preference and individual differences can have a larger role in physical navigation strategies while using distal pointing.

Evaluation of Distal Pointing Techniques

We demonstrated that the distal pointing model can successfully predict performance of distal pointing tasks in controlled environments. Also, we showed that we can analytically evaluate high-precision distal pointing techniques that are designed following the model guidelines and that the expected indexes of difficulties of different techniques may be used to relatively compare techniques.

We validated a predictive model of performance of SQUAD, which demonstrated that the time to perform a selection using progressive refinement is related to the number of refinements necessary, and does not depend on target size at all.

7.1 Future Work

The research presented in this thesis opens a wide range of future research in distal pointing interaction.

The model we proposed is successful in defining the important aspects of a distal pointing tasks. However, it could be improved in several ways. First, the model does not account for physical navigation, and that is a big part of distal pointing interaction. Future work should be done on how to integrate a component of physical navigation to the model, so that analysis can be made based on expected walking behaviors. A good direction to achieve a model of distal pointing which accounts for physical navigation is to perform a fully controlled study that has navigation strategy as an independent variable.

Other predictive models of distal pointing are also encouraged as future work. The model we derived successfully predicts the movement time for univariate tasks. A natural direction for continuing research on distal pointing understanding is to verify the existence of a model for tasks in which the target measurements vary in two dimension. Prior research on bivariate (Accot and Zhai, 2003) and probabilistic (Grossman and Balakrishnan, 2005b) models can serve as starting points for new models of distal pointing.

The model was shown to successfully guide the design of interaction techniques based on distal pointing that provide high precision. Future work should concentrate on novel distal pointing interaction techniques based on the guidelines. New distal pointing techniques can involve more user position-aware interfaces, taking advantage of display edges to increase effective target sizes and improved ways of increasing the visual size of objects to address visual acuity limitations.

The design space of selection by progressive progressive refinement opens up a wide range of techniques that afford rapid and accurate yet careless interaction. One important piece of future work is the design of novel selection by progressive refinement techniques that address SQUAD limitations. For such, in-context refinement is important, and targets can be selected based on their spatial location, rather than appearance. For example, a technique that reduces the set of selectable objects by progressively zooming to the area of interest addresses this issue. Further, progressive refinement is not limited to selection, but rather can be applied to other tasks, such as navigation in virtual environments, or data organization in very large high resolution displays. For example, target-based navigation can be achieved by using progressive refinement to specify the target area, and data can be organized into clusters by progressively specifying objects that have similar properties.

Appendix A

Appendix: Experimental Documents for Model Study

VirginiaTech

Office of Research Compliance
Institutional Review Board
1880 Pratt Drive (0497)
Blacksburg, Virginia 24061
540/231-4991 Fax: 540/231-0959
E-mail: moored@vt.edu
www.irb.vt.edu
FWA00000572(expires 7/20/07)
IRB # is IRB00000667.

DATE: October 23, 2006

MEMORANDUM

TO: Doug A. Bowman
Tao Ni
Yi Wang

Approval date: 10/23/2006
Continuing Review Due Date:10/8/2007
Expiration Date: 10/22/2007

FROM: David M. Moore

SUBJECT: **IRB Expedited Approval:** "Design and Evaluation of 3D Interaction Techniques for Large High-Resolution Displays" , IRB # 06-602

This memo is regarding the above-mentioned protocol. The proposed research is eligible for expedited review according to the specifications authorized by 45 CFR 46.110 and 21 CFR 56.110. As Chair of the Virginia Tech Institutional Review Board, I have granted approval to the study for a period of 12 months, effective October 23, 2006.

As an investigator of human subjects, your responsibilities include the following:

1. Report promptly proposed changes in previously approved human subject research activities to the IRB, including changes to your study forms, procedures and investigators, regardless of how minor. The proposed changes must not be initiated without IRB review and approval, except where necessary to eliminate apparent immediate hazards to the subjects.
2. Report promptly to the IRB any injuries or other unanticipated or adverse events involving risks or harms to human research subjects or others.
3. Report promptly to the IRB of the study's closing (i.e., data collecting and data analysis complete at Virginia Tech). If the study is to continue past the expiration date (listed above), investigators must submit a request for continuing review prior to the continuing review due date (listed above). It is the researcher's responsibility to obtained re-approval from the IRB before the study's expiration date.
4. If re-approval is not obtained (unless the study has been reported to the IRB as closed) prior to the expiration date, all activities involving human subjects and data analysis must cease immediately, except where necessary to eliminate apparent immediate hazards to the subjects.

Important:
If you are conducting **federally funded non-exempt research**, this approval letter must state that the IRB has compared the OSP grant application and IRB application and found the documents to be consistent. Otherwise, this approval letter is invalid for OSP to release funds. Visit our website at http://www.irb.vt.edu/pages/newstudy.htm#OSP for further information.

cc: File

Dear students:

We are looking for voluntary participants for our project “design and evaluation of 3D interaction techniques for very large, high-resolution displays. This project is in conjunction with our CS 6724 3D Interaction class.

By participating our study, you will have an opportunity to work with a large high-resolution display we have constructed at Virginia Tech, consisting of 50 LCD panels. Your participation will help us gather both quantitative and qualitative insights into user interface design for emerging display technologies.

The study will consist of performing a set of object selection and manipulation tasks on large display surface. More details will be provided on site. The whole experiment will last for about 1 hour.

Informed Consent for Participant of Investigative Project

Virginia Polytechnic Institute and State University

Title of Project: Design and Evaluation of 3D Interaction Techniques for Large High-Resolution Displays

Principal Investigator: Dr. Doug Bowman, Tao Ni, Yi Wang, Ryan McMahan, Regis Kopper, Mara Silver, Michael DellaNoce, Brian Badillo, and William McConnell.

I. THE PURPOSE OF THIS RESEARCH/PROJECT

You are invited to participate in a study of 3D interaction techniques for very large, high-resolution displays. This research studies how 3D interaction techniques we designed affect object selection and manipulation task performance on large displays, compared to traditional 2D techniques.

II. PROCEDURES

Upon arrival, you will be given an introduction of our experiment background, goals, facilities, and study procedures. After reading and signing this informed consent form, you will be given a short period of time to practice a set of interaction techniques we designed on a very large, high-resolution display, and familiarize yourself with input and output devices. When you feel ready to perform tasks, we will load the computer program and read you a set of predefined tasks. You are expected to finish each task as soon as possible without sacrificing accuracy. We will record the time you take to complete each task. Finally, you will fill out a post-experiment questionnaire and rate your experience with different interaction techniques. A free form interview is conducted if you have any additional comments not addressed by the questionnaire.

III. RISKS

There will not be more than minimal risks by involving in our study.

IV. BENEFITS OF THIS PROJECT

Your participation in this project will provide information that may be used to improve the user interface design for very large, high-resolution displays. No guarantee of benefits has been made to encourage you to participate. You may receive a synopsis summarizing this research when completed.

V. EXTENT OF ANONYMITY AND CONFIDENTIALITY

The results of this study will be kept confidential. Your written consent is required for the researchers to release any data identified with you as an individual to anyone other than personnel working on the project. The information you provide will have your name removed and only a subject number will identify you during analyses and any written reports of the research.

VI. COMPENSATION

Your participation is voluntary and unpaid.

VII. FREEDOM TO WITHDRAW

You are free to withdraw from this study at any time for any reason.

VIII. APPROVAL OF RESEARCH

This research has been approved, as required, by the Institutional Review Board for projects involving human subjects at Virginia Polytechnic Institute and State University, and by the Department of Computer Science.

IX. SUBJECT'S RESPONSIBILITIES AND PERMISSION

I voluntarily agree to participate in this study, and I know of no reason I cannot participate. I have read and understand the informed consent and conditions of this project. I have had all my questions answered. I hereby acknowledge the above and give my voluntary consent for participation in this project. If I participate, I may withdraw at any time without penalty. I agree to abide by the rules of this project

____ _______________
Signature Date

____ _________________________
Name (please print) Contact: phone or address or

____
Email address (OPTIONAL)

________ Not consent to be videotaped

________ Consent to be videotaped

Should I have any questions about this research or its conduct, I may contact:

Investigator: Dr. Doug A. Bowman Phone (540) 231-2058
Professor, Computer Science Department (231-6931)
email: bowman@vt.edu

Tao Ni, Yi Wang, Ryan McMahan, Regis Kopper, Mara Silva, Michael DellaNoce, Brian Badillo, and William McConnell.

Graduate students, Computer Science Department
Email: nitao@vt.edu, samywang@vt.edu, rymcmaha@vt.edu, kopper@vt.edu, mara@vt.edu, mdellano@vt.edu, bbadillo@vt.edu, wmcconne@vt.edu.

Review Board: Dr. David Moore Phone (540) 231-4991
Chair, Virginia Tech Institutional Review Board
For the Protection of Human Subjects
Email: moored@vt.edu

cc: the participant, Dr. Bowman, Tao Ni, Yi Wang, Ryan McMahan, Regis Kopper, Mara Silver, Michael DellaNoce, Brian Badillo, and William McConnell.

Participant #:
Date:

Background Questionnaire

Please help us to categorize our user population by completing the following items.

Gender (circle one): Male Female

Age: ____________

Do you wear glasses or contact lenses (circle one)?
No Glasses Contact Lenses

Are you (circle one)
Right-handed Left-handed Ambidextrous

Occupation (if student, indicate graduate or undergraduate):

____

Major / Area of specialization (if student): ________________________________

Rate your tiredness level today: (circle one)

•----------------------•----------------------•----------------------•

very tired somewhat tired a little tired not tired at all

Rate your expertise with computers: (circle one)

•----------------------•----------------------•----------------------•

beginner amateur intermediate advanced

How often do you use computers...

...for work? (circle the best answer)	...for fun? (circle the best answer)
a. not at all	a. not at all
b. once a month	b. once a month
c. once a week	c. once a week
d. several times a week	d. several times a week
e. daily	e. daily

Have you ever used a large display system? If so, please describe it (what type of display was used, what kind of application (e.g. game, virtual environments, geospatial) was running, how did you interact with the system, etc.).

Evaluation of pointing tasks on a very large tiled display

Welcome and thank you very much for your participation on our experiment. You will help us evaluate different aspects of selection using a pointing device on the Gigapixel display.

In this experiment, you will perform several selections of two objects on the screen. The objects size, distance between them and your distance from the display will vary.

To select an object, point the cursor on top of it and click with your index finger to select. Once selected, the object will turn to yellow. Each trial starts with the green object and ends with the blue object.

You will perform 5 trials per condition. Perform each trial (selecting the green and the blue icon) as fast as you can. Your performance times will be recorded and it is important that you do your best. Your errors (e.g. misclicks) will be recorded, so try to click as few times as possible to select each icon.

You will now have 5 minutes to practice and get used to the interface.

5 minutes . . .

Now, you will begin your trials. You will have a break after each set of trials. If you feel you need a break before that, let me know.

Participant #:
Date:

Post-Experiment Questionnaire

Please complete the following questions. Select a value that best describe your experience from 1 to 7 as directed.

1. Please rate how much arm fatigue you experienced while performing the trials (1 being very low, 7 being very high):

•-------------•-------------•-------------•------------•------------•------------•

1 2 3 4 5 6 7

2. Please rate how much back fatigue you experienced while performing the trials (1 being very low, 7 being very high):

•-------------•-------------•-------------•------------•------------•------------•

1 2 3 4 5 6 7

3. Please rate how much leg fatigue you experienced while performing the trials (1 being very low, 7 being very high):

•-------------•-------------•-------------•------------•------------•------------•

1 2 3 4 5 6 7

4. Please feel free to add any comments, suggestions, and concerns.

Appendix B

Appendix: Experimental Documents for High Precision Distal Pointing Techniques Study

VirginiaTech

Office of Research Compliance
Institutional Review Board
2000 Kraft Drive, Suite 2000 (0497)
Blacksburg, Virginia 24060
540/231-4606 Fax 540/231-0959
e-mail irb@vt.edu
Website: www.irb.vt.edu

MEMORANDUM

DATE: August 12, 2010

TO: Doug A. Bowman, Regis Kopper

FROM: Virginia Tech Institutional Review Board (FWA00000572, expires June 13, 2011)

PROTOCOL TITLE: An Evaluation of Five Distal Pointing Techniques

IRB NUMBER: 10-646

Effective August 11, 2010, the Virginia Tech IRB Chair, Dr. David M. Moore, approved the new protocol for the above-mentioned research protocol.

This approval provides permission to begin the human subject activities outlined in the IRB-approved protocol and supporting documents.

Plans to deviate from the approved protocol and/or supporting documents must be submitted to the IRB as an amendment request and approved by the IRB prior to the implementation of any changes, regardless of how minor, except where necessary to eliminate apparent immediate hazards to the subjects. Report promptly to the IRB any injuries or other unanticipated or adverse events involving risks or harms to human research subjects or others.

All investigators (listed above) are required to comply with the researcher requirements outlined at http://www.irb.vt.edu/pages/responsibilities.htm (please review before the commencement of your research).

PROTOCOL INFORMATION:
Approved as: **Expedited, under 45 CFR 46.110 category(ies) 6, 7**
Protocol Approval Date: **8/11/2010**
Protocol Expiration Date: **8/10/2011**
Continuing Review Due Date*: **7/27/2011**
*Date a Continuing Review application is due to the IRB office if human subject activities covered under this protocol, including data analysis, are to continue beyond the Protocol Expiration Date.

FEDERALLY FUNDED RESEARCH REQUIREMENTS:
Per federally regulations, 45 CFR 46.103(f), the IRB is required to compare all federally funded grant proposals / work statements to the IRB protocol(s) which cover the human research activities included in the proposal / work statement before funds are released. Note that this requirement does not apply to Exempt and Interim IRB protocols, or grants for which VT is not the primary awardee.

The table on the following page indicates whether grant proposals are related to this IRB protocol, and which of the listed proposals, if any, have been compared to this IRB protocol, if required.

Date*	OSP Number	Sponsor	Grant Comparison Conducted?

*Date this proposal number was compared, assessed as not requiring comparison, or comparison information was revised.

If this IRB protocol is to cover any other grant proposals, please contact the IRB office (irbadmin@vt.edu) immediately.

cc: File

Dear students:

We are looking for voluntary participants for our project "An evaluation of five distal pointing techniques."

Your participation will help us gather both quantitative and qualitative insights into user interface design for selection and placement techniques in a variety of environments, from home entertainment systems to virtual reality applications.

The study will consist of performing a set of object selection tasks on a very large high-resolution display. More details will be provided on site. The whole experiment will last for about 1 hour.

Informed Consent for Participant of Investigative Project

Virginia Polytechnic Institute and State University

Title of Project: **An evaluation of five distal pointing techniques.**

Investigators: Dr. Doug Bowman, Regis Kopper.

I. THE PURPOSE OF THIS RESEARCH/PROJECT

You are invited to participate in a study to compare five distal pointing selection and placement techniques. This research studies different tradeoffs that concern precise selection and placement from a distance.

II. PROCEDURES

Upon arrival, you will be given an introduction of our experiment background, goals, facilities, and study procedures. After reading and signing this informed consent form, you will be given a short period of time to gain practice in distal pointing selection tasks at the Gigapixel display, with varied object and target sizes, shapes and positions, and familiarize yourself with input and output devices. After practicing the tasks, we will load the computer program and read you a set of predefined tasks. **You are expected to finish each task as soon as possible without sacrificing accuracy.** Your performance will be determined by the time you take to complete a task, and each mistake you make will penalize your time. Finally, you will fill out a post-experiment questionnaire and rate your experience with different interaction techniques. A free form interview is conducted if you have any additional comments not addressed by the questionnaire.

III. RISKS

There will not be more than minimal risks by involving in our study.

IV. BENEFITS OF THIS PROJECT

Your participation in this project will provide information that may be used to improve the user interface design for applications that make use of distal pointing interaction. No guarantee of benefits has been made to encourage you to participate. You may receive a synopsis summarizing this research when completed.

V. EXTENT OF ANONYMITY AND CONFIDENTIALITY

The results of this study will be kept confidential. Your written consent is required for the researchers to release any data identified with you as an individual to anyone other than personnel working on the project. The information you provide will have your name removed and only a participant number will identify you during analyses and any written reports of the research.

VI. COMPENSATION

The participant with the best performance in each of the technique will be awarded $20, and the participant that has the best overall performance will be awarded $40. No participant shall receive more than one award.

VII. FREEDOM TO WITHDRAW

You are free to withdraw from this study at any time for any reason.

VIII. APPROVAL OF RESEARCH

This research has been approved, as required, by the Institutional Review Board for projects involving human subjects at Virginia Polytechnic Institute and State University, and by the Department of Computer Science.

IX. SUBJECT'S RESPONSIBILITIES AND PERMISSION

I voluntarily agree to participate in this study, and I know of no reason I cannot participate. I have read and understand the informed consent and conditions of this project. I have had all my questions answered. I hereby acknowledge the above and give my voluntary consent for participation in this project. If I participate, I may withdraw at any time without penalty. I agree to abide by the rules of this project

____ ____________

Signature Date

____ ________________________

Name (please print) Contact: phone or address or

____

Email address (OPTIONAL)

________ Not consent to be videotaped

________ Consent to be videotaped

Should I have any questions about this research or its conduct, I may contact:

Investigators: Dr. Doug A. Bowman Phone (540) 231-2058
Professor, Computer Science Department (231-6931)
email: bowman@vt.edu

Regis Kopper.
Graduate student, Computer Science Department
email: kopper@vt.edu

Review Board: Dr. David Moore Phone (540) 231-4991
Chair, Virginia Tech Institutional Review Board
For the Protection of Human Subjects
Email: moored@vt.edu

cc: the participant, Dr. Bowman, Regis Kopper.

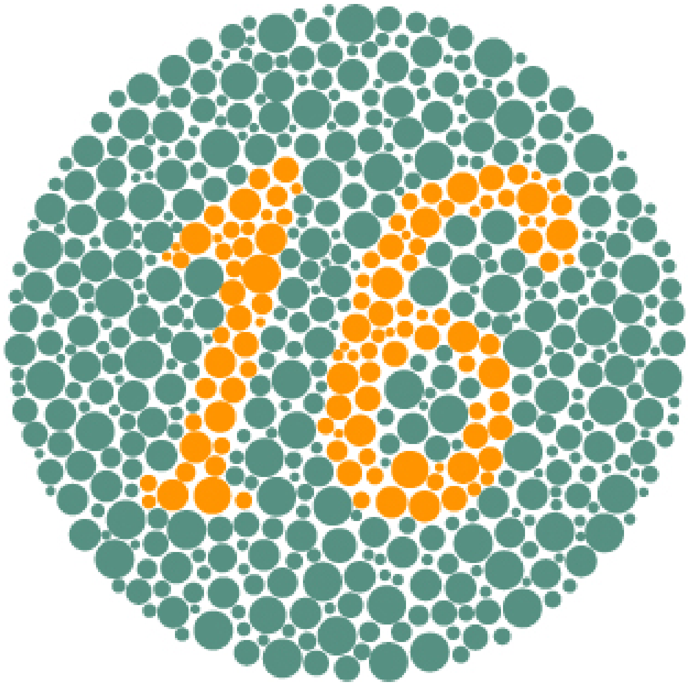

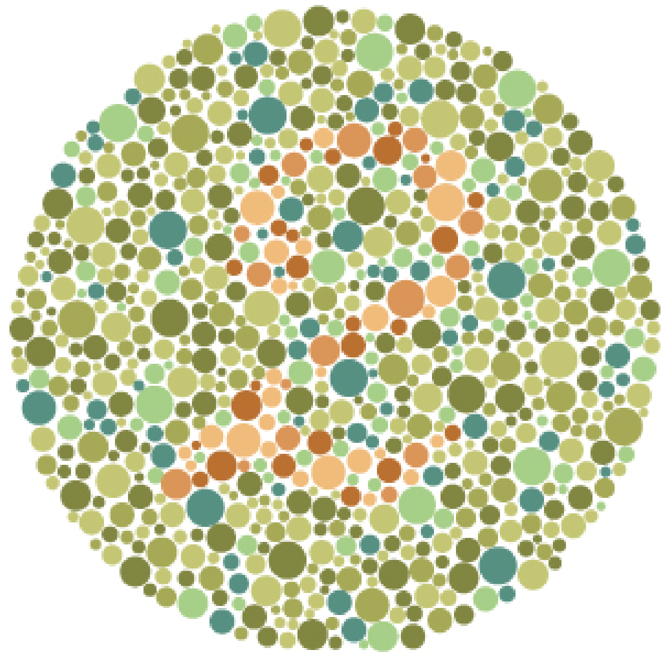

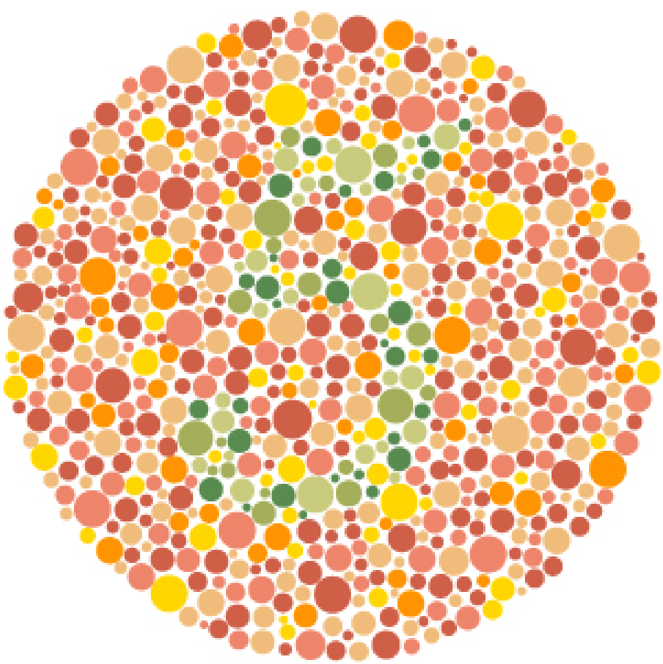

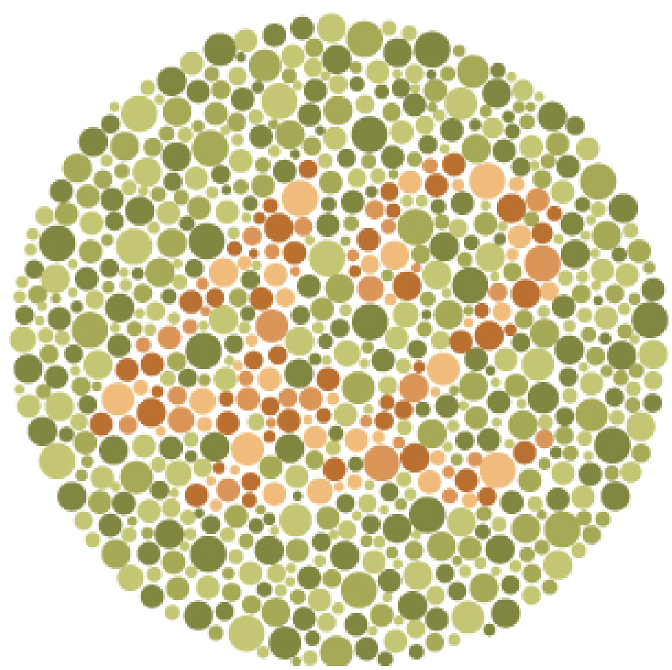

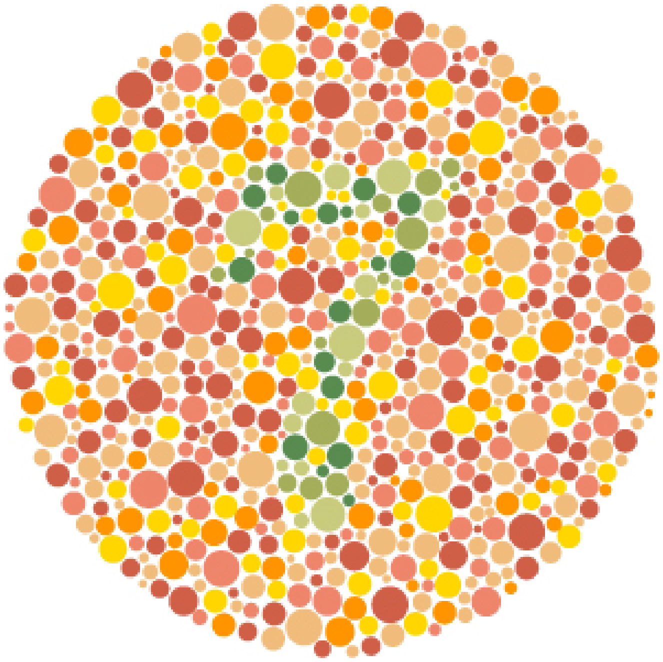

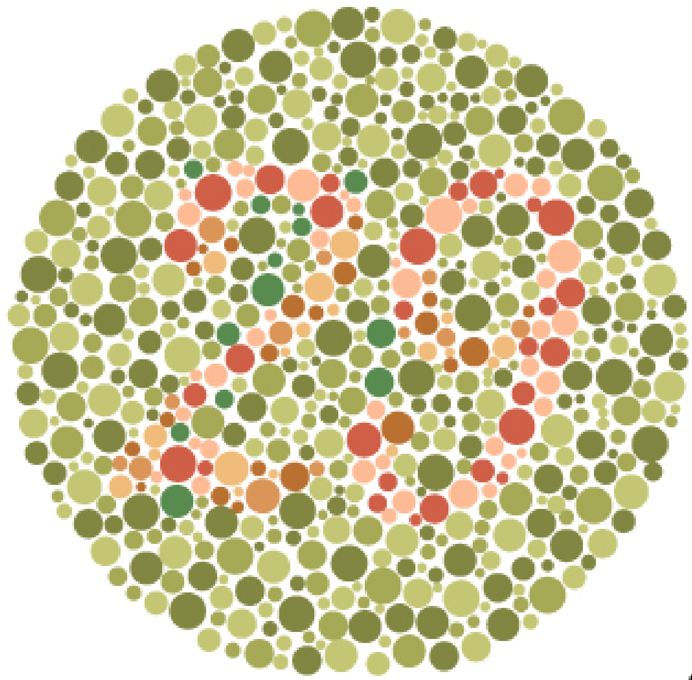

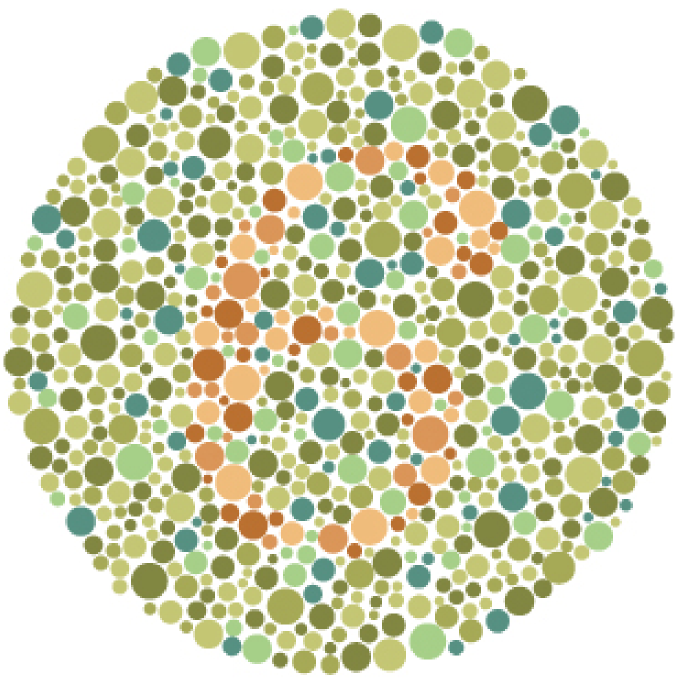

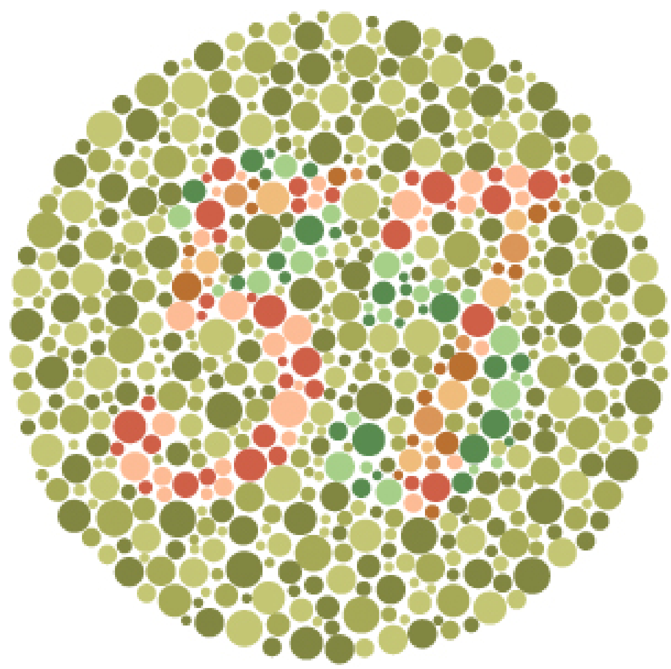

Participant #:
Date:

Background Questionnaire

Please help us to categorize our user population by completing the following items.

Gender (circle one): Male Female

Age: _____________

Do you wear glasses or contact lenses (circle one)? No Glasses Contact Lenses

Are you (circle one) Right-handed Left-handed Ambidextrous

Occupation (if student, indicate graduate or undergraduate):

____

Major / Area of specialization (if student): _________________________________

Rate your tiredness level today: (circle one)

•---------------------------•---------------------------•---------------------------•

very tired somewhat tired a little tired not tired at all

Do you have any experience with playing the Wii, Playstation Move or any other video game that uses motion tracking? Yes No

If yes, how many hours per week, on average, do you play?

•---------------•---------------•---------------•---------------•---------------•---------------•

less than 1 1 5 10 15 20 more then 20

Do you have any experience with playing video games other than the ones described above?
Yes No
If yes, how many hours per week, on average, do you play?

•---------------•---------------•---------------•---------------•---------------•---------------•

less than 1 1 5 10 15 20 more then 20

(continue behind)

What type of video games do you usually play (select all that apply)?

- ☐ First-person shooting games (e.g. Half Life, Quake, etc…)
- ☐ Sports games (e.g. Need for Speed, Madden, etc…)
- ☐ Massive multiplayer online games (e.g. World of Warcraft)
- ☐ Real-time strategy games (e.g. Starcraft)
- ☐ Mobile devices casual games (e.g. Iphone, Ipad, etc…)
- ☐ Motion capture casual games (e.g. Wii Sports, Wii Fit, …)
- ☐ Social games (e.g. Farmville, Mafia Wars, …)
- ☐ Third-person action games (e.g. Grand Theft Auto, Gears of War, etc…)
- ☐ Other: ______________

Do you have any experience with large high-resolution displays? Yes No

If yes, please explain:

An Evaluation of Six Distal Pointing Techniques

Procedural Overview

Techniques:
1 - Raycasting without zoom
2 - Raycasting with zoom
3 - ARM without zoom
4 - ARM with zoom
5 - Dynamic CD ratio without zoom
6 - Dynamic CD ratio with zoom

1. Pre-screening:
    a) Color-blindness test
    b) Confirm right-handedness
2. Consent Form
3. Background Survey
4. Initial Instructions
5. For each technique in the order that they are presented:
    a) Tutorial
    b) Trials
6. Exit Survey

Initial Instructions:

Start script with pure ray-casting, no zoom and no map.

[./usabilityEvaluationInitialTraining1] – One easy trial

Welcome and thank you for your participation in our experiment. You will help us evaluate different selection techniques using a pointing device resembling a laser pointer on the Gigapixel display.

In this experiment, you will perform several selections and placement of objects on the screen. The objects' shape, position and distance between them will vary.

The goal of this study is to compare different pointing techniques. In order for the techniques to be comparable, you need to perform the trials as quickly as possible, making as few mistakes as you can.

We need you to do your best. As an encouragement, we will award the participants with the best performance in each technique the amount of $20, and the participant with the best overall performance the amount of $40.

Now, I will explain how to complete the tasks. These instructions are general and are the same for all the six techniques you will use.

[give hat] Please keep wearing this hat throughout the experiment.

[show input device] This is the pointing device. Hold it with your right hand. The only button you will use is the trigger. Hold it in a comfortable position and use your index finger to click. During the experiment, you will not be allowed to use you other hand to steady the input device.

Before the start of each trial, you should locate both the orange and the red objects on the screen. Once you see it, point down with the input device and click. Then, point straight ahead and you will see the cursor. Sometimes you may need to click more than once. In that case, keep clicking until you see the black border.

Notice the orange object. This is the target. Point to the target and click the trigger button. Notice the orange target turns to green. This indicates that you have begun the task.

Now click on the target again. Notice that an error message has been displayed. This indicates that you have made an error. Anytime you make an error, the current task you are doing will reset and you will have to attempt the task again at the end of the set of trials. Click once to remove the error message. Click the target to begin the task again.

Now click anywhere in the blue space on the screen. Notice the error message indicating that you have made an error. You should not click in the open space for these tasks. Click once to remove the error message. Click the target to begin the task again.

Now click AND hold to drag the red object.

Now drag the object so that half is over the target and the other half is not. Drop the object here. Notice the error message. You must drop the object completely inside the target for these tasks. Click once to remove the error message. Click the target to begin the task again.

Remember that your errors will negatively affect your performance, so try to make as few errors as possible.

Now drag the icon completely over the target. Drop the icon here. Notice that there is not an error message displayed and the next task is ready for you to begin. This is the proper way to complete a task by clicking on the target and then dragging the icon completely inside of the target. Go ahead and complete this task.

Start script with pure ray-casting, no zoom and no map.
[./usabilityEvaluationInitialTraining2] – One medium trial

Now, see how this trial is more difficult. You don’t need to be standing at a certain position, you can move around freely in front of the display. If you think that moving around will improve your performance, you can do it. Remember that you should complete the task as fast as possible, trying not to make errors. Also, keep in mind that we are not evaluating you, but the techniques that we designed.

Go ahead and complete this task the most efficient way you can.

Tutorial for Raycasting no Zoom:

Start script. [./usabilityEvaluationRaycastingNoZoomTraining– Place cursor over window]

Now we're going to begin the first technique of the experiment. We are going to have you complete some tasks involving icons and targets.

Locate the orange target and the red object, then point down and click to activate the cursor. Observe the map. Your task is to place objects on top of buildings and secured areas on the map. It will be there for all the tasks, so it is important to locate the objects before you click to activate the cursor. The areas on the map that you need to place the red object on top of are assigned by the orange targets.

Notice the line that extends from the device defined by these three little spheres. This indicates where you are pointing with the device.

Go ahead and see how the cursor moves based on the ray and where you are pointing.

Go ahead and complete the tasks until the application quits. Let me know if you have any questions.

Session for Raycasting no zoom:

[application quits]
[load experimental environment] *Start script.* [./usabilityEvaluationRaycastingNoZoom – Place cursor over window]

Now you will perform two sets of 18 tasks with the technique you just tried. If you make any mistakes, you will have a chance to correct them at the end of each set.

After each set of tasks, a message will appear asking you to take a break. Take as long as you need and let me know when you're ready to continue.

Now, complete the trials as you were practicing them until the application finishes.

[application quits]
Please rate the technique you just tried.

Tutorial for raycasting with zoom:

Start script. [./usabilityEvaluationRaycastingZoomTraining– Place cursor over window]

Now we're going to begin the next technique of the experiment.

Point down and click to activate the cursor.

Point at the orange target and don't move for a second. Observe how the target zooms in. This is a feature of this technique. Whenever the cursor moves slowly, the zoom will appear as a visual aid. Once the cursor moves fast again, the zoom will disappear.

Go ahead and see how the cursor moves based on the ray and where you are pointing.

Go ahead and complete the tasks until the application quits. Let me know if you have any questions.

Session for raycasting with zoom:

[application quits]
[load experimental environment] *Start script.* [./usabilityEvaluationRaycastingZoom – Place cursor over window]

Now you will perform two sets of 18 tasks with the technique you just tried. If you make any mistakes, you will have a chance to correct them at the end of each set.

After each set of tasks, a message will appear asking you to take a break. Take as long as you need and let me know when you're ready to continue.

Now, complete the trials as you were practicing them until the application finishes.

[application quits]
Please rate the technique you just tried.

Please compare the last two techniques you just tried.

Tutorial for ARM no zoom:

Start script. [./usabilityEvaluationARMNoZoomTraining– Place cursor over window]

Now we're going to begin the next technique of the experiment.

In this set of trials, you will hold a mouse on your left hand.

Point down with your pointing device and click to activate the cursor.

Notice the left button on the device in your left hand. Remember that the left button is used for controlling the mode of movement for the cursor. When you press and hold the left button, you enter **relative pointing mode** and the movement of the cursor depends on your wrist rotation, and not anymore on the ray extending from the input device. This provides you more precision when needed. When you release this left button you go back to **absolute pointing mode** and the position of the cursor corresponds to the intersection of the ray extending from the pointing device with the screen.

Now, point anywhere at the display and activate relative mode by pressing the left mouse button in your left hand. Observe how the cursor moves slowly. This is a feature of this technique. Whenever you are in relative mode, your movement will be more precise. Once you release the button, the cursor moves fast again.

Go ahead and see how the cursor moves based on the ray and where you are pointing and experiment with changing the modes of movement.

Go ahead and complete the tasks until the application quits. Let me know if you have any questions.

Session for ARM no zoom:

[application quits]
[load experimental environment] *Start script.* [./usabilityEvaluationARMNoZoom – Place cursor over window]

Now you will perform two sets of 18 tasks with the technique you just tried. If you make any mistakes, you will have a chance to correct them at the end of each set.

After each set of tasks, a message will appear asking you to take a break. Take as long as you need and let me know when you're ready to continue.

Now, complete the trials as you were practicing them until the application finishes.

[application quits]
Please rate the technique you just tried.

Tutorial for ARM + zoom:

Start script. [./usabilityEvaluationARMZoomTraining– Place cursor over window]

Now we're going to begin the next technique of the experiment.

In this set of trials, you continue to hold a mouse on your left hand.

Point down with your pointing device and click to activate the cursor.

Notice the left button on the device in your left hand. Remember that the left button is used for controlling the mode of movement for the cursor. When you press and hold the left button, you enter **relative pointing mode** and the movement of the cursor depends on your wrist rotation, and not anymore on the ray extending from the input device. This provides you more precision when needed. When you release this left button you go back to **absolute pointing mode** and the position of the cursor corresponds to the intersection of the ray extending from the pointing device with the screen.

Now, point at the orange and activate relative mode by pressing the left mouse button in your left hand. Observe how the target zooms in. This is a feature of this technique. Whenever you are in relative mode, the zoom will appear as a visual aid. Once the cursor moves fast again, the zoom will disappear.

Go ahead and see how the cursor moves based on the ray and where you are pointing and experiment with changing the modes of movement.

Go ahead and complete the tasks until the application quits. Let me know if you have any questions.

Session for ARM + zoom:

[application quits]
[load experimental environment] *Start script.* [./usabilityEvaluationARMZoom – Place cursor over window]

Now you will perform two sets of 18 tasks with the technique you just tried. If you make any mistakes, you will have a chance to correct them at the end of each set.

After each set of tasks, a message will appear asking you to take a break. Take as long as you need and let me know when you're ready to continue.

Now, complete the trials as you were practicing them until the application finishes.

[application quits]
Please rate the technique you just tried.

Please compare the last two techniques you just tried.

Tutorial for DynCDRation no zoom:

Start script. [./usabilityEvaluationDynCDRatioNoZoomTraining– Place cursor over window]

Now we're going to begin the next technique of the experiment.

Point down and click to activate the cursor.

Now, point straight ahead. Observe how the cursor position coincides with the ray that projects from the input device. Now, slowly, move towards your right. Notice how much slower the cursor moves. This is a relative pointing mode, which gives you more precision. When you start moving fast, the cursor will catch up and will match the virtual ray extending the pointing device, going back to absolute mode.

If you ever want to bring the cursor back to absolute mode, just shake the input device for a little bit and the cursor position will match the ray.

Go ahead and see how the cursor moves based on the ray and where you are pointing.

Go ahead and complete the tasks until the application quits. Let me know if you have any questions.

Session for DynCDRatio no zoom:

[application quits]
[load experimental environment] *Start script.* [./usabilityEvaluationDynCDRatioNoZoom – Place cursor over window]

Now you will perform two sets of 18 tasks with the technique you just tried. If you make any mistakes, you will have a chance to correct them at the end of each set.

After each set of tasks, a message will appear asking you to take a break. Take as long as you need and let me know when you're ready to continue.

Now, complete the trials as you were practicing them until the application finishes.

[application quits]
Please rate the technique you just tried.

Tutorial for DynCDRatio + zoom:

Start script. [./usabilityEvaluationDynCDRatioZoomTraining– Place cursor over window]

Now we're going to begin the last technique of the experiment.

Point down and click to activate the cursor.

Point at the orange target and don't move for a second. Observe how the target zooms in. This is a feature of this technique. Whenever the cursor moves slowly, the zoom will appear as a visual aid. Once the cursor moves fast again, the zoom will disappear.

Now, point straight ahead. Observe how the cursor position coincides with the ray that projects from the input device. Now, slowly, move towards your right. Notice how much slower the cursor moves. This is a relative pointing mode, which gives you more precision. When you start moving fast, the cursor will catch up and will match the virtual ray extending the pointing device, going back to absolute mode.

If you ever want to bring the cursor back to absolute mode, just shake the input device for a little bit and the cursor position will match the ray.

Go ahead and see how the cursor moves based on the ray and where you are pointing.

Go ahead and complete the tasks until the application quits. Let me know if you have any questions.

Session for DynCDRatio + zoom:

[application quits]
[load experimental environment] *Start script.* [./usabilityEvaluationDynCDRatioZoom – Place cursor over window]

Now you will perform two sets of 18 tasks with the technique you just tried. If you make any mistakes, you will have a chance to correct them at the end of each set.

After each set of tasks, a message will appear asking you to take a break. Take as long as you need and let me know when you're ready to continue.

Now, complete the trials as you were practicing them until the application finishes.

[application quits]
Please rate the technique you just tried.

Please compare the last two techniques you just tried.

Participant #:
Date:

Raycasting ratings

Please answer the following questions based on the technique you just performed. Select a value that best describe your experience from 1 to 7 as directed.

1. How easy was it to learn the raycasting technique?

1	**2**	**3**	**4**	**5**	**6**	**7**
Very easy			Neutral			Very difficult

2. How easy was it to use the raycasting technique?

1	**2**	**3**	**4**	**5**	**6**	**7**
Very easy			Neutral			Very difficult

3. How did you like the raycasting technique?

1	**2**	**3**	**4**	**5**	**6**	**7**
Not at all			Neutral			Liked a lot

Participant #:
Date:

Raycasting with zoom ratings

Please answer the following questions based on the technique you just performed. Select a value that best describe your experience from 1 to 7 as directed.

1. How easy was it to learn the raycasting with zoom technique?

•--------------•--------------•--------------•-------------•-------------•-------------•

1	2	3	4	5	6	7
Very easy			Neutral			Very difficult

2. How easy was it to use the raycasting with zoom technique?

•--------------•--------------•--------------•-------------•-------------•-------------•

1	2	3	4	5	6	7
Very easy			Neutral			Very difficult

2. How easy was it to use the raycasting with zoom technique?

•--------------•--------------•--------------•-------------•-------------•-------------•

1	2	3	4	5	6	7
Very easy			Neutral			Very difficult

Raycasting comparison

For the raycasting techniques (the last two your performed), how do you rate the zoom feature as opposed to when zoom was not present?

•--------------•--------------•--------------•-------------•-------------•-------------•

1	**2**	**3**	**4**	**5**	**6**	**7**
Much worse with zoom			About the same			Much better with zoom

Participant #:
Date:

Relative pointing ratings

Please answer the following questions based on the technique you just performed. Select a value that best describe your experience from 1 to 7 as directed.

1. How easy was it to learn the relative pointing technique?

1	2	3	4	5	6	7
Very easy			Neutral			Very difficult

2. How easy was it to use the relative pointing technique?

1	2	3	4	5	6	7
Very easy			Neutral			Very difficult

2. How easy was it to use the relative pointing technique?

1	2	3	4	5	6	7
Very easy			Neutral			Very difficult

Participant #:
Date:

Relative pointing with zoom ratings

Please answer the following questions based on the technique you just performed. Select a value that best describe your experience from 1 to 7 as directed.

1. How easy was it to learn the relative pointing with zoom technique?

•--------------•--------------•--------------•-------------•-------------•-------------•

1	**2**	**3**	**4**	**5**	**6**	**7**
Very easy			Neutral			Very difficult

2. How easy was it to use the relative pointing with zoom technique?

•--------------•--------------•--------------•-------------•-------------•-------------•

1	**2**	**3**	**4**	**5**	**6**	**7**
Very easy			Neutral			Very difficult

2. How easy was it to use the relative pointing with zoom technique?

•--------------•--------------•--------------•-------------•-------------•-------------•

1	**2**	**3**	**4**	**5**	**6**	**7**
Very easy			Neutral			Very difficult

Relative pointing comparison

For the relative pointing techniques (the last two you performed), how do you rate the zoom feature as opposed to when zoom was not present?

•--------------•--------------•--------------•-------------•-------------•-------------•

1	2	3	4	5	6	7
Much worse with zoom			About the same			Much better with zoom

Participant #:
Date:

Adaptive pointing ratings

Please answer the following questions based on the technique you just performed. Select a value that best describe your experience from 1 to 7 as directed.

1. How easy was it to learn the adaptive pointing technique?

1	2	3	4	5	6	7
Very easy			Neutral			Very difficult

2. How easy was it to use the adaptive pointing technique?

1	2	3	4	5	6	7
Very easy			Neutral			Very difficult

2. How easy was it to use the adaptive pointing technique?

1	2	3	4	5	6	7
Very easy			Neutral			Very difficult

Participant #:
Date:

Adaptive pointing with zoom ratings

Please answer the following questions based on the technique you just performed. Select a value that best describe your experience from 1 to 7 as directed.

1. How easy was it to learn the adaptive pointing with zoom technique?

1	2	3	4	5	6	7
Very easy			Neutral			Very difficult

2. How easy was it to use the adaptive pointing with zoom technique?

1	2	3	4	5	6	7
Very easy			Neutral			Very difficult

2. How easy was it to use the adaptive pointing with zoom technique?

1	2	3	4	5	6	7
Very easy			Neutral			Very difficult

Adaptive pointing comparison

For the adaptive pointing techniques (the last two you performed), how do you rate the zoom feature as opposed to when zoom was not present?

•--------------•--------------•--------------•-------------•-------------•-------------•

1	2	3	4	5	6	7
Much worse with zoom			About the same			Much better with zoom

Participant #:
Date:

Post-Experiment Questionnaire

Please complete the following questions. Select a value that best describe your experience from 1 to 7 as directed.

1. How much arm/wrist fatigue did you experience while performing the trials?

•-------------•-------------•-------------•-------------•-------------•-------------•

1	**2**	**3**	**4**	**5**	**6**	**7**
No fatigue at all			A little fatigue			A lot of fatigue

2. How much back fatigue did you experience while performing the trials?

•-------------•-------------•-------------•-------------•-------------•-------------•

1	**2**	**3**	**4**	**5**	**6**	**7**
No fatigue at all			A little fatigue			A lot of fatigue

3. How much leg fatigue did you experience while performing the trials?

•-------------•-------------•-------------•-------------•-------------•-------------•

1	**2**	**3**	**4**	**5**	**6**	**7**
No fatigue at all			A little fatigue			A lot of fatigue

4. Please rate the techniques you used, based on your preference, 1 being most preferred technique and 6 being the least preferred technique?

( ) Raycasting
( ) Raycasting with zoom
( ) Adaptive Pointing
( ) Adaptive Pointing with zoom
( ) Relavtive Pointing
( ) Relative Pointing with zoom

5. Feel free to add any comments or suggestions.

Appendix C

Appendix: Experimental Documents for Selection by Progressive Refinement Study

VirginiaTech

Office of Research Compliance
Institutional Review Board
2000 Kraft Drive, Suite 2000 (0497)
Blacksburg, Virginia 24060
540/231-4606 Fax 540/231-0959
e-mail irb@vt.edu
Website: www.irb.vt.edu

MEMORANDUM

DATE: July 21, 2010

TO: Doug A. Bowman, Felipe Araujo e Silva, Regis Kopper

FROM: Virginia Tech Institutional Review Board (FWA00000572, expires June 13, 2011)

PROTOCOL TITLE: Comparison of Single Precise Selection and Multiple Coarse Selections in Distal Pointing Tasks

IRB NUMBER: 10-617

Effective July 21, 2010, the Virginia Tech IRB Chair, Dr. David M. Moore, approved the new protocol for the above-mentioned research protocol.

This approval provides permission to begin the human subject activities outlined in the IRB-approved protocol and supporting documents.

Plans to deviate from the approved protocol and/or supporting documents must be submitted to the IRB as an amendment request and approved by the IRB prior to the implementation of any changes, regardless of how minor, except where necessary to eliminate apparent immediate hazards to the subjects. Report promptly to the IRB any injuries or other unanticipated or adverse events involving risks or harms to human research subjects or others.

All investigators (listed above) are required to comply with the researcher requirements outlined at http://www.irb.vt.edu/pages/responsibilities.htm (please review before the commencement of your research).

PROTOCOL INFORMATION:
Approved as: **Expedited, under 45 CFR 46.110 category(ies) 7**
Protocol Approval Date: **7/21/2010**
Protocol Expiration Date: **7/20/2011**
Continuing Review Due Date*: **7/6/2011**
*Date a Continuing Review application is due to the IRB office if human subject activities covered under this protocol, including data analysis, are to continue beyond the Protocol Expiration Date.

FEDERALLY FUNDED RESEARCH REQUIREMENTS:
Per federally regulations, 45 CFR 46.103(f), the IRB is required to compare all federally funded grant proposals / work statements to the IRB protocol(s) which cover the human research activities included in the proposal / work statement before funds are released. Note that this requirement does not apply to Exempt and Interim IRB protocols, or grants for which VT is not the primary awardee.

The table on the following page indicates whether grant proposals are related to this IRB protocol, and which of the listed proposals, if any, have been compared to this IRB protocol, if required.

Date*	OSP Number	Sponsor	Grant Comparison Conducted?

*Date this proposal number was compared, assessed as not requiring comparison, or comparison information was revised.

If this IRB protocol is to cover any other grant proposals, please contact the IRB office (irbadmin@vt.edu) immediately.

cc: File

Dear students:

We are looking for voluntary participants for our project "Comparison of single precise selection and multiple coarse selections in distal pointing tasks." You are eligible to participate if you are older than 18 years, color sighted and right handed.

By participating in this study, you will have a chance to experience state of the art input and display devices, and you will help us gather both quantitative and qualitative insights into user interface design for selection techniques in a variety of environments, from home entertainment systems to virtual reality applications.

The study will consist of performing a set of object selection tasks on a projection screen. More details will be provided on site. The whole experiment will last for about 1 hour.

Please respond to kopper@vt.edu and fbacim@vt.edu if you wish to participate.

Informed Consent for Participant of Investigative Project

Virginia Polytechnic Institute and State University

Title of Project: **Comparison of single precise selection and multiple coarse selections in distal pointing tasks**

Investigators: Dr. Doug Bowman, Regis Kopper, Felipe Bacim.

I. THE PURPOSE OF THIS RESEARCH/PROJECT

You are invited to participate in a study to compare two distal pointing selection techniques. This research studies how a single precise selection compares to multiple coarse selections using distal pointing.

II. PROCEDURES

Upon arrival, you will be given an introduction of our experiment background, goals, facilities, and study procedures. After reading and signing this informed consent form, you will be given a short period of time to gain practice in distal pointing selection tasks at a projection screen, with varied target sizes and distractor densities, and familiarize yourself with input and output devices. After practicing the tasks, we will load the computer program and read you a set of predefined tasks. You are expected to finish each task as soon as possible without sacrificing accuracy. We will record the time you take to complete each task. Finally, you will fill out a post-experiment questionnaire and rate your experience with different interaction techniques. A free form interview is conducted if you have any additional comments not addressed by the questionnaire.

III. RISKS

There will not be more than minimal risks by involving in our study.

IV. BENEFITS OF THIS PROJECT

Your participation in this project will provide information that may be used to improve the user interface design for applications that make use of distal pointing selections. No guarantee of benefits has been made to encourage you to participate. You may receive a synopsis summarizing this research when completed.

V. EXTENT OF ANONYMITY AND CONFIDENTIALITY

The results of this study will be kept confidential. Your written consent is required for the researchers to release any data identified with you as an individual to anyone other than personnel working on the project. The information you provide will have your name removed and only a participant number will identify you during analyses and any written reports of the research.

VI. COMPENSATION

Your participation is voluntary and unpaid.

VII. FREEDOM TO WITHDRAW

You are free to withdraw from this study at any time for any reason.

VIII. APPROVAL OF RESEARCH

This research has been approved, as required, by the Institutional Review Board for projects involving human subjects at Virginia Polytechnic Institute and State University, and by the Department of Computer Science.

IX. SUBJECT'S RESPONSIBILITIES AND PERMISSION

I voluntarily agree to participate in this study, and I know of no reason I cannot participate. I have read and understand the informed consent and conditions of this project. I have had all my questions answered. I hereby acknowledge the above and give my voluntary consent for participation in this project. If I participate, I may withdraw at any time without penalty. I agree to abide by the rules of this project

____ ____________
Signature Date

____ _______________________
Name (please print) Contact: phone or address or

____
Email address (OPTIONAL)

________ Not consent to be videotaped

________ Consent to be videotaped

Should I have any questions about this research or its conduct, I may contact:

Investigators: Dr. Doug A. Bowman Phone (540) 231-2058
Professor, Computer Science Department (231-6931)
email: bowman@vt.edu

Regis Kopper.
Graduate student, Computer Science Department
email: kopper@vt.edu

Felipe Bacim.
Graduate student, Computer Science Department
email: fbacim@vt.edu

Review Board: Dr. David Moore Phone (540) 231-4991
Chair, Virginia Tech Institutional Review Board
For the Protection of Human Subjects
Email: moored@vt.edu

cc: the participant, Dr. Bowman, Regis Kopper, Felipe Bacim.

Participant #:
Date:

Background Questionnaire

Please help us to categorize our user population by completing the following items.

Gender (circle one): Male Female

Age: _____________

Do you wear glasses or contact lenses (circle one)? No Glasses Contact Lenses

Are you (circle one) Right-handed Left-handed Ambidextrous

Occupation (if student, indicate graduate or undergraduate):

____

Major / Area of specialization (if student): _________________________________

Rate your tiredness level today: (circle one)

•---------------------------•---------------------------•---------------------------•

very tired somewhat tired a little tired not tired at all

Do you have any experience with playing the Wii, Playstation Move or any other video game that uses motion tracking? Yes No

If yes, how many hours per week, on average, do you play?

•---------------•---------------•---------------•---------------•---------------•---------------•

less than 1 1 5 10 15 20 more then 20

Do you have any experience with playing video games other than the ones described above?
Yes No
If yes, how many hours per week, on average, do you play?

•---------------•---------------•---------------•---------------•---------------•---------------•

less than 1 1 5 10 15 20 more then 20

(continue behind)

What type of video games do you usually play (select all that apply)?

- First-person shooting games (e.g. Half Life, Quake, etc…)
- Sports games (e.g. Need for Speed, Madden, etc…)
- Massive multiplayer online games (e.g. World of Warcraft)
- Real-time strategy games (e.g. Starcraft)
- Mobile devices casual games (e.g. Iphone, Ipad, etc…)
- Motion capture casual games (e.g. Wii Sports, Wii Fit, …)
- Social games (e.g. Farmville, Mafia wars, …)
- Third-person action games (e.g. Grand Theft Auto, Gears of War, etc…)
- Other: _____________

Do you have any experience with virtual environments / virtual reality? Yes No

If yes, please explain:

Sphere-Casting Quad Menu versus Raycasting evaluation

Procedural Overview

Participants 1, 9 and 17

Techniques:
- Sphere-casting + Quad menu (SCQM)
- Raycasting (RC)

Tracking:
- Good (GT)
- Bad (BT)

Conditions:
- A: SCQM+GT
- B: SCQM+BT
- C: RC+GT
- D: RC+BT

This session: ABCD

1. Consent Form
2. Pre-screening:
    a) Color-blindness test
    b) Confirm right-handedness
3. Background Survey
4. Initial Instructions
5. For each technique in the order that they are presented:
    a) Tutorial
    b) Trials
    c) Post trials questionnaire
6. Exit Survey & Interview

Initial Instructions
[run Px.bat]

Welcome and thank you for your participation in our experiment. You will help us evaluate selection techniques for pointing at a projection screen.

In this experiment, you will perform several selections of targets on the screen. The targets' size and the number of distractors will vary.

The goal of this study is to evaluate various techniques for pointing at objects on the screen from a distance. In order for the techniques to be comparable, you need to perform the trials as quickly as possible, making as few mistakes as you can. We need you to do your best. Remember that we're not evaluating you, but the techniques, so do not feel pressured, but always try to do your best.

(*show input device*) This is the pointing device. Hold it with your right hand. The only button you will use is the trigger. Hold it in a comfortable position and use your index finger to click.

Now, stand at this point, such that the pointing device remains roughly over the mark on the ground. To complete the experiment, you will NOT be allowed to reach out with your arm or use your other hand for steadiness. You may have your arm relaxed and move your forearm to complete the tasks, just please do not extend and reach out with your arm.

[press i]

Imagine the line that extends from the device coming out of here (*point at tip of wand*) towards the display. The position of the cursor is determined by the intersection of this ray with the screen. The crosshair | blue circle] represents the cursor. Go ahead and move the cursor on the display to familiarize yourself.

Tutorial for SCQM:
[press n]

Now, I will explain how to complete the tasks for the first technique.

Before the start of each trial, you should locate the red target on the screen. Once you see it and are ready to begin the task, point at the large yellow object in the center of the screen with the input device and click to start the task. Notice that the display is now filled with grey objects indicating that the task has started.

Now, point and click on the blue circle anywhere on the display such that the red target *is not* inside it. Notice the error message. You should always have the red target *touching or inside* the blue circle to complete the task.

Notice that a different target is shown. You will always move on to the next task, even if you make a mistake. When you make a mistake, you will be able to try again at the end of the set of tasks.

Now, when you're ready to begin the task, click on the yellow object then point and click on the circle such that the target is touching or inside the blue circle. Notice the quad menu. This menu contains all the objects that were inside the blue circle when you clicked. The idea is to use the menu to narrow down the number of objects until you select the red target that you're after. For now, point at any quadrant that does *not* contain the red target and click. Notice the error message. You always need to click on the quadrant that contains the red object to complete the task.

Again, when you're ready to begin the task, do so by clicking on the yellow object and pointing the blue circle at the red target. Now, click *anywhere* in the quadrant that contains the red target. It is important to remember that you *do not* need to click *on* the target, but anywhere in the quadrant that contains it. That will help you be faster. Notice the new arrangement of quadrants. Go ahead and select the correct quadrant until the task finishes.

Notice the new task set up in the display. That means you completed the last task successfully.

[press n]
[press i]

Now you will practice the technique you just learned for a little while. Remember to perform as fast as you can. Go ahead and start when you're ready.

[press i]

Notice that the cursor is less stable and contains more jitter. This is normal, and is part of the experiment. Try to perform your best despite the jitter.

This is still practice. Go ahead and start when you're ready.

Session for SCQM:

Now you will perform eight sets of nine tasks with the technique you just tried.

[press i]

Notice that the cursor is more stable now.

Wait until I finish reading the instructions and complete the tasks as you were practicing them until the application finishes. If you need a break for any reason, please do so after you finish a task, when you see the yellow object in the center of the screen. Remember to be as fast as you can, but try not to make mistakes.

You may begin when you're ready.

[Waiting for instructions . . .]

Take a break and relax for a while. Let me know when you're ready to continue.
If participant rests for two minutes, call it time

You will now perform another eight sets of nine tasks with the technique you just tried.

Again, notice that the cursor is less stable and contains more jitter. Remember that this is normal, and part of the experiment. Try to perform your best despite the jitter.

After I finish reading the instructions, complete the tasks as you were practicing them until the application finishes. Remember to be as fast as you can, but without making any mistakes.

You may begin when you're ready.

[waiting for instructions]

[Give SCQM ratings questionnaire]

Take a break and let me know when you're good to continue.

Tutorial for Raycasting:
[press i]

Now, I will explain how to complete the tasks for the next technique.

Before the start of each trial, you should locate the red target on the screen. Once you see it and are ready to begin the task, point at the large yellow object in the center of the screen with the input device and click to start the task. Notice that the display is now filled with grey objects indicating that the task has started.

Now, point the cursor anywhere at the display *except* the red target and click the button. Notice the error message. You should always click with the center of the cursor *inside* the target to complete the task.

Notice that a different target is shown. You will always move on to the next task, even if you make a mistake. When you make a mistake, you will be able to try again at the end of the set of tasks.

Once you're ready to begin the task, click on the yellow object and select the target to complete the task. Notice there is no error message. That means the trial was successful.

[press n]

Notice that the cursor is less stable and contains more jitter. This is normal, and is part of the experiment. Try to perform your best despite the jitter.

Now you will practice the technique you just learned for a little while. Remember to perform as fast as you can. Go ahead and start when you're ready.

[press i]

Notice that the cursor is more stable now.

This is still practice. Go ahead and start when you're ready.

Session for Raycasting:

Now you will perform eight sets of nine tasks with the technique you just tried.

[press i]

Wait until I finish reading the instructions and complete the tasks as you were practicing them until the application finishes. If you need a break for any reason, please do so after you finish a task, when you see the yellow object in the center of the screen. Remember to be as fast as you can, but try not to make mistakes.

You may begin when you're ready.

[Waiting for instructions . . .]

Take a break and relax for a while. Let me know when you're ready to continue.
If participant rests for two minutes, call it time

You will now perform another eight sets of nine tasks with the technique you just tried.

[waiting for instructions]
[press i]

Again, notice that the cursor is less stable and contains more jitter. Remember that this is normal, and part of the experiment. Try to perform your best despite the jitter.

After I finish reading the instructions, complete the tasks as you were practicing them until the application finishes. Remember to be as fast as you can, but without making any mistakes.

You may begin when you're ready.

[waiting for instructions]

Pleaase rate the technique you just tried independently of the one you performed first.
[Give Raycasting ratings questionnaire.]

Please fill in this post-experiment questionnaire.
[give post-experiment questionnaire]

Sphere-Casting Quad Menu versus Raycasting evaluation

Quick reference guide

- You will learn, practice and perform tasks with 2 distal pointing techniques.
- For each of the two techniques, you will perform 2 sets of 90 seconds as practice, and 2 sets of 72 trials in the experimental session.
- Be as fast as you can, while not making mistakes.

Participant #:
Date:

Raycasting ratings

Please answer the following questions based on the technique you just performed. Select a value that best describe your experience from 1 to 7 as directed.

1. How easy was it to learn the technique you just used?

1	2	3	4	5	6	7
Very easy			Neutral			Very difficult

2. How easy was it to use the technique to perform the tasks you just did?

1	2	3	4	5	6	7
Very easy			Neutral			Very difficult

3. Comparing to the tasks with the smooth cursor, how much harder was it to complete the tasks when the cursor was jittery?

1	2	3	4	5	6	7
Not harder			A little harder			Much harder

4. Comparing to the tasks with larger targets, how much harder was it to complete the tasks when the targets were very small?

1	2	3	4	5	6	7
Not harder			A little harder			Much harder

5. Comparing to the tasks with fewer distractors (grey objects), how much harder was it to complete the tasks that contained a lot of distractor objects?

1	2	3	4	5	6	7
Not harder			A little harder			Much harder

(continue on the back)

6. How much arm/wrist fatigue did you experience while performing the trials? (1 being very low, 7 being very high):

1	2	3	4	5	6	7
No fatigue at all			A little fatigue			A lot of fatigue

7. How much back fatigue did you experience while performing the trials? (1 being very low, 7 being very high):

1	2	3	4	5	6	7
No fatigue at all			A little fatigue			A lot of fatigue

8. How much leg fatigue did you experience while performing the trials? (1 being very low, 7 being very high):

1	2	3	4	5	6	7
No fatigue at all			A little fatigue			A lot of fatigue

Participant #:
Date:

Sphere-casting plus quad-menu ratings

Please answer the following questions based on the technique you just performed. Select a value that best describe your experience from 1 to 7 as directed.

1. How easy was it to learn the technique you just used?

1	2	3	4	5	6	7
Very easy			Neutral			Very difficult

2. How easy was it to use the technique to perform the tasks you just did?

1	2	3	4	5	6	7
Very easy			Neutral			Very difficult

3. Comparing to the tasks with the smooth cursor, how much harder was it to complete the tasks when the cursor was jittery?

1	2	3	4	5	6	7
Not harder			A little harder			Much harder

4. Comparing to the tasks with larger targets, how much harder was it to complete the tasks when the targets were very small?

1	2	3	4	5	6	7
Not harder			A little harder			Much harder

5. Comparing to the tasks with fewer distractors (grey objects), how much harder was it to complete the tasks that contained a lot of distractor objects?

1	2	3	4	5	6	7
Not harder			A little harder			Much harder

(continue on the back)

6. How much arm/wrist fatigue did you experience while performing the trials? (1 being very low, 7 being very high):

•--------------•--------------•--------------•-------------•-------------•-------------•

1	**2**	**3**	**4**	**5**	**6**	**7**
No fatigue at all			A little fatigue			A lot of fatigue

7. How much back fatigue did you experience while performing the trials? (1 being very low, 7 being very high):

•--------------•--------------•--------------•-------------•-------------•-------------•

1	**2**	**3**	**4**	**5**	**6**	**7**
No fatigue at all			A little fatigue			A lot of fatigue

8. How much leg fatigue did you experience while performing the trials? (1 being very low, 7 being very high):

•--------------•--------------•--------------•-------------•-------------•-------------•

1	**2**	**3**	**4**	**5**	**6**	**7**
No fatigue at all			A little fatigue			A lot of fatigue

Participant #:
Date:

Post-Experiment Questionnaire

1. Which technique did you prefer, overall?

Precise single selection   Coarse multiple selections

Why?

2. When the cursor was jittery, which technique did you prefer?

Precise single selection   Coarse multiple selections

Why?

3. When the targets were very small, which technique did you prefer?

Precise single selection   Coarse multiple selections

Why?

4. When there were a lot of distractor objects, which technique did you prefer?

Precise single selection   Coarse multiple selections

Why?

5. Feel free to add any comments or suggestions.

Bibliography

Accot, J. and Zhai, S. (1997). Beyond Fitts' law: models for trajectory-based HCI tasks. In *CHI '97: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 295–302, New York, NY, USA. ACM Press.

Accot, J. and Zhai, S. (1999). Performance evaluation of input devices in trajectory-based tasks: an application of the steering law. In *CHI '99: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 466–472, New York, NY, USA. ACM.

Accot, J. and Zhai, S. (2003). Refining fitts' law models for bivariate pointing. In *CHI '03: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 193–200, New York, NY, USA. ACM Press.

Andrews, C., Endert, A., and North, C. (2010). Space to think: large high-resolution displays for sensemaking. In *Proceedings of the 28th international conference on Human factors in computing systems*, CHI '10, pages 55–64, New York, NY, USA. ACM.

Badillo, B., Bowman, D. A., McConnel, W., Tao, T. N., and da Silva, M. G. (2006). Literature survey on interaction techniques for large displays. Technical Report TR-06-21, Computer Science, Virginia Tech.

Bakker, N., Werkhoven, P., and Passenier, P. (1998). Aiding orientation performance in virtual environments with proprioceptive feedback. In *Virtual Reality Annual International Symposium, 1998. Proceedings., IEEE 1998*, pages 28 –33.

Ball, R. and North, C. (2005). Effects of tiled high-resolution display on basic visualization and navigation tasks. In *CHI '05: CHI '05 extended abstracts on Human factors in computing systems*, pages 1196–1199, New York, NY, USA. ACM Press.

Ball, R., North, C., and Bowman, D. A. (2007). Move to improve: promoting physical navigation to increase user performance with large displays. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 191–200, New York, NY, USA. ACM.

Bangerter, A. and Oppenheimer, D. (2006). Accuracy in detecting referents of pointing gestures unaccompanied by language. *Gesture*, 6(1):85–102.

Baudisch, P. (2006). Interacting with large displays. *Computer*, 39(4):96–97.

Baudisch, P., Cutrell, E., Czerwinski, M., Robbins, D. C., Tandler, P., Bederson, B. B., and Zierlinger, A. (2003a). Drag-and-pop and drag-and-pick: Techniques for accessing remote screen content on touch- and pen-operated systems. In *INTERACT*.

Baudisch, P., Cutrell, E., and Robertson, G. G. (2003b). High-density cursor: a visualization technique that helps users keep track of fast-moving mouse cursors. In *INTERACT*.

Beardsley, P., Baar, J. V., Raskar, R., and Forlines, C. (2005). Interaction using a handheld projector. *IEEE Comput. Graph. Appl.*, 25(1):39–43.

Bezerianos, A. and Balakrishnan, R. (2005). The vacuum: facilitating the manipulation of distant objects. In *CHI '05: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 361–370, New York, NY, USA. ACM Press.

Bi, X., Bae, S.-H., and Balakrishnan, R. (2010). Effects of interior bezels of tiled-monitor large displays on visual search, tunnel steering, and target selection. In *Proceedings of the 28th international conference on Human factors in computing systems*, CHI '10, pages 65–74, New York, NY, USA. ACM.

Bowman, D., Wingrave, C., Campbell, J., Ly, V., and Rhoton, C. (2002). Novel uses of pinch gloves™ for virtual environment interaction techniques. *Virtual Reality*, 6(3):122–129.

Bowman, D. A. and Hodges, L. F. (1997). An evaluation of techniques for grabbing and manipulating remote objects in immersive virtual environments. In *SI3D '97: Proceedings of the 1997 symposium on Interactive 3D graphics*, pages 35–ff., New York, NY, USA. ACM.

Bowman, D. A., Johnson, D. B., and Hodges, L. F. (1999). Testbed evaluation of virtual environment interaction techniques. In *Proceedings of the ACM symposium on Virtual reality software and technology*, VRST '99, pages 26–33, New York, NY, USA. ACM.

Bowman, D. A., Kruijff, E., Jr., J. J. L., and Poupyrev, I. (2004). *3D User Interfaces: Theory and Practice*. Addison-Wesley.

Buck, L. (1980). Motor performance in relation to control-display gain and target width. *Ergonomics*, 23:579–589.

Cao, X. and Balakrishnan, R. (2003). Visionwand: interaction techniques for large displays using a passive wand tracked in 3d. In *UIST '03: Proceedings of the 16th annual ACM symposium on User interface software and technology*, pages 173–182, New York, NY, USA. ACM Press.

Card, S. K., English, W. K., and Burr, B. J. (1978). Evaluation of mouse, rate-controlled isometric joystick, step keys, and text keys, for text selection on a crt. *Ergonomics*, 21.

Carl, J. R. and Gellman, R. S. (1987). Human smooth pursuit: stimulus-dependent responses. *Journal of Neurophysiology*, 57(5):1446–1463.

Carlton, L. (1981). Processing visual feedback information for movement control. *Journal of Experimental Psychology: Human Perception and Performance*, 7(5):1019.

Carroll, J. M. and Campbell, R. L. (1986). Softening up hard science: Reply to newell and card. *Human-Computer Interaction*, 2:227–249.

Chance, S., Gaunet, F., Beall, A., and Loomis, J. (1998). Locomotion mode affects the updating of objects encountered during travel: The contribution of vestibular and proprioceptive inputs to path integration. *Presence*, 7(2):168–178.

Connolly, J. D. and Goodale, M. A. (1999). The role of visual feedback of hand position in the control of manual prehension. *Experimental Brain Research*, 125:281–286. 10.1007/s002210050684.

Crossman, E. R. F. W. (1956). *The measurement of perceptual load in manual operations*. PhD thesis, University of Birmingham.

Dan R. Olsen, J. and Nielsen, T. (2001). Laser pointer interaction. In *CHI '01: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 17–22, New York, NY, USA. ACM Press.

Dasiyici, M. C. (2008). Multi-scale cursor: Optimizing mouse interaction for large personal workspaces. Master's thesis, Virginia Tech.

Davis, J. and Chen, X. (2002). Lumipoint: multi-user laser-based interaction on large tiled displays. *Displays*, 23(5):205–211.

Field, A. (2009). *Discovering Statistics Using SPSS*. Introducing Statistical Methods Series. SAGE.

Figueroa, P., Kitamura, Y., Kuntz, S., Vanacken, L., Maesen, S., Weyer, T. D., Notelaers, S., Octavia, J. R., Beznosyk, A., Coninx, K., Bacim, F., Kopper, R., Leal, A., Ni, T., and Bowman, D. A. (2010). 3dui 2010 contest grand prize winners. *IEEE Computer Graphics and Applications*, 30:86–96.

Fitts, P. M. (1954). The information capacity of the human motor system in controlling the amplitude of movement. *Journal of experimental psychology*, 47:381–391.

Fleishman, E. and Hempel Jr, W. (1956). Factorial analysis of complex psychomotor performance and related skills. *Journal of applied psychology*, 40(2):96–104.

Forlines, C., Balakrishnan, R., Beardsley, P., van Baar, J., and Raskar, R. (2005). Zoom-and-pick: facilitating visual zooming and precision pointing with interactive handheld projectors. In *UIST '05: Proceedings of the 18th annual ACM symposium on User interface software and technology*, pages 73–82, New York, NY, USA. ACM Press.

Forlines, C., Vogel, D., and Balakrishnan, R. (2006). Hybridpointing: fluid switching between absolute and relative pointing with a direct input device. In *UIST '06: Proceedings of the 19th annual ACM symposium on User interface software and technology*, pages 211–220, New York, NY, USA. ACM.

Foulkes, A. J. M. and Miall, R. C. (2000). Adaptation to visual feedback delays in a human manual tracking tasks. *Experimental Brain Research*, 131(1):101–110. 10.1007/s002219900286.

Frees, S., Kessler, G. D., and Kay, E. (2007). PRISM interaction for enhancing control in immersive virtual environments. *ACM Trans. Comput.-Hum. Interact.*, 14(1):2.

Gibbs, C. (1962). Controller design: Interactions of controlling limbs, time-lags and gain in positional and velocity systems. *Ergonomics*, 5(2):385–403.

Grossman, T. (2008). *Interaction with Volumetric Displays*. PhD thesis, University of Toronto.

Grossman, T. and Balakrishnan, R. (2004). Pointing at trivariate targets in 3d environments. In *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 447–454, New York, NY, USA. ACM Press.

Grossman, T. and Balakrishnan, R. (2005a). The bubble cursor: enhancing target acquisition by dynamic resizing of the cursor's activation area. In *CHI '05: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 281–290, New York, NY, USA. ACM Press.

Grossman, T. and Balakrishnan, R. (2005b). A probabilistic approach to modeling two-dimensional pointing. *ACM Trans. Comput.-Hum. Interact.*, 12(3):435–459.

Grossman, T. and Balakrishnan, R. (2006). The design and evaluation of selection techniques for 3d volumetric displays. In *UIST '06: Proceedings of the 19th annual ACM symposium on User interface software and technology*, pages 3–12, New York, NY, USA. ACM.

Guiard, Y., Blanch, R., and Beaudouin-Lafon, M. (2004). Object pointing: a complement to bitmap pointing in guis. In *GI '04: Proceedings of the 2004 conference on Graphics interface*, pages 9–16, School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada. Canadian Human-Computer Communications Society.

Guilford, J. (1958). A system of the psychomotor abilities. *The American Journal of Psychology*, 71(1):164–174.

Haan, G. D., Koutek, M., and Post, F. H. (2005). Intenselect: Using dynamic object rating for assisting 3d object selection. In *In Virtual Environments 2005*, pages 201–209.

Hoffmann, E. and Sheikh, I. (1994). Effect of varying target height in a fitts' movement task. *Ergonomics*, 37(6):1071–1088.

Jiang, H., Ofek, E., Moraveji, N., and Shi, Y. (2006). Direct pointer: direct manipulation for large-display interaction using handheld cameras. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 1107–1110, New York, NY, USA. ACM Press.

Jr, M. W. C. and Fullenkamp, S. C. (1988). Perceived contrast and stimulus size: Experiment and simulation. *Vision Research*, 28(6):695 – 709.

Kondraske, G. (1994). An angular motion fitt's law for human performance modeling and prediction. *Engineering in Medicine and Biology Society, 1994. Engineering Advances: New Opportunities for Biomedical Engineers. Proceedings of the 16th Annual International Conference of the IEEE*, pages 307–308 vol.1.

König, W., Gerken, J., Dierdorf, S., and Reiterer, H. (2009). Adaptive pointing – design and evaluation of a precision enhancing technique for absolute pointing devices. In Gross, T., Gulliksen, J., Kotzé, P., Oestreicher, L., Palanque, P., Prates, R., and Winckler, M., editors, *Human-Computer Interaction – INTERACT 2009*, volume 5726 of *Lecture Notes in Computer Science*, pages 658–671. Springer Berlin / Heidelberg.

Kopper, R., Bacim, F., and Bowman, D. (2011). Rapid and accurate 3d selection by progressive refinement. In *3D User Interfaces (3DUI), 2011 IEEE Symposium on*, pages 67 –74.

Kopper, R., Bowman, D. A., Silva, M. G., and McMahan, R. P. (2010). A human motor behavior model for distal pointing tasks. *International Journal of Human-Computer Studies*, 68(10):603 – 615.

Kopper, R., Silva, M. G., McMahan, R. P., and Bowman, D. A. (2008). Increasing the precision of distant pointing for large high-resolution displays. Technical Report TR-08-17, Virginia Tech.

Liang, J. and Green, M. (1993). Jdcad: A highly interactive 3d modeling system. In *3rd International Conference on CAD and Computer Graphics*, pages 217–222.

Liu, L., Nieuwenhuizen, C., and Martens, J. (2009). Comparing aimed movements in the real world and in virtual reality. In *Proceedings of the 2009 IEEE Virtual Reality Conference*, pages 219–222. IEEE Computer Society Washington, DC, USA.

Lucas, J. (2005). Design and evaluation of 3d multiple object selection techniques, vol. m.s. Master's thesis, Virginia Polytechnic Institute and State University, Blacksburg, VA, USA.

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

MacKenzie, I. S. (2003). Motor behavior models for human-computer interaction. In Carroll, J. M., editor, *HCI Models, Theories, and Frameworks - Toward a Multidisciplinary Science*. Morgan Kaufman.

Mackenzie, S. I. and Buxton, W. (1992). Extending fitts' law to two-dimensional tasks. In *CHI '92: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 219–226, New York, NY, USA. ACM Press.

Malik, S., Ranjan, A., and Balakrishnan, R. (2005). Interacting with large displays from a distance with vision-tracked multi-finger gestural input. In *UIST '05: Proceedings of the 18th annual ACM symposium on User interface software and technology*, pages 43–52, New York, NY, USA. ACM Press.

Matveyev, S. V. and Göbel, M. (2003). The optical tweezers: multiple-point interaction technique. In *VRST '03: Proceedings of the ACM symposium on Virtual reality software and technology*, pages 184–187, New York, NY, USA. ACM Press.

McGuffin, M. and Balakrishnan, R. (2002). Acquisition of expanding targets. In *CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 57–64, New York, NY, USA. ACM.

Mewhinney, M. and Fay, M. (2008). Nasa develops world's highest resolution visualization system [Press release]. Retrieved from http://www.nas.nasa.gov/News/Releases/2008/06-25-08.html.

Miall, R. (1996). Task-dependent changes in visual feedback control: A frequency analysis of human manual tracking. *Journal of Motor Behavior*, 28(2):125.

Mine, M. R. (1995). Virtual environment interaction techniques. Technical report, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA.

Mine, M. R., Frederick P. Brooks, J., and Sequin, C. H. (1997). Moving objects in space: exploiting proprioception in virtual-environment interaction. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 19–26, New York, NY, USA. ACM Press/Addison-Wesley Publishing Co.

Murata, A. and Iwase, H. (2001). Extending fitts' law to a three-dimensional pointing task. *Human Movement Science*, 20(6):791–805.

Myers, B. A., Bhatnagar, R., Nichols, J., Peck, C. H., Kong, D., Miller, R., and Long, A. C. (2002a). Interacting at a distance: measuring the performance of laser pointers and other devices. In *CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 33–40, New York, NY, USA. ACM Press.

Myers, B. A., Bhatnagar, R., Nichols, J., Peck, C. H., Kong, D., Miller, R., and Long, A. C. (2002b). Interacting at a distance: measuring the performance of laser pointers and other devices. In *CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 33–40, New York, NY, USA. ACM.

Newell, A. and Card, S. (1985). The prospects for psychological science in human-computer interaction. *Human-Computer Interaction*, 1(3):209–242.

Newell, A. and Card, S. (1986). Straightening out softening up: Response to carroll and campbell. *Human-Computer Interaction*, 2(3):251–267.

Ni, T., Schmidt, G. S., Staadt, O. G., Livingston, M. A., Ball, R., and May, R. (2006). A survey of large high-resolution display technologies, techniques, and applications. In *VR '06: Proceedings of the IEEE Virtual Reality Conference (VR 2006)*, page 31, Washington, DC, USA. IEEE Computer Society.

Nintendo Wii (2011). Retrieved from http://www.nintendo.com/wii. Last accessed on 05/20/2011.

Patel, S. and Abowd, G. (2003). A 2-way laser-assisted selection scheme for handhelds in a physical environment. *LECTURE NOTES IN COMPUTER SCIENCE*, pages 200–207.

Pausch, R., Proffitt, D., and Williams, G. (1997). Quantifying immersion in virtual reality. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '97, pages 13–18, New York, NY, USA. ACM Press/Addison-Wesley Publishing Co.

Peck, S. M., North, C., and Bowman, D. (2009). A multiscale interaction technique for large, high-resolution displays. In *3D User Interfaces, 2009. 3DUI 2009. IEEE Symposium on*, pages 31–38.

Pfeiffer, T., Latoschik, M., and Wachsmuth, I. (2008). Conversational pointing gestures for virtual reality interaction: Implications from an empirical study. In *Virtual Reality Conference, 2008. VR '08. IEEE*, pages 281 –282.

Phidgets (2011). Retrieved from http://www.phidgets.com/. Last accessed on 05/03/2011.

Pierce, J. S., Forsberg, A. S., Conway, M. J., Hong, S., Zeleznik, R. C., and Mine, M. R. (1997). Image plane interaction techniques in 3d immersive environments. In *SI3D '97: Proceedings of the 1997 symposium on Interactive 3D graphics*, pages 39–ff., New York, NY, USA. ACM.

Playstation Move (2011). Retrieved from http://us.playstation.com/ps3/playstation-move/. Last accessed on 05/20/2011.

Po, B. A. (2005). *A Representational Basis for Human-Computer Interaction*. PhD thesis, The University of British Columbia.

Ramos, G., Cockburn, A., Balakrishnan, R., and Beaudouin-Lafon, M. (2007). Pointing lenses: facilitating stylus input through visual-and motor-space magnification. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 757–766, New York, NY, USA. ACM.

Riviere, C., Rader, R., and Khosla, P. (1997). Characteristics of hand motion of eye surgeons. *Engineering in Medicine and Biology society, 1997. Proceedings of the 19th Annual International Conference of the IEEE*, 4.

Seifert, L. and Chollet, D. (2008). Inter-limb coordination and constraints in swimming: a review. *Physical Activity and Children: New Research*, page 65.

Shannon, C. and Weaver, W. (1949). *The Mathematical Theory of Communication*. University of Illinois Press, Urbana, Illinois.

Simon, R., Frohlich, P., and Grechenig, T. (2008). Geopointing: evaluating the performance of orientation-aware location-based interaction under real-world conditions. *Journal of Location Based Services*, 2(1):24 – 40.

Steed, A. (2006). Towards a general model for selection in virtual environments. In *3D User Interfaces, 2006. 3DUI 2006. IEEE Symposium on*, pages 103–110. IEEE.

Steed, A. and Parker, C. (2004). 3d selection strategies for head tracked and non-head tracked operation of spatially immersive displays. In *8th International Immersive Projection Technology Workshop*, pages 13–14.

Steed, A. and Parker, C. (2005). Evaluating effectiveness of interaction techniques across immersive virtual environmental systems. *Presence: Teleoper. Virtual Environ.*, 14:511–527.

Stefels, C. N., Kneissner, J., Aarnink, R. G., Kaufholz, P., Grimbergen, C. A., and Dankelman, J. (2007). Equipment control in a sterile environment using the gyromouse and a new interface, the user interface (ui) wand. *Minimally Invasive Therapy and Allied Technologies*, 16(3):163 – 172.

Stoakley, R., Conway, M. J., and Pausch, R. (1995). Virtual reality on a wim: interactive worlds in miniature. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '95, pages 265–272, New York, NY, USA. ACM Press/Addison-Wesley Publishing Co.

Tabachnick, B. G. and Fidell, L. S. (2006). *Using Multivariate Statistics (5th Edition)*. Allyn & Bacon, Inc., Needham Heights, MA, USA.

Tse, E., Hancock, M., and Greenberg, S. (2007). Speech-filtered bubble ray: Improving target acquisition on display walls. Technical Report 2007-867-19, Department of Computer Science, University of Calgary.

Vanacken, L., Grossman, T., and Coninx, K. (2007). Exploring the effects of environment density and target visibility on object selection in 3d virtual environments. *3D User Interfaces*, 0:null.

Vercher, J. L. and Gauthier, G. M. (1992). Oculo-manual coordination control: Ocular and manual tracking of visual targets with delayed visual feedback of the hand motion. *Experimental Brain Research*, 90:599–609. 10.1007/BF00230944.

Vetter, S., Jochems, N., Kausch, B., Mütze-Niewöhner, S., and Schlick, C. (2010). Age-induced change in visual acuity and its impact on performance in a target detection task with electronic information displays. *Occupational Ergonomics*, 9(2):99–110.

Vogel, D. and Balakrishnan, R. (2005). Distant freehand pointing and clicking on very large, high resolution displays. In *UIST '05: Proceedings of the 18th annual ACM symposium on User interface software and technology*, pages 33–42, New York, NY, USA. ACM Press.

Ware, C. and Balakrishnan, R. (1994). Reaching for objects in VR displays: lag and frame rate. *ACM Trans. Comput.-Hum. Interact.*, 1:331–356.

Welch, G. and Bishop, G. (1995). An introduction to the kalman filter. Technical Report 95-041, Department of Computer Science, University of North Carolina at Chapel Hill.

Wilkes, C. and Bowman, D. A. (2008). Advantages of velocity-based scaling for distant 3d manipulation. In *Proceedings of the 2008 ACM symposium on Virtual reality software and technology*, VRST '08, pages 23–29, New York, NY, USA. ACM.

Wilson, A. and Shafer, S. (2003). Xwand: Ui for intelligent spaces. In *CHI '03: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 545–552, New York, NY, USA. ACM.

Wingrave, C. and Bowman, D. (2005). Baseline factors for raycasting selection. In *Proceedings of HCI International*.

Wingrave, C. A., Bowman, D. A., and Ramakrishnan, N. (2002). Towards preferences in virtual environment interfaces. In *EGVE '02: Proceedings of the workshop on Virtual environments 2002*, pages 63–72, Aire-la-Ville, Switzerland, Switzerland. Eurographics Association.

Woodworth, R. (1899). *The accuracy of voluntary movement*. The Macmillan Company.

Worden, A., Walker, N., Bharat, K., and Hudson, S. (1997). Making computers easier for older adults to use: area cursors and sticky icons. In *CHI '97: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 266–271, New York, NY, USA. ACM Press.

Yost, B., Haciahmetoglu, Y., and North, C. (2007). Beyond visual acuity: the perceptual scalability of information visualizations for large displays. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '07, pages 101–110, New York, NY, USA. ACM.

Zhai, S., Kong, J., and Ren, X. (2004). Speed–accuracy tradeoff in fitts' law tasks–on the equivalency of actual and nominal pointing precision. *International Journal of Human-Computer Studies*, 61(6):823 – 856.

Zhai, S., Morimoto, C., and Ihde, S. (1999). Manual and gaze input cascaded (magic) pointing. *Proceedings of the SIGCHI conference on Human factors in computing systems: the CHI is the limit*, pages 246–253.

Zhai, S. and Woltjer, R. (2003). Human movement performance in relation to path constraint - the law of steering in locomotion. In *VR '03: Proceedings of the IEEE Virtual Reality 2003*, page 149, Washington, DC, USA. IEEE Computer Society.

Zhao, S., Agrawala, M., and Hinckley, K. (2006). Zone and polygon menus: using relative position to increase the breadth of multi-stroke marking menus. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, CHI '06, pages 1077–1086, New York, NY, USA. ACM.