

Exploiting the Cognitive and Social Benefits of Physically Large Displays

Thesis Proposal
March 2003

Desney S. Tan
School of Computer Science
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213
desney@cs.cmu.edu

Thesis Committee:
Randy Pausch (Chair)
Jessica Hodgins
Scott Hudson
Mary Czerwinski, Microsoft Research

Abstract

With recent advances in technology, large wall-sized displays are becoming prevalent. Although researchers have articulated qualitative benefits of large displays, little has been done to systematically quantify and exploit these benefits. Additionally, since most researchers have assumed that larger displays fill a greater percentage of the viewer's visual field, physical size is often confused with visual angle, or field of view. In my thesis, I will first isolate the effects that physical size has on the way users react to information. I will then show how we can exploit these reactions to construct display systems that make users more productive than they were on traditional systems.

In this proposal, I describe work I have done to compare the performance of users working on large wall-sized displays to that of users working on smaller displays, viewed at identical visual angles. In the first set of studies, I show that although there was no significant difference in performance on a reading comprehension task, users performed 26% better on a spatial orientation task done on a large display. Results suggest that these improvements can be attributed to the large display affording a greater sense of presence and biasing users into treating the spatial task as an egocentric rather than an exocentric rotation. In a second study, I apply an implicit memory paradigm to measure whether or not someone has read a particular passage of text. I show that people are more likely to read private text presented on a larger display, even when visual angles and legibility were constant. In a third study, I explore the pragmatics of using multiple display systems that combine small and large displays. I show that physical discontinuities introduced by bezels or by differences in depth between small and large displays had no effect on divided attention tasks. However, I also show that the separation in depth created by working on a small display that was nearby and a large display that was further away hurt task performance when information was placed at wider visual angles.

I discuss my general research approach as well as provide the detailed research agenda that I will pursue in proving my thesis, that:

Information elicits fundamentally different cognitive and social reactions when presented on large wall-sized displays as compared to smaller displays, even at identical visual angles. These reactions can be quantified and understood in controlled experiments and can be exploited to design display systems that make users more productive than they were on traditional systems.




Figure 1: Various stages of human perception - the distal object reflects light and forms an image on the retina, which is then interpreted to form the final perceived image.

INTRODUCTION

Most of us live in a mixed reality. Our daily activities are divided between the physical environment, or the real world, and a digital reality, or a virtual world. However, as observed by Ishii & Ullmer (1997), the two worlds remain largely disjoint and there exists “a great divide between the worlds of bits and atoms.” In their work, they identify input devices as bridges that serve to connect the two worlds. They focus on understanding how physical objects and architectural surfaces can be used to control digital objects in the virtual world. In my work, I assert that computer displays also serve as bridges that connect the two worlds. Displays exist in the real world and allow us to peer into the virtual world.

Traditional user interface design has focused largely on the information content that lies within the virtual world. This research is aimed at understanding the symbolic representation of information that most effectively communicates abstract ideas. Little effort has been spent on understanding the design of the physical computer and its associated display devices (Buxton 2001). As such, design principles have been applied uniformly across a variety of display devices that offer different cognitive and social affordances. In this thesis, I instead focus my attention on user reactions to physical properties of information. Specifically, I hypothesize that physical size, independent of visual angle, shapes the way users perceive and interact with information, and hence greatly affects their performance on tasks.

Understanding the role of physical display size is significant because of the emerging trend in the workplace towards multiple display systems, which has the potential to provide abundant display space distributed throughout the environment. Such workplaces typically include several types of displays, each with characteristics that may make it more or less suitable for certain tasks. Within these workplaces, large wall-sized displays are becoming prevalent. Researchers have previously realized that “when a display exceeds a certain size, it becomes qualitatively different” (Swaminathan & Sato 1997). Although these researchers have articulated qualitative benefits of working on physically large displays, little work has been done to systematically quantify or exploit these benefits.

In studying user reactions to physical properties of information, there are several stages of visual perception that we must recognize (see Figure 1). First, we assume that the physical world exists and that its existence is independent of the observer. Objects in the physical world are generally referred to as the distal or physical stimuli. Hence, physical size describes the actual size of an object in the environment. The physical world is governed by a well-defined set of physical laws. For example, objects reflect light in a predictable, though complex, manner. This world can also be described fairly completely with a homogeneous Euclidean geometry.

The observer, present in the physical world, views the physical stimuli when light reflected from objects in the environment stimulate receptors on the retinal surface of the eye. These impinging patterns

of light are referred to as the proximal or retinal stimuli. Retinal size, then, describes the size of the image cast on the retina. It is also sometimes measured as the visual angle, or field of view, of an object or scene. The retinal size is dependant both on the physical size of the object as well as the distance from which the object is being viewed.

Perception refers to the conscious experience of the physical object or scene. Physical objects and scenes can be observed directly. Retinal stimulus patterns can be observed by projecting light from physical stimuli onto a projection plane that represents the retinal surface. Perception, however, is a process internal to the human mind and cannot be directly observed. Instead, we observe behavioral responses to the stimuli in order to indicate that certain perceptions are occurring. In my work, I am interested in understanding and exploiting these behavioral responses by controlling the stimuli that shape visual perception.

The perceived image results from a complex relationship between physical and retinal cues, as well as cues internal to the observer, such as context and semantic knowledge. Unfortunately, most researchers have worked with the implicit assumption that the perceived image is dependant only upon the size and content of the retinal image. They have studied in detail how changing the retinal size by varying the visual angle or field of view affects perception and task performance. They have also developed many techniques to present information in a manner that is most easily processed by the human mind.

In this thesis, I recognize that perceptual space is shaped by more than just the retinal image. In fact, increasing the size of the display surface can fundamentally change the user perception and interaction with information. Even though a given image may have the same theoretical information content on a small or large display viewed at the same visual angle, it may elicit different cognitive and social reactions, causing responses that lead to different levels of productivity on different displays. In particular, I will explore how varying the physical size of displays, while keeping factors such as visual angle and information content constant, affects perception and task performance.

THESIS STATEMENT

Information elicits fundamentally different cognitive and social reactions when presented on large wall-sized displays as compared to smaller displays, even at identical visual angles. These reactions can be quantified and understood in controlled experiments and can be exploited to design display systems that make users more productive than they were on traditional systems.

GENERAL RESEARCH APPROACH

By combining anecdotal evidence, my initial findings, and established theoretical work in psychology, human-computer interaction, and computer science, I have identified several areas in which large displays are likely to have an impact on user performance. In my work I will focus on those areas that promise to have the greatest impact on human-computer interaction and display system design.

For each of these areas, I will design and conduct controlled experiments that allow me to isolate and understand effects more completely. I will then derive design principles and build real-world interfaces and systems that utilize knowledge gained about the various effects of physical display size. As a final step, I will evaluate these interfaces and systems with tasks for which they are designed.

One of the areas I have started to explore is spatial cognition. I discuss preliminary results that show users perform 26% better on a task requiring mental rotation when working on a large display, even with constant visual angles. Another promising area is that of multiple user or social environments, specifically on visual privacy. I have applied an implicit memory paradigm to measure whether or not a user has read a certain passage of text. Preliminary results show that this paradigm is useful and that users are indeed more likely to read private information presented on a physically large display than on a smaller one, again with constant visual angles. A third area is attention and memory. I show that wider visual angles hurt task performance, but only when information is further separated in depth, for example when it is

distributed across a small and large display in a multiple display environment. In my thesis work, I will continue to explore how large displays affect performance in each of these areas.

In addition to this work, I will conduct several experiments to isolate and understand the individual factors that make a large display fundamentally different from smaller ones. I will also explore subjective responses to physically large displays. This understanding will allow me to better formulate design principles for the construction of multiple and large display systems.

RELATED WORK

Physical size is an important cue to sensory and judgment processes in humans. For example, both infants and adults have been shown to exhibit preferences for larger objects in presentations (Silvera, Josephs, & Giesler 2002). Other research suggests that physical height in males is positively correlated with physical attractiveness, income, and occupational status (Jackson 1992). In fact, the tallest candidate has won all but 4 of the past 23 US presidential elections (Newsweek 1992). Wearing height and weight enhancing clothing and apparatus, including headdresses and shoulder pads, are instances of trying to manipulate this bias with size-inspired threat and power displays (Campbell 1976). Additionally, Josephs, Giesler, & Silvera (1994) found that a wide variety of judgments were strongly influenced by non-diagnostic physical size information through the application of a 'bigger is better' rule. They found that artificially increasing the physical size of a pile by attaching each sheet of actual work to an empty cardboard box drastically increased estimates of productivity and progress. It is thus not a new idea that physical size of information greatly affects the way humans respond to it.

This concern with human reaction to physical size has recently moved into the computing world, as new developments allow us not only to combine multiple displays within the workplace, but also to significantly vary the size of each display. In constructing complex workplaces, researchers have often pursued the use of large displays for collaborative tasks (Chou, Gruteser, Lai, Levas, McFaddin, Pinhanez, et al. 2001; Elrod, Burce, Gold, Goldberg, Halasz, Janssen, et al. 1992; Raskar, Welch, Cutts, Lake, Stesin, & Fuchs 1998; Streit, Geißler, Holmer, Konomi, Muller-Timfelde, Reischl et al. 1999; Tani, Masato, Kimiya, Koichiro, & Futakawa 1994). Large displays in these settings are easy for all users to see and interact with (Guimbretièrre 2002), providing a conduit for social interaction. Some researchers have begun to document performance increases for groups working on large displays (Dudfield, Macklin, Fearnley, Simpson, & Hall 2001). In projects such as Bishop & Welch's (2000) Office of "Real Soon Now" or Cardelli & Beckman's ClearDesk Projector, researchers have gone so far as to completely replace desktop displays with large-screen projection displays in order to study their effects.

While much work has focused on collaboration, less has been done to objectively measure and design for individual gains on large displays. In a series of three studies conducted to strengthen marketing claims for larger desktop displays, Simmons (2001) showed that users performed better on productivity tasks using large 21" monitors as compared to smaller 15", 17", and 19" ones. She also showed that users significantly preferred the 19" displays over the others due to tradeoffs on physical desk space imposed by current technologies. Although this work showed benefits of larger displays, Simmons explored only a small range of display sizes, each viewed at different visual angles and with different resolution. Also, she described the presence of effects without attempting to explain them.

Several researchers have explored the use of large displays as a means to provide contextual information to the individual. For example, Baudisch, Good, Belloti, & Schraedley (2002) provide a large low-resolution overview of the working context around a smaller high-resolution focal screen. MacIntyre, Mynatt, Volda, Hansen, Tullio, & Corso (2001) assist users in managing multiple working contexts by presenting montages of images on large peripheral displays. These serve as contextual cues into pending tasks and projects.

Other researchers have realized that large displays may afford users a greater sense of presence, which may benefit performance of certain tasks. Slater & Usoh (1993) define presence as "a state of consciousness, the (psychological) sense of being in the virtual environment." They distinguish it from immersion, which they define to be an objective description of the technology, describing "the extent to which computer displays are capable of delivering ... illusion of reality to the senses of the human

participant.” In most current models, the sense of presence is seen as a direct outcome of immersion. The more inclusive, extensive, surrounding, and vivid the display, the higher the potential of presence (Bystrom, Barfield, & Hendrix 1999). In fact, when users are present in Virtual Environments (VEs), the location of their physical bodies is often construed as being contained within that space rather than looking at it from the outside. It is in this state that users are most effective in VEs.

Tan, Steffanucci, Proffitt, & Pausch (2001) utilize large peripheral projection displays to show different scenes of distinct ‘places’ that the user can use as cues to remember more information. They hypothesize that the greater the sense of presence invoked in the user by the display, the better the memory for learned information. They do not, however, explore how the physical size of the display affects this sense of presence.

There exists a large body of research in media communications exploring the impact of screen size on viewer responses to media content. In their survey paper, Grabe, Lombard, Reich, Bracken, & Ditton (1999) summarize this work. They divide current work into the effects of screen size on the perception of reality and presence, viewer enjoyment and preference, viewing distance and picture quality, arousal, as well as attention and memory. Most of the studies they examined showed that larger screens generally evoke more intense responses from viewers, causing them to feel more present in the scenes. Reeves & Nass (1999) document similar results. In addition, they describe a study showing that larger displays influenced the subjective experience of arousal as well as memory for a variety of video segments. They asserted, as I do, that emotional responses are affected by the form of the message as well as its content. However, in many of these studies, researchers were interested primarily in the interaction of screen content with the display technology as it would be used in a situated setting. As such, visual angle was not carefully controlled. Even when it was measured, it was generally assumed that a larger image would fill a greater percentage of the viewer’s visual field, hence confounding physical size with visual angle.

In fact, in most current work, researchers have placed large displays at a distance that is not proportional to their increase in size over small displays. Due to space constraints, they are typically relatively closer, hence casting a larger retinal image and offering a wider field of view (FOV). It is generally agreed that wider FOVs can increase “immersion” in VEs (Lin, Duh, Parker, Abi-Rached, & Furness 2002; Prothero & Hoffman 1995). Researchers in the entertainment industry have reported that larger displays filling a wider FOV can increase the level of involvement experienced by users (Childs 1988). Czerwinski, Tan, & Robertson (2002) report evidence that a wider field of view offered by a large display leads to improved performance in 3D navigation tasks, especially for females. They document prior literature suggesting that restricting FOV leads to negative impacts on perceptual, visual, and motor performance in various tasks, possibly because users find it difficult to transfer real world experience and cognition into the VE. In subsequent studies, Tan, Czerwinski, & Robertson (2003) showed that this effect was due mainly to additional optical flow cues that were afforded by wider FOVs. Arthur (2000), in his doctoral dissertation, provides a comprehensive review of the effects of FOV on task performance, especially as carried out in head-mounted displays.

Despite the large amount of work done in comparing FOVs, few researchers have isolated the effects of physical size and distance with carefully controlled FOV on the sense of presence or on task performance. Media communication researchers have evaluated viewing distance preferences as a function of various display factors, including image size and resolution. Experimental evidence in this field generally shows that the ratio of preferred viewing distance to image height is a constant that varies based only on image resolution. For example, Fish & Judd (1991) showed that for standard NTSC video, the preferred viewing distance to height ratio is about 7. However, other researchers such as Lund (1993) argue that contrary to predictions, the ratio actually decreases as image size increases, and that ratios are relatively unaffected by resolution. He hypothesized that the non-linear curve is due to viewers selecting their viewing positions to optimize a sense of presence or reality. Regardless of the findings, all these studies propose different hypotheses to explain the different viewing distance preferences. However, none have concretely proven or disproven these hypotheses.

Dixon (1999) used the vertical-horizontal illusion, in which people overestimate the vertical but not horizontal extent of objects, to study the differences in perceived proportions between small and large

displays. She found that the difference in magnitude for this illusion was influenced by the physical object or image size and not the perceived depth of the display. Her results further suggest that the proportions of larger objects appeared more compressed when presented on smaller displays, but that this effect could be compensated for by stretching the vertical proportions of images on a small display.

To examine the psychophysical effects of distance and size, Chapanis & Scarpa (1967) conducted experiments comparing the readability of physical dials at different distances. They used dials of different sizes and markings that were proportional to the viewing distance so as to keep visual angles constant. Perhaps surprisingly, they found that beyond 28 inches away, dials adjusted to subtend the same visual angle were read more easily at greater distances. The effects they found were, however, relatively small.

In a more recent study, Patrick, Cosgrove, Slavkovic, Rode, Verratti, & Chiselko (2000) examined various display technologies, with comparable visual angles, and their effects on the spatial information users acquired by navigating through a VE. They found that while users performed significantly worse in forming cognitive maps and remembering the environment on a desktop monitor, they performed no differently using a head-tracked head-mounted display or a large projection display. They attributed part of this effect to a higher level of presence afforded by the size of the projection display, which compensated for the immersion afforded by the head tracking.

Despite the deep understanding the current body of literature offers, there seems to be a gap in work isolating the effects of display size and distance, given a constant visual angle, for performance on tasks. Because of the emergence of large displays in the workplace and in consideration of everyday desktop computing tasks, I have decided to examine the effects of holding visual angles constant and varying only physical size. I will examine subjective responses as well as cognitive and social reactions to information, and will show how they may be used to construct display systems that make people more productive.

DISPLAY GARDEN

To explore the benefits that physically large displays offer, I do my daily work in a system I created, called the Display Garden. This system is a rapidly configurable collection of physical display devices such as whiteboards and pin-up space, audio displays, mobile LCD panels, and large projection displays on various surfaces in the room (see Figure 2). In setting up the Display Garden, I have had to solve interesting technology-related problems that come with working on physically large displays. For example, because I was front projecting, users in the room often found themselves working between the projector and the display surface. This caused undesirable projection on the user as well as temporary blindness from looking into the bright light of the projector. To alleviate this problem, I developed Pre-emptive Shadows (Tan & Pausch 2002), a technique that uses a camera-projector system to detect and turn off pixels that would otherwise be needlessly cast upon users' bodies and faces.

Although I do not view the creation of this system as a significant intellectual contribution in and of itself, the Display Garden provides hardware infrastructure that allows me to rapidly prototype new ideas.



Figure 2: The Display Garden – A rapidly configurable collection of physical display devices.

It serves as a tool for me to apply and validate design principles I formulate for building information environments with large displays. Additionally, ideas garnered from working in the Display Garden form the basis of my thesis. I have already begun to explore ideas inspired by my work in the Display Garden. Below, I will describe a few of these areas: spatial cognition, information privacy in social environments, as well as pragmatics of using multiple display systems that include large displays.

WITH SIMILAR VISUAL ANGLES, LARGER DISPLAYS IMPROVE SPATIAL PERFORMANCE

I have run a set of three studies to explore the effects of large displays in the performance on spatial tasks (Tan, Gergle, Scupelli, & Pausch 2003). The first study included a reading comprehension task and a spatial task. Users performed both tasks on a large wall-projected display and on a smaller desktop monitor. Because I was interested in the effects of physical size, I kept the visual angle subtended from the user to each of the two displays constant (see Figure 3). Although I found no significant differences in reading speed or comprehension levels, I did find that users performed 26% better on a task requiring mental rotation when working on the large display. I hypothesized that this effect was due to a greater sense of presence afforded by the large display. This sense of presence allowed users to use an egocentric strategy, in which they imagined rotating their bodies within the environment, rather than an exocentric one, in which they imagined rotating the environment around themselves.

In order to validate this hypothesis, I ran a second study that presented the same task with two different representations, the original one that did not bias users to use either an egocentric or exocentric strategy, and the other which biased users toward an exocentric strategy. I found that the benefits of the large display extended only to the task that allowed easy access to egocentric representation and strategy.

I have recently run a third study using the same task, this time with identical representations. Users performed the task in one of three conditions, depending only on the type of instructions they were given: explicitly egocentric, explicitly exocentric, or the original unbiased instructions. Results from this study showed much more clearly that users performed better when they were provided with an egocentric strategy than when they used an exocentric one. In the absence of an explicit strategy, the small display biased users into an exocentric strategy, while the large display biased them into using the much more efficient egocentric strategy.

Results from the three studies together suggest that, given a constant visual angle, the size of a display affects perception and performance in spatial orientation tasks. The contribution of these findings is significant because I have begun to show that, along with traditional characteristics such as brightness, contrast, position, and field of view, physical size must be carefully understood in designing large-scale information display systems.

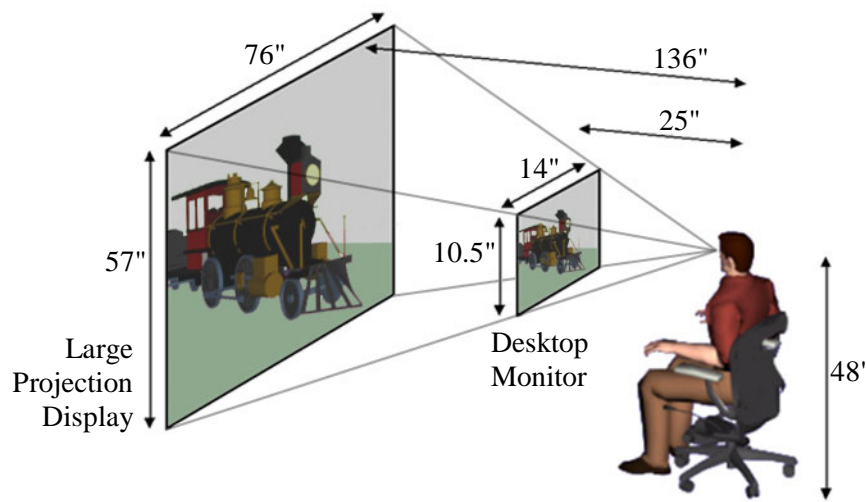


Figure 3: I maintained a constant visual angle and image for each of the two displays. Only size and distance to the user



Figure 4: Input device used (left); User working on the small (middle) and large (right) displays.

Experiment 1

Materials

I used two displays, an Eiki Powerhouse One LCD projector and an 18" Sony Trinitron E400 CRT monitor. Both displays ran at a resolution of 1024 x 768, and were calibrated to be of roughly equivalent brightness and contrast. I mounted the projector from the ceiling and projected onto a white plaster wall. The image on the monitor was 14" wide by 10.5" tall. The image projected on the wall was adjusted to be exactly 76" wide by 57" tall (see Figure 4). I set the two displays up so that when either display was viewed from a specific spot in the room, the visual angle and hence the size of the retinal image, would be identical (see Figure 3). I assumed a comfortable viewing distance of 25" for the monitor. In order to get an image of identical retinal size, the projection was set up to be 136" away from the user. The center points of both displays were set to be at eye-height, 48" above the ground. The position of the monitor was carefully marked so that it could be moved in and out as necessary. To mark the spot around which the user's eyes should be centered, I stretched fishing line from two stands, one on either side of the user. A mark in the center of the line indicated the exact spot in the room where the retinal images would be of identical size.

The study was run on a single 800 MHz Dell computer equipped with a dual-headed nVidia GeForce2 MX graphics card. I controlled the activation and deactivation of the displays using the Windows 2000 multiple monitor API so that only one display was active at any given time. The user provided input using an IBM USB numeric keypad with keys I had marked for the study (see Figure 4).


Procedure

After users filled out a background survey, I gave them the input device and had them sit comfortably in the chair. I adjusted the height and position of their chair so that the center of their eyes was as close to the marked fishing line as possible. Once they were viewing the displays from the spot in the room that provided retinal images of identical size, I removed the fishing line. Users were asked not to further adjust the chair or move it around.

Guilford-Zimmerman Spatial Orientation Task

One of the task domains I chose was that of mental rotations. Presentation, degree of immersion, and level of performance have been extensively measured for such tasks. In their work, Suzuki & Nakata (1988) had six students perform a mental rotation task similar to that of Shepard & Metzler (1971). Users were asked to judge whether pairs of figures, each of which had been rotated to different degrees, were identical in shape or not. They found, as Shepard & Metzler did, that mean reaction times increased linearly with the angular difference between figures. They also discovered that visual angle, which corresponds to retinal size, of the objects affected the speed of rotation. However, in this study, viewing distance, given constant visual angle, did not seem to affect reaction times.

Building on this work, Wraga, Creem, & Proffitt (2000) measured spatial knowledge by the time it took users to update their orientation after changing it. Results showed that users were faster at spatial



**Figure 5: Sample question from the Guilford-Zimmerman Spatial Orientation test.
Since the boat moved up and to the left, but did not rotate, the answer is 5.**

updating when they imagined rotating themselves in the environment rather than when rotating the environment around themselves. Carpenter & Proffitt (2001) extended these findings by examining egocentric rotations in each of the three possible rotation planes. They replicated the finding that egocentric rotation, or rotating one's self, was faster than exocentric rotation, rotating the environment. However, this was true only of planes in which users had experience rotating or locomoting. Tlauka (2002) found similar results by comparing rotations of images presented horizontally or vertically.

To evaluate the effects of display size on spatial performance, I utilized the Guilford-Zimmerman Spatial Orientation test (Guilford & Zimmerman 1948). This test has been well validated and researchers have shown that results from this test correlate highly with wayfinding ability (Infield 1991).

Each question in this test contained two pictures seen from the prow, or front, of a boat along with a multiple choice answer key (see Figure 5). The user was asked to imagine that each picture was taken with a camera fastened rigidly to the boat so that the camera bobbed up and down and slanted with the boat. First, the user looked at the top picture to see where the boat was initially heading. This heading is represented by the dot in the answer key. Next, the user looked at the bottom picture and determined the change in orientation of the boat. The line in each of the possible answers represents the new orientation of the boat relative to the previous heading. Finally, the user selected the answer with the number keys, confirmed the answer with the enter key, and proceeded to the next question.

I gave users the paper-based instructions that were provided with the standard Guilford-Zimmerman test. They then tried 3 practice questions on the large display. For these questions, the system provided users with immediate feedback explaining the correct answers. After they had performed the practice questions, users proceeded to perform the test on the small and the large display, which I will refer to as Display Size. They were not given feedback for these questions. 60 questions were randomized and broken into two sets. Users had 5 minutes to answer 30 questions in each of the two conditions, and were told to perform as quickly and accurately as possible. The order of Display Size was counterbalanced across users. Users had a 30 second rest interval between each condition.

Reading Comprehension Task

Another task I chose to evaluate was a reading comprehension task. As with mental rotations, presentation and level of performance for reading comprehension tasks have been well researched (Dillon 1992; Mills & Weldon 1987). Researchers have shown that many factors affect the readability of text, or ease with which meaning of text can be comprehended, on computer screens. Such factors include features of characters, formatting, contrast and color, as well as other dynamic aspects of the screens.

I tested subjects on a reading comprehension task in the two Display Size conditions. Based on normalized average scores for specific passages as well as pilot test data, I chose a suite of 7 passages

from practice GRE tests (Educational Testing Service 1994) that had relatively similar levels of difficulty. Each of these passages contained about 460 words, or 56 to 60 lines of text as laid out in paper-based GRE format. Each passage came with a set of 7 questions that the user answered after reading the passage.

I gave users verbal instructions on how to scroll through the passage and to answer questions. Then they performed the task with a practice passage on the large display with medium-sized text. I instructed them to work through the questions quickly but accurately. When they had finished the practice passage, they read the rest of the passages and answered questions in each of the 6 conditions, created by presenting text in a given Font Size (Small: 10 point vs. Medium: 14 point vs. Large: 18 point) on each of the Display Size conditions (Small vs. Large). Font Size and Display Size were counterbalanced separately. Users had a 30 second rest interval between passages.

Post-test Preference Questionnaire

After users completed the tests, they filled out a questionnaire indicating their preference for the conditions in each of the tasks. They were also encouraged to comment on their opinion of the displays.

Participants

Twenty-four (12 female) college students, who were intermediate to experienced computer users, participated in the study. I screened users to be fluent in English and to have normal or corrected-to-normal eyesight. The average age of users was 25.4 (25.5 for males, 25.3 for females), ranging from 19 to 32 years of age. Users were paid for their participation.

Results

I present the results from Experiment 1 in three parts. First I explore performance on the spatial orientation task, then the performance on the reading comprehension task, and finally I investigate preference measures collected at the end of the study.

Spatial Task Performance

I analyzed data for the spatial orientation task at the summary level. The dependent variable was the percentage of correct responses (number correct / number attempted)¹. Time differences between different Display Sizes were not significantly different and were therefore dropped from the final models (levels of significance did not change either way). I analyzed the percentage of correct answers with a 2 (Display Size) x 2 (Position) x 2 (Gender) repeated measures analysis of variance (RM-ANOVA). I analyzed Gender and Position as between-subjects factors and Display Size as a within-subject factor.

I found a significant main effect of Display Size ($F(1,20)=9.470$, $p=.006$) with the large display resulting in a higher percentage of correct responses on average (Small=43.8%, Large=55.4%; see Figure 6). I also observed a significant main effect of Gender ($F(1,20)=5.072$, $p=.035$), with males producing a higher percentage of correct responses than females on average (Female=38.7%, Male=60.4%). None of the 2-way or 3-way interaction effects was significant. While Czerwinski et al. (2002) suggest that females benefit significantly more than males in 3D navigation tasks using displays with wide fields of view, we saw no such effect for Display Size in our studies. The larger display seems to have benefited both males and females equally.

Since this was a timed task, an alternate explanation for these findings may include a speed-accuracy tradeoff. However, a separate analysis confirmed there was no difference in time spent per question in the two conditions. In addition, I implicitly controlled for the effect of time by taking the percentage of answered questions rather than raw scores (see Footnote 1).

Overall, I found a significant improvement in the percent of correct responses on the spatial task for users working on the larger wall display. Keep in mind that while the absolute size of the image was larger, the retinal image size was kept nearly constant regardless of Display Size.

¹ I use percentage of correct answers as the dependent variable since it is a straightforward and intuitive measure. However, I also examined the sum of correct responses, controlling for time, and found nearly identical results.

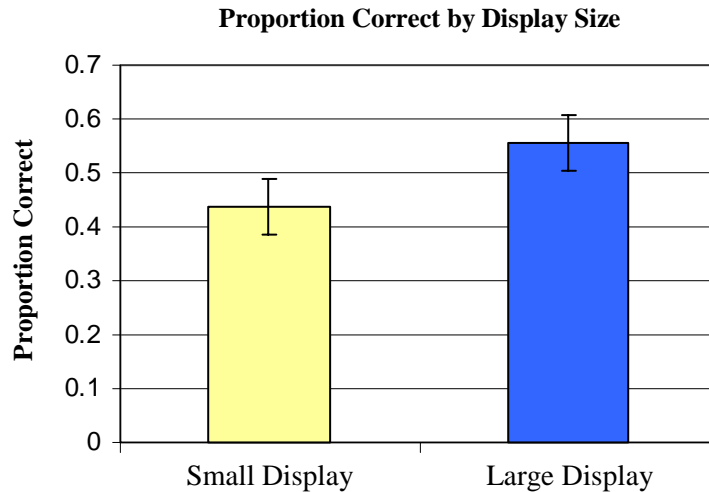


Figure 6: Main effect of Display Size. Large display aided users in the Guilford-Zimmerman test.

Reading Comprehension Performance

In the reading comprehension task, I again analyzed data at the summary level. I used the number of correct responses for each condition as the dependent variable. I performed an RM-ANOVA in which Position, Display Size, and Font Size were repeated and Time to complete the question was a covariate. I included all 2-way and 3-way interactions in the analysis. Because each user participated in multiple trials, within observations were not independent. User was modeled as a random effect.

Overall, performance in the reading comprehension task did not differ across the conditions. I found no difference between small ($M=3.86$) and large ($M=4.01$) Display Sizes ($F(1,106)=.367$, $p=.546$). Similarly, I found no difference between the small ($M=3.84$), medium ($M=3.96$) and large ($M=4.01$) Font Sizes ($F(2,106)=.176$, $p=.839$). The interaction between Display Size and Font Size was not significant, ($F(2,106)=1.159$, $p=.3178$).

I was unable to reject the null hypothesis that the displays were equal for performance on the reading comprehension task. Thus, while I did find differences on the spatial orientation task, I found no evidence to suggest that the performance on reading comprehension was different on either of the two Display Sizes, regardless of Font Size.

Preference Data

In addition to the performance data, I gathered preference data from users at the conclusion of the study. The questions were asked on a 5-point Likert scale of 1="Strongly prefer small display" to 5="Strongly prefer large display". Given the performance difference I found on the spatial task, I was primarily interested in user preference for this task. Users significantly preferred the large display for both 'Ease of Seeing,' ($M=3.61$, $p=.019$) and 'Overall Preference,' ($M=3.50$, $p=.045$). They marginally preferred the large display for their 'Confidence in the Rotation Task,' ($M=3.43$, $p=.066$). Users showed no significant preference for display in the reading comprehension task.

Summary

While I found a difference in performance on the spatial orientation task presented in this study, I did not know exactly what had caused the performance benefits. Was it because the larger display provided a greater absolute difference in scale, provided a more immersive environment that encouraged egocentric representation and strategy, or did it have something else to do with the surroundings? In order to further explore the reason behind the 26% improvement in performance, I decided to run a second study to investigate the difference.

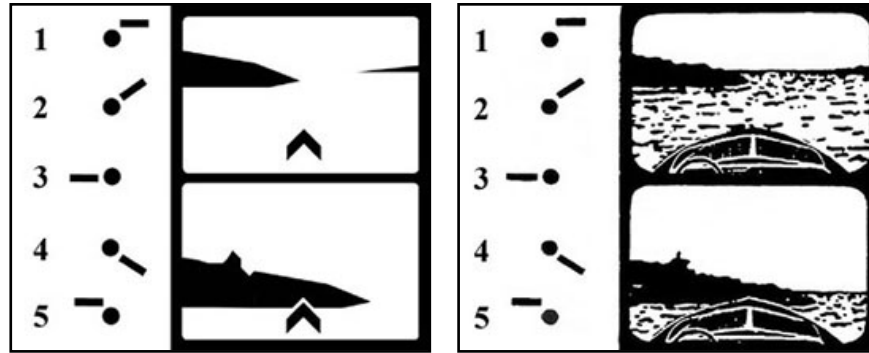


Figure 7: New Shape test with abstract representation (left); Original Guilford-Zimmerman test with boat theme (right).

Experiment 2

Materials and Procedure

I used the same hardware setup as in Experiment 1. For this study, I created a derivative of the Guilford-Zimmerman test, which I call the Shape test. The original Guilford-Zimmerman test was carefully crafted not to bias a user into any particular strategy. I designed the new Shape task to explicitly bias the user into an exocentric strategy. Although I could have used a test that already does this, such as the Shepard-Metzler test (Shepard et al. 1971), I decided on creating the Shape test to keep baseline performance levels between the two groups similar. To create this test, I did two things. First, I mimicked the scenes presented in the original test with abstract shapes (see Figure 7). Second I edited the instructions so that rather than asking users to imagine themselves looking through a camera mounted on a boat that was bobbing and rocking, I had them imagine themselves looking through a camera that was rigidly mounted on the ground in front of a painting that was moving and turning. After balancing for gender, each participant was randomly assigned to one of the two Task Types: original Guilford-Zimmerman test or my new Shape test. The procedure for each test was the same as in Experiment 1, minus the reading test.

Participants

Twenty-four (14 female) college students, who did not participate in the first study, participated in this study. As before, I screened users to have normal or corrected-to-normal eyesight. The average age of users was 24.1 (25.4 for males, 23.2 for females), ranging from 18 to 56 years of age. Users were paid for their participation.

Results

I present results from Experiment 2 in two parts. First I explore performance on the spatial orientation task; then I examine the preference data.

Spatial Task Performance

I modeled the data as I did for the spatial orientation task in the previous study. Again I examined data at the summary level and used the percentage of correct responses (number correct / number attempted) as the dependant variable. I analyzed the percentage of correct responses with a 2 (Display Size) x 2 (Task Type) x 2 (Position) RM-ANOVA. I analyzed Task Type and Position as between-subjects factors and Display Size as a within-subject factor.

Overall, I found a marginally significant effect of Display Size ($F(1,21)=3.296, p=.084$) with the large display resulting in a marginally higher percentage of correct responses, on average (Small=41.8%, Large=46.3%). I did not necessarily expect this main effect to be as strong as in Experiment 1, particularly if the Shape test showed no improvement in the large display condition. A more detailed investigation revealed that this was indeed the case. I found a significant interaction between the Display Size and Task Type ($F(1,21)=5.512, p=.0288$) demonstrating that the benefits of the large display were

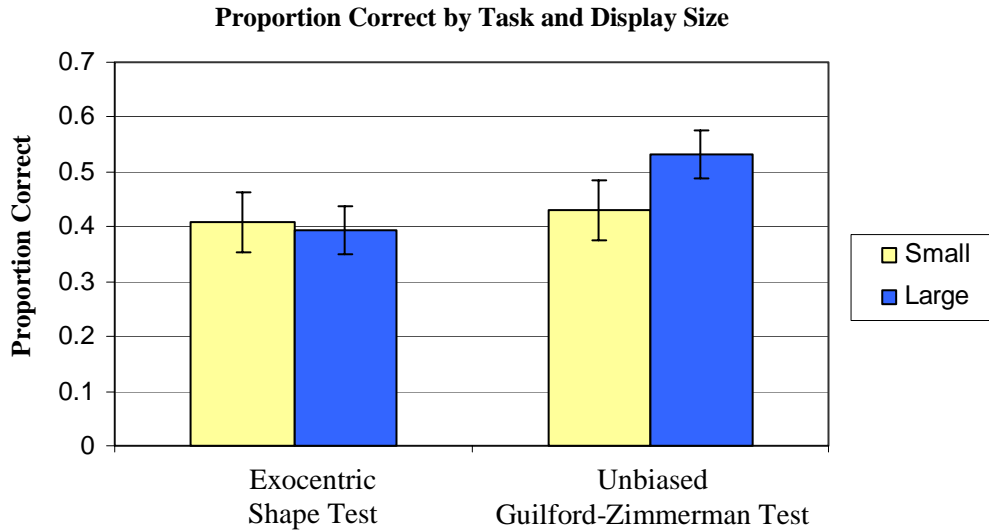


Figure 8: The large display was better for the Guilford-Zimmerman test but not the Shape test.

significantly greater for the Guilford-Zimmerman task (Small=42.9%, Large=53.2%) than for the Shape task (Small=40.8%, Large=39.5%). Thus, the large display provided for better performance in the Guilford-Zimmerman task, but it did not appear to help users with the Shape task (see Figure 8).

These results provide partial replication of the findings from Experiment 1, with additional insight into the hypothesis that large displays facilitate egocentric rotation.

Preference Data

As in Experiment 1, I gathered preference data at the conclusion of this study. The same questions were asked using the same 5-point Likert scale. The merged preference data for both the Guilford-Zimmerman and Shape groups were not significantly in favor of the large display. I explored whether or not the different task groups viewed the value of the displays differently (see Figure 9). I found that users performing the Guilford-Zimmerman task preferred the large display for ‘Confidence in the Rotation Task,’ ($t(22)=-2.250$, $p=.0348$) and marginally preferred the large display for ‘Overall Performance,’ ($t(22)=-1.773$, $p=.090$), in comparison to the users doing the Shape task. However, they did not rate the large display as any better than Shape test users did for ‘Ease of Seeing,’ ($t(22)=-1.089$, $p=.288$).

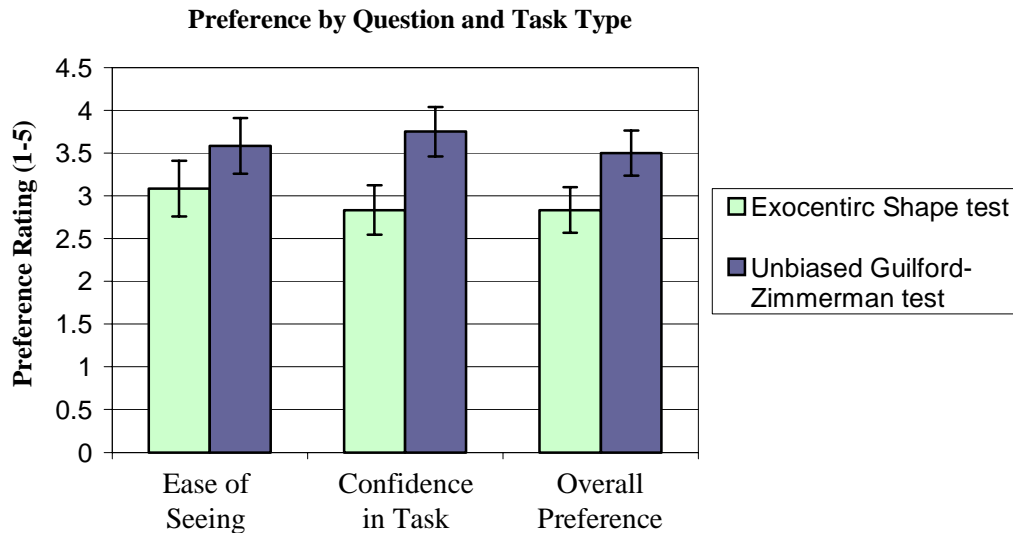


Figure 9: Users in the Guilford-Zimmerman test significantly preferred the large display (>3) for ‘Confidence in Task’ and ‘Overall Preference’, but not for ‘Ease of Seeing.’

Summary

In this study, I showed that users performing the Guilford-Zimmerman spatial orientation task on a small display performed only as well as users performing the Shape Test. Users performed the Shape Test equally on the small or large display. Recall that the Shape Test was a derivative of the Guilford-Zimmerman test that I created to intentionally bias users into an exocentric strategy. This study was a step towards proving my hypothesis that the size of displays affects the choice of strategy and that this choice in turn affects user performance. However, results did not decisively show that an egocentric strategy indeed accounted for the improved performance. Also, since I changed the visual stimuli along with the instructions in creating the Shape Test, there was some question as to how much the actual symbolic representation of the task mattered. I conducted a third study to further isolate the effects.

Experiment 3

Materials and Procedure

I used the same hardware setup as in Experiment 1. For this study, I used only the stimulus provided with the Guilford-Zimmerman test, as in Experiment 1. Recall that the instructions for the original Guilford-Zimmerman test were carefully crafted not to bias a user into any particular strategy. From this instruction set, I created two others, one that intentionally biased users into an egocentric strategy and another that biased users into an exocentric strategy. The egocentric instructions describe a scene in which users are asked to imagine themselves physically on the boat as it moves within the environment. The exocentric instructions describe the boat as a rigid prop mounted to the ground with the scene on a backdrop that is moving with respect to the boat. After balancing for gender, each participant was randomly assigned to one of the three Instruction Types: egocentric instructions, exocentric instructions, or original Guilford-Zimmerman instructions. The procedure was the same as in Experiment 1, minus the reading test.

Participants

Forty-two (18 female) college students, who did not participate in the first two studies, participated in this study. As before, I screened users to have normal or corrected-to-normal eyesight. The average age of users was 21.8 (21.7 for males, 22.2 for females), ranging from 18 to 35 years of age. Users were paid for their participation.

Results

Again, I modeled the data as I did for the spatial orientation task in the previous two studies. I examined data at the summary level and used the percentage of correct responses (number correct / number attempted) as the dependant variable. I analyzed the percentage of correct responses with a 2 (Display Size) x 3 (Instruction Type) x 2 (Position) repeated measures multivariate analysis of variance (RM-MANOVA). I analyzed Instruction Type and Position as between-subjects factors and Display Size as a within-subject factor.

Overall, I found a significant effect of Instruction Type ($F(2,37)=3.866$, $p=0.030$; see Figure 10). Paired comparisons using the Bonferroni technique showed a significant difference between the egocentric and the exocentric instruction sets ($p=0.01$), with users getting a higher percentage of questions correct with egocentric instructions (Egocentric=66.5%, Exocentric=47.2%).

I conducted several post-hoc tests to see if users who were explicitly instructed to use a given strategy performed any differently from users who implicitly chose a strategy due to the Display Size. I found no significant differences between users in the exocentric condition and the unbiased small display condition, which was assumed to elicit an exocentric strategy ($t(40)=-.079$, $p=.9371$). Similarly, I found no significant differences between users in the egocentric condition and the unbiased large display condition, assumed to elicit an egocentric strategy ($t(40)=0.953$, $p=.3463$). I also conducted additional tests comparing performance on the small display in the exocentric condition to the small display in the unbiased condition, as well as the large display in the egocentric condition to the large display in the unbiased condition. In both cases, there were no significant differences.

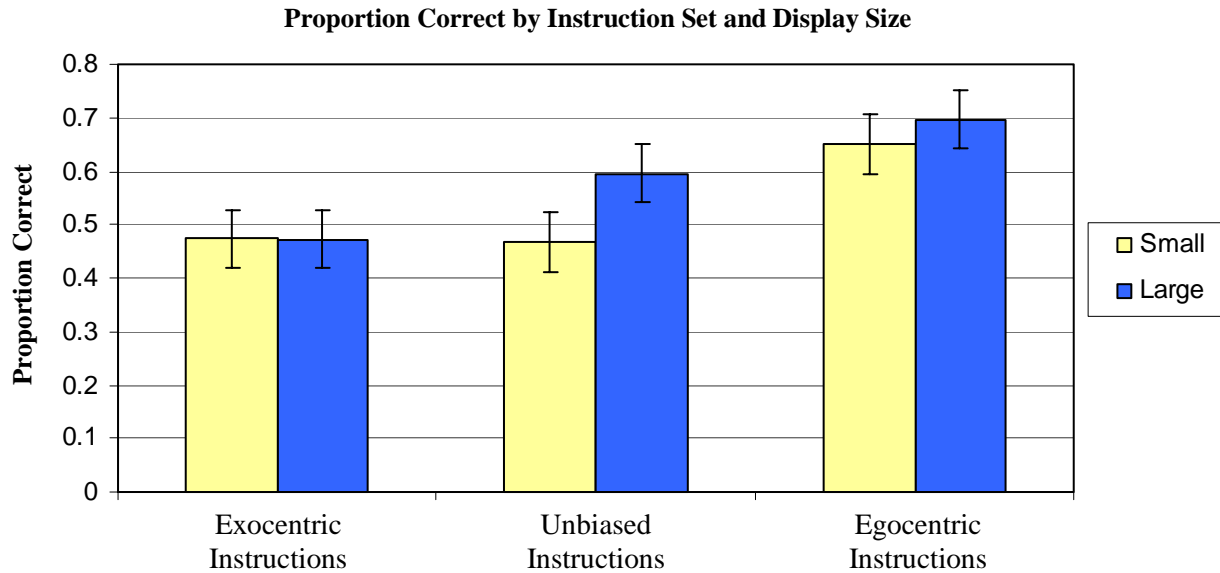


Figure 10: Main effect of Instruction Set.
Egocentric instructions led to increased performance over exocentric instructions.

These results replicate finding from the previous two studies as well as provide additional support into my hypothesis that large displays provide a greater sense of presence and bias users into using egocentric strategies.

Summary

Results from this study show much more clearly that users perform better when they are provided with an egocentric strategy than when they use an exocentric one. In the absence of an explicit strategy, users choose an exocentric one when working on the small display and the much more efficient egocentric one when working on the large display. Results from these three studies together suggest that, given a constant visual angle, the size of a display affects perception and performance in spatial orientation tasks. Future work includes determining whether or not user interaction affects the choice in strategy, as well as finding the class of tasks for which this finding generalizes.

SOCIAL IMPACT OF LARGE DISPLAYS ON INFORMATION PRIVACY

I have also explored effects of physically large displays on social reactions, specifically studying how large displays affect notions of information privacy (Tan & Czerwinski 2003b). One observation when working on physically large displays, such as wall-sized projection, is that a certain amount of information privacy is lost (Bishop et al. 2000). A common explanation for this loss in privacy is the higher legibility of information presented on large displays. Because large displays are typically viewed from a distance that is not proportionally scaled with the increase in display size, they often provide a larger visual angle, making them easier to see and read. While I agree that this contributes to the loss of privacy, I assert that there are more subtle social cues that may also contribute to this effect. I conducted a study to examine both the effectiveness of a new paradigm for measuring whether or not someone has read a particular passage of text, as well as to quantify the anecdotal observations and show that people are indeed more likely to read text presented on a larger display, even when legibility was constant.

Materials

In this study, I used two displays, an NEC MultiSync FE1250 22" monitor and a Sanyo PLC-XP30 LCD projector. Both displays ran at a resolution of 1024 x 768 and were calibrated to be of roughly equivalent brightness and contrast. The image on the monitor was 16" wide by 12.5" tall. The image projected on a wall-mounted screen was adjusted to be exactly 66" wide by 49.5" tall. I set the displays up so that when




Figure 11: Visual angle held constant between the small display and large display conditions.

either display was viewed from the participant's seat, the visual angle and the size of the retinal image would be identical (see Figure 11). As described above, I have already shown that reading performance did not significantly differ on two such displays. Additionally, I ensured that someone using the system would not occlude any part of either display from a participant.

I measured whether participants had read content on the displays with an implicit memory priming paradigm, usually employed to study learning without awareness (Schacter 1987). In this paradigm, participants are presented with target words and are later tested, for example with stem completion, on their implicit memory for these words. In stem completion, participants are given the beginning of a word (e.g. mil___ for military) and must complete it with the first appropriate word that comes to mind. Priming is reflected by an enhanced tendency to complete stems with target words.

In this study, I constructed seven e-mail subject lines and two e-mail messages that included a total of 30 target words selected from the Kucera & Francis (1967) norms. These words were selected to be between 5 and 12 letters in length and of medium frequency (mean: 87 per million). The initial three letters, or stem, of each word was unique to all other text given to the participant and each had at least 8 different completion possibilities. Additionally, 33 filler words were selected with the same criterion.

Procedure

Participants were seated in the armchair and handed a survey on multiple monitor preferences. They were informed that they would have exactly seven minutes to complete this survey and should look over their answers if they got done earlier. The survey, consisting of one open answer and ten multiple-choice questions, was designed to take less than seven minutes to complete.

I used a between-subjects design, with participants balanced by gender and assigned randomly to one of the two display conditions: small display vs. large display. While the participant completed the survey, the experimenter read the prepared target content on one of the two display setups. Each e-mail message was viewed in the Microsoft Outlook e-mail client for three and a half minutes. The seven subject lines remained visible in the inbox for the entire seven minutes. Although the experimenter reading the e-mail could not directly observe the participant, the participant was video taped during the experiment.

After the survey, the participant performed the stem completion test. The test consisted of 63 stem completions: 3 practice questions, 30 filler questions, and 30 target questions. The 3 practice questions were followed by the 60 filler and target questions presented in random order. Following the test, participants completed a questionnaire explicitly asking whether or not they had read content on the experimenter's display while doing the survey.

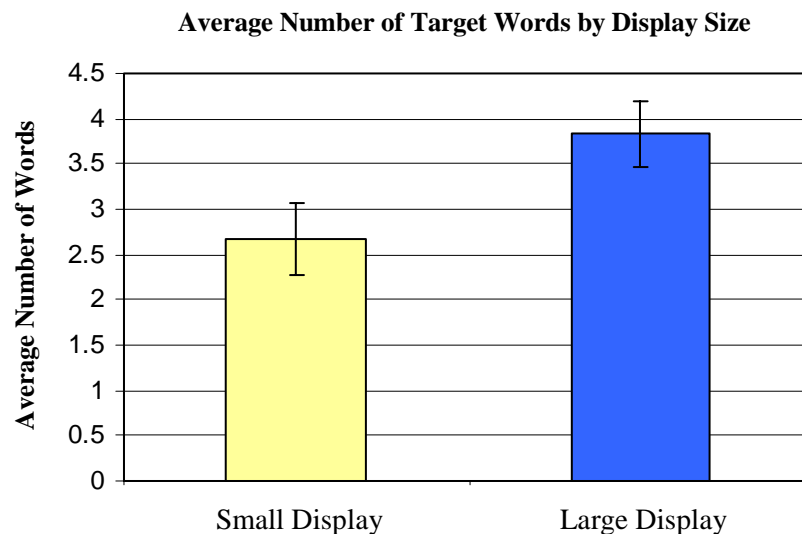


Figure 12: Users responded with more target words in the large display than the small display condition.

Participants

Twenty-four (12 female) intermediate to advanced windows users with normal or corrected-to-normal eyesight participated in this study. The average age of users was 36.9, ranging from 18 to 55 years of age. Users were given a software gratuity for their participation.

Results

I found significant differences between conditions in the number of stems completed with target words, suggesting that users had read more information displayed on the large display ($M=3.83$ words) than on the small one ($M=2.67$ words). This was true with both a loose metric that permitted different forms of the target words, as well as a concise one that allowed only exact forms that had been presented ($t(22)=2.0739$, $p=0.04$; see Figure 12).

On post-test surveys, more users admitted to having read text on the display in the large screen condition (7 of 12) than in the small one (3 of 12), marginally significant by Fisher's Exact Test ($p=0.089$). Additionally, video tapes showed users spending longer periods viewing material on the large screen ($M=19$ seconds) than the small ($M=14$ seconds), though this difference was not significant.

Summary

In this work, I designed a novel paradigm to measure whether or not a user has read certain content. I show that, even with constant visual angles and legibility, visitors are still more likely to glance over a user's shoulder to read information on a large wall-projected display than on a smaller traditional desktop monitor. I hypothesize that, in addition to legibility, there are more subtle social factors that may contribute to the loss of privacy on physically large displays. I believe that social convention prescribes that people have certain personal zones within which objects (information included) are deemed private. With few exceptions, any object outside of this zone is assumed to be public. Also, culturally, objects placed on walls are typically considered public. Large wall-sized displays exhibit both sets of public cues. In future research, I will further explore these hypotheses, as well as implement hardware and software ideas for exploiting these effects.

EFFECTS OF VISUAL SEPARATION AND PHYSICAL DISCONTINUITIES WHEN DISTRIBUTING INFORMATION ACROSS MULTIPLE DISPLAYS

I conducted a study designed to explore the effects of visual separation and physical discontinuities when distributing information across multiple displays in systems that include both small and large displays (Tan & Czerwinski 2003a). Having multiple displays enlarges the physical display area, allowing the system to present information across much wider visual angles from the user. Also, since displays are

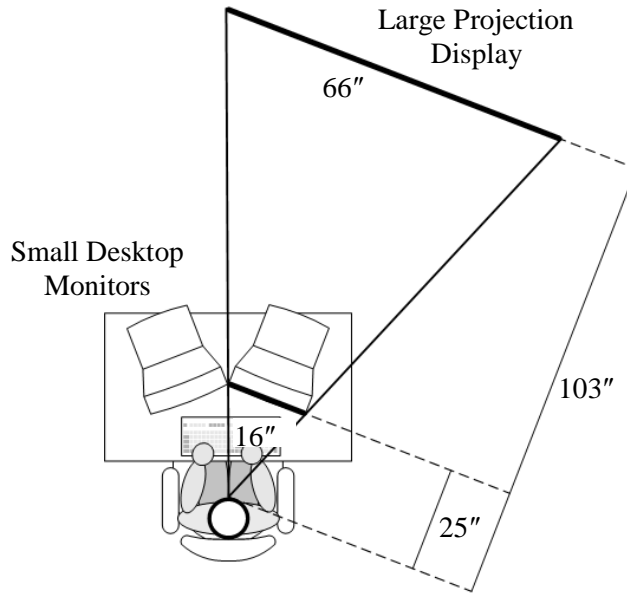


Figure 13: Visual angle held constant between the small display and large display conditions.

often placed at different depths or are framed by physical bezels, physical discontinuities are introduced in the presentation of information in these workspaces. Yet, relatively little is known about how to best present information to the user given these display characteristics.

To isolate individual factors, I created a display system that allowed me to carefully control the separation and discontinuities associated with multiple displays (see Figure 13). Specifically, I varied the visual angle at which information was separated (55 vs. 27 degrees), whether or not information was on a single display or separated across two displays, as well as the types of displays used together (small-small vs. small-large). I ran a preliminary study that utilized a divided attention paradigm across these different display conditions. The test included a primary task, done in conjunction with a secondary or tertiary task. In the primary task, users proofread and identified grammatical errors within a set of text articles. While doing this, users performed the secondary task, notification detection. In this task, users had to detect and act upon visual changes outside the focal region of the primary task. Upon detecting notifications, users performed the tertiary task, text comparison, in which they had to cross reference and compare content displayed in multiple locations on the displays. I picked these tasks to be representative of tasks carried out by information workers while multitasking. I ran twenty-four (12 female) users in this experiment.

Interestingly, for the tasks chosen in this study, I saw no effects of physical discontinuities, introduced either by monitor bezels or by the depth difference between the monitor and projection display. This was surprising, but implies that designers might have more freedom when splitting information across boundaries than I had anticipated. I do not doubt that there are tasks that will be hurt by splitting information across physical discontinuities, but this set of tasks do not seem to fall into that category.

However, I did find significant performance differences when the tasks were separated by a wider visual angle (55 vs. 27 degrees), but only when information was split between the desktop monitor and the projection display. This indicates that, even at similar visual angles, placing information further in the periphery on displays that are separated in depth is more detrimental to performance than the corresponding position at similar depths. However, it should be noted that for these tasks, effects seen were relatively small (about a 10% performance decrement), and designers, aware of the small differences present, can weigh the importance of the information to be displayed with this trade off in mind.

DETAILED RESEARCH PLAN

I present my detailed research plan as part of a proposed outline for my final thesis document. I will leave out chapters 1, 2, and 8, the Introduction, Related Work, and the Conclusion, respectively. For each of the

remaining chapters, I will first briefly discuss my initial results. I will then present a set of new hypotheses as well detailed approaches for exploring each of them. As with any research endeavor, this is merely meant to provide the direction in which I will begin my journey. Along the way, I hope to discover interesting detours and create new paths to pursue.

Chapter 3: Defining Large Displays

In my initial experiments, I have shown in several contexts that physically large displays elicit different cognitive and social reactions than smaller displays, even when they are viewed at identical visual angles. These reactions increase performance of certain tasks on large displays. However, these experiments were conducted with a specific set of small and large displays, namely, desktop monitors and wall-sized projection. In order for the principles derived in this thesis to be truly technology-independent, I will isolate the exact characteristics that allow people to perceive a display as being large and that allow for observed performance increases.

Hypothesis 3a: Users rely on environmental cues more than optical cues in perceiving physical size.

In order to isolate the effects of physical size in my initial experiments, I have varied only the size of displays, moving each to the corresponding depth in order to keep visual angles identical. I have kept all other factors as constant as possible. As such, I am able to conclude with confidence that any effects observed are due mainly to the change in the physical size of the display. However, there are several cues that may account for the user's perception of the change in physical size. I will explore two of these: visual optics and environmental cues. Visual optics provides cues to the user because it requires that the eyes physically converge and accommodate differently to focus on displays at different depths. Environmental cues are provided by relationships between the display and other objects in the environment. These relationships, including occlusion, relative size, linear perspective, and texture gradient, among others, allow users to infer the physical size of objects. I hypothesize that most workplaces, environmental cues are more important than visual optics in perceiving physical display size.

In order to prove this hypothesis, I will conduct experiments comparing various display conditions in the real world to completely synthetic display environments created in virtual reality. I will examine performance on some of the tasks for which we have seen differences. Using a head-mounted display, which is merely a set of small displays really close to the user's eyes, I will isolate and answer the question of how important visual optics is. Using virtual reality, I will carefully control environmental context that exists around displays and answer some of the questions of environmental cues.

Hypothesis 3b: The physical size of a display is optimal when it presents life-sized images.

Since I have been looking for main effects between large and small displays in my initial experiments, I have examined only displays of two different sizes, a 20" small display versus a 76" large display. A question that has been raised in my initial experiments is whether large displays are optimal when they present images at life-size. Does making the display even bigger than life serve to exaggerate effects observed and make users even more productive? Or will that break the illusion and dampen the sense of presence that users experience, thus hindering performance?

To explore whether or not the physical size of a display is optimal when it presents life-sized images, I will conduct an experiment to examine the effects of greater-than-life-sized displays as well as smaller-than-traditional displays in the tasks for which I have seen performance benefits. I will derive a simple curve relating the 'life-sizedness' ratio of the display to performance on the tasks.

Chapter 4: Spatial Cognition

In initial experiments, I have shown that users perform significantly better when they use an egocentric strategy over an exocentric strategy on the Guilford-Zimmerman mental rotation task. More interestingly, I have shown that in the absence of an explicit strategy, small displays appear to bias users into an exocentric strategy, while large displays bias them into using the more efficient egocentric strategy, even when both are viewed at identical visual angles.




Figure 14: This scene can be described many ways, depending on the frame of reference used (taken from Levelt 1999).

There exist several limitations of the current findings. First, the Guilford-Zimmerman task is a static task that requires user interaction only in selecting the multiple choice answers. As such, I have no evidence whether interacting more directly with the content would also bias users into an egocentric strategy, thus negating some of the beneficial effects of the large display. Second, the Guilford-Zimmerman task is a psychological test that has been carefully crafted such that it does not bias users into using a particular strategy. In order to apply these findings, I have to generalize these results and understand the characteristics of tasks that benefit from this effect. Finally, the current set of results does not immediately transfer to real-world application. I will have to carefully apply knowledge gained from these experiments in order to formulate design principles and to build applications that are more efficient on large displays.

Hypothesis 4a: Large displays better bias users into egocentric representations than small displays, even when users are interacting with content.

In order to address the first concern, interactivity, I will use interactive 3D environments to examine how large displays affect users' frames of reference in the world. These frames of reference can be egocentric or exocentric and necessarily influence how users react to and perform within these worlds. I will leverage spatial cognition paradigms to understand whether large displays are able to implicitly bias users into more egocentric representations, even with high levels of interactivity. I will also measure task performance to determine if these representations are more effective in these environments.

I will use spatial path or map memory tasks in order to test the efficiency of recalling a specific representation. I will look both at reaction time as well as task error rates to determine if users have most likely stored a particular representation or if they have to perform a perspective transform in order to get the information into the desired format. Another particular interesting paradigm is linguistic analysis of spatial representation (Levelt 1999). In this paradigm users verbally describe the spatial relationships of various objects within the world. The lexical structure of these descriptions can be analyzed to yield fairly robust indication of the mental representation. For example, in Figure 14, a user with an egocentric representation may look into the scene and say "the ball is to the right of the chair," whereas a user with an exocentric representation may say "the ball is to the left of the chair."

Hypothesis 4b: Large displays improve performance on 3D navigation tasks.

The Guilford-Zimmerman task has been shown to correlate well with wayfinding ability. I believe that current findings might generalize to 3D navigation tasks. I will conduct an experiment in which I train novice users to navigate through a 3D videogame-like environment on small and large displays. I will measure performance metrics such as ease of control, mental map formation, as well as spatial memory for the environment. I expect that users trained on the large display will be biased into an egocentric representation and will perform better, at least initially. By examining the training data over time, I will determine whether or not the benefits of large displays continue to hold for non-novice or expert users.

Hypothesis 4c: Egocentric representations afforded by large displays transfer more readily to the real-world and are more efficient for training tasks.

In order to address the applicability of these results, I will examine the transfer effects of learning spatial information on large displays. Specifically, I would like to know if the egocentric representation afforded by spatial training on large displays transfers more efficiently to real world situations than training on traditional desktop displays. I will also quantify possible detrimental effects, if any, that exist for users who are trained on large displays and tested on smaller ones. Using this knowledge I will make suggestions on the kinds of tasks which benefit from the use of large displays.

Chapter 5: Social Environments

Although researchers have noticed that people are more likely to treat information presented on large wall-sized displays as being public, this has been a difficult effect to quantify. In my work, I have used an implicit memory priming paradigm to measure whether or not a user has read information on a display. I have shown that users are indeed more likely to look at private information presented on a large wall-sized display than on a standard desktop display, even given identical visual angles. I hypothesize reasons for this effect, including distance of the display from the user as well as the fact that the large display was projected on a wall, usually deemed to be public. I will further these findings in a series of studies.

Hypothesis 5a: People treat information on large displays as being public and are more likely to interact with it.

When people treat information as being more public, there exists a heightened potential for ad hoc collaboration. I have noticed anecdotally that people entering someone else's workplace are willing to read and write on whiteboards, even if the information is not explicitly demarcated as being public. In fact, many people will even erase existing content to make space for their additions. Conversely, people are fairly uncomfortable peering over a user's shoulder to read content on their desktop display, and they certainly would not venture to edit or erase information on it. In this work, I will continue to explore hardware and software implications as well as mechanisms that will serve to make information intrinsically more public or private. I will couple large displays, appropriate input devices, as well as supporting software to create a system that is more likely to encourage ad hoc user interaction with content displayed. I will continue to construct new paradigms to measure these 'hard to measure' effects in controlled settings. Alternatively, I might utilize field experiments in which I observe users with these new systems in their real work environments.

Hypothesis 5b: People working on displays that are publicly visible are more likely to stay on task and be more productive.

In multiple user workplaces, having personal displays big enough that they are visible by others may motivate users to stay more on task. Social pressure may prevent such users from straying too far from the tasks that they are supposed to be doing, thus making them more productive. Surprisingly, such a system may also allow for more efficient use of space, since walls that separate small cubicles may be torn down and replaced by a much more open workplace in which personal spaces gracefully overlap. However, a user working in such a system also has the potential of creating distractions for others as well as feeling a fundamental lack of privacy. I will conduct studies to explore the tradeoffs that exist when personal displays are large enough that they become public.

Hypothesis 5c: Systems that are aware of their physical displays as well as their social impact can automatically lay information out based on social context.

One idea that I may pursue is the notion of having the system realize the social impact of each of its physical displays and to lay information out accordingly, depending on the context of the task at hand. This would integrate lessons learned from the other experiments into a larger system of multiple displays. For example, public information, or information that may be edited, might be automatically placed on the larger public displays, while private information should be kept on smaller ones. I will not only compare automatic layouts to ones carefully generated by users, but will also collect subjective satisfaction ratings.

(Optional) Chapter: Attention and Human Memory

I have described my initial experiment exploring the effects of large displays on divided attention tasks. In this experiment, I explored how attention can be divided across multiple displays (both small and large), even when they are visually separated either by visual angle, depth, or physical barriers such as monitor bezels. For the set of tasks picked, results show reliable, though small, detrimental effects when information is separated within the visual field, but only when coupled with an offset in depth (i.e. separated between the small and large display). Surprisingly, physical discontinuities such as monitor bezels and even separation in depth alone do not seem to affect performance on the set of tasks tested.

In a separate project, I have done work in using large displays to provide ambient contextual images that serve as cues for human memory (Tan et al. 2001). In that work, I have shown a 56% improvement in memory for information when users had these ambient cues along with spatial distribution of the information. I hypothesize that the ambient cues were effective because they allowed users to situate themselves in our synthetic environments. However, I did not thoroughly explore whether the size of these ambient cues was important in immersing users in these environments, thus supporting this effect.

Hypothesis 5c: Large displays arouse users more, better capturing their attention as well as allowing them to remember more information presented on the displays.

I have begun to explore attention and human memory in these two separate but related pieces of work. I would next like to answer the question of whether large displays are able to arouse users to a greater degree, hence better capturing their attention and allowing them to remember more information presented on these displays. To quantify and understand this effect, I will present content on both a small and a large display. I will later test users for their ability to remember the information presented. I would hypothesize that content invoking a larger emotional response, especially when it is physically large, would be more memorable. In addition, a user who feels more present within graphical or textual information may be able to more efficiently switch between such tasks. This would be useful, for example, in situations when a user has to monitor and control many tasks at once.

Chapter 6: Subjective Responses to Large Displays

For each of my initial experiments, I have collected preference data, most of which indicates that users prefer working on large displays over smaller ones. Although my thesis will be largely focused on cognitive and social reactions to large displays and their implications on user productivity, I will also explore the subjective responses users have to these displays.

Hypothesis 6a: Users prefer performing tasks requiring overviews of large areas on large displays; Users prefer performing tasks requiring detailed work on small displays.

Based on anecdotal evidence, I believe that users prefer working on large displays only when performing certain types of tasks. Specifically, I will compare tasks that require users to focus on large areas of the display to tasks that require more detailed views. An example of a task that requires focusing on a large area is spray painting a building in a paint program. A more detailed task might include painting in the windows along with their frames in the exact locations. Users working on smaller displays can move their heads and cause the image presented to grow or shrink much more easily. Conversely, physical movement of users on large display cause little change in their view of the information. I believe that users will prefer to perform overview tasks on large displays, but tasks that require more detailed work on small displays.

The setup for this experiment will consist of a system in which the user is always presented identical information on a small display and a large display. Users will be asked to perform a variety of tasks and can use whichever display they feel is most likely to help them complete the task efficiently. By observing their choice of displays and well as their physical movements while using these displays, I will be able to infer preference data for working on the various tasks. Subsequent studies will compare performance of the various tasks on each of the displays.

Hypothesis 6b: Even with user interaction, there exists a (possibly non-linear) function that describes viewing distance preferences as they relate to the size of displays.

Although there has been work done in media communications research on viewing distance preferences as they relate to the size of displays, most current studies have not included user interaction. I would like to understand subjective preference for viewing distance as it relates to the physical size of displays when users have to interact with content. This understanding will allow me to formulate viewing distance recommendations in the design of systems that comprise multiple and large displays. This experiment will entail having users move to the most comfortable position when interacting with displays of various sizes. Content presented will include text, images, as well as movies, on which users have to perform various actions. I will additionally test if the presence of sound affects preference for viewing distance.

Hypothesis 6c: Users find large displays more engaging and are more motivated to perform tasks on them.

Anecdotal evidence suggests that users working on large displays are more engaged in the tasks at hand and hence spend more time and effort working on them. Users that are motivated to do certain tasks should tend to be more productive. However, there is the danger that these observations have been made mainly of users who are new to working on large displays and are only more engaged because it is novel. Quantifying this effect is difficult and might entail a field study in which I observe users working both using traditional displays and then using large displays over an extended period of time.

Chapter 7: Lessons for System and Interface Designers

Eventually, I would like to build complex environments that comprise different types of displays. The Display Garden is a prototype of such an environment, but there are many questions to be answered before we can successfully leverage the advantages offered by each type of display and combine them in a coherent manner to form a unified display environment. As displays get larger and start filling a larger portion of the environment, it will also become important to consider how physical architecture can facilitate the construction of computer systems and workplaces.

In this final chapter, I will combine my experience and findings with results from other researchers to formulate design principles for building multiple display systems that include large displays. These principles will serve more as a guidebook than a recipe both for system designers as well as for researchers who would like to further this work.

CRITERIA FOR SUCCESS AND PROPOSED CONTRIBUTION

In this thesis, I will explore areas in which users experience significant benefits from using physically large displays. I have already shown spatial cognition, specifically mental rotations, to be one such area. I have also begun work in another, information privacy. Finally, I have worked on pragmatics of multiple display systems that include large displays. Successful completion of this thesis is contingent upon my finding other such areas and furthering the field, both in terms of our theoretical understanding of physically large displays, but also in terms of how we can use them to build effective systems. I expect my research to contribute:

- Fundamental understanding of physical size in the perception of and interaction with information, especially as it relates to subjective preference as well as cognitive and social benefits.
- New paradigms and metrics for quantifying the benefits of large displays. I expect that these new paradigms will then be applicable to larger classes of system and interface evaluation.
- Niche applications which benefit from the use of physically large displays, along with the design principles and interaction techniques that facilitate their use.
- A set of design principles for building systems and interfaces that use physically large displays.

EVALUATING THE RISKS

There exist several potential risks that might hinder the timely completion of my thesis work. The largest risk is that there are few significantly interesting effects associated with varying the physical size of displays. However, I have personally been working in the Display Garden environment for more than a year now and have anecdotal evidence suggesting several qualitative differences between standard desktop monitors and physically large displays. I have quantified some of these differences and will continue to design controlled studies that will help me understand the effects of physical size on the perception and action on information. Additionally, the Stage 3 Research Group, which comprises students from the Art, Drama, Design, and Computer Science departments at Carnegie Mellon University, as well as the Proffitt Perceptual Psychology Lab at the University of Virginia have been exposed to different display environments and continually provide stimulating discussion on the topic.

Another risk is that I will be unable to adequately reproduce and measure anecdotal evidence in a controlled experimental setting. However, working with psychologists Dennis Proffitt and Mary Czerwinski, I have already designed and run studies that apply existing paradigms in novel setups in order to quantify 'hard to measure' effects. I am confident that I can continue to creatively leverage our deep understanding of existing paradigms to devise novel experimental setups that will allow me to test and understand desired effects.

A final risk in my research approach is that I will create contrived or overly simple studies that bear little relationship to the complex problems that arise in real-world systems. In working with Mary Czerwinski and her Large Display User Experience group at Microsoft Research, I am constantly exposed to the process of transferring theoretical results to the development of software products. This has given me motivation and valuable experience in evaluating my results in light of the final end, real-world design principles and applications.

SCHEDULE

The proposed schedule for completion of this thesis is illustrated in Figure 15. Because of the timing of my departure for my internship at Microsoft Research, I did not get to propose before beginning a fairly sizable portion of my thesis work. At Microsoft Research, I continued to perform directly relevant thesis related experiments, as can be seen in the figure. The proposed completion time spans 15 months and reflects work already done.

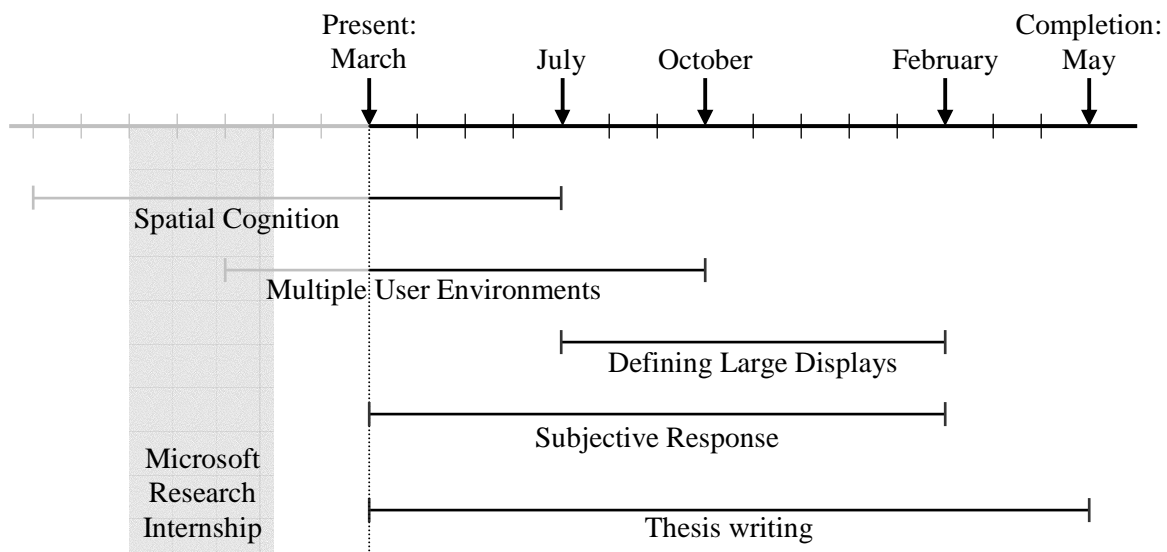


Figure 15: Proposed schedule for completion of my thesis.

CONCLUSIONS

Advances in display technology are moving us into a world in which displays are scattered throughout the environment. These displays range from very small personal displays to fairly large wall-sized displays. Most current research has focused on how content is displayed on traditional desktop displays. Recently, some attention has been shifted towards the change in representation required for much smaller personal displays. However, little work has been done to systematically explore the effects that physical size has on perception of information, especially for large displays.

In this proposal, I have discussed work I have done to show the benefits of using large displays. In fact, in one set of studies, I have shown that users on large displays are biased into egocentric strategies and perform better on spatial orientation tasks. In a separate study, I have shown that users treat information presented on large displays as being more public, and are hence more likely to read it. Finally, I have shown that wider separation of divided attention tasks across the visual field are detrimental to task performance, but only when coupled with an offset in depth introduced by coupling small and large displays.

I have also presented my general research approach as well as my tentative schedule and detailed agenda for proving my thesis, that:

Information elicits fundamentally different cognitive and social reactions when presented on large wall-sized displays as compared to smaller displays, even at identical visual angles. These reactions can be quantified and understood in controlled experiments and can be exploited to design display systems that make users more productive than they were on traditional systems.

ACKNOWLEDGEMENTS

I would like to thank my advisor Randy Pausch for his continuous guidance and supervision. Mary Czerwinski has also been instrumental in my development as a researcher. In her Large Display User Experience group at Microsoft Research, Mary has provided me with a research home away from home, as well as a much welcome escape from the harsh Pittsburgh winters. My other committee members, Scott Hudson and Jessica Hodgins, have given me invaluable advice in shaping this research project. I would also like to acknowledge Darren Gergle, Johnny Lee, Jeffrey Nichols, Peter Scupelli, and the Stage 3 Research Group at Carnegie Mellon University, Dennis Proffitt and his Perceptual Psychology Lab at the University of Virginia, as well as George Robertson and Eric Horvitz at Microsoft Research for their support and their stimulating discussions on multiple and large display systems.

REFERENCES

- Arthur, K.W. (2000). Effects of field of view on performance with head-mounted displays. *Dissertation Abstracts International*, 61(5B), 2614.
- Baudisch, P., Good, N., Bellotti, V., & Schraedley, P. (2002). Keeping things in context: A comparative evaluation of focus plus context screens, overviews, and zooming. *Proceedings of the CHI 2002 Conference on Human Factors and Computing Systems*, 259-266.
- Bishop, G., & Welch, G. (2000). Working in the Office of the "Real Soon Now". *IEEE Computer Graphics and Applications*, 20(4), 76-78.
- Buxton, W. (2001). Less is More (More or Less). In P. Denning (Ed.), *The Invisible Future: The seamless integration of technology in everyday life*. New York: McGraw Hill, 145-179.
- Bystrom, K.E., Barfield, W., & Hendrix, C. (1999). A conceptual model of the sense of presence in virtual environments. *Presence: Teleoperators and Virtual Environments*, 8(2), 241-244.
- Campbell, B.G. (1976). *Humankind Emerging*. Boston, MA: Little, Brown, & Co.

- Cardelli, L., & Beckman, K. The ClearDesk Project(or), 1990-1997. Retrieved October 22, 2002, from <http://www.luca.demon.co.uk/ClearDesk/ClearDesk.html>.
- Carpenter, M., & Proffitt, D. (2001). Comparing viewer and array mental rotations in different planes. *Memory & Cognition*, 29, 441-448.
- Chapanis, A., & Scarpa, L.C. (1967). Readability of dials at difference distances with constant viewing angle. *Human Factors*, 9(5), 419-426.
- Childs, I. (1988). HDTV-putting you in the picture. *IEE Review*, 34(7), 261-265.
- Chou, P., Gruteser, M., Lai, J., Levas, A., McFaddin, S., Pinhanez, C., Viveros, M., Wong, D., & Yoshihama, S. (2001). BlueSpace: Creating a personalized and context-aware workspace. *IBM technical report*, RC22281.
- Czerwinski, M., Tan, D.S., & Robertson, G.G. (2002). Women take a wider view. *Proceedings of the CHI 2002 Conference on Human Factors and Computing Systems*, 195-202.
- Detenber, B.H., & Reeves, B. (1996). A bio-informational theory of emotion: Motion and image size effects in viewers. *Journal of Communication*, 46(3), 66-84.
- Dillion, A. (1992). Reading from paper versus screens: A critical review of the empirical literature. *Ergonomics*, 35(10), 1297-1326.
- Dixon, M.W. (1999). Compensation for differences in perceived proportions between large and small displays. *Dissertation Abstracts International*, 60(10B), 5238.
- Dudfield, H.J., Macklin, C., Fearnley, R., Simpson, A., & Hall, P. (2001). Big is better? Human factors issues of large screen displays with military command teams. *Proceedings of People in Control 2001*, 304-309.
- Educational Testing Service. (1994). *Practicing to take the GRE test*. New Jersey: ETS.
- Elrod, S., Bruce, R., Gold, R., Goldberg, D., Halasz, F., Janssen, W., Lee, D., McCall, K., Pederson, E., Pier, K., Tang, J., & Welch, B. (1992). Liveboard: A large interactive display supporting group meetings, presentations and remote collaboration, *Proceedings of the CHI 1992 Conference on Human Factors and Computing Systems*, 599-607.
- Fish, R.S., & Judd, T.H. (1991). A subjective visual quality comparison of NTSC, VHS, and compressed DS1-compatible video. *Proceedings of the Society for Information Display*, 3212, 157-164.
- Grabe, M.E., Lombard, M., Reich, R.D., Bracken, C.C., & Ditton, T.B. (1999). The role of screen size in viewer experiences of media content. *Visual Communication Quarterly*, 6, 4-9.
- Guilford, J.P., & Zimmerman, W.S. (1948). The Guilford-Zimmerman Aptitude Survey. *Journal of Applied Psychology*, 32, 24-34.
- Guimbretière, F. (2002). Fluid interaction for high resolution wall-size displays. *Dissertation Abstracts International*, 63(1B), 349.
- Infield, S.E. (1991). An investigation into the relationship between navigation skill and spatial abilities. *Dissertation Abstracts International*, 52(5-B), 2800.
- Ishii, H., & Ullmer, B. (1997). Tangible bits: Towards seamless interfaces between people, bits and atoms. *Proceedings of the CHI 1997 Conference on Human Factors in Computing Systems*, 234-241.
- Jackson, L.A. (1992) *Physical appearance and gender*. Albany, NY: State University of New York Press.
- Josephs, R.A, Geisler, R.B., & Silvera, D.H. (1994). Judgment by quantity. *Journal of Experimental Psychology*, 123, 21-32.

- Kucera, H., & Francis, W.N. (1967). *Computational analysis of present-day American English*, Providence, RI: Brown University Press.
- Levelt, W.J.M. (1999). Perspective taking and ellipsis in spatial descriptions. In P. Bloom, M.A. Peterson, L. Nadel, & M.F. Garrett (Eds.), *Language and Space*. Cambridge, MA: MIT Press, 77-107.
- Lin, J.J., Duh, H.B.L., Parker, D.E., Abi-Rached, H., & Furness, T.A. (2002). Effects of field of view on presence, enjoyment, memory, and simulator sickness in a virtual environment. *Proceedings of IEEE Virtual Reality Conference 2002*, 164-171.
- Lund, A.M. (1993). The influence of video image size and resolution on viewing-distance preferences. *Society of Motion Picture and Television Engineers Journal*, 102, 406-415.
- MacIntyre, B., Mynatt, E.D., Volda, S., Hansen, K.M., Tullio, J., & Corso, G.M. (2001). Support for multitasking and background awareness using interactive peripheral displays, *Proceedings of the 2001 Symposium on User Interface Software and Technology*, 41-50.
- Mills, C.B., & Weldon, L.J. (1987). Reading text from computer screens. *ACM Computer Surveys*, 19(4), 329-357.
- Newsweek. (1992). 9 March, 8.
- Patrick, E., Cosgrove, D., Slavkovic, A., Rode, J.A., Verratti, T., & Chiselko, G. (2000). Using a large projection screen as an alternative to head-mounted displays for virtual environments. *Proceedings of the CHI 2000 Conference on Human Factors and Computing Systems*, 478-485.
- Prothero, J.D., & Hoffman, H.D. (1995). Widening the field of view increases the sense of presence within immersive virtual environments. *Human Interface Technology Laboratory Technical Report*, University of Washington, Seattle, WA, R-95-4.
- Raskar, R., Welch, G., Cutts, M., Lake, A., Stesin, L., & Fuchs, H. (1998). The office of the future: A unified approach to image-based modeling and spatially immersive displays. *Proceedings of SIGGRAPH 1998 International Conference on Computer Graphics and Interactive Techniques*, 179-188.
- Reeves, B., & Nass, C. (1999). *The Media Equation: How people treat computers, television, and new media like real people and places*. CLSI Publications: Stanford, CA.
- Schacter, D.L. (1987). Implicit Memory: History and Current Status. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13(3), 501-518.
- Shepard, R.N., & Metzler, J. (1971). Mental rotations of three-dimensional objects. *Science*, 171(3972), 701-703.
- Silvera, D.H., Josephs, R.A., & Giesler, R.B. (2002). Bigger is Better: The influence of physical size on aesthetic preference judgments. *Journal of Behavioral Decision Making*, 15, 189-202.
- Simmons, T. (2001). What's the optimum computer display size? *Ergonomics in Design*, Fall 2001, 19-25.
- Slater, M., & Usoh, M. (1993). Presence in Immersive Virtual Environments. *Proceedings of the IEEE Conference - Virtual Reality Annual International Symposium*, 90-96.
- Streitz, N.A., Geißler, J., Holmer, T., Konomi, S., Müller-Tomfelde, C., Reischl, W., Rexroth, P., Seitz, P., & Steinmetz, R. (1999). i-LAND: An interactive landscape for creativity and innovation. *Proceedings of the CHI 1999 Conference on Human Factors and Computing Systems*, 120-127.
- Suzuki, K., & Nakata, Y. (1988). Does the size of figures affect the rate of mental rotation? *Perception & Psychophysics*, 44(1), 76-80.
- Swaminathan, N., & Sato, S. (1997). Interaction design for large displays. *Interactions*, 4(1), 15-24.

- Tan D.S., & Czerwinski, M. (2003). Effects of Visual Separation and Physical Discontinuities when Distributing Information across Multiple Displays. *Submitted to INTERACT 2003 ninth IFIP International Conference on Human-Computer Interaction.*
- Tan, D.S., & Czerwinski, M. (2003). Information Voyeurism: Social Impact of Physically Large Displays on Information Privacy. *Short Paper at CHI 2003 Conference on Human Factors in Computing Systems.*
- Tan, D.S., Czerwinski, M., & Robertson, G.G. (2003). Women go with the (optical) flow. *Proceedings of the CHI 2003 Conference on Human Factors and Computing Systems.*
- Tan, D.S., Gergle, D., Scupelli, P.G., & Pausch, R. (2003). With similar visual angles, larger displays improve spatial performance. *Proceedings of the CHI 2003 Conference on Human Factors and Computing Systems.*
- Tan, D.S., & Pausch, R. (2002). Pre-emptive Shadows: Eliminating the blinding light from projectors. *Interactive poster at the CHI 2002 Conference on Human Factors in Computing Systems.*
- Tan, D.S., Stefanucci, J.K., Proffitt, D.R., & Pausch, R. (2001). The Infocockpit: Providing location and pace to aid human memory. *Workshop on Perceptive User Interfaces 2001.*
- Tani, M., Masato, H., Kimiya, Y., Koichiro, T., & Futakawa, M. (1994). Courtyard: Integrating shared overview on a large screen and per-user detail on individual screens. *Proceedings of the CHI 1994 Conference on Human Factors and Computing Systems*, 44-50.
- Tlauka, M. (2002). Switching imagined viewpoints: The effects of viewing angle and layout size. *British Journal of Psychology*, 93, 193-201.
- Wraga, M., Creem, S.H., & Proffitt, D.R. (2000). Updating displays after imagined object and viewer rotations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 26(1), 151-168.