Large Displays Enhance Optical Flow Cues and Narrow the Gender Gap in 3-D Virtual Navigation

Desney S. Tan, Mary P. Czerwinski, and **George G. Robertson,** Microsoft Research, Redmond, Washington

Objective: Existing reports suggest that males significantly outperform females in navigating 3-D virtual environments. Although researchers have recognized that this may be attributable to males and females possessing different spatial abilities, most work has attempted to reduce the gender gap by providing more training for females. In this paper, we explore using large displays to narrow the gender gap within these tasks. **Background:** While evaluating various interaction techniques, we found that large displays affording wider fields of view seemed to improve virtual navigation performance in general and, additionally, to narrow the gender gap that existed on standard desktop displays. **Method:** We conducted two experiments (32 and 22 participants) exploring the individual contributions of display and geometric fields of view to the observed effects as well as isolating factors explaining performance increases seen on the large displays. **Results:** We show that wider fields of view on large displays not only increase performance of all users on average but also benefit females to such a degree as to allow them to perform as well as males do. We further demonstrate that these benefits can be attributed to better optical flow cues offered by the large displays. **Conclusion:** These findings provide a significant contribution, including recommendations for the improved presentation of 3-D environments, backed by empirical data demonstrating performance benefits during navigation tasks. **Application.** Results can be used to design systems that narrow the gender gap in domains such as teleoperation and virtual environments for entertainment, virtual training, or information visualization.

INTRODUCTION

Three-dimensional virtual environments (VEs) are becoming a popular medium for applications such as training, modeling, and entertainment. Because most useful VEs encompass more space than can be viewed from a single vantage point, users have to be able to move around the environment in order to obtain different, optimal views of a scene. In fact, a 3-D world is only as useful as the user's ability to efficiently navigate and interact with the information within it. Users benefit greatly while navigating if they possess a cognitive map of the environment. This cognitive map allows them to stay oriented within the VE and to find information more easily.

Much of our work began with efforts to design and evaluate interaction techniques that would allow users not only to efficiently move around 3-D virtual environments but also to easily construct cognitive maps crucial to effective navigation (Tan, Robertson, & Czerwinski, 2001). In doing so, we found that wide fields of view afforded by large displays provided better optical flow cues that aided users in the formation of cognitive maps and improved 3-D virtual navigation. Interestingly, even though such cues improved performance for all users, females were aided significantly more than were males with these additional cues. Thus, even though both males and females performed better, gender differences that existed on standard desktop displays with narrow fields of view were drastically reduced when users worked on displays affording wider fields of view.

Address correspondence to Desney S. Tan, Microsoft Research, One Microsoft Way, Redmond, WA 98052; desney@microsoft.com. ***HUMAN FACTORS***, Vol. 48, No. 2, Summer 2006, pp. 318–333.

In this paper, we present several pilot studies and two formal experiments that we conducted to explore these effects. Although the pilot studies were aimed at evaluating the effectiveness of various interaction techniques for cognitive map formation, we found interesting results suggesting that large displays with wider fields of view benefited all users, especially females. Experiment 1 carefully explored the effects that field of view had on the performance of males and females. Experiment 2 provided evidence suggesting that it is the additional optical flow cues provided by wider displays that improves performance in general and that narrows the gender gap in cognitive map formation while navigating 3-D virtual environments. Results from this experiment also suggest that a 100° field of view is sufficiently wide, at least for the tasks explored in this study.

RELATED WORK

In our work, we extend existing results in 3-D navigation and cognitive map formation, especially as they relate to field of view and optical flow cues. Specifically, we examine how some of these factors influence the performance of males and females differently in computer-generated environments. Although many researchers have studied gender differences in navigating virtual environments, there is little research explicitly exploring gender biases induced by optical flow cues. One exception to this can be found in Cutmore, Hine, Maberly, Langford, & Hawgood (2000), whose work we describe in more detail later in this section.

Spatial Knowledge and 3-D Navigation

There exists a large body of work on general principles of 3-D wayfinding and navigation. For example, Thorndyke and Hayes-Roth (1982) studied the differences between spatial knowledge acquired from maps and from exploration. They defined several forms of navigational knowledge: landmark knowledge, or orientation using highly salient landmarks; route knowledge, or navigation from one landmark to the next; and survey knowledge, or navigation using broader bearings and a cognitive map of the environment. Their research demonstrated that when people learn new environments, they encode the information using one or more of these strategies. These principles have also been extensively studied in virtual environments (e.g., Czerwinski, Tan, & Robertson, 2002; Darken & Sibert, 1996; Richardson, Montello, & Hegarty, 1999; Tan et al., 2001). Hunt and Waller (1999) provided a review of research on orientation and wayfinding. They summarized work exploring the contribution of artifacts such as maps as well as the different strategies that people use to acquire spatial information. Additionally, they discussed individual differences, such as age and gender, which have been shown to be related to spatial abilities and navigation effectiveness.

Gender and Spatial Ability

Many researchers have examined the effects of gender on spatial ability. In fact, there exist many summaries of the known gender differences in spatial abilities and navigation strategies (Halpern, 2000; Kimura, 1999; Prestopnik & Roskos-Ewoldsen, 2000). Most reports document male advantages in spatial tasks. For example, Devlin and Bernstein (1995) presented results indicating that males made significantly fewer errors and were significantly more confident in finding their way around in a computer-simulated campus tour. They also reported that males utilized visual-spatial wayfinding information more than did females in these tasks.

Some authors have argued that these differences may be a result of many years of evolution (e.g. Crook, Youngjohn, & Larabee, 1993; Kimura, 1999; Lawton, Charleston, & Zieles, 1996). Others have hypothesized that the differences may be attributed to shorter term experience with the tasks involved. For example, in their meta-analysis of cognitive gender differences, Voyer, Voyer, and Bryden (1995) found partial support for the decrease in magnitude of gender differences in recent years. These effects were studied over a period that was much too short for evolutionary explanations and are most reasonably attributed to the changing roles and experiences of both genders.

Recently, researchers have begun to study gender differences in utilizing visual cues while performing computer-based spatial tasks. Hubona and Shirah (2004) suggested developing “gender neutral” interfaces by either adding meaningful landmarks to decrease user reliance on spatial presence and mental rotation ability or by converting spatial information into textual content. Although we believe that these suggestions might reduce the gender gap in some scenarios, we are interested in exploring techniques that also increase performance for all groups of users. In his dissertation work, Waller (1999) suggested that the large gender differences observed when users acquired spatial knowledge in virtual environments could be partially attributed to females being less familiar and proficient with computer interfaces. As a result, researchers have attempted to narrow the gap by providing additional training for females. Such efforts have had limited success. In our work, we propose instead that some of these gender differences may be attributed to the use of different strategies by different genders and that the reduction of spatial cues when virtual environments are viewed through displays with narrow fields of view is detrimental, especially to females. Alleviating these display problems, and improving performance for all users, may be as easy as increasing the field of view with larger displays.

Field of View

When considering field of view (FOV), it is important to define precisely what characteristics are being referred to. There are two FOV angles that must be considered: the display field of view (DFOV) and the geometric field of view (GFOV). The DFOV is the physical angle subtended from the eye to the left and right edges of the display screen. For a 16-inch (40.64 cm) wide display placed 24 inches (60.96 cm) from the user's eyes, the DFOV is approximately 37°. This angle is dependent on the display width and the distance of the user from the screen. It can be manipulated only through physical adjustments to the setup (i.e., making the display width narrower or wider or moving the user nearer to or farther from the screen). Alternatively, the GFOV is the horizontal angle subtended from the virtual camera to the left and right sides of the viewing frustum, or the visible part of the virtual environment. This angle is under control of the virtual environment designer and can be adjusted by zooming the virtual camera in or out. For example, zooming the camera in shows less of the environment on the screen (though at a larger size) and hence reduces the GFOV.

In the real world, DFOV and GFOV can usually be treated as one and the same, since people cannot usually zoom in and out at will. Hence, most reported literature does not make a distinction between DFOV and GFOV. Because it is possible to manipulate DFOV and GFOV independently and therefore isolate their individual effects on performance when experimenting with virtual environments, we explicitly state which display characteristic we are referring to when it is not clear from the study context.

It has recently been reported that it is harmful to deviate from a 1:1 ratio of GFOV and DFOV (Draper, Viirre, Furness, & Gawron, 2001). Large deviations can cause either magnification or miniaturization of items in the virtual world, possibly leading to discrepancies between studies as well as contributing reliably to simulator sickness. Our findings demonstrate that this ratio is important but is not necessarily the variable most responsible for good performance on navigation tasks.

There has been much evidence that restricting the user's FOV leads to perceptual, visual, and motor decrements in various kinds of performance tasks (e.g., Alfano & Michel, 1990; Hosman & van der Haart, 1981; Patrick et al., 2000; Piantanida, Boman, Larimer, Gille, & Reed, 1992), though there is some debate about what FOV parameters are optimal in designing computing tasks. For example, Dolezal (1982) described the effects of restricting FOV to 12°, including disorientation, dizziness during rapid head movements, difficulty in tracking objects, and difficulty forming a cognitive map of unfamiliar places. He observed that hand-eye coordination is impaired in smaller FOV conditions and that there was greatly reduced ability to integrate visual information across successive views. Note that the inability to form a cognitive map of unfamiliar places coincides with the decrement in the overlap of visual information across successive views. Chambers (1982) reported that increasing the amount of peripheral information in cockpit displays by increasing the FOV (up to 90°) allowed users to construct an overlapping sequence of spatial map fixations in memory, which led to faster cognitive map construction.

In summary, it appears that wider FOVs provide more spatial cues to users and are important aids for many spatial tasks, helping especially with cognitive map construction when the visual complexity of a display or the demands of a task increase. However, we have found no reports in the literature suggesting that FOV restrictions are more or less harmful based on gender. In our work, we explore how FOV as well as the spatial cues changed by varying the FOV affect the specific performance of males and females.

Optical Flow

One of the cues affected by varying the FOV is optical flow. Gibson (1966, 1979) initiated a new field of psychological study called *ecological optics.* According to his theory, the pattern of light falling on the retina changes constantly as one moves around in the environment, producing optical flow in the optical array. This optical flow, coupled with proprioceptive and kinesthetic perception of motion, allows people to perceive not only the structure of the environment but also their movement within it (Duffy, 2000; Klatzky, Loomis, Beall, Chance, & Golledge, 1998). For example, the fixed point, or singularity, in the flow field specifies the observer's direction of self-motion. Warren, Morris, and Kalish (1988) provided an overview as well as competing theories on how people derive their "translational heading," or movement through space, from the information available through vision. They argue that humans rely heavily on optical flow cues for navigation in the real world.

Various researchers have claimed that given optical flow, only the central visual field is necessary for accurate judgments of heading and velocity (Atchley & Andersen, 1999; Crowell & Banks, 1993; Warren & Kurtz, 1992). However, because these researchers were interested mainly in showing that optical flow cues were not constant across the retina, most of the studies relied mainly on standard computer displays with relatively small FOVs and did not document the utility of peripheral optical flow cues in cognitive map formation tasks. However, Richman and Dyre (1999) used larger FOVs (up to 90°) to provide optical flow in both active navigation and passive viewing. Their results suggest that optical flow in the periphery benefits heading perception, particularly during active navigation. We further explore this to determine if optical flow presented with wider fields of view affects cognitive map formation in males and females differently during active navigation.

Loomis, Klatzky, Golledge, and Philbeck (1999) described path integration, the process of navigation by which travelers update their movement to provide a current estimate of position and orientation within a larger spatial framework or cognitive map. However, because they did not differentiate between discrete and continuous movement, they did not explore the effects of optical flow on the acquisition of survey knowledge. In their work, Kirschen, Kahana, Sekuler, and Burack (2000) found that in the absence of other cues, such as landmarks, optical flow was a significant aid in wayfinding. Users performed better learning a maze with fluid optical flow cues than one with choppy cues. Similarly, Kearns, Warren, Duchon, and Tarr (2002) found that although vestibular and proprioceptive cues seemed to dominate when present, path integration can be performed solely by integrating optical flow cues. Interestingly, they also found that females were less hurt from the lack of optical flow than were males. They hypothesized that this could be attributable to females relying on alternate strategies, such as timing or static information. Unfortunately, a high attrition rate for females suffering from simulator sickness may have filtered their female population for users that were intrinsically less visually dependant. As such, the authors could not draw explicit conclusions regarding the effect of optical flow on gender. In our studies, we embrace the notion of path integration and explore how the absence or presence of optical flow affects spatial learning in males and females.

In order to further explore the influence of factors such as gender, passive versus active navigation, cognitive style, optical flow, and brain hemisphere activation on the acquisition of route and survey knowledge in a 3-D virtual environment, Cutmore et al. (2000) performed a series of experiments using maze traversal tasks. They demonstrated significant gender differences in initial studies on navigation strategies and performance and hypothesized that optical flow cues might have been driving the results. In a final experiment, they focused on the hypothesis that spatial ability, not gender, was driving the effectiveness of optical flow cues. However, to maintain a homogeneous population while exploring this issue, the authors included only female participants in their sample. Thus, although they found that optical flow significantly benefited low spatial ability users in navigating their 3-D virtual environment, they could not report whether gender effects were also influential with these cues. In fact, we have found few reports in the literature explicitly exploring whether optical flow cues are more or less helpful based on gender. In our second experiment, we extend the paradigms used by Cutmore et al. (2000) to more carefully explore the gender effects of optical flow.

Our Preliminary Studies

Much of our current work grew out of surprising results we found while evaluating novel interaction techniques for navigating 3-D virtual environments (Tan et al., 2001). In the first of these studies, 4 out of 17 users tried our navigation techniques on an experimental large screen display that increased both DFOV as well as the GFOV over traditional desktop displays. Although the overall pattern of the data did not differ significantly with regard to the navigation conditions studied, we found that users performed trials about 30% faster, on average, when they worked on the large display as compared with the desktop display. Given the set of tasks performed, this result suggested that the wider FOVs offered by the large display were helping users better form and remember cognitive maps of the environment. Although it has been shown that large projection screens may be effective substitutes for immersive displays such as head-mounted displays (Patrick et al., 2000), we found little work done to quantify the differences between large semi-immersive displays and regular nonimmersive displays.

Hence, we explicitly manipulated display size as a factor of interest in a follow-up study with 13 users (6 females and 7 males). Results from this study indicate that the large display does indeed aid users and allows them to perform the navigation tasks more effectively. This was seen in the significant reduction not only in the overall trial times but also in the errors that users amassed as they navigated through the environment.

Perhaps the most interesting and unexpected finding from this study was the interaction we observed between gender and display size for overall trial times. Although the large display increased performance for all users on average, females improved so much so that the significant gender differences seen on the desktop display were drastically reduced on the large display that offered wider FOVs. We explore this effect more thoroughly in the experiment described in the next section. Additionally, because we coupled the display and geometric fields of view in this study, manipulating them in tandem, we could not conclude whether one or the other contributed more strongly to the observed gender effect. Therefore, we also explicitly manipulate display and geometric fields of view independently in the following experiment.

EXPERIMENT 1

Many reports of field of view effects in the literature do not distinguish between display field of view (DFOV), the angle subtended from the eye to the left and right edges of the display screen, and geometric field of view (GFOV), the horizontal angle subtended from the virtual camera to the sides of the viewing frustum. In fact, when we manipulated display size in our prior studies, we actually manipulated both DFOV and GFOV in tandem. In this experiment, we explore the individual contributions of each of these manipulations to the observed effects. Thus,

> Hypothesis 1a: The large display with its wider display field of view provides important navigation information in the optical periphery and allows users to perform better on navigation tasks.

> Hypothesis 1b: Wider geometric fields of view provide more useful information on the display and help the user perform navigation tasks more effectively.

Because prior reports assert that it is harmful to deviate from a 1:1 GFOV:DFOV ratio (Draper et al., 2001), and we know of no work that has specifically reported our observed gender effects in more traditional (geometric) field of view work, we claim

> Hypothesis 1c: It is the wider display coupled with the geometric fields of view that increases performance by females and narrows the gender gap that exists in other display conditions.

Materials and Procedure

We ran this experiment on a 450 MHz Pentium II Dell computer using Arcturus, a prototype display comprising two projectors that rear-project onto a semicurved, tinted Plexiglas surface (see Figure 1). With careful calibration, the seam between the two projections can be made arbitrarily

Figure 1. User working on experimental Arcturus display.

small, creating a virtually seamless 2048 × 768 pixel (8:3 aspect ratio) display surface. The display was about twice as wide as a regular monitor and provides a ~74° display field of view within its 36-inch (91.44-cm) frame. We controlled the display using standard Windows 2000 multimonitor support and provided users with a Microsoft Internet Keyboard Pro and an Intellimouse.

We evaluated two levels of display field of view (41°, or small, vs. 74°, or large) and two levels of geometric field of view (32.5°, or narrow, vs. 75°, or wide). We controlled the display as well as the geometric fields of view in software. For the small display field of view, we reduced the width of the display to 18 inches (45.72 cm) by setting the outer parts of the projection to be black. For the geometric field of view, we controlled the horizontal angles by zooming the virtual cameras accordingly. Each of the four conditions (DFOV: Small or Large × GFOV: Narrow or Wide) corresponded to the following DFOV:GFOV ratios: small-narrow = ~1:1, small-wide = ~1:2, large-narrow = ~2:1, and large-wide = ~1:1. We included these conditions to verify earlier published results that 1:1 ratios are more effective. The experiment was a within-subjects design, with each user performing all four conditions in an order that was fully counterbalanced across users.

The software setup and procedure for each trial in the main task was similar to that used in our preliminary studies. We used the Alice 3-D authoring system (Carnegie Mellon University, n.d.) to create two 3-D virtual worlds. The first world, which we refer to as the *tutorial world,* was 300 × 300 m large and contained 4 objects for navigation and manipulation purposes. The second world, the *experimental world,* was 500 × 500 m large and contained 23 objects, most of which consisted of carnival-themed structures, such as tents, roller coasters, and rides (see Figure 2). Each world contained cubes and drop pads that were dual coded to match each other via color and numeric coding. The user's task was to select each cube, numbered on only one of its faces, and move it to the matching numbered drop pad. The tutorial world contained only one cube and two pads for each trial. The tutorial consisted of the user finding the cube and moving it to its corresponding pad once for each of the five navigation conditions. In the experimental world, there were four cubes and four pads for each trial. The user had to find each cube, in any order, and move it to its respective drop pad. Each trial was completed once each of the four cubes was placed on its respective drop pads.

We utilized a navigation technique called *speed-coupled flying* (Tan et al., 2001). With this technique, dragging the mouse forward/backward moves the camera forward/backward, and dragging the mouse left/right turns the camera left/right. The farther the user drags in a particular direction, the faster the camera moves. Additionally, the user's forward speed is coupled to the

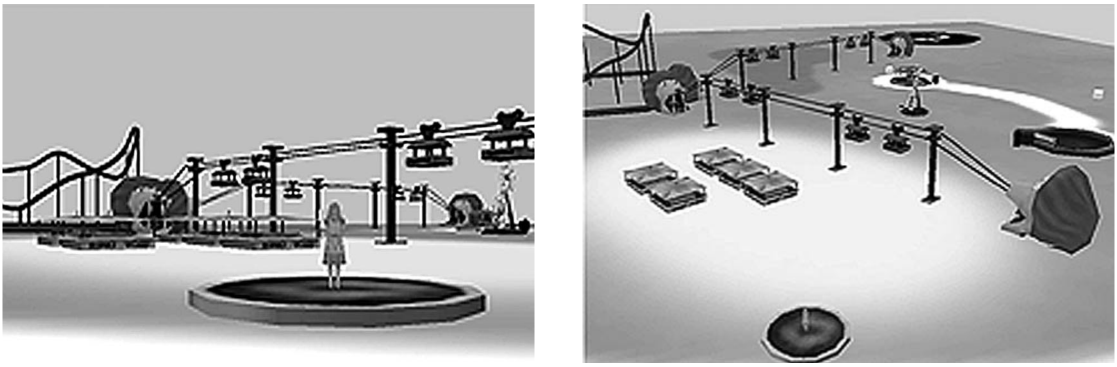

Figure 2. Left: Local view of carnival scene used in the first two experiments. Right: Overview of scene.

camera's viewing height and angle. The faster the user moves forward, the higher the camera moves, getting a zoomed-out overview of the environment. This gives the user the ability to transition seamlessly between local and global views of the virtual environment.

Prior to beginning the experiment, all users completed the Map Memory (MV2 and MV3) subtests of the Kit of Factor-Referenced Cognitive Tests (Eckstrom, French, Harman, & Dermen, 1976). These subtests evaluate users' abilities to remember not only the position of objects on a street map (MV2) but also parts of a map such that they could recognize them again (MV3), two skills critical to our navigation task and dependent measures.

After each condition, users performed a pointing task. In this task, we removed two objects and two drop pads from each of the worlds. First, we showed the user an object for 5 s. We then placed users in the world and asked them to turn and point to the location where they thought the object had previously resided. To do this, users dragged the mouse to move a virtual pointer to the proper position, just as if they were using their pointer finger in the real world. Users pointed from three positions, each 60° away from the others, for each of the objects. Because the drop pads were on the outskirts of the world, users only did this once for each pad. We measured performance errors as the distance between the closest part of the object and the projected pointing ray. When all trials were completed, the user filled out a final preference survey.

We recorded the following dependent measures on the user's computer: overall task time, travel distance, travel height while traveling ("flying" is a more efficient travel mode), and user satisfaction for each condition as well as overall preference. We also collected the error measures for each of the pointing tasks.

Participants

Thirty-two intermediate to experienced computer users (17 women, 15 men) from the greater Puget Sound area volunteered to participate in the experiment. None of the users had participated in previous, related studies, and all played less than 1 hr of 3-D games per week. The average age was 41 years (38.6 for men and 43.1 for women), and ages ranged from 19 to 60 years. Each session took about 2 hr, and users were given a software gratuity for participating.

Results and Discussion

Map memory. We submitted scores on the two map memory cognitive subtests to a paired t test, assuming unequal variances, comparing scores of men and women. We obtained no significant difference based on gender on this measure, $t(29) = -0.29$, $p = .77$. Hence, we assert that observed effects cannot be simply attributed to specific map memory abilities measured by this instrument.

Performance data. We found no significant effects in the percentage correct data so we removed it from further analysis. We submitted the average trial time, pointing error, distance traveled, and height of travel to a 2 (gender) × 2 (DFOV) × 2 (GFOV) repeated measures multivariate analysis of variance (RM-MANOVA). We tested gender between subjects and all other variables within subjects.

We observed significant main effects of DFOV, $F(4, 27) = 15.5, p < .001$, and GFOV, $F(4, 27) = 4.2, p = .009$. On average, the large DFOV condition resulted in improved performance, as evidenced by less pointing error (15.4 vs. 14.8 m error for the small vs. large DFOV, respectively), greater distance traveled (5461 vs. 6918 m on average), more flying (average heights of 14.9 vs. 15.5 m), and faster trial times (213.5 vs. 205 s). The wide GFOV had a similarly beneficial effect on average performance for pointing error (15.3 vs. 14.8 m error for the narrow vs. wide GFOV), distance traveled (6601.7 vs. 5777.4 m), travel height (14.6 vs. 15.8 m), and trial time (218.7 vs. 199.85 s).

The between-subjects tests for gender reached significance for trial time, $F(1, 30) = 5.9, p = .02$, and borderline significance for travel height, $F(1, 30) = 3.3, p = .07$. In both cases, men were faster than women in their average trial times (192.9 vs. 225.5 s) and had higher average travel heights (16.5 vs. 13.8 m).

Additionally, we found a borderline significant interaction between gender and GFOV, $F(4, 27) = 2.3, p = .08$. Across three of the measures (trial time, travel height, and pointing error), women benefited more than men from the wider GFOV, but the interaction reached significance only for the distance traveled metric. Figure 3 demonstrates how wider GFOV brings out markedly different strategies between men and women. Women traveled less distance (concurrent with shorter trial times) in the wide GFOV, whereas men traveled farther (also with shorter trial times) in this condition. Both genders "flew" higher in the wide GFOV condition. Why men flew further distances remains unclear from these data and will be the subject of further experiments that will more closely examine 3-D navigation strategies. However, it should be noted that the flight height data reveal that something about female navigation strategies was supported by the wide GFOV condition, allowing for shorter travel distances and faster travel methods. We found no other significant interactions, including that of DFOV and GFOV.

Motivated by the results from our initial studies, we used a planned comparison analysis to determine whether or not there was a significant difference between men and women in the Large DFOV × Wide GFOV condition for trial times. The difference between men and women was not significant at the $p = .05$ level, $t(28) = -1.32, p = .19$, indicating that the gender gap had been significantly reduced in this condition. Because the Large DFOV × Wide GFOV condition in this experiment used the same display parameters as did our preliminary studies, we consider this result a replication of those findings. Figures 4 and 5 show trial time data as well as the differences (female-male difference) between men and women under the various conditions. The difference graph shows the reduction in gender bias in the Large DFOV × Wide GFOV condition.

User satisfaction. At the end of the session, we asked users which condition provided more information for performing the tasks. Eighteen users (9 women and 9 men) chose the Large DFOV × Wide GFOV condition, followed by 8 (5 women and 3 men) choosing the Small DFOV × Wide GFOV. In other words, 12 out of 15 men and 14 out of 17 women chose the wide GFOV condition

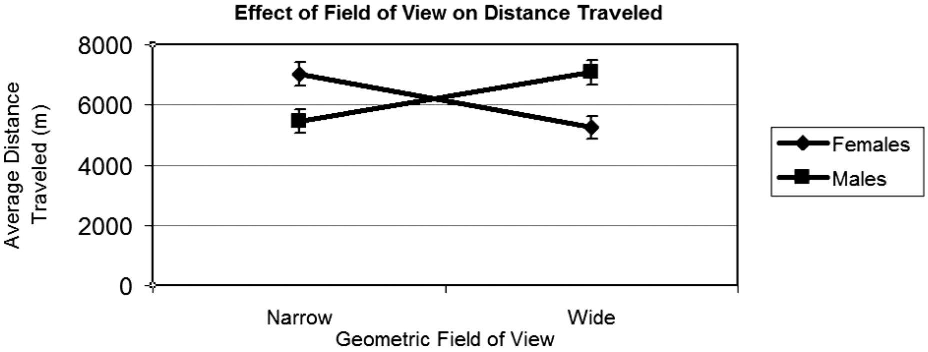

Figure 3. Wider fields of view bring out different strategy differences for males and females. Error bars represent standard error.

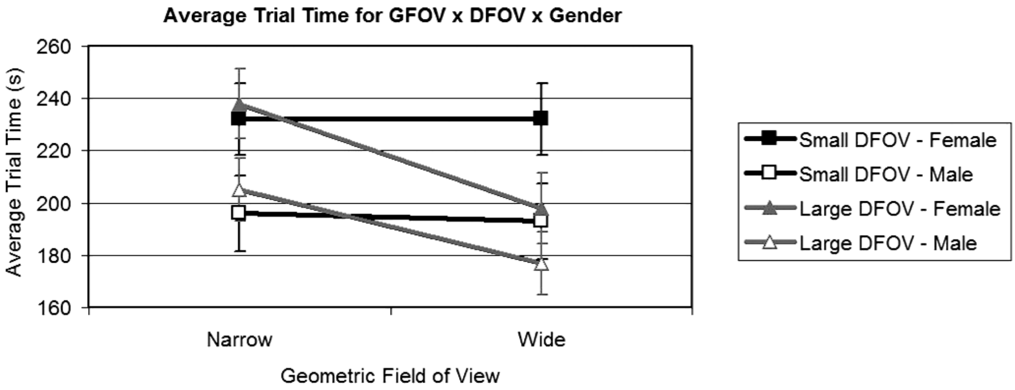

Figure 4. Users performed better in the wide field of view condition, but only with the large display. Error bars represent standard error.

as providing the most information, both significant results by binomial tests.

Summary

Results from this experiment supported our first two hypotheses, that both DFOV and GFOV increased navigation performance. Results also revealed the typical overall male superiority in navigating within a 3-D virtual environment, both for travel times and travel height. That is, when conditions were collapsed, men performed significantly better than did women. The experiment further revealed opposing gender strategies for dealing with wider GFOV, with women choosing to navigate shorter distances in those conditions than did men.

Although all users seem to benefit from the large DFOV and the wide GFOV, women improved so much as to narrow the gender gap, especially in terms of travel times in the Large DFOV × Wide GFOV condition. Hence it appears as if something about the combination of a wider display field of view and a wider geometric field of view is causing this effect. It should also be noted that the 1:1 DFOV:GFOV ratio in the Large DFOV × Wide GFOV condition cannot account for this finding, given that our Small DFOV × Narrow GFOV condition also provides a 1:1 DFOV: GFOV ratio.

From these results, we hypothesized that wider display and geometric fields of view allow better tracking of environmental information and spatial orientation via head/eye movements, offloading the cognitive map development task to the perceptual system. As this is typically an easier cognitive task for males than for females, females

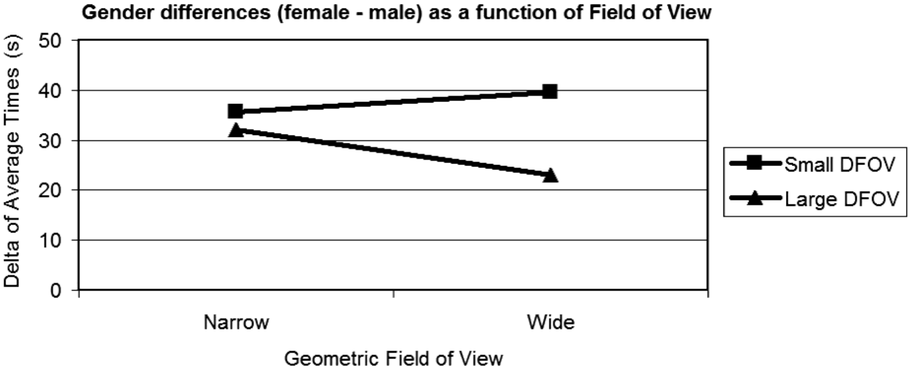

Figure 5. Results show a reduction in the gender gap in the large Display × Wide Field of View condition. Because error bars on repeated-measure independent variables are not relevant for inferences about differences, we do not show error bars on this graph.

may benefit more from the wide GFOV conditions, at least on large displays. The good news is that these benefits come without a concomitant decrement in male performance. In Experiment 2, we focus specifically on why wider fields of view on large displays were helping female participants to navigate more efficiently.

EXPERIMENT 2

In this experiment, we explore optical flow cues as one of the factors differentiating large displays from smaller displays and as being at least partially responsible for closing the gender gap that exists on smaller displays (Tan, Czerwinski, & Robertson, 2003). As reported in related work, Cutmore et al. (2000) reported a significant benefit for smooth optical flow cues for participants with lower spatial abilities, but gender was not a variable included in their final experiment examining those cues. We chose to replicate Cutmore et al.'s (2000) experimental paradigm and to explore the gender variable explicitly. If benefits from optical flow exist, removing optical flow cues from the virtual worlds should return the gender gap in 3-D navigation performance, even on large displays. Hence,

> Hypothesis 2a: Additional optical flow cues offered on large displays benefit female users more than male users and narrow the gender gap that exists on traditional small displays.

Additionally, we wanted to see if we could exaggerate any effects observed by making the display even larger and, possibly, continue to increase performance for all users. We were also curious to see if female participants would continue their improvements and perhaps start to outperform male participants. Thus,

> Hypothesis 2b: Larger displays and wider fields of view provide even better cues and continue to increase performance for all users, especially female users.

Task and Procedure

Because it would have been difficult to examine the presence of optical flow cues with the previous tasks used, we decided to create a slightly more constrained navigation and map memory task. Hence, we designed our tasks, derived from Cutmore et al. (2000), to examine not only the absence or presence of optical flow cues while navigating but also the optimal field of view for active navigation.

We used the Alice 3D authoring system to construct a 3-D virtual maze. The user was positioned to start at an interior room position and could then make constrained movements (controlled by pressing the right, up, or left arrow keys) in order to find the exit from the maze. In each room, users always saw three doors through which they could travel (see Figure 6). There were always exactly

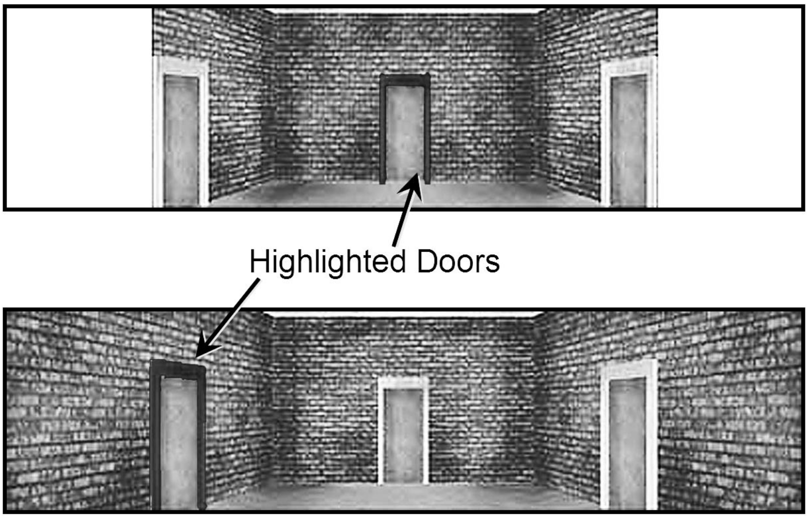

Figure 6. User views of the maze with narrow (top) and wide (bottom) fields of view. The user should follow the center door above and the left one below.

three turn options. Each correct turn resulted in a door being raised and the user being moved to the middle of the next room.

When optical flow was present, the user saw animated movement of the virtual camera from the center of the current room through the door to the center of the next room. When optical flow was absent, the user simply saw instantaneous transportation from the center of the current room to the center of the next room (no animated movement). We controlled the time required for each move between rooms and kept transitions constant at 2.5 s regardless of optical flow condition.

Each path through the maze involved a randomly selected path through 14 rooms. There were exactly eight turns (left or right) and six straight movements in each path. Additionally, paths were allowed to cross back over themselves. Several example paths through the maze are provided in Figure 7.

Because we were interested in examining the effects of optical flow cues, and not of the traversal of sequences of distinct landmarks, we created the rooms with a brick texture that afforded maximal optical flow cues but removed distinct landmarks. Thus, learning routes through the maze required users to create some internal representation of the layout. There were several ways users could have encoded this path through the maze. First, users could encode their actions into a sequence, forming a symbolic representation of some kind. For example, a sequence may be of the form "up, left, right, left, up, left, left…" (more easily, "u, l, r, l, u, l, l…") or even made up of the mapping of the directions to the door numbers, such as "2, 1, 3, 1, 2, 1, 1…." Alternatively, users could generate a spatial representation of the environment, as we instructed them to do in this experiment. With this strategy, users would employ some kind of path integration to form a cognitive map of the environment. In order to encourage users to use the latter strategy, we implemented an additional test performed after each move. For each room, we asked users to determine if they had previously been in that room (i.e., if the path had crossed itself). Only by building a cognitive map of the space could they have accomplished this task correctly.

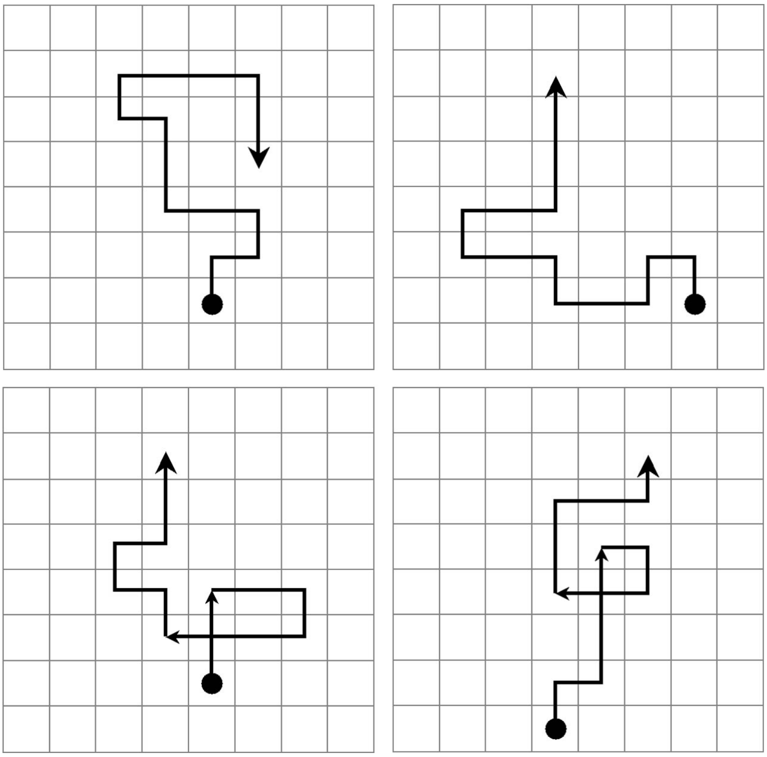

Figure 7. Example paths through the maze. Users never saw these overview maps.

Each user completed five trials. The first was a practice trial representative of the four experimental trials. Each trial in the experiment consisted of three phases: learning, forward test, and backward test. During the learning phase of a trial, the user navigated through the maze while the computer displayed (via green highlighting of the door frame, as shown in Figure 6) which turn (key press) to take. If a user hit the wrong key, the incorrectly chosen door would flash red and the user would not move until the correct direction was chosen. During the forward and backward test phases, the green highlighting was removed and the user had to remember how to navigate the maze without the turn cues. In the forward test they were placed at the start and would navigate in the same direction in which they had learned the path, and in the backward test they were placed at the end and had to reverse the path to find the start. Again, if the user attempted an incorrect turn, that door flashed red, after which the user could choose a different turn. The system kept track of the number of doors correctly and incorrectly opened in addition to the time it took to complete the learning and two test phases of each trial. After each trial, the user provided a satisfaction rating for that set of optical flow and field of view settings on a scale of 1 (*frustrated*) to 5 (*satisfied*). There was one trial for each combination of display settings: Display Size (100° vs. 120°) × Optical Flow (absent vs. present). Note that we opted to keep a 1:1 DFOV:GFOV ratio and that display size represents a combined display and geometric field of view manipulation for simplicity. After experiencing all of the display conditions, users were allowed to alter their satisfaction ratings to better reflect their overall preference for the display settings.

Prior to the start of the practice trial, users performed Parts 1 and 2 of the VZ2 "paper folding" subtest from the Eckstrom et al. (1976) Kit of Factor-Referenced Cognitive Tests. This test has been used to evaluate spatial ability skills and has been widely validated and used in 3-D navigation studies. Although each part of the test takes only 3 min, administration of the test with instructions took approximately 10 to 15 min.

Equipment and Design

In this experiment, we used Dsharp, a novel 43-inch (109.22-cm) wide display, created by rear-projecting three displays onto a curved Plexiglas panel (see Figure 8). We used the Windows XP™ multiple monitor software to "stitch" the three desktops into one large, curved, display surface. Each projector displayed at a resolution of 1024 × 768, for an equivalent of a 3072 × 768 resolution display. The straight-line distance from left edge to right edge of the display area is 43 inches (109.22 cm). The actual distance along the curve is 46.5 inches (118.11 cm) at the top and 47 inches (119.38 cm) at the bottom. The height of the

Figure 8. User working on the experimental Dsharp display.

display is 11 inches (27.94 cm). The distance from eye to screen is 20 inches (50.8 cm) in the center and 24 inches (60.96 cm) at the edges. The display field of view for a user seated in this position is about 120°. The system, an 800 MHz Pentium III Dell computer, maintained a frame rate of about 45 frames/s in all conditions. We used a Microsoft natural keyboard, allowing only the arrow keys and the spacebar for input.

We used a 2 (gender) × 2 (display size: small vs. large) × 2 (optical flow: absent vs. present) × 2 (test direction: forward vs. backward) design (all within subjects except for the gender variable, thereby a "mixed design"). Again, we manipulated the DFOV and GFOV in tandem with display size. In fact, both were about 100° on the small display and about 120° on the large display. We balanced display size and optical flow conditions using a Latin square design, and the order of tests was fully counterbalanced. Dependent measures included spatial abilities scores, overall task times, number of doors opened correctly on the first try, and user satisfaction ratings for each condition.

Participants

Twenty-two intermediate to expert Windows users (11 female and 11 male) from the greater Puget Sound area volunteered to participate in this experiment. All users played fewer than 5 hr of 3-D video games per week. The average age was 38.5 years (for the female participants the average was 36.7 years, for the males 40.2), and ages ranged from 13 to 50 years. The entire session lasted approximately 1.5 hr, and the participants were provided with a software gratuity for their participation.

Results

Spatial abilities. We performed a split-mean division of the paper-folding test data so that any score higher than 11.8 (the average score for all of our users) was labeled as "high spatial ability" and any score below that was labeled "low spatial ability." We observed no main effects or interactions for this measure. Though the male participants did score slightly higher than the female participants, with average scores of 12.0 vs. 11.6, respectively, the difference was not statistically significant. Thus, we assert that effects observed from this study cannot be simply attributed to standard measures of spatial ability, but rather to some other difference between genders.

Overall MANOVA. We submitted the data to a 2 (gender) × 2 (spatial ability: low vs. high) × 2 (test direction: forward vs. backward) × 2 (display size: small vs. large) × 2 (optical flow: absent vs. present) RM-MANOVA. The first two variables were between subjects and the rest were within sbjects. The two dependent measures submitted to the analysis were task reaction time and the number of doors chosen correctly on the first attempt. We discuss each dependent measure separately in terms of main effects and interaction with other variables.

Task times. We observed a main effect of test direction for overall task time, $F(1, 18) = 11.5$, $p = .003$, with average task times in the forward test significantly faster than those in the backward test (77.9 vs. 85.8 s for forward vs. backward, respectively). In addition, we observed a significant main effect for the optical flow manipulation, $F(1, 18) = 15.22$, $p = .001$. Having optical flow cues present during the 3-D maze navigation task significantly shortened average maze traversal times (86.9 s without optical flow vs. 76.8 s with optical flow). We found a borderline significant interaction between direction and optical flow, $F(1, 18) = 12.6$, $p = .066$, with optical flow benefiting users more in the forward direction. There was also a significant three-way interaction among gender, optical flow, and test direction, $F(1, 8) = 12.63$, $p = .002$. Follow-up post hoc analyses revealed a significant Gender × Optical Flow interaction in the forward, $F(1, 20) = 8.20$, $p = .01$, but not the backward direction. Female participants benefited significantly more from the optical flow cues, but only in the forward condition. In fact, in the forward direction, male participants significantly outperformed the female participants in the absence of optical flow cues, $t(42) = 3.29$, $p = .002$, but not when optical flow cues were present, $t(42) = 0.42$, $p = .67$. These data are shown in Figure 9.

Number of correct doors opened on first attempt. We observed a significant main effect of direction for the number of correct doors opened on the first attempt, $F(1, 18) = 11.5$, $p = .003$, with the forward direction resulting in more correct turn choices, on average (8.6 vs. 7.5 for forward vs. backward, respectively). No other main effects or interactions were significant at the $p = .05$ level for this measure.

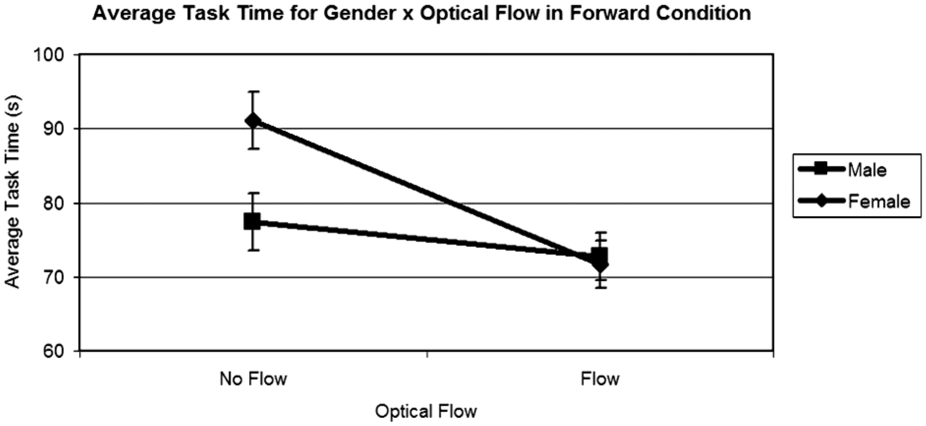

Figure 9. In the forward condition, although males significantly outperformed females when there was no optical flow, there was no significant difference between the two groups when optical flow was present. Error bars represent standard error.

User satisfaction. We used a 2 (gender) × 2 (display size) × 2 (optical flow) analysis of variance to analyze the user satisfaction ratings for each display condition. We obtained a significant main effect of optical flow, $F(1, 20) = 6.7, p = .017$, with conditions incorporating optical flow cues being rated as more satisfactory by users, on average. In addition, we observed a significant interaction between gender and optical flow, $F(1, 20) = 5.6, p = .025$, with female participants rating conditions with optical flow significantly more highly than did the male participants, relative to no-flow conditions.

Summary

Although there was no significant effect of optical flow on number of correct doors opened on the first attempt, users were able to recall mazes significantly faster with optical flow present. After piloting the experiment, we picked the number of rooms (14) and turns (eight) to tax working and long-term memory. From debriefings with the users, we observed that given a particular strategy, users performed as well as they were able. Task time, or the time users took to recall and navigate the paths, is therefore a good indicator of how well the spatial information was encoded and retrieved during the test.

One of the reasons we provided the forward and backward tests was to distinguish between users who might have encoded the paths sequentially as opposed to spatially. On one hand, we expected that the former group would have more trouble flipping the path backward than would the latter. As an analogy, most people can articulate the alphabet from A to Z effortlessly but find it extremely difficult to recite it from Z to A. This is, presumably, because the alphabet is stored in a unidirectional fashion. On the other hand, users who utilize a spatial encoding and form a cognitive map of the environment may have an advantage in backward navigation because only one mental rotation is involved. It should be noted that although users may dominantly use one or the other, these encoding schemes are not mutually exclusive. We did not, however, observe the effect we expected. Users were able to more quickly and more accurately recall paths going in the forward direction than in the backward direction. This might be explained by the fact that regardless of encoding strategy, navigating backward required users to perform an extra cognitive step in reversing the path. This added cognitive load may account for the slower and less accurate performance on the backward test, regardless of the display condition. However, it should be noted that the display size and optical flow manipulations did not differentially affect the performance of male and female participants on the backward test.

That being said, we would expect that the presence of optical flow would help users strengthen their spatial encoding of the paths. This hypothesis is supported in the result that optical flow helped the female participants more than the male participants on the forward test, because female participants have been reported to be "landmark" navigators on average (Sandstrom, Kaufman, & Huettel, 1998). User satisfaction ratings support these performance results in that all users significantly preferred having optical flow cues present, and the female participants rated them as significantly more satisfying than did the male participants. In addition, given that the paper-folding test revealed no significant differences in spatial ability between our gender samples, we assert that the effects reported here are not likely related to spatial ability, as measured by this particular test. Of course this conjecture requires replication, most notably with a sample population showing the more typical gender gap in spatial abilities.

Previous studies have demonstrated performance advantages for larger fields of view. In our studies, we were hoping to learn more about the limits on increasing fields of view. The fact that there was no reliable performance difference between the small (100°) and large (120°) display conditions indicates that there may be no advantage to increasing field of view beyond 100° for the particular class of navigation tasks we examined. This was contrary to our initial hypothesis, that performance would continue to increase as the screen became larger.

Taken together, the experiments suggest that large screen displays offering wider fields of view provide important navigation information in the optical periphery and allow users to perform better in certain 3-D virtual navigation tasks. We assert that much of this additional information comes from optical flow cues, which are particularly important in tasks requiring cognitive integration of various viewpoints, such as forming and remembering cognitive maps while navigating virtual environments. Furthermore, and perhaps more interestingly, although larger fields of view benefit all users on average, female participants benefit to such a degree that the gender gap that exists when using traditional desktop displays seems to be drastically narrowed when users use large displays. Hence, large screen displays may be a simple, low-cost method to assist in reducing the gender gap in 3-D virtual navigation tasks.

One limitation of the current work is the scope of tasks and environments tested. Because the experiments were targeted at understanding the effects of large displays on both genders, we carefully eliminated factors that could have made the results difficult to interpret. We believe that more work examining the interactions with some of these factors, such as distinct landmarks or a less structured environment, is required in order to further generalize findings. We are continuing this research as well as examining the effects of different fields of view for more traditional desktop productivity tasks to see if this new finding generalizes to other task domains.

CONCLUSION

In this paper, we presented work demonstrating that wider fields of view provided by larger displays afford better optical flow cues and improve performance for all users navigating 3-D virtual environments. We have shown that this improvement is so large for female participants that it narrows the gender gap, in which researchers have repeatedly observed males outperforming females in such tasks. We have also shown that providing the cues that females rely on during 3-D navigation may allow them to perform as well as males, even on the spatial tasks that have traditionally exhibited strong gender biases against them. Recognizing this, we believe that we are now much better equipped to design and build systems that allow all users to perform effectively in 3-D virtual navigation tasks. Designing systems based on the principles we have derived may be a desirable alternative or a complement to existing attempts to narrow the gender gap by requiring increased training for female users. We are currently exploring similar design principles for large and multiple display systems that make users more effective in more traditional productivity tasks.

ACKNOWLEDGMENTS

We gratefully acknowledge Earl Hunt and Maryam Allahyar at the University of Washington, Barbara Tversky at Stanford University, and Steven Ellis at the NASA Ames Research Center for their stimulating discussion and insightful comments. Brent Field at Microsoft helped with user testing and data analysis on the initial studies. Gary Starkweather at Microsoft Research designed and built the large screen displays, without which this research would not have been possible. Dennis Cosgrove at Carnegie Mellon University assisted greatly in developing software used in the experiments. We also thank our anonymous reviewers for their time and feedback.

REFERENCES

Alfano, P. L., & Michel, G. F. (1990). Restricting the field of view: Perceptual and performance effects. *Perceptual and Motor Skills, 70,* 35–45.

Atchley, P., & Andersen, G. J. (1999). The discrimination of heading from optic flow is not retinally invariant. *Perception and Psychophysics, 61,* 387–396.

Carnegie Mellon University. (n.d.). Alice 3D authoring system [Computer software]. Available from http://www.alice.org

Chambers, W. W. (1982 April). *Visual technology research simulator: Flight simulation-avionic systems and aero medical aspects.* Presented at the Royal Aeronautical Society's Flight Simulation Symposium, London.

Crook, T. H., Youngjohn, J. R., & Larrabee, G. J. (1993). The influence of age, gender, and cues on computer-simulated topographic memory. *Developmental Neuropsychology, 9*, 41–53.

Crowell, J. A., & Banks, M. S. (1993). Perceiving heading with different retinal regions and types of optic flow. *Perception and Psychophysics, 53,* 325–337.

Cutmore, T. R. H., Hine, T. J., Maberly, K. J., Langford, N. M., & Hawgood, G. (2000). Cognitive and gender factors influencing navigation in a virtual environment. *International Journal of Human-Computer Studies, 53,* 223–249.

Czerwinski, M., Tan, D. S., & Robertson, G. G. (2002). Women take a wider view. In *Proceedings of the CHI 2002 Conference on Human Factors in Computing Systems* (pp. 195–202). New York: Association for Computing Machinery.

Darken, R., & Sibert, J. (1996). Navigating in large virtual worlds. *International Journal of Human-Computer Interaction, 8,* 49–72.

Devlin, A. S., & Bernstein, J. (1995). Interactive wayfinding: Use of cues by men and women. *Journal of Environmental Psychology, 15,* 23–38.

Dolezal, H. (1982). *Living in a world transformed: Perceptual and performatory adaptation to visual distortion.* New York: Academic Press.

Draper, M. H., Viirre, E. S., Furness, T. A., & Gawron, V. J. (2001). Effects of image scale and system time delay on simulator sickness within head-coupled virtual environments. *Human Factors, 43,* 129–146.

Duffy, C. J. (2000). Optic flow analysis for self-movement perception. *International Review of Neurobiology, 44,* 199–218.

Eckstrom, R. B., French, J. W., Harman, H. H., & Dermen, D. (1976). *Kit of factor-referenced cognitive tests.* Princeton, NJ: Educational Testing Service.

Gibson, J. J. (1966). *The senses considered as perceptual systems.* Boston: Houghton Mifflin.

Gibson, J. J. (1979). *The ecological approach to visual perception.* Boston: Houghton Mifflin.

Halpern, D. F. (2000). *Sex differences in cognitive abilities* (3rd ed). Mahwah, NJ: Erlbaum.

Hosman, R., & van der Haart, J. (1981). Effects of vestibular and visual motion perception on task performance. *Acta Psychologica, 48,* 271–281.

Hubona, G. S., & Shirah, G. (2004). The gender factor performing visualization tasks on computer media. In *Proceedings of the 37th Annual Hawaii International Conference on System Sciences* [CD-ROM]. Los Alamitos, CA: IEEE Computer Society Press.

Hunt, E., & Waller, D. (1999). *Orientation and wayfinding: A review* (ONR Tech. Rep. N00014-96-0380). Arlington, VA: Office of Naval Research.

Kearns, M. J., Warren, W. H., Duchon, A. P., & Tarr, M. J. (2002). Path integration from optic flow and body senses in a homing task. *Perception, 31,* 349–374.

Kimura, D. (1999). *Sex and cognition.* Cambridge, MA: MIT Press.

Kirschen, M. P., Kahana, M. J., Sekuler, R., & Burack, B. (2000). Optic flow helps humans learn to navigate through synthetic environments. *Perception, 29,* 801–818.

Klatzky, R. L., Loomis, J. M., Beall, A. C., Chance, S. S., & Golledge, R. G. (1998). Spatial updating of self-position and orientation during real, imagined, and virtual locomotion. *Psychological Science, 9*, 293–298.

Lawton, C. A., Charleston, S. I., & Zieles, A. S. (1996). Individual and gender-related differences in indoor wayfinding. *Environment and Behavior, 28*, 204–219.

Loomis, J. M., Klatzky, R. L., Golledge, R. G., & Philbeck, J. W. (1999). Human navigation by path integration. In R. G. Golledge (Ed.), *Wayfinding: Cognitive mapping and other spatial processes* (pp. 125–151). Baltimore, MD: Johns Hopkins.

Patrick, E., Cosgrove, D., Slavkovic, A., Rode, J., Verratti, T., & Chiselko, G. (2000). Using a large projection screen as an alternative to head-mounted displays for virtual environments. In *Proceedings of the CHI 2000 Conference on Human Factors in Computing Systems* (pp. 478–485). New York: Association for Computing Machinery.

Piantanida, T. P., Boman, D., Larimer, J., Gille, J., & Reed, C. (1992). Studies of the field-of-view/resolution tradeoff in virtual reality. In *Proceedings of the SPIE: Human Vision, Visual Processing and Digital Display III* (Vol. 1666, pp. 448–456). Bellingham, WA: SPIE-The International Society for Optical Engineering.

Prestopnik, J. L., & Roskos-Ewoldsen, B. (2000). The relations among wayfinding strategy use, sense of direction, sex, familiarity, and wayfinding ability. *Journal of Environmental Psychology, 20,* 177–191.

Richardson, A. E., Montello, D. R., & Hegarty, M. (1999). Spatial knowledge acquisition from maps, and from navigation in real and virtual environments. *Memory and Cognition, 27,* 741–750.

Richman, J. B., & Dyre, B. P. (1999). Peripheral vision stimulation benefits heading perception during active control. In *Proceedings of the Human Factors and Ergonomics Society 43rd Annual Meeting* (p. 1411). Santa Monica, CA: Human Factors and Ergonomics Society.

Sandstrom, N. J., Kaufman, J., & Huettel, S. A. (1998). Males and females use different distal cues in a virtual environment navigation task. *Cognitive Brain Research, 6,* 351–360.

Tan, D. S., Czerwinski, M., & Robertson, G. G. (2003). Women go with the (optical) flow. In *Proceedings of the CHI 2003 Conference on Human Factors in Computing Systems* (pp. 217–224). New York: Association for Computing Machinery.

Tan, D. S., Robertson, G. G., & Czerwinski, M. (2001). Exploring 3D navigation: Combining speed-coupled flying with orbiting. In *Proceedings of the CHI 2001 Conference on Human Factors in Computing Systems* (pp. 418–425). New York: Association for Computing Machinery.

Thorndyke, P., & Hayes-Roth, B. (1982). Differences in spatial knowledge acquired from maps and navigation. *Cognitive Psychology, 14,* 560–589.

Voyer, D., Voyer, S., & Bryden, M. P. (1995). Magnitude of sex differences in spatial abilities: A meta-analysis and consideration of critical variables. *Psychological Bulletin, 117*, 250–270.

Waller, D. (1999). An assessment of individual differences in spatial knowledge of real and virtual environments. *Dissertation Abstracts International, 60*(4), 1882B. (UMI No. DA9924139)

Warren, W. H., & Kurtz, K. J. (1992). The role of central and peripheral vision in perceiving the direction of self-motion. *Perception and Psychophysics, 51,* 443–454.

Warren, W. H., Morris, M. W., & Kalish, M. (1988). Perception of translational heading from optical flow. *Journal of Experimental Psychology: Human Perception and Performance, 14,* 646–660.

Desney S. Tan is a researcher in the Visualization and Interaction group at Microsoft Research. He received a Ph.D. in computer science at Carnegie Mellon University in 2004.

Mary P. Czerwinski is a principal researcher and manager of the Visualization and Interaction Research Group at Microsoft Research. She received a Ph.D. in cognitive psychology from Indiana University in 1988.

George G. Robertson is a principal researcher at Microsoft Research. He received a Ph.D. equivalent in computer science (via a faculty appointment) from Carnegie Mellon University in 1977.

Date received: January 5, 2005
Date accepted: May 29, 2005