Move to Improve: Promoting Physical Navigation to Increase User Performance with Large Displays

Robert Ball, Chris North, and Doug A. Bowman
Department of Computer Science
Virginia Polytechnic Institute and State University
Blacksburg, VA 24060

ABSTRACT

In navigating large information spaces, previous work indicates potential advantages of physical navigation (moving eyes, head, body) over virtual navigation (zooming, panning, flying). However, there is also indication of users preferring or settling into the less efficient virtual navigation. We present a study that examines these issues in the context of large, high resolution displays. The study identifies specific relationships between display size, amount of physical and virtual navigation, and user task performance. Increased physical navigation on larger displays correlates with reduced virtual navigation and improved user performance. Analyzing the differences between this study and previous results helps to identify design factors that afford and promote the use of physical navigation in the user interface.

Author Keywords

large displays, physical navigation, virtual navigation, embodied interaction.

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

Navigating in large virtual information spaces such as virtual environments (VEs) or visualizations can be difficult for users. Virtual navigation techniques, such as using a joystick control or pan & zoom widgets, are often disorienting and confusing. In response, information visualization researchers have developed virtual navigation aids such focus+context techniques [20]. In VEs, researchers employ wayfinding aids, but also augment virtual navigation with physical navigation (e.g. [23]).

We define physical navigation as bodily movement, such as walking, crouching, head rotation, etc., for the purpose of controlling the virtual camera that produces views of the information space. We view physical navigation as a specific type of embodied interaction [8]. Embodied interaction promotes the better use of humans' physical embodied resources such as motor memory, peripheral vision, optical flow, focal attention, and spatial memory to enhance the experience, understanding, or performance of the user.

Physical navigation is used in VEs and visualization in conjunction with a variety of display technologies such as CAVEs, head-mounted displays, projectors, wall-sized displays (e.g. Figure 1), and even desktop displays. Each of these display technologies has its own benefits and affordances for physical navigation.

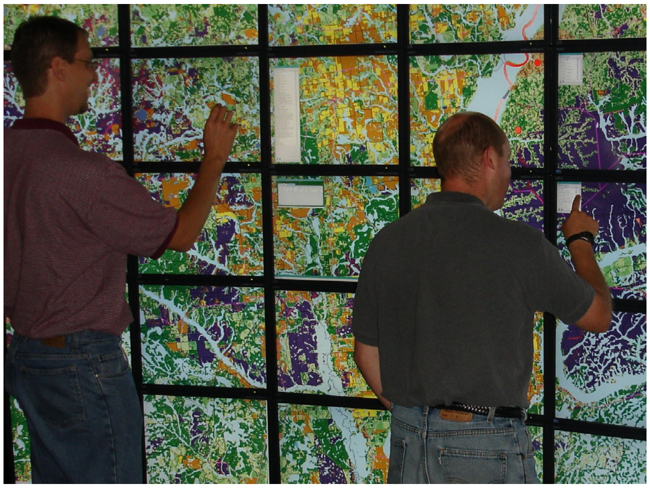

Figure 1. Example large, high-resolution display being used with physical navigation.

For example, in a CAVE (a VE display made up of multiple surrounding projection screens) head tracking is used to afford physical navigation, so that users can move around (within the confines of the physical CAVE) to adjust the 3D viewpoint. Most CAVEs, however, do not completely surround the user. Head-mounted displays also use head tracking, but also offer a 360-degree surrounding view, and do not take up as much real space as a CAVE. Large, high-resolution displays allow users to see large amounts of the information at amplified scales and degrees of detail. Users

CHI 2007, April 28–May 3, 2007, San Jose, California, USA.

can then step forward to see details (Figure 1) or step back to obtain an overview.

When navigating in information spaces with such displays, users must manage the tradeoff between physical navigation and virtual navigation (Table 1). For instance, where a user maintains a higher degree of spatial orientation with physical navigation, virtual navigation is often required to significantly change the viewpoint.

Table 1. Tradeoffs between physical and virtual navigation. The positive side of each tradeoff is denoted by italics.

	Physical navigation	**Virtual navigation**
Spatial understanding	*Higher*	Lower
Directness	*More direct*	Less direct
Navigation interface	*No explicit UI; body provides input*	Requires a dedicated navigation UI (button, widget, mode, etc.)
Generality	Not always sufficient	*Can always be used*
Fatigue	Higher	*Lower*
Input devices	Must be mobile	*Any device can be used*

Based on embodiment theory, we hypothesize that physical navigation should outperform virtual navigation and should be preferred by users. For example, physical navigation should help users better maintain spatial orientation. Indeed, some empirical evidence does indicate performance benefits for physical navigation in VEs, but other studies and anecdotal evidence show that virtual navigation is usually preferred by users (these results are described in detail in the section on Related Work section).

However, it also appears that although physical navigation may be more efficient in terms of performance, it is often not chosen by users in CAVEs and head-mounted displays. In fact, it appears that preference of physical navigation over virtual navigation is an exception rather than the norm.

We believe that large, high-resolution displays provide better affordances than other displays for encouraging physical navigation. This paper seeks to answer the following questions:

- Do users prefer physical navigation with large, high-resolution displays? Why?
- If so, does this result in improved user performance? Is physical navigation truly more beneficial than virtual navigation in terms of performance time?
- If physical navigation is more beneficial than virtual navigation, how can users be encouraged to physically navigate?

RELATED WORK

A review of the literature reveals that there have been a relatively large number of studies related to physical navigation, especially in the context of three-dimensional VEs.

In a VE, we must distinguish between two types of movements: rotations (turns) and translations. Either of these types can be physical or virtual, resulting in four possible combinations. Most desktop and single-screen VEs make use of virtual rotation and translation of the viewpoint (e.g. first-person shooters). With a tracked head-mounted display (HMD), users can perform physical turns, but most translations are done virtually due to limited tracking range. Locomotion devices such as a treadmill [11] allow (simulated) physical translations but require virtual turns. Finally, wide-area tracking systems [12] or specialized devices like the omni-directional treadmill [6] allow both physical turns and translation.

Displays like the CAVE [5] afford an interesting mix of both physical and virtual movements. Physical turns can be used, but virtual rotation is also necessary if the display does not completely surround the user; physical translation is also possible, but limited to a very small area. Informal observations of CAVE users indicate that they tend to prefer virtual rotation and translation (standing near the center of the CAVE, facing the front wall). Bowman et al. [3] showed that users of a CAVE with a missing back wall chose virtual rotations more often than HMD users for the same task (maze traversal), and that HMD users tended to outperform CAVE users.

The trend towards better performance with physical navigation has been confirmed by a number of researchers. The use of head tracking in an immersive information visualization was preferred by users and also appeared to improve comprehension and search [15]. Similarly, Pausch et al. [14] showed that users of a head-tracked HMD took less time to indicate that a target was not present in a visual search task as compared to users of the same display when the viewpoint was controlled by a handheld tracker. Chance et al. [4] demonstrated that when users physically turn and translate, they maintain spatial orientation better than when they virtually turn and translate. Bakker et al. [2] found that subjects could more accurately estimate the angle through which they turned if provided with vestibular feedback.

Although not as common, some research has also investigated physical navigation with 2D data displays. Ball et al. [1] investigated visual search performance on fairly large, high-resolution displays. Although users were seated, they observed some physical navigation (head turning, leaning, standing up) even though virtual navigation controls (pan & zoom) were also provided. In a follow-up study, Shupp et al. [19] also observed some physical navigation with larger tiled displays, and found that more physical movements occurred with the largest display size. However, users were reluctant to move too much because the tasks in this study required the use of a keyboard placed on a table in front of the display.

Other related work with large displays has shown general performance and accuracy improvements. For example, Tan et al. [21] show how women can improve their 3D navigation with larger displays. Czerwinski et al. [6] report on a study that shows general performance improvement with multi-tasking with multiple monitors. Sabri et al. [17] show how strategies and heuristics can change or be improved in spatial environments with large displays.

In summary, previous research has shown that most displays do not adequately afford physical navigation. In VEs, however, when users are required to turn or translate physically, performance improvements often result. In the following study, we wanted to investigate whether these performance improvements might also be measurable in 2D display settings. Since our previous work indicated that display size and tethering affected the amount of observed physical navigation, we used the largest display available to us and developed tasks in which the user could move freely in front of the display.

EXPERIMENTAL DESIGN

The goal of this experiment was to determine if large high-resolution displays afford physical navigation, to examine the resulting performance impacts, and to learn whether users preferred physical or virtual navigation in an un-tethered 2D information space.

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 2. Figure 2a 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 2b), and further zooming resulted in the display of square footage, number of bedrooms, and number of bathrooms, in addition to price (Figure 2c).

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 (approximately 131,072,000 pixels).

We used a modified version of the NCSA TerraServer Blaster [20], 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.

Figure 2. 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.

Display Used

The display used for the experiment was made up of twenty-four seventeen-inch LCD monitors in an 8×3 matrix (Figure 3). 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.

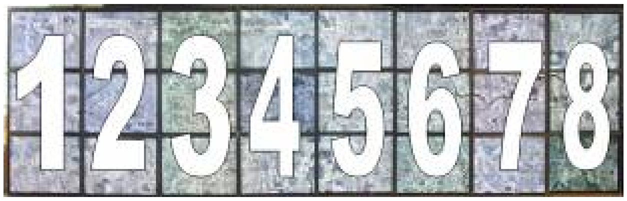

Figure 3. The display was separated into eight different columns. The total resolution of the display is 10240 X 3072 (31,457,280 pixels). The physical dimensions of the display were roughly 9 feet (2.7 m) by 3.5 feet (1 m).

In order to simplify the experiment, participants were tested on different widths of the display by column number (Figure 3). For example, in the four-column condition (15,728,640 pixels) only the first four columns would be used, and columns five through eight would be left unused. In the eight-column condition (31,457,280 pixels) all columns, one through eight, would be used.

Figure 4. a) Participant using the wireless mouse with the display. b) The hat used to track users' position.

Each task began with the overview/best-fit of the map always showing the same area of Houston. The aspect ratio of the base map was preserved so that each display width condition initially showed the same total overview area, but with different amounts of detail. Hence, the larger display width conditions with more pixels show more detail at startup. This offers the opportunity for more physical navigation, since users can examine more data without virtually navigating the display.

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 (see Figure 4a). 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 4b), 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.

Tasks

The participants performed four tasks: navigation to a target, search for a target, pattern finding for a group of targets, and open-ended insight for a group of targets. In order to measure only performance time and not accuracy for the first three tasks, participants were asked to keep working until the task was completed correctly. For instance, in the pattern task participants searched for the correct pattern until they reported it correctly.

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. This was done 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). The task was complete when the participant had spoken aloud the correct corresponding attribute of the house. This might require navigating (zooming) to the house to see the textual attributes.

The search tasks involved finding houses that had particular attributes (e.g. find a house priced between $100,000 and $110,000). There was not a unique correct answer per task as several houses fit each criterion. Approximately the same numbers of houses were potential correct answers for each search task.

Pattern finding tasks required participants to identify 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?" Each pattern finding task had a unique correct answer; participants did not have any difficulty arriving at this answer once the correct information was in view.

The open-ended insight task followed Saraiya's method of evaluating information visualizations based on insights gained [16]. For this task participants were given a rolling lecture stand on which to write insights (see Figure 4b). No performance time was recorded as all participants had ten minutes to write down as many insights as possible.

Prior to the first task, all participants were given at least five minutes to familiarize themselves with the wireless mouse and the different tasks. More time was given if it was felt more time was needed for a baseline.

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 in the university. The ages of the participants ranged from 24 to 39 with an average age of 28.

Design and Protocol

The independent variables for the experiment were viewport size (i.e. display width) and task type. The dependent variables were performance time (for the first three tasks), physical navigation (i.e. participant's 3D position), and virtual navigation (i.e. mouse interaction). For the insight task, the papers were graded for depth of insights by two graders that were familiar with the data.

The first two tasks, basic navigation and search, used a within subject design in which all 32 participants performed tasks on all eight display width conditions. We used a Latin Square design to determine the order in which participants used the display widths. The second two tasks, pattern finding and insight finding, used between-subject designs. Only the 1, 3, 5, and 7 column conditions were used for these tasks to increase statistical power.

Each of the first three tasks required a range of levels of detail, hence requiring a range of zooming navigation. As a result the navigation task was repeated twice and the search and pattern tasks were repeated three times.

EXPERIMENT RESULTS

This section reports the results of the experiment. We found no significant results based on the level of insight for the fourth task, so we focus on results for the first three tasks in this section.

Performance Time Analysis

In order to analyze performance results we ran a two-way ANOVA on performance times with display width as a continuous variable, and tasks as a discrete variable. We found main effects for both display width (F(1,1324)=20.56, p<0.01) and task type (F(2,1324)=77.05, p<0.01).

Table 2. Statistical performance time results.

Task	Main effect of display width
navigation	(F(1,508) = 118.9, p<0.01)
search	(F(1, 762) = 38.18, p<0.01)
pattern finding	(F(1, 90) = 3.53, p=0.06)

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 the others we performed individual ANOVAs for each of the tasks (Table 2). There was a significant effect of display width for the navigation and search tasks, but only a near-significant trend for the pattern finding task.

Figure 5 shows mean performance results for of the navigation and search tasks. For the navigation and search tasks, the smaller displays (one and two columns) performed significantly worse than the larger displays (seven and eight columns).

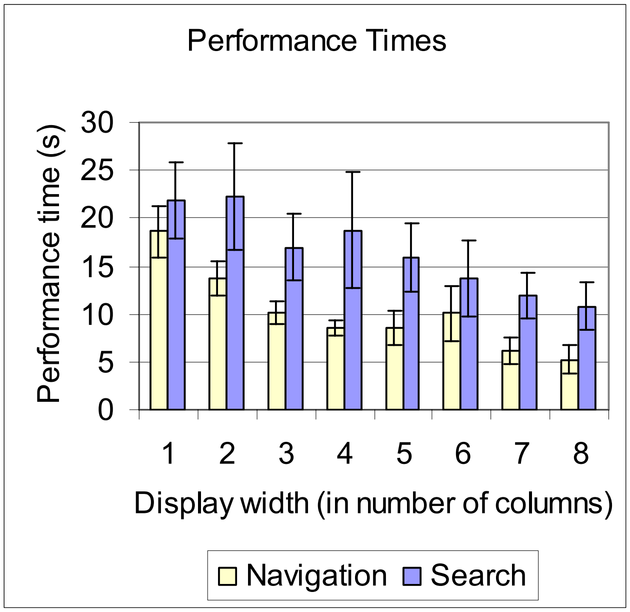

Figure 5. Performance averages for the navigation and search tasks on different width displays.

In summary, larger viewport sizes caused faster performance. For example, on the navigation task, performance time was reduced by a factor of 3.5, from 18.6 seconds on the one column condition to 5.2 seconds on the eight column condition. In the search task, performance was reduced by a factor of 2, from 21.9 seconds on the one column condition to 10.8 seconds on the eight column condition.

Virtual Navigation Analysis

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 zoom level to which participants had to navigate to see the necessary detail (e.g. price of the houses). Second, participants would sometimes pan to see different geographical areas at a particular zoom level.

We performed two-way ANOVAs on display width and task type for both the number of zooms and the number of pans. For the number of zooms, we found a main effect of task type (F(3,1400)=416.2, p<0.01), a main effect of display width (F(1,1400)=34.8, p<0.01), and a near-significant interaction of task type and display width (F(3,1400)=2.4, p=0.06).

The second analysis was the number of pans performed. The reader should note that the number of pans is only mouse movement that actively moved the viewport in space. It is not inactive mouse movement that was used to reposition the cursor without moving the viewport. The ANOVA showed a main effect of task type (F(3,1400)=301.3, p<0.01), a main effect of display width (F(1,1400)=63.86, p<0.01), and a significant interaction of task type and display width (F(3,1400)=17.22, p<0.01).

Table 3. Statistical results of the virtual navigation data for the different tasks.

Tasks - Metrics	**Main effect of display width**
navigation - zooms	(F(1,508)= 144.6, p<0.01)
navigation - panning	not significant
search - zooms	(F(1,762) =114.1, p<0.01)
search - panning	(F(1,762) = 26.7, p<0.01)
pattern - zooms	not significant
pattern - panning	(F(1,90) = 7.8, p<0.01)

As with the time data, we performed separate ANOVAs for each task (Table 3). Figure 6 and Figure 7 show the corresponding graphs.

Figure 6 shows that, in general, the number of zooms decreases as the display size increases, for all three tasks. This trend of number of zooms closely matches that of performance time.

We found a significant difference in the number of zooms based on display width for the navigation and search tasks. Display width did not have a significant effect on the number of zooms for the pattern task due to a high variance.

Another thing that separates the pattern task as different from the other tasks is that participants were observed to virtually zoom out to better see the overall pattern. In the other tasks, participants were only observed to virtually zoom in. However, the seven-column condition started out showing more details than were needed for an overall pattern task. As that particular 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 more easily see *only* the geospatial pattern.

Figure 7 shows the corresponding amount of panning for the different tasks and display widths. Again, the number of pans is seen to generally decrease as display size increases. Display width had a significant effect on the number of pans for the search and pattern finding tasks, but not the navigation task, as panning was not typically necessary for the navigation task.

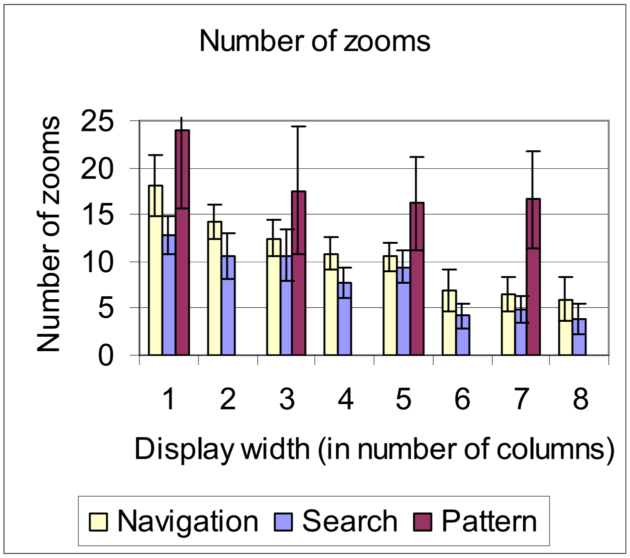

Figure 6. Average number of zooms (virtual navigation) for each task and display width.

Interestingly, for a number of tasks at certain scales there was not any zooming or panning performed. There were four different task conditions where all 32 of the participants chose not to perform *any* virtual navigation. For example, for one of the navigation tasks in the eight-column condition all the participants chose to use only physical navigation to complete their task. Zero virtual navigation also occurred for one of the search tasks in the eight-column condition, and for one of the pattern finding tasks in the three- and five-column conditions.

When virtual navigation is not required users have a choice to either virtually navigate or physically navigate. We found that when there is a choice, physical navigation is preferred over virtual navigation. For instance, on another search task, 90% (29 out of 32) of the participants did not zoom and 100% of the participants did not pan in the eight-column condition. This pattern continued for all such choices.

Physical Navigation Analysis

We analyzed the physical navigation of participants based on head movements relative to X, Y, and Z axes in the area in front of the display where the users physically navigate. Figure 8 shows an illustration of how the three axes map to the large display. The X-axis runs parallel to the display and corresponds to horizontal movements; the Y-axis runs perpendicular to the display and corresponds to moving closer or farther from the display; the Z-axis is vertical and corresponds to crouching or standing up straight. In effect, X- and Z-axis movement is physical panning while Y-axis movement is physical zooming.

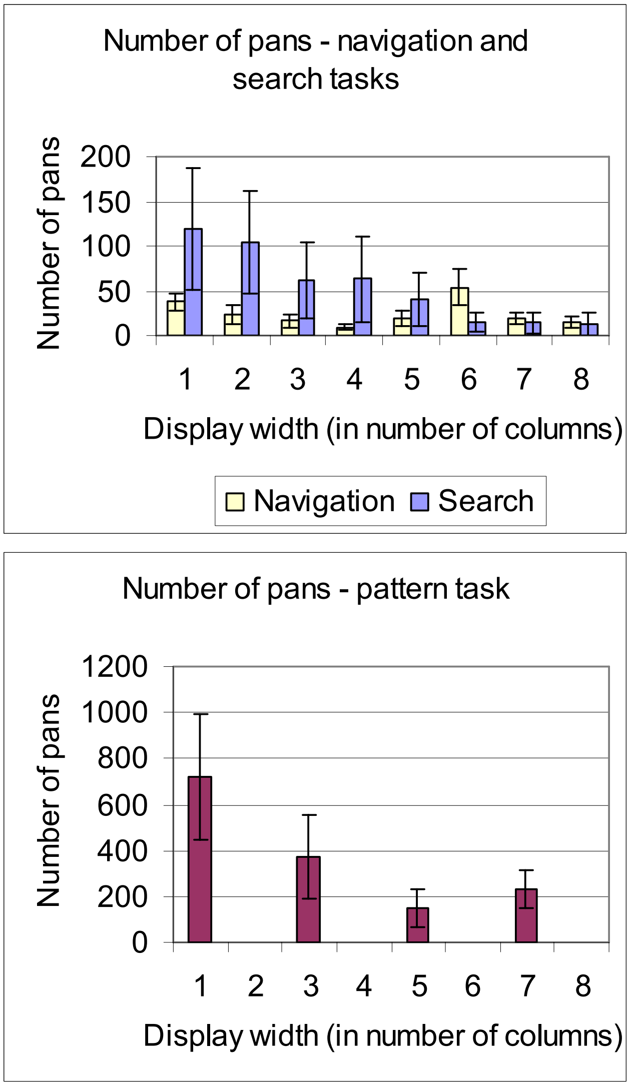

Figure 7. Average number of panning actions (virtual navigation) for the navigation and search tasks (top) and pattern task (bottom).

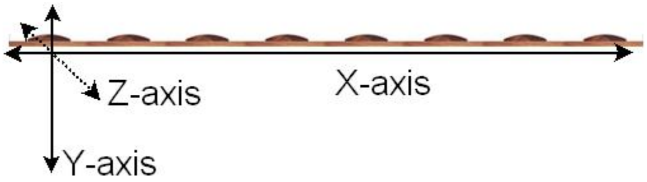

Figure 8. Illustration of the X, Y, and Z axes relative to the display (overhead view).

Physical movement distance was calculated by using a modified Douglas-Peucker algorithm [3]. The algorithm helps to guarantee that what we were analyzing was actual movement from one physical location to another and not jitter from the tracking system.

Table 4. Statistical analysis of the total X distance moved for the different tasks.

Task	Main effect of display width
navigation	not significant
search	(F(1,762) = 4.52, p=0.03)
pattern finding	(F(1,84) = 16.62, p<0.01)

We performed a two-way ANOVA on display width and task type for the total X distance. Total X distance takes into account moving back and forth over the same positions. We found a main effect of task type (F(3,1400)=75.1, p<0.01), a main effect of display width (F(1,1400)=24.1, p<0.01), and a significant interaction of task type and display width (F(3,1400)=4.0, p<0.01). Separate ANOVAs for each task resulted in main effects of display width for only the search and pattern finding tasks (Table 4). The non-significance for the navigation task can be explained by the low need to move in the X direction, similar to the virtual navigation result.

Figure 9 and Figure 10 show the average total distance covered in the X direction for the search and pattern finding tasks. There is a clear preference for more physical navigation in the wider display conditions.

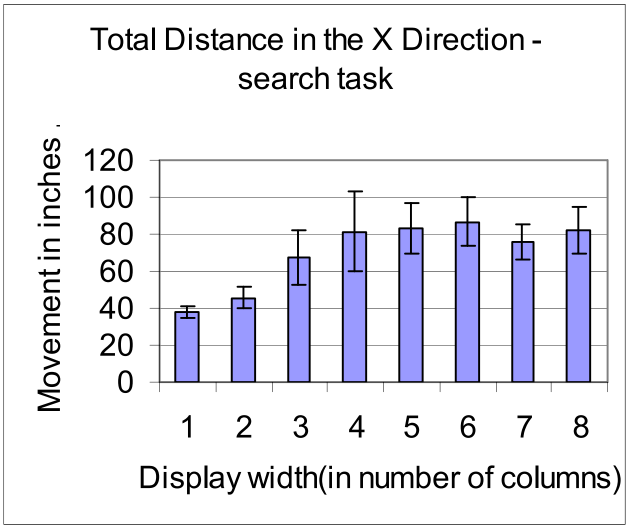

Figure 9. Average total X distance of participants in the search task.

There is also a difference between Figure 9 and Figure 10. In Figure 9 there appear to be diminishing returns or leveling off of physical navigation, while in Figure 10 there appears to be more of a linear increase in physical navigation.

However, the search task indicates that participants' physical navigation did not always increase as display size increased. As Figure 5 shows, performance time for the search task continued to improve as display size increased even though the amount of physical navigation did not increase. Particularly, participants were observed to make better strategic decisions based on being able to see more overview and details at once.

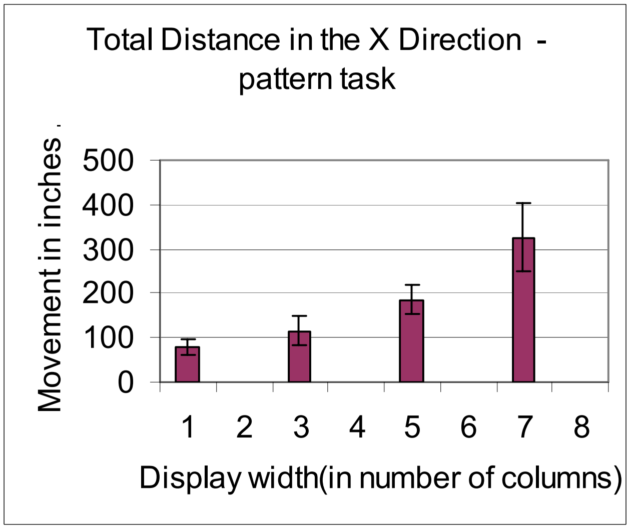

Figure 10. Average total X distance of participants in the pattern task.

For example, on the one column condition of the search task participants were generally seen to randomly select areas of Houston to look at in detail. They would then search the area at a detailed zoom level, and then if they failed in finding a house that met the search criteria in that area they would randomly search another area of Houston until they succeeded.

However, on the larger display widths participants were able to see general overview and detail trends in the data at the beginning of the task. As more information was visually presented participants were able to navigate less to complete the task. They were able to visually see more information and were generally observed to make more intelligent navigation decisions.

For example, instead of randomly navigating to an area to look at in more detail, participants would visually scan the display then narrow their focus on an area that appeared to have more promise. Then, participants would navigate (e.g. walk) to that part of the display for further detail. For more information on improved strategies and heuristics with large displays see [17].

Visual Representations

Figure 11 is an example of physical movement for the pattern finding task in different display width conditions for different participants. The top image corresponds to an overhead view of the participant. It shows where in the space participants' head locations were at different times.

The bottom image shows the head orientation of the participants projected onto the display. In other words, what is shown is the approximate gaze position – where the participants were looking on the display. Head gaze can predict eye gaze with an 87-89% degree of accuracy [13].

One can see in Figure 11 that 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.

In the experiment we gave participants a 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 12 shows the physical navigation visualizations for the insight task for all the participants on seven columns (Figure 12a) and for the pattern finding task for all the participants on seven columns (Figure 12b). Clearly there participants were more physically constrained in the insight task; we claim this is due to tethering.

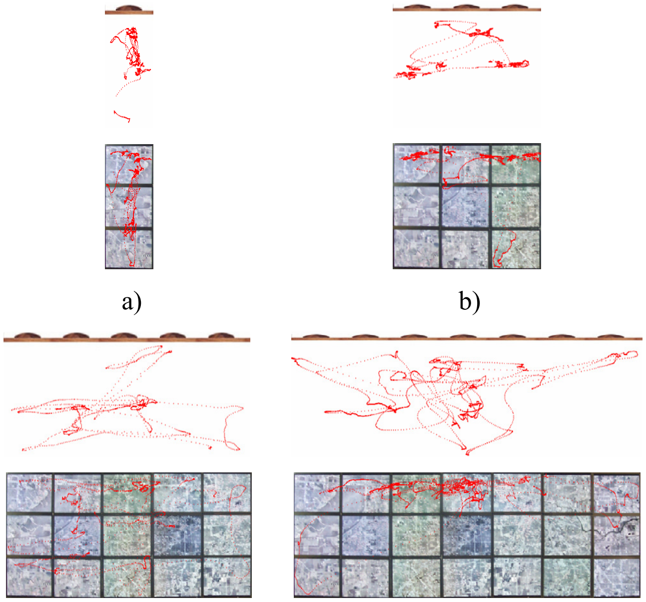

a) b)

c) d)

Figure 11. Visualizations of four different participants' movement for four different display-width conditions. For all image pairs (a-d) the top image corresponds to an overhead view while the bottom image corresponds to a projection of head orientation onto the display (approximating gaze direction). All four data visualizations are for a pattern finding task.

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.

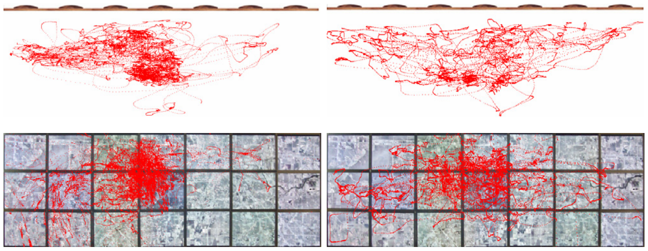

a) b)

Figure 12. Comparison of the insight finding task (a) to the pattern finding task (b) for all participants showing the effects of tethering on the insight task.

Experiment Conclusions

There are a number of important findings in this analysis.

First, 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.

Second, we found that as displays sizes increased, virtual navigation decreased, and performance time also decreased. For example, with the number of zooms recorded for the search task, the number of zooms decreased 3.2 times from the one-column condition to the eight-column condition. The corresponding performance decreased 3.5 times from the one-column condition to the eight-column condition.

There were two exceptions to the rule of decreased virtual navigation with increased display width. The first exception was that people zoomed out to see fewer details for an 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 other exception is with the insight task. Since bodily movement was impaired, tethering participants to the table had a large negative effect on their physical navigation, which affected their amount of virtual navigation and likely affected their resulting performance.

Third, our experiment showed that, in general, the larger the display, the more physical navigation. Combined with the decreased performance time on large displays, we see a strong suggestion that physical navigation was also more efficient. However, larger displays did not always lead to increased physical navigation (as seen in the search task), as participants were observed to use better strategies and heuristics with the larger displays as they could see more overview and details at once. In essence, the larger view helped to guide physical navigation and hence less virtual navigation as well.

Fourth, physical navigation was preferred over virtual navigation. When possible, participants preferred to physically navigate to understand their data. We observed that participants first physically navigated as much as possible before virtually navigating. After virtually navigating they would then repeat the behavior of attempting to complete the task with physical navigation before relying on virtual navigation.

Finally, it appears that larger displays are a critical factor in producing these effects. For example, we show that larger displays promote more physical navigation with several instances where 100% of the participants chose only to physically navigate.

ENCOURAGING PHYSICAL NAVIGATION

This study suggests significant benefits of physical navigation over virtual navigation, similar to earlier results. In contrast to previous work, however, it also demonstrates a clear preference by users to take advantage of these benefits by choosing physical navigation over virtual navigation when using large displays. Why? What are the key differences between this study and previous studies that caused this preference to occur? Can we identify the important factors to better promote physical navigation in the design of future systems, and reduce dependency on virtual navigation?

Several key factors emerge:

1. Non-tethered users: The use of the wireless handheld input device in this study gave users more freedom to physically navigate. On the other hand, with the use of the keyboard in the insight task and in a previous study [19], less physical navigation was observed. Other forms of tethering, such as wired HMDs, may have similar effects.
2. Large physical space for range of motion: There was a great deal of open space in front of our display. In contrast, enclosing CAVE walls and limited range trackers can constrain users' movement.
3. Increased display resolution: The large, high-resolution displays afforded users the ability to scan a large amount of information at multiple levels of scale through physical navigation. Smaller display conditions do not offer such advantages. The low resolution of CAVEs causes information to become less clear as the user physically translates nearer to the CAVE wall. HMDs provide a constant resolution regardless of physical navigation. The near-infinite resolution of the real world is a goal.
4. Body and physical world are visible: In our setup, users could see both themselves and the physical environment. A common problem in HMDs is that users lose track of where they are in the physical world. Fearing that they will crash into a physical wall or trip over a wire, they avoid physical movements.
5. Physical and virtual match-up: In 3D virtual environments, sometimes the goal is to immerse the user entirely in a virtual world and completely hide the physical world. Thus, a disconnect arises when users must physically navigate in the real world in order to move in the virtual world. The real world and virtual world do not match and physical navigation becomes an overloaded operator. Physical navigation would have to be virtualized to match the virtual world, and this is difficult to fully achieve. A successful example is a car or flight simulator that uses an actual cockpit, where the 'display' becomes physical.

Together, these factors suggest that the display is a physical real-world object that users directly interact with. In this study, users perceived the display as an object in their interaction space and that they could physically navigate with respect to it. The display became like a large physical map hanging on a wall, but also provided dynamic virtual features. Perhaps this is evidence for embodied interaction theory, in which physical resources are fully exploited. If these factors are considered in the designs of large information spaces, it is likely to encourage physical navigation over virtual navigation, and improve performance.

CONCLUSION

This work offers several important results. The study identifies definite relationships between display size, user performance time, amount of physical navigation, and amount of virtual navigation. For the spatial visualization tasks we explored, larger displays lead to more physical navigation, which reduces the need for virtual navigation, which offers improved user performance.

Is physical navigation beneficial? Yes, physical navigation is indeed an efficient and valuable interaction technique that reduces dependency on less-efficient virtual navigation.

Is physical navigation preferred by users? Yes, we found that in the right conditions, physical navigation was also preferred over virtual navigation by users, leading to improved performance times. In situations where either physical or virtual zoom-in navigation could be used to fully complete the task, physical navigation was chosen 100% of the time.

Why was physical navigation preferred? Can physical navigation be promoted in other system designs? By examining the broader context of this study within the literature, several key design factors are identified that make a difference in affording and promoting physical navigation. These factors can be broadly applied to improve acceptance and user task performance.

This work has been conducted solely on spatial visualizations. As a result, would the results extrapolate to non-spatial, more abstract visualizations? In addition, what are the long term affects of physical navigation with large displays? How do the results extrapolate w multiple views?

ACKNOWLEDGEMENTS

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.

REFERENCES

1. Ball, R. and North, C. “Effects of Tiled High-Resolution Display on Basic Visualization and Navigation Tasks,” In *Extended Abstracts CHI’05*, 1196-1199.
2. Bakker, N. Werkhoven, P., and Passenier, P. “Aiding Orientation Performance in Virtual Environments with Proprioceptive Feedback,” In proceedings of *IEEE Virtual Reality Annual International Symposium*, 1998, pp. 28-35.
3. 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.
4. Chance, S., Gaunet, F. Beall, A., and Loomis, J. “Locomotion Mode Affects the Updating of Objects Encountered During Travel,” In *Presence: Teleoperators and Virtual Environments*, vol. 7, pp. 168-178, 1998.
5. Cruz-Neira, C., Sandin, D., DeFanti, T. “Surround-screen projection-based virtual reality: The design and implementation of the cave.” In proceedings of *ACM SIGGRAPH, 1993*.
6. Czerwinski, M., Smith, G., Regan, T., Meyers, B., Robertson, G., and Starkweather, G. “Toward characterizing the productivity benefits of very large displays,” In proceedings of *Interact 2003*, 2003.
7. Darken, R., Cockayne, W., and Carmein, D. “The Omni-directional Treadmill: A Locomotion Device for Virtual Worlds,” In proceedings of *UIST ‘97*, 1997, pp. 213-221.
8. Dourish, P. (2004) Where the Action Is: The Foundations of Embodied Interaction. MIT Press.
9. Douglas, D., Peucker, T. “Algorithms for the reduction of the number of points required to represent a digitized line or its caricature,” In *The Canadian Cartographer,* 10(2), 1973, pp. 112-122.
10. Mackinlay, J. and Heer, J. “Wideband displays: Mitigating multiple monitor seams.” In proceedings of *CHI, 2004.*
11. Hollerbach, J. “Locomotion Interfaces,” in *Handbook of Virtual Environments*, K. Stanney, Ed.: Lawrence Erlbaum, 2002, pp. 239-254.
12. Interrante, V., Anderson, L., and Ries, B., “Distance Perception in Immersive Virtual Environments, Revisited,” In proceedings of *IEEE Virtual Reality*, 2006, pp. 3-10.
13. Nickel, K., Stiefelhage, R. “Pointing Gesture Recognition on 3D-Tracking of Face, Hands and Head Orientation.” In proceedings of *the Fifth International Conference on Multimodal Interfaces*, 2003.
14. Pausch, R., Proffitt, D., and Williams, G. “Quantifying Immersion in Virtual Reality,” In proceedings of *ACM SIGGRAPH*, 1997, pp. 13-18.
15. Raja, D., Lucas, J., Bowman, D., and North, C. “Exploring the Benefits of Immersion in Abstract Information Visualization,” In proceedings of *Immersive Projection Technology Workshop*, 2004.
16. Robertson, G., Czerwinski, M., and van Dantzich, M. “Immersion in desktop virtual reality,” In proceedings of *UIST ’97*, 11-19.
17. Sabri, A., Ball, R., Bhatia, S., Fabian, A., and North, C. “High-Resolution Gaming: Interfaces, Notifications and the User Experience,” In *Interacting with Computers Journal*, 19(2), 2007.
18. Saraiya, P., North, C., and Duca, K. “An insight based methodology for evaluating bioinformatics visualization.” In *IEEE Transactions on Visualizations and Computer Graphics*, 11(4), July/August 2005.
19. Shupp, L., Ball, R., Yost, B., Booker, J., and North, C. “Evaluation of Viewport Size and Curvature of Large, High-Resolution Display,” In proceedings of *Graphics Interface (GI) 2006*, 123-130.
20. Spence, R. Information Visualization. Published by Addison-Wesley, 2001.
21. Tan, D.S., Czerwinski, M., and Robertson, G. “Women go with the (optical) flow,” In proceedings of *ACM SIGCHI ’03*, 209-215.
22. TerraServer Blaster. http://brighton.ncsa.uiuc.edu/prajlich/wall/tsb.html
23. Usoh, M., Arthur, K., Whitton, M., Bastos, R., Steed, A., Slater, M., and Brooks, F. “Walking > Walking-in-Place > Flying, in Virtual Environments,” In proceedings of *ACM SIGGRAPH ‘99*, 359-364.