A Multiscale Interaction Technique for Large, High-Resolution Displays

Sarah M. Peck

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

Master of Science
in
Computer Science and Applications

Dr. Chris L. North, Chair
Dr. Doug A. Bowman
Dr. Francis Quek

23 May 2008
Blacksburg, VA

Keywords: design, embodiment, evaluation, high-resolution, interaction technique, large tiled display, physical navigation

A Multiscale Interaction Technique for Large, High-Resolution Displays

Sarah M. Peck

Abstract

The decreasing price of displays has enabled exploration of ever-larger high-resolution displays. Previous research has shown that as the display grows larger, users prefer to physically navigate, which has proven benefits. However, increasing the display size so radically creates a new difficulty in interaction. The paradigm has changed from sitting at a desktop computer to taking users' physical navigation into account and designing more mobile interactions.

Currently, when users move, they change the scale at which they are viewing information without changing the interaction scale. This is a problem because tasks change at different levels of visual scale. Mulitscale interaction aims to exploit users' movement by linking it to interaction, changing the interaction scale depending on users' distance from the display.

This work accomplishes three things: first, we define the design space of multiscale interaction; secondly, through a case study, we explore the design issues for a specific area of the design space; lastly, we evaluate one application through a user study that compares it to two other interaction types. We wanted to know, do users in fact benefit from the linkage of physical navigation with interaction?

Results show a trend of a natural link between user distance and interaction scale, even with the other techniques that did not enforce this link. In addition, multiscale interaction benefits from the link by having more consistent performance. They also show that while participants using multiscale interaction tend to move more, they benefit from this additional movement, unlike with the other interaction types.

Acknowledgements

I would like to thank my advisor Dr. Chris North for his continued support and guidance throughout the difficult process of completing my thesis work. More than anything, he has taught me to have faith in my abilities. Additionally, I would like to thank the other members of my committee, Dr. Doug Bowman and Dr. Francis Quek for all the comments, suggestions, and additional support they have offered me along the way.

Big thanks go to the members of the Blacklab committee, Christopher Andrews, Ryan McMahan, and Chreston Miller. They have given of their time to make the lab equipment more easily available to everyone, including me.

I would also like to extend thanks to other members of my lab, both past and present who have assisted me in various ways. Thank you to Beth Yost and Lauren Shupp for being welcoming to a newcomer, and for the friendship that has followed. Thank you to Christopher Andrews for his initial collaboration with me on the Storyboard project, from which my thesis topic came, and his extensive python knowledge, on which I drew many a time. Thank you to Mike DellaNoce and Dave Machaj, whose multiscale gaming research I mention in my design space.

Finally, I would like to thank my family and friends who have supported me in many ways. Thanks to one of my best friends and my sister, Amanda, who has been a superb cheerleader, and who has reminded me on several occasions to thank God for all my blessings and to take it one day at a time. Thanks to Pardha Pyla for being a sounding board, for being a good listener in general, and last but not least for all the chocolate. Lastly, thanks to Jeremy Archuleta for his invaluable emotional support.

Table of Contents

List of Figures

List of Figures

List of Tables

Chapter 1: Introduction

1.1 Motivation

The decreasing price of displays has enabled the exploration of ever-larger and higher-resolution displays. Previous research has quantified benefits from both the increased size and the increased resolution. Several studies have shown that with large datasets, such as those found in geospatial analysis, the larger viewport size improves users' performance time, as well as decreases their frustration [4, 5, 33]. Another study looked at increasing the amount of information on such displays until the information was beyond visual acuity, showing that even though users cannot see all the information at once, they are still more efficient and accurate because of additional data that can be displayed [40].

In addition, these studies also found a correlation between faster user performance time, a decrease in virtual navigation and an increase in physical navigation. As display size increases, instead of using the keyboard and mouse to zoom and pan around an information space, users preferred to do these interactions physically. Users moved closer to the display to zoom in on the details in an area, backed away from the display to zoom out for an overview, and walked along the length of the display to pan between areas of the display at various levels of detail. The researchers concluded that, because this increase in physical navigation (and a consequent decrease in virtual navigation) was correlated to faster user performance, users were more efficient at physical navigation.

Another study looked at how changes in display size and resolution affect user performance with information-rich virtual environments [28]. The research showed that increased display size and resolution improved user performance time on navigation tasks. In addition, the researchers observed users tended to both move their heads and use their peripheral vision more, indicating that they felt more present when using the larger display sizes.

Increasing the display size has enormous benefits for user performance, but also creates a new difficulty – how do users interact with it? The paradigm has changed from sitting at a desktop computer to somehow interacting with a wall-sized display. While acceptable for interacting with a desktop computer, it is clear that using a stationary keyboard and mouse, as they currently are, is not advantageous for interaction with large displays. These devices tether users, discouraging them from physically navigating, which they prefer and which makes them more efficient. For this reason, researchers encourage interaction designers to take physical navigation into account and design more mobile interaction. In addition the mouse is a relative pointing device, meaning that the movement of the cursor is relative to some starting point, and can be repositioned. Other direct pointing techniques, such as those using a raycasting technique, have an absolute mapping between the device and the cursor, meaning the device matches up exactly with the cursor. These types of techniques are seen as the most natural [9].

1.2 Problem

We know that users are navigating within the large display information space by moving. When they move in and out from the display, they zoom into details or out to an overview, essentially *changing the scale* at which they are viewing the information. With mobile interaction, as users move around in the space in front of a large display, they are able to interact with the display from anywhere in that space but this interaction is static.

Users' visual scale is changing, but their scale of interaction never changes. This is a problem because users are doing different types of tasks at different levels of visual scale. As they step out to see an overview, they are still interacting with it on the detail level, even though they are no longer able to see any of the details. For example, consider the storyboarding scenario described in section 3.4. An intelligence analyst stands close to the display while reading evidence tidbits, writing hypotheses, posting them to the display and linking one to the other (Figure 1.1a). Detail interaction works well for all of these things. However, once the analyst steps back to see how the subhyphothesis he has just created fits in with the overarching story (Figure 1.1b), he no longer wants to work on the detail level. He wants to be able to select the entire tree structure that represents this subhypothesis at once.

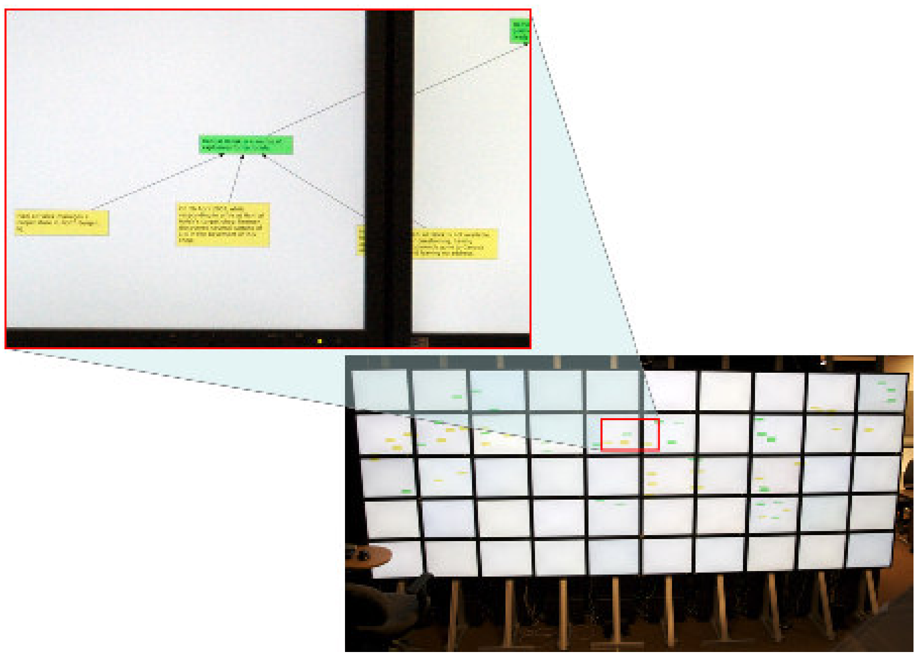

Figure 1.1: An analyst works on (a) individual evidence tidbits that are part of the (b) larger overarching story.

Furthermore, other controls that currently enable some sort of multi-select introduce extra confusion, visual and interactive clutter and difficulties with accuracy.

Going back to the intelligence analyst, if he were to use a typical multi-select control, such as a lasso tool from far away, he would have to draw very precisely around everything he wished to select, most likely with a shaky hand. There should be a less difficult way to interact on both the detail and overview levels when necessary.

There is evidence that people naturally interact on different levels of scale in the literature on pointing. Kranstedt and Wachsmuth define two functions of pointing, based on Lyons' description of pointing [27]: object pointing and region pointing. Object pointing indicates one object amongst the rest, while region pointing narrows the focus to an area or group of objects [23]. Related closely to this is that by its nature, pointing is ambiguous, and varies according to distance from the object being pointed to. Research has shown that people indicate nearby objects using object pointing alone, while adding referential descriptions when indicating a faraway object [26].

Therefore, the way people gesture is related to distance. For example, as pointing becomes more ambiguous, people compensate with longer, more complicated verbal descriptions. Up close, pointing is used to define location and almost no verbal description is used, while further back, location had to be described, often in unnecessary detail [6, 35]. In addition, pointing to closer objects is more precise and static, while pointing to further objects is more imprecise and tends to be dynamic, as people tend to circle or fill in area around referred object with movement. People also point more when they are referring to faraway objects, as they also refer to surrounding objects to help define the particular object in question. While looking at the effect of target size on pointing, researchers found that it is much easier to point to smaller things up close [35]. This change in how people point to a target based on their distance from it can be modeled approximately by a cone extending out from the finger [23]. The cone represents the degree of pointing accuracy with which people perceive they are pointing, with the cone denoting object pointing having a narrower angle at the apex than the cone denoting region pointing.

In summary, two things have been shown:

1. People use physical navigation to efficiently perceive information on large displays at different levels of scale [5], and
2. people naturally tend to point at different degrees of detail based on their distance from the target(s), and this can be approximately modeled by a cone extending from the finger [23].

Thus, *is it possible to combine these behaviors to provide an efficient and explicit multiscale pointing capability on large displays?* Furthermore, can the conical degree of detail model of pointing be applied to the scale of selection (or interaction in general), analogous to the scale of perception? Can physical navigation be exploited, using it as an explicit operator to control the scale level of interaction, in addition to perception?

1.3 Solution

We suggest a novel interaction technique called mulitscale interaction, which links the two behaviors by *changing the users' scale of interaction depending on their distance from the current object(s) of interest*. This solves the intelligence analyst's problem from earlier by increasing the size of the cursor as the analyst moves away from the display. His cursor is still small while he is doing more detailed work close to the display, but as he backs away, he can now select the entire tree at once, because his cursor has grown large enough to encompass it (Figure 1.2). This is very easy for the analyst, since it only involves him making one button click, and does not require him to point very accurately. Changing the intelligence analyst's interaction scale here is interpreted as varying the size of this cursor, but a change in interaction scale can be implemented in several other ways (discussed in Chapter 3:).

The goal of the work described here was threefold. First, we defined the design space for multiscale interaction. Secondly, we explored some areas of the design space in greater detail with a case study. Finally, we evaluated an implementation of multiscale interaction through a user study.

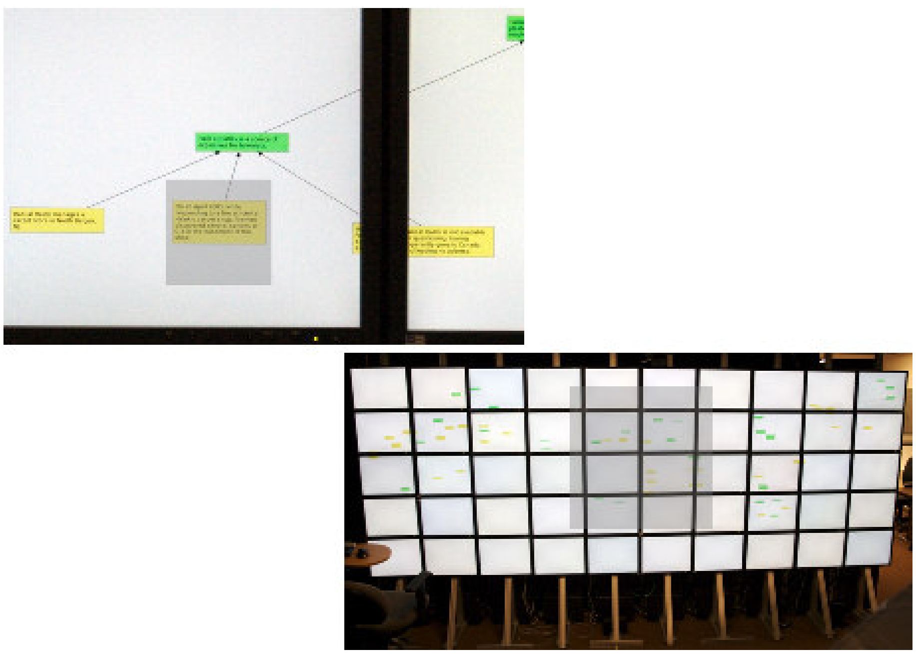

Figure 1.2: The analyst can select on both the (a) detail and (b) overview scales.

1.3.1 Design Space

With a general description of the technique in hand, we defined the design space of multiscale interaction. To what kinds of situations can it be applied? Through this definition, we explored several research questions:

1. How can multiscale interaction be applied to various domains? For example, a change in scale might have a different meaning for spatial data as opposed to temporal, or for hierarchical data as opposed to unordered.
2. How can multiscale interaction be applied to the various aspects of interaction? For instance, there are things beyond simple selection, including positioning and navigation.

Each of the possible interpretations that we have defined is discussed in greater detail in Chapter 3:.

1.3.2 Case Study

With a design space defined, we explored one of the identified situations in greater detail. For each of these applications, what are the possible design issues and tradeoffs? The scenario we chose to explore further dealt with a storyboarding application for intelligence analysis on large displays. Here, multiscale interaction was implemented by changing the size of the cursor according to the analyst's distance from the display. When the analyst was close to the display, his cursor was small, still enabling him to interact with individual tidbits of evidence or hypotheses. When the analyst moved further from the display, his cursor continuously grew in size, enabling him to select whole subtrees at once. This case study is discussed in greater detail in section 3.4.

1.3.3 Evaluation

In addition, we chose another of the interpretations to implement and evaluate with a user study, discussed in more detail in Chapter 4:.

Through the course of this study, we wanted to know, is multiscale interaction an effective technique? To this end, we wished to examine several things. Multiscale interaction enforces a link between users' scale of perception, controlled by their physical navigation, and their scale of interaction. Is exploiting physical navigation to control the scale of interaction as well as the scale of perception natural for users? Do users in fact benefit from the linkage of physical navigation with interaction? We hypothesized that they would, for several reasons:

1. Multiscale interaction exploits the fact that users are already physically navigating, perhaps even encourages it. Because of this, it will see the same time performance benefits as physical navigation.
2. There is a reason why users prefer physical navigation over virtual. This is because it is more natural. We believe this "naturalness" will manifest itself in both a reduced mental workload and tendency towards similar interaction in the other, non-multiscale interaction techniques.
3. Multiscale interaction is more natural because it reinforces for the user a sense of embodiment, leveraging his preexisting knowledge about the world around him. In addition it reduces both the visual and interactive clutter found in other techniques. Making these aspects easier frees up cognitive resources users may then apply to better problem solve, manifesting itself in higher accuracy.

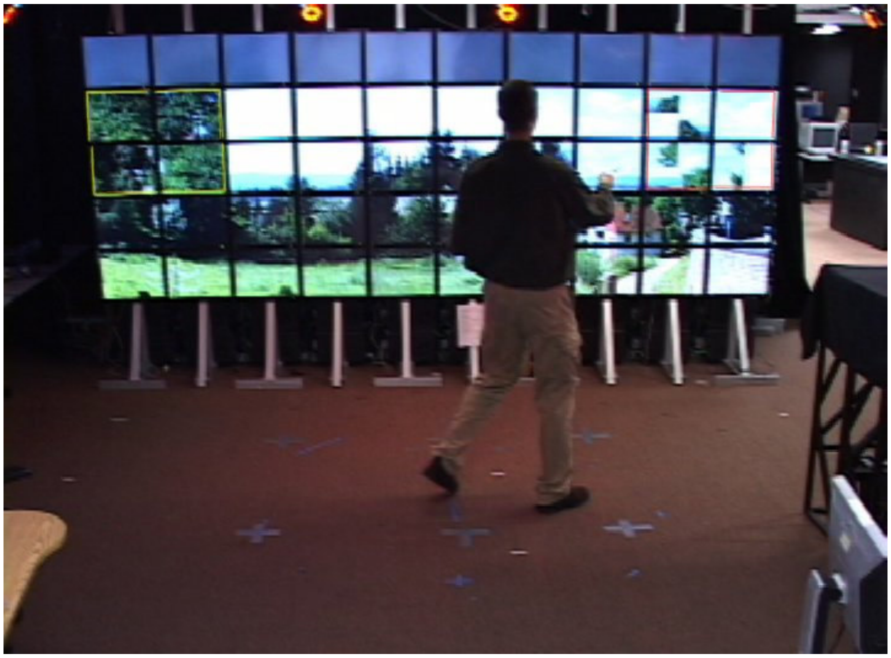

Figure 1.3: Interacting with the display at overview level using multiscale interaction

During the study, we asked participants to solve a multiscale puzzle using either the implementation of multiscale interaction or one of two other techniques. As with the storyboarding application, multiscale interaction was again implemented by automatically changing the size of the cursor according` to users' distance from the display. However, as they moved further away, the cursor grew, making discrete size jumps so as to select increasingly larger level size pieces, as in Figure 1.3. The other techniques included one which still discretely changed the cursor size, but allowed users to do this explicitly by choosing the interaction scale from a menu and a typical marquee technique in which users drew a box selection around the piece they wished to select.

1.4 Summary

This work explores the link between users' physical navigation, specifically their distance from their current object(s) of focus, and their interaction scale. We have defined a new interaction technique, multiscale interaction, which links users' scale of perception and their scale of interaction. The technique exploits users' physical navigation, using it to explicitly control scale of interaction, in addition to scale of perception. No other interaction technique for large displays has previously considered physical navigation to this degree (discussed further in Chapter 2:). Because this is a completely novel interaction technique, we have defined its design space, which other researchers can continue to explore and build on, as well as explored some design issues through a case study in Chapter 3:. Lastly, we have evaluated one implementation of multiscale interaction, to begin to quantify the benefits of the technique, as discussed in Chapter 4:. We show evidence that a natural psychological link between scale of perception and scale of interaction exists, and that exploiting it as an explicit control in the user interface is beneficial to users in problem solving tasks. In addition, we show that designing against this philosophy can be detrimental.

Chapter 2: Related Work

Due to the decrease in the price of displays, recent research related to high-resolution displays has sought to quantify the benefits of using increasingly larger displays. Several studies have shown that with large datasets, such as those found in geospatial analysis, the larger viewport size improves users' performance time, as well as decreases their frustration [4, 5, 33]. Researchers have sought to explain these benefits of large displays by turning to embodiment theory.

Previous work has also been done developing area cursors and other interaction techniques for large, high-resolution displays. However, these techniques were only intended to make selection of one object easier, and aimed at making interaction from a distance possible. They neither consider a change in interaction scale nor users' movement.

2.1 Physical Navigation with Large Displays

Not only has research been pushing the increase in physical size of displays, it has also been pushing an increase in the amount of data shown. A recent study looked at using large, high-resolution displays to show information that is beyond visual acuity, meaning there is no single spot where a user can stand and distinguish all the displayed information. Even though task times increased with increasing amounts of information, they were less than proportional time increases, indicating that user performance is more efficient and accurate because of additional data that can be displayed [40].

When visualizing these large datasets on single monitor-sized displays, the amount of data will always outstrip the amount of space in which to show it; so in order to see all of the data, users must navigate away from the original view. Users may achieve this through *virtual navigation* by panning around the space to view different areas of the visualization or by zooming in or out to change the scale.

When the available viewing window of the data is increased (as with large displays), there is less of a need for virtual navigation and more opportunity for *physical navigation*. Movement back and forth across the length of the display can be likened to panning and movement closer to and further away from the display can be likened to zooming.

In fact, on large displays with datasets that show information beyond visual acuity, users must physically navigate, for example, stepping forward to see details in an area and then stepping back to obtain an overview of all the information [5].

Previous studies found for basic visualization tasks that increasing the viewport has enormous benefits for user performance, but also found that with larger viewports, users tended to physically navigate more and virtually navigate less [5, 33]. This indicates that the better performance is both due to the additional information available and the ability to physically navigate [32]. When given the choice, users tended to physically navigate to understand the data, possibly because they wished to avoid losing context [5].

Users were also much less frustrated when using larger displays. The authors of the studies surmised this was because physical navigation is more natural, allowing for better use of innate human abilities, such as recalling spatial properties based on muscle memory and that larger displays grant greater use of the human visual capacity, allowing users to build mental models of the data using the additional information in the periphery [32].

The authors of one study went so far as to suggest several design factors to encourage physical navigation, including using mobile interaction, providing open space, using high-resolution displays, making sure the body is visible and that it matches up with the virtual world [5].

2.2 Embodied Interaction

In the studies just mentioned, the researchers found that physical interaction appeared to be more natural for the participants. The researchers turned to embodiment theory for an explanation as to *why* this type of interaction was more natural. Embodiment theory is an emerging topic in human-computer interaction, and as such is not yet clearly defined. However, embodiment deals with an individual's feeling of presence within the surrounding world and his or her participation with what is going on in the world. It claims that people interact with the world around them to increase their understanding of it [13].

Human thought processes are tightly coupled with the world around them, because a human's body exists physically within the world, interacting with it, and thought not only happens in the brain, but in the body as well [37]. Therefore, embodiment is relevant to interaction because interaction is intimately connected with the settings in which it occurs [12]. In short, *doing* helps people think.

An example of this is tangible interfaces. These types of interaction use the knowledge that people already have about the world and apply it, creating familiar (and therefore natural) mappings from user actions to outcomes [13, 22]. Multiscale interaction attempts to leverage knowledge about the world – that moving closer to an object will allow you to see the details and moving further away will allow you to see the object as a whole – to create a more natural interface with the information on the large display.

In her 2002 paper, Wilson discusses six views of embodied cognition, several of which are applicable in thinking about multiscale interaction [37]. The first view states that cognition is situated. Essentially, this means that context is important when considering how humans will behave and interact. In the case of multiscale interaction, the context in the broader sense is that people are physically navigating. Users must move in order to see all of the information presented on the display. Cognition is also situated in that different cognitive tasks happen at various distances. Users are thinking about the information differently when they are close to it (and viewing details) as opposed to further from it (and seeing the overview).

Another of these views states that people offload cognitive work onto the environment. Rather than remembering something themselves, humans place information out in the world and retrieve it later, when needed. An example of this is epistemic actions, like leaving mail by the door, so as to remember that it needs to be placed in the mailbox [37]. Multiscale interaction seeks to do this by offloading the cognitive activity of changing the interaction scale onto the environment, automatically changing the scale as the user changes his or her distance from the display.

2.3 Area Cursors

The idea for the form of the cursor aspect of the multiscale interaction technique is based on the idea of area cursors. An area cursor is a "cursor that has larger than normal activation area" [39].

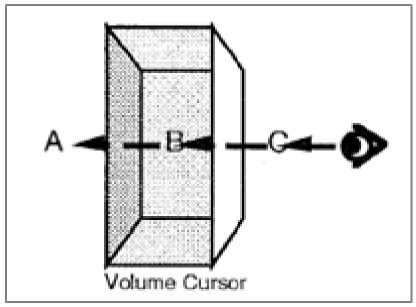

Figure 2.1: The "Silk" cursor. A volume cursor is covered with "silk" to give perception of depth based on occlusion. Reproduced from [41].

Zhai, Buxon and Milgram first introduce a 3D form of the area cursor, called the "Silk cursor" [41]. This technique uses a selection volume rather than a point (Figure 2.1). The volume is indicated by translucent sides, which better indicates depth than wireframe counterparts. Kabbash and Buxton follow up on this idea by introducing what they call the "Prince technique" [21]. The idea is that increasing the selection area would make it easier to select small targets. The authors also suggest that you should switch between Prince and point-cursor techniques, depending on *task context*, and indicated that the Prince technique was not optimal for fine positioning tasks. Worden et al improved upon the "Prince technique" by introducing a disambiguation technique to help indicate a single object for selection when multiple objects fell within the area cursor [39]. Their implementation added another point hot spot that would take effect with multiple targets present.

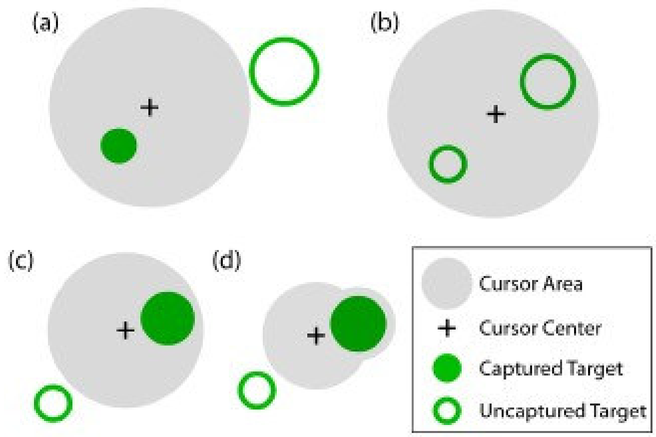

Figure 2.2: The Bubble Cursor. (a) An illustration of a typical area cursor. (b) This can lead to selection of multiple possible targets. (c) Bubble Cursor solves this problem by resizing. (d) If resizing is not enough, the Bubble Cursor can also morph to encompass a target. Reproduced from [16].

Another technique known as the Bubble Cursor (Figure 2.2), is based off of Worden et al's area cursor, but additionally disambiguates selection by dynamically resizes the cursor depending on the proximity of other targets [16]. The shape of the Bubble cursor is also changed from previous area cursors. The circle shape ensures that the target closest to the center of the circle is always what is captured.

While this cursor implementation for multiscale interaction is based on the idea of area cursors, it differs from this previous work in several ways. These techniques are intended for single object selection, and are therefore only dealing with one level of scale, the detail level. Bubble cursor is the only technique that really considers the possibility of having a cluttered information space with lots of data; however, it only acknowledges the situation as an impediment to single object selection and uses disambiguation strategies to cope.

2.4 3D Interaction Techniques

Much work has already been done developing pointing and selection techniques for large display environments, both virtual and not. These techniques can be classified by the degree to which they incorporate physical navigation.

1. The baseline level of integration with physical navigation is interactions that either ignore or avoid it. This interaction would mainly be concerned with how to point quickly and accurately in a large information environment, perhaps assuming that the user will not be moving.
2. The next level of integration with physical navigation includes techniques that support it, but go no further. This interaction would be concerned about enabling pointing from any position in space or while moving.
3. The third level of integration is interaction that exploits physical navigation. This includes techniques that take advantage of the fact that users are physically navigating, integrating it into the interaction itself.

Previously, as in levels 1 and 2, interaction has remained constant, regardless of the context of the user's position relative to the display or the semantics of his current visualization task. The 3D interaction techniques summarized here all fall into either of these two levels. *Multiscale interaction*, however, falls into level 3, because it seeks to take this context into account, changing the scale at which users interact with an information space according to their distance from the current object(s) of interest.

4. The hierarchy could possibly continue here; the final level would include interaction that exploits multi-user physical navigation. This would focus on allowing users to collaboratively point. Multiscale interaction could be used in a

collaborative setting, as in the multiscale gaming example mentioned in section 3.3.3, but this application is not the focus of this body of work.

2.4.1 Ignoring Physical Navigation

Direct pointing is a cognitively direct interaction technique which uses ray-casting to point to and select objects. Jiang, et al's Direct Pointer was designed for situations when only remote interaction with large displays is possible [20]. This technique seeks to enable interaction from afar while completely discounting close interaction. It uses simple handheld digital cameras, such as those found in cellphones to do direct, laser pointer-like interaction. The Direct Pointer indicates cursor position with a red circle of diameter 48 pixels.

The Multi-Scale Cursor takes large display interaction in a different direction and seeks to alter typical mouse interaction to better fit the scale of large displays [11]. The technique is intended for use with large personal workspace scenarios, where the user is seated. The main problems with mouse interaction are "clutching" for long distance movements and visibility of the cursor, which the multi-scale cursor solves by increasing the speed and size of the cursor based on users' mouse movements.

2.4.2 Supporting Physical Navigation

Other direct pointing techniques, however, seek to enable interaction from multiple locations in a space. Cao and Balakrishnan's VisionWand technique (Figure 2.3) passively tracks a wand in 3D space that is used to interact with the display [10]. The technique focuses less on direct pointing and more on enabling coarser interactions from anywhere in the camera-tracked space. These coarser interactions include a set of wand gestures developed that match to specific commands. Several of the gestures incorporate movement in relation to the display, but only small arm movements as part of the gesture only. The VisionWand technique provides several visual feedbacks: two differently colored circles indicating the orthogonal position of the wand markers on the screen, as well as a set of black crosshairs indicating the intersection point of the ray.

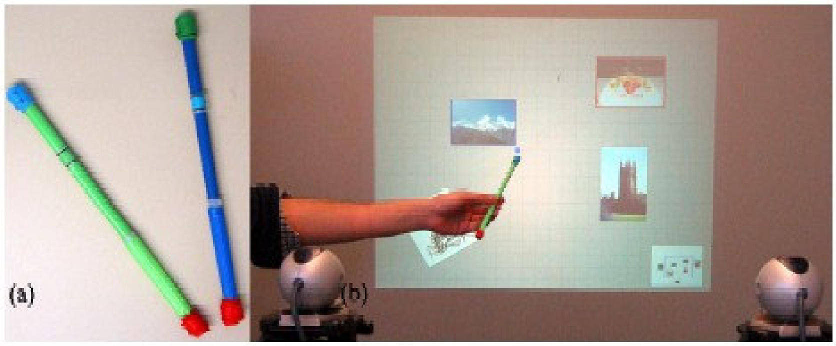

Figure 2.3: VisionWand. (a) The wand itself. (b) System setup. Reproduced from [10].

Vogel and Balakrishnan describe a set of pointing and clicking techniques for large, high-resolution displays that use the human hand as the implement [36]. The authors emphasize the need for mobile interaction techniques that can transition smoothly from close to far interaction. The techniques involve tracking hand gestures for both direct and relative pointing (Figure 2.4), as well as several clicking techniques. The cursor is indicated by the normal arrow, but must be augmented with additional visualizations to compensate for lack of tangible feedback from interactions such as clicking.

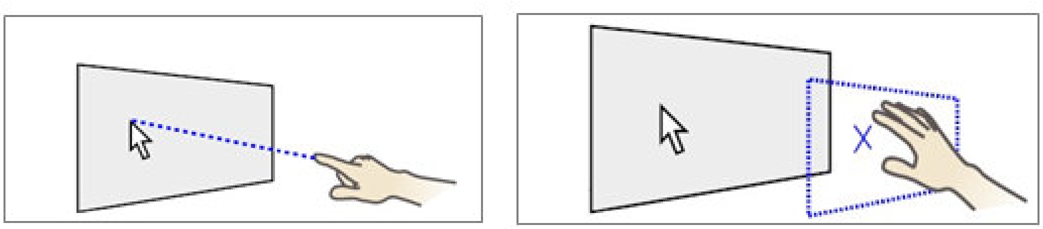

Figure 2.4: Interacting remotely using direct pointing (left) and relative pointing (right).

Liang and Green first describe a 3D selection technique that extends the ray-casting concept, called the "spotlight" technique [24, 25]. This technique shoots a cone out of a pointing implement, with the apex of the cone at the end of the pointer. Objects that fall inside the cone are candidates for selection. Only one object may be selected at a time; object selection is disambiguated by selecting the one closest to the origin of the cone. The selection area is visualized by a transparent cone, with the intensity of light indicating closeness to origin.

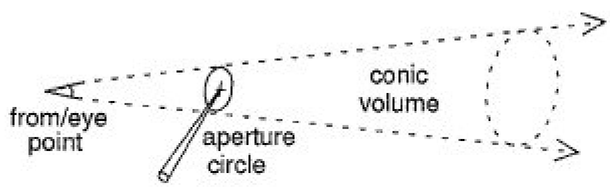

Figure 2.5: Aperture based selection. Effective selection volume is based on eye point and aperture geometry. Reproduced from [14].

Aperture based selection is based on the "spotlight" technique, and is designed to alleviate problems associated with aiming and selection of an object from a distance [14]. Much like the original "spotlight" technique, things inside the cone are candidates for selection; however, the user is additionally able to control the size of the selection cone by moving the pointer away from/towards a fixed point called the eye (Figure 2.5). Again, the technique is intended only to select one object at a time. Selection area is shown using the aperture cursor, reducing the clutter from visual feedback about what is selected.

2.5 Multiscale interaction on single-screen or low resolution displays

The technique described in this thesis differs from other "multiscale navigation" techniques for several reasons. These navigation strategies focus on use in zoomable user interfaces (ZUIs) such as Pad++ [8], which were designed with single monitor displays in mind.

Previous research concerning multiscale navigation has always looked at moving through information spaces much larger than the single monitor viewport [17, 18]. In order to acquire the target, the user must first point the view (which can include zooming) so as to make the target visible, then point the cursor at the target. In these cases, multiscale navigation is essentially view pointing.

2.6 Summary

Motivated by the benefits we have seen from using large, high-resolution displays, we have therefore developed a novel interaction technique that takes advantage of the physical navigation that is already being done by users of large displays.

The implementation of the described multiscale interaction technique is based on several prior techniques in the areas of both standard and 3D interaction, including area cursors, remote pointing, and "spotlight" selection. However, these techniques either only seek to enable accurate pointing from a distance or only consider movement inasmuch as to provide a technique that works anywhere in the space. The direct pointing techniques can select multiple objects by drawing a lasso or marquee around the objects. However, this is done using a point cursor. The "spotlight" technique and its derivatives change the size of the selection area, but are still only intended for single-object selection. In both cases, the interaction is still on the detail level and nothing more.

Multiscale interaction aims to exploit users' movement, incorporating physical navigation into the interaction itself, thus making the interaction more natural. It also seeks to enable interaction on different levels of scale by allowing the user to select individual objects as well as groups or hierarchies.

Chapter 3: Design of Multiscale Interaction

3.1 A case for multiscale interaction

As displays have grown larger, both in the physical sense and the resolution sense, the amount of information available on them at one time has become astonishing, reaching the point where there is no single place where a user could view every piece of information on-screen. Users will somehow have to change their view of the information space in order to see all aspects of such visualizations, and physically navigating is one way to do so. This is not a negative aspect of such displays; we have presented other research which has proven the benefits of physical navigation on several aspects of user performance.

When users are up close to the display, they are able to see all the details about the area nearest to them. However, the edges are so far away and the angle between them and the edges is so shallow that it becomes difficult for them to see the details outside of their close vicinity. As they move further from the display, more of the area becomes visible to them, but the detail information becomes more and more obscured. At the furthest position out, they can see the entire display at an overview level. For example, consider the storyboarding scenario described in section 3.4. An intelligence analyst stands close to the display while reading evidence tidbits, writing hypotheses, posting them to the display and linking one to the other. Once the analyst steps back, he can see how the subhyphothesis he has just created fits in with the overarching story, but can no longer make out the text of the subhypothesis itself.

Essentially, the large display users are changing the scale at which they are viewing the displayed information. Much like with a ZUI, different amounts of detail and different percentages of the overall visualization are available, depending on the distance from the display. Through all this, the interaction remains constant, which is a problem because user tasks vary at different levels of visual scale. Detail interaction works well for working with detail information, but when users are no longer able to see detail information, they do not want to be interacting at that level. Considering the intelligence analyst, when he has stepped back, he no longer wants to work on the detail level. He wants to be able to select the entire tree structure that represents this subhypothesis at once.

3.2 Hypothesized Benefits

We hypothesize that users will benefit from multiscale interaction due to improvements in several areas. Multiscale interaction:

1. exploits users' physical navigation to enable simpler interaction at multiple levels of scale. Unlike other techniques mentioned in Chapter 2:, which at best enable interaction from any position, multiscale interaction links physical navigation with interaction scale, changing the users' scale of interaction depending on their distance from the display.
2. considers human embodiment. By leveraging users' preexisting knowledge about the world around them and exploiting their physical navigation, multiscale interaction is much more natural than those techniques that do not.
3. simplifies the users' cycle of interaction (as defined by Donald Norman [29]) by reducing the visual and interactive clutter as compared to other typical multi-selection techniques, such as the lasso tool.

We hypothesize that these benefits will manifest themselves in measurable effects on users' performance, which will be discussed in more detail in the following chapter.

3.2.1 From physical navigation

Users of large displays might be physically navigating out of necessity, but this is not the only reason for multiscale interaction to consider it. As previously mentioned, prior experiments have shown that physical navigation is advantageous over virtual navigation for many reasons, including increasing users' efficiency and decreasing their frustration. By exploiting physical navigation, multiscale interaction should see the same time performance benefits.

As users move away from the display, their perception of the data changes. Visual aggregation comes into play, and they begin to lose individual details and start to see patterns and groups. With this change in perception comes a change in the tasks that are appropriate to perform on the data. By designing the interaction to change as a function of the distance from the display, we are supporting users' changing perceptions with equivalently changing interactive control.

3.2.2 From embodied interaction

Times when users are already physically navigating it may make sense to link physical navigation with interaction by "giving meaning to distance." In this manner, we seek to create a more natural interaction technique for large display information spaces.

Movement closer to and further away from an object to change the amount of visual information available is nothing new. This phenomenon is already present in the world around us, and is therefore intuitive. For example, consider a patron at an art gallery. While he is up close to a particular painting, he can see one aspect of the painting or perhaps individual brush strokes. When he backs up, he can no longer see the brush strokes because they have faded together to form the entire painting. This is a change in visual scale. Interaction also changes. One art student might say to another while pointing very close to a painting, "Look at the type of brush strokes the artist uses here." Once the two students step back, the other might point to entire paintings and say, "I prefer that scene to this one."

Large display users already expect this from the world around them, and will transfer this expectation to their interaction with large displays. They already know that when they move closer to the display, they will be able to see details, so they know that it makes sense to do detail-oriented tasks while they are close to the display. Users know that when they are far from the display, they will no longer be able to make out details, but instead will be able to see the display in its entirety. Therefore they know that it would make sense to do overview-related tasks while they are far from the display. Multiscale interaction uses distance from the current object(s) of focus to change the scale of interaction, which corresponds to the task contexts at each location, thus leveraging knowledge and assumptions users already have about the world around them.

3.2.3 From visual and interaction design

Donald Norman's cycle of interaction theory indicates two opportunities for a breakdown in a user's interaction with the computer: the *gulf of evaluation* and the *gulf of execution* [29]. The gulf of evaluation refers to how effortful for the user it is to interpret the state of the system and figure out the degree to which the system has met his expectations. In this vein, there is a perceived visual advantage to linking physical navigation with a change in interaction scale – it semantically "overloads" the selection cursor. In other words, in addition to indicating the item for selection, the multiscale cursor also informs the user of the current level, and gives the user information about the levels in general, without having to additionally clutter the display with more cues. In the user study discussed in Chapter 3:, the lasso interaction technique represents the absence of both a multiscale aspect and a link to physical navigation. Because of this, the user must be informed of the levels of hierarchy by some sort of additional cue. The advantage here is less visual clutter, distracting the user from the dataset less and making it easier for him to interpret what is currently being pointed at.

The gulf of execution looks at how directly the user can perform actions with the system and avoid extra effort. In this vein, linking physical navigation with change in interaction scale has another benefit: there is less "interactive clutter." The user does not have to think about operating a selection control *and* an interaction scale control; the latter is taken care of automatically. This allows the user to get to the same state (selecting his intended object at the intended scale) in fewer steps.

In addition, multiscale interaction provides a solution for the accuracy problem inherent in direct pointing. As users move further from a display, it becomes more difficult to accurately point at any one pixel. This is due to both uncertainties present in the pointing device's tracking system and human physiology. However, with multiscale interaction, as users move further from a display, the interaction scale increases. In the case of the technique used in the user study, the size of the cursor increases to select larger puzzle pieces. What this means as far as pointing goes is, as the user moves further from the display and pointing accurately becomes more difficult, the selection area grows larger and users no longer need to point as accurately to select the larger pieces.

3.3 Defining the Design Space

There are several subtasks within interaction: manipulation, including selection and positioning, and navigation [9]. Once an object of a particular scale is selected, we assume the positioning technique will remain the same, so we focus on applying multiscale interaction to selection and navigation; to the former by defining *what* gets selected and to the latter by defining zoom or speed scale. In addition, multiscale interaction can also be applied to multi-user (collaborative) or multi-step interaction. Subcategories are further explained in Figure 3.1.

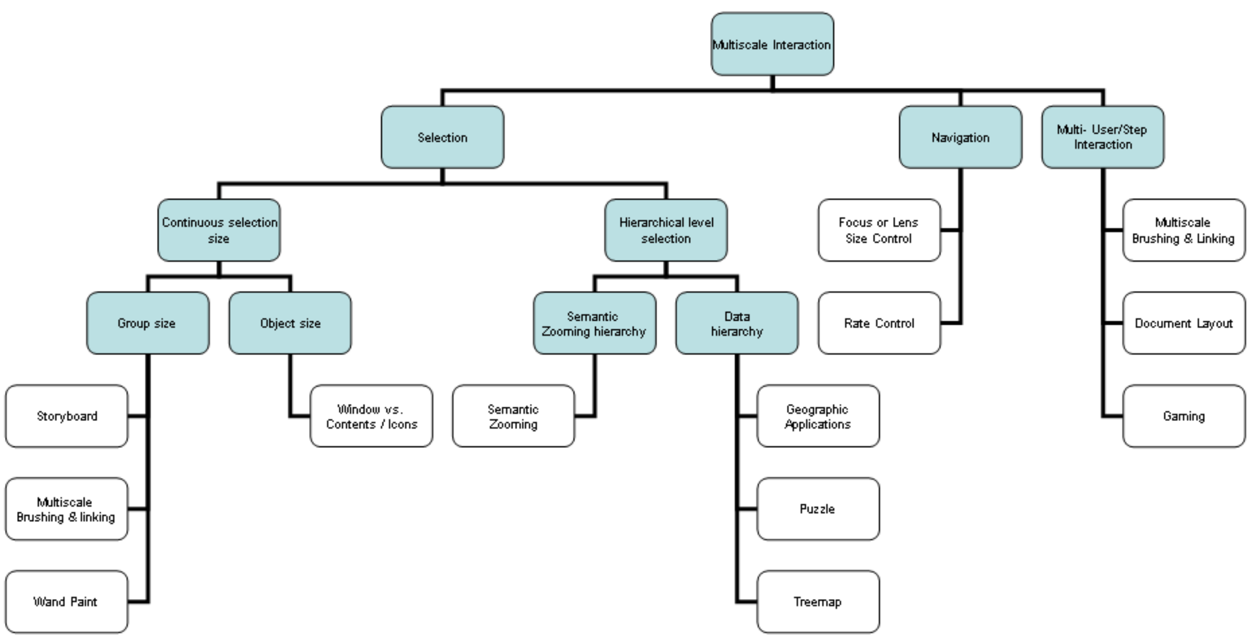

Figure 3.1: Design Space for multiscale interaction. Areas of the taxonomy are represented by blue, while specific examples are represented in white.

3.3.1 Selection

Continuous Selection Size

This interpretation is very similar to the "spotlight" technique introduced by Liang and Green in that the size of the cursor changes smoothly based on the user's distance from the display. However, there is a distinct difference in semantic use – unlike the "spotlight" technique, which focuses on making a single target easier to select from a distance, the change in cursor size changes the scale at which the user is interacting with the display. At every scale, the cursor allows the user to select *everything* that falls inside its area. This interpretation varies the size of the cursor in a continuous (or smooth) manner, according to a predefined equation, such as the intersection of a conic section or an exponential function. Because of this, it might be good for large datasets that do not have predefined hierarchical groupings. This can be interpreted two ways. The cursor can grow or shrink to select either (a) larger or smaller groups of objects or (b) larger or smaller objects themselves.

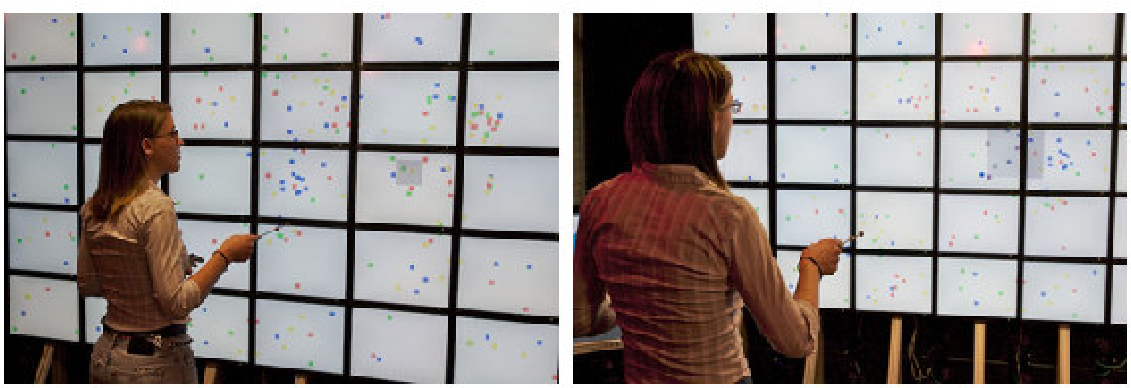

Figure 3.2: Interaction with a graph layout application using multiscale selection

Group Size

Graph layout. We developed a simple graph layout application as a case study that explores this interpretation of multiscale interaction, as seen in Figure 3.2. When users are close to the screen, the cursor is small enough to pick out a single node of the graph. As they move back, the cursor's selection area increases exponentially, allowing users to manipulate whole sub-graphs at a time. The storyboard case study, discussed in further detail in section 3.4, is based off of this initial application

Wand Paint. One of the earliest applications in which we first experimented with wand-based interaction was a simple painting program. The wand behaves like a virtual paintbrush; wherever the user points the wand, "paint" appears on the screen (Figure 3.3). To provide additional input beyond basic pointing, we use a normal wireless mouse. One tap of the mouse button starts painting and a second one stops it. The mouse's scroll wheel cycles through a predefined collection of colors, allowing the user to rapidly change the color of the brush. Breaking the metaphor slightly, but in a way that all of our users have grasped instantly, flipping the wand around backwards gives the user an eraser.

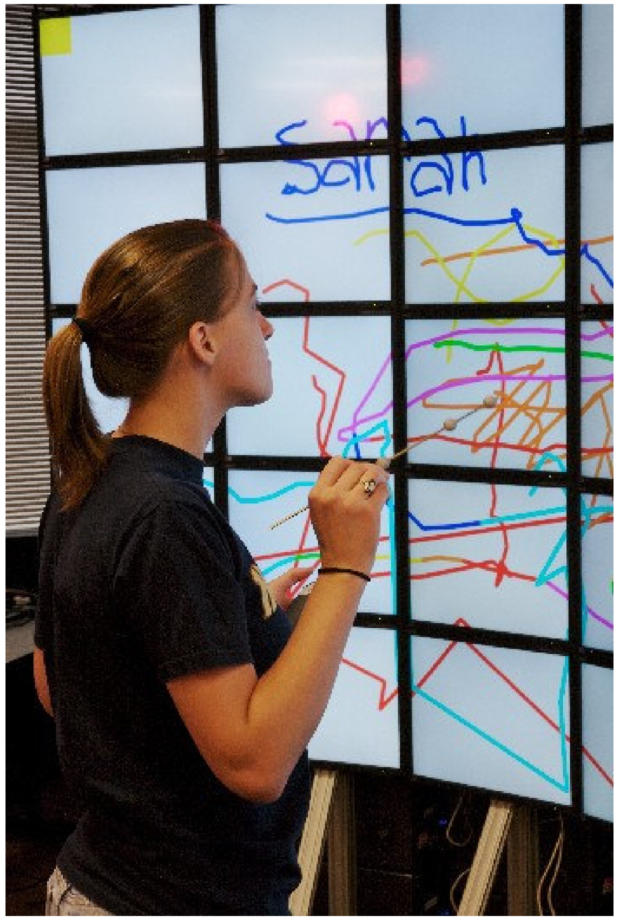

Figure 3.3: A user interacting with the Wand Paint application

Users control the size of the virtual brush by varying their distance from the display, bringing a multiscale aspect to the interaction of this application. The size of the brush varies smoothly with distance, so this is still an instance of continuous multiscale interaction. However, instead of varying the selection area of the cursor, this implementation varies the pointer's area of influence by varying the amount of paint. As users move away from the display, the size of the mark left by the virtual brush increases. Thus, users can move to the back of the room and paint in broad sweeping strokes or move forward for smaller, more detailed work.

Multiscale brushing and linking. When visualizing large heterogeneous information sets, the use of multiple coordinated views enables users to gain different simultaneous perspectives and exploit advantages of multiple visual representations. Brushing and linking enables users to explore relationships between the representations by selecting data items in one view to see corresponding items highlighted in other views [7]. However, on large high-resolution displays containing large views of millions of data values, users need to explore relationships across scale. Hence, there is a need for scalable interactions that enable users to rapidly select subsets of data points at different levels of scale.

When users are up close to the display, working at the detail level, they can select individual data points to see their locations highlighted in other views. In addition, with multiscale interaction, users can also step back to rapidly brush larger sets of data points over broader regions of the space to find larger scale patterns between views.

Object Size

Continuous variance of selection size can also be used to select larger and smaller objects themselves. This is similar to hierarchical level selection, but the data items are not hierarchically organized. A simple example of this is selection of windows versus selection of icons, while a more complex example is layout (of a newspaper, poster, presentation, or other large document).

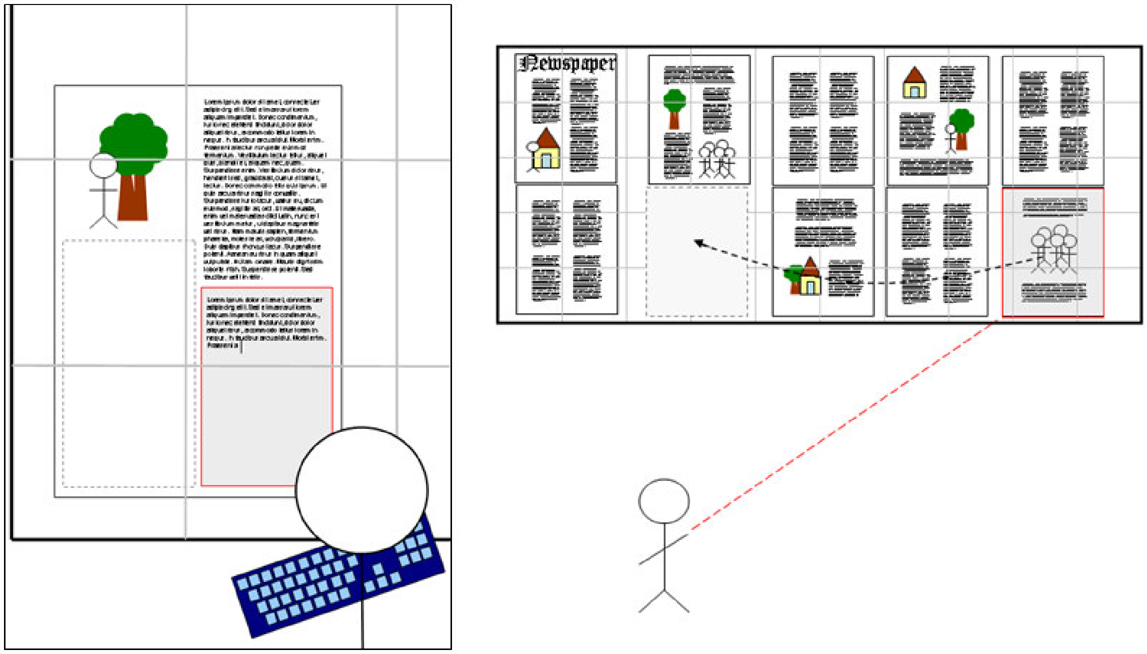

Figure 3.4: Using the example of document layout, with multiscale interaction, a user would be able to move from detail work, such as writing an article (left), to overview work, such as arranging the order of pages in a document (right).

While standing close to the display, the user would be doing very detail-oriented work, such as typing or editing an article. Standing a bit further back, the user would be able to see a larger area of information, such as a whole page of the document. A typical task done at this level would be arranging articles, photographs, advertisements, etc., within a page or between pages. Standing farthest back, the user can no longer read individual articles, but can now see the overall layout of the entire document. An appropriate task for this level of scale would be arranging pages within the document as a whole. Here, the user would be selecting objects of different sizes, so the cursor would act like a "timorous" or "void-phobic" cursor, snapping to the nearest appropriately sized object [19].

Hierarchical Level Selection

Instead of smoothly varying the effective selection area of the cursor, another interpretation of multiscale interaction would be to vary the hierarchical level of interaction of the cursor. Essentially, this would still change the size of the cursor, but by discrete intervals, according to a predefined hierarchy within the data. Objects on the same level of hierarchy may not be the same size, for example states in a geographic visualization (think New Jersey vs. Texas), so the cursor would resize to encompass the object being directly pointed to.

Data Hierarchy

Hierarchical level selection can also be used with data sets that already have some inherent structure. Users' physical navigation would be exploited to select objects on different levels of hierarchical scale.

Puzzle. The puzzle task used in the study discussed in Chapter 4: is also an example of hierarchical selection. The pieces of the puzzle are arranged in a hierarchy, and smaller scale level pieces are only shuffled within the bounds of their parent piece. Multiscale interaction automatically adjusts the cursor size to select on different scale level size pieces. In order to solve the smallest scale level size pieces in one area, users step forward to focus in on that area. As they step further away from the display, users see an overview of the puzzle; simultaneously, the cursor size grows, so users may select on larger scale size pieces.

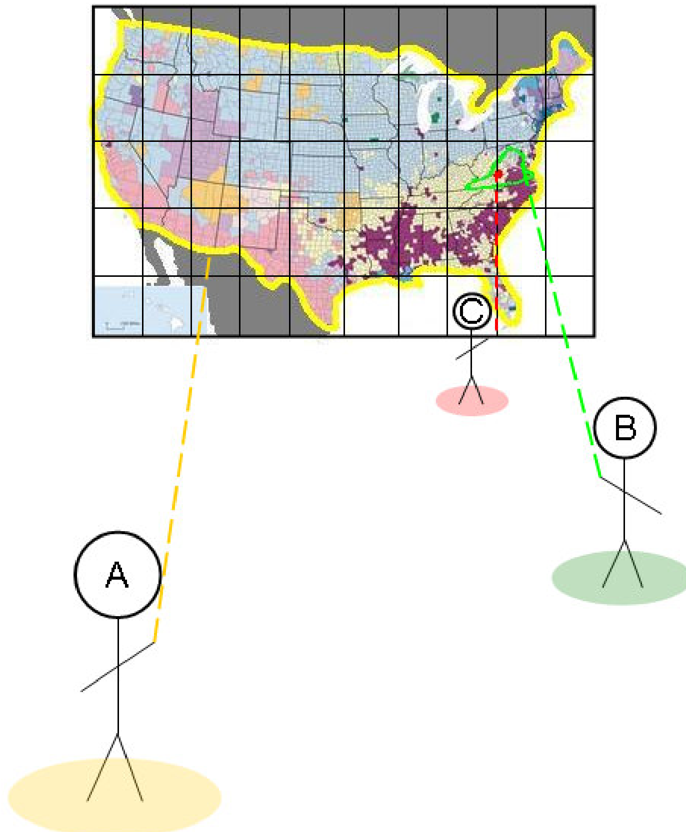

Figure 3.5: Multiscale interaction can discretely vary selection area according to hierarchical levels

Geographic Applications. Applying this to geographic information visualization, a frequent problem is specifying the scale of selection in a hierarchy of political boundaries. Multiscale interaction can automatically adjust selection scale to different levels of boundary hierarchy (city, district, county, state, country). For instance, in a visualization of US demographics, stepping back from the display would enable a user to select state level aggregations (Figure 3.5b), while stepping closer to the display would enable selection of county level aggregations (Figure 3.5c).

Treemap Interaction. This interpretation would also work well with treemap visualizations. There does seem to be an existing problem with indicating selection with these sorts of visualizations, because it is unclear whether one means to choose the

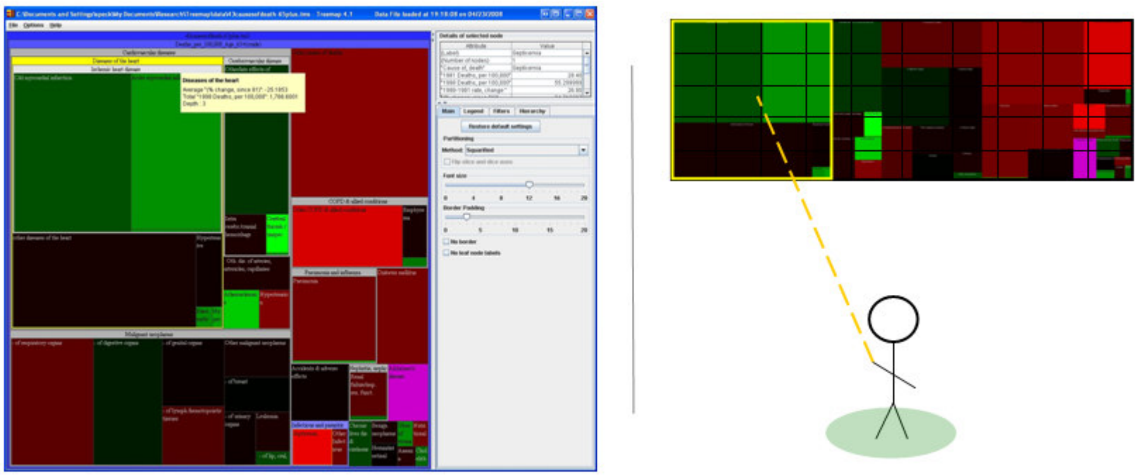

Figure 3.6: University of Maryland's Treemap (left). Similar selection done with multiscale interaction (right).

innermost node or one of its parents. Treemap, by the University of Maryland [3], puts borders around each level, and one must click in the border to explicitly select parent or interior nodes (Figure 3.6, left). In this case, using multiscale interaction to discretely change the cursor size improves the visualization by removing clutter due to the visual cues that indicate hierarchy (Figure 3.6, right). Another implementation, SequoiaView [2], removes these borders to save space for more nodes, but compromises the interaction, since this makes it impossible to select parent nodes. At best, it indicates the path from the root to the leaf node which has been selected. There is a tradeoff here between visual clutter and the existence of an affordance for interaction. Multiscale interaction solves this by not requiring the visual clutter of the parent borders while still allowing users to select on those higher levels of scale.

Semantic Zooming Hierarchy

Semantic zooming. Semantic zooming is a visualization technique that changes the representation of the displayed information according to the virtual zoom factor of the application. It is often used as a form of details on demand, allowing users to see varying amounts of detail as they zoom in and out. For example, semantic zooming could be applied to a map visualization. At the overview level, users may only be able to see state boundaries and the largest roads. As they zoom in, smaller roads and labels for smaller towns may appear. This extends the idea of using multiscale interaction for data hierarchies, by still allowing selection on various levels of hierarchical scale, but in addition, only showing levels of detail when users are close enough for it to be of use to them.

Semantic zooming has usually been applied to virtual navigation of information spaces, but it could be linked to users' physical navigation in a similar manner. Additional detail information can be used to augment the already existing information space as users move in closer to the display. For example users may plan their route for a long distance trip. Standing further back, they can select the large highways they will take to get from one state to another. As they move in closer, more information about smaller roads becomes available, and they can plan out how to get from the highway to their final destination.

3.3.2 Navigation

The idea of multiscale interaction may also be applied to preexisting visualization navigation strategies.

Focus or lens size control. Focus+Context strategies, such as fisheye views, seek to show detail information within the already present overview information by distorting the overview. Fisheye views use a Degree of Interest (DOI) function to assign relative sizes to areas of a visualization, with the current focus area of the view assumed to have the largest DOI and therefore the largest size allotment [15]. Multiscale interaction can be applied here to vary the size of the fisheye view focus with users' distance from the display, allowing users to focus a larger fisheye lens as they move away from the display, and narrowing the focus as they move closer.

A Magic Lens filter is an arbitrarily shaped area, movable by the user, that allows the user to generate a selective altered view of the underlying data [34]. These lenses have been used many ways, including augmenting visualizations with further details or additional context information, simultaneously showing coordinated alternate views and highlighting chosen features. Multiscale interaction can, in general, be applied here to vary the size of the lens with users' distance from the display. More specifically, distance from the display can be related to other application specific properties of the lens. For example, if it is a magnifying lens, distance can be related somehow to the magnification factor.

<u>Rate control in 2D and 3D navigation.</u> In large-scale 2D and 3D navigable information spaces, rate of navigation motion should be tailored to the user's visual perspective. For example, consider a ZUI (zoomable user interface) such as a space-scale zoom+pan navigation interface for large 2D satellite imagery analysis (e.g. TerraBlaster [30]). Note that a virtual navigation UI is usually needed even in the presence of physical navigation because total image size typically exceeds total display size.

Multiscale interaction can provide automatic rate adjustment for navigation controls, such as panning and zooming, based on users' physical position relative to the display. For example, when users are up close, they would have fine movement control and when they are further back, they would make more coarse movements. Hence, objects move across the visual field at a constant visual-angular rate regardless of the distance between user and object.

3.3.3 Multi-User or Multi-Step Interaction

Multiscale interaction can also be applied to collaborative or multi-step interaction. This would involve multiple individuals working on a single dataset simultaneously, possibly at different levels of scale; or it could involve a single user making interactions within interactions on various levels of scale.

Multiscale gaming. This interpretation can be applied to gaming, specifically real-time strategy games such as Stratagus (Figure 3.7). These types of games currently expect the player to manage all resources in detail. With a multiscale technique, players would be able to step back and make more vague commands that would affect the campaign as a whole, not specifying any particular unit. The player would then also have the option of stepping forward and issuing more detailed unit-specific orders. Multiple players may also be interacting at once, one individual playing the role of general while another individual plays the part of unit captain.

Figure 3.7: A player stands back from the display and gestures to western defenses to prepare for an attack (left). A player stands close to the display to issue specific orders to an engineering unit (right).

Multi-step brushing and linking. The multiscale brushing and linking application mentioned in section 3.3.1 can be extended to incorporate multiple interaction steps on different levels of scale. Users could brush multiscale patterns – patterns within patterns – by brushing broad patterns first and then stepping forward to brush local patterns with a different color.

3.4 Case Study: Storyboard

We performed a case study to further explore one area of the design space. In particular, the application we developed uses an implementation of continuous selection size, enabling users to select larger and smaller groups of objects. First, the domain of the application is described, and then some design issues are discussed.

3.4.1 Domain

From the basic graph layout application discussed in section 3.3.1, we developed a more complex system called Storyboard, designed to assist intelligence analysts in building scenarios out of a large collection of unorganized intelligence information. Our basic approach is based on Wigmore charts, a structured tool for evidence marshalling that builds, in essence, a tree that illustrates the analyst's chain of reasoning from evidence to hypotheses [31]. The leaves of the tree are made from facts that have been gathered in the field. At the root of the tree, the analyst places a possible scenario that ties together a subset of the gathered facts into a plausible story. Larger trees are built by combining several smaller trees together to arrive at some more encompassing scenario, which ties together the suppositions represented by the smaller trees.

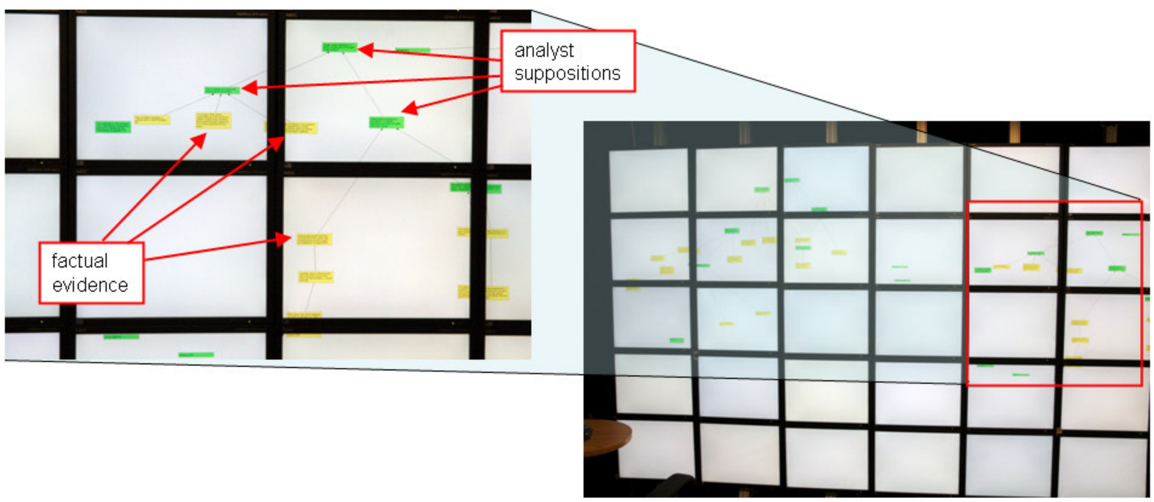

Figure 3.8: The Storyboard tool. The analyst builds subtrees from facts and suppositions (left) that can fit into a larger scenario (right).

A typical usage scenario for an analyst may begin with an evidence marshalling session standing up close to the display. He begins the session by looking through his available tidbits of evidence. Moving to various areas of the display, he arranges the evidence into groups, looking for similarities in the evidence such as shared locations or people. Once he has formed these groups, he begins forming his suppositions, or hypotheses about what he believes the evidence at hand means. He posts these suppositions to the display and links the relevant evidence tidbits to them. Once the analyst has several subhypotheses, he may move away from the display in order to arrange the subtrees and link them together to form larger, more encompassing hypotheses. Multiscale interaction supports these changes in selection scale.

3.4.2 Design Issues

Pointing Technique

The current interaction device for storyboard is a PDA with an attached wand. The wand provides pointing information, while the PDA provides a rich set of interaction tools, from hardware buttons to basic GUI widgets like menus and text entry boxes. The basic tasks that the system supports are adding evidence from the database onto the workspace, creating suppositions, drawing links between objects and rearranging objects within the workspace.

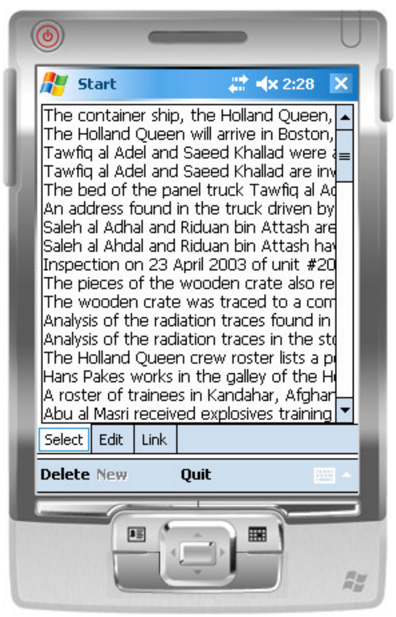

Figure 3.9: The PDA interface for Storyboard

The combination of PDA and tracked wand worked well for several reasons. GUI development for PDAs is very fast, and the Storyboard system was intended as a prototype. We struggled with finding a simple solution for mobile text input, and the

PDA worked well for this while also keeping well with the post-it note metaphor we used.

Pointing technique used can vary with domain. Many things can be tracked for pointing, including wand alone or the human hand. Instead of using buttons on a mobile device like a PDA, gestures can be developed for selection or other actions.

Cursor Shape

To incorporate multiscale interaction, rather than thinking of the wand as pointing a ray with a single intersection point towards the display, we think of the wand as behaving like the tip of a large pyramid. The intersection with the display is no longer a point, but a square slice of the pyramid. The square shape of the cursor better fits the form factor of the evidence tidbits, represented as post-it note sized objects.

Other selection techniques like spotlight use conic sections to form circular shaped cursors. Other shapes may be used, such as ellipses. Shape could also vary with distance if one shape is more appropriate for detail tasks than for overview tasks, for instance; or it could vary with a secondary user input, such as wand distance from eye used in the aperture technique.

Cursor Size

As the user moves away from the display, the surface area of the intersection slice increases. In practice here, we use an exponential function to grow the cursor size, as we found that it better matched the transition between tasks at different distances from the screen. We tested interaction using a linear function, but as we reached the back of our tracked space the cursor had not grown fast enough to perform the overview level tasks of large subtree selection.

We also defined the distance from the display as the length of the ray between the tip of the wand and the intersection point on-screen. However, it may also make sense to define distance as the perpendicular distance from the wand to the display, for instance if users do not want to select on larger levels of scale when pointing at oblique angles to the display.

The cursor size function should be calibrated according to the size of the space available in front of the display, as well as the variance in scale needed between detail and overview tasks. As with cursor shape, cursor size can also be tweaked using a secondary user input.

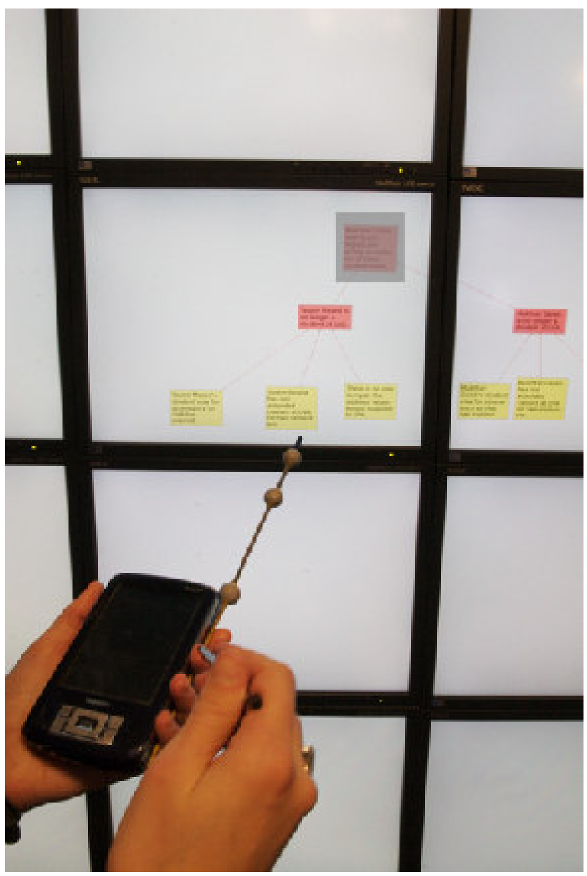

Figure 3.10: An analyst selecting a supposition using the multiscale cursor

Feedback

Storyboard's cursor was represented by a filled medium gray translucent square. The filled grey square stood out well against the white background, and could be distinguished easily from the colored evidence tidbits and suppositions. Because the organization of the evidence tidbits, suppositions and links were not predefined, the filled cursor worked well in showing the selection area and whether items fell inside it. Because the cursor was filled, it was also translucent so as to not obscure on-screen objects from view. Object color was not crucial; the colors of the evidence tidbits and the suppositions differed enough that overlaying them with the translucent cursor did not affect the color so much that the two types of object could not be differentiated, and text could still be read.

Selection level feedback should be given with the nature of the dataset in mind. Color should be chosen that is easy to distinguish from the background and the data itself. Data should not be seriously obscured while being pointed to. If color is crucial, for instance as with the puzzle solving example, a cursor with no fill color that outlines the selection area should be considered.

Enabling Other Actions

Selection of individual notes enables users to perform actions on those notes, such as deleting them, linking them and editing them. We implemented linking and editing as different modes, representing linking mode with a crosshair cursor and editing mode with an arrow cursor, as we thought both actions would only be done on the detail level. Another action that could be performed on the notes which we did not implement would be to create groups of notes, allowing users to apply a hierarchical structure of notes within groups and groups within other groups to a previously unstructured dataset. If this were possible, it would make sense to enable linking on different levels of scale, allowing users to both link one note to another and one group of notes to another group of notes.

3.4.3 Results

There was one great difficulty with using intelligence analysis as the domain -- truly helpful feedback on the tool and the interaction necessitated bringing in actual analysts to evaluate it, which we were unable to do. Analysis of evidence is difficult to do well, and analysts must be trained in order to hone this skill. None of us having prior experience in intelligence analysis, we evaluated the Storyboard tool to the best of our abilities by performing realistic tasks ourselves, using the tool. We began standing close to the display, pointing at it and posting all evidence tidbits randomly throughout the space. We then selected each tidbit individually, read it, and moved them into related groups in different locations around the space. At this point in the process, we remained close to the display to do these detail interactions. The majority of our movement was to the right and left across the length of the display. As the groups grew in size, we began to move back to select entire groups at once to move them further away from one another.

After all the tidbits were grouped, we began to use the on-screen keyboard of the PDA to type in suppositions and post those to the display. We then linked evidence tidbits to these suppositions, and moved all the notes around to form a tree-like structure. As we built larger trees, we stepped back to select entire subsections at once to move them around. The interaction supported all of these tasks, some better than others. Text input via PDA was somewhat awkward. The size of the cursor fit the detail tasks very well, though in some cases, the square shape did not fit the desired selection on the overview level, the subtrees having one dimension larger than the other. For these cases, we used another button on the PDA as a “control-click”; the button is held down, and additional selections are added to the original selection group.

3.5 Summary

Multiscale interaction has its basis in the idea that users must already be moving in order to see all the information present in large display environments. Not only that, multiscale interaction will benefit users because it improves on other multi-selection techniques in three ways. First, it exploits the physical navigation that users are already doing by linking it to the interaction scale. Secondly, multiscale interaction leverages the users' sense of embodiment, enabling them to apply knowledge they already have about the world around them. Finally, multiscale interaction simplifies the user interaction cycle by reducing visual and interactive clutter.

Multiscale interaction can be applied to the two subtasks within interaction: selection and navigation, in addition to collaborative or multi-step interactions. We discovered many important design issues to consider by further exploring one area of this design space through a case study on a storyboarding application. These issues include: pointing technique, cursor size and shape, and selection area feedback.

Chapter 4: Experiment

4.1 Motivation

This implementation of multiscale interaction was developed to further explore one aspect of the design space. The motivation of this study is to begin to evaluate whether multiscale interaction is an effective technique. As discussed in Chapter 3: multiscale interaction can be applied in many different ways and it is probable that it affects these various applications differently. Therefore, we wished to determine if the particular multiscale interaction technique used would show improvements over other techniques that do not exploit users' physical navigation. To this end, we wished to examine several things. Multiscale interaction enforces a link between users' physical navigation and their interaction scale. Is this natural? Do users tend toward such interaction anyway? If so, is this something we should actually encourage? Do users in fact benefit from the linkage of physical navigation with interaction? We hypothesized that they would, for several reasons:

1. Multiscale interaction exploits the fact that users are already physically navigating, perhaps even encourages it. We hypothesize that the same time performance benefits seen from physical navigation with perception tasks will hold for interaction tasks as well.
2. There is a reason why users prefer physical navigation over virtual. This is because it is more natural. We believe this "naturalness" will manifest itself in both a reduced mental workload and tendency towards similar interaction in the other, non-multiscale interaction techniques.

3. Multiscale interaction is more natural because it reinforces for the user a sense of embodiment, leveraging his preexisting knowledge about the world around him. In addition it reduces both the visual and interactive clutter found in other techniques. Making these aspects easier frees up cognitive resources users may then apply to better problem solve, manifesting itself in higher accuracy.

To evaluate these hypotheses, we compared the multiscale interaction implementation to two other techniques, one representing the basic multi-selection tool found, for example, in Windows operating systems and another which allowed interaction on different scales, but did not link it to physical navigation.

Because the technique was designed to take advantage of several aspects, including users' physical navigation and sense of embodiment, we knew it would be tricky to define what user performance is in this sense. We hypothesized that this 'improved performance' is due to a more natural technique, but how do you measure 'naturalness'? We hypothesized that this improved performance would manifest itself in several areas. First, we designed the technique to take advantage of users' physical navigation, which we have already seen improves users' time performance versus using virtual navigation. Therefore, we hypothesized that we would see a similar increase in performance time with multiscale interaction as compared to the other interaction techniques. Secondly, we both leveraged the users' preexisting knowledge of the world around them to create an embodied technique and reduced the complexity of the visual display of the cursor. This, we hypothesized, would reduce the cognitive load on the user, freeing up resources and allowing the user to better problem solve, evidencing itself in increased accuracy. Finally, we defined 'naturalness' as the degree to which users tended to such interaction, measured by the correlation between their distance from the screen and their current interaction level.

Designing a study to capture both the quantitative performance aspects and the "naturalness" aspect also posed some difficulties. We included both short and long tasks in the study, the former being focused on recording quantitative performance measurements for each technique, while the latter collected longer-term performance measures, but also recorded user movements to help evaluate "naturalness".

4.2 Method

During the experiment, users were asked to solve a very large multiscale puzzle on a nearly 100 megapixel display using one of three different interaction techniques. One technique implements multiscale interaction while the other two are common approaches to interaction. They include the explicit technique, which allows users to select interaction level via menu, and the lasso technique, which requires users to draw a box around the area they wish to select. Users performed two short tasks and one long task, the former beginning with the majority of the puzzle presolved. For both tasks, we measured time and number of piece swaps made. In addition, for the long tasks, we recorded data about users' movement within the space.

4.2.1 Hardware and Software Used

Participants performed the tasks on our large, high resolution display, nicknamed the "Gigapixel" display (Figure 4.1). This display consists of 50 tiled 1600x1200 pixel LCD monitors, arranged in ten columns of five monitors each. The columns are freestanding so that the display is reconfigurable; however, for the purpose of this experiment, the display remained in a flat configuration. The 5x10 grid of monitors provides a single display surface with a total screen dimension of 16000x6000 pixels, or 96 million pixels (96 megapixels). The display itself is run by 25 Linux computers, with an additional Linux machine running the head node. The size of the entire display is roughly 4.4 meters wide by 1.7 meters tall.

Figure 4.1: The Gigapixel display

The Gigapixel display is located so that the area in front of the display is open, allowing users of the display to move around freely. This space, which is roughly 3.5 meters wide (parallel to the display) by 3 meters deep, is tracked by a Vicon

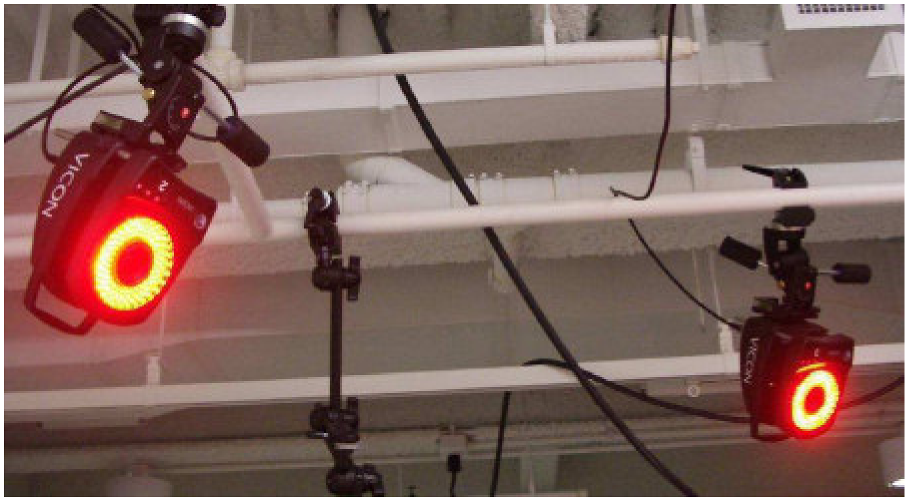

Figure 4.2: Vicon infrared motion tracking cameras

MX motion-capture system [1], consisting of eight cameras. The Vicon system is a six degree-of-freedom, near-infrared vision based tracking system, which uses retro-reflective balls as positional markers. For the purpose of this experiment, we used a simple metal rod instrumented with three reflective markers as a pointing device, or wand (Figure 4.3). We similarly marked the display, allowing us to track its plane. Using this plane and the silhouette formed by the Base and Tip markers of the wand, we used a simple ray-casting technique [9] to project a line out of the tip of the wand to the plane of the display, therefore calculating the intersection point on the display.

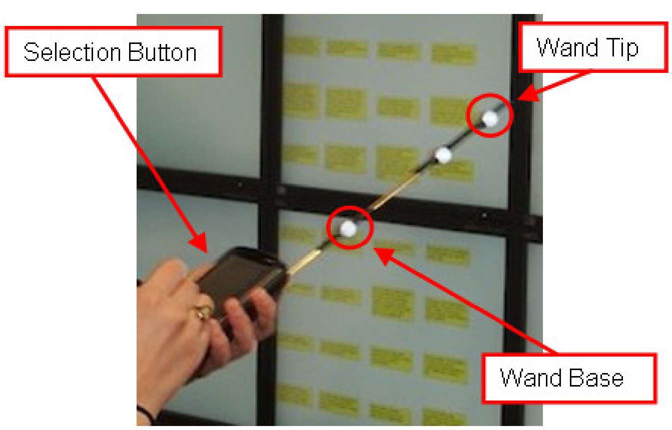

Figure 4.3: Interaction system used for the experiment

In addition to the wand (used for pointing), participants also carried an HP iPAQ personal digital assistant (PDA), which was used for selection. Depending on the interaction type condition, participants saw different on-screen controls (Figure 4.9); however, all participants used the side button to select items during the experiment tasks.

4.2.2 Experimental Design

The independent variable was interaction technique type. There were three interaction techniques: physical navigation, explicit and lasso. These technique types are described in more detail below. This was a between-subject variable to avoid any learning effects that might occur with task repetition, and because we wanted to ask participants to complete a long puzzle-solving task, with an average completion time of around 20 minutes. If participants were asked to complete three of these in succession, they would become fatigued and task performance would likely quickly degrade.

Each participant was asked to complete two short and one long puzzle solving tasks. The order of the two short puzzle-solving tasks was alternated so that half of the participants performed one short task first and the other half performed the second short task first.

For each interaction technique type we had eight participants, for a total of 24 participants. All participants were either undergraduate or graduate students. Nearly all participants were computer science majors with the one exception being an industrial and systems engineering major. The average age of the participants was 28, with a minimum age of 21 and a maximum age of 37. Nineteen of the participants were male and five were female. Participants had a wide range of prior experience working with large display environments, from no prior experience to having conducted their own experiments with large display environments.

4.2.3 Tasks

When we were designing this experiment, we decided that the task participants would complete must be multiscale in both the visual sense and the task sense, in order to properly evaluate the multiscale interaction technique. By multiscale in the visual sense, I mean at different levels of zoom, different amounts of information are available. For example, details present in the visualization fade or become indistinct when far away, but are clear when close up. By multiscale in the task sense, I mean there are different types of applicable tasks at each hierarchal level or information scale of the visualization. In addition, we wanted the tasks to deal with a hierarchical dataset, because the data is already ordered, and the selections are well defined. For these reasons, hierarchical selection is much easier to evaluate.

To achieve this we designed a hierarchical puzzle that we would ask the participant to solve. In order to solve a puzzle, the participant must swap various pieces, placing them so they form a final image. The pieces in the puzzle are arranged hierarchically, so that smaller scale level pieces may only be swapped within the confines of the next highest up scale level "parent piece."

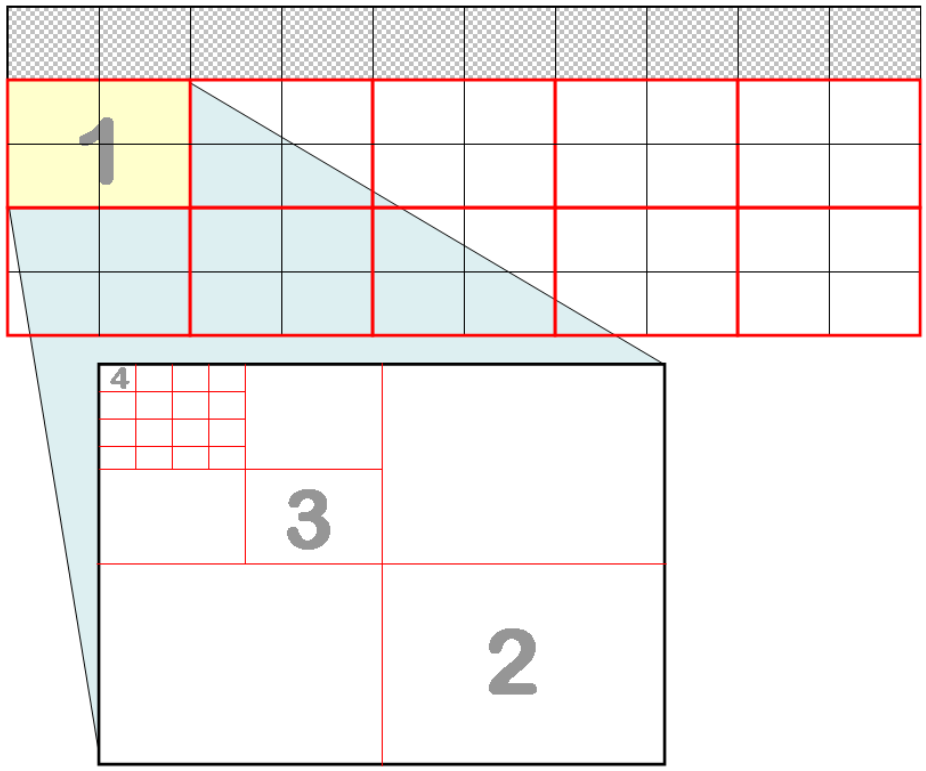

Figure 4.4: Hierarchy of puzzle piece sizes

This task was designed to represent other typical scenarios that could utilize multiscale interaction, specifically the subset in which the dataset would be pre-organized hierarchically, and where the size of the cursor influence area would vary according to those predefined hierarchical levels. The actually scenarios tend to be too complex for evaluation purposes. We needed a testbed with a shorter task completion time and a defined goal with an optimal solution in which we could perform specific measures.

Level 1 size pieces are the largest, encompassing a 2x2 monitor square of the display. There are 10 of these size pieces in a puzzle, and can be moved anywhere. Within each of these Level 1 size pieces, there are four Level 2 size pieces, encompassing one monitor each. There are 40 of these size pieces in a puzzle; however, each group of four may only be moved around within its parent Level 1 size piece. Within each Level 2

Figure 4.5: The beginning state of one of the short tasks. Out of place pieces have been highlighted.

size piece, there are four Level 3 size pieces, so that each is the size of a quarter of a monitor. There are 160 of these size pieces in a puzzle; however, each group of four may only be moved around with its parent Level 2 size piece. Within each Level 3 size piece, there are 16 Level 4 size pieces, the smallest scale level. There are 2560 of these size pieces; however, each group of 16 may only be moved around within its parent Level 3 size piece.

Figure 4.6: The (a) beginning and (b) ending states of the long puzzle task

Participants were asked to complete a series of these puzzle-solving tasks. The first two were short tasks, requiring the participant to only make two swaps on the largest scale level and two swaps on the smallest scale level. The last task was longer, requiring the user to make many swaps on all scale levels, in order to solve a puzzle that began with fewer pieces in place. Because of the difficulty and time commitment required to solve a puzzle from scratch, none of the puzzle tasks began at this state, always starting with the puzzle partially solved to some extent. Each of these tasks required the participant to solve a puzzle with a different final image and a different beginning piece shuffle state. The individual task final images and beginning piece shuffle states remained constant across all participants.

The short puzzle tasks were designed specifically to evaluate users' efficiency at switching among the various interaction scale levels. The long task also looked at these things (over the more long-term), but it was also designed to capture richer movement data, to help evaluate the "naturalness" of the technique.

The dependent variables for these tasks were time, accuracy (measured in total number of swaps), and total movement, a combined measure of in/out and sideways movement (taken from Vicon tracking data).

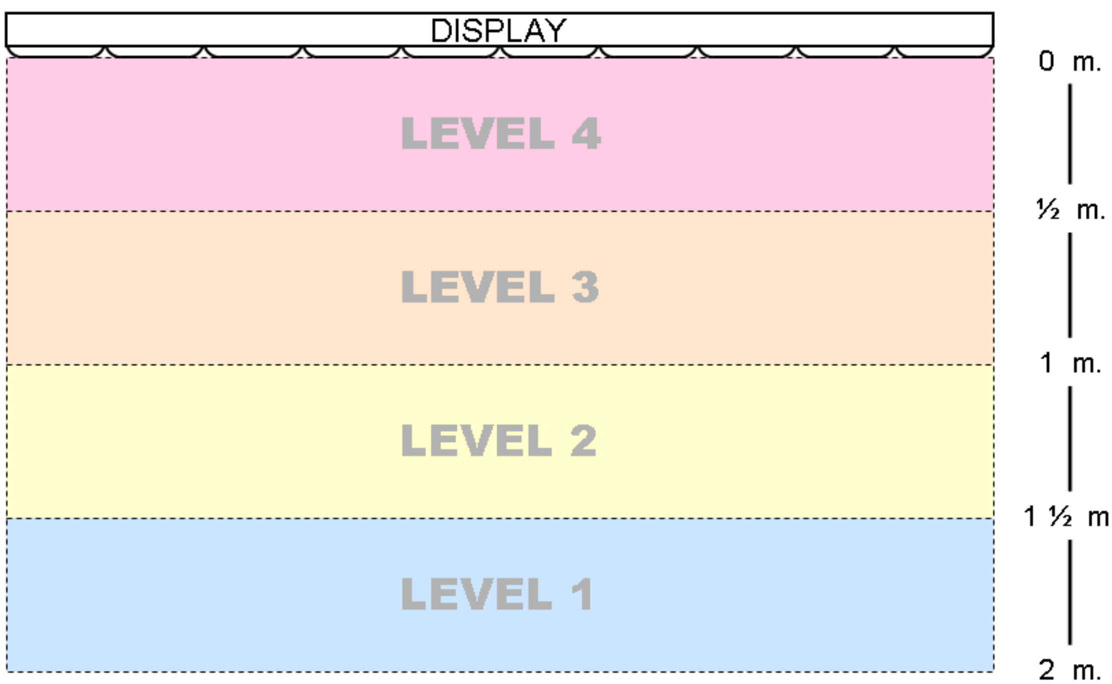

Figure 4.7: Interaction scale levels and distance from display for physical navigation technique (Note: figure is not to scale)

4.2.4 Interaction Technique Types

As previously mentioned, we varied the interaction technique the participants used to complete the tasks. There were three techniques used: physical navigation, explicit and lasso.

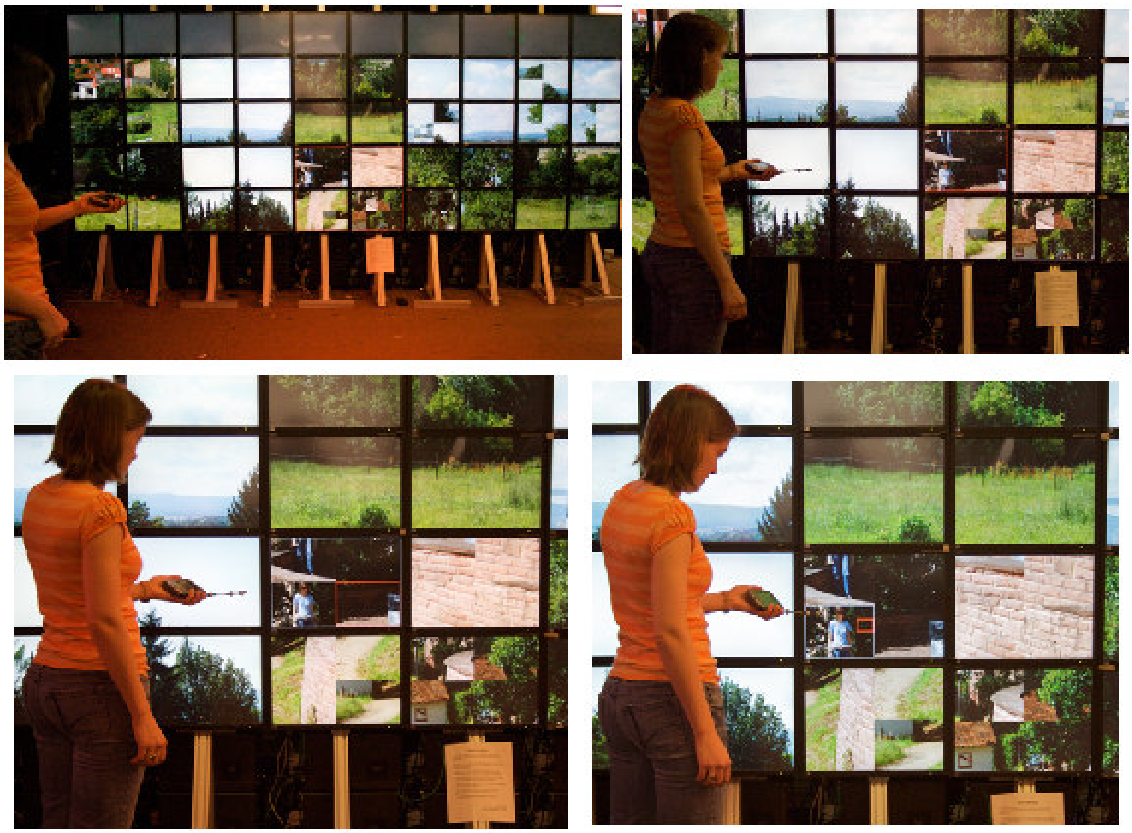

Figure 4.8: The four hierarchical selection scale levels

Physical Navigation: The *physical navigation* technique refers to the implementation of multiscale interaction being evaluated in the study. The users' cursor automatically changes size (and therefore the scale at which the users are interacting) depending on their distance from the display. Figure 4.7 shows the interaction scale levels used during the study. A portion of the available space was split into four equal parts. When the users were within ½ meter of the display, they were able to select the smallest size puzzle pieces. Likewise, if they were 1.5 meters or more from the display, they were able to select the largest size pieces.

Explicit: Much like the physical navigation technique, the cursor for the explicit technique can discretely change size, allowing the user to interact with information on the display at different levels of scale. However, changing between the various scales must explicitly be done by the user via a menu on the PDA (Figure 4.9). This allows users to switch from any one level to any other level, without requiring them to visit each of the levels in between. This technique differentiates between the effects of the multiscale nature of the cursor and its coupling with user movement on performance.

Figure 4.9: On-screen PDA controls. Participants using the physical navigation and lasso interaction types had no on-screen controls (left). Participants under the explicit condition used the on-screen menu to switch among interaction levels (right). All participants used the side button (marked with an arrow) to select items.

Lasso: The *lasso* technique keeps the cursor size constant, and requires the user to manually select all items by drawing a box around them. This technique is reminiscent of more conventional multi-selection tools found in operating systems like Windows.

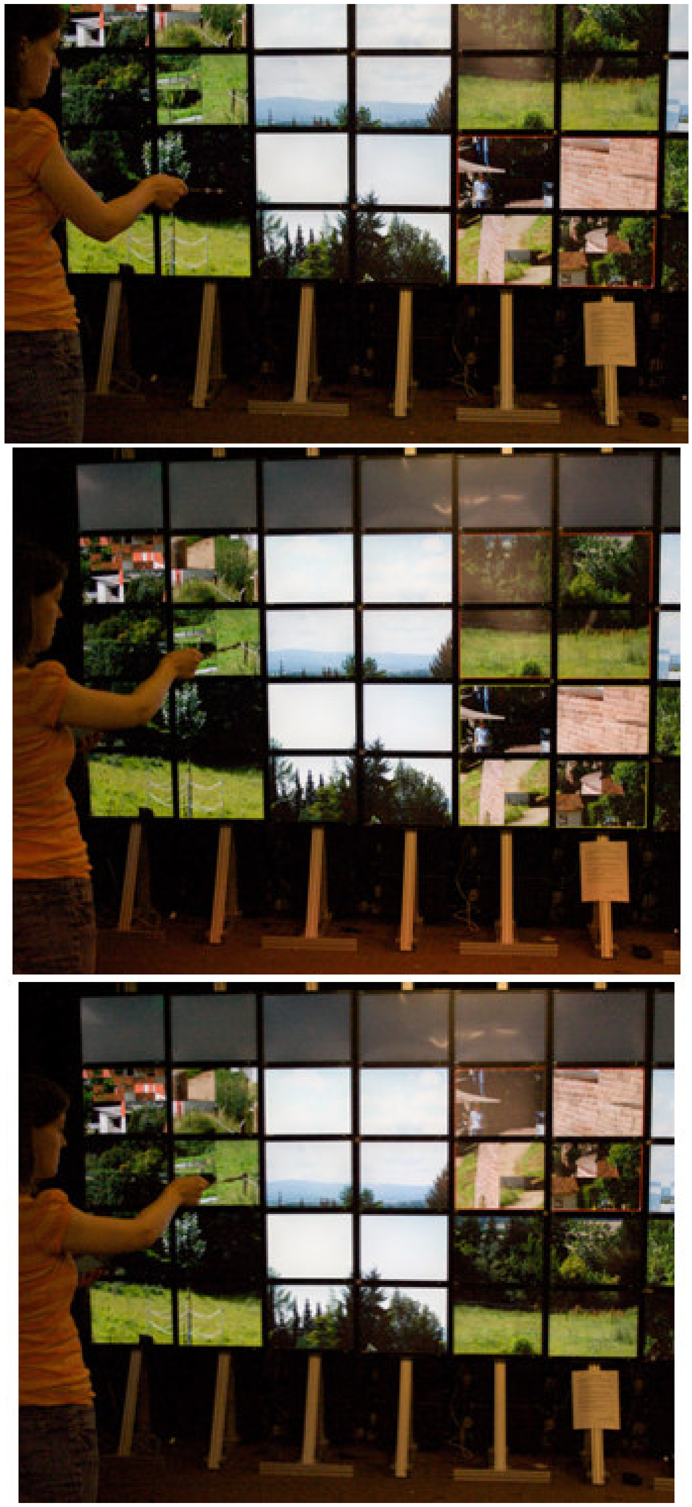

Figure 4.10: Swapping two pieces. The user (a) selects the first piece and it is (b) highlighted in yellow to indicate selection. The user (c) selects the second piece to complete the swap.

For the physical navigation and explicit techniques, the cursor can change among four discrete sizes. It appears on the display as a thick, red box, and jumps from item to item when the 'actual' intersection point crosses the boundary from one item into another. If the current scale is not the largest, the next highest up scale level in the hierarchy is indicated by a thin, white box. Once the chosen item is enclosed by the cursor, the user clicks the PDA button to select the item. Selection is indicated by a thick, yellow box.

For the lasso technique, the cursor is represented by a 30 by 30 pixel square drawn at the 'actual' intersection point. Every scale level of piece currently being pointed at is indicated by a thin, white box. To select an item, the user must click and hold the PDA button, dragging a box around at least 56 percent of the chosen item. When the user wishes to select a level 4 size piece, this translates to merely clicking inside the

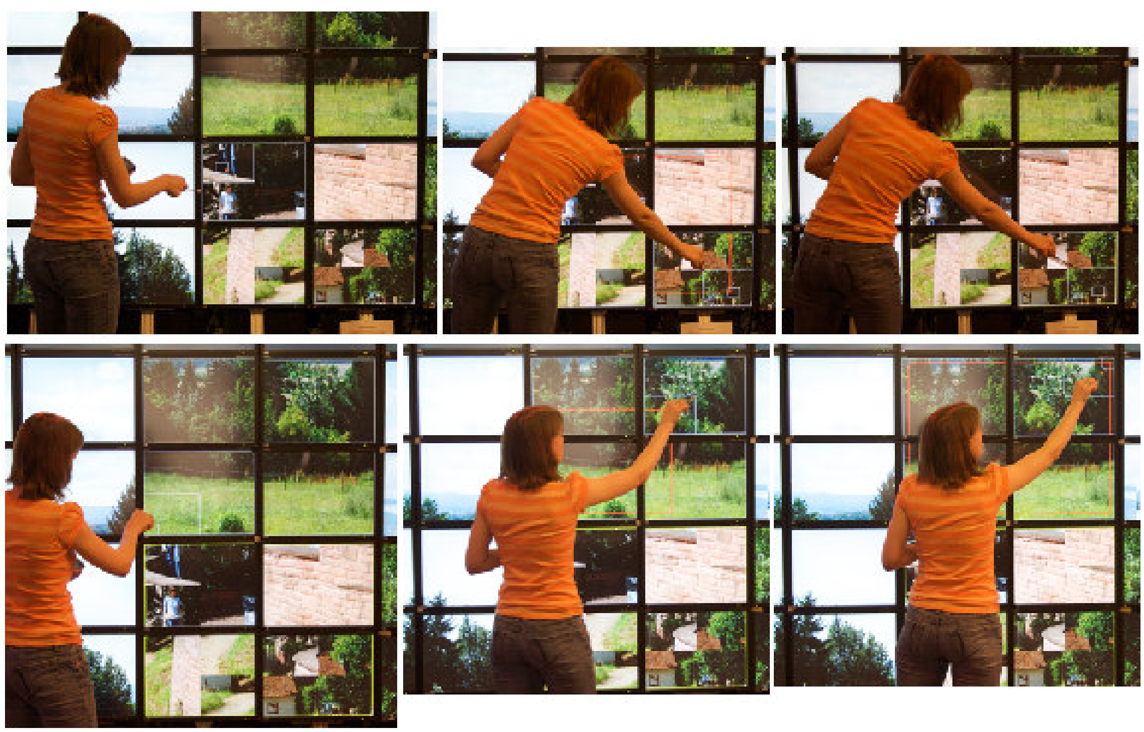

Figure 4.11: Swapping two pieces using lasso.

piece. Again, selection is indicated by a thick, yellow box.

With all techniques, if the user mistakenly selects one item, it may be unselected by "reselecting" the item in the same manner. Two selections must be made in order to perform a swap of the two items. To swap two items, the user must first select one item, and then select a second, and the two are then immediately swapped.

4.2.5 Procedure

Each participant took no more than an hour and a half to complete the experiment. The short tasks took no longer than five minutes each and the long task took no longer than 45 minutes, as these were the timeouts for each task. Few participants approached these timeouts.

At the beginning of the experiment, participants were asked to complete an embedded figures test (EFT) [38]. The EFT is designed to assess a subject's perceptive ability to pick out features within a complex pattern. Results from this test were used to gauge a participant's puzzle-solving ability. Perfect score on the test was 18, with the median being 17 and the lower quartile being 15. Because the scores were so tightly clustered on the high end of the scale, they were not considered when evaluating the other experiment results. Participants were also asked to fill out a short demographic questionnaire. They then were given instructions explaining the task, the equipment and the interaction technique and allowed a practice session (which resembled the long task) to solve a sample puzzle so that they may become comfortable with all three. Participants were encouraged to ask any necessary questions during this practice session. Instructions were to complete each task as quickly as possible. Participants were given a picture of the final image for this practice session and the long task.

After the participants announced they felt comfortable with everything, the experiment began. They first completed two short puzzle solving tasks and then the long puzzle solving task. After all tasks were completed, participants were asked to complete the NASA Task Load Index (NASA-TLX) rating workload.

4.3 Results

There were interesting findings related to completion times, accuracy (swaps), workload measures and various movement measures, including total movement, correlation with interaction scale and user strategies. Because this study evaluates an entirely new kind of interaction technique, there was little else to build on and compare it to, and it was thus exploratory in nature. In addition, embodiment is a very difficult aspect to measure. For these reasons, we feel it is still important to report non-significant results or trends. Even though some of the results are not strong statistically, they raise some interesting potential issues.

4.3.1 Completion Times

Task completion time was measured for both the short tasks and the long task.

Long task

We performed a one-way ANOVA on interaction type. There are no significant differences in completion time between any of the interaction types.

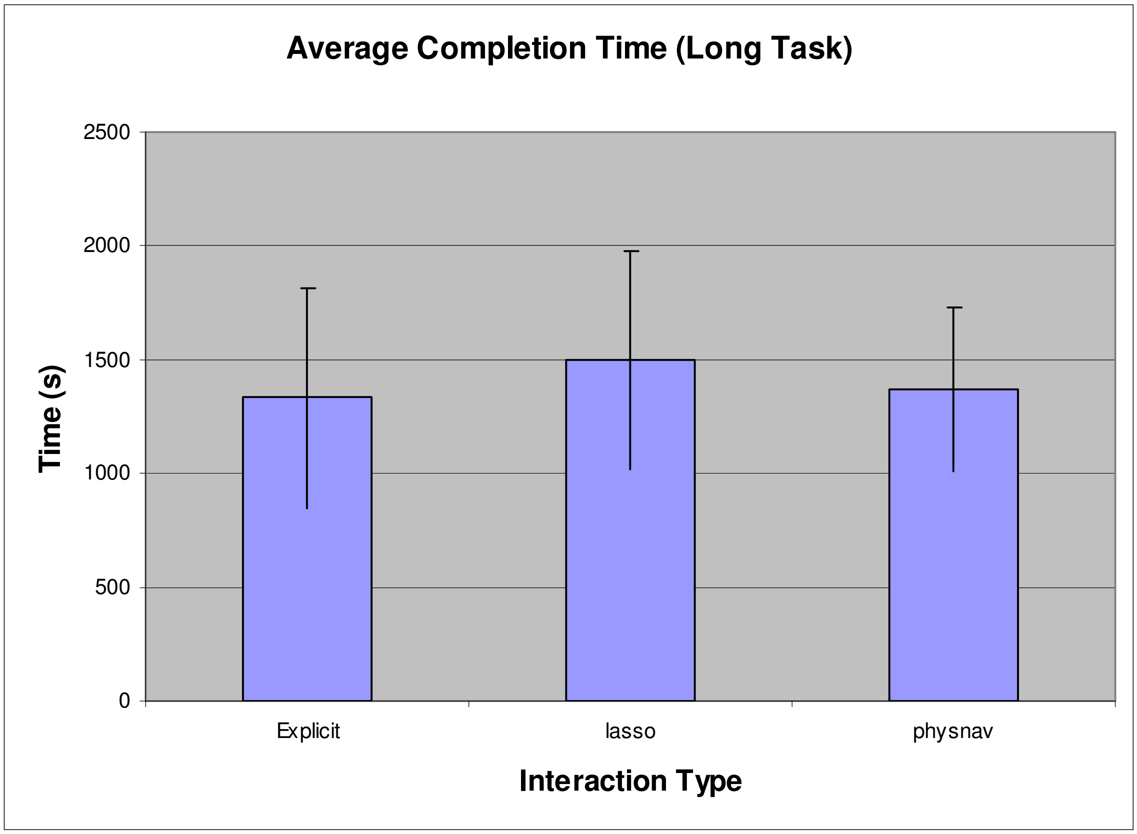

Figure 4.12: Average Task Completion Time (Long Task) by Interaction Type

Interaction Type	Mean	Std Dev
Explicit	1332.5	481.385
Lasso	1499.5	478.138
Physical navigation	1375.0	359.538

Table 4.1: Means and standard deviations for long task completion times (in seconds)

However, in closer inspection of the individual participants, we can look in detail at the spread of completion times. Unlike the explicit and lasso types, the physical navigation type appears to have escaped having outlier participants with unusually high completion times. This may indicate that the physical navigation type allows users to perform a task more consistently in the same amount of time.

If we color the markers according to user experience with large displays, we can see that the users with extensive experience (yellow) have consistently lower times, which is not surprising. However, we can also see that the outliers in the explicit and lasso types are participants who have lesser experience with large displays. This seems to indicate that more specifically, the physical navigation type allows less experienced users to perform more consistently on tasks as those with more experience.

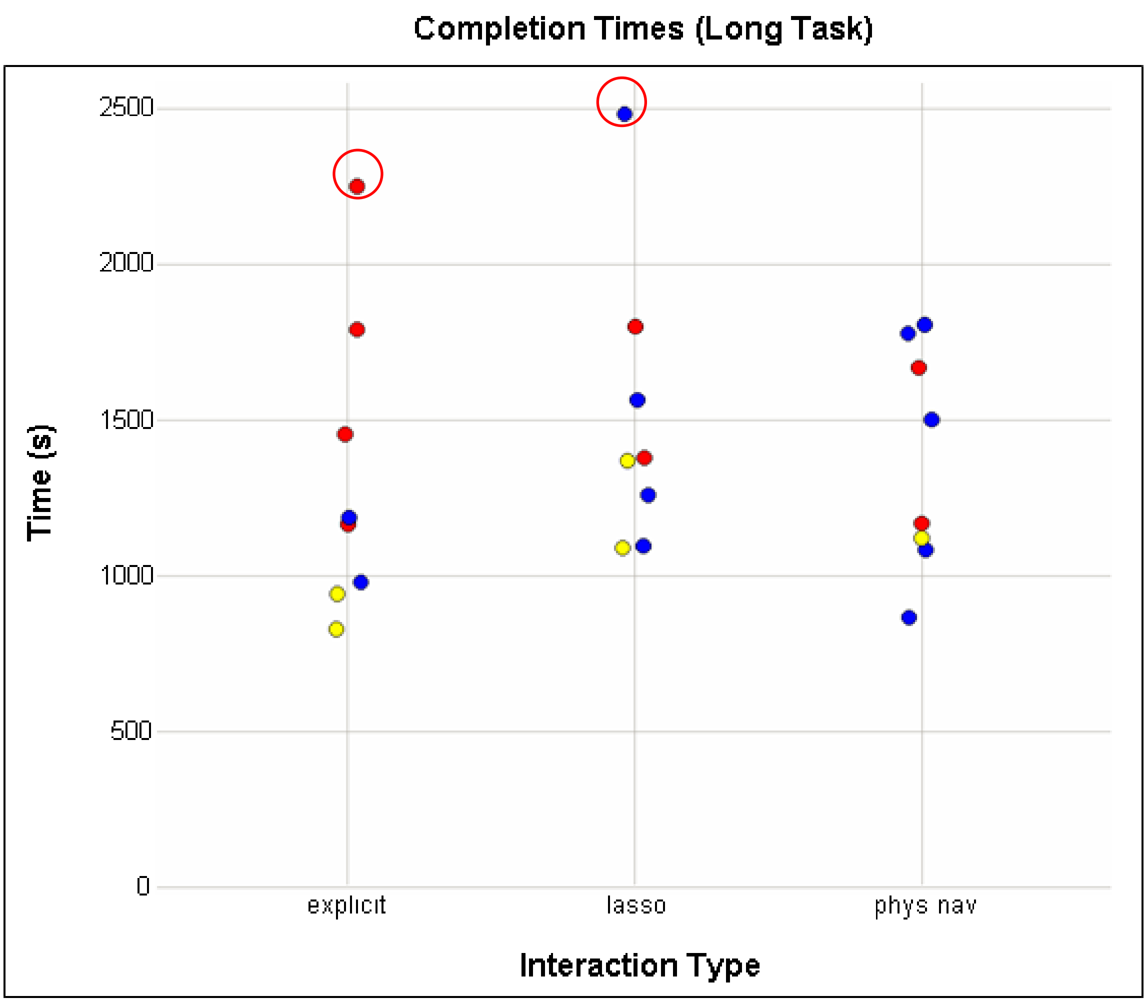

Figure 4.13: Task Completion Time by Interaction Type (Long Task). Markers are colored by large display experience (none – red, med – blue, expert – yellow). Interesting outliers are indicated in red.

Short Tasks

We performed a one-way ANOVA on interaction type. Again, there are no significant differences in completion time between any of the interaction types.

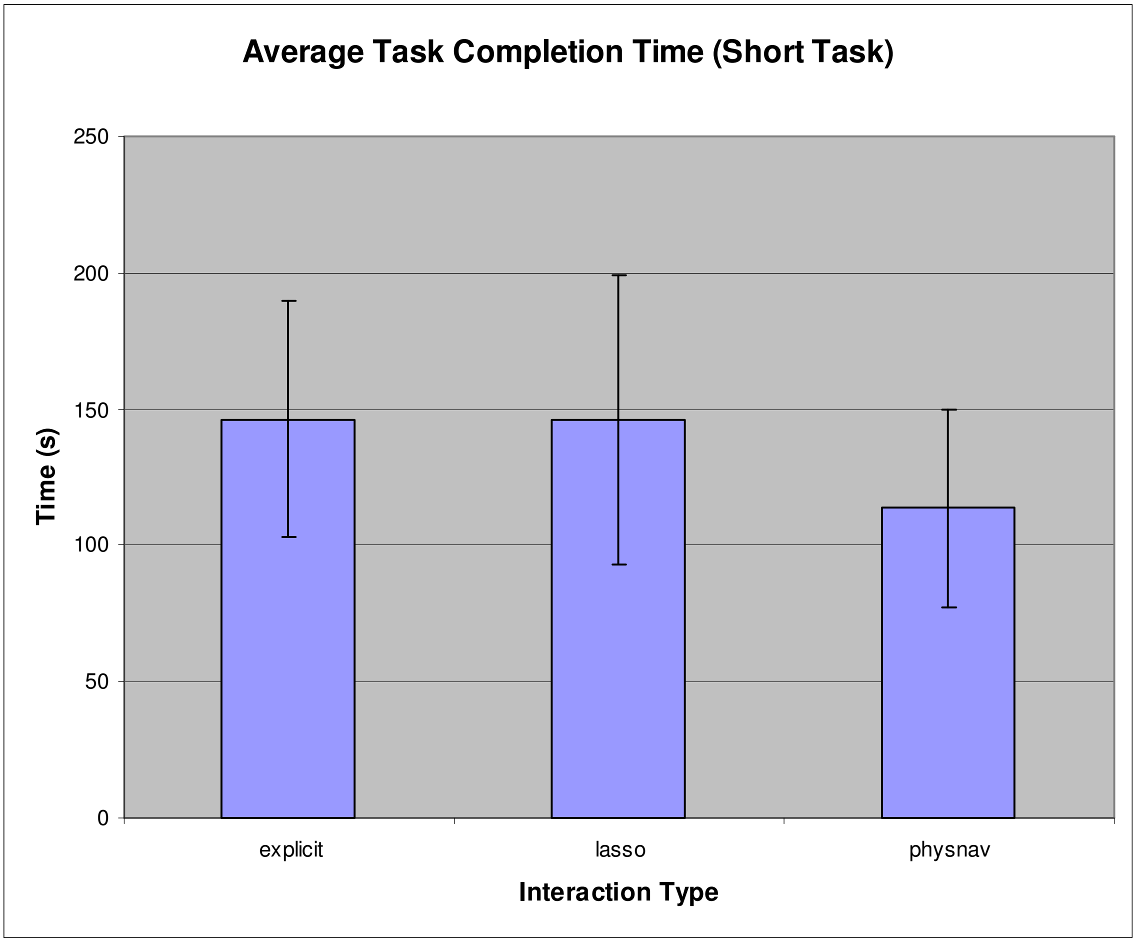

Figure 4.14: Average Task Completion Time (Short Task) by Interaction Type

Despite the fact that the differences in time are not significant, there seems to be a trend that the physical navigation type might be different from the other two types, ($F(2, 21)=1.4207$, $p=0.2638$). Further analysis using the student's T test reinforces this possible trend of physical navigation being different from explicit ($p=0.1573$) and lasso ($p=0.1611$).

We performed an analysis of variance (ANOVA) on interaction type that showed that total distance moved was weakly significant ($F(2,18)=2.6739$, $p=0.0962$). Further analysis using the student's T test showed a significant difference between the explicit and lasso types ($p=0.0377$).

The short tasks were not designed to be difficult, having only two swaps on both the largest and smallest scale levels, but instead were designed to really test the performance difference between interaction types with respect to switching among scale levels. However, because there were so few Level 4 pieces out of place, it was difficult to find them. Despite the fact that the researcher informed the participants where these pieces were prior to the task beginning, some participants still had difficulty finding the out of place pieces. We believe that this might have skewed the short task completion times.

4.3.2 Accuracy (Swaps)

Total swaps and swaps per scale level were counted for both the long task and the short tasks, with more piece swaps considered less accurate.

Long Task

We performed a one-way analysis of variance (ANOVA) on interaction type that showed that overall piece swaps was weakly significantly different ($F(2,21)=3.21$, $p=0.0606$). Further analysis using the student's T-test showed a significant difference between the explicit and lasso types ($p=0.0224$), as well as a weakly significant difference between the physical navigation and lasso types ($p=0.0962$). Participants using the lasso type made statistically significantly fewer piece swaps than either explicit or physical navigation participants.

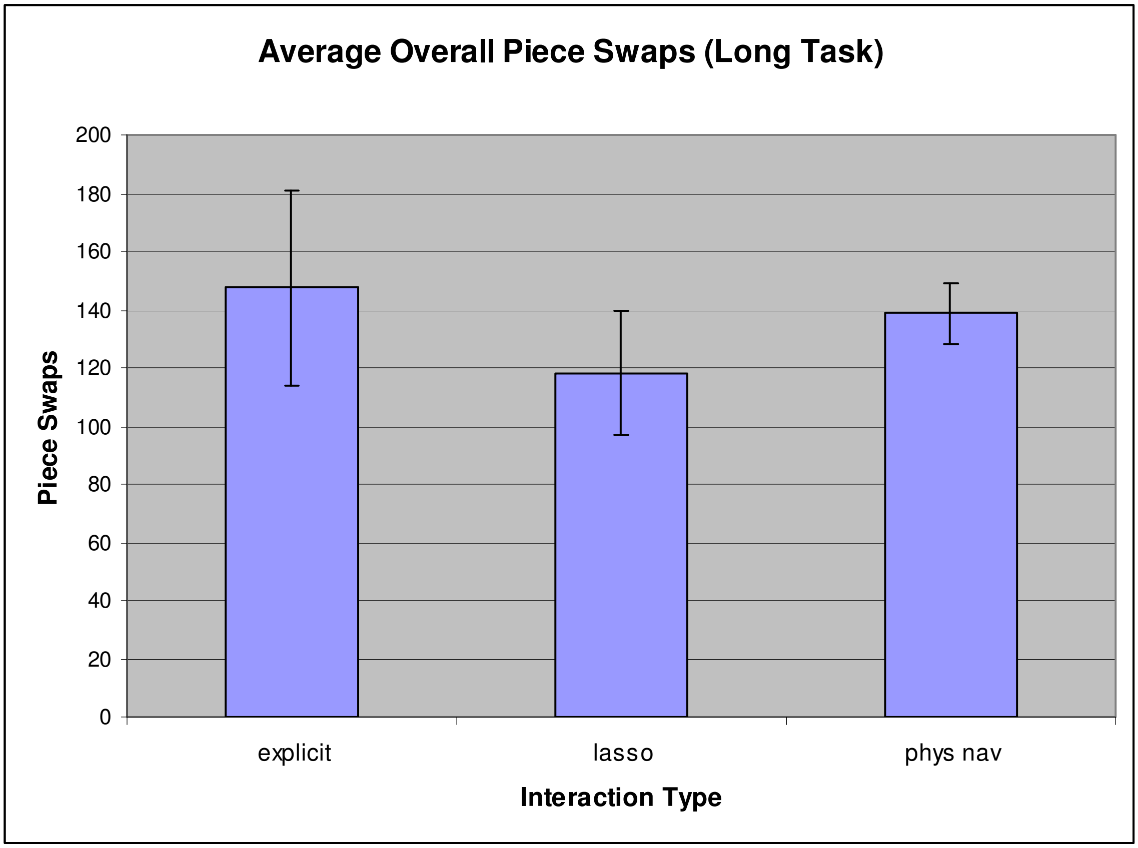

Figure 4.15: Average Overall Piece Swaps by Interaction Type.

If we look more closely at data for the individual participants, there are several interesting findings. First, there appears to be a large amount of variance in the number of swaps made by explicit participants, and to a lesser extent, the lasso participants as well. The physical navigation participants seem to be very consistent in their swap performance.

Within both the lasso and explicit types, there are some subgroups. Within the explicit type, there are two outliers who have made a distinctly larger number of swaps than any of the other participants. From looking at swap counts from the short tasks below, the explicit type had the largest mean number of swaps made. Many of these were “mistake” swaps, due to pointing inaccuracies as explained later. It is possible this effect is also being seen here

There is also a sub-group in the lasso type with a very small number of swaps. This is perhaps because the lasso technique does not resize the cursor according to preset hierarchical levels, so it is more difficult to select on different scale levels; therefore, lasso participants avoid swapping as much as possible.

If we color the markers according to experience with large displays, we can see that the outliers in the explicit type both have no prior experience. We can also notice that both of the most experienced users are in the lasso type’s low-swapping sub-group. The physical navigation technique appears to equalize participants across experience level.

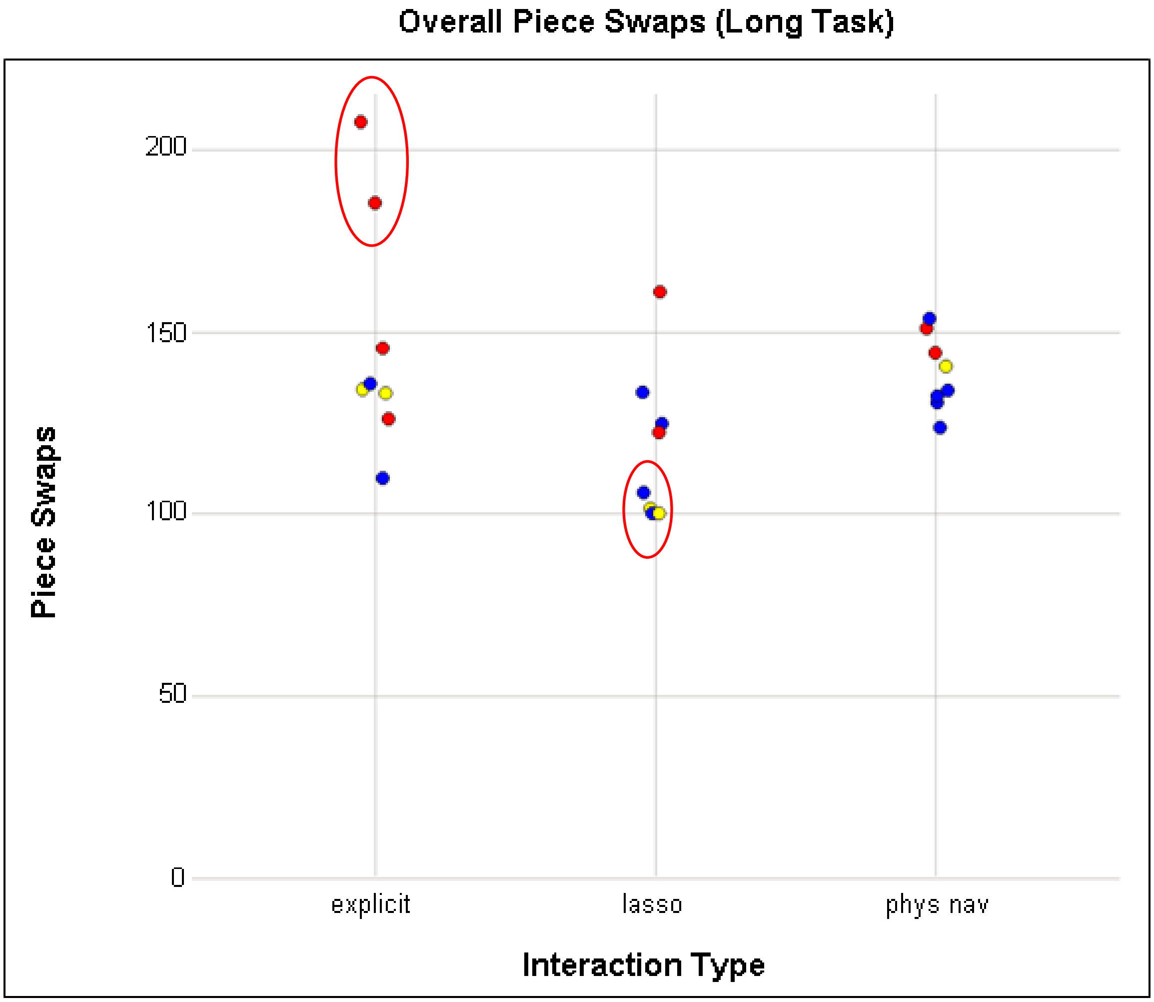

Figure 4.16: Overall Piece Swaps by Interaction Type. Markers have been colored according to large display experience (none – red, med – blue, expert – yellow). Interesting sub-groups indicated in red.

Short Tasks

We performed a one-way analysis of variance (ANOVA) on interaction type. There were no statistically significant differences between any of the types.

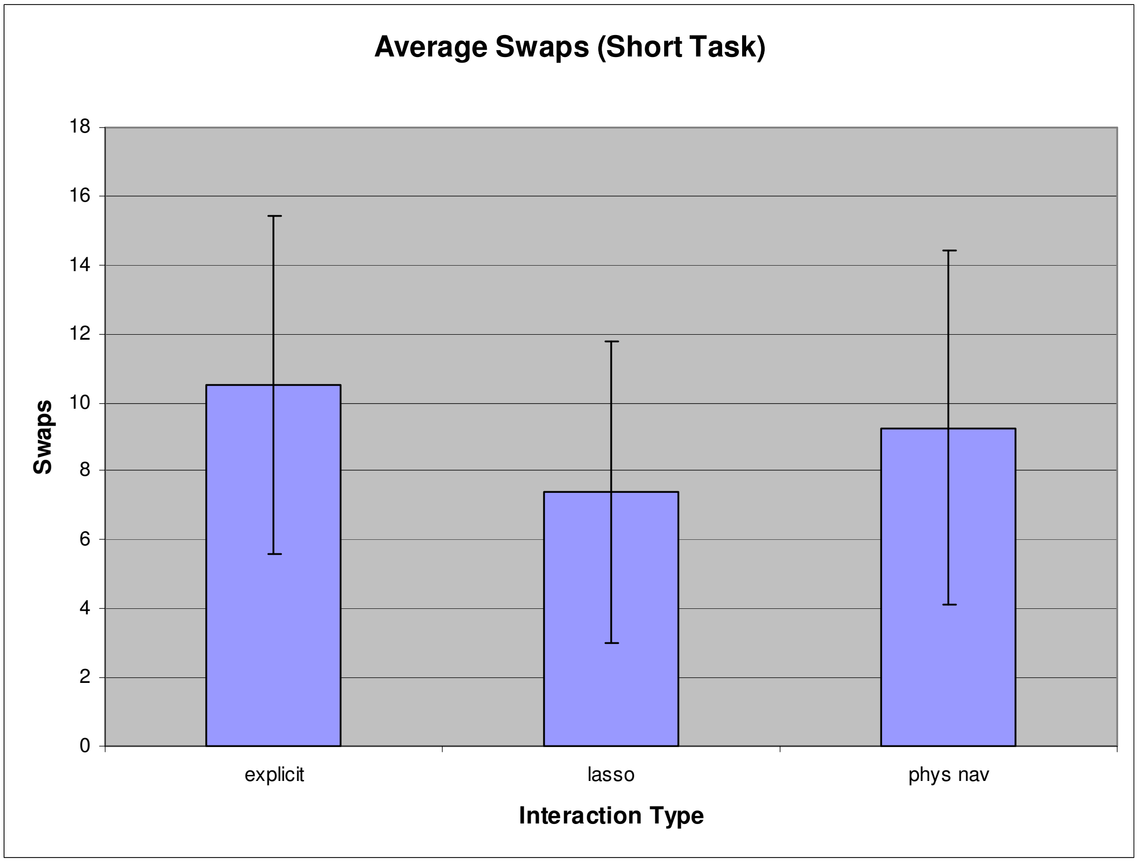

Figure 4.17: Average swaps (short task) by interaction type.

If we look at individual participant data, we can see that the desire by lasso participants to swap as little as possible is even more apparent. This appears to be advantageous for these participants, because they are making fewer mistakes. For these short tasks, participants needed only to make four swaps total, two on the largest level and two on the smallest. However, only two participants achieved this, one from the lasso type and one from the physical navigation type. This is most likely because of selection mistakes due to either user jitter or inaccuracies with the Vicon motion tracking system.

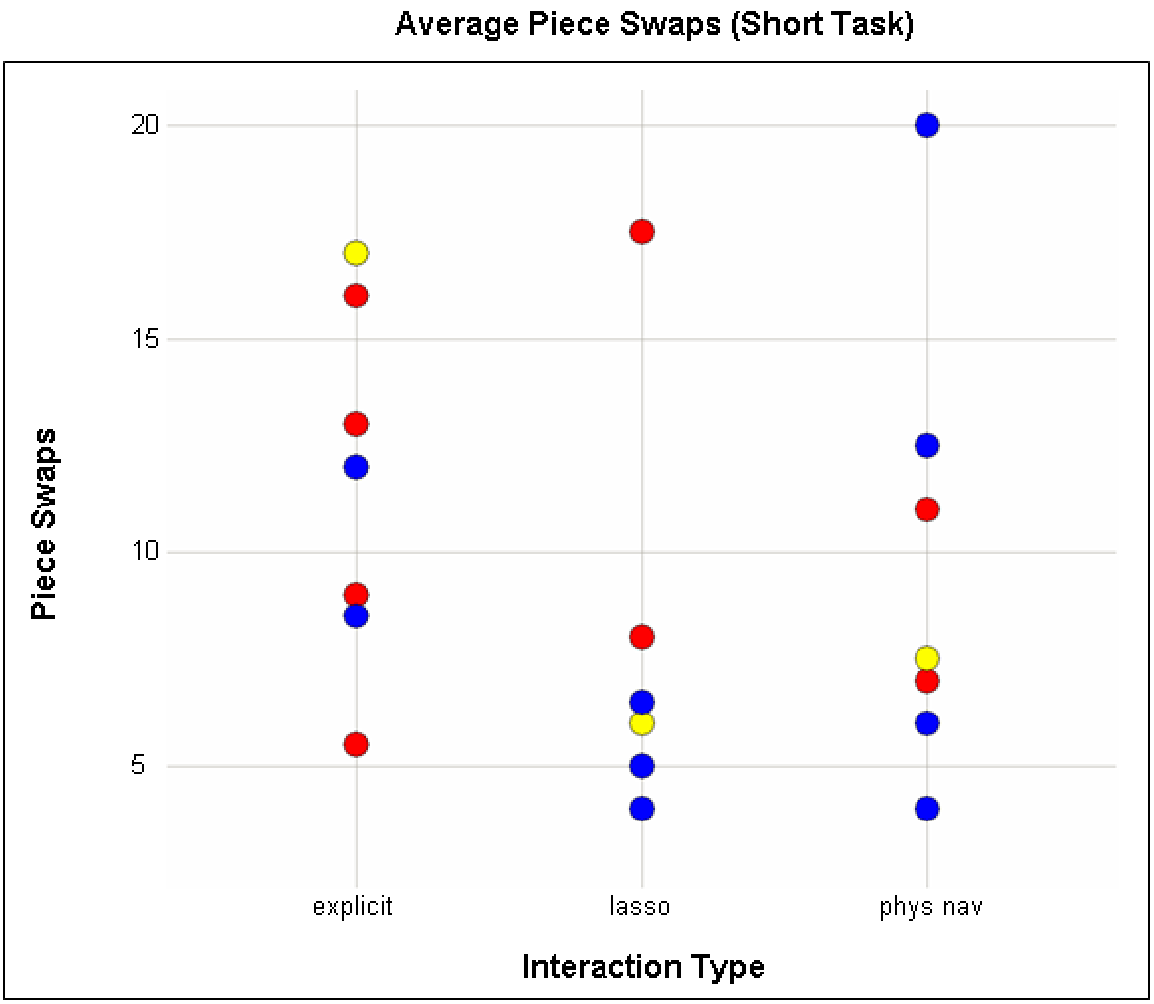

Figure 4.18: Individual participant Average Swaps (Short Task) by Interaction Type. Markers are colored by large display experience (none – red, med – blue, expert – yellow).

Post-experiment, we noted that number of swaps made may not be the best measure of accuracy. We considered more swaps to be less accurate; however, more swaps may be a positive thing, indicating that users were comfortable enough with the technique to make swaps that were, in the overall scheme of the puzzle, "mistakes" in order to test a piece in a particular location, because they were confident they could easily swap the pieces back to their prior locations. With this in mind, we still believe the majority of the mistake swaps seen with the explicit technique were due to pointing inaccuracies, because explicit users did not move closer to select details. However, we can expand our analysis to surmise that the lasso technique discouraging swaps may be a negative effect, while the physical navigation technique is "just right." It still enables the ability to easily swap pieces, as with the explicit technique, but enforces movement that makes it easier for users to point accurately on all levels of scale, thus cutting down on "mistake" swaps due to pointing inaccuracies.

4.3.3 Movement

We recorded information about users' position in the space in front of the display during the long task, using the Vicon motion tracking cameras. For three participants, one from each interaction type, the position data did not correctly record and were removed from movement analysis.

Correlation

We matched this position data with the user's current interaction scale, for every second in task time. We then calculated the correlation between the measured interaction scale, using the Pearson product-moment correlation. Coefficients for the physical navigation type were all close to -1.0, since interaction scale was automatically changed, based on user distance from the screen. Interaction scale did not match up exactly with user distance (with all correlations being exactly -1.0) due to inaccuracies in the motion tracking equipment. For the remainder of the discussion, only coefficients for the explicit and lasso types will be considered.

All correlation coefficients were consistent in sign, except for those of participant number 19. Participant 19 had a correlation coefficient of +0.4857. All correlation coefficients relating measured scale and distance (except for 19) were negative, indicating an inverse relationship: if scale went up, distance went down. This makes sense – at scale level 4, when users are interacting with the smallest pieces, they should be close to the display, and vice-versa. All correlation coefficients relating measured scale to expected scale (except for 19) were positive, indicating a linear relationship: if measured scale went up, so did expected scale.

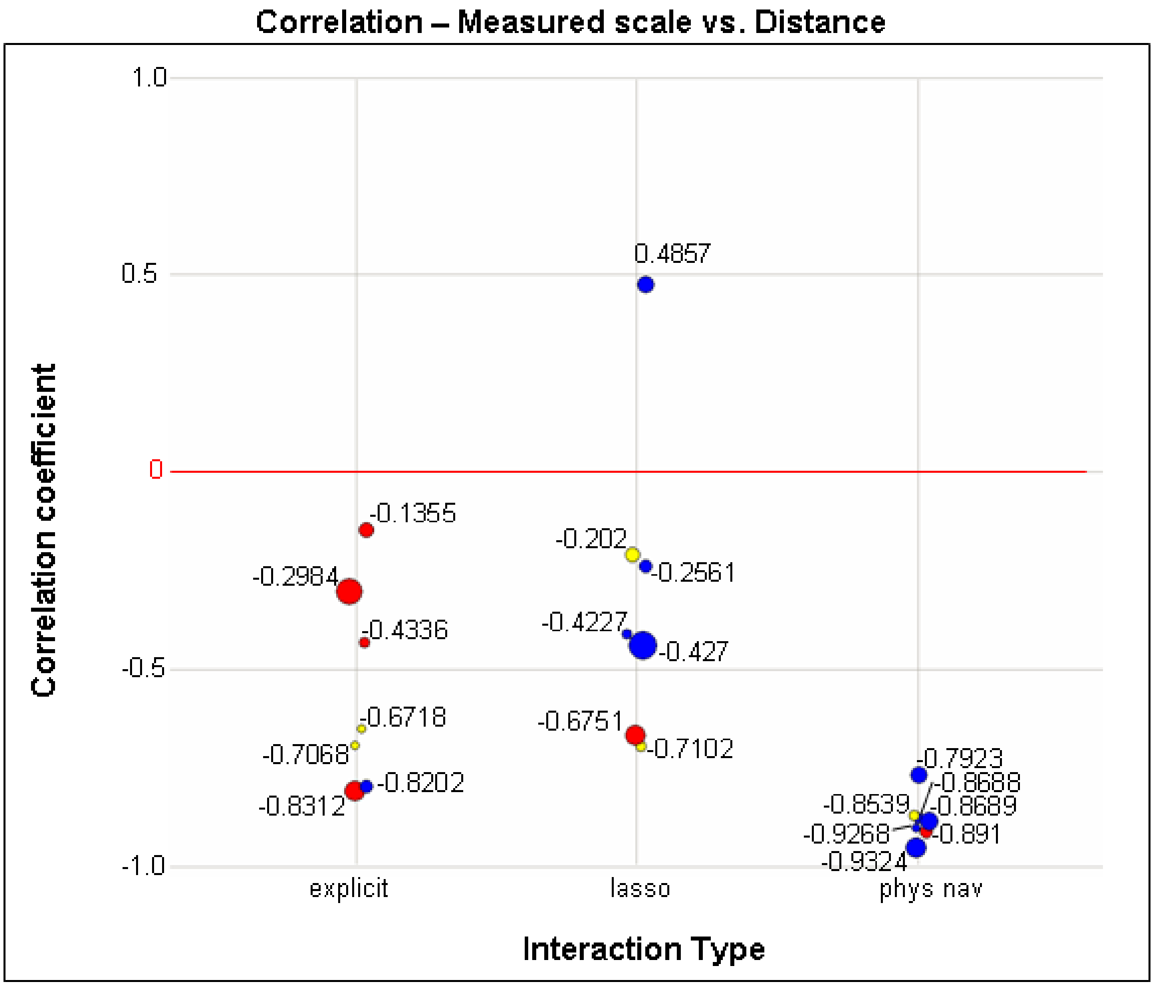

Figure 4.19: Correlation between measured scale and distance, by interaction type. Markers are colored by experience and sized by long task performance time.

Despite the wide range of correlation coefficients, this indicates that there is a natural link between a user's distance from the display and his interaction scale. This begs the question: will this link grow stronger as tasks push the multiscale aspect further and pack even more detail into the information space?

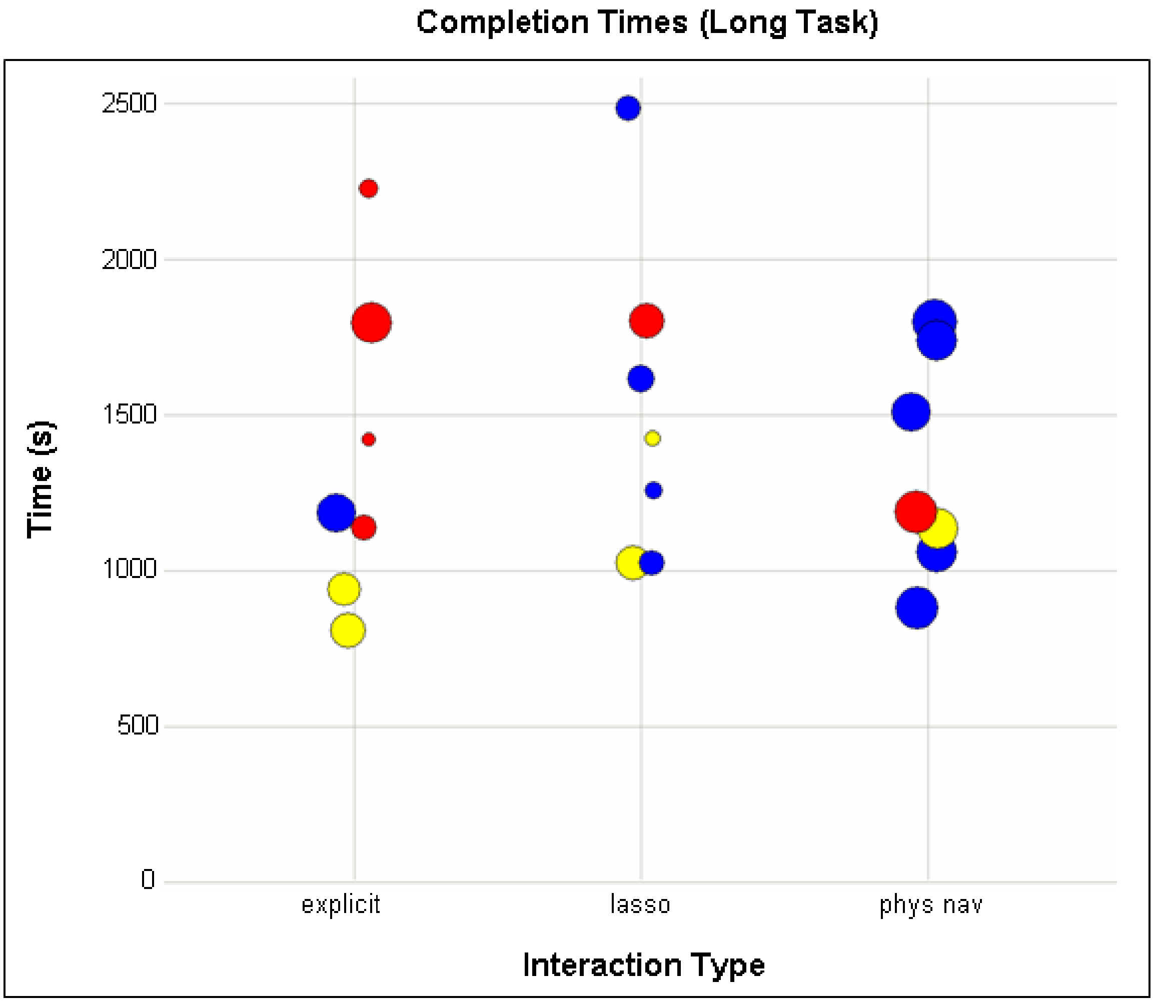

Figure 4.20: Completion Time by Interaction Type (Long Task). Markers are colored by experience and sized according to correlation. Three data points are missing due to errors in movement data.

Returning to long task completion times, if we size the markers by correlation coefficient, we see that the two outliers have smaller correlations between distance and interaction level scale. These correlations are not the lowest overall, so the connection is weakened, but this may indicate that linking distance with interaction scale level improves user completion time performance. Further investigating this using Pearson's product moment correlation, the correlation between the physical navigation / interaction scale correlation and performance time is -0.2228. This is not very strong; however, the two are negatively correlated, indicating additional evidence that linking distance with interaction scale improves user performance time. It is also interesting to note the puzzle strategies used by these participants (discussed in further detail later). The explicit outlier (participant 13) tends to stand further back from the display, lunging in close to peer at details but making all selections standing back. The lasso outlier (participant 28) starts out making an effort to stand further back, but slowly moves in for accuracy while making increasingly frequent movements backward to see the overview.

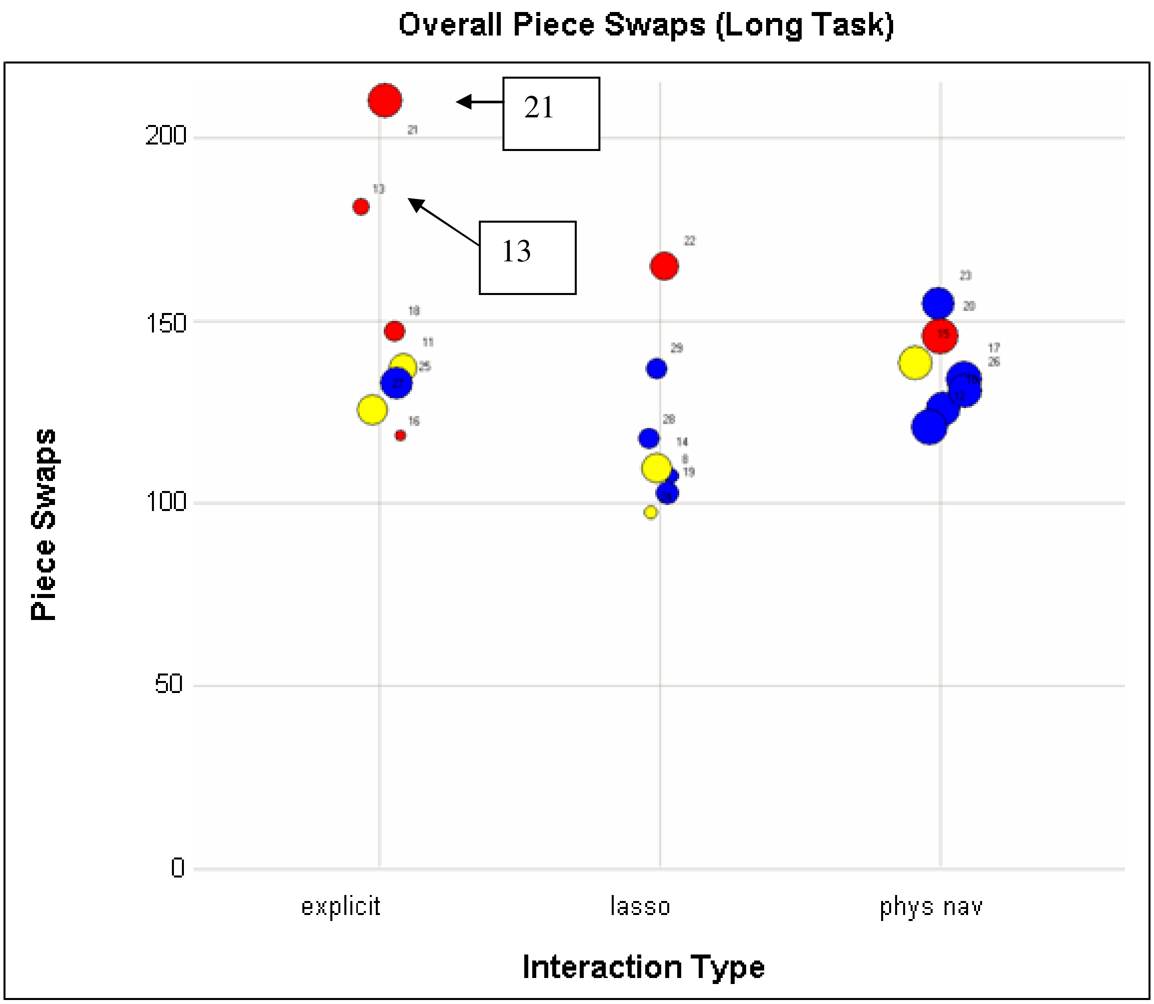

Figure 4.21: Overall Piece Swaps (Long Task) by Interaction Type. Markers colored as in Figure 4.16. Markers now sized by correlation. Three data points left off due to errors in movement data.

Returning to long task piece swaps, if we size the markers by correlation coefficient, this is somewhat inconclusive. One outlier's movement is highly correlated with his interaction scale level changes, while the other is not. Participant 13, in fact, makes a concerted effort to stay further away from the display. It is also interesting to note here that participant 21 is an outlier in puzzle solving strategy – he begins by solving detail pieces first and moving outward as time elapses (Figure 4.28).

Total Distance Moved

We used the position information from each user to calculate his total distance moved throughout the course of the long task. User position was noted every second, removing some jitter from the data, as the motion tracking system samples position information 60 times a second. The distance between each consecutive pair of position samples was calculated, and these distances were summed together to calculate the total distance moved. We performed an analysis of variance (ANOVA) on interaction type that showed that total distance moved was weakly significant ($F(2,18)=2.6739$, $p=0.0962$). Further analysis using the student's T test showed a significant difference between the explicit and lasso types ($p=0.0377$). Explicit participants moved significantly less than lasso participants. Though not significant, the difference in means between explicit and physical navigation is also fairly high ($p=0.1246$).

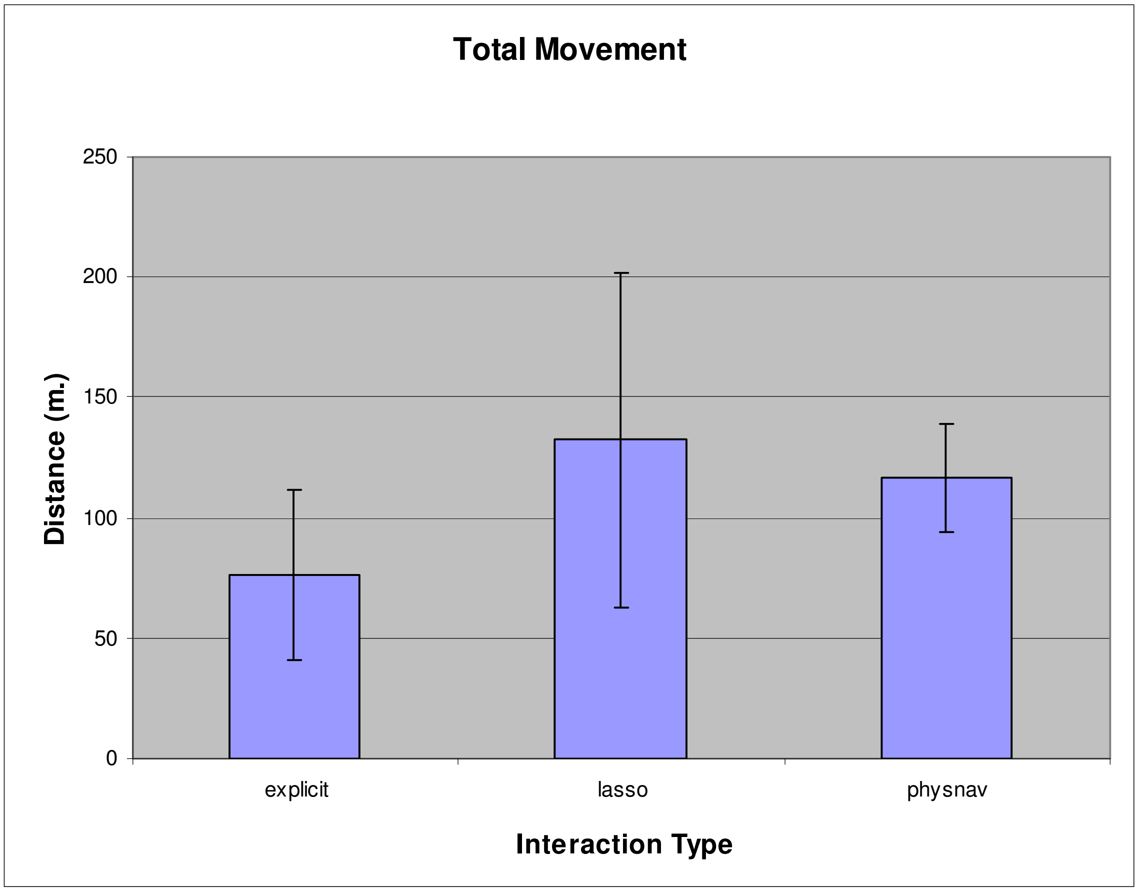

Figure 4.22: Average Distance Moved by Interaction Type

Though there is no difference in distance moved between the lasso and physical navigation types, if we look at the individual participant data, we see that the averages are the same for very different reasons.

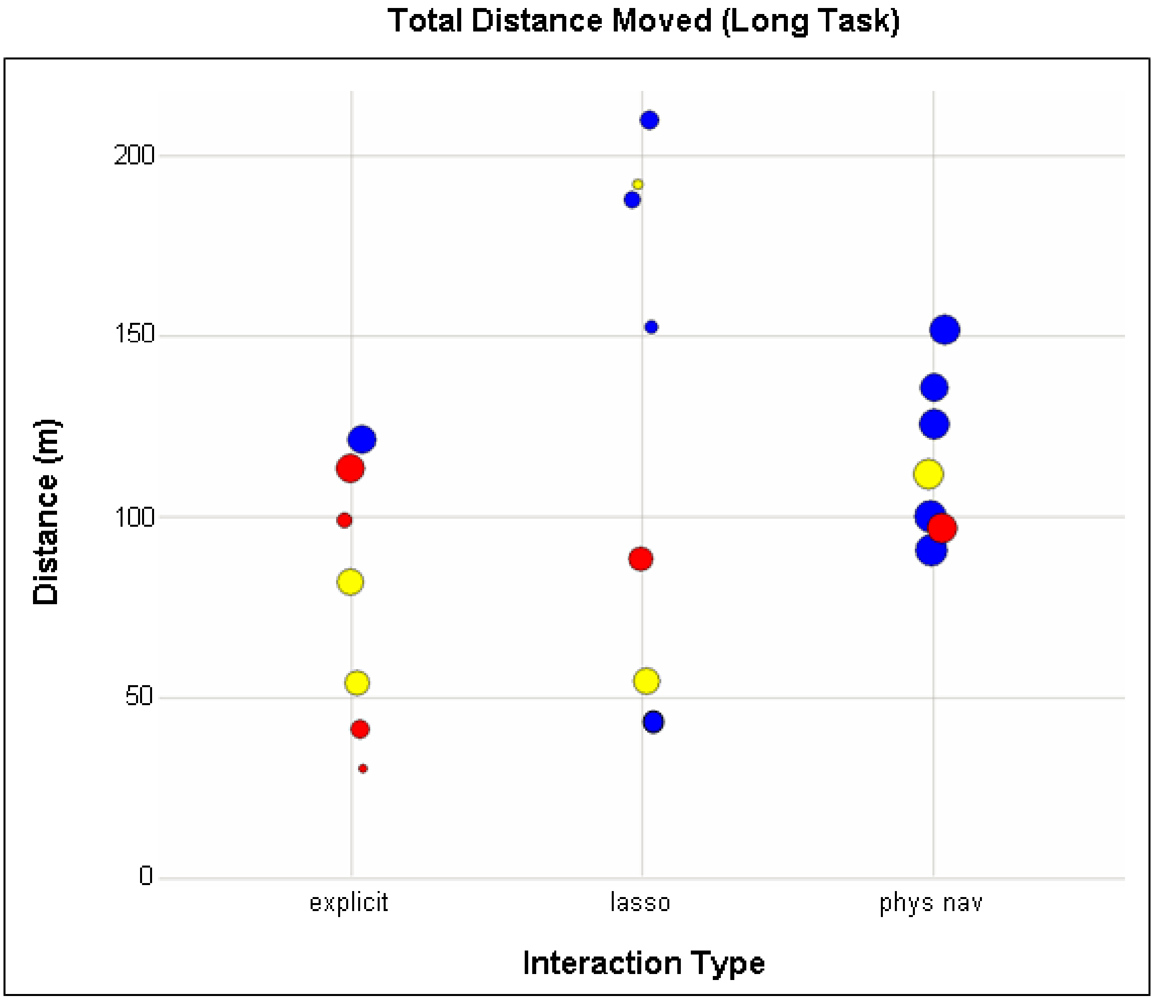

Figure 4.23: Total Distance Moved by Interaction Type. Markers are colored by experience and sized by correlation.

While the total distance moved is fairly consistent across all participants for the physical navigation type, it varies greatly for the participants in the lasso type. There are two distinct sub-groups: those who tended not to move, and those who moved quite a bit. For those lasso participants in the latter group, there was an interesting trend. These participants tended to move up very close to the display to make all selections, regardless of piece size. For example, to select a piece of Level 1 size, these users would bring the wand within an inch of the display and physically drag it diagonally across the entire piece to select it. We think this is because participants in this group had great difficulty pointing with accuracy using this technique, since they had to precisely interact at level 4 scale, even to specify a level 1 scale selection. In order to plan their next swap, these users would move away from the display to gain some amount of overview, and then dive back in close to the display to actually make the swap. This resulted in a large amount of "thrashing". An example of this can be seen in Figure 4.24. This participant is an extreme case of the scenario just described. Known selections are marked with a star, and most of the selections occurred when the user was very close to the display. This pattern of movement is clearly very inefficient.

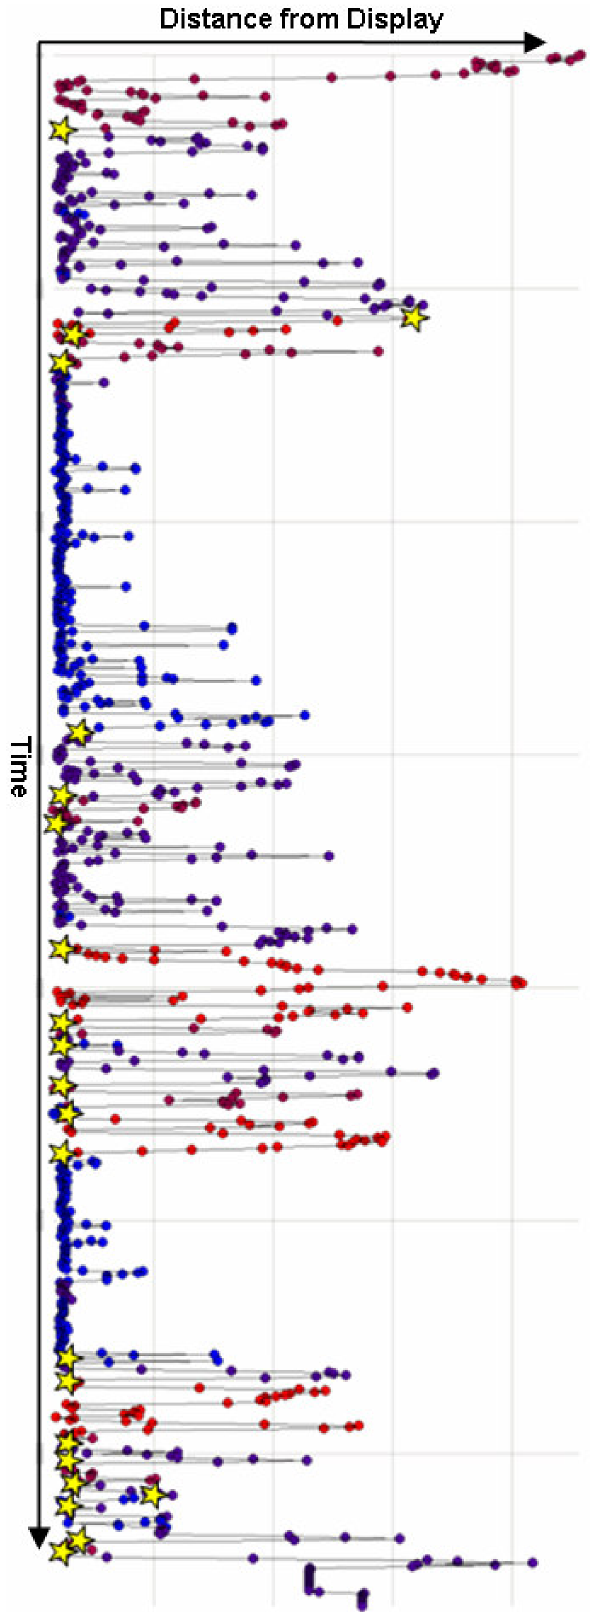

Figure 4.24: A lasso participant consistently moves up close to select. Known selections are marked with stars.

In addition, if we look at how distance moved correlates with the correlation between physical navigation and interaction scale, we see two very interesting trends. For the lasso participants, the correlation coefficient is -0.6971, meaning there is a fairly strong correlation that as the total distance moved increases, the physical navigation / interaction scale correlation decreases. This makes sense after examining how the lasso participants have changed their movement strategy to compensate for accuracy issues. On the other hand, for the explicit participants, the correlation coefficient is 0.6265, meaning there is a fairly strong correlation that as the total distance moved increases, the physical navigation / interaction scale correlation becomes stronger.

Naturalness

For the former group of lasso participants, there is a different story to tell. As in Figure 4.25, there is evidence that the user thought about solving the puzzle at different visual levels of scale, changed his distance from the display accordingly and tended to change his interaction scale as well. The correlation between physical navigation and interaction scale for this participant is -0.7102, the strongest correlation for lasso participants.

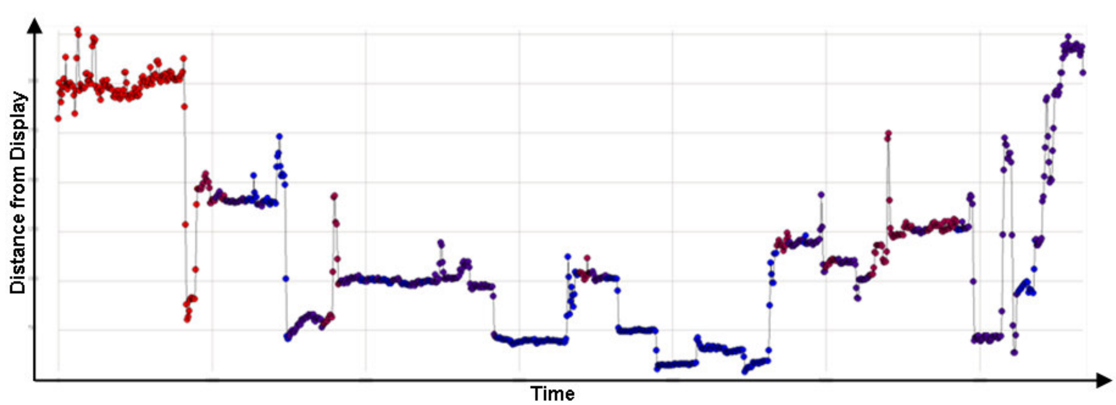

Figure 4.25: Distance and interaction scale for lasso participant 14. Distance is along the x-axis. Markers are colored according to scale level (red is Level 1 and blue is Level 4).

There is evidence that this movement with interaction scale helped performance as this participant performed on par with physical navigation participants for time and swap measures.

Evidence of "naturalness" of the technique is also present in participants using the explicit technique. As seen in Figure 4.26, the user tended to change his interaction scale according to his distance. The correlation between physical navigation and interaction scale for this participant is -0.6718, which is fairly strong. Again, performance measures for time and swaps are on par with other physical navigation participants.

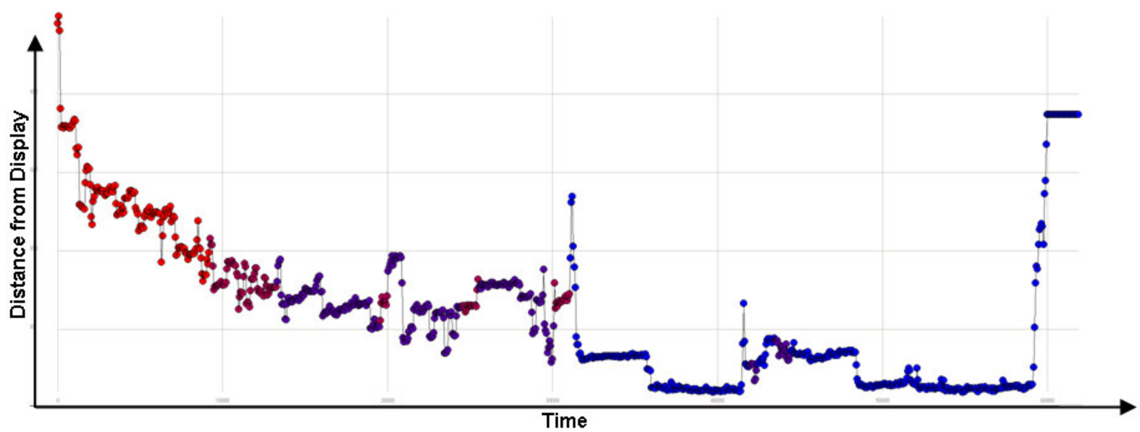

Figure 4.26: Distance and interaction scale for explicit participant 11. Distance from display is along the x-axis. Markers are colored according to scale level (red is Level 1 and blue is Level 4).

Strategies

We found that participants using the physical navigation technique used a fairly consistent strategy to solve the puzzle. The majority of these participants starts further out from the display and arranges Level 1 size pieces, selecting them and swapping them by swiveling their bodies to "pan." Some place all the Level 1 pieces, while others place one, and drill down, solving progressively smaller pieces within that Level 1 area. As they continue on with the task, they progressively move closer to the display and work with smaller scale pieces, while also tending to dynamically change their panning strategy from swiveling to shifting left and right across the length of the display. They also tend to leave the smallest scale level pieces (Level 4) for last. An example of this can be found in Figure 4.27, participant 15.

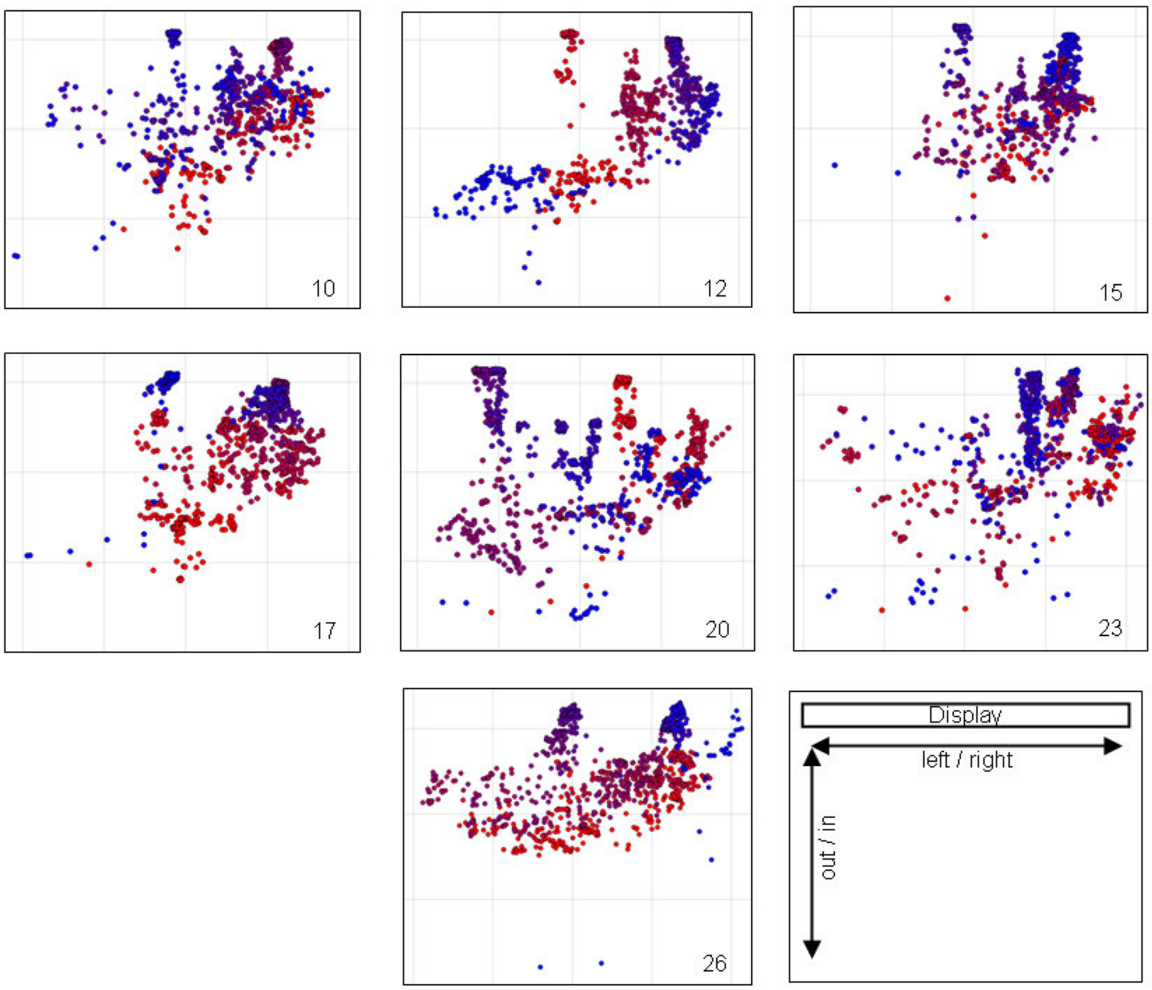

Figure 4.27: Birds' eye view of movement patterns for physical navigation participants. Markers are colored by elapsed time, from red to blue.

There is one exception to this, participant 20. This participant begins by immediately delving down to the smallest scale level and solving the Level 4 pieces first. He then moves out and solves the rest of the puzzle. Both sub-groups of physical navigation participants stand directly in front of the Level 4 pieces they are currently working on; this is not the case for some of the participants from the other interaction types.

Another phenomenon noticed with physical navigation participants is their uncanny ability to move in and out from the display so as to stand directly on the border between scale levels. Once participants were thusly situated, they would often skew their elbows backwards to move up a scale level, or stretch their arm out further to move down a scale level.

There was a similar trend with the explicit participants to start with Level 1 pieces, delve down within that piece and then move on to another focus area, but it was not as widespread. Only 3 of the 7 users seemed to follow the same patterns as physical navigation users, slowly delving down, saving the Level 4 pieces for last, and solving them while standing directly in front of them, close to the display. However, another 3 users seemed stubborn to move, sometimes lunging close to peer at details but then stepping back out to select (see Figure 4.28, participants 13, 16 and 18), and only inching closer to the display when they had problems pointing accurately. Many of these participants had noted problems with pointing accurately, especially with Level 4 size pieces.

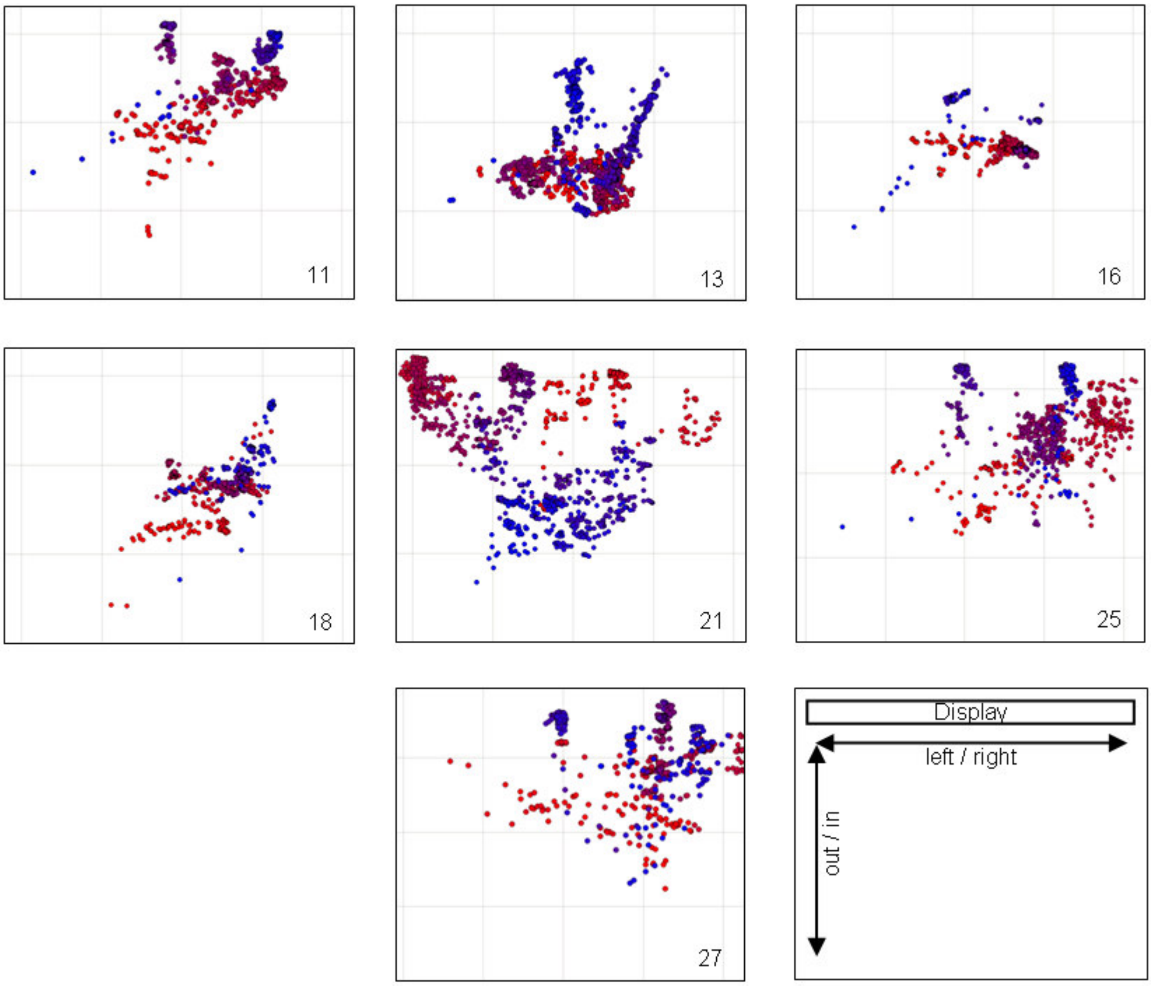

Figure 4.28: Birds' eye view of movement patterns for explicit participants. Markers are colored by elapsed time, from red to blue.

For the lasso participants, the only strategy they all had in common was that they all worked within one focus area (Level 1 sized piece), solving it completely before moving on to another. Four out of the 7 participants had moderate to severe accuracy problems when selecting pieces; many of these participants coped with this by moving close to the display, some especially for detail work and some for all selections (as previously mentioned in Figure 4.24.) There were few commonalities outside these; there were no strong strategies on how participants solved the puzzle within focus areas.

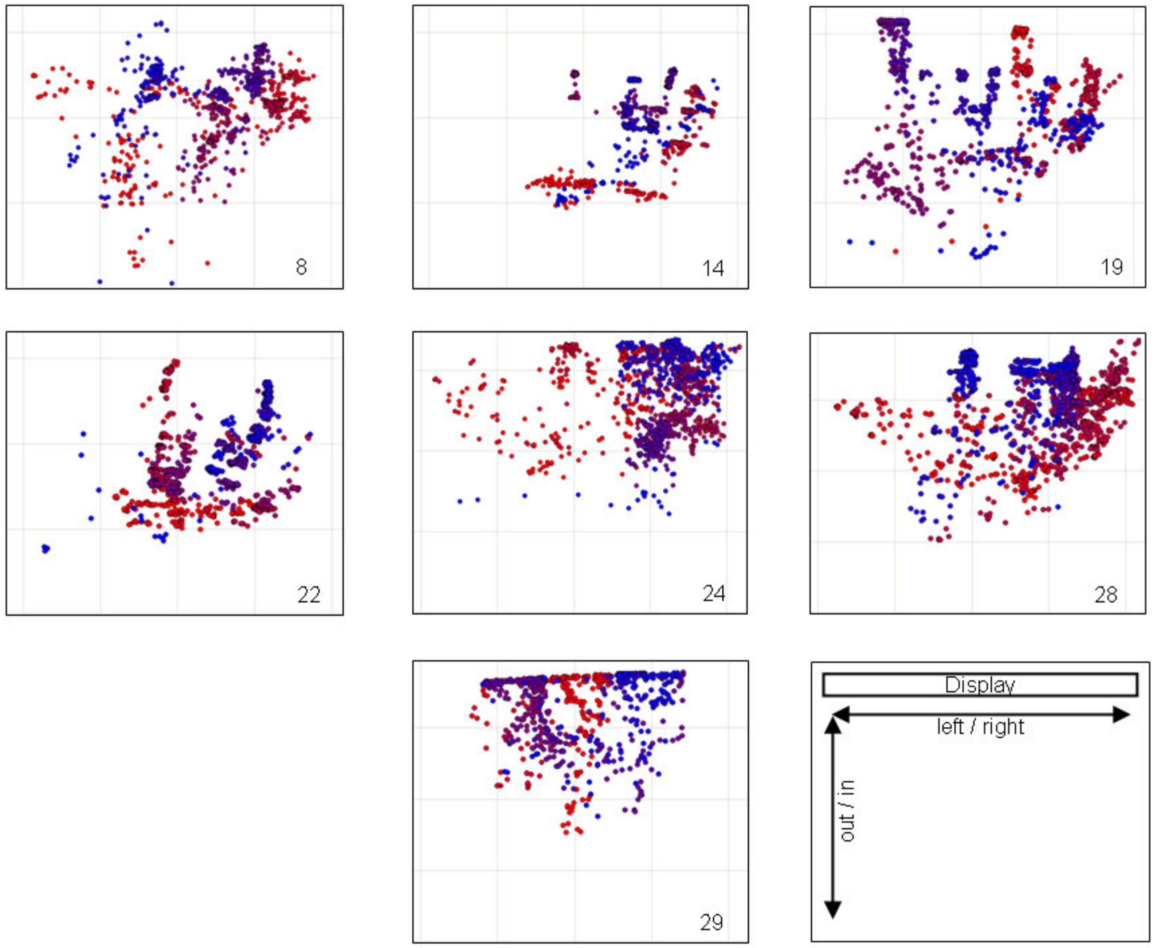

Figure 4.29: Birds' eye view of movement patterns for lasso participants. Markers are colored by elapsed time, from red to blue.

4.3.4 Subjective Measures

After all tasks were completed, users were asked about their overall workload using NASA-TLX. Within overall workload, we looked at mental demand, physical demand, performance, effort and frustration.

We performed an analysis of variance (ANOVA) on interaction type for each measure and found that effort was weakly significant ($F(2,21)=2.9070$, $p=0.0768$).

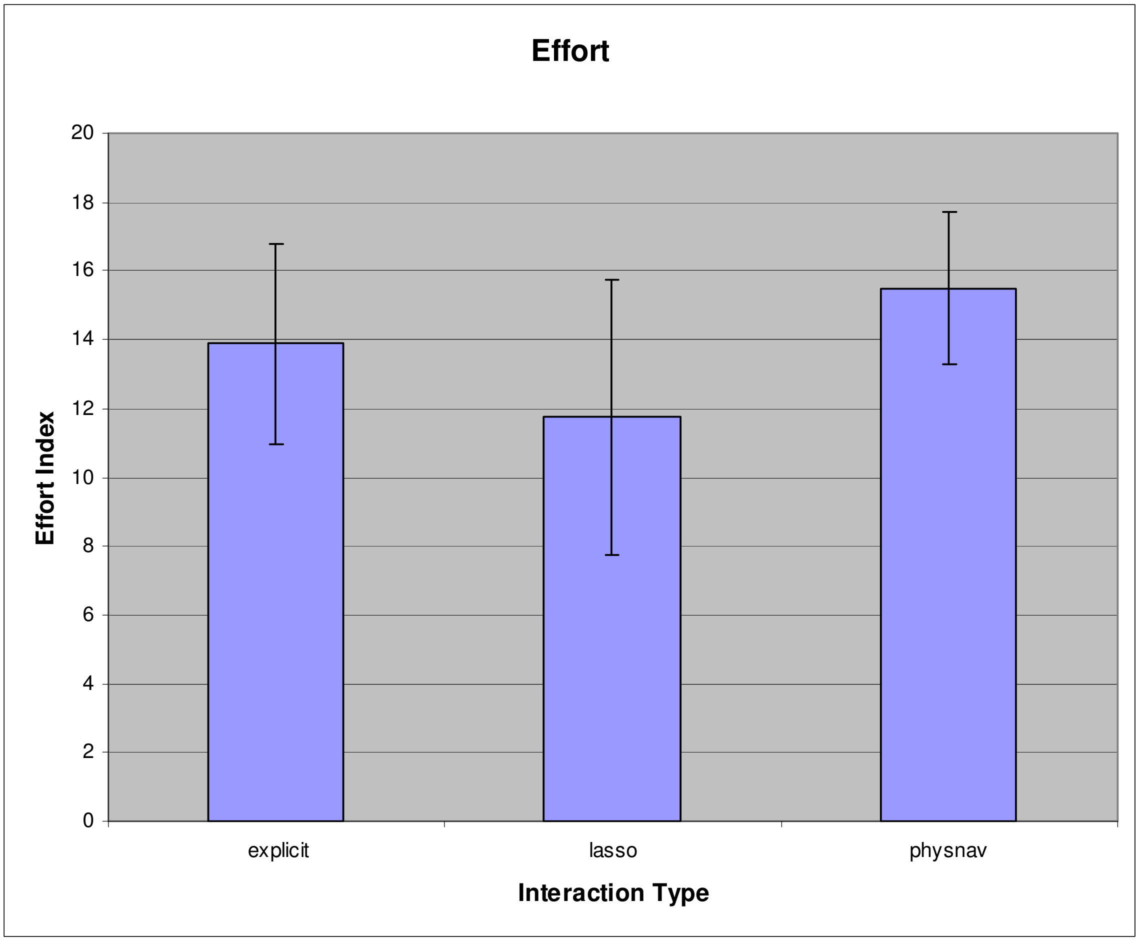

Figure 4.30: Effort by Interaction Type.

The physical navigation type is significantly higher than the lasso type (p=0.0255). This is surprising, as we hypothesized that the physical navigation technique would be more natural for people to use. However, despite some measured performance benefits, users have perceived this technique to be more effortful.

Effort is not correlated to total movement (p=0.2405); however, there are some interesting findings if we look at other subjective measure correlations with movement. Table 4.2 shows the correlations coefficients for mental demand, physical demand, frustration and overall workload versus movement measures. None of the correlations are strong, but it is interesting to note that for frustration and overall workload in particular, all of them are negative. For example, this indicates a trend that the more users move, the *less* frustrated they are.

	Mental demand	Physical demand	Effort	Frustration	Overall workload
In/out mvmt	-0.0176	-0.1983	0.2296	-0.2782	-0.2124
L/R mvmt	0.1090	0.0171	0.2583	-0.2543	-0.0601
Total mvmt	0.0488	-0.0920	0.2565	-0.2733	-0.1362

Table 4.2: Correlation coefficients between various subjective and movement measures.

4.4 Summary

Again, this study was exploratory in nature. Future studies might more directly measure the benefits of multiscale interaction by changing several things. First, the tasks could be made shorter, enabling the interaction type variable to be applied across subjects. This way, user preference data could be collected. Secondly, the tasks could be made more difficult by pushing the multiscale aspect even further, packing in even more details. Lastly, the task domain could be changed altogether. The puzzle solving task was meant to be typical, but the benefits of multiscale interaction for this task could be confounded by things like varying user strategies.

While not as strong as originally hypothesized, we did find evidence in favor of multiscale interaction.

The correlation between interaction scale and distance from the display for physical navigation participants was understandably very strongly correlated, since scale was automatically changed according to a user's distance. For both the explicit and lasso types, there were participants whose movement in and out from the display also strongly correlated with their interaction scale, but there were also participants whose movements did not correlate. However, correlation coefficients for participants were consistently negative, meaning as interaction scale number got larger (from Level 1, the largest piece size to Level 4, the smallest piece size), user distance from the display got smaller. This indicates that there is indeed a natural link between a user's distance from the display and his scale of interaction. Even with the lasso technique, which imposed no interaction scale on the user whatsoever, users still to some degree moved forward to interact on a smaller scale and moved further away to interact in a larger scale.

In addition, across all measures on the long task, the physical navigation type has remained very consistent. For completion time, it has appeared to avoid outlier participants with unusually long times. While lasso and explicit types had high variance in piece swap counts and total distance moved, physical navigation participants were again clustered very consistently. Lastly, 6 of the 7 physical navigation participants used a very similar puzzle solving strategy. Often, the outlier participants, such as those for completion times and puzzle swaps, were those with less experience with large display environments. The physical navigation type appeared to act as an equalizer, allowing less experienced users to perform more consistently as well as those with more experience.

Explicit participants moved significantly less than those using the lasso technique, and while not significant, also appeared to move less than physical navigation participants. While there was no difference between total movement in the lasso and physical navigation participants, this additional movement seemed to be beneficial for the latter. However some lasso participants had distinct trouble selecting pieces accurately and compensated by moving up close to the display to select any piece, regardless of level, and then immediately moving back out to survey their work. This "thrashing" movement added up but did not appear to be beneficial to the participants in any other way.

The explicit participants' lack of movement could be detrimental to their performance. While some users were not affected, others suffered in performance, notably with piece swaps, making more "mistake" swaps than either of the other two types. So in a sense, the physical navigation technique helped users avoid these "mistake" swaps by enforcing a strategy to overcome the accuracy problems inherent in the system.

Chapter 5: Conclusions

5.1 Summary

Current interaction with large displays is missing something. Users are physically navigating in the space in front of the display in order to see all of the displayed information, essentially changing the scale at which they are viewing the information. However, the users' scale of interaction never changes, which is a problem because users are doing different types of tasks at different levels of visual scale. This prompted our initial research question: how can we exploit users' physical navigation to enable simpler interaction at multiple levels of scale?

Our solution to this is multiscale interaction, a technique that exploits users' physical navigation by linking it to interaction scale, changing the users' scale of interaction depending on their distance from the current object(s) of interest. To begin our exploration of multiscale interaction, we defined the design space. Here, we asked, to what kinds of situations can multiscale interaction be applied? The two largest aspects of interaction are selection and navigation, and multiscale interaction can be applied to both, for multiple types of datasets. In addition, multiscale interaction has applications in multi-user and multi-step interaction.

We further explored several areas of the design space through a case study and a user study. Through the case study, we identified several design issues and tradeoffs. Designers should keep the dataset in mind when designing the various aspects of the multiscale interaction techniques, such as pointing technique, cursor size and shape, and selection area feedback.

Through the study, we evaluated an implementation of multiscale interaction, to answer another of our research questions, is multiscale interaction an effective technique? Do users benefit from the linkage of physical navigation with interaction? Do they tend toward such interaction anyway (is it natural)?

While we did not find the strong performance benefits that we hypothesized, we still found evidence supporting multiscale interaction.

- Across all performance measures, the physical navigation type remained very consistent in its results. While the explicit and lasso types reported outliers with higher performances in completion time, piece swap counts and total distance moved, physical navigation results were always clustered very consistently.
- Often, the outliers in the other techniques were participants with little or no experience with large displays. The physical navigation technique appeared to act as an equalizer, allowing less experienced user to perform more consistently as well as those with more experience.
- Physical navigation participants also used an overwhelmingly consistent puzzle solving strategy, where they progressively moved in towards the display while also solving pieces on progressively smaller scale levels. Few participants in the other interaction types used a similar strategy, but most strategies varied widely. This variation in puzzle solving strategy did not appear to affect other performance measures, such as completion time or piece swaps.

There is also evidence that there is a link between a user’s distance from the display and his scale of interaction, and that this movement is to some degree natural to users.

- The link between interaction scale and distance enforced in the physical navigation technique appears to hold across the other interaction types. Despite the fact that there not all participants' movement was highly correlated with their interaction scale, there was a strong trend of the two being negatively correlated. This means, to some degree across all participants, as they changed interaction scale to work with smaller scale pieces (in this study, interaction scale level number increased from Level 1 to Level 4), they also moved forward to be closer to the display.
- Explicit participants moved significantly less than lasso participants and though not significant, also appeared to move less than physical navigation participants. There is a possibility that this lack of movement on the part of the explicit participants could be detrimental to their performance. Some users were not affected; however, others suffered notably in their performance. The key example of this is found in explicit users' piece swaps, where there are several participants who swapped pieces much more frequently than either of the other types. These additional piece swaps appear to be "mistake" swaps.
- While the mean distances moved for the lasso and physical navigation types were similar, this extra movement as compared to the explicit type seemed to only be beneficial to the physical navigation participants. As mentioned previously, the performance for the physical navigation type was very consistently clumped. However, the additional movement for the lasso participants did not appear to be beneficial. The participants had distinct trouble selecting pieces accurately and compensated for this by moving up close to the display to select any piece,

regardless of scale level. They would then immediately move back out to survey their work. This resulted in a large amount of "thrashing" movement in and out from the display that did not benefit the participants' performance in any other way.

5.2 Implications for Future Design and Theory

The results of this study can inform future work on developing interaction techniques for large, high-resolution displays. Understanding the link between user movement, specifically the movement in and out from the display, and interaction scale can help others create more natural techniques for interaction with such displays.

So what does this say about embodiment theory? The multiscale interaction implementation produced consistent performance with its participants, across all long task measures. Is this consistency of performance a good thing – does it mean that we have created a technique which successfully taps into the users' sense of embodiment – or is it merely the result of forcing the users into over conforming? The other results do show that changing interaction level according to movement is natural to people and that, while forcing users to move more, this extra movement with the physical navigation technique was beneficial. Combined with these results, the consistent performance indicates that we have connected with the users' sense of embodiment.

What about the other techniques? What happens when techniques are designed against embodiment? The lasso technique forces users to interact on the detail level, and they must compensate for this by selecting up close while still viewing at a distance. Users' physical navigation is used against them; they must move to perform interactions, but this movement has no benefit for them. The explicit technique, while still allowing for interaction on different levels of scale, is an interesting case. Some, but not all users were affected by not having the interaction scale linked to their movement, tending to cause them, at the extreme, to stand still. Why? These users might have been overloaded, in an interaction sense, by trying to physically navigate and simultaneously manage their interaction scale via menu; however, further investigation is needed to better understand this relationship.

Users might even benefit from a combination of the physical navigation and explicit techniques. On several occasions, the physical navigation users would stretch out their arm or skew it backwards to move up or down a scale level. This indicates that there is more to changing scale than just visual distance – there is also a measure of the hand's distance from both the eye and the display. Users unconsciously combine hand-eye coordination into their interaction with the display. Would users benefit from a physical navigation technique that controlled the coarse scale level changes which then used hand distance as an explicit fine tuning of scale?

How else did people benefit from multiscale interaction, other than quantitative performance measures? We looked at puzzle solving strategy as a way to gauge how it helped people think, reason about the data and solve problems. It was interesting to note that the physical navigation participants were consistent even in their puzzle solving strategy. This might indicate that the technique helped the users better understand the structure of the data and have a more complete mental model of the puzzle, which in turn allowed them to plan a more organized solution to the problem. It would be interesting to test this on a more complex dataset.

All of this strengthens the argument for embodiment theory. In his 2007 paper, Robert Ball showed that physical navigation and embodiment improved user performance in the perception of data from visualizations [5]. Visualization has two parts: the users' perceptions of the information and their interactions with the visualization. This work has shown that physical navigation and embodiment benefits visualization interaction, in addition to perception.

5.3 Future Work

This work has discovered some benefits of multiscale interaction, but how will these benefits hold as the dataset is changed? The task in this study incorporated 4 levels of hierarchy. However, the multiscale aspect could be pushed to the extreme, and many more levels of hierarchy could be packed into the information space (imagine a detail level of 8 point font text that users could manipulate). It would be interesting to evaluate this technique as the hierarchy of the information space became more and more dense. Would the benefits of multiscale interaction hold or become even more apparent?

Much work can still be done exploring the remaining areas of the design space of mutliscale interaction. As enumerated in Chapter 3:, this study only implements one of the many interpretations of multiscale interaction, namely an automatic resizing of the cursor on discrete levels that is useful for datasets in which there is already a predefined hierarchy. Are the same benefits found in other selection interpretations, such as those for more unorganized datasets?

Other interpretations affect other aspects of interaction apart from selection. Does the theory of embodiment still hold for these other interpretations? Are these other relations, such as linking distance with rate control or lens size, still natural interactions for users? Do the same benefits discovered for hierarchical selection hold for these other aspects of interaction, or is there a different story going on there?

The relationship between physical navigation, embodiment and interaction is not yet completely understood. Further investigation is needed into the other areas of multiscale interaction to design more natural interaction for the future.

References

1. *Motion Capture Systems from Vicon*. [cited 2008; Available from: http://www.vicon.com.

2. *SequoiaView*. [cited 2008; Available from: http://w3.win.tue.nl/nl/onderzoek/onderzoek_informatica/visualization/sequoiaview//.

3. *Treemap*. 5 August 2003 [cited 2008; Available from: http://www.cs.umd.edu/hcil/treemap.

4. Ball, R. and C. North, *Effects of tiled high-resolution display on basic visualization and navigation tasks*, in *CHI '05 extended abstracts on Human factors in computing systems*. 2005, ACM Press: Portland, Oregon, USA.

5. Ball, R., C. North, and D.A. Bowman, *Move to improve: promoting physical navigation to increase user performance with large displays*, in *SIGCHI conference on Human factors in computing systems*. 2007, ACM Press: San Jose, California, USA. p. 191-200.

6. Bangerter, A., *Using Pointing and Describing to Achieve Joint Focus of Attention in Dialogue*. Psychological Science, 2004. **15**(6): p. 415-419.

7. Becker, R.A. and W.S. Cleveland, *Brushing scatterplots*. 1987, American Society for Quality Control and American Statistical Association. p. 127-142.

8. Bederson, B.B. and J.D. Hollan, *Pad++: a zoomable graphical interface system*, in *Conference companion on Human factors in computing systems*. 1995, ACM: Denver, CO, United States.

9. Bowman, D.A., E. Kruijff, J. Joseph J. LaViola, and I. Poupyrev, *3D User Interfaces: Theory and Practice*. 2005, Boston: Addison-Wesley.

10. Cao, X. and R. Balakrishnan, *VisionWand: interaction techniques for large displays using a passive wand tracked in 3D*, in *Proceedings of the 16th annual ACM symposium on User interface software and technology*. 2003, ACM: Vancouver, BC, Canada.

11. Dasiyici, M., *Multi-Scale Mouse: Optimizing Mouse Interaction for Large Personal Workspaces*, in *Computer Science*. 2008, Virginia Polytechnic Institute and State University: Blacksburg, VA.

12. Dourish, P., *Seeking a Foundation for Context-Aware Computing.* Human-Computer Interaction, 2001. **16**(2): p. 229 - 241.

13. Dourish, P., *Where the Action Is: The Foundations of Embodied Interaction*. 2004: MIT Press. 233.

14. Forsberg, A., K. Herndon, and R. Zeleznik, *Aperture based selection for immersive virtual environments*, in *9th annual ACM symposium on User interface software and technology (UIST)*. 1996, ACM Press: Seattle, Washington, USA. p. 95-96.

15. Furnas, G.W., *Generalized fisheye views*, in *Proceedings of the SIGCHI conference on Human factors in computing systems*. 1986, ACM: Boston, Massachusetts, United States.

16. Grossman, T. and R. Balakrishnan, *The bubble cursor: enhancing target acquisition by dynamic resizing of the cursor's activation area*, in *SIGCHI conference on Human factors in computing systems (CHI 2005)*. 2005, ACM Press: Portland, Oregon, USA. p. 281-290.

17. Guiard, Y. and M. Beaudouin-Lafon, *Target acquisition in multiscale electronic worlds.* International Journal of Human-Computer Studies, 2004. **61**(6): p. 875-905.

18. Guiard, Y., M. Beaudouin-Lafon, J. Bastin, D. Pasveer, and S. Zhai, *View size and pointing difficulty in multi-scale navigation*, in *Proceedings of the working conference on Advanced visual interfaces*. 2004, ACM: Gallipoli, Italy.

19. Guiard, Y., R. Blanch, and M. Beaudouin-Lafon, *Object pointing: a complement to bitmap pointing in GUIs*, in *Proceedings of Graphics Interface*. 2004, Canadian Human-Computer Communications Society: London, Ontario, Canada.

20. Jiang, H., E. Ofek, N. Moraveji, and Y. Shi, *Direct pointer: direct manipulation for large-display interaction using handheld cameras*, in *SIGCHI conference on Human Factors in computing systems*. 2006, ACM Press: Montreal, Quebec, Canada. p. 1107-1110.

21. Kabbash, P. and W. Buxton, *The "Prince" Technique: Fitts' Law and Selection Using Area Cursors*, in *SIGCHI conference on Human Factors in Computing Systems (CHI'95)*. 1995, ACM Press/Addison-Wesley Publishing Co.: Denver, Colorado, USA. p. 273-279.

22. Klemmer, S.R., B. Hartmann, and L. Takayama, *How bodies matter: five themes for interaction design*, in *ACM conference on Designing Interactive systems*. 2006, ACM Press: University Park, Pennsylvania, USA. p. 140-149.

23. Kranstedt, A. and I. Wachsmuth. *Incremental Generation of Multimodal Deixis Referring to Objects*. in *Proceedings of European Workshop on Natural Language Generation (ENLG2005)*. 2005. Aberdeen, UK.

24. Liang, J. and M. Green, *Geometric modeling using six degrees of freedom input devices*, in *3rd International Conference on CAD & Computer Graphics (CAD/Graphics '93)*. 1993: Beijing, China.

25. Liang, J. and M. Green, *JDCAD: A highly interactive 3D modeling system.* Computers & Graphics, 1994. **18**(4): p. 499-506.

26. Lucking, A., H. Rieser, and J. Stegmann. *Statistical Support for the Study of Structures in Multimodal Dialogue: Inter-rater Agreement and Synchronisation*. in *Proceedings of the 8th Workshop on the Semantics and Pragmatics of Dialogue*. 2004.

27. Lyons, J., *Semantics*. 1977, Cambridge, UK: Cambridge University Press.

28. Ni, T., D.A. Bowman, and J. Chen, *Increased display size and resolution improve task performance in Information-Rich Virtual Environments*, in *Proceedings of Graphics Interface 2006*. 2006, Canadian Information Processing Society: Quebec, Canada.

29. Norman, D.A., *The Design of Everyday Things*. 1988, New York: Basic Books.

30. Rajlich, P. *NCSA TerraServer Blaster*. 2003 [cited 30 March 2007]; Available from: http://brighton.ncsa.uiuc.edu/prajlich/wall/tsb.html.

31. Schum, D., *A Wigmorean Interpretation of the Evaluation of a Complicated Pattern of Evidence*. 2005.

32. Shupp, L., *The Effects of Curving Large, High-Resolution Displays on User Performance*, in *Computer Science*. 2006, Virginia Polytechnic Institute and State University: Blacksburg, VA. p. 95.

33. Shupp, L., R. Ball, B. Yost, J. Booker, and C. North, *Evaluation of viewport size and curvature of large, high-resolution displays*, in *2006 conference on Graphics interface (GI'06)*. 2006, Canadian Information Processing Society: Quebec, Canada. p. 123-130.

34. Stone, M.C., K. Fishkin, and E.A. Bier, *The Movable Filter as a User Interface Tool*, in *Proceedings of CHI 94*. 1994, ACM Press: Boston, MA, United States.

35. van der Sluis, I. and E. Krahmer. *The Influence of Target Size and Distance on the Production of Speech and Gesture in Multimodal Referring Expressions*. in *Proceedings of the 8th International Conference on Spoken Lanugage Processing*. 2004.

36. Vogel, D. and R. Balakrishnan, *Distant freehand pointing and clicking on very large, high resolution displays*, in *Proceedings of the 18th annual ACM symposium on User interface software and technology*. 2005, ACM: Seattle, WA, USA.

37. Wilson, M., *Six views of embodied cognition.* Psychonomic Bulletin & Review, 2002. **9**: p. 625-636.

38. Witkin, H.A., P.K. Oltman, E. Raskin, and S.A. Karp, *Group embedded figures test*. 1971, Palo Alto, CA: Consulting Psychologists Press.

39. Worden, A., N. Walker, K. Bharat, and S. Hudson, *Making computers easier for older adults to use: area cursors and sticky icons*, in *SIGCHI conference on Human factors in computing systems*. 1997, ACM Press: Atlanta, GA, USA. p. 266-271.

40. Yost, B., Y. Haciahmetoglu, and C. North, *Beyond visual acuity: the perceptual scalability of information visualizations for large displays*, in *SIGCHI conference on Human factors in computing systems*. 2007, ACM Press: San Jose, California, USA. p. 101-110.

41. Zhai, S., W. Buxton, and P. Milgram, *The "Silk Cursor": investigating transparency for 3D target acquisition*, in *SIGCHI conference on Human factors in computing systems*. 1994, ACM Press: Boston, Massachusetts, USA.

Appendix A: Participant Instructions

The task I would like to have you do today is to solve a multiscale puzzle. You will be shown an image that has been sliced up into pieces, and the pieces have been shuffled. The pieces of the puzzle are arranged in 4 hierarchical levels. (Show each level & name.) So within level 1, the largest level, for instance, there are 4 level 2 size pieces which are shuffled – these 4 pieces only need to be solved within that larger level 1 piece. The same goes for the other levels.

You should solve the puzzle as quickly and as accurately as possible.

You will have a printout of the final image (that you are trying to solve) to refer to throughout the task.

To point to the display, you will use this wand. Your cursor appears on the display at the place where the wand is pointing. To click on a piece, use the side button on the PDA. To swap two puzzle pieces, first click on one, and then on another. Those two pieces will be swapped. You are not allowed to swap pieces on different levels.

FOR PHYSICAL NAVIGATION:

Your cursor will change sizes according to the different hierarchical levels of pieces. The cursor changes sizes automatically, depending upon your distance from the display. The cursor is indicated by the thick red line. The next lowest level is also indicated by a thin white line. A selected piece is indicated by a thick yellow line.

FOR EXPLICIT:

Your cursor can change sizes according to the different hierarchical levels of pieces. You should use the menu on the PDA to switch from level to level. The cursor is indicated by a thick red line. The next lowest level is also indicated by a thin white line. A selected piece is indicated by a thick yellow line.

FOR RUBBER BAND:

Your cursor appears as a small red square. You must draw a box around the hierarchical level of pieces that you wish to select. To select, click the button, hold it down and drag. When the box, indicated by a thick red line, encompasses the hierarchical level of pieces you wish to select, release the button. As you are pointing at the display, in addition to your small red square cursor, each of the hierarchical levels being pointed at are indicated by thin white boxes. Once a hierarchical level of pieces is selected, it is indicated by a thick, yellow line.

First, you I will give you time to practice on a sample puzzle image. Please solve this puzzle to become accustomed to the interaction technique and the task. If you have any questions about either, please ask.

Appendix B: Informed Consent

I. Purpose of this Research/Project

The purpose of our study is to evaluate a novel interaction technique to be used with large, high-resolution displays. We are looking at several aspects of the new technique to see how they affect user performance. We believe that use of the technique will result in improved performance. The findings from this experiment will inform research on ideal interaction techniques for large displays.

II. Procedures

This experiment will take place in the Blacklab area of Knowledgeworks II.

You will first be asked to complete a short Embedded Figures Test, which will assess your ability to solve puzzles. Then you will be asked to complete a puzzle solving task on the Gigapixel display, using one version of the interaction technique. The experimenter will describe in more detail the task and how to use the interaction device/technique. Before the task begins, you will be given time to practice with the interaction device/technique on an alternate puzzle task. You will also be shown the final image of the solved puzzle, and you will be able to refer to this image throughout the task. You should work on the puzzle until it has been solved.

After the puzzle task has been completed, you will be asked to fill out a short questionnaire concerning your background and workload during the task.

The total time commitment is expected to be 1.5 hours.

III. Risks

This experiment is an evaluation of an interaction technique for large displays using a mobile interaction device. Participation involves standing, moving around the area in front of the large display, and manipulating a handheld interaction device. The physical components of the task previously described may result in leg or arm fatigue.

If you experience any leg or arm fatigue during the task, please inform the experimenter. You are free to rest, during which the stopwatch measuring your time on the task will be paused, and may continue when you are ready. If the fatigue becomes uncomfortable, you will be allowed to leave with no penalty.

IV. Benefits

Through your participation in this project, your are assisting in research that may help improve interaction techniques for large displays. No promise or guarantee of benefits has been made to encourage you to participate. If you wish, you may receive a synopsis summarizing this research when completed. Please contact the experimenter at speck@vt.edu.

V. Extent of Anonymity and Confidentiality

The results of this study will be kept strictly 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 an only a subject number will identify you during analyses and any written reports of the research. The data will be stored in a locked cabinet, and will be destroyed in three years. It is possible that the Institutional Review Board (IRB) may view this study's collected data for auditing purposes. The IRB is responsible for the oversight of the protection of human subjects involved in 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. Subject's Responsibilities

I voluntarily agree to participate in this study, and I know of no reason I cannot participate. As a participant, I may withdraw from this experiment at any time without penalty. I agree to abide by the rules of this project.

IX. Subject's Permission

I have read the Consent Form and conditions of this project. I have had all my questions answered. I hereby acknowledge the above and give my voluntary consent:

____ Date____________________
Subject signature

Should I have any pertinent questions about this research or its conduct, and research subjects' rights, and whom to contact in the event of a research-related injury to the subject, I may contact:

Sarah Peck	864-982-6276 / speck@vt.edu
Investigator	Telephone/e-mail
Chris North	540-231-2458 / north@vt.edu
Faculty Advisor	Telephone/e-mail
David M. Moore	540-231-4991/moored@vt.edu
Chair, Virginia Tech Institutional Review Board for the Protection of Human Subjects Office of Research Compliance 2000 Kraft Drive, Suite 2000 (0497) Blacksburg, VA 24060	Telephone/e-mail

Appendix C: User Questionnaire

A Multiscale Selection Technique for Interaction with Large, High-Resolution Displays

Gender (circle one):   Male Female
Age:_______________
Occupation (if student, indicate graduate or undergraduate):_______________________
Major / Area of specialization (if student): ______________________________________

Please describe your previous experience with large display environments:

NASA Task Load Index

For each category, mark on the scale at the point that best indicates your experience of the task.

Mental Demand   *How mentally demanding was the task?*

|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|

Very Low   Very High

Physical Demand   *How physically demanding was the task?*

|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|

Very Low   Very High

Performance   *How successful were you in accomplishing what you were asked to do?*

|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|

Failure   Perfect

Effort   *How hard did you have to work to accomplish your level of performance?*

|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|

Very Low   Very High

Frustration   *How insecure, discouraged, irritated, stressed and annoyed were you?*

|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|___|

Very Low   Very High

For each pair, mark next to the factor that represents the more important contributor to workload for the task.

☐ Physical demand	☐ Frustration
☐ Mental demand	☐ Physical demand
☐ Performance	☐ Mental demand
☐ Frustration	☐ Mental demand
☐ Effort	☐ Physical demand
☐ Frustration	☐ Effort
☐ Performance	☐ Frustration
☐ Physical demand	☐ Performance
☐ Mental demand	☐ Effort
☐ Effort	☐ Performance