5-2012

Effects of Immersion on Spatial Updating in Virtual Panoramas

Phillip Napieralski
Clemson University, pnapier@clemson.edu

EFFECTS OF IMMERSION ON SPATIAL UPDATING IN VIRTUAL PANORAMAS

A Thesis
Presented to
the Graduate School of
Clemson University

In Partial Fulfillment
of the Requirements for the Degree
Master of Science
Computer Science

---

by
Phillip E. Napieralski
May 2012

Accepted by:
Dr. Sabarish V. Babu, Committee Chair
Dr. Larry F. Hodges
Dr. Timothy A. Davis

ABSTRACT

An experiment was conducted to examine the effects of immersion on a common spatial knowledge acquisition task in a virtual panoramic environment. A virtual panorama, such as Google Streetview, is an environment that is comprised of 3D (spherical) images taken at regular intervals in a real world setting. Participants navigated the National Mall area of Washington DC in Google Streetview panoramas using either a keyboard and mouse or a head-mounted display (HMD) with a head orientation tracker. In an exploration phase, participants were asked to navigate and observe landmarks. Then, in a testing phase, participants were asked to look in the direction of the perceived landmarks, their gaze direction was recorded. We measured the angular difference between participants' gaze direction and the landmark direction. We found no significant difference between the immersive and desktop conditions on participants' accuracy of direction to landmarks as well as no difference in their presence scores. We found that some participants in the HMD condition viewed the environment from an egocentric perspective, while the participants in the desktop condition tended to form cognitive maps to inform spatial orientation to landmarks. These findings suggest the visual cues in the virtual panoramic environment, such as street signs, buildings and trees, seem to have a stronger influence in both viewing conditions in determining direction to landmarks, than the egocentric cues such as first-person perspective and natural head-coupled motion.

ACKNOWLEDGMENTS

I would first like to thank my advisor, Dr. Sabarish V. Babu for his guidance and support. I would also like to thank Dr. Christopher Pagano, Dr. Tim Davis, Bliss Altenhoff, Lindsay O. Long and Lauren Cairco for statistics assistance and feedback on the project. I would also like to thank everyone who participated in the experiment. Finally, I would like to thank my parents and family for their undying support.

TABLE OF CONTENTS

Table of Contents (Continued)

LIST OF TABLES

LIST OF FIGURES

CHAPTER ONE

INTRODUCTION

Virtual panoramic environments (VPEs) are now widespread. In fact, VPEs such as Google Streetview and Bing Streetside can be used in any modern browser with Flash, WebGL or Silverlight. When used in conjunction with head-coupled motion, the VPE becomes an immersive virtual panoramic environment (IVPE). However, no experiment to date has studied the effect of immersion on spatial knowledge in an IVPE. The present experiment uses an IVPE with a consumer level head-mounted display (HMD) and provides a cost effective way for the average person to navigate an environment almost anywhere in the world.

Many users tend to find VEs confusing and difficult for wayfinding training when compared to the real world or map [Waller et al. 1998]. In Waller et al, participants needed to be exposed to a VE up to five times longer than the real world to obtain accurate route and configurational knowledge.

In order to assess spatial knowledge acquisition ability of users in IVPEs, a thorough understanding of perceptual and cognitive influences, such as proprioceptive cues, is crucial. A common way for people to keep track of their positions in space, for instance, is by keeping track of where landmarks are in relation to themselves. This tracking of landmarks in relation to oneself is called *spatial updating*. Since this relies heavily on keeping an accurate representation of an internal perceived heading, [Klatzky 1998], it seems likely that an immersive experience, one with natural head rotation and a first person perspective, would benefit the user's spatial updating ability. In this paper, we attempt to bridge the gap in understanding the differences in wayfinding ability, specifically spatial updating, between an immersive and non-immersive VPE using a simple direction to landmark task.

Viewing and Navigating a Virtual Environment

VEs may be viewed with a variety of different devices such as monitors, CAVEs, HMDs and with a stereoscopic view with screen displays or HMDs. Bowman et al. [2001] talks about the five different ways one can travel through a VE. Of those, the most commonly employed technique is called "steering," in which participants' travel in their gaze or pointing direction. This can be employed from an egocentric perspective using an HMD, or from a more exocentric perspective using a monitor.

These different viewing modalities and different travel techniques may influence the speed and accuracy of cognitive map construction and spatial knowledge acquisition in a VE or VPE. This cognitive process of creating a cognitive map and acquiring spatial knowledge from an unfamiliar environment is called wayfinding. Thorndyke & Hayes-Roth [1982] categorized spatial knowledge into landmark, procedural and survey knowledge.

The first type of knowledge that may help a person in wayfinding is *landmark knowledge*, which is information gained from visual details and unique landmarks in an environment. Another type of spatial knowledge gained is *procedural knowledge* or route knowledge. People may try to remember the steps and turns they have taken to get to a particular place in space. This can be combined with landmark knowledge (e.g. at the

Washington Monument, take the first left) to begin creating a cognitive map. This knowledge can only be accurately obtained when navigating an environment. Finally, *survey knowledge* is acquired when a person begins creating a cognitive map of the environment [Thorndyke and Hayes-Roth 1982]. This knowledge can be accurately obtained from studying a map or from extended exposure to the environment. This map contains topological information of the environment. Distances between landmarks and routes are encoded in terms of a fixed frame of reference. This frame of reference may be egocentric, if the environment was viewed from a first-person perspective, or geocentric, if a map of the environment was studied. Survey knowledge is the most useful knowledge to gain for wayfinding [Darken and Silbert 1996].

There are two primary bases for obtaining the spatial knowledge necessary for successful wayfinding when navigating an environment are piloting and path integration [Gallistel 1990; Mittelstaedt and Mittelstaedt 1980]. Piloting is a term that refers to navigating an environment based on tracking the location of distinct objects in the environment and relating that to a cognitive map. Path integration consists of using kinesthetic and sensory cues to develop an accurate map of the environment. Previous research suggests that path integration is necessary to integrate a disjoint environment into a unified and coherent representation of an environment [Galistel 1990; Poucet 1993].

Much work has been done to further understand path integration. Klatzy et al. [1998] showed that users will not sufficiently update their internal heading, a cue that assists in path integration, without physical rotation. In the paper, participants walked along a two segment path. Each segment was connected by a turn of some degrees. Participants did this either by physically walking the path (while blindfolded), by viewing an experimenter walking the two segments, having the two segment path described, sitting on a stool with HMD and having the stool rotated by an experimenter when the turn came (this is known as the *real-turn* condition), and finally, some participants were placed in a *visual-turn* condition in which they both translated and rotated through the two-segment path by using a joystick. It was shown that participants in the real-turn and walking conditions performed equally well. The others performed quite poor as the turn angle increased. This alludes to the importance of natural head motion on a simple turn to origin task.

A person's ability to update their internal heading when rotating, whether passive or active, is known as *spatial updating*. Riecke et al. [2012] showed that a rotating sound field ("auditory vection") can be utilized to facilitate updating of a person's spatial orientation. Wang et al. [2006] showed that tracking more objects can make spatial updating more difficult.

The Use of Virtual Environments for Navigation Training

There are three main ways that have been studied in the past to assess spatial knowledge acquisition and their carryover effect to the real world: real world navigation, studying maps and virtual navigation.

Bliss et al. [1997] studied the use of VEs in training firefighters. In the experiment, participants were trained in the building layout using a VE, blueprints or not at all. After training, participants were tasked with navigating the building and rescuing a life-sized doll. Dependent measures included the number of wrong turns and the time to complete the rescue. Both types of training were superior to no training. Performance when training in the VE did not differ significantly from performance when training with a blueprint. However, the system used by the authors ran at only 20 frames/sec and at a relatively slow translation pace.

Darken and Banker [1998] considered the effect of training with a VE and a map on a large-scale outdoor environment. Participants were tasked with locating nine control points. Three methods for training were employed: studying an orienteering map, using a VE and a map, or walking in the real world and using a map. Subjects were trained in all conditions for one hour. The authors found that subjects with intermediate orienteering ability performed best when trained the VE with map. The VE system used by the authors was rendered with high-fidelity, however the VE ran at only 3 frames/sec. This may indicate why the difference between the map and VE conditions was not stronger in people with beginner and expert orienteering ability.

Both Bliss et al. [1997] and Darken and Banker [1998] hint at a fundamental flaw of training in a VE. Creating a VE costs money and time to create the assets. Further, the cost of rendering a realistic environment can be even greater, resulting in low frame rates as realism increases. An easy fix to this is to employ a VPE. That is, instead of modeling all the components to create a realistic and expensive looking VE, simply take panoramic imagery at discrete locations and have participants navigate in that way. This would save greatly on rendering and creation costs.

Related Work

Wayfinding in VEs has been studied extensively. The most common techniques used to evaluate wayfinding performance and spatial knowledge acquisition in a VE are point to landmark or point to origin tasks [Klatzky et al. 1998; Chance et al. 1998], relative distance between landmarks [Patrick et al. 2000] and sketch maps [Zanabaka et al. 2004]. The point-to-landmark or point-to-origin tasks allow us to test a participant's spatial updating ability. This task involves walking along a path, while keeping track of objects seen and origin. At the end of the path, participants are ask to point or look in the direction of previously seen landmarks or the point of origin [Klatzky et al. 1998; Chance et al. 1998]. Individual differences in this type of task tend to be high [Chance et al. 1998; Riecke et al. 2012]. Relative distance between landmarks also helps test landmark knowledge, but further helps us test their survey knowledge. Finally, sketch maps give us a clear picture of the accuracy of a person's overall spatial knowledge [Zanbaka et al. 2004].

To the best of our knowledge, this is the first perceptual wayfinding experiment performed in an IVPE. The closest research performed in an immersive VE (IVE) is by Chance et al. [1998] and Sigurdarson et al. [2012]. In the second experiment performed by Chance et al., participants were instructed to walk through a maze using a joystick and desktop monitor, joystick and HMD or using a HMD with natural walking. After they reached the end of the maze, participants were instructed to point to two different, now non-visible, objects located somewhere in the maze. The absolute value of the angle subtended between the landmark direction and the gaze direction was the dependent measure. Over 24 between-subject participants in all conditions, they found there was no significant difference between the desktop and joystick/HMD condition. There was also no significant difference between joystick/HMD condition and the natural walking condition. However, the natural walking condition led to a significant improvement over the desktop condition. This can be explained by the additional spatial updating cues that natural walking provides.

In Sigurdarson et al. [2012], an experiment was performed similar to that of Klatzky et al. [1998]. Participants were either placed in a REAL TURN condition, where they physically rotated when moving through the environment, or a VISUAL ONLY condition, where participants did not physically rotate. The REAL TURN condition used a joystick and HMD with natural head movement, while the VISUAL ONLY condition used a joystick and HMD with no natural head movement, the joystick controlled both rotation and translation through the environment. The environment used was much more realistic than Klatzky et al. [1998] to determine if the difference Klatzky et al. found was due to an unrealistic environment. Over 10 different turning angles, and 12 within-subject participants, no difference in pointing accuracy was found between the conditions. This suggests that physical rotation may not help disorientation in a VE if the VE is realistic enough.

In this paper, an experiment was conducted to examine a subject's spatial updating ability in a large-scale VPE using a gaze to landmark task similar to that of [Chance et al. 1998]. Our contributions include utilizing a low-cost VR system to provide an immersive experience to a user. This experiment is also one of the first, if not the first, to examine spatial updating and presence in a large-scale VPE as compared to an IVPE.

CHAPTER TWO

EXPERIMENT DESIGN

The primary aim of this study was to compare spatial updating in a VPE versus an IVPE. The following research questions were specifically asked:

1. What techniques are used when navigating a large-scale VPE?
2. Are participants better able to update the position of landmarks relative to themselves, when given an immersive setup?

First, we wanted to determine what techniques participants used to recall the location of landmarks in relation to themselves in a VPE, much like what Bowman et al. [1999] did for a similar experiment.

Second, to determine if an immersive setup allowed for more accurate spatial updating, participants' gaze judgments were logged. Participants made gaze judgments, either by naturally looking or clicking and dragging with the mouse, to where they perceived previously seen landmarks to be in a VPE. Perceived gaze angles were measured in two ways: the angle subtended between the north axis and the participant gaze angle, and the angle subtended between the actual landmark position and the participant gaze angle. In the following subsections, the participants, experiment design, apparatus, and procedure will be detailed.

Participants

There were 27 participants total between conditions. Participants 2, 11 and 13 were excluded for not following directions or technical difficulties. Further, unless otherwise stated in the results section, 7 participants were excluded from any data analysis because they had extensive prior exposure to the Washington DC area as measured by a pre-questionnaire. This left 17 participants between conditions: 10 participants in the HMD condition and 7 in the Desktop condition. The average age in the Desktop condition was 26. The average age in the HMD condition was 25.1. There were 5 females and 5 males in the HMD condition; 3 males and 4 females in the Desktop condition.

Variables and Design

The experiment used a between subjects design, where participants were assigned one of two conditions (HMD or Desktop). Each participant was asked to point at landmarks in three different tests. From each testing point, the landmarks were located at angles (measured from the north axis): 326º, 273º, and 305º in the first test; 128º, 151º, 180º and 213º in the second test; and finally 213º, 224º and 190º in the final test. Landmarks viewed by participants included: the Marriott Hotel, Washington Monument, Department of Commerce Building, Gallery of Art, Space Center, Sculpture Garden, Smithsonian Center, Old Post Office, US Post Office and the IRS building. The primary dependent variable was the angle subtended between participants' gaze direction and the north axis. The angle subtended between gaze direction and landmark direction was also recorded.

Apparatus and Materials

Figure 1 - The apparatus in the HMD condition (left). An example of what the participant would see in the virtual panoramic environment (right).

Figure 1 depicts the apparatus used. Participants in the Desktop condition were seated in a chair, viewing a 17 inch monitor at 800x600 pixel resolution. Participants in the HMD condition were standing during the experiment, viewing the same image as in the desktop condition, but viewing biocularly (same image viewed in each eye) through an eMagin Z800 HMD at 800x600 resolution and 32ºx42º field of view. The HMD condition uses the same setup as described in Hu et al. [2011], however with no treadmill and the addition of a WebGL layer. Participants were teleported to adjacent panoramas in their gaze direction using a button press on a Wiimote or by clicking on the mouse. The view was rotated by clicking and dragging the mouse, or by naturally looking in the direction to move in the HMD condition. Participants in the HMD condition were standing, and physically rotated their bodies in the direction of travel.

System Goals

There were three primary goals of the system required for this experiment:

1. Allow a head-mounted display to communicate with Google Streetview.
2. Annotate a route for participants to follow.
3. Annotate specific landmarks along the route that participants follow.

The architecture for allowing the head-mounted display and Wiimote to interact with a webpage is described in Hu et al. [2011]. In summary, the server program continually polls the head-mounted display for yaw and pitch data, and continually polls a Nintendo Wiimote to determine if a button has been pressed. This data is continually sent, using UDP, from the server to the client web page.

WebGL and Streetview Layers

To annotate Google Streetview, a WebGL layer was added to the system from Hu et al. [2011]. The system viewed by participants consisted of two layers: a Google Street View layer to view panoramas and a WebGL layer on top that allowed 3D arrows and landmark markers to be denoted on specific panoramas. The browser used to view these two layers and run the entire experiment was Google Chrome version 13.0.782.55 beta-m. The browser was made fullscreen during the experiment. Google Street View's version 2 API was used to load the panoramas. THREE.JS was used to create the WebGL Layer. Rodrigues' rotation formula [Belongie 2012] was used to generate the rotation matrix around the camera's up and right axes; the current look vector was then pre-multiplied by this rotation matrix to obtain a new looking direction. This technique was executed to update the WebGL layer every time the HMD's yaw and pitch was updated via the server and client system [Hu et al. 2011]. The Google Streetview layer was also updated at this time using an API call.

Procedure

Pre-experiment

Each participant first read and signed an Informed Consent Form. Next, each participant completed a demographic and computer user questionnaire which included the questions "Have you spent more than 2 weeks (total throughout your life) in Washington DC?" and "Have you been to Washington DC in the past 5 years?" Participants were then given a standard Guilford-Zimmerman (GZ) spatial ability questionnaire [Guilford and Zimmerman 1948]. The participant was given the overall instructions for the GZ test and given time to thoroughly read the instructions. Each participant was then given 10 minutes to complete as many questions as possible. The score at the end of 10 minutes was determined to be the number correct minus one-fourth the number incorrect. The pre-questionnaire is available in Appendix B. Afterwards, participants were either seated at the computer, or donned with a HMD and Wiimote in their right hand with the training environment loaded. The complete system describing how the HMD and Wiimote were used in the system is described in Hu et al. [2011]. Participants wearing the HMD were asked to read the text on the back of a nearby bus in the first panorama to ensure they could see the training environment clearly. Once the participant was situated, the instructions for the task were explained.

Training Environment

In both conditions, participants were first placed in a training environment to become accustomed to the system and the task. Participants navigated Google Streetview in an area of Washington DC that was separate from the testing environment. The area was annotated with orange arrows that the participant was required to follow. Three landmarks were annotated with yellow billboard arrows. Participants were told to observe and remember the location of these landmarks as they will be asked to look where they perceive the landmark to be later on. Once the participant reached the end of the route (denoted by orange arrows), they entered a “testing” phase. In the testing phase, the participants were shown a picture of the first landmark they had observed and asked to “look in the direction where you believe the landmark to be.” Participants logged their guesses using a button press on the Wiimote or on the keyboard and the next landmark was shown. After the three landmarks were shown and answers logged, the participant was asked if they had any questions. The data from the training environment was not used in any analysis.

Testing Environment

After the participant finished the training environment and understood the task, the testing environment was loaded. The participant started near the Marriot Hotel on 1331 Pennsylvania Avenue in the Washington DC area. As in the training environment, they were instructed to view, and remember, the location of landmarks that were marked with a yellow arrow as in Figure 1 (right) while following a path denoted by arrows on the

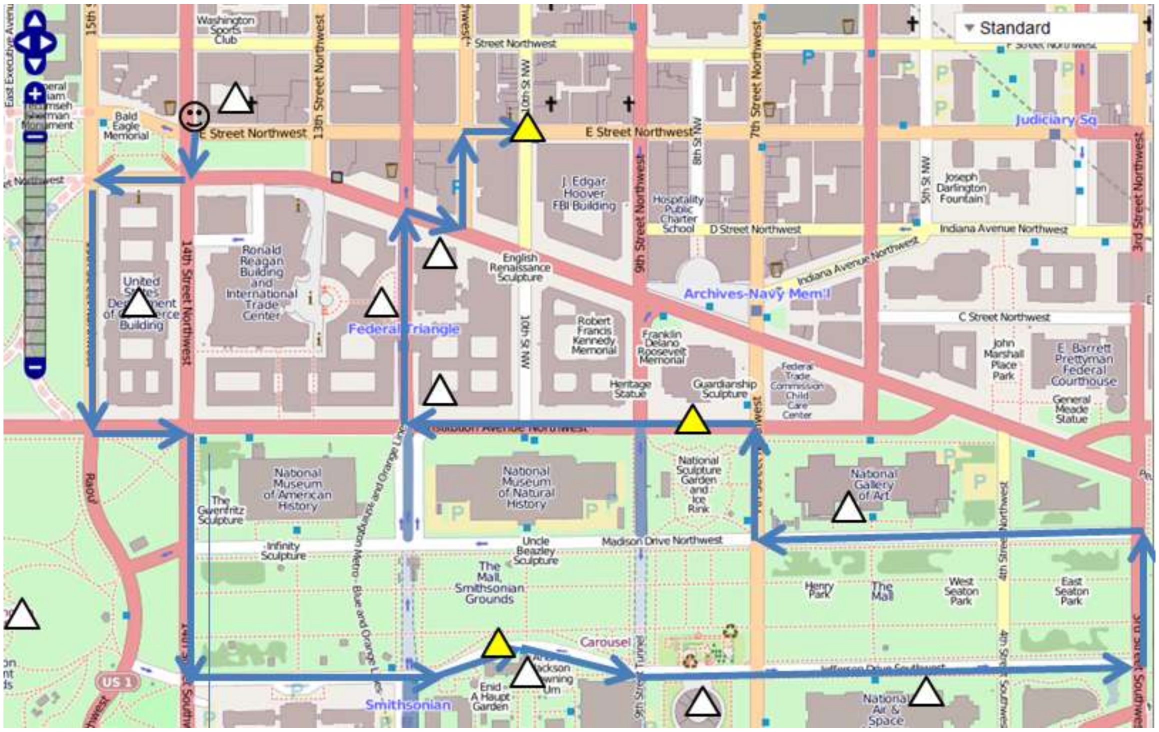

Figure 2 – The path taken by participants in the testing environment. Yellow triangles are the testing points. White triangles are the landmark positions. The smiley face is where participants started. Map image © OpenStreetMap contributors, CC BY-SA.

ground (also shown in Figure 1). The route taken by participants in this environment is shown in Figure 2. Instead of only being tested once, as in the training environment, participants were given three separate tests. Participants viewed different landmarks for each test. The participant was allowed as much time as necessary to think about where the landmark was located. When the participant thought they had the correct direction to the landmark, a key press on the keyboard or Wiimote logged their answer and prompted an image of the next landmark to look at. Participants were allowed as much time as needed to make a guess. Trials in which the participant did not recall seeing the landmark were marked and not used in any data analysis.

Post-experiment

After the participant logged their guess of the final landmark, they were prompted with a thank you message, the HMD was removed (if necessary), post-questionnaire was administered, and finally a NASA Task Load Index [Hart and Staveland 1988] was administered online. The post-questionnaire included a Steed-Usoh-Slater presence questionnaire [Usoh et al. 2000]. Participants were also asked about their confidence in performing the task as well as the techniques they used in performing the task. This questionnaire is available in Appendix C.

CHAPTER 3

RESULTS

The first thing necessary was to prep the data for analysis. If this was not done, a simple difference between the participant's angle judgment and the landmark angle answer would not necessarily suffice. For example, consider a participant's judgment of 1º from the north axis compared to an angle of 359º for the landmark answer (to the north axis). If we simply subtract the two, there is a difference of 358º, but in reality, the difference is only 2º. Later in this section, a linear regression analysis is performed with the actual landmark's angle on the x-axis and the participant's guess on the y-axis. A guess of 1º when compared to an actual answer of 359º would grossly underestimate the correlation coefficient of the linear regression. To overcome these difficulties, angles can be "wrapped" by adding or subtracting 360, or circular statistics can be employed [Batschelete 1965, 1978; Turvey 1992].

After the data for each participant was wrapped, the first thing to verify was whether those with prior experience to the Washington DC area performed better than those without prior experience. This was analyzed using a multiple regression. Then, an analysis of task performance between conditions was analyzed (on those who did not have prior exposure to Washington DC). To maximize the chance of finding a difference, the data was analyzed in two different ways. First, using circular statistics, and second, using multiple regression over all participants guesses. Finally, the techniques used in the study are summarized based on those who responded to the "Techniques used" question in the post-questionnaire.

Wrapping the Data

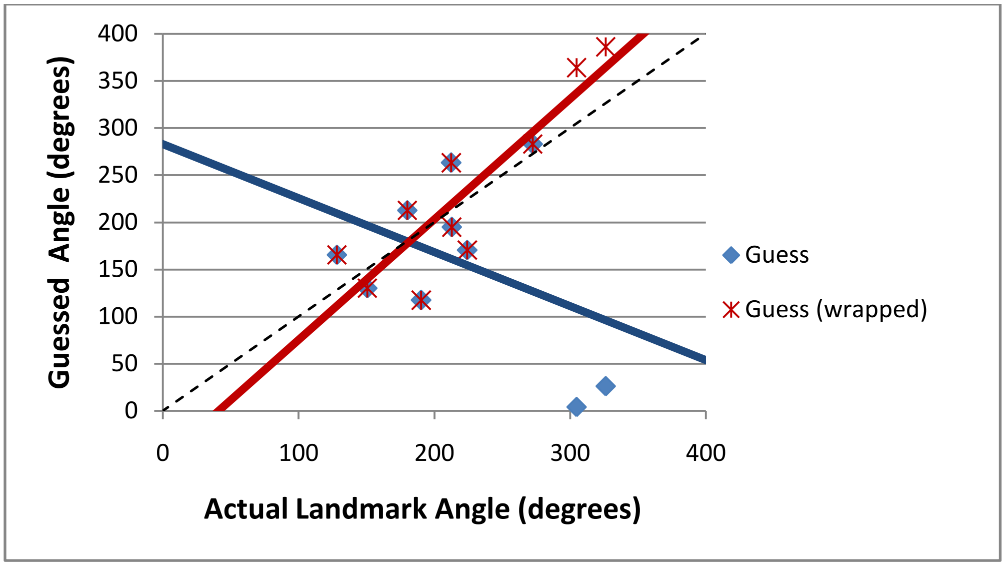

Figure 3 – This graph demonstrates why "wrapping" each participant's data by adding or subtracting 360º is necessary

For each participant, regression plots were created to determine the variation per participant. After computing each participant's linear regression, it seemed that many participants did very poor on those landmarks whose answers were near 360º. This was not the case, however. The data was falling into the trap previously mentioned. Figure 3 shows an example of this with participant 4's data. The blue dots show the participant's original guesses, compared to the landmarks actual angle. The red dots show the participant's "wrapped" guesses. Based on the original regression line, this participant did quite poor (r^2 = .162). However, after 360º was added to the 2 data points on the right side of the graph (shown in red on the same figure), the r^2 value improved substantially (r^2 = .774). In general, if the difference between any participant's guess and the actual landmark angle was greater than 180º, 360º was either added or subtracted. If the r^2 value improved, the new value was kept. All participant data was wrapped in this way.

<u>Does prior exposure to the area affect performance?</u>

In the pre-questionnaire, participants were asked if they had been to Washington DC in the past 5 years and if they spent more than 2 weeks in the area over their whole life. To test for differences between slope and intercept of those with prior DC exposure and those without, a multiple regression was performed. The multiple regression was initially performed with an actual landmark angle X exposure interaction term, yielding an r^2 = 0.643 (n = 220), with a partial *F* of 292.1 for actual landmark angle ($p < 0.0001$). The partial *F* for the exposure was 5.81 for effect of prior exposure (p = 0.0170). The partial *F* for the interaction term was 6.656 (p = .0104). The partial *F* for prior exposure condition decreased slightly to 5.71 with the removal of the interaction term (p = 0.0175).

The partial *F* for actual landmark angle assesses the ability that the actual landmark angle predicts the variation in the gaze angle responses after variation due to other terms (prior exposure and interaction) have already been accounted for. The partial *F* for actual landmark angle tests for a main effect of actual landmark angle. The partial *F* for prior exposure assesses the degree that the intercepts of the prior exposure conditions differ, or a main effect of prior exposure. The partial *F* for the interaction term assesses the degree that the slopes for the two conditions differ from each other. Thus, the multiple regression revealed a statistically significant main effect for actual landmark angle, as well as a main effect for prior exposure. There also was an interaction. Therefore, the slopes and intercepts of the functions predicting guessed angle from actual landmark angle did indeed differ between prior exposure conditions. The angle guesses were much less accurate for those without prior exposure to Washington DC. The difference was 52.8º on average. See Table 1 for a summary of the data obtained from the regression analysis.

	n	**Slope**	**Intercept**	**r^2**
Prior Exposure	66	0.916	-2.86	0.473
No Exposure	154	1.24	-55.67	0.701
Overall	220	1.14	-38.57	0.622

Table 1 – This table demonstrates the difference between those with prior extended exposure to the Washington DC area and those without exposure

Given the vast difference in performance between those who have prior exposure to the Washington DC area, further analysis excluded participants with prior exposure.

Spatial Ability Scores Between Conditions

Before comparing the performance of participants between conditions, it is important to remove one major confound – participants' innate spatial ability. To do this, a one-way ANOVA was performed to determine if there was a difference in overall GZ score between the conditions. The mean and standard deviation of the desktop condition participants' GZ scores were: 17.89 and 7.08 (n = 7). The mean and standard deviation for the HMD condition was: 20.75 and 9.93 (n = 10). The ANOVA showed no significant difference between condition GZ scores (p = 0.34).

Circular Statistics

The four circular statistics used in the data analysis are: α, which can be thought of as the *average* angle guess for a particular landmark; *r*, which is the *coherency* of the guesses; Rayleigh's *z*, which is used to assess the significance of the *r* value; *h*, the *homing* coefficient, that measures how well the participants' guesses "home in" on the actual landmark angle. The *h* value is similar to the *r* value except it includes information about the accuracy of the guess to the landmark. See [Turvey 1992] for a more detailed explanation of circular statistics.

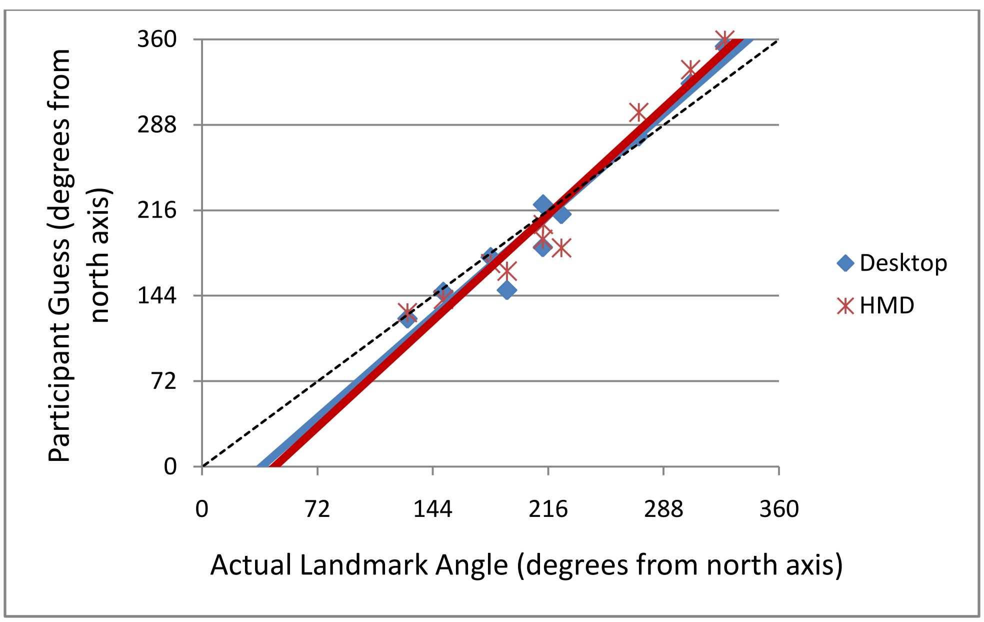

Figure 4 – Participant performance for both conditions, using circular statistics. Theta (Θ) is represented on the x-axis. Alpha (α) is presented on the y-axis.

Landmark	n	ϴ (degrees)	α (degrees)	r	h	Z
0	14	326	357	0.51	0.44	3.67
1	14	273	288	0.76	0.73	8.09
2	17	305	330	0.69	0.62	8.03
3	16	128	127	0.91	0.91	13.2
4	17	151	143	0.95	0.94	15.22
5	17	180	173	0.86	0.85	12.22
6	17	212	189	0.63	0.58	6.79
7	16	212	210	0.82	0.82	10.87
8	17	224	195	0.66	0.58	7.43
9	17	190	159	0.66	0.57	7.48

Table 2 – Circular statistics per landmark (for both conditions). Theta (ϴ) is the actual angle to the landmark (from the north axis), α is the *average* angle guess to the landmark (from the north axis). The *r* and *h* columns determine the accuracy over all participants' guesses (0 being not accurate, 1 being exact) and the Rayleigh's *z* determines the significance of α.

Table 2 shows the circular statistics for each landmark for both conditions. Recall that some participants did not notice all the landmark indicators; this explains why the *n* column varies. Circular statistics were computed for each condition. Then, the given alphas were placed into a linear regression against the actual landmark angle, theta. Figure 4 shows the plots between conditions and their associated linear trendlines.

The intercepts of the desktop and HMD conditions were -43.90 and -56.34, respectively. The slopes were 1.18 and 1.25. A multiple regression predicting participants' average angle judgment (α) from actual landmark angle (Θ) was performed with an actual landmark angle X condition interaction term, giving an $r^2 = .95$ ($n = 10$), partial F of 0.1869 ($p = 0.6713$) for interaction and a partial F of 0.012 ($p = 0.9156$) for condition with the removal of the interaction term. These results seem to indicate that any difference between conditions was purely to chance. However, the small n in circular statistics led us to perform a multiple regression over all the data.

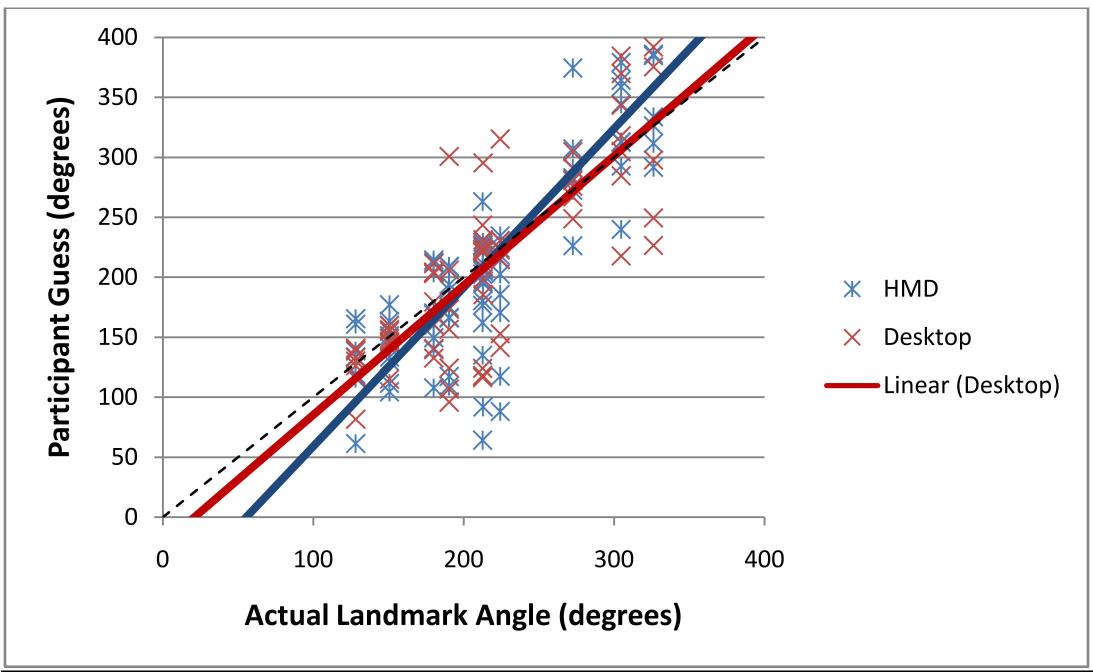

Figure 5 – Interaction between conditions over all data. The result is the same as with the circular statistics analysis (no significant difference).

<u>Multiple Regression Over All Guesses</u>

Figure 5 depicts the relation between actual landmark angle (from the north axis) and the angles reported by participants' gaze direction (again from the north axis) over all

Number of Turns	N	Mean	Standard Deviation
1	14	31.64	38.55
2	32	21.10	25.63
3	34	43.27	38.02
4	68	33.73	32.41
5	14	58.08	40.25

Table 3 – Summary of statistics based on number of turns taken between seeing the landmark and the corresponding test

participants between conditions. Each point in Figure 5 represents the judgments made by an individual subject for a particular landmark.

The intercepts of the desktop and HMD conditions were -31.49 and -71.7, respectively. The slopes were 1.08 and 1.31. A multiple regression predicting participants angle judgment from actual landmark angle was performed with an actual landmark angle X condition interaction term, giving an $r^2 = 0.70$ (n = 162), partial *F* of 0.20 (p = 0.16) for interaction and a partial *F* of 1.8 (p = 0.18). The partial *F* for effect of condition drops to 0.03 (p = 0.97) with the removal of the interaction term. These results suggest there is a strong chance any difference between conditions is purely by chance.

Effect of Number of Turns

We wanted to know if the numbers of 90º turns made since viewing the landmark had an effect on performance. A one-way ANOVA was performed. This ANOVA compared the absolute value of the angle difference between gaze direction and landmark direction within each of the number of turn categories: 1, 2, 3, 4 or 5 turns before testing. A summary of the mean and standard deviation for each level is displayed in Table 3. The ANOVA showed a significant difference (p = 0.0083). A post hoc Tukey HSD test showed a strong difference between 2 turns and 5 turns (p = 0.0071) as well as a trend between 2 turns and 3 turns (p = 0.0638). No other findings were significant.

A multiple regression was performed to determine if there was an interaction between number of turns and the condition. This regression was performed with a condition X number of turns interaction term. The partial *F* for condition was 0.68 (p = 0.41). The partial *F* for number of turns was 4.86 (p = 0.028). Which indicates, as in the ANOVA analysis above, that "number of turns" alone was a factor in performance. The partial *F* for the interaction term was 0.37 (p = .54), indicating no significant interaction.

Demographic Differences

A one-way ANOVA was performed with a dependent measure of angular difference between gaze direction and landmark direction and independent measure as gender. Accuracy between gender was not significant (*F* = 0.05, p = 0.82).

A linear regression was performed to determine if estimated hours of playing video games per week was an indicator of performance. The partial *F* was 2.49 (p = 0.1168), indicating a possible trend, but there was no significant effect.

Effect of Spatial Ability

To determine if a participants' innate spatial ability, as measured by the Guilford-Zimmerman spatial ability test, affected performance, a multiple regression was performed. The dependent variable was the angle difference between gaze direction and landmark direction. The independent variables were GZ score and condition. The regression was initially performed with a condition X GZ score interaction term. The partial *F* for the interaction term was 0.26 (p = 0.61). The partial *F* for GZ score was 4.56 (p = 0.034). The partial *F* for condition was 1.212 (p = 0.27). After the removal of the interaction term, the partial *F* for GZ score was 7.65 (p = 0.006) while the partial *F* for condition remained almost unchanged at 1.29 (p = 0.257). The r2 for the multiple regression with no interaction was 0.034. The r^2 for the simple regression of condition versus performance was 0.003. This means that 3.2% of the variance is accounted for by the GZ score.

Presence Between Conditions

A one-way ANOVA on the sum of the presence scores was performed on participants between conditions. The mean and standard deviation of the Desktop and HMD conditions were: 20.46 and 6.89; and 23.00 and 7.07, respectively. While presence scores were, overall, higher in the HMD condition, this difference was not significant (p = 0.3836). A possible reason for this is discussed in the next section.

Questionnaire Responses

The post-questionnaire asked participants the questions: "Did you use any non-marked objects (trees, cars, people, etc…) to determine the position of previously seen labeled landmarks during the test?" and "Did you use any other techniques? If so, what techniques?" Eight participants overall mentioned using street signs and other non-marked objects to navigate the environment. When removing the participants with prior exposure four participants mentioned using trees, cars and signs to keep track of their position.

Responses to the question, "what other techniques were used", were categorized in three ways: *procedural technique*, in which a participant would try to remember the landmark position based on previous turns; *cognitive map technique*, where participants would describe drawing a mental map from a "bird's eye" view; and finally, *egocentric technique*, where some participants described either physically or mentally pointing, or "lining up" themselves to remember previously seen landmarks. Participants in both conditions utilized a procedural and cognitive map technique when remembering previously seen landmarks. However, only participants in the HMD condition mentioned using a more egocentric technique. A summary of techniques used can be found in Table 3.

Condition	***Procedural***	***Cognitive Map***	***Egocentric***
Desktop	3	5	0
HMD	2	2-3	3-4

Table 4 – Techniques used. Note: one comment left by a participant could be interpreted as either cognitive map or as an egocentric technique. To account for this, a range was provided.

CHAPTER 4

DISCUSSION

Lack of Significance Explained

The hypothesis that participants would perform better in the HMD condition was not supported. Results showed no significant difference between HMD and Desktop conditions. However, this can be explained with two factors: the realism of the panoramas, and the jumping between panorama waypoints.

The realism of the panoramas facilitates the use of *piloting*. That is, many participants remarked using trees, signs and cars to determine their position more often than relying on their perceptual processes: spatial orientation and kinesthetic cues provided by the HMD. These piloting cues are freely available in both the Desktop and HMD conditions. This suggests that participants were using visual cues more often than the perceptual cues offered in the immersive HMD condition, though some participants did remark using an egocentric technique of keeping track of landmark positions (mentally or physically pointing to previous landmarks). Since the visual cues, the salient landmarks, in each of the panoramas are quite numerous; the additional perceptual cue is essentially "washed" out by the excess of visual information provided.

Further, the original goal of this experiment was to determine if there is an effect on spatial updating ability in an immersive condition. However, Google Streetview does not provide spatial updating or path integration cues since a participant will *jump* from one panorama to the next. They then lose the ability to keep track of how fast they are moving and, thus, the distance they have traveled. Consequently, once the landmark is removed from view, participants may be updating their memory of the location of landmarks in relation to themselves, but it quickly becomes inaccurate due to the waypoint jumping. The jumping can also explain the lack of difference in presence scores between the conditions. A number of participants in both conditions remarked:

- “The leaps from one view to another seemed of inconsistent distance…”
- “Discontinuous nature of navigation very jarring. Frustrating to move slowly through space…”
- “The jump-forward when clicking the mouse always made it seem as if I had locomoted farther than I actually had…”
- “It was very hard for me to estimate relative distances of the paths…”

This feedback suggests that a jumping metaphor between panoramas is, perhaps, the primary reason for the lack of significant difference between conditions in both task accuracy and sense of presence.

Number of Turns

From looking at Table 3, it seems that having an even number of turns between viewing and test shows the best performance. Figure 6 attempts to explain this phenomenon. Participants can generally be split up into two categories: turners and non-turners [Riecke 2007]. Turners, participants who consistently update their internal heading (green arrows in the figure), will get the “correct” answer when pointing to landmarks. Non-turners, or participants who consistently do not update their internal heading (red arrows

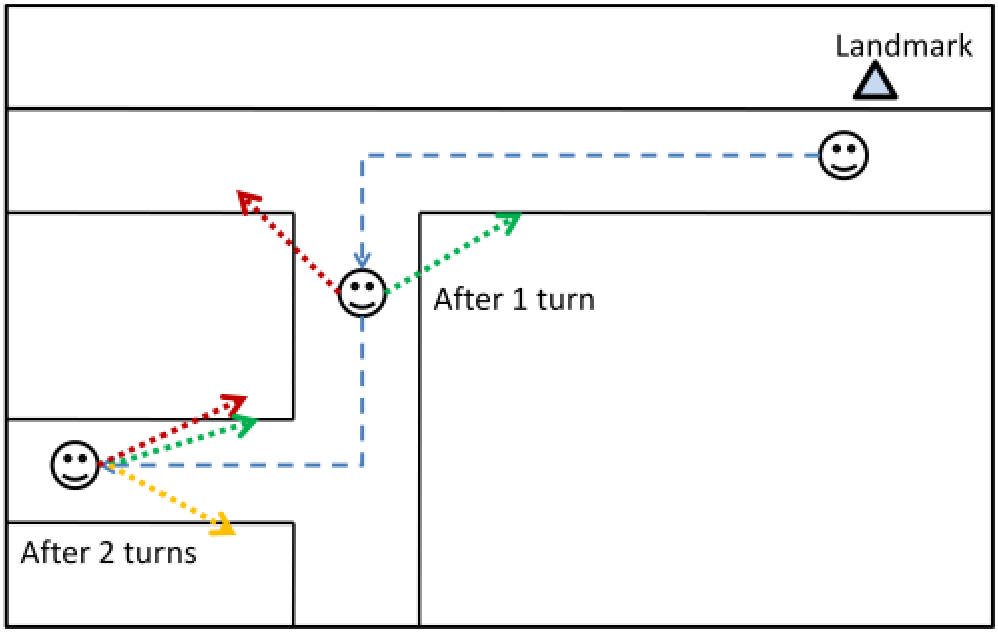

Figure 6 – This figure demonstrates why landmarks tested after turning only twice before testing shows better accuracy than a single turn before testing. The green arrows represent a “turner” internal heading. The red arrows represent a “nonturner” internal heading. The yellow arrow represents a turner that did not update his heading after the second turn, but did after the first turn.

in the figure) when turning, will be more correct after two turns than participants who only updated their internal heading after the first turn (yellow arrow in the figure). Thus, it appears more participants are likely to be in the correct range of the landmark. This may explain why participants overall performed better after two turns, versus one turn.

CHAPTER 5

CONCLUSION AND FUTURE WORK

In this research, we have successfully compared the spatial updating ability of participants in an IVPE, where participants wore an HMD, and a non-immersive VPE using a "look-to-landmark" task and system based on Hu et al. [2011]. Participants' angular judgments showed no difference between conditions. However, participants overall performed better when the landmark was recently seen as shown by a "number of turns from viewing to test" ANOVA analysis. Further, we found that participants were generally unsatisfied with the discontinuous nature of the virtual panoramic environment and that, along with the large number of salient landmarks in each panoramic image, may suggest a reason why there was no difference in spatial updating accuracy between conditions.

One of the reasons suggested for the lack of significance was that *jumping* from panorama to panorama brought the user out of the experience and removed necessary path integration and spatial updating cues. One way to mitigate this may be to interpolate between panoramas. Google Earth actually provides this feature. Comparing the performance of participants in this study with the same study but with interpolation between panoramas, would be interesting to compare.

A limitation of our study is that we have not investigated the carryover effect this task may have on wayfinding/spatial updating performance in the real world. It is unclear if a virtual environment, a map, or a VPE would be best for transferring knowledge to the real world, though much research has focused on the former two options [Thorndyke and Hayes-Roth 1982; Satalich 1995].

Future work may focus on this idea. A spherical camera, such as the Point Grey Ladybug2 [Point Grey 2012], could be used to capture panoramic images in a smaller scale environment. Participants could then perform a training task in the VPE representation of the environment, by viewing a map of the environment, or by training in a virtual representation of the environment. Finally, participants could then perform a searching or wayfinding task in the real world. Task performance in the real world could be compared between the conditions.

APPENDICES

Appendix A

Streetview Camera Source Code for THREE.JS - Revision 41/ROME

```
/**
 * SVCamera.js
 */

THREE.SVCamera = function(fov, aspect, near, far, target) {
    THREE.Object3D.call(this);

    this.maxPitchAngle = 0.0;
    this.currentPitchAngle = 0.0;

    this.maxYawAngle = 0.0;
    this.currentYawAngle = 0.0;

    this.fov = fov || 50;
    this.aspect = aspect || 1;
    this.near = near || 0.1;
    this.far = far || 2000;

    this.target = target || new THREE.Object3D();
    this.useTarget = true;

    this.matrixWorldInverse = new THREE.Matrix4();
    this.projectionMatrix = null;

    this.updateProjectionMatrix();
}

THREE.SVCamera.prototype = new THREE.Object3D();
THREE.SVCamera.prototype.constructor = THREE.SVCamera;
THREE.SVCamera.prototype.supr = THREE.Object3D.prototype;

THREE.SVCamera.prototype.translate = function ( distance, axis ) {
       this.matrix.rotateAxis( axis );
       this.position.addSelf( axis.multiplyScalar( distance ) );
       this.target.position.addSelf( axis.multiplyScalar( distance ) );
}

THREE.SVCamera.prototype.updateProjectionMatrix = function () {
       this.projectionMatrix = THREE.Matrix4.makePerspective(this.fov,
this.aspect, this.near, this.far );
}

THREE.SVCamera.prototype.updateMatrix = function () {
       this.update( undefined, true );
}


THREE.SVCamera.prototype.rotateAroundVector = function (angle, axis, absolute ) {
       var currLook = new THREE.Vector3(0,0,0);
       var newLook = new THREE.Vector3(0,0,0);
```

```
        var cosTheta = Math.cos(angle);
        var sinTheta = Math.sin(angle);

        var x = axis.x;
        var y = axis.y;
        var z = axis.z;

        // Use Rodrigues' Rotation Formula to compute the new target point
        // http://mathworld.wolfram.com/RodriguesRotationFormula.html

        var b_absolute = absolute || 0;
        if( b_absolute == 0 )
               currLook.sub( this.target.position, this.position );
        else
               currLook.sub( (new THREE.Object3D()).position, this.position );

        newLook.x  = (cosTheta + (1 - cosTheta) * x * x) * currLook.x;
        newLook.x += ((1 - cosTheta) * x * y - z * sinTheta) * currLook.y;
        newLook.x += ((1 - cosTheta) * x * z + y * sinTheta) * currLook.z;

        newLook.y  = ((1 - cosTheta) * x * y + z * sinTheta) * currLook.x;
        newLook.y += (cosTheta + (1 - cosTheta) * y * y) * currLook.y;
        newLook.y += ((1 - cosTheta) * y * z - x * sinTheta) * currLook.z;

        newLook.z  = ((1 - cosTheta) * x * z - y * sinTheta) * currLook.x;
        newLook.z += ((1 - cosTheta) * y * z + x * sinTheta) * currLook.y;
        newLook.z += (cosTheta + (1 - cosTheta) * z * z) * currLook.z;

        this.target.position.x = this.position.x + newLook.x;
        this.target.position.y = this.position.y + newLook.y;
        this.target.position.z = this.position.z + newLook.z;
}

THREE.SVCamera.prototype.yaw = function ( angle, absolute ) {
        var b_absolute = absolute || 0;

        if( b_absolute == 0 )
               this.currentYawAngle += angle;
        else
               this.currentYawAngle = angle;

        this.rotateAroundVector( angle, this.up, b_absolute );

}

THREE.SVCamera.prototype.pitch = function ( angle, absolute ) {
        var b_absolute = absolute || 0;

        var right = new THREE.Vector3(0,0,0);
        var lookVec = new THREE.Vector3(0,0,0);


        if( b_absolute == 0 ) {
               lookVec.sub( this.target.position, this.position );
```

```
            this.currentPitchAngle += angle;
        }
        else {
            this.currentPitchAngle = angle;
            lookVec.sub( (new THREE.Object3D()).position, this.position );
        }

        right.cross( lookVec, this.up );
        right.normalize();

        if( this.maxPitchAngle > 0 ) {
            if( this.currentPitchAngle > this.maxPitchAngle )
                this.currentPitchAngle = this.maxPitchAngle;
            else if( this.currentPitchAngle < -this.maxPitchAngle )
                this.currentPitchAngle = -this.maxPitchAngle;
            else
                this.rotateAroundVector( angle, right, b_absolute );
        } else {
            this.rotateAroundVector( angle, right, b_absolute );
        }
}

THREE.SVCamera.prototype.pitchTo = function ( angle ) {
        this.pitch(angle, 1);
}

THREE.SVCamera.prototype.yawTo = function ( angle ) {
        this.yaw(angle, 1);
}

THREE.SVCamera.prototype.moveTo = function ( p ) {
        var deltaP = new THREE.Vector3(0,0,0);
        deltaP.sub(p, this.position);

        this.position = p;
        this.target.position = this.position;
        this.target.position.addSelf(deltaP);
}

THREE.SVCamera.prototype.moveBy = function ( p ) {
        this.position.addSelf(p);
        this.target.position.addSelf(p);
}

THREE.SVCamera.prototype.moveForward = function ( howMuch ) {
        var dir = new THREE.Vector3(0,0,0);
        dir.sub(this.target.position, this.position);
        dir.multiplyScalar(howMuch);
        this.position.addSelf(dir);
}

THREE.SVCamera.prototype.update = function ( parentMatrixWorld, forceUpdate,
camera ) {
        if ( this.useTarget ) {
            // local
```

```
		this.matrix.lookAt( this.position, this.target.position, this.up );
		this.matrix.setPosition( this.position );


		// global
	if( parentMatrixWorld ) {
			this.matrixWorld.multiply( parentMatrixWorld, this.matrix );
		} else {
			this.matrixWorld.copy( this.matrix );
		}
		THREE.Matrix4.makeInvert( this.matrixWorld, this.matrixWorldInverse
);

		forceUpdate = true;
	} else {
		if ( this.matrixAutoUpdate ) {
			forceUpdate |= this.updateMatrix();
		}
		if ( forceUpdate || this.matrixWorldNeedsUpdate ) {
			if ( parentMatrixWorld ) {
				this.matrixWorld.multiply( parentMatrixWorld,
this.matrix );
			} else {
				this.matrixWorld.copy( this.matrix );
			}

			this.matrixWorldNeedsUpdate = false;
			forceUpdate = true;

			THREE.Matrix4.makeInvert(this.matrixWorld,
this.matrixWorldInverse );
		}
	}

	// update children
	for ( var i = 0; i < this.children.length; i ++ ) {
		this.children[ i ].update( this.matrixWorld, forceUpdate, camera );
	}
};
```

Appendix B

Pre-questionnaire Administered to Participants

Participant ID _____

Pre Questionnaire

- Date/Time: _______________
- Age: __________
- Gender (*circle*): Male Female
- Profession: ____________________________
- Highest Education level (or current level if still in school): ___________________
- Major: ________________________
- Handedness (*circle*): Right handed Left handed Ambidextrous
- Do you wear (*circle all applicable*): Contact Lenses Glasses
- Which are you wearing today (*circle*): Contact Lenses Glasses
- Average estimated computer use per week: _________ hours
- Do you own a game console (*circle*): Yes No
- Average estimated time playing video games per week? ________ hours
- Are you confident when navigating with the mouse (*circle*): Yes No
- Have you spent more than 2 weeks (total throughout your life) in Washington DC? Yes No
- Have you been to Washington DC in the past 5 years? Yes No

Appendix C

Post-questionnaire

Participant ID __

1. Please rate *your sense of being in the* Washington DC environment, on the following scale from 1 to 7, where 7 represents your normal experience of being in a place.

   I had a sense of "being there"…
   Not at all 1 – 2 – 3 – 4 – 5 – 6 – 7 Very much

2. To what extent were there times during the experience when the Washington DC environment was the reality for you?

   There were times during the experience when the Washington DC environment *was the reality for me…*
   Not at all 1 – 2 – 3 – 4 – 5 – 6 – 7 Almost all the time

3. When you think back about your experience do you think of the Washington DC Environment more as *images that you saw,* or more as *somewhere that you visited?*

   The environment seems to me to be more like …
   Images that I saw 1 – 2 – 3 – 4 – 5 – 6 – 7 Somewhere that I visited

4. During the time of the experience, which was strongest on the whole, your sense of being in the Washington DC environment, or of being elsewhere?

   I had a stronger sense of…
   Being elsewhere 1 – 2 – 3 – 4 – 5 – 6 – 7 Being in the Washington DC environment

5. Consider your memory of being in the Washington DC environment. How similar in terms of the *structure of the memory* is this to the structure of the memory of other *places* you have been today? By "structure of the memory," consider things like the extent to which you have a visual memory of the Washington DC environment, whether that memory is in color, the extent to which the memory seems vivid or realistic, its size, location in your imagination, the extent to which it is panoramic in your imagination, and other such *structural* elements.

I think of the Washington DC environment as a place in a way similar to other places that I've been today…

Not at all 1 – 2 – 3 – 4 – 5 – 6 – 7 Very much so

6. During the time of the experience, did you often think to yourself that you were actually in the Washington DC area?

During the experience I often thought that I was really standing in the Washington DC environment…

Not very often 1 – 2 – 3 – 4 – 5 – 6 – 7 Very much so

7. How accurate do you think you were in predicting the landmark positions?

Not accurate 1 – 2 – 3 – 4 – 5 – 6 – 7 Very accurate

8. Did you use any non-marked objects (trees, cars, people, etc…) to determine the position of previously seen labeled landmarks during the tests (*circle*)? Yes No

9. Did you use any other techniques? Yes No

10. If so, what techniques? ______________________________

11. Comments?

REFERENCES

Batschelet, E. 1965. Statistical methods for the analysis of problems in animal orientation and certain biological rhythms. *American Institute of Biological Science,* Washington D.C.

Batschelete, E. 1978. Second-order statistical analysis of directions. *Animal migration, navigation and homing*, 1-24.

Bowman, D. A., Tech, V., Davis, E. T., Hodges, L. F., and Badre, A. N. 1999. Maintaining spatial orientation during travel in an immersive virtual environment. *Presence: Teleoperators and Virtual Environments*, 8, 6, 618-631.

Bowman, D. A, Kruijff, E., LaViola Jr, J. J., and Poupyrev, I. 2001. An introduction to 3D user interface design. *Presence: Teleoperators and Virtual Environments*, *10, 1, 96-108.*

Chance, S. S., Beall, A. C., and Loomis, J. M. 1998. Locomotion mode affects the updating of objects encountered during travel: The contribution of vestibular and proprioceptive inputs to path integration. *Presence: Teleoperators and Virtual Environments*, *7*, 2, 168-178.

Darken, R. P., and Sibert, J. L. 1996. Navigating large virtual spaces. *International Journal of Human-Computer Interaction*, *8*, 1, 49-71.

Gallistel, C. R. 1990. The organization of learning. Cambridge, MA: MIT Press.

Guilford, J. P., and Zimmerman, W.S. 1948. The Guilford-Zimmerman Aptitude Survey. *Journal of Applied Psychology*, 32, 24-34.

Hart, S. G., and Staveland, L. E. 1988. Development of NASA-TLX (task load index): results of empirical and theoretical research. *Human mental workload*, 1, 139-183.

Hu, J., Fadel, G. M., Wood, I., Napieralski, P., and Babu, S. 2011. Prototyping process of a virtual-reality treadmill system for exploration of real world panoramic environments. In *Proceedings of the 5th international conference on advanced research in virtual and rapid prototyping*, Leiria, Portugal.

Klatzky, R. L., Loomis, J. M., Beall, A. C., Chance, S. S., Reginald, G., and Golledge, R. G. 1998. Spatial updating of self-position and orientation during real, imagined, and virtual locomotion. *Psychological Science*, 9, 4, 293-298.

Mittelstaedt, M. L., and Mittelstaedt, H. 1980. Homing by path integration in a mammal. *Naturwissenschaften,* 67, 566-567.

Poucet, B. 1993. Spatial cognitive maps in animals: new hypotheses on their structure and neural mechanisms. *Psychological Review,* 100, 2, 163-182.

Point Grey. 2012. Ladybug2. Available at http://www.ptgrey.com/products/ladybug2/ladybug2_360_video_camera.asp

Riecke, B. E., Feuereissen, D., Rieser, J. J., and McNamara, T. P. 2012. Self-motion illusions (vection) in VR – are they good for anything? In *Proceedings of the IEEE Virtual Reality,* Orange County, CA.

Riecke, B. E., and Wiener, J. M. 2007. Can people not tell left from right in VR ? Point-to-origin studies revealed qualitative errors in visual path integration. In *Proceedings of IEEE Virtual Reality,* Charlotte, North Carolina.

Satalich, G. A. 1995. Navigation and wayfinding in virtual reality: finding the proper tools and cues to enhance navigational awareness. Thesis for University of Washington.

Sigurdarson, S., Milne, A. P., Feuereissen, D., and Riecke, B. E. 2012. Can Physical Motions Prevent Disorientation in Naturalistic VR? In *Proceedings of the IEEE Vitrtual Reality,* Orange County CA.,

Thorndyke, P. W., and Hayes-Roth, B. 1982. Differences in spatial knowledge acquired from maps and navigation. *Cognitive Psychology*, 14, 4, 560-589.

Turvey, M. T., Burton, G., Pagano, C. C., Solomon, H. Y., and Runeson, S. 1992. Role of the Intertia Tensor in Perceiving Object Orientation by Dynamic Touch. Journal of Experimental Psychology: Human Perception and Performance, 3, 714-727.

Usoh, M., Catena, E., Arman, S., and Slater, M. Using Presence Questionnaires in Reality. Presence: Teleoperators and Virtual environments, 9, 5, 497-503

Waller, D., Hunt, E., and Knapp, D. 1998. The transfer of spatial knowledge in virtual environment training. *Presence: Teleoperators and Virtual Environments*, 7, 2, 129-143.

Wang, R. F., Crowell, J. A., Simons, D. J., Irwin, D. E., Kramer, A. F., Ambinder, M. S., Thomas, L. E., Gosney, J. L., Levinthal, B. R., and Hsieh, B. B. 2006. Spatial Updating on an Egocentric Representation of Space: Effects of the Number of Objects. *Psychonomic Bulletin and Review*, 13, 2, 281-286.

Belongie, S. "Rodrigues' Rotation Formula." MathWorld--A Wolfram Web Resource, created by Eric W. Weisstein. Retrieved April 2012. http://mathworld.wolfram.com/RodriguesRotationFormula.html

Zanbaka, C., Lok, B., Babu, S. V., Xiao, D., Ulinski, A., and Hodges, L. F. 2004. Effects of Travel Techniques on Cognition in Virtual Environments. In *Proceedings of the IEEE Virtual Reality,* Chicago, IL.