Author’s Accepted Manuscript

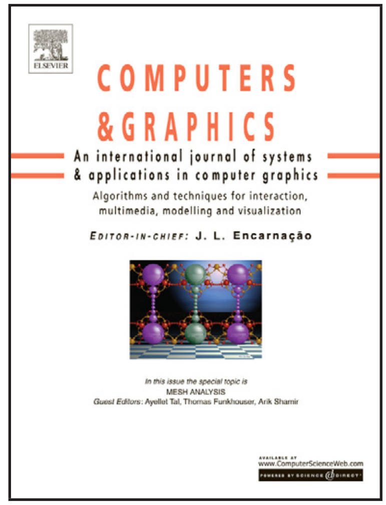

www.elsevier.com/locate/cag

Realizing embodied interaction for visual analytics through large displays

Robert Ball, Chris North

PII: S0097-8493(07)00060-X
DOI: doi:10.1016/j.cag.2007.01.029
Reference: CAG 1763

To appear in: *Computers & Graphics*

Cite this article as: Robert Ball and Chris North, Realizing embodied interaction for visual analytics through large displays, *Computers & Graphics* (2007), doi:10.1016/j.cag.2007.01.029

Realizing Embodied Interaction for Visual Analytics through Large Displays

Robert Ball and Chris North
Center for Human-Computer Interaction
Department of Computer Science
Virginia Polytechnic Institute and State University
Blacksburg, VA 24060

Abstract

Visual analysts are engaged with the arduous task of scrutinizing increasingly larger datasets. Where conventional desktop displays are reaching their limits in terms of performance efficiency with large datasets, analysts can turn to larger displays. In a world of extensive multi-scale datasets, *large high-resolution displays* have the potential to show both more overview and detail for a given dataset than their smaller counterparts. In addition, people are able to use their *embodied resources*, such as spatial memory, proprioception, and optical flow to help them maintain orientation and improve performance times on analytic tasks when using larger displays.

This paper looks at how *physical navigation*, physically interacting with large scale visualizations (e.g. walking, crouching, moving the head), affects user performance times on analytic tasks, such as finding patterns in geospatial data. The paper extends the space-scale diagram to take into account physical navigation and explains the theoretical repercussions. The paper then explains an empirical study performed for the purpose of further understanding how physical and virtual navigation affect performance times of tasks on varying size displays.

In general, we found that large displays can decrease performance time of basic visualization tasks by more than ten times. In addition, we found overwhelming evidence from the empirical study that participants preferred physical navigation over virtual navigation (e.g. mouse interaction). Specifically, we found that for a number of tasks 100% of the participants chose to physically navigate - physically moving to different areas on the display instead of using virtual navigation to manipulate the view on the display.

Keywords: embodied interaction, large displays, space scale, visual analytics, space scale

1. Introduction

Visual Analytics is concerned with collecting data and finding patterns in large scale data through a visual medium. Our approach to visual analytics includes viewing data through a large visual medium and doing so in a more efficient, embodied approach. By expanding visualizations to human scale, we can potentially increase the scalability of visual analytics. This approach takes advantage of not only the eyes and cognition but the entire body.

Embodied interaction is the theory that the cognitive mind is not separated from the physical body. Exploiting this idea involves using the body's already usable functionality to enhance performance and insight in analytic tasks. Embodied interaction makes better use of physical embodied resources such as motor memory, peripheral vision, optical flow, focal attention, and spatial memory [10].

This approach to visual analytics using embodied interaction is to use large, high resolution[1] displays such as in Figure 1. Such displays allow people to physically interact with their data in a way that is not possible with smaller displays. In addition, by having a higher resolution (more pixels) than projector-based displays more data is shown and, therefore, more insight and knowledge about the data can be garnered at once.

Figure 1. Example large display that may allow for a more embodied interaction approach to visual analytics of large scale data.

With small displays many of the body's built-in functionality, such as peripheral vision, are wasted. Gone is the ability to make use of movement in the periphery or to maximize spatial memory. Spatial memory has been shown to be highly effective for categorization and memory (e.g. [22] and [25]). Optical flow is the continuous stream of input to the eye that naturally exists in real environments, but is often lacking with many virtual environments such as with smaller displays and has been shown to improve navigation (e.g. [7]).

As opposed to desktop displays, larger displays allow more use of the human body's resources to interact and *physically navigate* with large displays. Physical navigation is changing where the user's physical eyes are looking. Any physical motion that affects the user's view, such as moving the eyes or head, walking, crouching, standing, sitting are all forms of physical navigation. On the other hand, *virtual navigation* is moving from one point to another in space and scale through the use of external devices, such as a mouse, that manipulate the underlying view shown on the display.

When the user's viewpoint does not encompass the data domain in its entirety, virtual navigation is required [8]. Being unable to see all of the data at once, the user is forced to integrate the information shown on the display into a mental representation often called a cognitive map. The user then must use their cognitive map to navigate the data to gain insight [24]. This can be problematic for two reasons. First, instead of using one's cognitive resources to understand the data, much of the time and effort is spent navigating the data. Second, such cognitive maps are often incorrect and rely on landmarks, particular patterns or pieces of data, and result in distorted cognitive map [25],[32].

Real environments typically do not represent data. As a result, navigating in a virtual environment to better understand one's data demands greater accuracy of the user's cognitive map than do real environments [33]. The motor, peripheral vision, and proprioceptive cues that come from

[1] The word resolution historically means the density of pixels on the screen, usually in terms of dots per inch (DPI) [1]. However, it is becoming common practice to refer to resolution as the number of pixels on a display, especially when people use the term "high-resolution displays".

walking and turning one's body and head that help in forming a cognitive map are often absent from small display environments [25].

In addition to allowing people to use more of their physical resources to navigate in space and scale, by using large, high-resolution displays, visual analytics is hypothesized to be more scalable (e.g. [37]). For example, people can see larger overviews and more details about their data at once.

However, will larger display sizes enable more efficient visual analytics? Theoretically larger displays should enable better performance; however, how well embodied interaction actually helps in real environments needs to be determined. If large displays do allow more efficient visual analytics, how much more do they help? And, possibly the most important question, why do they help?

In order to answer these research questions we present a theoretical extension of the space-scale model to further understand how physical movement and positioning plays a role in space and scale. Second, we present empirical proof to suggest that embodied interaction with large displays can greatly improve performance time when visually analyzing large scale data. In addition, we show empirical proof that people prefer physical navigation over virtual navigation when given a choice.

In summary, we present initial evidence to suggest that when people are given the opportunity to use their body and mind together (embodied interaction) to accomplish an analytic task, they are able to perform at greatly enhanced performance times with the same degree of accuracy. Instead of being forced to construct cognitive maps, which are error prone (e.g. [25]), people can use their bodily resources to help free up precious cognitive resources to better accomplish the analytic task at hand.

For the visual analytics community this translates into improved performance and comprehension. In the end, large displays provide the potential for better and faster understanding of large datasets which is one of the primary goals of visual analytics.

2. Related Work

Embodied interaction is "interaction with computer systems that occupy our world, a world of physical and social reality, and that exploit this fact in how they interact with us" [10]. A relevant example of such interaction comparing physical to virtual interaction was presented by Bowman, et al. [4]. They found interesting results forcing virtual navigation in a CAVE environment. They report that after a few episodes of forcing participants to use virtual navigation that participants continued using virtual navigation even when it was not required. In the situation where virtual navigation was not forced, more physical navigation was seen, and performance was higher than when virtual navigation was forced on the users.

Large displays also allow for greater use of embodied interaction and consequently better performance. For example, Czerwinski et al. [6] explain the current state of performance measurements and explain that their own study showed conclusively that participants using a multi-monitor configuration affording increased resolution (3 monitors wide) performed better than on a single monitor.

Tan, et al. [30] also show how retention can be increased by using extra screen space to display different images in the user's periphery to help recall more from a particular task session using their prototype called Infocockpit. In addition, Shupp, et. al [28] explored how performance of large displays varies with display size and display curvature.

A few longitudinal studies showing benefits of multiple monitors have also been performed. Bishop and Welch [3] created a "desktop" environment that used projections on the wall to alleviate bezel and ergonomic issues. They report improvement in everyday work and an increase in physical interaction. Ball and North [2] performed a similar study but with multiple LCD monitors for a six month period of time with multiple users. They report a number of benefits in perceived increase in productivity and problems with bezels, adaptation to the display, and interaction problems.

As more studies show the usefulness of large displays, different interactive techniques have followed. A number of different types of techniques, from using less traditional input techniques to different ways of interacting with the mouse have been developed. Large displays and multiple monitor displays are inherently different from smaller displays and logically should be interacted with differently [29].

Touch screens and camera-based touch gestures have been implemented with large displays (e.g. Ringel, et al. [20]). Other well-known interaction techniques with large displays also exist, such as laser pointers (e.g. [17]), head-tracking (e.g. [1]), and hand gesture tracking (e.g. [34]).

Khan, et al. [13] created an interface for physically larger displays that allows a user to see through a "telescope," similar to a porthole, to another part of the display. The user then can manipulate the other part of the display through the telescope similar to remote computing. In addition, Microsoft Research has been active in the area of interaction for large displays. Their work is summarized in [21].

The work on large displays is starting to reach a more mature level. As a result we present theoretical and empirical evidence to help guide the future direction of research with large displays; we show how large displays take advantage of the entire physical body, not just the cognitive mind and how it impacts space-scale navigation in information visualization.

Previous work has largely measured the potential for user performance benefits. As a result, the question becomes not if large displays have performance benefits, but why they have such benefits and how they can be harnessed.

3. Expanding the Space-Scale Model

This section expands upon the basic space-scale model from Furnas and Bederson [12]. Specifically, it expands the theoretical ideas of space and scale to include the idea of a changing viewport size and the idea of a difference in behavior between virtual and physical navigation.

In order to explain these concepts we present a few definitions. *Scale* refers to the zoom level, or magnification level, of the data. *Space* is the total amount of area that the information takes up at the given zoom level. *Viewport* is some subset of the area of the space that a user sees. It is fixed in size regardless of scale. As opposed to space whose area is defined by scale, the viewport size is constant and is based on the hardware of the display.

3.1. Overview versus Detail

When dealing with large amounts of data, analysts are often concerned with different scales of that data from various perspectives [37],[38]. For example, one task might be concerned with the general overview of a data set while another task might be concerned with only certain detailed parts of the data.

In space scale there exists a tradeoff of overview versus details as viewport size increases, or in other terms a tradeoff of visible data space versus data scale. To generalize the idea, if the viewport size increases then either detail, overview, or a combination of detail and overview must increase as well. On the other hand, if the viewport size decreases the opposite must happen. Figure 2 shows a visual representation of how viewport is associated with detail and overview and how the number of pixels is associated with visible data space and visible data scale.

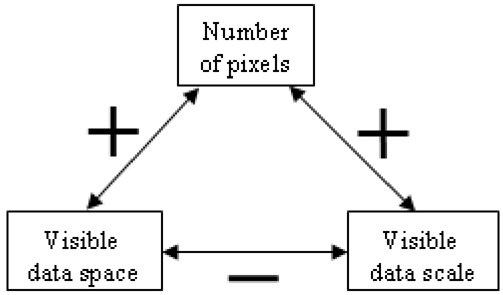

Figure 2. Visual representation on how number of pixels, visible data space, and visible data scale relate.

Figure 3 shows how the relationship can be applied to the space-scale diagram with a constant viewport size. The left image of the map shows a certain amount of overview of the data. As one zooms into a deeper scale (the image on the right) one loses more of the overview. One cannot maintain the same amount of overview while increasing scale with a constant viewport size. This is true even with the use of visualization techniques such as focus plus context.

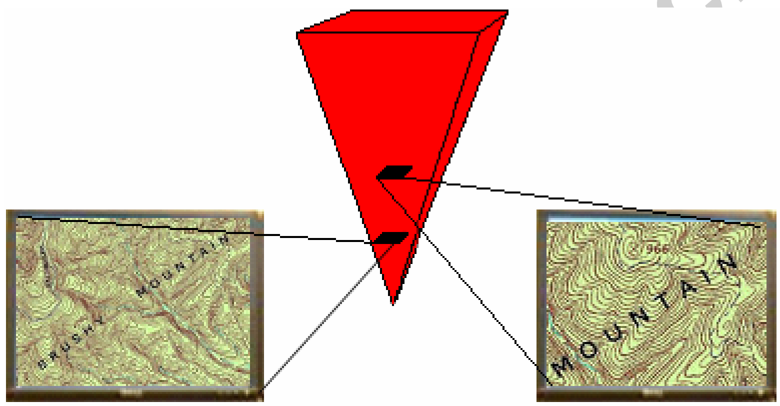

Figure 3. Space-scale diagram showing viewports with greater overview (left) and greater detail (right).

Figure 4 is a two-dimensional space-scale diagram that shows three different viewports. The bottom viewport is a small display. The second viewport is at approximately the same scale but shows more space as it is larger. The third viewport (the top viewport) shows the same amount of space as the small viewport but with greater scale (more detail).

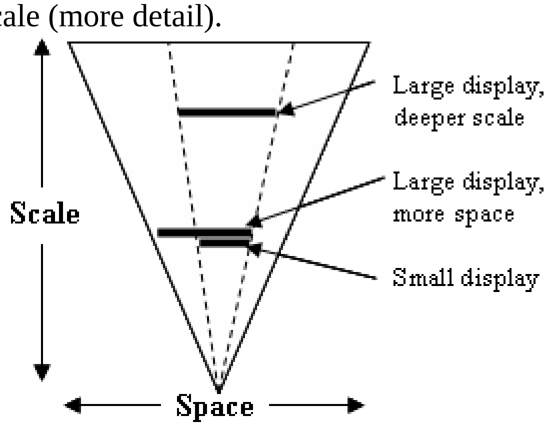

Figure 4. Augmented space-scale diagram augmented to showing how a larger display, a display with more pixels, can show more overview and details than a smaller display.

3.2. Space-Scale Navigation

3.2.1. Physical Navigation Extension

If one augments the space-scale model with physical navigation, then one is able to *physically* zoom and pan in space scale as well. Figure 5 shows two space-scale models that are augmented with physical navigation. In the model there is a black rectangle that represents the current viewport of the data. Based on that viewport, users can physically navigate the area in blue, a subset of the space-scale model. To change the viewport to see the areas in red, users must virtually navigate.

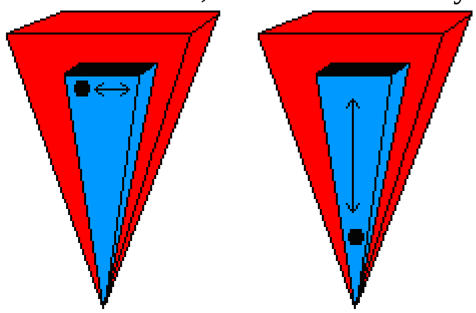

Figure 5. Two space-scale models augmented with physical navigation. The left model shows how a person (depicted as a black dot) can physically pan. The right model shows how a person can physically zoom.

In Figure 5 the model on the left shows a black dot that represents a person in 3D space. This person is able to physically "pan" by moving parallel to the display. The model on the right shows a person that is able to physically "zoom" by physically getting closer or further away from the display. The blue area of physical navigation is *within* the virtual navigation space. As long as one is interested in only the data represented in the blue area, no virtual navigation is required.

An example of someone physically panning is moving from one point of the display to another point while maintaining the same distance away from the display. This could be accomplished by moving the body, such as walking, to another part of the display.

Physical panning is similar to virtual panning in that a person does not change the virtual zoom level, but is able to see different data in the same space. However, the difference between virtual panning and physical panning is that physical panning does not change the viewport location in space. One can only look at what is currently being shown on the display. Virtual panning actually changes the view by virtually moving the viewport.

The larger the viewport (e.g. large high-resolution displays) the larger the space that can be explored with physical navigation and the less virtual navigation is required. For instance, if the space is smaller than the viewport size then virtual navigation is not required. If the space is larger than the viewport then virtual navigation is required to view the entire space.

Physical zooming and virtual zooming are fundamentally different from each other. Physical navigation does not change the zoom level, only the distance away from the physical display. Physical zooming involves visual aggregation, visual acuity, and visual perception. Virtual zooming is usually performed at a constant physical distance (e.g. sitting at a chair) and manipulates the virtual zoom level. With virtual navigation changing zoom level may affect the view of the data through geometric zooming or computation aggregation (e.g. semantic zooming).

3.2.2. Computational and Visual Aggregation

There are a number of types of aggregation techniques. For example, one common type of aggregation is mathematical aggregation that shows one or more object attributes averaged. Another example shows more or less data attributes at each zoom level; this is commonly called semantic aggregation or semantic zooming. These types of aggregation techniques are *computational aggregation* as they are calculated by a computer. Computational aggregation is a common technique for displaying large datasets on smaller displays.

Computational aggregation is helpful in getting precise overview statistics such as finding out what the exact average of a particular attribute. Also, it is helpful in hiding non-relevant details. For example, computational aggregation can be helpful in hiding unnecessary details that are not important at the time such as showing all the streets and roads in the United States when one is only interested at the state level through the use of semantic zooming.

However, computational aggregation has drawbacks. First, finding out the average of a particular attribute falls away from the reason of visual analytics: one can query a database and find the same answer. However, one cannot query a database to find the trends and patterns that one sees with visual aggregation. Second, sometimes it is best not to hide information. Often one does not know what details are missing due to computational aggregation and as a result misinterpretations or misunderstanding can occur in analysis.

Visual aggregation on the other hand is aggregation that is performed by users' physical eyes. For example, if a person were looking at a visualization and stands back (zooms out via physical navigation) then that person is not able to see as much detail and visually aggregates the details through visual perception. Visual aggregation is helpful in finding the trends, or patterns, of data at a detailed level. As people are able to see all of the details for a particular zoom level they are able to more fully see all the detail at once and mentally aggregate the data themselves.

However, visual aggregation introduces problems of color aggregation, visual distortions, etc. For example, if color is meaningful to a visualization then colors that are not present may appear when the user is physically distant from the display. For example, a series of blue and green color placed close together may appear to be cyan (neither blue or green) when standing ten feet away from the display. This non-present color may indicate to the user a value that does not exist in the underlying data. For more information on visual distortions see [35] and [11].

In summary, physical zooming causes visual aggregation which allows people to see more patterns and trends but may lead to visual distortions if the visualization is not designed correctly. Virtual zooming causes computational aggregation and is helpful in showing precise overviews and hiding irrelevant data but may lead to misjudgments or misconceptions.

3.3. Physical/Virtual Tradeoff

In addition to a space and scale tradeoff there exists a physical navigation to virtual navigation tradeoff as well. Figure 6 shows two example space-scale diagrams. The left image represents a smaller viewport (e.g. a smaller display) than the right image.

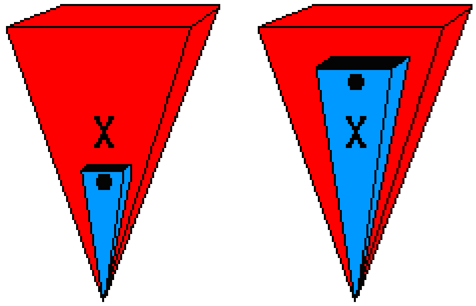

Figure 6. Two images showing the same target in the space scale. The black "X" indicates the target that one is interested in.

On each diagram there is a destination data point, or target, indicated by a black "X." In order to access the target at that point on the smaller display (the image on the left) one must virtually zoom in. However, for the larger display one is given an option to physically navigate (e.g. stand back), to virtually navigate (e.g. zoom out), or both (e.g. stand back a little and virtually zoom out).

So, if the target is within the blue area the user may decide to physically or virtually pan and zoom. This opportunity allows the person to decide between visual or computational aggregation. If physical zooming is chosen then a more complete mental model may be created of the target with the details surrounding it.

In addition, performance time may be affected. Standing back may be quicker if the person is already standing, but may be much slower if sitting. Also, the dataset may be large and a certain amount of computing time must be given to change the view for the correct zoom level requiring a non-trivial computation.

In summary, the larger display on the right of Figure 6 allows a choice between physical and virtual navigation where the smaller display on the left does not. The larger display on the right allows a person to perform both physical *and* virtual zooming to get different perspectives of the data which is impossible with the smaller display on the left.

In summary:

1. Large displays allow more data to be shown at once allowing for greater potential visual perception and data scalability.
2. Large displays allow users to have larger visual aggregations of their data through physical zooming. This visual aggregation allows for different perspectives that may not be possible with smaller displays.
3. Larger displays allow for more choice between virtual or physical navigation. This choice affords an opportunity for users to improve their performance times by choosing between which navigation will allow the best results for their tasks.
4. If the given target is entirely within the physical navigation subset of the space-scale diagram (the blue part of the pyramid) then the user can maximize his use of embodied resources and not need to use virtual navigation at all.
5. Together, with users being able to perceive more data and choose between virtual and physical navigation, they are able to perform their tasks potentially faster and potentially gain more insight into their data.

4. Physical and Virtual Navigation Experiment

The extension of the space-scale diagram to include physical navigation (e.g. see Figure 5 and Figure 6) is the realization, or the theoretical link between embodied interaction and visualization. This theoretical link provides a better model for understanding how data and users relate in space and scale.

However, as it is still only a theoretical link, we ran an experiment to better understand empirically how people are able to use their embodied resources to facilitate better performance times. We wanted to see how performance, physical navigation, and virtual navigation are affected by different display sizes. In essence, how do people's physical and virtual behavior change as display size changes?

The independent variables for the experiment were as follows:

1. Viewport size (e.g. display size)
2. Task type
3. Task scale (scale/detail level of tasks)

The dependent variables for the experiment were as follows:

1. Performance time (or number of insights for the insight task)
2. Physical navigation (i.e. participant's movement in 3D space)
3. Virtual navigation (i.e. mouse interaction)

4.1. Data and Visualization Explanation

We created a visualization of 3,500 houses for sale in Houston, TX. The visualization displayed data about the houses on a map of the Houston area, and used semantic zooming, as shown in Figure 7. Figure 7.a shows only the geospatial position and bar charts of the prices of the houses. When the user zoomed in, prices were shown as text (Figure 7.b), and further zooming resulted in the display of square footage, number of bedrooms, and number of bathrooms, in addition to price (Figure 7.c).

Figure 7. a) Image showing only a bar chart of normalized price values and geospatial position. b) Image showing the houses at a deeper scale - text values are also shown. c) Image showing all the details about a house.

In our semantic zooming scheme, zooming only resulted in more information being displayed. To see all of the houses with all the details shown would require about a 100-monitor display (131,072,000 pixels).

We used a modified version of the NCSA TerraServer Blaster [31], an application that views images from US Geological Survey. Specifically, we modified the application to zoom and pan via direct mouse manipulation instead of using a control panel, and by adding superimposed data visualizations to the base map.

4.2. Display Used

The display used for the experiment was made up of twenty-four seventeen-inch LCD monitors in an 8×3 matrix (Figure 8). Each monitor was set to the highest resolution of 1280×1024. We removed the plastic casing around each monitor to reduce the bezel size (gap) between monitors. Twelve Linux-based computers drove the display.

In order to simplify the experiment participants were tested on different widths of the display by column number (Figure 8). For example, in the four-column condition only the first four columns would be used, and columns five through eight would be left unused. In the eight-column condition all columns, one through eight, would be used.

It is important to note that the viewport size condition controlled the number of pixels of the display, but we report in terms of the number of columns for simplicity. A seamless display would have been ideal, but we find that high pixel density is more important. To give an idea of the number of pixels per column, one column had 3,932,160 pixels so that eight columns had 31,457,280 pixels.

Figure 8. Image showing how the display was artificially separated into eight different widths. The total resolution of the display is 10240 X 3072. The physical dimensions of the display were roughly 9 feet (2.7 m) by 3.5 feet (1 m).

Each task began with the overview/best-fit of the map (see Figure 9). This preserved the aspect ratio of the base map so that each column showed the same amount of overview but with different amounts of detail. In other words, the same area of Houston was shown each time but with greater amounts of detail for larger viewport sizes.

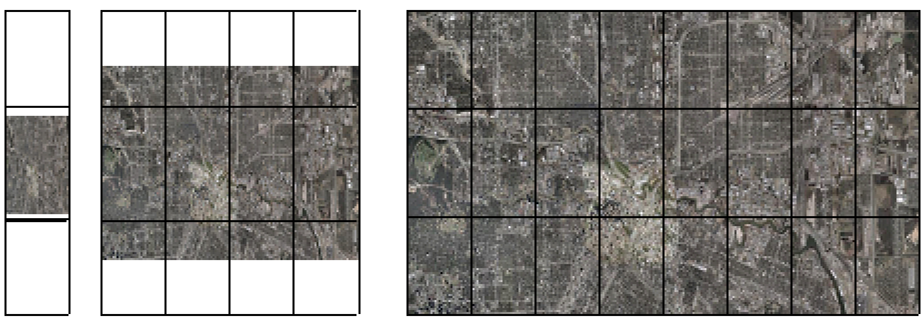

Figure 9. Example of best-fit of the geographic base map for the one column, four column, and eight column conditions without the houses visualized.

As participants zoomed in on display sizes that did not completely fill the display vertically (e.g. the one column condition) that more of the viewport was filled with the map. The left image of Figure 10 shows how the one column condition might start out by only seeing a small overview and not filling the display. However, the right image of Figure 10 shows greater detail and uses more of the display as participants zoom in.

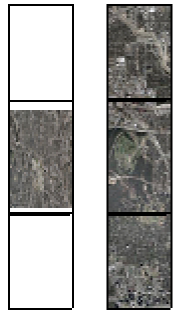

Figure 10. Example of why the one column condition does not use the entire display at the beginning of each task with the "best-fit" view. The left image corresponds with the starting point "best-fit" and the right image corresponds with a zoomed in view of more detail.

4.3. Tasks

The search and pattern finding task had three levels of scale: high, medium, and low; the navigation task had two: medium and high; the insight task had none - all scales were important. A "high" task was an overview task in which only geospatial position of the data was taken into account. A "medium" task had some higher-level details required to complete the task, such as the price of a house. A "low" task required all the lower-level details, such as the number of square feet, to complete the task.

Figure 11 visually shows how much detail was shown at the beginning of every task with best fit and what the visualization might look like at different scales (e.g. high, medium, and low). For example, column conditions one through six started out only seeing the geospatial positions of each house and a task bar approximating the house's price; this is shown as the bottom image on the right. However, column conditions seven and eight were started at a scale that crossed the semantic threshold (shown as a dashed line); this is shown as the middle image on the right.

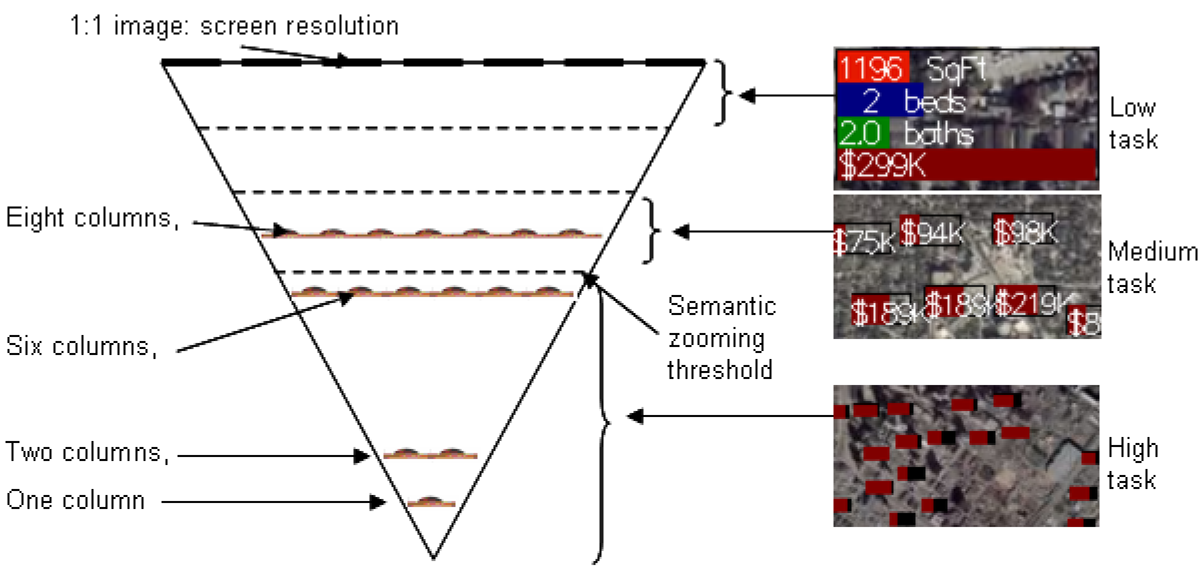

Figure 11. Space-scale representation of display sizes and semantic zooming thresholds.

In order to see all of the details of a house a participant would have to virtually zoom in to the lowest semantic zoom threshold; this is shown as the top image on the right. The thick dashed line at the top of the space scale diagram indicates a one-to-one mapping of detail. Zooming in beyond that point would not show any more detail to the user but only enlarge the details (e.g. pixilated).

The navigation task was created as a benchmark against the other tasks. Its purpose was to determine the minimal time it would take a person to access different levels of details for different display sizes.

For the navigation task, a single house was shown on the display. The participant was asked to verify that he could see the house before proceeding. The reason for this verification was to ensure that the participant was not being asked about their ability to find the house. After verifying the presence of the house he was then asked for an attribute about the house (e.g its price). No overview task of indicating geospatial position was used because the participant was required to see the geospatial position of the house before the task began in order not to test perception and only navigation. The task was complete when the participant had spoken aloud the correct corresponding price or square feet of the house.

The search tasks involved searching, or finding, particular houses that had particular attributes (e.g. find a house between $100,000 and $110,000 inclusively). There was not one correct answer per task as several houses fit each criterion.

Pattern finding tasks found patterns for all the displayed houses. For example: “Where is the largest cluster of houses?” “What is the pattern of the prices of the houses?” “What is the pattern of the number of bedrooms of the houses?”

In order to measure only performance time and not accuracy for the first three tasks participants were asked to continue until the task was completed correctly. For instance, in the pattern task participants searched for the correct pattern until they reported it correctly.

The open-ended insight task followed Saraiya, et al.’s [26] model of evaluating different information visualizations based on the depth of insights. However, instead of evaluating different information visualizations, different display sizes were evaluated. For this particular task participants were given a mobile lecture stand with wheels on which to write insights. Figure 12.b shows a participant using the mobile stand. Unlike the other tasks that each recorded performance time, the open-ended insight task involved participants writing as many insights about the data as possible in ten minutes.

4.4. Interaction

All interaction with the display was performed using a wireless Gyration GyroMouse. The wireless mouse was used so as to not encumber participants as they walked around [28] (see Figure 12.a). Zooming used the scroll wheel on the mouse and was performed relative to the mouse cursor; the position of the cursor became the center of zooming. Panning was performed by holding down a mouse button and then dragging the map.

To track physical navigation in 3D space, we used a VICON vision-based system to track the users’ head (Figure 12.b), but head movements did not change what was shown on the display. All participants stood during the experiment to allow for physical navigation. A chair was provided during breaks between tasks.

Figure 12. a) Image showing a participant using the gyro mouse with the display. The gyro mouse is enlarged in the red square. b) An image showing the hat used to track users' position.

4.5. Participants

The experiment had 32 participants (10 females and 22 males). Approximately half the participants were from the local town and the other half from a variety of majors from the university. The ages of the participants ranged from 24 to 39 with an average age of 28.

4.6. Protocol

The experiment consisted of four tasks: basic navigation, search, pattern finding, and insight finding. The first two tasks, basic navigation and search, were a with-in subject design in which all 32 participants performed on all eight column widths using a Latin Square design.

The second two tasks, pattern finding and insight finding, were between-subject designs. Only the 1, 3, 5, and 7 column conditions were used to increase statistical power by having eight participants in each cell instead of only four.

A general tutorial time of about five minutes was given for each participant before they began. After the tutorial the participant would perform the navigation and search tasks, on each column condition. Then, participant would then be randomly assigned to a single column condition to perform the pattern and insight tasks.

5. Experiment Results

For the first three tasks (navigation, search, and pattern finding) the performance times were analyzed. For the insight task, the participant's insights were graded for depth by domain experts. However, after analysis of variance was performed on the insight grades non-significant results were found due to high variance in the answers.

5.1. Performance Time Analysis

A number of related research has shown that large displays provide performance benefits over smaller displays which usually compares a single small display to a single large display. However, what do the performance curves look like for a series of display sizes?

In order to analyze performance results we ran a two-way ANOVA on performance times with column widths as a continuous variable, and tasks as a discrete variable (i.e. navigation, search, and pattern tasks). Our results found a main effect for column widths ($F(1,1324)=20.56$, $p<0.01$) and task type ($F(2,1324)=77.05$, $p<0.01$). The scale of the tasks was not included as the different tasks used different scales (e.g. the navigation task only had a medium and low task).

In other words, we found that there was a statistical significance in column widths and with task type. With the task type we performed a post-hoc Tukey HSD analysis that showed that the different task types were all in different groups. As each task type was statistically different from each other we performed individual ANOVA's for each of the tasks (see Table 1).

Table 1. Statistical results on performance time.

	main effect of column width	**main effect of task scale**	**interaction effect**
navigation	F(1,508)=118.9, p<0.01	F(1,508)=98.1, p<0.01	F(1,508)=4.09, p=0.04
search	F(1,762)=38.18, p<0.01	F(2,762)=130.13, p<0.01	F(2,762)=9.34, p<0.01
pattern finding	F(1,90)=3.53, p=0.06	F(2,90)=89.65, p<0.01	F(2,90)=3.22, p=0.04

Figure 13 shows the general trend of performance results. It should be pointed out that the navigation low task (highly detailed task) on the six column condition is an outlier due to the target being displayed across a bezel. Targets were placed in the space randomly. However, only the target on the low task on the six column condition occurred across the bezel. The main reason for the increase in performance time can be attributed to the additional time it took participants to pan the viewport so that they could clearly read the text that crossed the bezels. Additional information about bezels and their problems and solutions is discussed by Mackinlay, et al.[14].

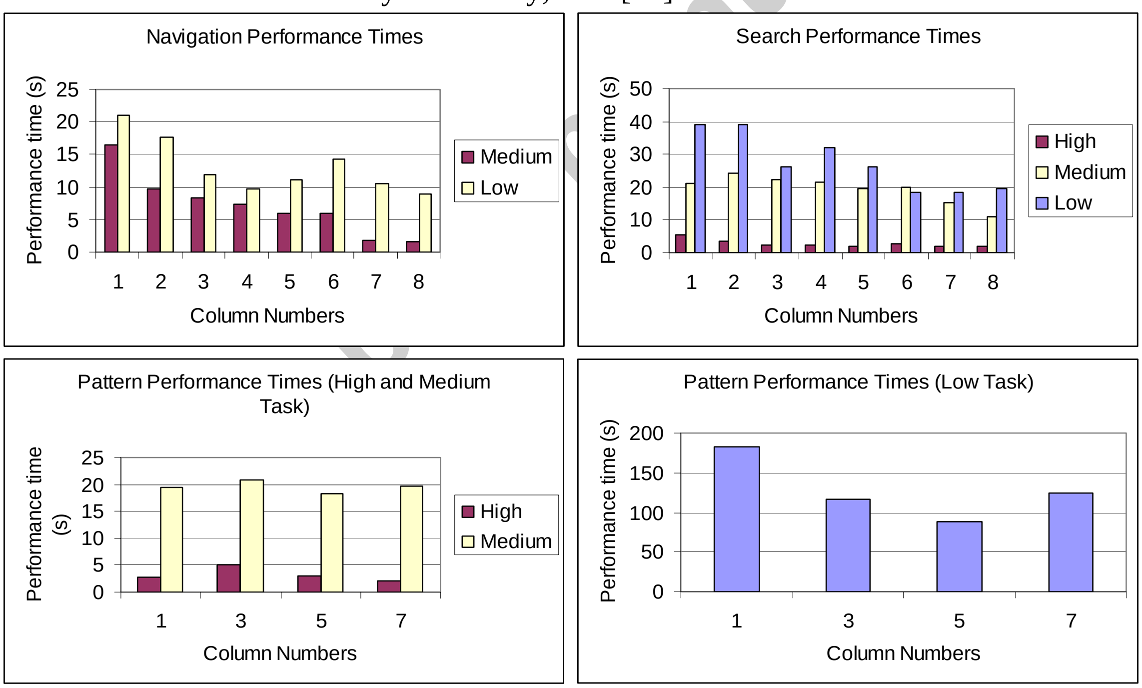

Figure 13. Performance averages for the navigation, search, and pattern finding tasks. Performance averages for the pattern task were separated to help with readability.

In addition, Figure 13 shows that the performance times largely depended on both the task scale and the column width. For instance, with the search high task (overview task) after the first column condition, the performance time appears roughly uniform. However, for the medium task (some detail) and the low task (highly detailed) there are different patterns.

The low task for the pattern task in Figure 13 appears to be slightly different from the other tasks. Instead of a linear decrease in performance there appears to be a slight increase in performance time at the seven column condition. Later on in this paper it is explained that there was a drastic increase in physical navigation for that condition. Therefore, it is possible that for more complex tasks there exits a point where performance time increases instead of decreases as display sizes pass a certain threshold. It is also possible that the slight increase in performance for that condition is also due to noise or random variation among the participants.

In summary, larger viewport sizes can drastically decrease performance times. For example, on the medium navigation task (some detail), performance time was reduced more than ten times from 16.3 seconds on the one column condition to 1.5 seconds on the eight column condition. Another example is the low search task (highly detailed) where performance was reduced more than two times from 39 seconds on the one column condition to 19 seconds on the eight column condition.

5.2. Virtual Navigation Analysis

Again, it is evident from related work that less virtual navigation is performed with larger displays (e.g. [26]). However, it is unclear what the virtual navigation curves look like. Also, why do such decreases occur?

In understanding the virtual navigation results it is important for the reader to understand why participants needed to virtually navigate. First, for each task there was a particular scale or zoom level that participants had to navigate to see the necessary details for the tasks (e.g. price of the houses for a medium task). Second, the participants would sometimes pan to move around the space. Panning was never required as moving around space can also be accomplished by a series of zoom movements - see [12].

To understand how virtual navigation differed generally, we performed a series of two-way ANOVA's on column widths and task types. First, we wanted to see how the number of zooms that a person performed was affected by column widths and task types. We found a main effect of task type ($F(3,1400)=416.2$, $p<0.01$), a main effect of column width ($F(1,1400)=34.8$, $p<0.01$), and an interaction of task type and column width ($F(3,1400)=2.4$, $p=0.06$). Post-hoc Tukey HSD analysis shows that the different tasks were all in different groups.

Another analysis of interest is the number of pans performed. The reader should note that the number of pans is only mouse movement that actively moves the viewport in space. It is not inactive mouse movement that is used to reposition the cursor without moving the viewport. The resulting ANOVA showed a main effect of task type ($F(3,1400)=301.3$, $p<0.01$), a main effect of column width ($F(1,1400)=63.86$, $p<0.01$), and an interaction of task type and column width ($F(3,1400)=17.22$, $p<0.01$). Post-hoc Tukey HSD analysis shows that the insight task was in a different group from the other tasks.

Table 2 shows the summary of statistical results for the various ANOVA's performed for the different tasks. Particularly, it shows results for analysis of the number of zooms and the number of pans that participants performed. Figure 14 and Figure 15 show the corresponding visual charts.

Table 2. Statistical results of the virtual navigation data for the different tasks.

	main effect of column width	**main effect of task scale**	**Interaction effect**
navigation task - number of zooms	$F(1,508)=144.6$, $p<0.01$	$F(1,508)=198.8$, $p<0.01$	$F(1,508)=9.5$, $p<0.01$

navigation task - amount of panning	not statistically significant	F(1,508)=10.1, p<0.01	not statistically significant
search task - number of zooms	F(1,762)=114.1, p<0.01	F(2,762)=270.0, p<0.01	F(2,762)=16.5, p<0.01
search task - amount of panning	F(1,762)=26.7, p<0.01	F(2,762)=23.9, p<0.01	F(2,762)=16.3, p<0.01
pattern finding task - number of zooms	not statistically significant	F(2,90)=72.9, p<0.01	not statistically significant
pattern finding task - amount of panning	F(1,90)=7.8, p<0.01	F(2,90)=29.9, p<0.01	F(2,90)=7.6, p<0.01
insight task - number of zooms	not statistically significant	n/a	n/a
insight task - amount of panning	not statistically significant	n/a	n/a

Figure 14 shows two things. First, in general, the larger the viewport size the fewer the number of zooms performed. Second, task scale is also important in understanding the number of zooms. The level of detail shown increased as the viewport size increased (due to starting each task at a best-fit overview position) to the point that no zooms were necessary for some tasks, such as the navigation medium task on the seven and eight columns. Clearly, when not all the detail that is necessary is shown then one must zoom in to see it.

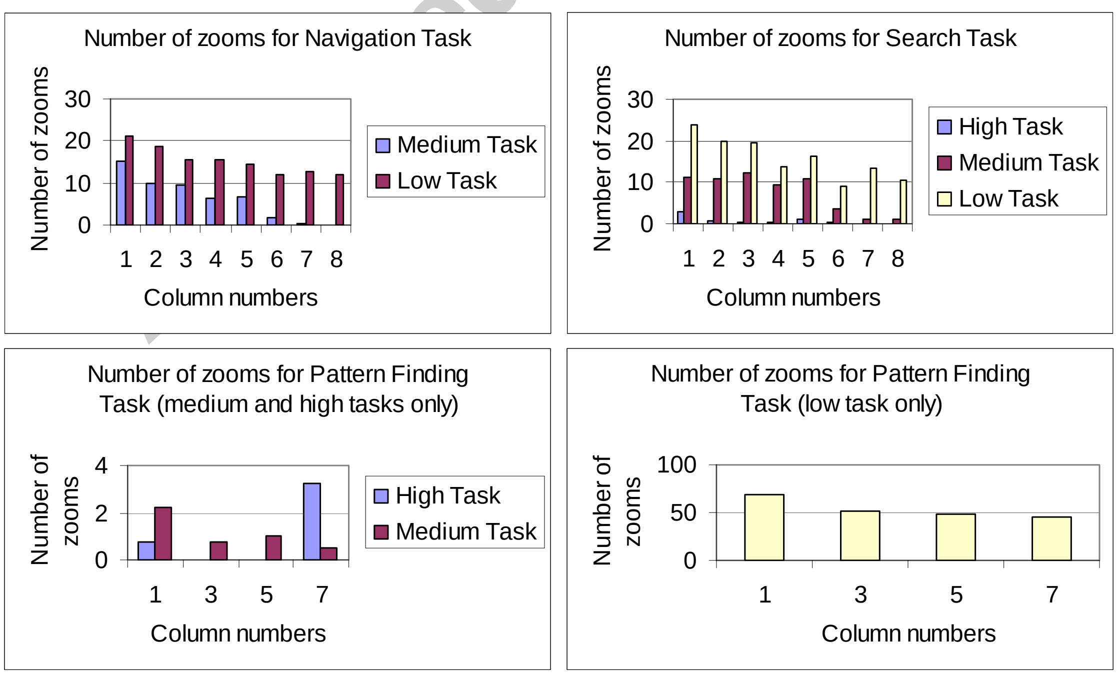

Figure 14. Average number of zooms for the navigation task.

The exception to the normal trend that we have seen is in the pattern high task on the seven column condition where participants were observed to virtually zoom out to better see the overall pattern. Previously participants were only observed to virtually zoom in. However, the seven column condition started out showing more details than were needed for the high task. As the task involved only finding the pattern of the geospatial positions of the houses, the additional details of the houses was a distraction. As a result, participants were observed to first physically zoom out (step back) to get a better overview of the data. However, as the additional details were a distraction, participants would then virtually zoom out to go to a higher semantic view to more easily see *only* the geospatial pattern.

The implications of this finding are that more details are not always preferred. Semantic zooming was created for the very reason that too many details at once can be distracting. Therefore, it is logical to conclude that understanding virtual navigation does not simply mean how much people might zoom in, but how much they might zoom out as well. This would be particularly important when doing multi-scale comparisons.

Figure 15 shows the corresponding amount of panning for the different tasks and scales. Comparing Figure 14 and Figure 15 shows that there was not any zooming or panning performed on the navigation medium task on the eight column condition. In other words, for that task participants chose not to virtually navigate at all but to use 100% physical navigation to complete their task.

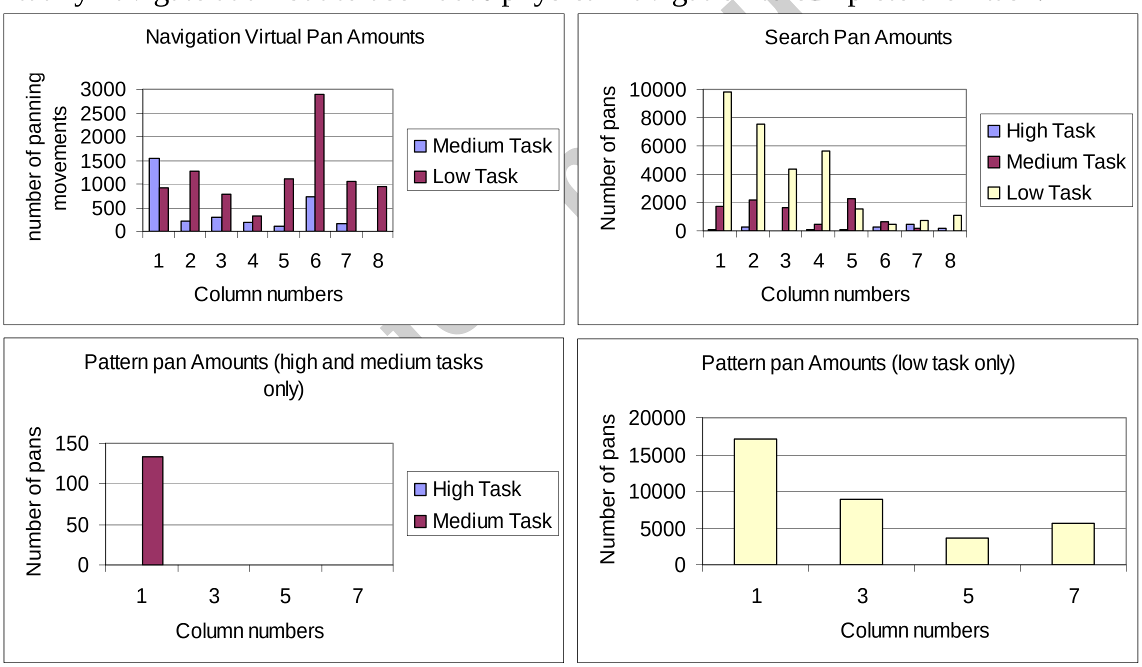

Figure 15. Average number of virtual panning for the navigation task. Average number of virtual pan movements for the pattern task. The high task did not have any recorded pan movements and the medium task only had recorded movements for the one column condition.

This trend of no virtual navigation also appears for the eight column conditions for the high search task and the three and five column conditions for the high pattern finding task. In other words, 100% of the participants (32 out of 32) chose to physically navigate rather than virtually navigate for these tasks. For each of these situations participants chose to physically reposition their bodies to view a particular spot on the display rather than manipulate the viewport and virtually change the view.

As explained previously in the extension of the space-scale model section, when virtual navigation is not required people have a choice to either virtually navigate or physically navigate. However, we found that when there is a choice that physical navigation is preferred over virtual navigation. For example, even with situations where not all of the participants chose to perform 100% of physical navigation the majority of participants chose not to virtually navigate. For instance, on the medium search task 90% (29 out of 32) of the participants did not zoom and 100% of the participants did not pan on the eight column condition.

As a side note, the six column low task was an outlier in pan amounts due to the target being located across a bezel. In general, most participants chose to pan the viewport so that they could more easily see the target.

5.2.1. Insight Task Virtual Navigation Analysis

The insight task was performed differently than the other tasks. First, it did not have specific task scales that we were testing for. Second, the reporting mechanism was different. Instead of having participants verbally speak their answers, participants wrote down their insights on paper. The rational was that more complete insights would be generated if written down than if spoken verbally. As a result of writing on paper, the participants were given a mobile lecture stand to write their answers on.

However, the resulting statistics did not differentiate column widths. In other words, there was not a statistical difference in virtual navigation between the different columns tested.

5.2.2. Semantic Zooming

In order to better understand how performance is related to viewport size and task scale we present Figure 16. Figure 16 shows how participants saw different views of the data based on the semantic thresholds and viewport size. As explained earlier, the visualization was started at the beginning of each task as a best-fit overview for every column condition. In other words, the larger the viewport size, the more of the visualization could be seen at once, the deeper the zoom level presented at the beginning of the task, and consequently the more detail shown (see Figure 11).

Figure 16 shows how the view of the visualization changed and at what points with the corresponding performance time. For the one column condition only the house positions were shown initially. However, the houses were shown as small squares that were hard to see. Thereafter the houses became easier to see until the price of the house appeared after the six column condition due to a semantic zooming threshold.

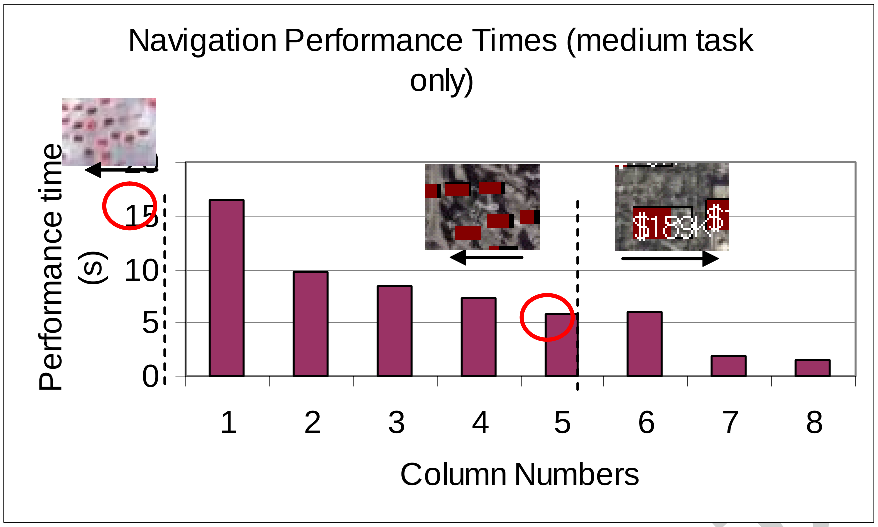

Figure 16. An illustration of what the geospatial visualization looked like for participants at different semantic zooming thresholds.

Comparing the semantic zooming thresholds in Figure 16 with Figure 14 (number of virtual zooms) shows that semantic zooming is the reason for the number of zooms. If the column condition started beyond the semantic threshold, such as with the seven and eight column condition, then no zooming was required. Looking again at Figure 16 shows that the decreased amount of virtual navigation corresponds to the decreased performance times of the task.

5.2.3. Virtual Navigation Analysis Conclusions

There are a number of things that the virtual navigation analysis shows. First, physical navigation is preferred over virtual navigation. When possible, people prefer to physically navigate to understand their data rather than virtually navigate. This find is especially important for visual analytics in that faster performance times can come as a direct result of physical navigation [4]. In addition, we show that larger viewports (larger displays) promote physical navigation where we have shown several instances where 100% of the participants chose to physically navigate on their own accord.

Second, semantic zooming is a key factor in understanding how much virtual navigation will be required for a particular task scale (see Figure 16). Knowing that particular tasks require a particular level of detail, the semantic zooming of the visualization dictates how much zooming in will be needed from a particular starting location.

The result is a series of linear step-wise performance curves. Specifically, it appears that there is a linear decrease in performance as display size increases within a semantic zoom threshold, but different linear performance curves between.

Third, the larger the viewport size, in general (with two exceptions), the less virtual navigation is performed. For example, with the number of zooms recorded for the low search task (highly detailed), the number of zooms decreased 2.25 times from an average of 24 zooms for the one column condition to 10.6 zooms for the eight column condition.

The first exception was where people zoomed out to see fewer details for an high (overview) pattern task – from 0.8 average zooms on the one column condition to 3.3 average zooms on the eight column condition. This confirms the need for semantic zooming, that all details all the time are not always helpful.

The second exception is with the low (highly detailed) pattern task – 5591 average pans on the seven column condition compared to 3636 average pans for the five column condition. More panning was seen on the seven column condition (the largest column condition tested for that task) which appears to have influenced the performance time in a negative way.

As smaller viewports have to zoom in more to reach distant targets, they also have the disadvantage of more disorientation per zoom. Each time a person zooms in they have to reorient themselves with the new view. The smaller the viewport, the more difficult it is to reorient themselves thus taking increasing performance time (e.g. [12]).

5.3. Physical Navigation Analysis

This subsection analyzes the physical navigation of the participants. Intuitively more physical navigation may result with larger displays. However, when will physical navigation increase? When physical navigation does increase, why does it do so? In order to answer these question we first analyze head rotation and head gaze of the participants then analyze the physical 3-dimensional positions (x,y,z positions) of the participants.

5.3.1. Head rotation

A normal healthy person has head rotation along three axes: x, y, and z. *Yaw* is the side to side motion of the head (e.g. looking left or right). This might be performed when a participant wants to look from side of the display to another. *Pitch* is the up and down movement of the head. A participant might perform a pitch movement when trying to look from the top to the bottom of the display. A *roll* movement of the head is tilting the head closer to one shoulder and farther away from the other. In general this kind of movement does not benefit participants in looking at a display.

Performing a 2-way ANOVA for all the head pitch data with task type and column width as variables resulted in non-significance. However, running a similar ANOVA for head yaw (side to side movement) found a main effect of column width ($F(1,1400)=4.6$, $p<0.01$), a main effect of task type ($F(3,1400)=1.67$, $p<0.01$), and an interaction of column width and task type ($F(3,1400)=3.7$, $p=0.01$). Post-hoc analysis shows that the navigation task was in a different group from pattern and search and the insight task was in both groups.

In other words, we did not find that pitch movement was statistically significant, but we did find that yaw movement was. The yaw results are intuitive in that one would expect a general increase in side to side motion for larger viewport sizes and harder (deeper scale) tasks as we varied the width of the display, not the height in the experiment.

In addition to analyzing physical head rotation movement, we performed analyses of where participants were looking. As we knew where the display was, what the display dimensions were, what the participant's physical location was at any particular time, and what their head rotation was, we accurately estimated where on the display the participants head gaze was. According to research on head gaze analysis, head gaze can be attributed to between 87-89% accuracy of eye gaze direction [16].

We performed a two-way ANOVA comparing the resulting total distance with task type and column width as variables. We found a main effect of task type ($F(3,48)=34.6$, $p<0.01$), and a main effect of column width ($F(1,48)=5.3$, $p=0.024$). Post-hoc analysis of the task types shows that the insight task was in a different group than all the other tasks.

Note that the "total gaze distance" for a particular task is the total length in terms of pixels that participants gazed at. So, a gaze distance of 200 pixels would be a sum of the distances of different pixels that participants looked at. Table 3 shows the statistical results from individual tasks of the head gaze data. Figure 17 shows corresponding visual charts.

Table 3. Statistical results of the head gaze data for the different tasks.

	main effect of column width	**main effect of task scale**	**interaction**
navigation task	not statistically significant	F(2,46)=14.4, p < 0.01	F(1,12)=4.2, p = 0.06
search task	not statistically significant	F(2,18)=31.26, p < 0.01	not statistically significant
pattern task	F(1,6)=8.87, p=0.025	F(2,6)=9.07, p=0.014	F(2,6)=5.17, p = 0.049
insight task	F(1,2)=50.25, p = 0.02	n/a	n/a

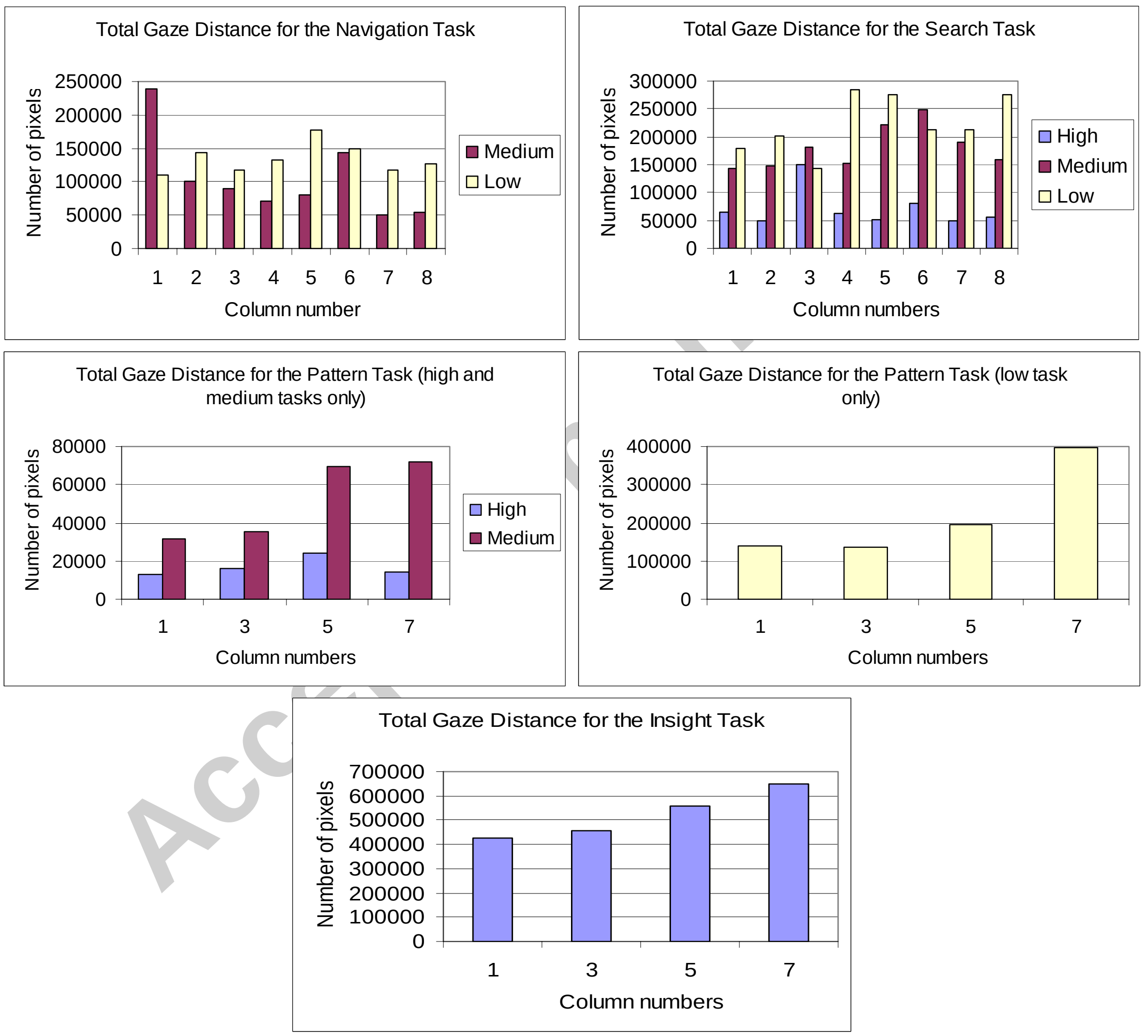

Figure 17. Average total gaze distances for the different tasks.

The head rotation and head gaze analyses show a number of things. First, they show that different head movement behaviors were exhibited at different task scales. This is intuitive as different task scales require different amounts of effort from participants.

For example, the high pattern task shows only a 93% increase of saccades for the seven column condition. In contrast, the low pattern task shows a 286% increase for the one column condition to an average of 397,802 saccades for the seven column condition.

Second, different tasks exhibit different behaviors. For example, a navigation task is different from a pattern task in a number of ways. One way is that a navigation task is a task where one must simply navigate to, or go to, the target. However, a pattern task relies on more perceptual cues as it includes many more targets. Second, the pattern task contains a number of comparisons of targets for the person to get a general idea of the data.

With these additional comparisons, additional viewport size can better used. In other words, the complex tasks (pattern and insight tasks) had a lot more of looking back and forth. This makes even more sense when the virtual navigation part is compared. It is intuitive to presume that a certain amount of data must be seen for the task to be finished. If less virtual navigation is performed then it is intuitive that more physical navigation must make up for it. For example, the additional comparisons in the pattern task are shown either in additional virtual navigation or additional physical navigation; regardless of which form of navigation is used, they must be performed to complete the task.

To better understand how physical navigation and virtual navigation have a relationship with performance we performed a correlation of performance to virtual and physical navigation. It appears that virtual navigation has a greater negative effect on performance than physical navigation. We found that the number of zooms correlated with performance with a correlation coefficient of 0.69, and the number of pans correlated with performance with a correlation coefficient of 0.68, while physical distance traveled did not significantly correlate with performance (correlation coefficient 0.46). In other words, increased virtual navigation correlates with increased performance time.

For the visual analytics community the impact of this empirical evidence is that if a task can be completed using either physical or virtual navigation then physical navigation will result in less virtual navigation and faster performance. Intuitively, if a person needs to view a certain amount of targets then they have to do it by moving themselves or moving the view virtually. However, moving one's body takes little thought in comparison to moving the virtual display. Therefore, we hypothesize that the improved performance time is due to how physical navigation makes better use of embodied resources than virtual navigation.

So, if a person can accomplish a task by physically navigating or virtually navigating, then physically navigating is often faster because of embodied resources. Also, virtually navigating is also slower because more cognitive demand must be spent on manipulating the virtual environment.

5.3.2. Physical Bodily Movement

In addition to head rotation and head gaze we analyzed participants' physical bodily movement. Although more bodily movement is expected, does it correspond exactly with head gaze? Also, does physical bodily movement increase linearly with the linear increase of display sizes from the experiment?

We performed our analysis by mapping the X, Y, and Z axes to the display. Figure 18 shows an illustration of how the three axes map to the large display. The illustration is a simulated top shot of looking at the display from above. The brown line forms the back of the display stand while each semi-circle represents the back of each of the individual column stands.

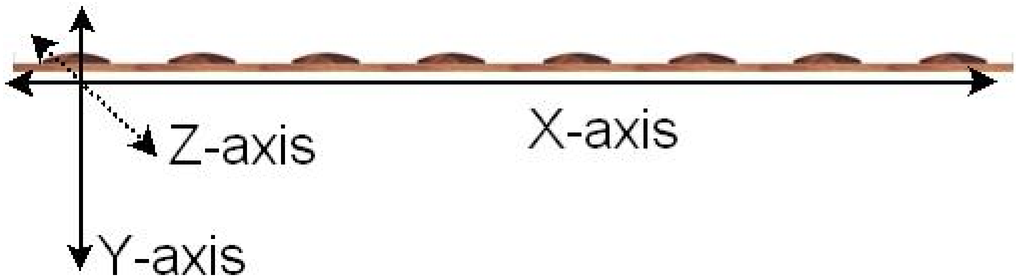

Figure 18. Illustration of the x, y, and z axes with relation to display.

The X-axis runs parallel to the display; a participant moving along the X-axis moves beside the display. The Y-axis runs perpendicular to the display; a participant moves along the Y-axis moves closer to or further away from the display. The Z-axis is the third dimension of the display, running along the height of the display; a participant moving along the Z-axis is getting closer to the floor (e.g. crouching) or getting closer to the ceiling (e.g. standing up). In effect, X and Z axes movement is physical panning while Y axis movement is physical zooming.

We analyzed participants' physical bodily movement in two ways: Range of position and total physical movement distance. Range of position is the actual range of physical area usage. This was measured by taking the maximum position and subtracting the minimum position. For example, taking the maximum position along the X-axis (maxX) and subtracting it from the minimum position along the X-axis (minX): maxX – minX.

Physical distance was calculated the by using a modified Douglas-Peucker algorithm [9]. By using the algorithm we guaranteed that what we were analyzing was actual movement from one physical location to another and not jitter of people standing at one location.

In this subsection we present a few of the more interesting results that best relates to this paper. In particular we only describe results of movement along the X-axis. Table 4 summarizes the statistical results of the X range of position for the four tasks. Figure 19 shows the trends that as the viewport size increases the range of X position increases as well.

Table 4. Statistical results of the X range of position data for the different tasks.

	main effect of column width	**main effect of task scale**	**interaction**
navigation task	$F(1,508) = 78.35, p<0.01$	not statistically significant	not statistically significant
search task	$F(1,762) =82.3, p<0.01$	$F(2,762)=27.18, p<0.01$	$F(2,762)=3.31, p=0.036$
pattern finding task	$F(1,84)=39.4, p<0.01$	$F(2, 84)=23.7, p<0.01$	$F(2, 84)=12.95, p<0.01$
insight task	$F(1,30)=5.6, p = 0.024$	n/a	n/a

An obvious outlier to the trend in the X position data is for the navigation task at the seven column condition. As explained earlier, the navigation targets being placed randomly; it so happens that the target for the seven column condition was randomly placed almost exactly in front of the starting position of the participants. As a result little movement parallel to the display was needed.

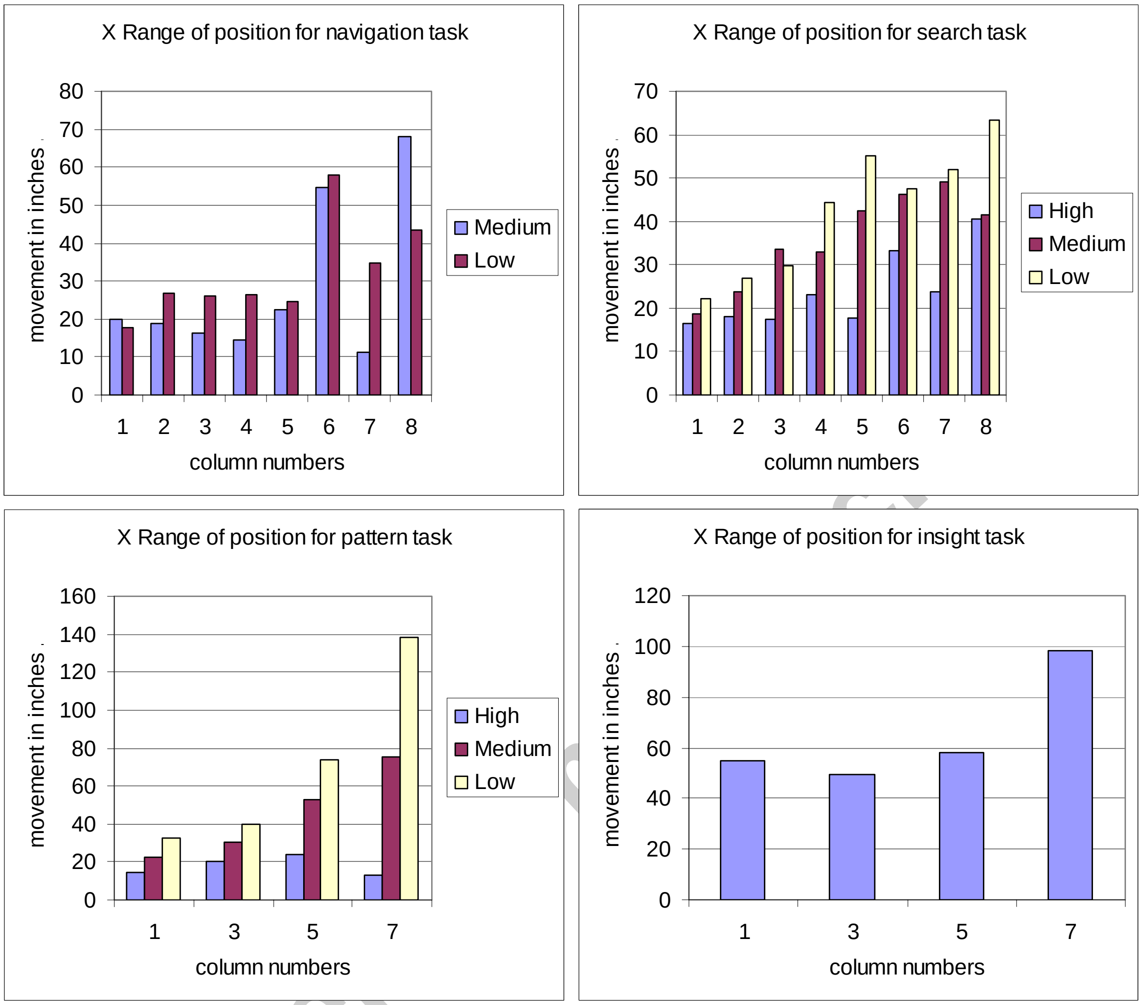

Figure 19. Average amount of range of movement in the X, Y, and Z axes for the navigation task.

Overall, there is a general trend that both column width and task scale make a difference for the X range of position. First, there is a trend that more detailed tasks (e.g. low tasks) use more of the display. Second, the larger the display, the more of the display is generally used; as viewport size increases participants take advantage of it and use the additional size.

Table 5 summarizes the statistical results of the X range of position for the four tasks. Figure 20 shows the trends of total distance covered. As opposed to X range of position, total distance takes into account moving back and forth over the same positions.

Table 5. Statistical results of the X total distance data for the different tasks.

	main effect of column width	**main effect of task scale**	**interaction**

navigation task	not statistically significant	not statistically significant	not statistically significant
search task	F(1,762)=4.52, p=0.03	F(2,762)=24.7, p<0.01	not statistically significant
pattern finding task	F(1,84)=16.62, p<0.01	F(2,84)=44.21, p<0.01	F(2,7.24), p<0.01
insight task	not statistically significant	n/a	n/a

Once again one can see that the amount of total physical navigation is different for different task scales. Although more detailed tasks took longer to complete, the reader should *not* understand that there naturally is a greater total physical distance because of longer tasks. Participants could have chosen not to move from a particular position and to only virtually navigate. However, this was not the case thus showing that people prefer to move around when given the opportunity and once again showing preference of physical navigation over virtual navigation.

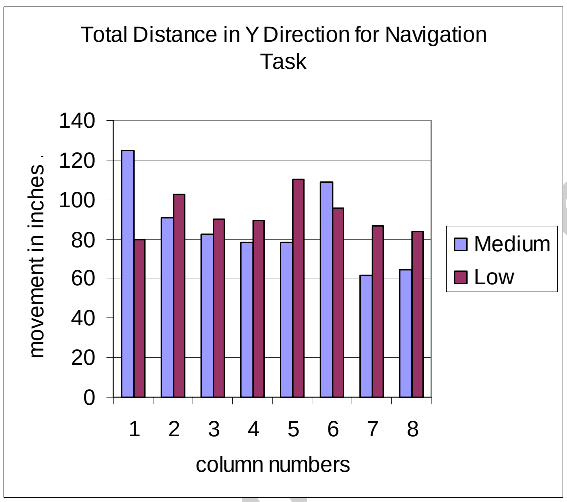

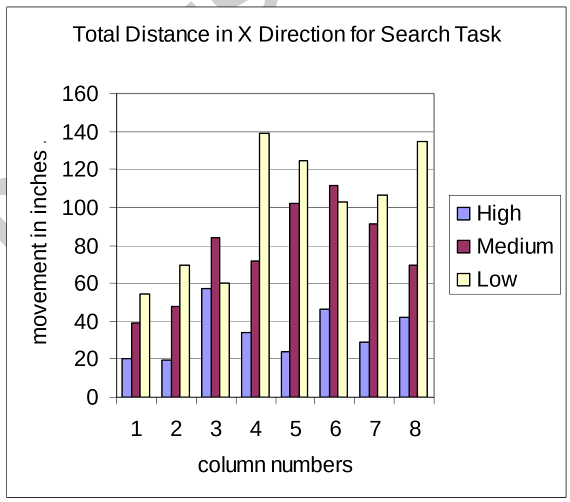

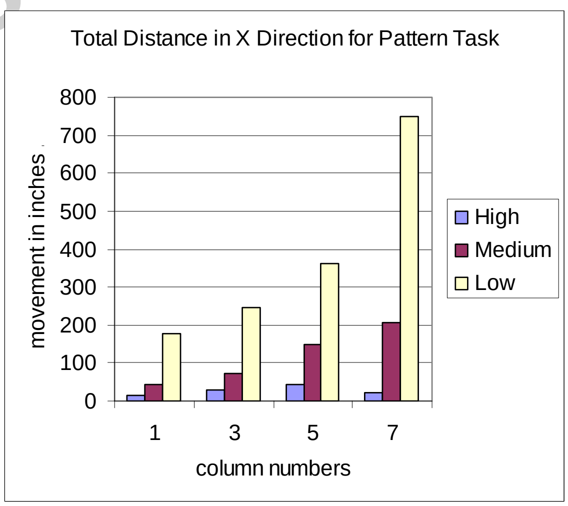

Figure 20. Average total distance of participants in the Y and X axes for the navigation task.

Analysis of the insight data led to non-significance. We argue that one of the reasons is due to the “tethering” effect. By inadvertently requiring participants to stay at one location they could not effectively make use of the entire large display. The mobile stand had the effect of keeping participants at the same location which did not allow them to move around freely. As they were not able to move around freely, they were not able to resolve (i.e. see accurately) as many pixels of the large display. As a result, the larger display did not help them as much as with the other tasks as the participants could not perceive the data far from them.

5.3.3. Physical Navigation Visualization

This subsection answers the question of where on the display participants looked at. In order to better understand the effects of large displays, we created visual representations of head/eye gaze projected onto an image of the display. Figure 21 is an example of physical movement for the pattern finding task at different column conditions for different participants.

Each set of images is for a single participant. The top image is an “overhead camera shot” of the participant involved in the task. The brown line with semicircles represents the stands that held the monitors in place. The bottom image is the approximate position of where the participants were approximately looking.

Figure 21 shows four different participants at four different column conditions – one column, three columns, five columns, and seven columns – all for the low pattern task. One can see as the viewport size increases that people naturally take advantage of the additional space. Although each participant had slightly different physical navigation patterns, looked at as a whole, the participants adapted to the larger displays and correspondingly increased their range of physical movement.

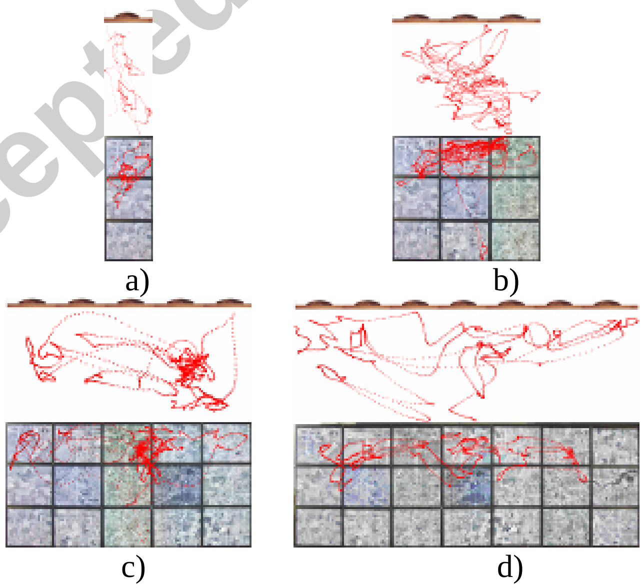

a) b)

c) d)

Figure 21. Four different participant data visualizations for four different column conditions. For all image pairs (a-d) the top image corresponds to an "overhead shot" while the bottom image corresponds to where the participant was looking at with an accuracy of about 88%. All four data visualizations are for the low pattern finding task.

In the experiment we gave participants a 3D wireless mouse specifically so that participants did not feel tethered to any particular location. However, for the insight finding task participants were given a mobile lecture stand to write their answers on. Figure 22 shows the physical navigation visualizations for the insight task for all the participants on seven columns (Figure 22.a) and for the pattern finding task for all the participants on seven columns (Figure 22.b). Clearly there was a smaller range of physical navigation in the insight task; we claim this is due to tethering.

As participants physically navigated less for the insight task they also virtually navigated more. The insight task was the only task where display width had no effect on virtual navigation.

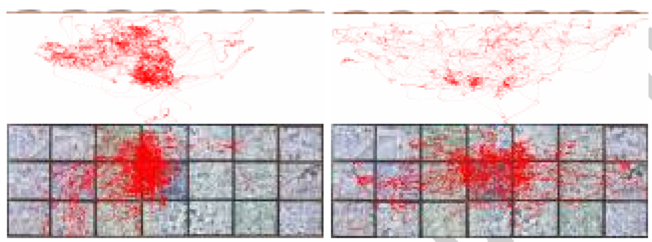

a) b)

Figure 22. Comparison of the insight finding task (a) to the pattern finding task (b).

5.3.4. Physical Bodily Movement Conclusions

The analysis of physical bodily movement shows a number of things. First, in general, the larger the display, the more X-axis range of position and total distance was found. In other words, the larger the viewport size the more participants used it.

For example, the low pattern task shows an increase in X-axis range of position of 428% increase in range of position from the one column condition the seven column condition. The total X-axis distance showed a similar result.

Second, tethering participants to the mobile table had a large effect on their physical navigation which likely affected their performance. Both head gaze and bodily movement were impaired.

5.4. Experiment Conclusions

The experiment showed a number of things. First, it validated with empirical data that there is a definite correlation of virtual navigation to performance time with a correlation coefficient of 0.69 for the number of zooms and with a correlation coefficient of 0.68 for the number of pans. Second, there was no such relationship with physical navigation as it did not significantly correlate with performance (correlation coefficient 0.46).

Second, it also showed that semantic zooming and where the semantic zooming thresholds are for the visualization also plays an important role in understanding both virtual and physical navigation. This seems to be especially important in understanding the relationship of performance curves: there appears to be a step-wise linear relationship of performance curves as a direct result of semantic thresholds.

Third, as viewport size influences how much virtual navigation is needed in conjunction with semantic zooming we found that viewport size also plays an important role in determining performance time. For example, we saw that performance time was reduced 10 times (1096%) in the navigation task by using larger displays and no virtual navigation.

Fourth, tethering, or being tied to one physical location plays an important role as to how much physical navigation can take place. Regardless of how large the viewport size is the more one is tethered to a location the less of the display one can use. This directly affects physical navigation and therefore affects performance. Specifically, we found no statistical significance in performance for the insight task; likewise we found no virtual navigation or physical navigation statistical significance either (except for the X-range of position).

Fifth, task types and task scales exhibit different virtual and physical navigation behaviors and thus exhibit different performance times. Different tasks and different task scales require different semantic zooming levels. As a result, larger viewport sizes will be more beneficial for some tasks and task scales than others.

Lastly, we conclusively found that participants prefer physical navigation over virtual navigation. Relating back to the extended space-scale diagram, we found that participants first physically navigated within the physical subset of the space-scale diagram before moving the viewport. In other words, participants repeatedly physically navigated in the blue pyramid in Figure 23 before virtually moving the viewport in the entire space-scale diagram (depicted in red).

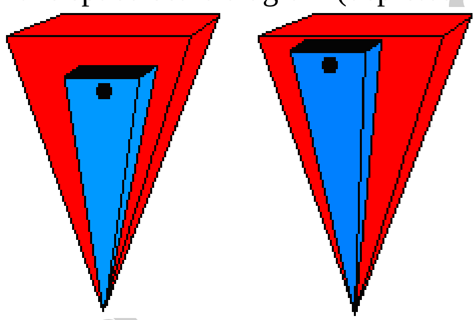

Figure 23. Two example space-scale diagrams showing the typical behavior of participants. Participants would try to only physically navigate in the blue area on the left. If they failed to accomplish the task they would then resort virtual navigation to change the viewport position then return to physical navigation.

6. Conclusion

This research has a number of impacts for the field of visual analytics. The following is a summary:

1. **Extended space-scale model with physical navigation**: The extended space-scale model realizes the theory of embodied interaction for visualization, by integrating the concepts of physical navigation and virtual navigation with large high-resolution displays. This model offers *theoretical* hypotheses that large displays should enable better user performance on visualization tasks, according to the theory of embodied resources. Extending the space-scale diagram to include physical navigation conceptualizes how physical navigation plays a role in perceiving data.
2. **Empirical evidence of the effects of large displays**: The *empirical* results demonstrate user behavior in the extended space-scale, provide evidence for the validity of the embodied resources theory, and quantify the actual effects of large, high-resolution displays.
    a. **Physical navigation preference over virtual navigation**: The space-scale diagram extension with physical navigation shows that with larger displays users have a choice to physically or virtually navigate. However, we found that for many instances 100% (32 out of

32) of the participants chose physical navigation over virtual navigation when given the choice.

b. **People use larger displays effectively**: In addition to showing a preference to physical navigation, we found that as the display sizes increased participants used the additional space provided. The larger displays generally resulted in faster user performance on visualization tasks by making better use of embodied resources. For example, this research found more than a 10 times improvement in performance in basic navigation tasks and a 2 to 3 times performance improvement for more difficult tasks such as search and pattern finding.

c. **There is a correlation between virtual navigation and performance**: We found that the number of zooms correlated with performance with a correlation coefficient of 0.69, and the number of pans correlated with performance with a correlation coefficient of 0.68, while physical distance traveled did not significantly correlate with performance (correlation coefficient 0.46). In other words, increased virtual navigation correlates with increased performance time while increased physical navigation does not necessarily correlate with increased performance time.

d. **Design of semantic zooming with visualizations is important**: We found that performance time was clearly affected by semantic zooming thresholds. For example, if insufficient details were presented then zooming in was required and if too much detail were presented then zooming out was required. Zooming requires computational power and time and disorients people (by not having an optical flow), both of which hurts performance time. Larger displays were able to mitigate some of these problems by broadening the view beyond initial semantic thresholds.

In conclusion, this paper extends the space-scale diagram to include the concept of physical navigation of large datasets with large, high-resolution displays, and demonstrates that people prefer physical navigation over virtual navigation and that they perform better when they do so. These results provide a concrete grounding for the theory of embodied interaction and its impacts on visual analytics. By better utilizing embodied resources such as spatial memory, proprioception, and optical flow, people can more efficiently navigate large information spaces with less disorientation, thus enhancing performance by alleviating the cognitive resources to focus on the analytic task at hand.

7. Future Work

During the course of this study we have presumed a constant pixel density of displays. However, how would people better benefit from higher or lower density displays? Would a higher pixel density display afford better performance results?

There are also a number of other factors that this study did not address. First, how do bezels affect performance with large displays? Second, how could people with disabilities still take advantage of large displays? Third, with additional display space there are numerous opportunities for having multiple views of data. What types of techniques and paradigms would lead to the best performance?

In addition, there are a number of questions of how the extension of physical navigation with space scale would work with multiple viewports of the same data. Similarly, how do interactive techniques such a focus+context or overview plus detail change the space scale model?

Lastly, how would the results from this research differ with abstract visualizations? Would there be differences in physical and virtual navigation? Are large displays more effective for spatial visualizations than purely abstract visualizations?

8. Acknowledgments

This research is partially supported by the National Science Foundation grant #CNS-04-23611. This study was also supported and monitored by the Advanced Research and Development Activity (ARDA) and the National Geospatial-Intelligence Agency (NGA) under Contract Number HM1582-05-1-2001. The views, opinions, and findings contained in this report are those of the authors and should not be construed as an official US Department of Defense position, policy, or decision, unless so designated by other official documentation.

We would like to thank Paul Rajlich from NCSA for writing the base to the software application that we used for experimental purposes.

9. References

[1] Ashdown, M., Oka, K., and Sato, Y. (2005) Combining head tracking and mouse input for a GUI on multiple monitors. Extended abstracts of CHI ’05, pp. 1180–1191.

[2] Ball, R. and North, C. (2005) An Analysis of User Behavior on High-Resolution Tiled Displays. Proceedings of INTERACT ‘05, Rome Italy, pp. 350-364.

[3] Bishop, G. and Welch, G. (2000) Working in the office of ‘real soon’ now.’ IEEE CG&A, 20(4).

[4] Bowman, D., Datey, A., Ryu, Y., Farooq, U., and Vasnaik, O. Empirical comparison of human behavior and performance with different display devices for virtual environments. In *Proceedings of Human Factors and Ergonomics Society Annual Meeting*, 2002, pp. 2134–2138.

[5] Bowman, D., Kruijff, K., LaViola, J., Poupyrev, I. *3D User Interfaces: Theory and Practice*. Peason Education, Inc., 2005.

[6] Czerwinski, M., Smith, G., Regan, T., Meyers, B., Robertson, G., Starkweather, G. (2003) Toward characterizing the productivity benefits of very large displays. In Proc. of Interact ’03.

[7] Czerwinski, M. Tan, D., Robertson, G. Women take a wider view. In *Proc. of CHI ’02*, p. 195 – 201.

[8] Darken, R. and Sibert, J. Navigating large virtual spaces. International Journal of Human-Computer Interaction, vol. 8, no. 1, pp. 49–71, 1996.

[9] Douglas, D., Peucker, T. Algorithms for the reduction of the number of points required to represent a digitized line or its caricature. The Canadian Cartographer. 10(2), 1973, 112-122.

[10] Dourish, P. Where the action is: the foundations of embodied interaction. MIT Press; 2001.

[11] Enns, J.T. The Thinking Eye, The Seeing Brain: Explorations In Visual Cognition. Published by W.W. Norton & Company, 2004.

[12] Furnas, G., Bederson, B. Space-scale diagrams: Understanding multiscale interfaces. In *Proceedings of CHI ’95 Human Factors in Computing Systems*, p. 234 – 241.

[13] Khan, A., Fitzmaurice, G., Almeida, D., Burtnyk, N., Kurtenbach, G. A Remote Control Interface for Large Displays. In *UIST 2004*, 127-136.

[14] Mackinlay, J. and Heer, J. Wideband displays: Mitigating multiple monitor seams. In *CHI, 2004.*

[15] Mackinlay J., Robertson G., Card S., The information visualizer: A 3D user interface for information retrieval. In Advanced Visual Interfaces, AVI, 1992, pp. 173–179.

[16] Nickel, K., Stiefelhage, R. Pointing Gesture Recognition on 3D-Tracking of Face, Hands and Head Orientation. Proceedings of *the Fifth International Conference on Multimodal Interfaces*, 2003.

[17] Olsen, D. and Nielsen, T. Laser pointer interaction. In CHI ‘01, pp. 17 – 22, 2001.

[18] Perlin, K., Fox, D. Pad: An alternative approach to the computer interface. *Proc. of Siggraph ’93*, 57-64.

[19] Plaisant, C., Grosejean, J., Bederson, B. SpaceTree: Supporting exploration in large noted link tree, design evolution and empirical evaluation. In *Proceedings of InfoVis ’02*, 57-64.

[20] Ringel, M., Berg, H., Jin, Y., and et al. (2001) Barehands: Implement-free interaction with a wall-mounted display. CHI Extended Abstracts.

[21] Robertson, G., Czerwinksi, M., Baudisch, P., Meyers, B., Robbins, D., Smith, G., and Tan, D. The Large-Display User Experience. IEEE Computer Graphics and Applications. 25(4), 2005, pp. 44-51.

[22] Robertson, G., Czerwinski, M, Larson, K., Robbins, D., Thiel, D., van Dantzich, M. Data Mountain: Using Spatial Memory for Document Management. In *Proceedings of UIST ’98*. p. 153-162.

[23] Robertson, G., van Dantzich, M., Robbins, D., Czerwinski, M., Hinckley, K., Risden, K., Thiel, D., and Gorokhovsky, V. The task gallery: a 3d window manager. In *CHI 2000*, pp. 494-501.

[24] Rovine, M. and Weisman, G. Sketch-map variables as predictors of way-finding performance. Journal of Environment Psychology, vol. 9, 1989, pp. 217–232.

[25] Ruddle, R., Payne, S., Jones, D. Navigating buildings in ‘desktop’ virtual environments: Experimental investigations using extended navigational experience. Journal of Experimental Psychology: Applied, vol. 3, no. 2, pp. 143–159, 1997.

[26] Sabri, A., Ball, R., Bhatia, S., Fabian, A., and North, C. High-Resolution Gaming: Interfaces, Notifications and the User Experience. *Interacting with Computers Journal* Volume 19, issue 2.

[27] Saraiya, P., North, C., and Duca, K. An insight based methodology for evaluating bioinformatics visualization. *IEEE Transactions on Visualizations and Computer Graphics*, 11(4), July/August 2005.

[28] Shupp, L., Ball, R., Yost, B., Booker, J., and North, C. (2006) Evaluation of Viewport Size and Curvature of Large, High-Resolution Display. Graphics Interface (GI) 2006.

[29] Swaminathan, K and Sato, S. (1997) Interaction Design for Large Displays. Interactions, 4(1), pp. 15-24.

[30] Tan, D., Stefanucci, J., Proffitt, D., Pausch, R. The infocockpit: providing location and place to aid human memory. In Proc. of PUI 2001. pp. 1-4.

[31] TerraServer Blaster. http://brighton.ncsa.uiuc.edu/prajlich/wall/tsb.html

[32] Tversky, B. Distortions in cognitive maps. Geoforum, vol. 23, no. 2, pp. 131–138, 1992.

[33] Vinson, N. Design guidelines for landmarks to support navigation in virutal environment. In Proceedings of CHI ’99, Pittsburgh, PA, 1999, pp. 278–285.

[34] Vogel, D. and Balakrishnan, R. Distant freehand pointing and clicking on very large, high resolution displays. In UIST ‘05.

[35] Ware, Colin. Information Visualization: Perception for Design. Published by Morgan Kaufmann Publishers Inc., 2000.

[36] Wickens, C. and Hollands, J. Engineering Psychology and Human Performance, third edition. Prentice Hall, 2000.

[37] Yost, B. and North, C. The Perceptual Scalability of Visualization, *IEEE Transactions on Visualization and Computer Graphics*, vol.12, no.5, Sept.-Oct. 2006.

[38] Zhang, X., Convertino, G., Ganoe, C., Schafer, W., Yost, B. Space-scale animation: enhancing cross-scale understanding of multiscale structures in multiple view. *Coordinated and Multiple Views in Exploratory Visualization (CMV ’05)*, 109-120.