Evaluation of Locomotion Techniques in Room and Standing Scale Tracked Spaces

By

Richard Paris

Dissertation
Submitted to the Faculty of the
Graduate School of Vanderbilt University
in partial fulfillment of the requirements
for the degree of

DOCTOR OF PHILOSOPHY
in
Computer Science

May 8th 2020

Nashville, Tennessee

Approved:

Bobby Bodenheimer, Ph.D.

Tim McNamara, Ph.D.

Mary C. Whitton

Shilo Anders, Ph.D.

Matt Berger, Ph.D.

ACKNOWLEDGMENTS

I'd like to thank my each member of my committee for all of their insight, comments, and help in the preparation of this dissertation. I would like to thank Dr. Bodenheimer in particular for his tremendous help in completing all of the work herein. Lauren and Haley, thank you for being my cohorts, my colleagues, my sounding board, my... friends. Miti, Joe, and Sara, thank you each for your help in completing this work. Paul, Conner, Ruida, and Hunter, thank you for your help in my other academic pursuits. You each allowed me to spend so much more time on my dissertation work. To my family, thank you for your love and support. In particular my father who always was there to offer help and support no matter the situation. Finally to every one of the friends I made along the way. I say thank you. You made grad school not just bearable but so much fun and an experience that I'll treasure forever.

TABLE OF CONTENTS

Page

ACKNOWLEDGMENTS ii

LIST OF TABLES vii

LIST OF FIGURES ix

CHAPTER

1 Introduction 1

- 1.1 Specific Aims 3
- 1.2 Overview 4

2 Background 6

- 2.1 Overview 6
- 2.2 Locomotion Methods and Evaluation 8
  - 2.2.1 Body-Based Cues 9
    - 2.2.1.1 Experience and Timing 11
  - 2.2.2 Curvature Manipulation 11
    - 2.2.2.1 Implementation 11
  - 2.2.3 Rotational Manipulation 13
    - 2.2.3.1 Implementation 13
  - 2.2.4 Positional Manipulation 14
  - 2.2.5 Evaluation 16
- 2.3 Spatial Cognition 18
  - 2.3.1 Spatial Cognitive Microgenesis 18
  - 2.3.2 Acquisition of Knowledge 19
    - 2.3.2.1 Components of Working Memory 20
    - 2.3.2.2 Distinguishing Route and Survey Knowledge 22
    - 2.3.2.3 Transference 24
- 2.4 Individual Differences 25
  - 2.4.1 Physiological 25
  - 2.4.2 Self-Report 26
  - 2.4.3 Individual differences in skills 26
  - 2.4.4 Individual differences in strategy 27
- 2.5 Summary 29

3 Specific Aim 1 31

- 3.1 Introduction 31

3.2 Research Design 33
    3.2.1 Methods of Exploration 33
    3.2.2 Path 37
    3.2.3 Performance Measures 37
3.3 Experiment 1 38
    3.3.1 Methods 38
        3.3.1.1 Participants 38
        3.3.1.2 Equipment 38
    3.3.2 Environment 39
    3.3.3 Spatial Awareness Task 39
    3.3.4 Results 42
        3.3.4.1 Spatial Awareness 42
        3.3.4.2 Presence 43
    3.3.5 Discussion 44
3.4 Experiment 2 45
    3.4.1 Methods 45
        3.4.1.1 Participants 45
        3.4.1.2 Equipment 45
    3.4.2 Environment 46
    3.4.3 Navigation and Wayfinding Task 47
    3.4.4 Results 48
    3.4.5 Discussion 52
3.5 General Discussion 53

4 Specific Aim 2 55

4.1 Introduction 55
4.2 Research Design 57
    4.2.1 Techniques of Walking 57
        4.2.1.1 Resetting 57
        4.2.1.2 Walking in Place 57
    4.2.2 Environment 60
    4.2.3 Testing Protocol 63
    4.2.4 Metrics 63
        4.2.4.1 Initial Angular Error 63
        4.2.4.2 Estimated Path Length 64
        4.2.4.3 Recall Time 64
4.3 Experiment 3 64
    4.3.1 Hypotheses 64
    4.3.2 Participants 65
    4.3.3 Equipment 65
    4.3.4 Procedure 66
    4.3.5 Results 67
        4.3.5.1 Initial Angular Error 67
        4.3.5.2 Path Length 68

4.3.5.3 Recall Time 70
4.3.6 Conclusions 70
4.4 Experiment 4 72
4.4.1 Hypotheses 72
4.4.2 ID Tests 73
4.4.3 Participants 73
4.4.4 Equipment 73
4.4.5 Procedure 74
4.4.6 Results 75
4.4.6.1 Path Length 76
4.4.6.2 Initial Angular Error 76
4.4.6.3 Individual Differences 76
4.4.7 Discussion 78
4.4.7.1 Testing Method Effects 78
4.4.7.2 Learning Method Effects 79
4.5 General Discussion 79
4.5.1 Future Work 81

5 Specific Aim 3 82

5.1 Introduction 82
5.2 Research Design 84
5.2.1 Locomotion Methods 84
5.2.2 Sample Size 84
5.2.3 Metrics 85
5.2.4 Tracked Spaces 85
5.3 Methods 85
5.3.1 Participants 86
5.3.2 Equipment 86
5.3.3 Individual Difference Measures 86
5.3.4 Environment 87
5.3.5 Navigation Task 90
5.3.6 Metrics 91
5.3.7 Simulator Sickness 91
5.3.8 Post Test 92
5.4 Hypotheses 92
5.5 Results 93
5.5.1 Angular Error 93
5.5.2 Simulator Sickness 94
5.5.3 Questionnaire Responses 95
5.6 Discussion 96
5.6.1 Future Work 96

6 Conclusion 98

6.1 Future Directions 102
    6.1.1 Individual Differences 102
    6.1.2 Larger and Irregular Spaces 103
    6.1.3 Other Locomotion Affordances 103
    6.1.4 Walking in Place Improvements 104

Glossary 105

BIBLIOGRAPHY 109

LIST OF TABLES

Table Page

2.1 This table shows the availability of the various translational and rotational cues in a number of locomotion modes. This list is not exhaustive, but represents a majority of the methods in recent research. "A" denotes that a cue is available; "*" denotes that a cue is available but not always correct; and "x" denotes that a cue is always incorrect or incompatible because of method. [1] Only Curvature manipulation. 30

3.1 Mean high responses for presence questionnaire of Experiment 1. 44

3.2 Shows a summary per subject of the distance walked by the subject compared to the ideal distance. interventions are the total number of times the subject had to reorient. 51

4.1 Virtual reality based phases of Experiment 4 for the four between subjects conditions. Twenty-four subjects completed each of the above conditions. 74

4.2 Mean and median SBSOD and MRT scores among all participants 77

4.3 Comparison of improvement in angular error between those with weak and strong mental rotation abilities. 78

4.4 Correlations between dependent measures and the individual differences measured in this experiment. 78

5.1 The four conditions in this study as well as the locomotion method (and tracked space if applicable) in each of the three VR phases of the experiment. 85

5.2 Post test questionnaire presented to each subject. Questions 3-6 were masking questions and all 6 questions were presented in a random order. 92

5.3 Post test interview presented to each subject. Followup questions were asked based on responses to each question in this table. 92

5.4 Mean and median SBSOD and MRT scores among all participants 93

5.5 Correlations between dependent measures and the individual differences measured in this experiment. 95

5.6 Number of dropouts due to simulator sickness occurring in each of the four conditions. 95

LIST OF FIGURES

Figure Page

3.1 Real (left) and virtual (right) paths walked during Experiment 1. The paths from top to bottom are reorientation short (7 meters, 0 ROTs), reorientation medium (14 meters, 2 ROTs), reorientation long (21 meters, 4 ROTs), and redirected walking long (21 meters). S represents the start of each path, F the finish, T the turn, and each reorientation is marked by R. Subjects start at S, walk along the marked path to T and perform a 135° turn. Finally, subjects finish the path at point F. 34

3.2 Top-down view of the environment used in Experiment 1 showing all objects, a randomized subset of which were memorized in each trial. 40

3.3 An example of a starting view of a subject. The column marks the target walking location and the mailbox and streetlight were potential markers to be memorized. . 41

3.4 Absolute turning error (and the standard error of the mean) associated with each locomotion method by distance walked. 44

3.5 Top-down view of the environment used in Experiment 2, illustrating both distance and location of the storefronts. The blue line represents a wall boundary around the environment. A sample path starting from the center and the reorientations necessary to walk that path are shown as well. Each arrow represents one length before a reorientation is required. 46

3.6 This figure shows an example storefront and the target column placed in front of the store. Each storefront has an equivalent column. 47

3.7 Total time spent (and the standared error) in the system by condition for each trial. . 49

3.8 Time spent searching, i.e., total time minus time spent during interventions, (and the standared error) by condition for each trial. 50

3.9 Average distance traveled (and the standared error) by condition for each trial. . . . 50

3.10 Average distance traveled by condition for each trial divided by the ideal distance for that trial (and the standared error of that measure). 51

4.1 A demonstration of WiP. Cardboard is placed on the floor to prevent subjects drifting as they walk in place. 58

4.2 Top-down view of the environment used in the practice phase of each experiment. There are four objects for the subject to find and four landmarks (paintings). 60

4.3 Top-down view of the environment used in the learning phase of each experiment. There are eight objects for the subject to find and four landmarks (paintings) to facilitate learning. 61

4.4 This figure shows the telephone booth as would be seen by subjects during the orientation phases of each experiment. 62

4.5 At the beginning of the testing phase subjects are informed of the target object via a heads up display. This disappears shortly so as not to distract the subject during walking. 62

4.6 This figure shows the mean initial angular error of subjects by condition and blocks of 10 trials. 67

4.7 This figure shows the mean distance subjects walked as a percentage of true distance by condition and blocks of 10 trials. 69

4.8 This figure shows the mean time subjects spent determining the relative direction of the target object. Time is plotted by blocks of 10 trials and collapsed across condition. 69

4.9 This figure shows the distance overwalked when testing in each condition. 77

5.1 Top-down view of the environment used in the practice phase of Experiment 1. There are four objects for the subject to find and four landmarks (paintings). 87

5.2 Top-down view of the environment used in the learning phase of Experiment 1. There are eight objects for the subject to find and four landmarks (paintings) to facilitate learning. 88
5.3 This figure shows the telephone booth as would be seen by subjects during the orientation phases of the experiment. 89
5.4 At the beginning of the testing phase subjects are informed of the target object via a heads up display. This disappears shortly so as not to distract the subject during walking. 89
5.5 This chart shows the estimated marginal mean angular error in each of the four conditions and the standard error of the mean. These means take into account the individual differences measured in this experiment. 94

Chapter 1

Introduction

This dissertation evaluates a number of methods that allow for movement in a virtual environment, called locomotion methods, when real world physical space in which to move is limited. Common alternatives to real walking (e.g., teleportation and joystick) have been shown to yield reduced spatial knowledge [105] and/or require the use of a subject's hands. To eliminate those concerns, this work focuses on locomotion methods that are either bipedal (i.e., walking) or an emulation of real walking. Physical (tracked) space can vary due to a number of factors and this work will strengthen our understanding of how differences in the size of the physical space, along with an individual's abilities, affect one's ability to navigate in a virtual environment.

Virtual Reality (VR) has reached the consumer level. People want compelling, interactive, and high quality experiences and those experiences can finally be delivered by commodity VR systems such as those provided by Oculus and HTC. There has been a proliferation of commodity level virtual reality devices in recent years [138]. This availability has brought with it a large demand for interesting and interactive VR experiences. These experiences include things such as museum tours, city/site tours, and games. But these VR systems (e.g., Oculus Rift) lack realism as they are still in their infancy and the goal is to maximize the realism these VR systems invoke.

Navigation is the ability to learn and efficiently explore an environment through direct experience. Navigating is something that occurs naturally when exploring or moving through an environment. It can be accomplished through walking or driving. Requiring that navigation occurs in VR as it does in the real world is an important part of ensuring realism in VR. Our problem lies in that one's ability to navigate is dependent on the locomotion method employed.

Since navigation is the natural result of direct experience, this dissertation will evaluate a number of locomotion methods based on how well they permit users to navigate. Specifically, we are interested in how quickly and correctly users acquire spatial knowledge. Spatial knowledge is what allows a person to recall the location of objects or landmarks relative to themself or another remembered location. In this work we measure spatial knowledge using the framework of "Spatial Cognitive Microgenesis (SCM)" first proposed by Siegel and White [118]. This framework consists of three stages: the highest and most complete of which is called *survey knowledge*. Survey knowledge is sometimes called configural knowledge as it represents the ability to recall the configuration of an environment. We employ survey knowledge as our metric because it is part of a structured framework and has two independent components against which to test. The two components are distance and directional information, each of which can be acquired through visual or body-based information. Visual information provided by any locomotion method is almost completely unchanged, so we focus on body-based information or cues.

Body-based cues are those that are self-generated and do not come from the surrounding environment. We classify body-based cues as either rotational or translational. These cues work in conjunction to permit the acquisition of spatial knowledge. To learn the environment both of these cues are required. Visual cues are provided naturally by nearly all VR systems, however, as we present in this dissertation, the visual and body-based information does not always align. Translational cues give knowledge of how far away things are either relative to one's self or some third object. Rotational cues are those which provide knowledge of the relative directions among objects. While it is difficult to disentangle these two, prior work from the navigation literature indicates that different forms of body-based information are instrumental in the acquisition of the three levels of spatial knowledge. [48, 15, 18]

There has been much research to show the importance of body-based information for navigation [105, 114]. This dissertation will be limited to locomotion methods that provide at least some level of both positional and rotational information. This work will also consider that the physical real world space is limited and varied between users. The locomotion method types that fit our criteria of being similar to real walking, provide both types of information, and work in limited space are redirected walking [99], reorientation [149], and walking in place (WiP) [90]. Redirected walking is a technique that constantly manipulates the user's perception to redirect them away from obstacles. Redirected walking steers a user away from obstacles in an undetectable way by changing their real world heading without affecting their virtual heading. Reorientation is a method of reorienting the user in the real world while, again, leaving their virtual heading unchanged. The final method, WiP, induces virtual linear translation using a gesture similar to real walking. In this way it does not require the subject to physically change position.

1.1 Specific Aims

This work will bring together the areas of space requirements, tracking systems, and individual differences to determine the "best" locomotion method. There is not a one size fits all solution, but we present different solutions that work for different sets of space and tracking availability.

To investigate how navigation is affected by locomotion we have three specific aims.

1. To assess methods for locomoting in large scale virtual environments such as reorientation and redirected walking in terms of how they affect presence, spatial awareness, and spatial knowledge. This assessment involves gathering a comparative measure of how the manipulations required by reorientation and redirected walking affect navigation, way-finding, and spatial orientation. Two tasks will demonstrate the direct effects and the size of those effects: a path integration task and a more complex exploration task. Determining which of these methods has the smaller effect on navigation ability will reveal a walking technique that can be used as a baseline moving forward.

2. To develop and evaluate a standing scale navigation technique. The evaluation of this technique will center around measurements of survey knowledge to determine how well a user is able to learn their environment. Specifically this work will compare the room scale navigation technique of resetting to a WiP method which was developed with mobile VR in mind. We will test whether subjects can acquire and demonstrate interpoint distance and directional components of survey knowledge in the two techniques as well as the extent to which those can be acquired.

3. To assess the effect of room size on the resetting locomotion mode as it relates to spatial awareness, simulator sickness, and individual differences. Resetting differs from walking in place in two ways: it grants additional idiothetic information and requires an intervention. As tracking space shrinks more interventions are required making idiothetic information less reliable. A substantial amount of work on individual differences shows that navigation ability, strategy, and working memory affect the acquisition of the different stages of knowledge. By ascertaining how navigation performance using resetting can be predicted by individual differences in ability and strategy we can begin to study how navigators are affected by resetting.

1.2 Overview

This work will investigate five specific locomotion methods within the three classes of reorientation, redirected walking, and WiP. It will begin to answer the question of which is best in room-scale spaces (i.e., 5x5 m^2 and smaller). Many labs have virtual spaces at or above 10x10 m^2, and second/third generation VR headsets are beginning to support such large spaces. This does not extend to the typical home where we anticipate much use of VR will take place in the near future. In this dissertation we also evaluate the lower ends of consumer level tracked space (i.e., room-scale), including VR without positional tracking.

In order to compare the chosen locomotion methods, this dissertation utilizes spatial knowledge frameworks developed by Kahana and colleagues [71, 87] and Chrastil and Warren [16, 17]. We use these two frameworks to measure how well each locomotion method affords learning and navigating of a virtual environment. When testing against navigation subjects will exhibit a large amount of variation in their navigation performance, so in this work it is important to consider individual differences. The Chrastil and Warren [16] framework allows us to account for individual differences while testing against a standard framework of navigation.

We intend to look further at the importance of the reliability of body-based cues. We intend to determine how the aspects of body-based information and individual differences in ability

[137, 29, 39, 12] affect navigation. In order to test the importance of body-based information we utilize five methods of walking under the three classes: two methods of WiP, two methods of reorientaion, and redirected walking. The two methods of WiP differ in how they track the subjects, one uses accelerometers and one uses an external tracking system (Kinect). Reorientation methods allow 1:1 walking during a majority of exploration. However, these methods require an intervention when approaching a boundary and the two reorientations methods differ in how that intervention occurs. The intervention reorients subjects away from the boundary without changing their overall virtual orientation. Redirected walking continuously manipulates a subject's perception of curvature. These methods of walking have different tracking and space requirements that we will analyze in this paper.

This dissertation is organized into six chapters. Chapter 2 provides an overview of the literature considered important to the development of this dissertation. It explains the relevant background that helped motivate and shape the work presented herein. Chapter 3 is the work completed for Specific Aim 1 to assess the current methods of navigation that are feasible and permit real walking in the room-scale. Chapter 4 examines the differences between the room scale technique of resetting and the two implementations of WiP. This Chapter also presents work that helps distinguish the cause of the differences in survey knowledge measured between WiP and resetting. Chapter 5 examines room size and its effect on navigation performance in virtual environments. Finally, Chapter 6 provides some conclusions and future directions.

Chapter 2

Background

To evaluate locomotion and navigation this dissertation will bring together three expansive areas of research: locomotion methods, spatial cognition, and individual differences. Section 2.1 introduces each research area explains how we will leverage each in accomplishing the specific aims introduced in Chapter 1. In Section 2.2 of this chapter, we discuss the current literature on methods of locomotion in large virtual environments and techniques used to evaluate them. Next Section 2.3 introduces work on SCM and how spatial knowledge is acquired. Specifically we look at the components of active memory used in acquiring knowledge, then we look at body-based cues and transference of spatial knowledge between virtual environments and the real world. Section 2.4 then looks at the individual differences literature on the use of a number of ways for categorizing individual differences. By looking at how individual differences affect the acquisition of spatial knowledge, we can determine which of the various techniques of locomotion discussed best fit an individual, and how to adapt these techniques to each individual. Finally we look at all three areas of research and explain how they will be used in this dissertation. Additionally we provide a table (Table 2.1) to aid in that explanation.

2.1 Overview

A large body of literature exists detailing methods to navigate in large virtual environments, some of which facilitate natural walking [149, 99, 47, 27]. Other implementations simulate walking [135, 157, 150] or permit movement through more abstract metaphors (e.g., [36, 56, 60, 136]). All of these methods have their advantages and disadvantages, so determining the best method requires a number of considerations. Factors such as room size and layout [4], tracking and input technology, the virtual environment [64], performance metrics, etc. dictate which locomotion methods are appropriate. There are various performance metrics used to judge the success of locomotion methods, such as breaks in presence [93], simulator sickness [85, 31, 34], and judgments of relative direction [149].

To assess the performance of various locomotion methods this dissertation examines navigation and wayfinding in virtual environments, particularly looking at spatial cognitive microgenesis (SCM) [118] as an evaluative tool. Of interest is how well locomotion methods afford the acquisition of spatial knowledge of an environment. Quick, accurate, all-encompassing cognitive maps [48] afford us the ability to navigate in large virtual spaces in the way we would the real world. This makes the ability to acquire these maps of great interest. Determining which of the numerous methods of locomotion affords an individual the ability to build the strongest cognitive map for a given situation will allow us to assign the appropriate locomotion method. As this work will detail, spatial cognition is heavily influenced by individual differences which cause large variation in navigation performance [39, 48]. There are large individual differences in how we acquire spatial memory [74, 80, 79, 63, 17, 11], utilize spatial cues [76, 84, 101, 52], and integrate this information into a full cognitive model [83, 82, 78, 72, 81, 48, 39]. While a great deal of work has been done in real world environments and to a lesser extent in virtual environments, no work has been done to examine individual differences in these cognitive aspects as it pertains to locomotion affordances in large virtual environments.

In order to understand and account for the presence of large variation in navigation performance in VR [48, 39], we use quantifiable individual differences that can be used to predict and influence the acquisition, utilization, and integration of spatial knowledge. Biological characteristics such as age and sex have been linked to navigation ability [68, 67, 12, 84, 89, 37, 46, 86, 76, 45]. More advanced cognitive measurements have also been linked to individual differences (e.g., Mental Rotation Test (MRT), Water Level Test (WLT) [65]). In addition there are several metrics which can be measured via questionnaire or by self report. These metrics can be used to categorize good or bad navigators [39], egocentric or allocentric navigators [79, 144, 66], preference for relative or absolute cues [66], and preference for directional or positional cues [141, 12].

By leveraging the work on individual differences and applying it to methods of locomotion, individual differences in how strongly a walking method permits acquisition of spatial knowledge may arise. Locomotion methods that use walking must make some manipulation in order to enable the exploration of an arbitrarily large VE. This information can include positional cues [122, 69, 28, 157, 135, 151, 150, 133], reliable rotational cues [9, 43, 43, 85, 99, 125], and/or proprioceptive information [136, 31, 95] which affects how the user is able to acquire spatial knowledge. Hence, it is important to understand each of these areas of research in order to fully investigate locomotion in large immersive VEs.

2.2 Locomotion Methods and Evaluation

Locomotion methods that provide people the ability to explore large virtual environments are not new. Many methods exist which have various philosophies on the appropriate way to enable extended navigation even when the physical space is limited. To permit exploration some manipulation must be employed to allow for the virtual space to extend beyond the physical. This manipulation can be classed into three types: manipulation of curvature [98], orientation [149, 93], or position[154, 34, 150, 136]. Manipulation of curvature techniques typically fall under a class of locomotion methods called “redirected walking” [98], which involves making the subject’s real world and virtual world curvature different to steer them away from obstacles. Rotation manipulation techniques scale the rotation of a subject to allow for their self generated turning to redirect them away from obstacles at, typically, discrete instances [149, 93]. Finally, translation manipulation involves scaling [154, 34] or inducing [150, 136] translation based on the subject’s self motion to afford for a greater covering of the virtual space.

When considering which locomotion method is appropriate it is important to first decide what criteria will effect that decision. Some factors such as room size and layout [4], tracking and input technology, and the virtual environment itself [64], further complicate this question. For sufficiently sized rooms, redirected walking can be implemented [99, 127], but the question still remains how individuals will respond to these methods. In the remainder of this section we review the state of the art in techniques for navigating large virtual environments which have been developed to fit mainly the unique setups and geometry of a few labs. For example, in our lab, redirected walking is not feasible due to the limited space. With so many commodity level solutions to VR becoming available, it becomes increasingly important to consider that many of these methods can be used but there are many factors to be considered.

Many of the methods presented here are designed to provide as many spatial cues as reasonably possible in the environment given that some trade-off between the real and virtual worlds must be made. The literature on cues is typically framed around body-based cues and what spatial information they can provide. In this work we look at vestibular and proprioceptive body-based cues. Other cues include visual and audio cues, but these cues are generally not affected by the chosen locomotion method. However, the level of convincing vestibular and proprioceptive information provided is dependent on the locomotion method. For this reason we look at body-based cues as the primary means of evaluating and categorizing navigation techniques. Because we are interested in providing these cues, as best we can, we look only at techniques that provide these cues in part.

2.2.1 Body-Based Cues

Extensive research has been compiled related to body-based cues in VR [114, 110, 113, 111, 105, 103, 102, 17, 18, 19], and what they provide. Much of this research has shown that in large or complex virtual environments, the vestibular and proprioceptive cues provided by physically translating is critical in acquiring spatial knowledge [114, 110], and in particular survey knowledge [16, 17, 18]. This effect seems to hold only in large or complex environments, as Riecke et al. [105] looked at acquiring configural knowledge in a small environment and only found rotation to be of any benefit. This section presents the relevant work on how body-based cues can affect spatial knowledge acquired through direct experience.

Ruddle et al. [113] was a pioneering study into the effect of body-based information on navigating in large scale (or complex) virtual environments. In this experiment subjects were to explore a complex environment and were evaluated on their acquisition of survey knowledge in terms of direction and interpoint distance. Subjects were given either visual information only, rotational vestibular cues, or full body-based information. The results of the first experiment showed that translation based cues were important for acquiring survey knowledge. Subjects performed 25% better in the first search and 50% better in the second search indicating that the effect continued to grow, and they were acquiring and continually refining their survey knowledge. Both components of survey knowledge (directional and positional) were superior when full body-based information was provided. The question remained as to whether this effect was due to the transitional cue provided, or if both translation and rotation were necessary. In a second experiment the authors looked at providing either translational or rotational cues. Of note in this experiment is that translation only cues were only necessarily correct in proprioception, but not vestibular, as subjects were on a treadmill to allow for translation to occur. Despite this, the results still indicate that survey knowledge requires translation-based motion, but not rotation-based motion. This however, depends on the extent of the space, as we will see later in this subsection.

In virtual environments there are two ways we can measure the size of the space. We can use scale as it is typically thought of, and measure the length between boundaries in the environment, or we can measure it as an analog to complexity. For example, a parking lot is large in scale but not complexity, while an office with many cubicles may be small in scale but large in complexity. When examining simple environments, where every landmark and decision point exist in the same reference frame, different body-based information is required. Riecke et al. [105] looked at an environment that was simple in terms of geometry, yet complex as a searching task. Subjects were placed in an environment with 15 possible targets and required to search for a hidden object. By removing any directional, orientation, or positional cue, subjects had to rely on their ability to form a map as best as possible. This work showed that only body-based rotation was empirically beneficial to searching in such an environment. This argues against the findings of Ruddle and Lessels [111] and Ruddle and Lessels [112] which found that translational motion was beneficial. Of greater interest to this work is that subjects adjusted their navigation strategy when forced to translate physically. This finding is important because it suggests the ability to enforce upon a subject a particular strategy, and, as we will see later, shifts in strategy can lead to different rates and levels of acquisition of spatial knowledge among differing individuals [29, 59, 37, 142].

2.2.1.1 Experience and Timing

Ruddle et al. [115] conducted a longitudinal study to determine if people were experienced in their interface for travel, and how long it took to acquire that experience. The chosen interfaces involved varying levels of body-based information, similar to Ruddle et al. [113]. The authors further looked at how to evaluate proficiency in the various interfaces. They examined three metrics, traversal time, time stationary, and collisions; all of these had similar patterns. In general performance, as expected, subjects with full body-based information performed better than those with only rotation or no information. This was in large part due to a start and stop pattern of motion that resulted in the poor performance of the conditions with less body-based information. Even with as much training as necessary to reach asymptotic performance, the rotation information and visual information groups had worse performance than those who were given full body-based information. Further it took over twice as long to reach asymptotic performance. Of additional interest is that those who reported high game usage had better performance on all metrics which continued through reaching asymptotic performance. In a second experiment there were four conditions that varied body-based information: none, rotation only, translation only, and both. The time taken for those in the translation only condition was much larger due to the unnaturalness of the device. The general finding of this work is that some methods of navigating require different levels of training to become proficient and that this varies with the type of body-based information available. The authors stress that training should be based on individual performance rather than some nominal amount of time.

2.2.2 Curvature Manipulation

2.2.2.1 Implementation

Manipulating the virtual camera, as discreetly as possible, is the typical method of exploring and navigating in large virtual environments. One class of methods accomplishes this using curvature gain and manipulation. These are called techniques for redirected walking (RDW). Research has shown that people are not sensitive to small deviations in curvature, and in fact naturally will walk in a circle absent any external cues [124]. Leveraging this research Razzaque et al. [99] developed the method of redirected walking and examined its effect on both navigation and presence, finding it to be a convincing method. People did not notice any strange occurrences in the simulation. Redirected walking did not induce any undue simulator sickness. The standard implementation of redirected walking involves three manipulations to the curvature and rotation of the user. To allow some curvature gain, the user is continuously redirected from their current heading, and instead steered toward a different virtual position. The rate of curvature is dependent on the speed of the user. The second modification is to manipulate a user's head rotational gain to rotate the environment around the user to match their virtual heading to a more desirable real world path. The final aspect is to imperceptibly rotate the environment at a fixed small rate when no motion is occurring to allow for a more desirable forward path.

Commodity level room scale VR is typically smaller than the usual space available to researchers[4, 34, 41, 9, 85, 99, 126], which means that the implementation of redirected walking is not practical in its general form for unbounded exploration. Engel et al. [27] looked at "meandering" through a virtual environment. As with redirected walking, the virtual camera slowly turned based on rotation and curvature as a user walking through the environment. In this work, the path was artificially made longer to allow more time for redirection. The major problem with this approach is that it required a great deal of unnecessary walking in the virtual environment, which can result in increased fatigue and simulator sickness. The additional time cost and small amount of space covered during nominal times in this method can also lead to the poor acquisition of spatial knowledge. Langbehn et al. [64] looked at improving the path by allowing for a significantly less circuitous route and implementing curvature gain as users walked along an already curved path. This allowed for greater than normal curvature and the system could be contained in the HTC Vive's limited 4m x 4m tracked space.

2.2.3 Rotational Manipulation

2.2.3.1 Implementation

Overt manipulation of rotation has been shown to be tolerable to subjects [61, 97], which has allowed researchers to develop methods of reorienting users away from obstacles and boundaries. Maintaining orientation in a virtual environment can be difficult [109, 1, 96] and doing so requires subjects to utilize various visual, proprioceptive, and vestibular cues. Rieser et al. [106] and Kuhl et al. [62] have shown that visual cues can dominate over the proprioceptive and vestibular cues when in conflict; this finding allows for reorientation to be believable and natural. Nitzsche et al. [88] presented a method that rotated the environment 180 degrees at the boundary irrespective of a subject's bodily input.

Expanding on this work, Williams et al. [149] developed two methods that affect rotational gain when an obstacle (or boundary) is reached. Two methods presented in the paper fall into a category of discrete overt manipulations to the environment, which means that they require overt action at discrete occurrences. The first system presented is a freeze-turn in which the viewpoint would be locked to the heading when a reorientation was necessary. The subject would then turn 180 degrees at which time the viewpoint would unlock. This method was found to be the worst in their evaluation as it put the cues in conflict. A second method which places those cues in only small conflict called 2:1 turn would continue to update the viewpoint, but scale the rotation about the body by 2. This allowed for the same process of a 180 degree turn, but with the added benefit that some visual feedback occurred to reinforce the turn. In the virtual world the turn was 360 degrees leading subjects to believe that they had no variation in heading beyond turning around. The results of this paper showed that there is some cognitive cost associated with resetting and updating spatial information. Maintaining spatial heading is important in the acquisition of spatial knowledge.

Similar research was performed in this field by Razzaque [98], who looked at rotating the environment while subjects slowly looked left and right. This had the effect of reorienting the environment which the subject had to follow. Peck et al. [93] extended this work by adding a distractor to give a natural metaphor as a reason to turn around. In this work they placed an object in front of the user that rotated left and right about the user's head to force rotation upon them; this allowed the environment to rotate imperceptibly. Distractors were shown to increase presence, as there was a reason for the interruption.

The distractors present in Peck et al. [93] are still somewhat unnatural, as they have no reason to be there other than to distract. Yu et al. [156] looked at modifying the environment and dividing it into cells to give a more natural metaphor for reorientation. By placing bookshelves in every room they limited each room to the size of the tracked space and forced reorientation to occur based on user action. Whenever a user wanted to leave the room, he would move to a bookshelf in the room which would, in the virtual world, rotate the user to the next room. In the real world, however, the user has not moved, but must now turn around to face the new room (or cell).

2.2.4 Positional Manipulation

Full rotation is simple and can be done in standing space. Much of the work in the prior section only manipulates rotational cues as a pathway for granting additional translation. In this section we look at ways in which we can affect the translation directly to achieve the same effect. In general we can either induce translation based on user input/system controls or we can manipulate the rate of optic flow of the system. Direct input to afford translation (e.g., joystick locomotion) has been examined more in the literature [95, 55, 136].

The simplest solution to this issue is to give direct control in the form of a controller or keyboard to the user. Joysticks have been utilized to induce translation [95], but this is not a strong method of acquiring spatial knowledge [105]. This may be because of a mismatch in proprioceptive cues in walking and the expectation that walking should afford those cues. Flying, however, has no such expectation and is therefore used as a metaphor to explain how you are able to move freely without walking. This can be combined with more body-based motion to place controls such as speed into the system. Thus, the metaphor becomes more complete and body-based control is given a more salient and controllable cue.

Other more advanced methods of inducing translation from body-based motion have been developed. Leaning was shown to be useful here as well. Kitson et al. [56, 55] have developed several techniques that allow translational motion from leaning. This mode is expected to provide additional vestibular information that is indirectly associated with translational motion. In fact these cues are more akin to a car accelerating. Therein lies their use as large virtual environments may require a different locomotion method, like a car, to explore. The development of these leaning based interfaces typically involves using a chair as a locomotion interface.

Another technique that gives a smaller range of translation is walking in place. This technique involves turning the proprioceptive cue of real walking, but not the vestibular, into translation. To detect real walking we take information from three sources: head motion [135], arm motion [152], and leg motion [150, 151, 136]. Leg motion is the most natural as real walking requires this motion, and thus a great amount of research deals with how to capture this information and how to use it to infer real walking. Less research into how to use arm motion in particular exists, as Wilson et al. [152] have shown this method to be inferior in measures of spatial cognition. However, head motion shows promise as Tregillus and Folmer [135] and Paris et al. [90] have shown that this motion can be as good as leg motion. Additionally, head motion has the benefit that it is natively detectable in mobile VR using accelerometers similar to how pedometers capture walking information.

In spaces slightly larger than standing scale, real walking still dominates [105], and so methods that increase scale by facilitating real walking show promise. In these methods mid-scale translations are afforded by controls of the system, while more fine-tuned motions are still performed by real walking. This can be done instantaneously as in Freitag et al. [31] and [138] but this can be disorienting and confusing. The optic flow present in non instantaneous teleportations have been shown to grant better spatial awareness. Beyond moving continuously through the environment Yu et al. [156] have developed a naturalistic metaphor to explain this generated motion in an effort to increase presence and plausibility. In their work a bird carries you from cell to cell, increasing the multiplicity of the real world tracked space. This requires that you re-position yourself in the virtual world.

A method orthogonal to those presented in the prior paragraph is to add translational gain. This expands the virtually explorable space before some external intervention is required. Several studies (e.g., Rieser et al. [106]) have shown a lack of sensitivity to increasing translation gain, i.e., that users can re-calibrate to increased optic flow and believe they are inducing all motion. There still remains the question of how much gain can be placed in the system. Williams et al. [148] tested various levels of gain and found that subjects could even tolerate 50 times normal translation and still maintain their spatial orientation.

2.2.5 Evaluation

Some investigations have evaluated the relative merits of these locomotion modalities. This section looks at how various studies have evaluated locomotion methods that implement positional manipulations (e.g., Paris et al. [90]); these studies have evaluated these locomotion methods based on factors such as spatial cognition, spatial orientation, presence, and simulator sickness. There are two components to consider when evaluating navigation abilities. We can compare modalities and which modality affords greater spatial knowledge or we can look at how to improve and refine a single modality to grant certain aspects of spatial knowledge or reduce the cognitive cost associated with them.

First we look at comparing across navigation types; several studies have looked at how walking in place compares to various forms of real walking. Paris et al. [90] found that reorientation provides a better sense of scale than does walking in place. Peck et al. [94] found that the cognitive cost of distractors did not detract from spatial knowledge and ability to form a cognitive map as much as walking in place did. Several studies have found walking in place to be superior to joystick-based motion in terms of presence [93, 92, 120, 121] and spatial abilities [94, 105, 114, 115]. Redirected walking has been shown to be similar to real walking under certain conditions [127, 32]. These studies show a hierarchy that roughly correlates to various levels of idiothetic information providing certain levels of spatial cognition.

In addition to cross method comparisons, we can look at refining the existing methods to allow increased virtual space coverage, heightened presence, better spatial knowledge, etc. There is debate and research into how to implement redirected walking; the targets to steer towards or away from [43], the amount of curvature [85], and even how to detect when the curvature is too much [127]. Hodgson et al. [43] looked at various ways to steer individuals in search of a generalized redirected walking method. They found that steering to an orbit resulted in the fewest boundary events. This follows conceptually as it keeps the subject steering away from the boundary and toward the center. Room size is an important consideration in single user redirected walking, where the detection threshold is large compared to the size of standard rooms. Azmandian et al. [4] looked at how to steer people when multiple users are present. Their paper presents some guidelines on how to implement multi-user redirected walking on factors such as room size and shape.

Studies show detection thresholds have large variation based on criteria and testing methodology ranging from curvature radii of 3m to 22m [4, 34, 41, 9, 99, 126]. Neth et al. [85] looked at how speed affected this measurement and found that people are significantly less sensitive towards walking on a curved path when walking more slowly. In their second experiment they looked at implementing a velocity dependent gain controller and found it to cause fewer interventions. Looking at other ways to give more space in redirected walking, Steinicke et al. [127] found that distance can be subtly compressed by 14% or expanded by 26% while unnoticed by the user to allow for longer virtual distances to be traveled before the necessity of a reorientation. Rotation can be scaled by 50% greater or 20% smaller without the user noticing to allow for quicker reorientation towards the desired heading. This would occur either during a boundary reorientation or the baseline rotation occurring while stationary.

2.3 Spatial Cognition

2.3.1 Spatial Cognitive Microgenesis

In the SCM framework, a person first obtains landmark knowledge, which is an understanding of large landmarks (objects that have high visibility and distinctiveness) in the environment. This stage is merely the ability to recognize these landmarks. This first stage of knowledge is thought to develop quickly. However, it does not provide much in the way of path integration – which would allow for shortcuts from landmark to landmark– because, in this stage, one only knows the order of landmarks along previously traveled paths.

The second stage of knowledge, route knowledge, is the stage at which one begins to develop a sense of direction relative to landmarks. This form of knowledge allows for description such as "after passing the water tower" or "turn west when you reach the windmill." This knowledge is described as sequential or ordinal, meaning that people know the relative order of these landmarks on a route, but do not have a good sense of their relative locations or inter distances. During this stage of knowledge acquisition one is thought to have no concept of metric knowledge, which is developed in the next stage.

The final stage of knowledge is survey. This stage implies the development of a survey map of all landmarks (i.e., one with metric information rather than simply direction information between known landmarks). One attains this multi-part metric knowledge through experience. One first gains the ability to recount the distances along traveled paths and later, through the process of path integration, develops the ability to determine direction and distance between landmarks even in the absence of prior direct travel between the two landmarks. At this point one should be able to accurately describe the entire environment with appropriate distance and direction. Survey knowledge can be encoded and recalled either egocentrically or allocentriclly with the latter being more flexible [144]. Egocentric knowledge is a self-to-object representation whereas allocentric is an object-to-object representation [117].

2.3.2 Acquisition of Knowledge

With direct experience people are able to acquire, encode, store, and later recall knowledge of an environment. In this section, we discuss some theories related to the acquisition of knowledge. Siegel and White [118] posited that the stages of knowledge were sequential building blocks requiring the previous stage to be fully developed before the next stage of knowledge can begin to form. This hypothesis has been shown to be false, at least in some individuals. A later proposal by Montello [77] suggested the continuous framework in an attempt to explain how people were able to report accurate metric information before having fully developed route knowledge of a given route. The continuous framework is one in which individuals acquire all three stages of knowledge simultaneously. Using this framework as a guide Ishikawa and Montello [48] showed that individuals could acquire all three types of knowledge simultaneously (with the latter stages, route and survey, developing more slowly). Ishikawa and Montello [48] also showed that individual differences, rather than aggregate error, were the driving force of high variance. They performed a dis-aggregate analysis and found large individual differences in the accuracy and developmental pattern of spatial knowledge. Half of the subjects improved in performance after repeated exposure to the environment while the other half did not, indicating the need for analyses that look at individual differences in performance. Even when examining how well subjects improved, individual differences were seen. Some subjects acquired their maximum level of survey knowledge within the first few exposures while others continued to improve after each exposure.

Acquisition of knowledge also depends on the avenue of learning [83]. Münzer et al. [83] investigated computer-assisted navigation and its effect on the acquisition of route and survey knowledge. In their study, subjects navigated through a zoo along a predefined route of 16 decision points and 15 segments. In one group, an unoriented map fragment instructed subjects as to the correct direction, meaning said subjects had to consider orientation when determining which direction to turn. The other three groups were given a personal digital assistant (PDA) that did one of three things. Each of the three PDA conditions displayed a picture of the intersection. The PDA either displayed the correct turning direction and a map of the two path segments, verbally commanded the correct direction and displayed a map of the two segments, or verbally commanded the correct direction. The results showed that the navigation assistance users had good route knowledge but poor survey knowledge. The map users, however, had significantly better survey knowledge and nearly perfect route knowledge.

2.3.2.1 Components of Working Memory

Meilinger et al. [74] looked at how verbal, spatial, and visual subcomponents of working memory are used in one's encoding process of spatial information. To determine the effect each has on working memory when acquiring spatial knowledge, the authors utilized one of three tasks to suppress each of the subcomponents during the acquisition of spatial knowledge. Subjects were passively taken through a route with one or none of the suppression tasks occurring. They then had to use a joystick to recreate the route. The primary measure here was when the subject was lost in the environment. The results of this paper showed markedly different performance in the four conditions, which indicate that each of these subsystems has some effect on spatial cognition. The absence of the verbal subcomponent of working memory led to the greatest number of incidences of getting lost. These effects are investigated more deeply in further research [74, 63, 142, 29].

Substantial literature has shown that both verbal and spatial subcomponents of working memory are instrumental to the acquisition and encoding of spatial memory [63, 83, 39, 144]. To determine which of these two components of working memory permit stronger learning, Labate et al. [63] investigated spatial learning and suppressed either the verbal component or spatial component of working memory. Subjects learned a route covering two floors and eight landmarks. To complete the given route, subjects were instructed to follow the experimenter and pay attention to the landmarks indicated along the way. To suppress spatial memory subjects tapped a given pattern on an Android device. Researchers suppressed verbal memory by having subjects repeat a series of five syllables: ba-be-bi-bo-bu. A third group of control subjects had no dual task and learned without intervention. The results showed that spatial suppression interfered with the acquisition of spatial survey knowledge in all three measures (pointing, shortcut, map completion). The effect of suppressing verbal memory was more mixed, but showed some instances of degraded survey knowledge. In particular, the harder spatial tasks showed worse performance for the verbally suppressed group.

With any measure of spatial knowledge, care must be taken to account for individual differences. Wen et al. [144] looked at how good and bad navigators differed in the sub-components employed to encode of spatial knowledge. Subjects learned a route by watching a video of one of four routes each over 1 km in length. While learning the routes, subjects were alerted to each landmark. To elucidate the effect each of the three components (verbal, spatial, visual) had on the acquisition of spatial knowledge, subjects were given one of three concurrent interference tasks. To mask verbal memory, subjects heard two syllables and stated if those two formed a word. The visual interference condition required subjects to imagine the face of a clock at a time stated by the observer. Finally, the masking of spatial memory involved locating the direction of a noise from one of three directions. A final condition with no interference served as a control group. Subjects with good sense of direction acquired egocentric survey knowledge in verbal and spatial memory and used all three working memory components to transform this knowledge into allocentric survey knowledge. The authors found that distances were processed in verbal and spatial memory, whereas directions were in visual and spatial memory. Poor navigators, i.e., people with low sense of direction, relied on verbal working memory and lacked spatial processing meaning that they never acquired accurate survey knowledge.

In the previous work, researchers blocked specific components of memory to determine which components different navigators used. Wen et al. [145], however, sought to determine the effect of inducing good and poor navigators to use one of two strategies for spatial navigation. In this study, subjects learned a route by watching a video taken from a car driving along a route. To induce verbal learning, subjects were asked to mention things they noticed or remembered. To induce spatial learning, subjects were instructed to rotate a toy car in line with the turns of the video and place a small item near the toy car in the same relative location of the landmark they were learning. The authors found that, in line with Wen et al. [144], good navigators' landmark learning used both spatial and verbal components of memory, whereas that of poor navigators used only the verbal component. Verbalization showed a disruptive effect on survey learning, which the authors suggested may be due to the subject's verbalizing features of the buildings rather than configural knowledge of the environment. This work further investigated individual differences in the suppression of these subcomponents. In survey learning, good navigators were more affected by the concurrent spatial operation; poor navigators, however, were unaffected by forcing verbalization upon them. Those findings provide further results into how good and bad navigators are able to encode spatial information.

2.3.2.2 Distinguishing Route and Survey Knowledge

Buchner and Jansen-Osmann [10] further called into question the dominant framework as it relates to route knowledge and its lack of metric information. In their paper they performed two experiments. The first experiment presented landmarks to subjects in either dynamic or static format, i.e., subjects either moved through the environment or they were presented with landmarks in serial fashion. To control for length of path segment, each subject was presented with each landmark in the serial condition for the same amount of time they would have seen the landmark in the corresponding dynamic condition. The second factor was context; subjects experiencing both the dynamic and static presentations either had no context, meaning that the only landmarks were visible, or with context in which subjects could see the ground, walls, and sky. These two conditions formed a 2x2 study with four groups. From this experiment they found that neither dynamic presentation nor spatial context in isolation, but only their combination afforded superior route knowledge. In their second experiment the authors further examined only the dynamic with context and the static without context groups but made all path lengths equal. In this experiment no difference between the groups was found. The authors concluded that this finding implies that metric information is a part of route knowledge and is developed in very early exposures to a route.

The acquisition of spatial knowledge can vary based on the activeness in learning [14] and is largely related to the body-based cues present during learning [114, 14, 17, 18, 105]. Chrastil and

Warren [17] looked at this issue, specifically examining the role of vestibular and proprioceptive information as it relates to decision making. The authors had six conditions, three in which the participants made decisions about where to go in a maze and three in which subjects followed the path of a corresponding subject in the active condition. Subjects either watched a video, were moved around in a wheel chair, or walked naturally giving either visual; visual and vestibular; or visual, vestibular, and proprioceptive information respectively. Subjects traversed a maze with eight salient landmarks and a series of interconnected routes to test their acquisition of survey knowledge. The results showed that the presence of a proprioceptive cue yielded significantly better survey acquisition, while the presence of a vestibular cue did not.

Adapting to this result that route knowledge is dependent on metric information, Chrastil [15] introduced a fourth form of knowledge, situated between route and survey, which attempts to explain this phenomena. Graph knowledge is hypothesized to be a mental representation of the environment's multiple routes in a graph like structure. It is a form of knowledge that allows people to take novel shorter paths between landmarks. In their 2015 work, Chrastil and Warren [18] examined the acquisition of graph knowledge in active versus passive learning. Subjects were either guided through a maze with eight landmarks and many interconnected routes, or they were allowed to explore freely the maze. The authors found that active decision making improved the acquisition of graph knowledge, but not survey knowledge. Taken with previous results, this finding indicates that different cues are important for acquiring different types of spatial knowledge. The importance of this result is that, as we will see later, different people acquire different forms of knowledge at different rates. By looking at what type of knowledge an individual needs to be a successful navigator and what type of knowledge they will struggle with obtaining, we can decide which method of navigation may be appropriate based on which cues it provides.

Further delineation between route and survey learning can be seen in Meilinger et al. [75], which examined how navigators encoded route information and survey information. The authors performed a correlation between measures of route and survey knowledge and found little to no correlation. It is important to note that this result was for a highly familiar space, as subjects walked through their own virtual town. The work demonstrated that route knowledge is not a single allocentric north-up frame like survey knowledge [30]. Route knowledge likely consists of a number of different reference frames integrated together. This finding is further support for the different skills and information needed to acquire, encode, and recall spatial knowledge.

2.3.2.3 Transference

When exploring virtual environments, it is important to consider transference of skills and learning. To this end, Waller et al. [139] examined the acquisition of the three stages of knowledge in a virtual environment with varying factors such as fidelity. Waller et al. [139] looked at providing varying amounts of information and exposure to a maze-like environment to investigate the effects of acquiring spatial knowledge from this information. The control condition was given no exposure to the maze, some subjects experienced the real maze, some explored a virtual maze, and lastly some were given a top down map of the maze. The results showed that people were able to develop useful representations of the environment mentally. Additionally, while short exposures led to training in an immersive virtual environment that was no better than a map or desktop VR, longer exposures allowed for acquisition of route knowledge that was indistinguishable from real exposure. However, survey knowledge was not, in general, acquired in this setup; later research has since been performed to show that the acquisition of survey knowledge is possible and likely dependent on the use of vestibular cues present in real walking and its analogs [90, 73, 16].

There are many factors that can influence the acquisition of spatial knowledge in virtual environments. Meijer et al. [73] examined how visual fidelity affects this acquisition. The authors placed subjects in a photorealistic environment and one of reduced realism. The results showed that virtual realism increased spatial cognition of the environment; however, this result was somewhat limited to route knowledge, as the authors suggested that survey knowledge was less likely to be present in a subject's cognitive model.

2.4 Individual Differences

Hegarty et al. [39] looked at characterizing individual differences in environmental spatial abilities to determine if they reflect a single underlying ability or a disparate set of abilities. The authors found that measures of environmental learning define separable factors that were characterized by whether the environment was experienced directly or using some visual medium (e.g., a virtual environment). They suspected that direct experience and visual media load on different factors due to the lack of body-based information. All of the above then indicate that there are individual differences in how subjects use visual cues that are not related to how subjects use other sensory based cues. Hegarty et al. [39] also looked at how large-scale spatial abilities were affected by small-scale spatial abilities, spatial updating ability, and verbal ability. Of interest is that small-scale spatial ability is a strong predictor of large-scale learning and is even more so for subjects who experienced the environment visually (i.e., in the virtual environment). Sense of direction, while also a strong indicator, was not as strong an indicator for the VR condition. This work was some of the first to examine the large individual differences in performance and indicate that some people take much longer to acquire any level of spatial knowledge.

2.4.1 Physiological

At this stage we delve further into ways of categorizing individual differences as a way of better understanding why there is such variance among performance. We have discussed some individual differences such as cultural [58, 155] and touched briefly on reported levels of difference in ability to navigate. For the remainder of this discussion we will be looking at several classes of individual differences. The first and easiest to determine are physiological differences. For example, scores of papers have found a significant sex bias. Men tend to perform better than women in many of the sense of direction tasks and metrics [12]. The gender differences discussed earlier also apply to young children. Some work has been done in the development of cognitive maps of children that shows that the differences exist as early as age eight [45].

2.4.2 Self-Report

Additionally there are several questionnaires that rely on self reporting navigation ability. Many of these surveys are to determine if the navigators believe themselves to be good or bad navigators (SBSOD [38], PSAS [40], PVAS [40]). Others, such as a survey of video game experience or occupation, can give insight into how much each person makes use of their ability to create these cognitive maps [132]. These are ways of differentiating experience. We also have surveys that reveal directly how people think about directions. For example the Spatial Updated Heading test[6] can be used to determine if a person updates their theoretical heading. Finally we have various cultural differences that can indirectly lead to individual differences. For example certain languages have words for uphill and downhill but nothing for the lateral direction of the hill [143]. Still other languages do not make use of left versus right and rely primarily on cardinal directions [24].

2.4.3 Individual differences in skills

We can also examine differences in other cognitive abilities. Several different factors of spatial ability have been identified, and these tests task the user with imagining certain actions be taken and report the result. For example, the water level test presents slanted containers of water and one must state the water level that would be present if the container were to be level. These tests may also measure the ability to imagine the shape of objects for other viewing angles. The mental rotation task, for example, asks people to report which two of a number of objects are the same object [137]. These similar objects differ only through rotation. As they are conceptualized, tests of spatial visualization and tests of spatial orientation involve different types of mental spatial transformations. These transformations require the viewer to update relations between three different spatial frames of reference; the intrinsic reference frame of objects, the egocentric reference frame, and the reference frame of the environment.

Fields and Shelton [29] looked at individual differences in how people apply the same skills to acquire spatial knowledge through different strategies of learning. The authors looked at route and survey processing and found that, while the two strategies used the same set of skills, the two methods utilize them in different ways. This work showed that the two perspectives of learning, route and survey, were affected by individual spatial skills. This indicates that while survey learning (i.e., using a map) generally leads to better performance, we may be able to predict based on these skills the strength of that effect. From these results we can see that different people are affected differently by the inclusion of a global map, and that while this may be beneficial to some, others may exhibit more benefit from other techniques.

Arnold et al. [2] looked at how five measures of spatial knowledge correlated with one another. Their first finding was that performance on the cognitive map formation task was highly related to all other four orientation tests. The authors further found that some individuals were able to successfully form survey representations with very little exposure. Due to the high degree of correlation between path integration ability and cognitive map formation ability, the majority of survey knowledge is due to path integration. Further, these results suggest that only visual information is necessary for path integration. The major important result from this paper is we can more accurately predict spatial navigation ability from various other orientation skills. This gives us the ability to predict good and bad navigators more accurately, an important ability as it provides a good deal of power as various studies have found links between strong and poor navigation skills and cues which can be utilized.

2.4.4 Individual differences in strategy

When encoding a large virtual environment strategy is just as important as skill. Work by Kraemer et al. [59] looked at how applying either a visual or verbal strategy predicted the level and type of spatial knowledge acquired. In their first experiment the authors had subjects watch videos of routes through four virtual cities and tested them on their landmark and survey knowledge. The authors used the Object/Spatial and Verbal Questionnaire [5] to assign subjects to a verbal style or visual style of learning. The results of the first experiment showed that those who utilized the verbal style of learning had stronger landmark knowledge, whereas those who had more of a visual style of learning exhibited higher levels of survey knowledge. To test if this difference in performance was due to strategy or individual differences, the authors performed a second study to coax subjects into using one of the two strategies. The results of this second experiment showed that the strategy used is of greater importance to how the knowledge is encoded. The authors find that this means that the employed strategy can be manually adapted to facilitate the type of learning desired. The author further stress that this enhances the need to consider multiple sources of individual differences when considering spatial cognition.

The previous study examined passive learning relative to spatial cognition; however, as discussed in an earlier section, the strategy employed in active learning plays an important role in individual differences in spatial cognitive ability. Lawton [66] looked at the role of orientation in wayfinding and the specific role that absolute cues versus relative cues plays in the context of individual differences. She examined two strategies of wayfinding, route (relative) and orientation (absolute), and how these two strategies could be employed in either outdoor or indoor environments. Prior work had shown that these two strategies were the dominant methods of wayfinding. Using two questionnaires Lawton categorized subjects into one of the two strategies for both indoor (i.e., inside buildings) and outdoor (e.g., driving through a city) environments. Further, there was a trend towards using similar strategies for both indoor and outdoor navigation. There was, however, some correlation between the two strategies indicating that individuals either switch between strategies depending on the context or may employ both strategies in conjunction to complete a task.

Serino and Riva [116] looked at how direct experience, interactive viewpoint dependent maps, and static maps interact when encoding and retrieving spatial information. The novelty being that this paper examines this during real time presentation. Two groups of subjects were permitted to explore a virtual environment with or without an interactive aerial view. The testing phase then had three crossed conditions where subjects searched for a missing target with the interactive aerial view, without the aerial view, and on a map. The results showed that the presence of an interactive aerial view facilitated the retrieval of spatial information. The results further showed that individuals who preferred to use an allocentric reference frame tended to be less precise in retrieving the spatial locations of the objects in the absence the aerial map. They were more accurate when retrieving the object on a map. This indicates that there are individual differences in how people are able to use additional spatial information to aid encoding and retrieval of information.

2.5 Summary

This chapter provided an in depth review of both locomotion and navigation. In our review of navigation we pay particular attention to individual differences as they have been shown to be an important cause of variation in performance [39]. The differences we will focus on are confidence (SBSOD), mental rotation ability (MRT), and working memory (CORSI). The work will also look at measuring strategies for exploration, navigation, learning, and recall.

The goal of this dissertation is to discuss and describe how to provide satisfying locomotion that affords maximal spatial knowledge to the end user. The navigation and individual differences literature discussed in this chapter covered how we will assess locomotion and what factors we will consider to reduce variation in navigation performance, but this dissertation is primarily focused on comparing a number of locomotion options. This chapter covered body-based cues and how we can use them to categorize the methods presented here. Table 2.1 reviews the methods and cues covered in this chapter. We choose to evaluate methods that provide as many of those cues as possible and fit into each of the manipulations covered.

Table 2.1 divides the available cues first into translational and rotational. Translational cues provide information about how much a subject has translated, while rotational cues provide information about how much the subject has rotated. These cues are with respect to one's self but can be used to learn the layout of a complex environment [105]. Each locomotion method provides these cues in some way and in this chapter we have described the manipulation required to allow each of them. The remaining chapters of this dissertation will investigate a selection of those methods in depth. These selected methods were chosen to provide the greatest amount of body-based information for each of the three manipulations.

Locomotion Method	**Translational Cue**			**Rotational Cue**		
	Vestibular	Visual	Proprioceptive	Vestibular	Visual	Proprioceptive
Redirected Walking[1] [99]	A	A	A	x	A	x
Resetting [149]	A	A	A	*	A	*
Translational Gain [47]	x	A	x	A	A	A
WiP [90]		A	x	A	A	A
Controller [91]		A			A	
Controller w/Body-Based Turns [21]		A		A	A	A
Arm Cycling [152]	A	A	A	x	A	x
Grappling Hook [91]	A	A	A	x	A	x
Teleporting [31]				A	A	A

Table 2.1: This table shows the availability of the various translational and rotational cues in a number of locomotion modes. This list is not exhaustive, but represents a majority of the methods in recent research. “A” denotes that a cue is available; “*” denotes that a cue is available but not always correct; and “x” denotes that a cue is always incorrect or incompatible because of method.

[1] Only Curvature manipulation.

Chapter 3

Specific Aim 1

3.1 Introduction

Walking through virtual environments (VEs) is a natural mode of locomotion and provides tangible benefits over other forms of locomotion, such as using a motion controller [13, 105, 112] or walking in place [134, 150, 136, 28, 146, 90]. For example it provides increased spatial awareness and wayfinding ability [136]. Broadly speaking, there are two sources of spatial information available when moving through an environment: vision-based and proprioceptive[1]. Walking uses these sources in a natural way. Other methods of movement through a VE may not, but may be necessary when the physical space requirements of the virtual reality equipment do not match the space requirements of the VE itself. Usually, such a mismatch occurs because the scale of the VE is substantially larger than the scale of the tracked physical space. Motion controllers like joysticks offer a ready solution to the movement problem, but suppress body-based information critical for navigation; a large body of literature shows that way-finding and navigation performance is impaired in this type of system [13, 105, 112, 128]. Omnidirectional treadmills [44, 49, 123] allow for full use of vision and some use of body-based information, but are usually quite expensive, as consumer-level versions of these devices have proven elusive.

Of specific interest to this paper are locomotion modes involving walking that distort or suppress body-based information to allow locomotion through large virtual environments. Such modes rely on the fact that when vision-based information conflicts with body-based information, vision often dominates. One class of methods manipulate the virtual camera in the translational direction, scaling the optical flow [126, 140, 148] or the “seven league boots” method of Interrante et al. [47]. Since space is only scaled in these methods, unfortunately they cannot accommodate arbitrarily large virtual spaces. Alternatively, redirected walking and reorientation techniques al-

---

[1] Acoustic information can also be important [3], but is a significantly weaker cue [104].

low users to walk continuously (with possible interruptions) through large virtual environments by changing the users' views (often orientation) into the VE so that their actions in the physical world are mapped onto different actions in the VE [8, 42, 88, 94, 99, 149].

This paper compares three specific techniques of the redirected walking and reorientation type in their consequences for spatial learning and way-finding in VEs that are significantly larger than the physical space that a user can walk in. We implement versions of the "steer to circle" technique of Hodgson and Bachmann [42] first implemented in Razzaque et al. [98], the "reorientation using rotational gain" or "resetting" technique of Williams et al. [149] and Xie et al. [153], and the "reorientation with distractors" technique of Peck et al. [94, 93, 92]. Our implementations are described in Section 3.2. One significant note is that our physical tracking space is 5m x 4m, which is considerably smaller than the space used by Hodgson and Bachmann.

When choosing a locomotion method it is important to consider the available space. In large spaces redirected walking with small curvature gains is possible [98] and allows one to maintain spatial awareness [129]. Its performance in smaller spaces is relatively unknown in the current literature. Other studies [127, 85] have shown that the curvature required to fit within our space is much greater than the minimum detection thresholds. Given this, we would expect subjects to immediately notice the manipulation but we still wish to understand how redirected walking affects spatial awareness in such a small space.

The remaining two methods (resetting and distractors) are much better suited to our available lab space. These two methods are considered overt discrete methods of redirection [131] and, while similar in effect, require very different interventions to complete. The intervention in resetting is designed to be short and unnoticed without concern for breaking presence. Resetting also requires a subject to make a full virtual rotation. The distractor intervention, while long, does not cause a major deviation to virtual heading during the intervention. Each of these methods, then, will likely be detrimental in some way to the aspects of spatial cognition. The motivation behind this work was to determine which of those will cause a stronger hindrance to navigation.

We compared these methods on two specific spatial learning and way-finding tasks. The first was how well people could maintain and update spatial relations as they moved along a path in a VE. We assessed this measure using classic measures of path integration [70] and people's ability to judge relative directions to remembered objects [72]. The second task was modeled on a virtual task that has been previously used to study way-finding and navigation in the context of developing cognitive models of spatial learning and way-finding, the virtual taxicab of Kahana and colleagues [71, 87]. By using these tasks, we attempted to get a comparative measure of how methods of locomoting through a large virtual environment might interfere with navigation and way-finding. While locomotion methods have been compared before [42, 93], to our knowledge they have not been assessed for their ability to preserve or interfere with people's ability to naturally navigate and way-find while walking. We also measured presence in the virtual environment under these three techniques.

The goal of this study is to assess methods for navigating in large scale virtual environments and determine which locomotion method is appropriate for use in future chapters. This chapter is organized as follows. Section 3.2 discusses our implementation of the three modes of locomotion. Section 3.3 presents the first experiment and Section 3.4 presents the second. We conclude with a general discussion in Section 3.5.

3.2 Research Design

In this section we discuss the design of our experiments from the perspective of what we manipulate (independent variables) and what we measure (dependent variables).

3.2.1 Methods of Exploration

In both of our experiments, we manipulate the method by which subjects walk through a large virtual environment while constrained within a smaller physical tracking space in the real world. As mentioned in Section 1, we implemented three different methods, reorientation using rotational gain (*resetting*) [149], reorientation using distractors (*distractors*) [93], and redirected walking while steering to a circle (*steer to circle*) [42, 98]. The methods of resetting and distractors are

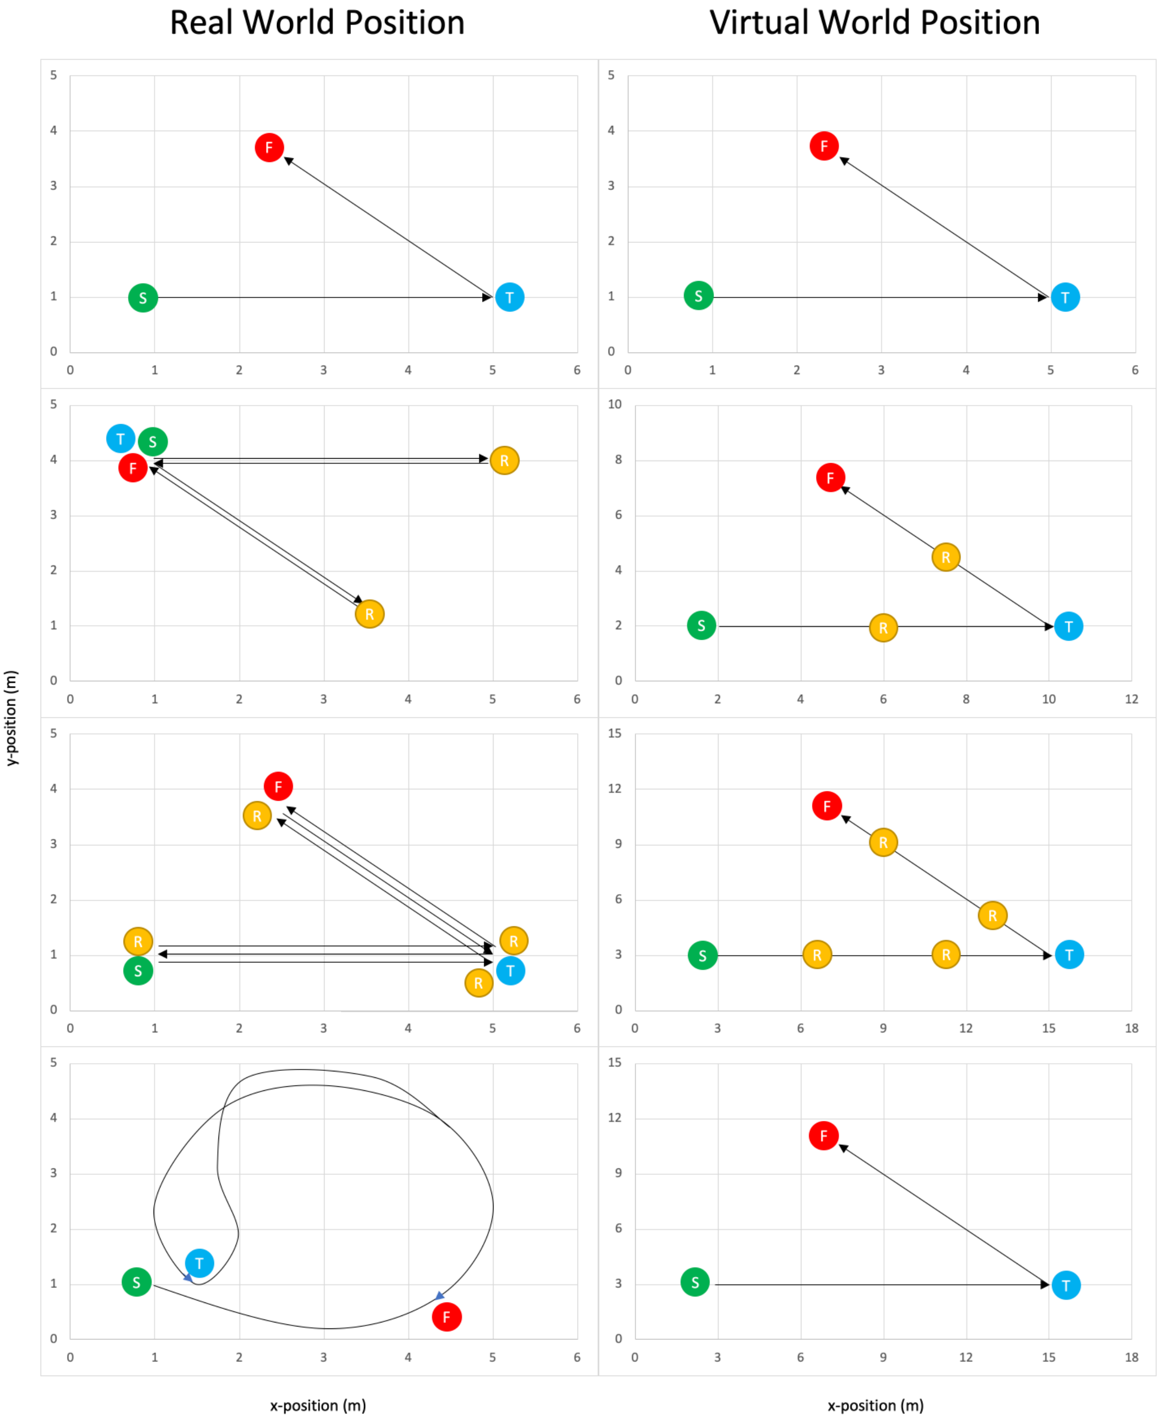

Figure 3.1: Real (left) and virtual (right) paths walked during Experiment 1. The paths from top to bottom are reorientation short (7 meters, 0 ROTs), reorientation medium (14 meters, 2 ROTs), reorientation long (21 meters, 4 ROTs), and redirected walking long (21 meters). S represents the start of each path, F the finish, T the turn, and each reorientation is marked by R. Subjects start at S, walk along the marked path to T and perform a 135° turn. Finally, subjects finish the path at point F.

similar, but there is a key difference in that the distractor method typically requires more time to complete the intervention. Further, the goal of the method is to distract subjects from seeing the rotation injected into the environment. Figures 3.1C and D illustrate the same real world path (top) and virtual path (bottom) walked by a subject using the methods. Figure 3.1C shows the resetting and distractor methods and Figure 3.1D shows the redirected walking. In the virtual world, subjects walk 12 meters in the positive x direction, turn to their left 135° and then walk 9 meters. This is the longest of the three virtual paths (a total of 21 meters). The real world path is the path taken in the physical space in order to walk the virtual path.

When exploring with resetting, the subjects' virtual motion matches their real motion until they are near the border and need to be reset. The reset is done by modifying their rotational gain and instructing them to turn around. By turning around in the virtual world subjects are under the impression that they face the same, or nearly the same, direction as before the reset. However, the altered gain causes the subjects to face backwards. It works because people are not sensitive to slight mismatches in the rotational gain in the presence of visual cues [62]. This modification allows for further travel in the intended direction. Figures 3.1A, B, and C show the typical real and virtual paths taken during reorientation using this method. The difference between the virtual paths of Figures 3.1A, B, and C is their length.

The second method explored is distractors [93], itself a modification of [98] which disguises the intervention using a distractor. Just as with the first method, the physical and virtual motions match exactly until a intervention is required. An intervention of this type requires subjects to follow a distractor object with their eyes, moving their body as necessary, while the world rotates around them. In these experiments the distractor object was matched as closely as possible to the work of Peck et al. [93]. We use a feather model placed 0.5 meters in front of subjects virtual camera at their own height to indicate when an intervention was required. This feather followed a circular arc with sinusoidal motion moving at 8° per second. To cause this mismatched rotation a rotational gain of 0.67 is applied when subjects move their heads against the intended rotation and 1.5 when subjects move their heads in the other direction. This mismatch induces rotation in the real world, and after enough rotation the intervention is complete and the subject is again facing exactly 180° away from their previous real path as in the first method. Later Peck papers [92] used a variation of this method which reoriented subjects toward the center or some ideal location [94]. Figures 3.1A, B, and C show the real and virtual position of the subject during this type of intervention; note that it is the same as the previous method, but a distractor aids the subject in reorienting at points of intervention.

For both these methods, the overall implementation varied slightly between Experiments 1 and 2. In Experiment 1, as discussed below, subjects walked in a pre-specified path where the points that intervention were going to occur at were known. In this experiment, we implemented distractors as it was implemented in the original papers [93, 149], so that the overall rotation in the physical world amounted to a turn of 180°. However, in Experiment 2, subject could freely explore, and piloting revealed that this allowed subjects to become "cornered" in an area of the tracking space where they keep bumping up against boundaries. When this phenomenon occurs, an unusually high number of interventions can occur for very little gain in traversed space. To remedy this problem, we implemented a variant of the solution suggested by Xie et al. [153]. In this variant, when an intervention is required the rotational gain is manipulated until the injected rotation has caused the subject to be facing back toward the center of the tracked space, while still facing in the intended virtual walking direction. Thus, the overall rotational gain may vary from 1.8 to 2.2 during an intervention, well within the bounds of what people are sensitive to [62]. To prevent extraneous intervention a region was created under which the system both abandons the intervention and which subjects must reenter for additional interventions to occur. This region extends inward from the intervention boundary (the boundary of the tracked space) by about 0.3m. Subjects may abandon an intervention by simply stepping backwards. To force subjects back into the tracked space when an intervention occurred they were given instructions to walk two steps forward (this was necessarily towards the center of the tracked space); stepping forward was typically what subjects intended to do, and this method worked seamlessly.

The third method explored here is redirected walking with steering to a circle, and is fundamentally different from the first two. It involves a continuous induced rotation rather than discrete breaks. The full method as well as the governing mathematics can be found in Hodgson and Bachmann [42]. The intention of this method is that it reduces the need for intervention by continually steering the subject to an orbit of given size. This orbit keeps the subject away from walls by effectively pulling them towards the center of the room. Figure 3.1D shows the paths taken in the real and virtual world of a typical path. Since our tracking space is 5m x 4m, the circle we implemented for this method had a radius of 2m. This radius is significantly smaller than the one recommended by Steinicke et al. [125] for redirected walking methods, and we note that Hodgson et al. [41] had difficulty with subjects experiencing simulator sickness with a circle of radius 7.5m.

3.2.2 Path

In the first experiment, the paths that subjects walked were specified, and thus they walked along a route of a given length with given turns in it. If interventions occurred, then the locations of those interventions were known *a priori*. In the second experiment, however, the path was not specified, and subjects explored freely. Unlike Experiment 1, the number of interventions that occurred was unspecified by the experimenters.

3.2.3 Performance Measures

Our primary goal in this paper was to understand people's performance in wayfinding and navigation tasks under the locomotion methods described in Section 3.2.1. Nonetheless, we were interested in standard measures of presence as well [120, 121]. We measured presence using a modified SUS presence questionnaire.

One way to determine the spatial awareness of a subject is blind pointing, where a subject must turn and face a remembered object without visual feedback [107]. Experiment 1 recreated this blind pointing task to measure spatial awareness in each locomotion method. Higher accuracy is indicative of greater spatial awareness. Additionally, we are interested in how easily and quickly subjects could acquire route knowledge of an environment [70], possibly employing landmarks and other features of the environment. So in a second experiment, we employed a navigational task in which subjects are asked to learn an environment and locate targets not within line-of-sight [87]. We measured how accurately and quickly subjects could do this under the various modes of locomotion.

3.3 Experiment 1

3.3.1 Methods

3.3.1.1 Participants

The participants in the study consisted of college age students from Vanderbilt University between the ages of 18 and 25. Each subject volunteered and was compensated $10 for an hour of their time. The subjects were informed of the purpose of the experiment but were not told how the system allowed them to explore the environment until the conclusion of the experiment. Fourteen subjects participated and two subjects dropped out during the steer to circle portion of the experiment, complaining of simulator sickness. Their data are not included in the analysis below.

3.3.1.2 Equipment

The experiment was performed using the NVIS nVisor SX-60 head mounted display (HMD). The SX-60 provides stereo and full color. The resolution in each eye is 1280x1024 and provides 60° field of view diagonally. The physical space used is roughly 6x5 meters; the available tracked space is 5x4 meters. The tracking system used is the WorldViz PPT Precision Motion Tracking System, which used four cameras to track the position of two LEDs attached to the HMD; this tracked both position and orientation with the aid an Intersense tracking module.

Subjects were required to indicate when they intended to begin the walking task and when they were facing the intended object. To accomplish this subjects were given a Wii-Mote motion controller, which interfaced with the environment.

3.3.2 Environment

Figure 3.2 shows a top-down view of the environment being used; the environment contained seven objects that were placed throughout the space, each 20 meters from the starting position of the subject. Additionally the world contained a model of a city, which was used as background to the experiment and some of which can be seen in Figure 3.3. Figure 3.2 labels each of the items that the subjects were asked to remember. They are (clockwise from the top) a well, a parking meter, a fire hydrant, a trashcan, a basketball goal, a mailbox, and a streetlight.

3.3.3 Spatial Awareness Task

Experiment 1 uses a spatial orientation task where subjects walk a path then turn to face a remembered direction, similar to the pointing task of Rieser and others [107, 57, 53, 149]. To test a subject's ability to maintain spatial awareness of the environment the following task was performed six times under each walking condition (resetting, distractors, and steer to circle). Walking condition was blocked with the order of blocks randomized. The ordering of the three path lengths was randomized for each subject, but consistent across blocks. Subjects were instructed to remember three target objects selected randomly from among the seven listed in Section 3.3.2, and given as much time to memorize the location of the objects as needed. The subject then walks a path through the environment consisting of two straight path segments with a 135° turn in between the two segments. To mark the subject's target location a column (see Figure 3.3) was placed in the environment. Upon reaching this first way-point, the subject was instructed to find the second way-point (again marked by a column) and proceed to this new location. Reaching this second column marked the conclusion of the path at which time the display went black in preparation for the task. The subject was then instructed to face toward a given object from among the three originally given and memorized. Subjects indicated they had faced the object by pressing a button on the WiiMote controller. For each trial we recorded turning error, the difference in heading angle and angle to the correct object, and the trial concluded.

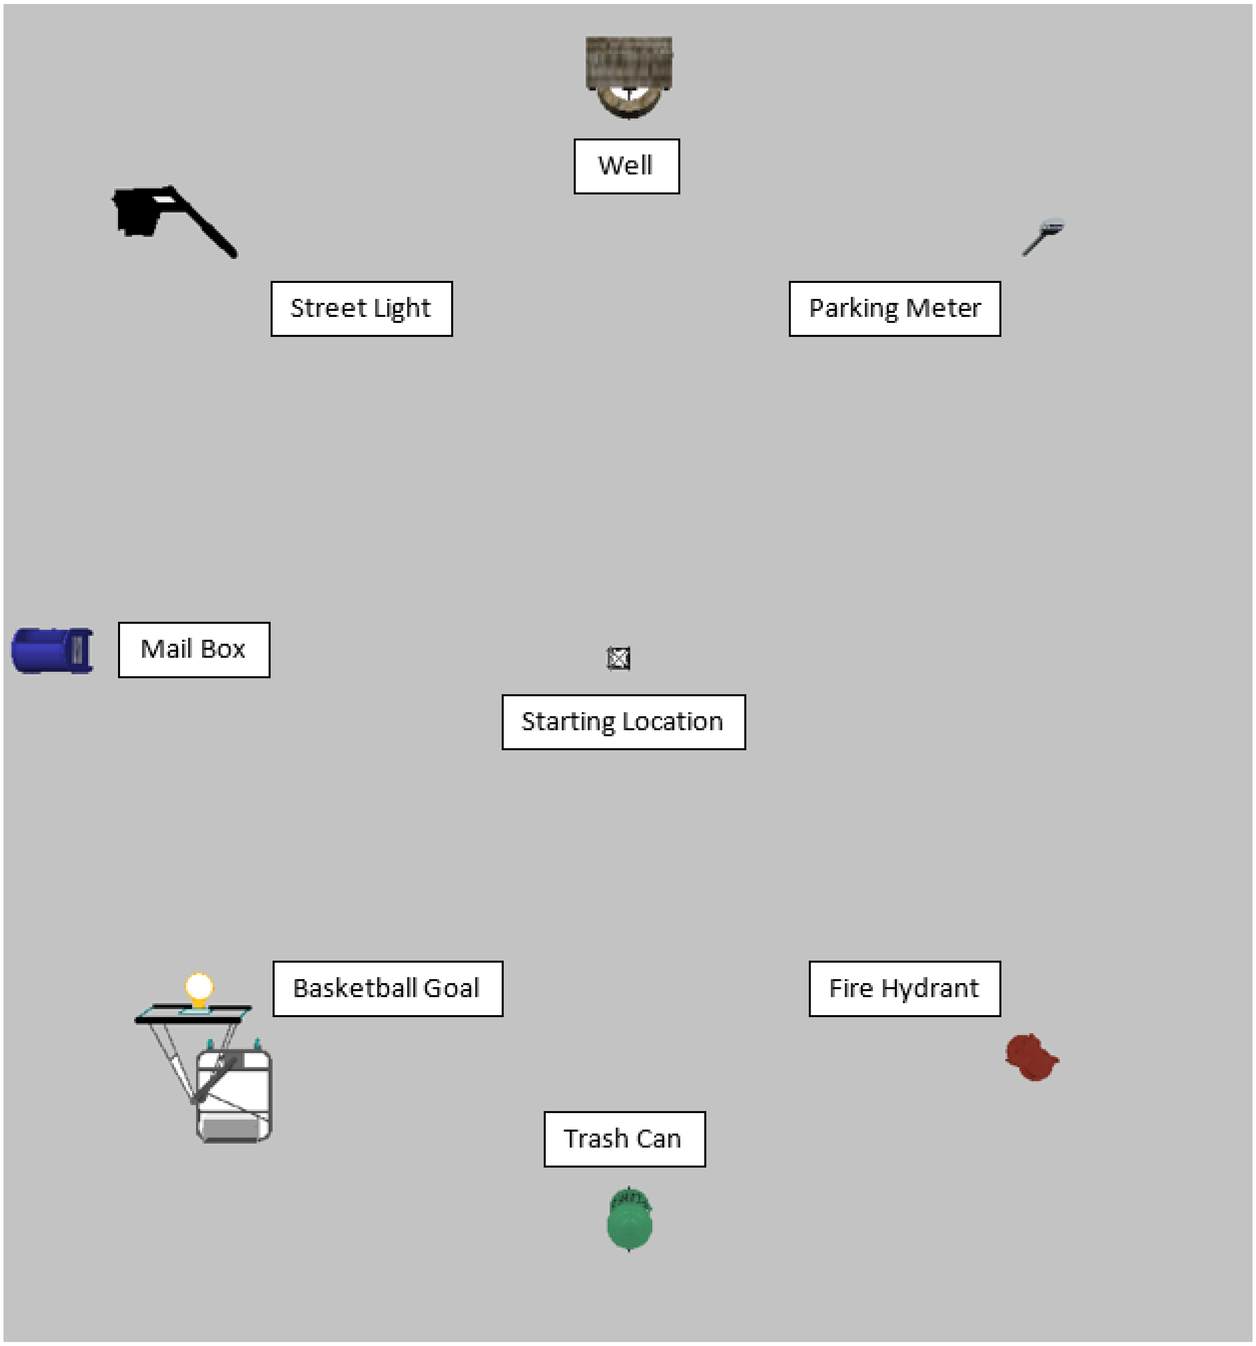

Figure 3.2: Top-down view of the environment used in Experiment 1 showing all objects, a randomized subset of which were memorized in each trial.

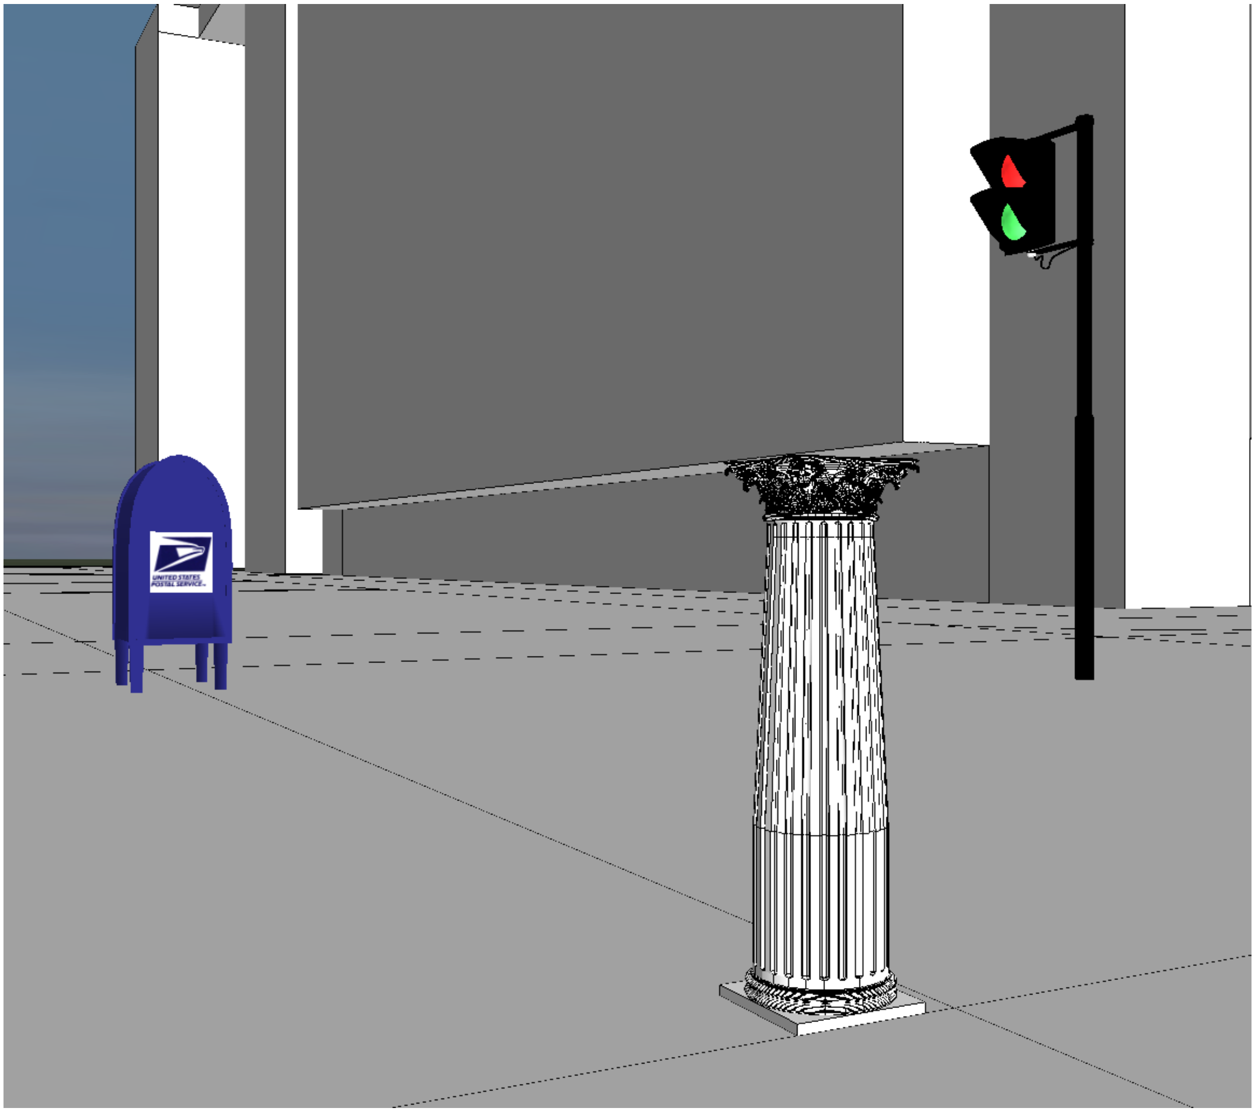

Figure 3.3: An example of a starting view of a subject. The column marks the target walking location and the mailbox and streetlight were potential markers to be memorized.

In a given trial during the resetting and distractor locomotion methods, a subject experienced either 0, 2, or 4 interventions; the number corresponded to the length of the segments walked in the path — a higher number of interventions indicating a longer path. The two path segments were either 4 and 3 meters (0 interventions), 8 and 6 meters (2 interventions), or 12 and 9 meters (4 interventions). intervention occurred every 4 meters during the initial segment and every 3 meters during the final segment. Figure 3.1 shows the 3 real and virtual paths walked during the experiment. During the redirected walking condition the same segment length was used, however there was no need for an intervention. Starting position was randomized to a corner of the physical room for each trial and subject, however, in the virtual world the starting position was fixed. As seen in Figure 3.1 each path had a turn of 135° which could be right or left as the room geometry permitted.

The design of the experiment was within subjects. The trials of each locomotion method were blocked together and the order of the three conditions is counterbalanced to eliminate any ordering effects among conditions. The subjects repeated each path length twice, yielding six trials within each condition. For each trial the amount and direction of the subjects turn as well as the error associated with the intended object was recorded.

To test the experience of the user within the environment each subject completed a modified SUS [93, 120, 121] presence questionnaire after the blocks of trials for each condition. The questionnaire consisted of six questions each of which measured the presence experienced by the subject on a Likert scale from 1 to 7.

3.3.4 Results

3.3.4.1 Spatial Awareness

There were two independent variables in this experiment: the locomotion method and distance travelled, 7m, 14m, or 21m. The 7m path could be performed within the tracked space of the laboratory without intervention in the case of the first two locomotion methods. The 14m path required two interventions and the 21m path required four. Each path length was walked twice.

Thus, each locomotion method consisted of six trials, which the subject performed consecutively. It should be noted that the 7m path length trial had zero resets making it identical for the resetting and distractor conditions.

Figure 3.4 shows the mean absolute turning error with respect to condition and distance as well as the standard errors of the mean. A repeated measures analysis of variance (ANOVA) with locomotion method and distance as independent variables found no main effects or interactions. We were interested in whether the experimental findings indicated that the locomotion methods were equivalent in the present experiment (i.e., the null hypothesis was true). However, a well-known limitation of traditional statistical tests is that they cannot provide evidence for the null hypothesis [108]. Increasing the sample size does not solve this problem. To assess the evidence in favor of the null-hypothesis, one can use Bayes factors [50, 51]. The use of Bayes factors provides a principled method for calculating the evidence in favor of the null hypothesis and expressing that evidence as an odds ratio.[2] Bayes factors depend on sample size and therefore intrinsically adjust for power. Our analyses used the methods developed by Rouder et al.[108]. We set the prior odds to 1, which neither favors the null nor the alternative. Comparing locomotion methods then, the Jeffrey-Zellner-Siow (JZS) Bayes factors favoring the null hypothesis were as follows: 7.72 for distractor vs. resetting (i.e., the null hypothesis is 7.72 times more likely than the alternative; this is conventionally considered strong evidence); 1.34 for steer to circle vs. distractor; and 1.02 for steer to circle vs. resetting. The latter two comparisons were ambiguous, providing roughly equivalent evidence for the null and the alternative.

3.3.4.2 Presence

Following each method, the subject was asked to fill out a modified SUS questionnaire as described previously. The responses were transformed into a binary value [93, 120, 121] in which responses of 5, 6, and 7 were considered high presence; all other values were considered low presence. The results of this transformation are summarized in Table 3.1. A pairwise logistic re-

---

[2] Online calculators for these statistics are available at http://pcl.missouri.edu/bayesfactor.

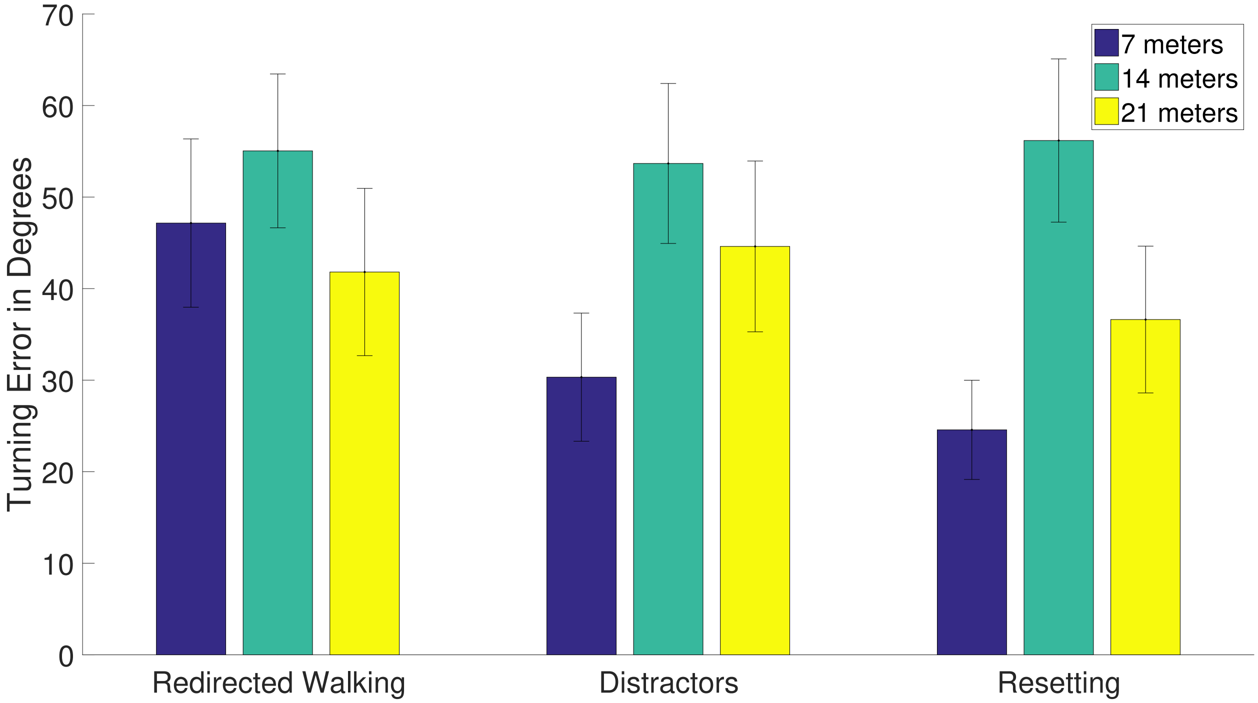

Figure 3.4: Absolute turning error (and the standard error of the mean) associated with each locomotion method by distance walked.

gression analysis between locomotion methods was performed. We found a statistically significant difference in presence between the steer to circle method and the distractor method ($\chi^2(1) = 7.22$, $p < 0.01$), with the distractor method having significantly higher presence, and an approaching significant difference between the steer to circle method and the resetting method ($\chi^2(1) = 2.87$, $p = 0.09$), with resetting having marginally higher presence.

	High Responses	Mean High
Steer to Circle	25	0.3472
Distractor	41	0.5694
Resetting	35	0.4861

Table 3.1: Mean high responses for presence questionnaire of Experiment 1.

3.3.5 Discussion

Our test was unable to distinguish any difference between the three locomotion methods in terms of spatial awareness, but we can assert strong odds for equivalence between the distractor and resetting methods. All subjects performed at much better than chance (i.e., below 90°) in terms of performance, indicating that spatial awareness is maintained using all locomotion methods. We found differences in presence measures between steer to circle methods and the other two, and experienced a moderate dropout rate (2 out of 14) with the steer to circle method. While the dropout rate was foreshadowed by experiences of prior work [41, 125], we wanted to explore the use of a redirected walking method in a smaller tracking space, a size more likely to be supported in commodity level setups.

3.4 Experiment 2

In Experiment 2, we tried a different task, that involved way-finding and navigation. In this experiment subjects explored the environment freely and walked significantly further than those in Experiment 1. Upon piloting, pilot subjects complained about the steer to circle method. Given our prior experience with dropouts, we chose to compare only the locomotion methods of resetting and distractors.

3.4.1 Methods

3.4.1.1 Participants

This study consisted of 16 college age students at Vanderbilt University between the ages of 18 and 26 (10 male/6 female). Each subject was compensated $12 for participation of no more than 75 minutes. As in Experiment 1, the subjects were not told how the system would manipulate their rotation to allow them to explore the environment.

3.4.1.2 Equipment

All equipment was identical to that used in Experiment 1 (e.g., HMD, tracking system, Wii-Mote). In this experiment, subjects learned the geographic layout of a virtual environment. We employed a between groups design, with 8 subjects per group (5 male/3 female).

3.4.2 Environment

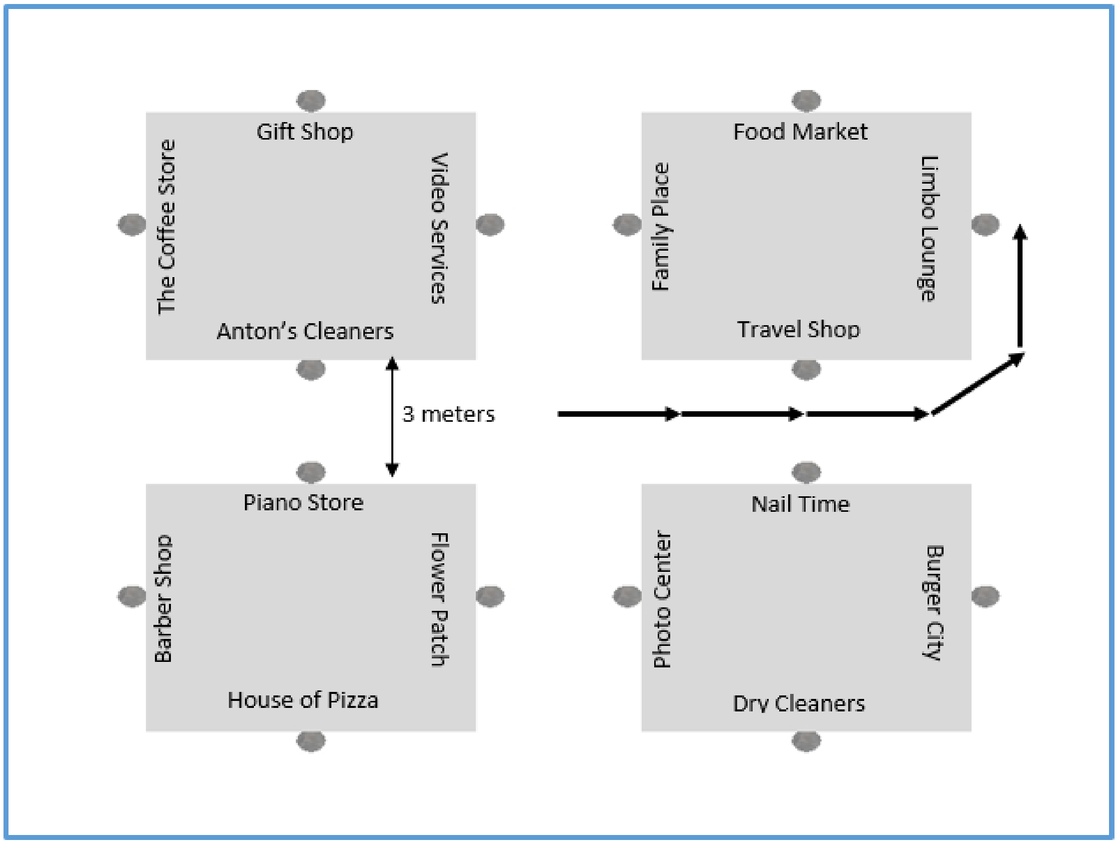

Figure 3.5: Top-down view of the environment used in Experiment 2, illustrating both distance and location of the storefronts. The blue line represents a wall boundary around the environment. A sample path starting from the center and the reorientations necessary to walk that path are shown as well. Each arrow represents one length before a reorientation is required.

Figure 3.5 shows a top-down view of the environment used in this experiment. There are 16 storefronts located on four buildings (each 6 by 6 meters) with each building spaced 3 meters apart. Surrounding the environment are four identical textured walls. Subjects started in the center of the four buildings and always started facing north (towards the top of Figure 3.5). From the starting position subjects could see (without translating) eight storefronts. Figure 3.6 shows one example of a storefront, in this case the "Family Place," which the subject might be asked to locate. There is a column at the center of this storefront; as with all other storefronts. These columns mark the exact target location for the subjects.

Figure 3.6: This figure shows an example storefront and the target column placed in front of the store. Each storefront has an equivalent column.

3.4.3 Navigation and Wayfinding Task

This task was modeled on the yellow taxicab task described in Newman et al. [87]. Subjects were given instructions on the task beforehand; they were also informed of the basic layout of the environment.

For this task each subject was assigned one random storefront from each of the four buildings shown in Figure 3.5 (e.g., House of Pizza, Family Place, Gift Shop, and Photo Center). Each of four trials consisted of a random ordering of those four assigned storefronts. For the first trial subjects started at the center of the environment and were asked to visit each of the four storefronts in randomized order indicated by an on screen display (these four paths made up one trial). Each subsequent trial started at the ending point of the previous trial, and subjects then visited the four storefronts once more in a new randomly shuffled order. Thus subjects necessarily visited each of four selected storefronts four times. Additionally, each trial necessarily consisted of one visit to each of the four buildings. Subjects began the next trial by pressing a button on the Wii-Mote. The number of interventions that occurred was completely dependent on the path the subject walked. For this task an intervention was needed whenever a subject approached the boundaries of the tracking space as described in Section 3.2.1.

Subjects were randomly assigned to one of the reorientation locomotion methods, either the distractor or resetting method. For each trial the time taken, time spent reorienting, distance traveled, and number of interventions were recorded. From this data, we calculated time spent walking, average speed, and ratio of walked distance to ideal distance to the target. Here ideal distance in the minimum possible walking distance. This was calculated assuming subjects would walk directly toward the target object when in sight, and toward the optimal corner of the building when a turn was required. In many of these situations the ideal path hugged one or more of the buildings that would be passed in the ideal path.

3.4.4 Results

Figure 3.7 shows the total time spent in the environment (including interventions). The results of this analysis are unsurprising: a 2 (locomotion method) x 4 (trial) repeated measures ANOVA shows a significant effect of trial, $F(3,42) = 4.68$, $p = 0.03$, and a main effect of method, $F(1,14) = 218$, $p < 0.001$. This latter result is because the distractor method takes more time. However, if we consider only the time spent searching, that is, subtract out the difference in intervention times for each method, a somewhat different picture emerges. Figure 3.8 shows the average time spent searching across subjects by condition in each of the four trials. That is, this figure shows the total time a subject spent in the environment minus the time spent reorienting per trial averaged across subjects. A 2 (locomotion method) x 4 (trial) repeated measures ANOVA on time spent searching shows a significant effect of trial, $F(3,42) = 9.84, p < .001$, $\eta_p^2 = 0.64$. Planned contrasts show that subjects get faster progressing from trial 1 to trial 2 but plateau at trials 3 and 4.

Figure 3.9 shows the average distance traveled across subject by condition by trial, Figure 3.10 shows the same for ratio of distance walked to ideal distance, that is, the shortest possible path between targets. Table 3.2 shows this data on a per subject basis, showing the total distance walked, the ideal distance, the ratio between the two, and the number of interventions each subject experienced over the course of the experiment. A 2 (locomotion mode) x 4 (trial) repeated measures

ANOVA for both distance and ideal distance show main effects of trial, $F(3,42) = 8.25$, $p < 0.001$, and $F(3,42) = 12.12$, $p < 0.001$, respectively. These results are similar to the results for time. In all of the preceding ANOVAs, tests for normality and homogeneity of variances were met. Note that the number of interventions for the distractor method is, on average, higher than the number for resetting, 112 interventions (8.6) vs. 101 interventions (9.7), but this difference is not significant.

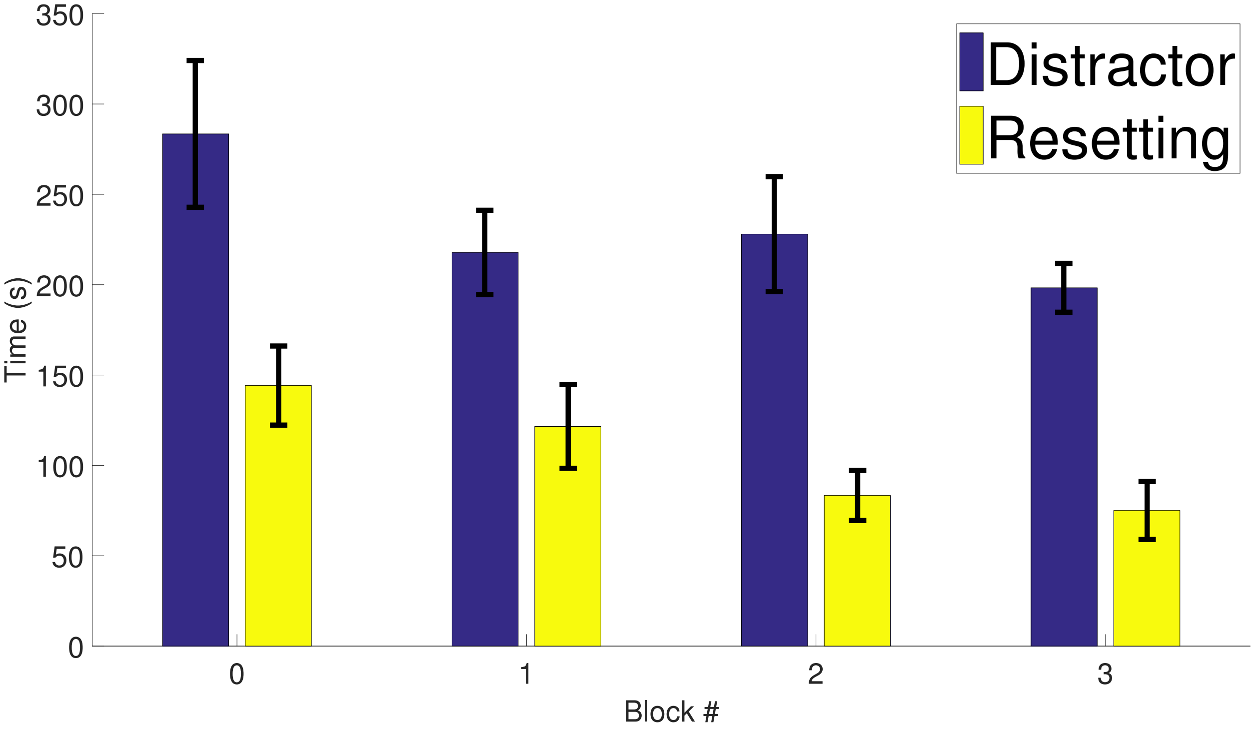

Figure 3.7: Total time spent (and the standared error) in the system by condition for each trial.

As in Experiment 1, we wished to confirm that there was no systematic difference between the locomotion methods of distractors and resetting with respect to our primary measures of interest: time searching (with intervention time removed), distance traveled, and ratio of distance to ideal distance traveled. Thus we again employed a Bayes factor analysis [108]. We again set the prior odds to 1, which neither favors the null nor the alternative hypothesis. In terms of time spent searching, the JZS Bayes factor favoring the null was 2.35, that is, the null is 2.35 time more likely than the alternative. For number of interventions encountered, the JZS Bayes factor was 2.57. Both of these odds ratios are ambiguous, and do not give us any clear indication of whether the hypotheses are equivalent or not. However, for distance traveled, the JZS Bayes factor was 3.17,

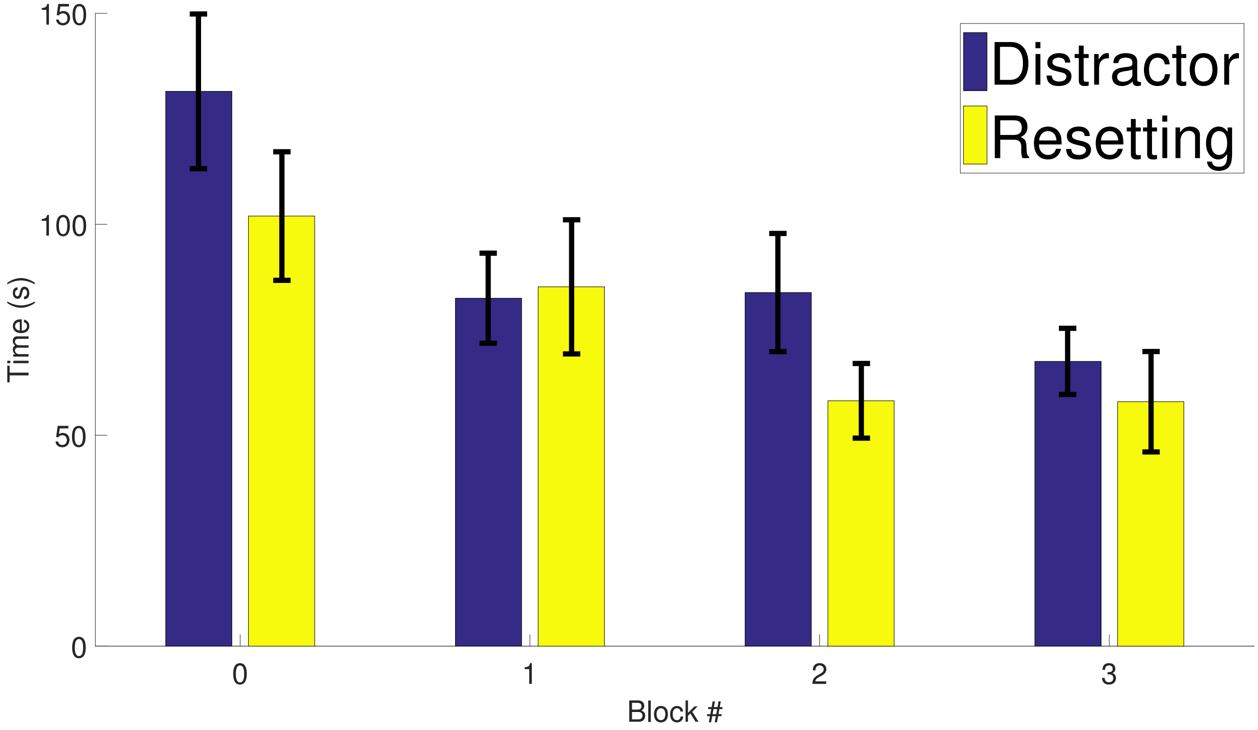

Figure 3.8: Time spent searching, i.e., total time minus time spent during interventions, (and the standared error) by condition for each trial.

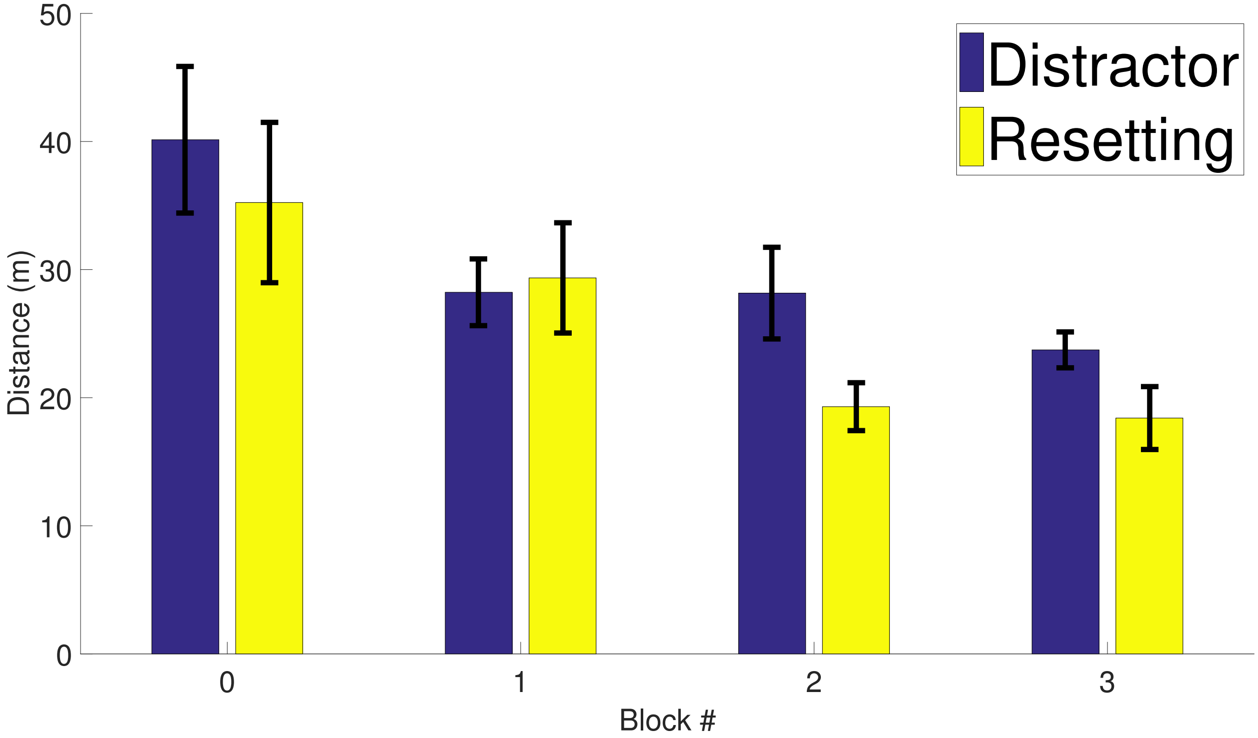

Figure 3.9: Average distance traveled (and the standared error) by condition for each trial.

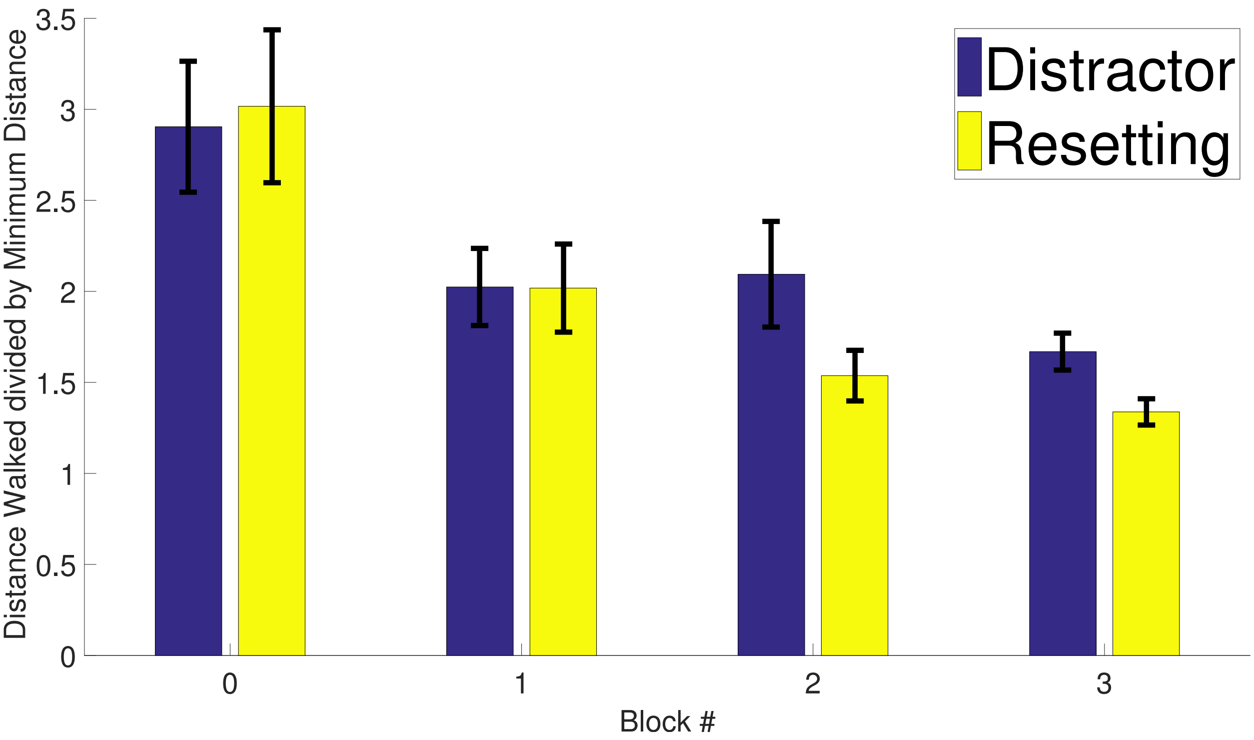

Figure 3.10: Average distance traveled by condition for each trial divided by the ideal distance for that trial (and the standared error of that measure).

Subject #	Condition	Total Distance Walked (m)	Ideal Distance (m)	Ratio	ROTs
0	Distractor	368.7	236.9	1.56	88
1	Resetting	356.7	210.6	1.69	88
2	Distractor	565.7	252.7	2.24	137
3	Resetting	408.2	166.8	2.45	100
4	Distractor	464.9	201.3	2.31	115
5	Resetting	202.3	157.9	1.28	48
6	Distractor	460.8	193.6	2.38	99
7	Resetting	403.9	210.2	1.92	96
8	Distractor	551.5	225.2	2.45	130
9	Resetting	533.3	253.1	2.11	132
10	Distractor	648.9	255.8	2.54	150
11	Resetting	523.5	259.5	2.02	133
12	Distractor	360.3	214.7	1.68	86
13	Resetting	496.6	205.3	2.42	116
14	Distractor	427.6	195.4	2.19	94
15	Resetting	349.0	204.4	1.71	91

Table 3.2: Shows a summary per subject of the distance walked by the subject compared to the ideal distance. interventions are the total number of times the subject had to reorient.

and for the ratio of distance to ideal distance, the JZS Bayes factor was 6.74. These odds ratios are higher and generally considered [108] to give strong evidence for an equivalence between resetting and distractors on these measures.

3.4.5 Discussion

Subjects were able to do the task in Experiment 2, improved at it, and learned the geographic layout of the environment despite the large number of interventions involved. This task involved learning route knowledge and we have shown that there is good evidence to support that one method of intervention does not interfere with the acquisition of route knowledge more than another. Note that the intervention with distractors takes a longer amount of time than the intervention with rotational gain (resetting) method by design: the distractor takes time to accomplish the intervention. When this is accounted for in the search statistics (by subtracting intervention time out), there is no difference in how long it takes to navigate an environment.

In this experiment, the ideal distance of trials was variable. We designed the experiment this way to ensure that the trials covered many of the possible routes from any two storefronts. This design means that not every subject had the same ideal distance, which caused some subjects to have much shorter distances and hence large variation in interventions performed by each subject. For example subject 5 had only 48 interventions in large part because that subject's total distance walked was much shorter. This behavior does not therefore represent outlier behavior on the part of the subject, but the design of the experiment and natural variance in navigation. This variation raises the issue of whether our experiment had appropriate power to detect differences between the two methods of intervention [20]. First, we note that we did find significant differences between conditions in our experiment, for example, total time is significantly less using the resetting method than the distractor method. Next, we note that a Bayesian analysis, as we have performed, provides relative support for the equivalence of these methods, and it incorporates the the sample size in its calculation of this support. This type of analysis strengthens the rationale for using some correction (which we have in figure 3.10) for ideal distance for a given subject. Of particular note here, the ideal distance ratio metric yields the strongest Bayes factor in favor of the null hypothesis among those measured, that ignoring the reorienting time, the methods are not systematically different.

3.5 General Discussion

This paper examined three methods of traversing large virtual environments while walking when the physical tracking space is limited. Our focus in comparing these methods was on how well they would perform in tasks of spatial awareness and navigation. Our first comparison was on a classical method of spatial awareness involving path integration, and our second comparison involved a task of learning a route in a town via active navigation and free exploration. All methods proved reasonable at performing the tasks for which we tested them, and we were unable to distinguish between them on these measures. Nonetheless, we can offer some guidelines for the choice of a reorientation system based on these tasks.

First, our tracking space is significantly smaller than one that is typically used for a method of redirected walking such as the steer to circle. As has been noted previously in the literature, with a tracking space so small, the perceptual effects of redirected walking will be quite noticeable [125], and can lead to simulator sickness [41], which we experienced in a non-negligible portion of our subject pool. That and pilot testing caused us to not employ the method in the second experiment, where we expected subjects to walk hundreds of meters in a virtual environment. So, while the steer to circle method works, and seems to work as well as the reorientation methods in terms of spatial awareness, we would recommend that it be used in a larger tracking space than we have employed.

Our results provide strong evidence that resetting and distractors are equivalent on both spatial awareness and learning a geographic layout in this paper. The distractor method had higher ratings of presence than did resetting in Experiment 1, and prior literature has tested this more thoroughly than done here [93]. If presence is a quantity that it is desirable to maximize, then the distractor method may be the method of choice. However, the drawback of the method with distractors over the resetting method is that it takes significantly more total time to traverse and navigate through an environment. This may be a factor in designing a virtual experience that needs to be considered, and presence may need to be traded off against this.

There is potential for future exploration of wayfinding and navigation with locomotion modes in virtual environments. There are other attributes of navigational and wayfinding knowledge that people acquire in the real world while traversing routes as was done in Experiment 2 that we did not test; for example, we did not evaluate how well subjects acquired a knowledge of total layout of the environment, whether they could find routes to novel locations, and so on [48, 17, 18]. Also, there are other modalities than walking that it it may be profitable to explore, either independently or in combination with walking, such as teleportation, or some of the alternative methods mentioned in Chapter 2. With the advent of commodity level VR devices (e.g., the HTC Vive), that permit walking in a slightly smaller tracked area than the one used here, many game designers are exploring novel methods of traversing large virtual environments. How these combine with walking in the tracked space to allow navigation and way-finding will be an important future concern.

Chapter 4

Specific Aim 2

4.1 Introduction

Virtual reality (VR) provides engaging experiences and allows for training and simulation in a controllable environment. Virtual worlds, as with the real world, can differ vastly in size and scale. Large virtual environments are necessary to provide analogs to large real-world environments. Exploring these large environments in smartphone-based VR systems (e.g., the Samsung Gear VR) introduces a difficult challenge. There is no native position tracking for these devices, which is disappointing given the freedom they provide. The rendering and display systems of smartphone-based VR are entirely on-board, making the system tetherless. With the broad adoption of smartphones and the ease of acquiring a cheap housing unit (e.g., the Google cardboard [33]), smartphones provide an enticing platform for VR. In this paper we explore three naturalistic methods of navigating in large environments using smartphone-based VR systems: resetting [149] and two types of walking in place (WiP) [135, 150, 122]. Resetting utilizes real walking and allows for unbounded space by reorienting subjects towards the center of the room when a boundary is encountered, all the while maintaining virtual heading. WiP induces translation based on a subject's stationary motion. The study here is, to our knowledge, the first evaluation of spatial cognition using mobile VR and the first to utilize a real walking metaphor in mobile VR.

While many methods have been developed for navigation in large virtual environments [99, 47, 149, 27], the best method of doing so depends on many factors. These factors include room size and layout [4], technology, the virtual environment [64], performance metrics, etc. For example, if the room size is sufficiently large, then a redirected walking technique might be employed [125, 127]; if the room size is small, however, some of these techniques can induce simulator sickness or require pre-computed trajectories, e.g., the methods of Langbehn et al. [64]. If the environment is very large, then some method of locomotion beyond normal bipedal locomotion may be appropriate, e.g., the methods of Williams et al. [149] or Interrante et al. [47]. In this paper, we will consider some of these factors and make assumptions about others. In particular, for resetting, we assume a reasonably sized open room (roughly 4m x 4m) is available for the real world enclosing space, and that the virtual environment is easily navigable by ordinary locomotion. WiP methods, of course, do not have any space requirements beyond standing space.

There are various performance metrics used to judge the success of locomotion methods [147], such as breaks in presence [93], simulator sickness [85, 31, 34], and judgments of relative direction [149], etc. Because we are interested in training and simulation, we believe the acquisition of spatial knowledge to be a valuable metric. Navigation methods have been linked to the acquisition of spatial knowledge [113, 115, 151, 152] with walking methods outperforming other methods, e.g., joystick, teleportation, and flying. In this paper we compare how well three methods of navigation afford the acquisition of spatial knowledge, specifically survey knowledge [48, 15].

The four components of spatial knowledge are landmark, route, graph, and survey knowledge [48, 15]. Of most interest is survey knowledge, the acquisition of which provides subjects with knowledge of the straight-line distances and directions between places defined in a common frame of reference. To test the acquisition of survey knowledge we employ three metrics: initial angular error [17, 107], estimated path length [17], and orientation time, a measure of how quickly subjects recall the relative direction of objects. These measurements are collected to determine how well spatial knowledge is acquired in the three modes of navigation. The VR interfaces provide varying levels of body-based information during both learning and testing. In Experiment 3, we find that there is a difference in the expressed spatial knowledge. Experiment 4 shows that the difference was an effect of how we tested spatial knowledge.

4.2 Research Design

4.2.1 Techniques of Walking

We implemented three techniques for locomotion in this work: the reorientation method called resetting [149], which uses real locomotion with interventions at the boundaries of the tracked space, and two techniques of WiP [90], which require only the space to stand up.

4.2.1.1 Resetting

Our implementation of resetting was originally developed by Williams et al. [149] and Xie et al. [154], involves real walking, and grants full idiothetic cues to self motion. Resetting consists of two phases, the traditional walking phase in which no modification is done to either the orientation or position of the subject, and a reorientaion phase that occurs when a collision with the boundary of the tracked space occurs. During the first phase resetting is functionally equivalent to real walking. When a subject reaches a boundary of the tracked space the system initiates the second phase of resetting. During this phase the actual reset takes place and the rotational gain of the system is modified so that a virtual turn of 360 degrees is equivalent to a real world turn towards the exact center of the room. Thus, subjects believe themselves to have turned completely around and maintained their headings, while they have actually been reset away from the boundary. In the original Williams et al. [149] work, the gain was modified by a factor of 2, but, consistent with later work, our implementation dynamically adjusts the gain to point subjects to center of the room. This manipulation results in fewer resets. Resetting provides full translational and idiothetic cues to self-motion, but the process may cause some degradation in the acquisition of rotational information because of the body rotations inserted into the path.

4.2.1.2 Walking in Place

Our WiP techniques are body-based turning methods, with turning indicated by head rotation. The direction of motion is taken from the direction of the virtual camera. This means that all linear

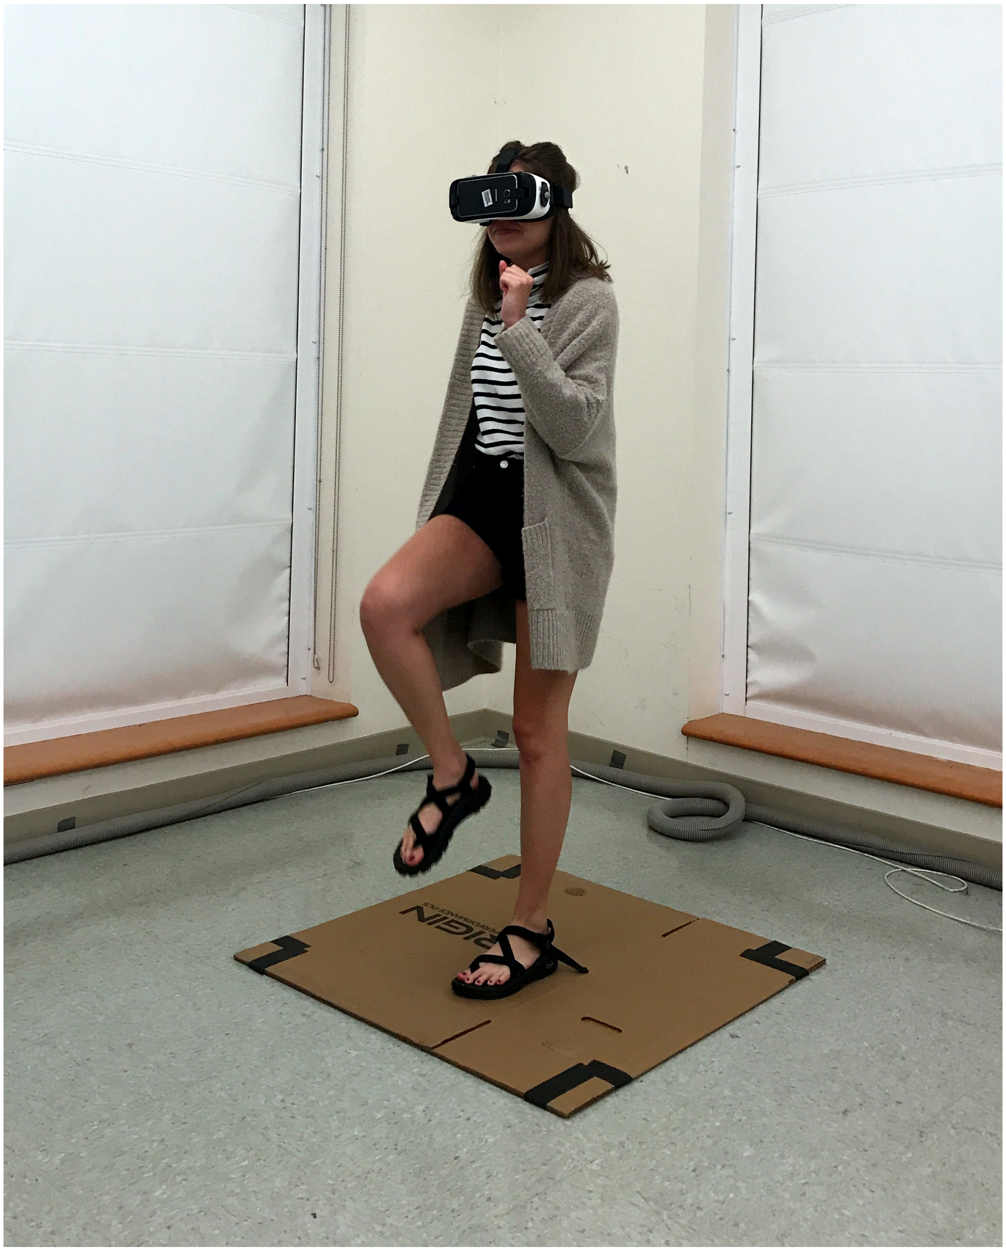

Figure 4.1: A demonstration of WiP. Cardboard is placed on the floor to prevent subjects drifting as they walk in place.

motion is in the direction the subject is looking in the virtual environment, which prevents looking around when walking. The first WiP technique (Gear Only) we use takes inertial measurement unit (IMU) data directly from the smartphone of a Gear VR to determine head motion. This method is similar to how pedometers in smartphones and smartwatches work. Pedometers use pattern analysis techniques to search for repeating motions that are assumed to be steps. The repetition requirement of pedometers leads to either unreliability when the number of repetitions is low or severe lag when the number of repetitions is high. For this reason, we chose not to implement a standard pedometer. Instead, we extract the up and down acceleration of the head to impart motion.

For our method walking is divided into two repeating states. We say the subject is stepping when the magnitude of the upward acceleration is greater than 0.1 m/s/s and not stepping otherwise. We assume an average walking pace of 1 m/s with each step taking 0.5s. The true speed of the system — that is, the rate at which the environment visually flows by — is the ratio of the user's stepping rate to the average stepping rate (2 steps per second) multiplied by the average walking speed (1 m/s). This allows for speed to be controlled by taking faster or slower steps. When a step begins we exponentially decay in an increasing form, (i.e., $1 - e^{-kt}$) to that maximum speed and remain there until we detect the current step has stopped, at which time we decay to half of the maximum speed. This decay upward and downward repeats for as long as the user is considered to be walking. A user is no longer considered walking if enough time (50% of a user's rate of stepping) has passed since the end of the previous step. If this cessation is detected we decay to a speed of zero. The time constant for each of these decays is 0.2 seconds to make the change in speed subtle but noticeable. Note that while optic flow and motion do not stop immediately, they do stop within roughly 0.5 seconds of a user stopping.

The Microsoft Kinect v2.0 [54], which directly tracks users' feet and legs, offers a more direct method of measuring WiP. This second WiP technique utilizes the Kinect v2.0 to track the angle formed (for both the left and right leg) by the hip, the knee, and the ankle joint. We assume the user is currently stepping if the angle formed is less than 145 degrees. We add an additional modification to the algorithm to allow for quicker detection of cessation of steps. We assume that if both leg angles are over 165 degrees for 0.1 seconds then stepping has stopped. The Kinect produced errors with step detection in certain orientations. Tracking was most stable when users faced either towards or directly away from the Kinect. When facing to the side and obscuring the back leg from the camera the system was prone to missing steps and causing unintended slow downs. This occurred primarily during the initial practice phase as subjects learned to adjust rapidly and controlled their motion effectively. In both methods we placed a 1m x 1m piece of cardboard on the ground and instructed subjects to walk in place, only on the cardboard; the cardboard, in effect, prevented drift.

4.2.2 Environment

Figure 4.2: Top-down view of the environment used in the practice phase of each experiment. There are four objects for the subject to find and four landmarks (paintings).

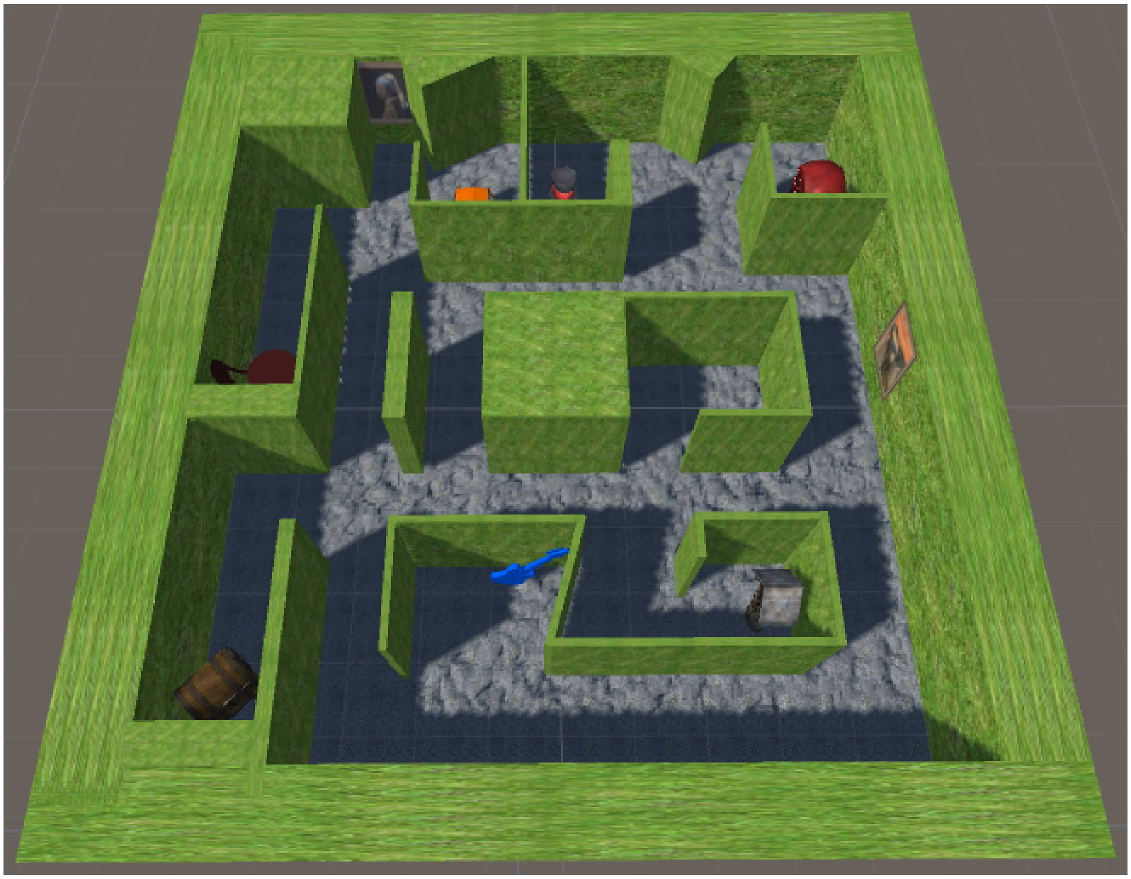

Figure 4.3: Top-down view of the environment used in the learning phase of each experiment. There are eight objects for the subject to find and four landmarks (paintings) to facilitate learning.

In both experiments every subject was presented with three distinct environments. The first was a training maze shown in Figure 4.2. Subjects were placed in a practice maze in order to train in their assigned technique(s) of walking. The second was the learning maze shown in Figure 4.3; subjects were instructed to learn the spatial relations among the eight objects contained in this maze. Four landmarks were present to aid in learning the overall layout. A first person view of the learning maze is shown in Figure 4.4, with one of the eight objects. Due to the geometry of the maze, a subject could not see any two objects at the same time. The final environment is presented during the actual testing phase. Figure 4.5 shows a first person view while the subject is in a sparse environment with a Voronoi textured ground plane to give the subject some ocular flow for feedback on distance traveled.

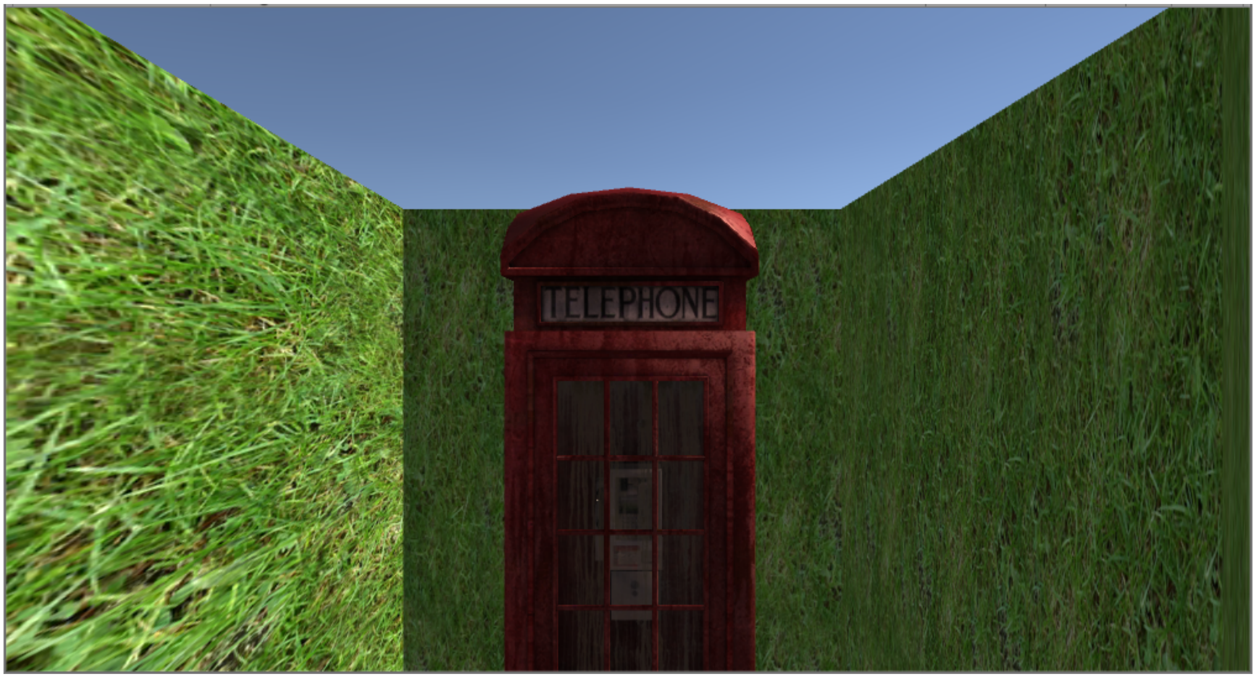

Figure 4.4: This figure shows the telephone booth as would be seen by subjects during the orientation phases of each experiment.

Figure 4.5: At the beginning of the testing phase subjects are informed of the target object via a heads up display. This disappears shortly so as not to distract the subject during walking.

4.2.3 Testing Protocol

To measure how well spatial knowledge had been acquired each of the experiments in this study followed the same basic testing protocol and will be described in complete detail in Sections 4.3.4 and 4.4.5. The testing protocol consists of an orientation portion (Figure 4.4) and a walking portion (Figure 4.5). Subjects were instructed to walk directly in a straight line from the object presented in the orientation portion to an object given via an on screen display in the walking portion [17]. During the walking portion no part of the maze was visible allowing for the subject to walk in a single straight line.

4.2.4 Metrics

Since we are interested in how well our tested methods of walking facilitate the acquisition of spatial knowledge, in particular survey knowledge, we employ two metrics that allow us to measure how well survey knowledge has been acquired. As stated in Chapter 2 there are two important aspects of survey knowledge: interpoint distance information and direction information. Another important aspect in any of the types of spatial knowledge is recall time, and for this reason we look at orientation time, the time a subject takes to determine the direction of the target object and begin walking.

4.2.4.1 Initial Angular Error

To measure how well a subject is able to judge relative direction we employ a task similar to the classic pointing task of Rieser et al. [107] and Loomis et al. [70]. Subjects are instructed to walk to a target object per the testing protocol (Section 4.2.3) and the direction they walk measures the strength of their configural knowledge without respect to scale. Specifically, we measure the difference in angle between the direction of the actual target and the direction of walking after 1 meter [17].

4.2.4.2 Estimated Path Length

The other component of survey knowledge is interpoint distance information. We test how well subjects have acquired this metric information by employing a simple walking task. By controlling for the amount of time spent searching through the environment we can analyze how quickly survey knowledge is obtained, in particular its metric component. Note that metric information in these experiments is gained primarily through path integration and is therefore not likely subject to biases in perception that have generated a significant body of work in the virtual environments community (see Creem-Regehr et al. [23] for a recent review).

4.2.4.3 Recall Time

We can also measure the strength of subjects' survey knowledge by testing their recall time in determining the relative direction. Before subjects can begin walking they must know in what direction the target object is. By measuring the time the starting object is presented to the time subjects begins walking we can measure how well subjects can recall the relative direction of objects and orient themselves in the maze.

4.3 Experiment 3

4.3.1 Hypotheses

Given the fundamental differences in the proposed locomotion methods we developed two hypotheses for this experiment. First, given that the subjects are only actually translating during the resetting condition, we expect users in the resetting condition to exhibit better performance on the path length metric than users employing the other methods to walk. Our hypotheses derives from prior literature [113, 17] that demonstrates the importance of locomotion in acquiring spatial knowledge.

Regarding initial angular error, we have two WiP methods which do not manipulate rotation-based cues, and resetting, which does. We hypothesize that resetting's rotational manipulation will result in degraded performance for initial angular error.

4.3.2 Participants

For this study we recruited college age students from our institution between the ages of 18 and 25 (mean=20.7, median=21). Twenty subjects participated (9 male, 11 female), gave written consent, and were compensated $10 for participating in the experiment, which was approximately one hour in duration. Subjects were informed of their method of walking and the metrics we would collect (initial angular error and path length). Two subjects were excluded from data analysis due to system malfunctions, and therefore six subjects remained in each of the three conditions.

4.3.3 Equipment

The environment was developed in Unity and based on the maze developed by Chrastil and Warren [17]. A Samsung Gear VR head-mounted display (HMD) provided visual information to subjects. Subjects used either a Samsung S6 or S7 phone as the rendering device. The resolution in each eye was 1280x1440. The field of view of the Gear VR varied somewhat depending on the phone used; we did not measure this, but Samsung reports it as 96°; however, online reports place it at about 90°. Subjects' motion was tracked in one of three ways. The first was using a Vicon Motion Capture system in which body data were transmitted directly to the phone. Our system used 8 MX40 cameras to track the position of 6 optical markers and reconstruct the orientation and position of each subject's head. The physical space used was roughly 6x5 meters; the available tracked space was 5x4 meters. The second utilized the built-in IMU of the Gear VR to detect steps and the final method used a Microsoft Kinect v2.0, which used KinectVR [54] to transmit data using a Node JS server to the phone. All data were transmitted over a LAN using a NETGEAR WNR3500 router. Subjects provided input using the Gear VR's touchpad.

4.3.4 Procedure

Experiment 3 consisted of three between subjects conditions. The three conditions were resetting, WiP using the Kinect v2.0, and WiP using the Gear VR IMU. The exact implementation of each of these conditions can be found in Section 3.2. Subjects were randomly assigned to a locomotion method so that six subjects (3 men, 3 women) completed testing in each condition.

Subjects were first given instructions on how to move in their technique of walking. The first phase of the experiment was to place the subject in the training maze (Figure 4.2) and give them five minutes to freely explore the maze to learn how to walk around in their condition. Subjects were required to complete the full five minutes of this phase to ensure they were confident and competent at navigating. After this phase subjects were taken out of the headset and given instructions on what to do in the next maze. Subjects were also told to try and remember relative locations of objects as they would be tested later on them.

Next they were placed in the learning maze (Figure 4.3) and given 10 minutes to freely explore and learn the layout. They were not able to walk through the walls. At the end of 10 minutes subjects began the assessment phase and were placed in the Voronoi textured environment where they were given their next set of instructions.

In this phase subjects experienced a series of trials to find objects from various locations within the maze. To begin, subjects pressed the touchpad on the Gear VR and were placed back in the learning maze directly in front of an object, which allowed subjects to orient themselves. They were instructed not to walk around to prevent seeing any more of the maze. When oriented, subjects pressed the touchpad again placing them back in the Voronoi environment. The time taken to orient themselves in the environment was recorded as orientation time. Upon being placed in the Voronoi environment subjects were given another object in the maze to walk to via a heads up display (see Figure 4.5). Subjects were instructed to walk directly to the target object in a straight line. A second measurement of time (target acquisition time) was taken here to denote the time taken to recall the direction of the target object. This straight line condition ensured their walked path was a novel shortcut. The position in the maze, time, acceleration (for the Gear only condition), and orientation of the headset were recorded at every frame for potential reconstruction. Subjects indicated the conclusion of a trial by pressing the touchpad a final time. Each subject completed 40 trials in total. Beginning and end objects were randomly selected for each subject.

4.3.5 Results

To simplify analysis and remove the variability we divide our 40 trials into four blocks of 10 trials each. We present three measures of how well subjects have acquired spatial knowledge and, in particular, survey knowledge. All ANOVAs were performed using SPSS and the tests for normality and homogeneity of variances were met. Error bars in all figures denote standard errors of the mean.

4.3.5.1 Initial Angular Error

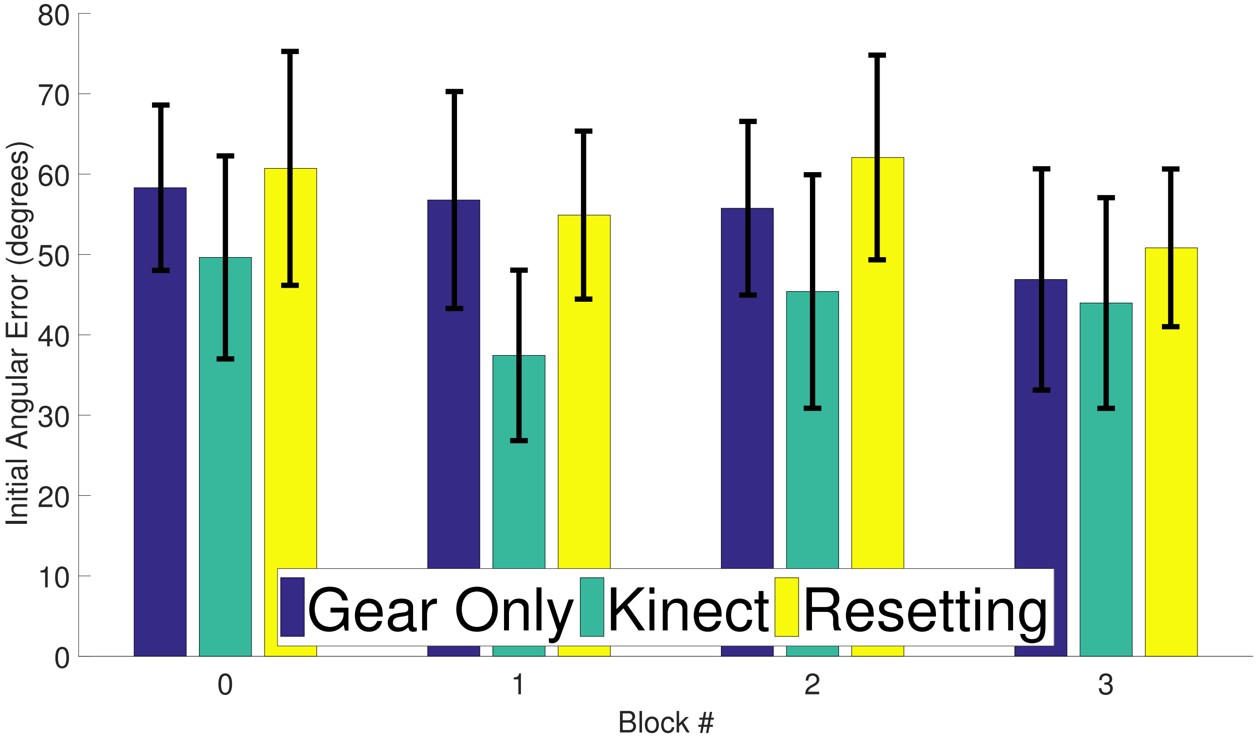

Figure 4.6: This figure shows the mean initial angular error of subjects by condition and blocks of 10 trials.

Figure 4.6 shows the initial angular error by condition and four blocks of 10 trials. A 3 (navigation method) x 4 (block) repeated measures ANOVA on absolute angular error found no effect of condition (F(2,15) = .338, p = .718) or block (F(3,45) = 1.171, p = .331). The lack of a significant effect differs from our second hypothesis, a surprising result given that the WiP conditions do not manipulate orientation whereas resetting does. In all conditions, the mean absolute angular errors are in line with those found by Chrastil and Warren[17].

To further explore this result we turned to Bayes factors, an analysis method that can provide support for the null hypothesis and expresses that evidence in an odds ratio. The following analyses use the methods of Rouder et al. [108] that, because they account for sample size, adjust for power. We set the prior odds to 1 as this favors neither the null nor the alternative. We first compare the methods of Gear only and Resetting which gives us a Jeffrey-Zellner-Siow (JZS) Bayes factor of 4.47 indicating strong evidence in favor of the null hypothesis. Comparing the Gear only and Kinect conditions gives us a JZS Bayes factor of 1.45 and the Kinect vs Resetting conditions gives us .47, both of which are marginal and do not strongly support either the alternative nor the null.

4.3.5.2 Path Length

The second type of information needed to accurately walk to a target object is interpoint distance information. Figure 4.7 shows how accurately subjects were able to walk the true distance. A 3 (navigation method) x 4 (block) repeated measures ANOVA on the error in relative path length shows a main effect of condition (F(2,15) = 4.923, p = 0.023). A post hoc Tukey HSD revealed that there is a significant difference between the Kinect and Resetting conditions (p = 0.022) and a marginally significant difference (p ¡ .1) between the Resetting and Gear only conditions. In both the Gear only and Kinect conditions subjects significantly overwalked the distance to the target (M=168% and M=181%, respectively) whereas in Resetting the distance was more accurate (M=111%).

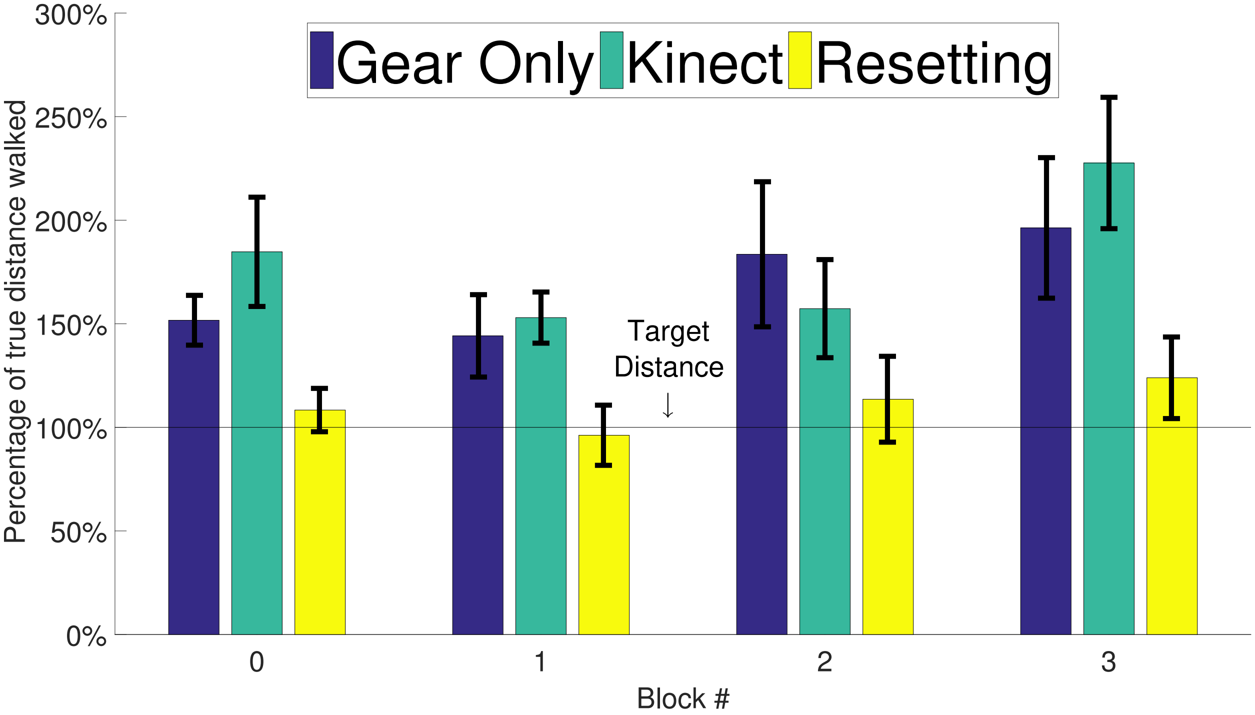

Figure 4.7: This figure shows the mean distance subjects walked as a percentage of true distance by condition and blocks of 10 trials.

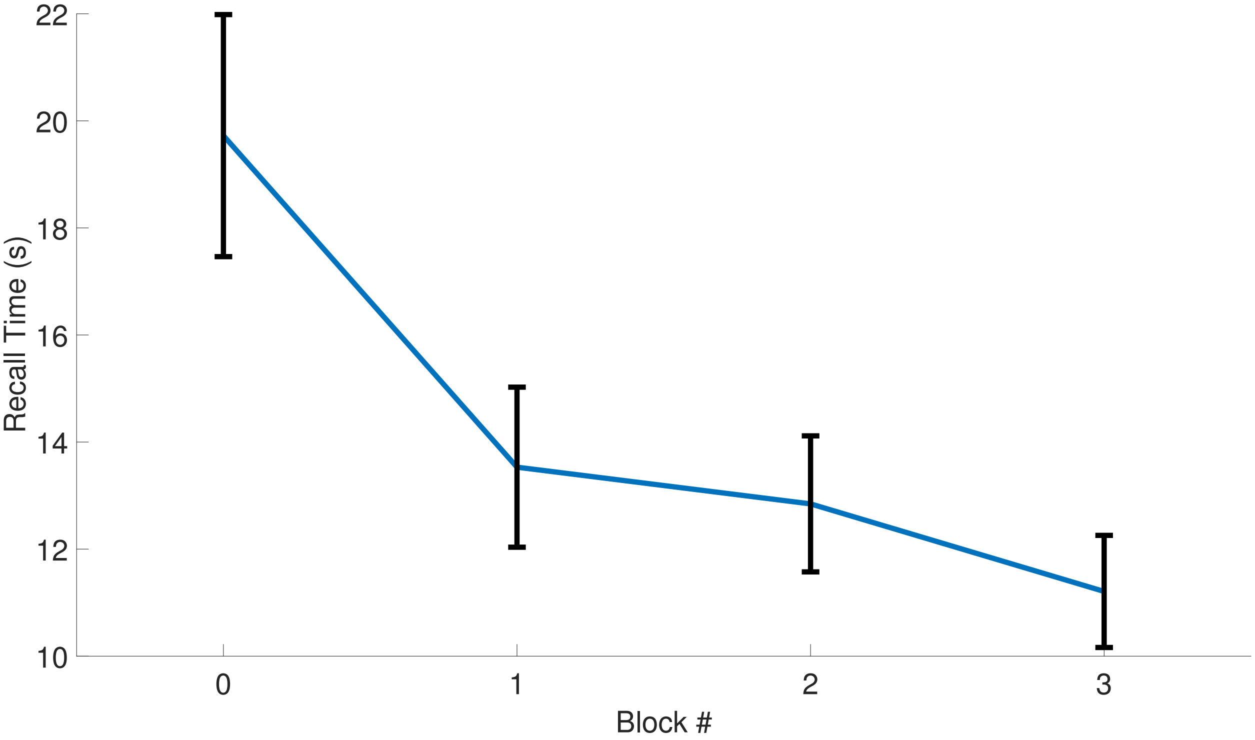

Figure 4.8: This figure shows the mean time subjects spent determining the relative direction of the target object. Time is plotted by blocks of 10 trials and collapsed across condition.

4.3.5.3 Recall Time

Finally, we analyze how quickly subjects are able to recall the maze and the target object. To do this we look at total recall time, i.e., the total time between being presented with the starting object and beginning to walk to the target object. A 3 (navigation method) x 4 (block) repeated measures ANOVA on this total recall time shows a significant main effect of block ($F(3,45) = 21.833, p < .001$, $\eta_p^2 = 0.59$). There is also a large drop-off between the first and second block of 6.2 seconds (a 31.4% drop). Given this large dropoff and the significance of the effect we removed the first block to see if this effect continued throughout the testing procedure. We performed a 3 (navigation method) x 3 (block) repeated measures ANOVA on total recall time for the last three blocks of data and we still find a main effect of block ($F(2,30) = 6.256, p = 0.005$, $\eta_p^2 = 0.29$).

There are two components to total recall time. The first component is the time subjects take to orient themselves within the maze (orientation time). The second component is the time subjects spend determining the direction of the target object (target acquisition time). Exploring further into which of these was the driving force in the previous result, we performed a 3 (navigation method) x 4 (block) repeated measures ANOVA on orientation time and target acquisition time and found significant main effects of block ($F(3,45) = 16.525$, $p < 0.001$, $\eta_p^2 = 0.52$, and $F(3,45) = 13.305$, $p < 0.001$, $\eta_p^2 = 0.47$, respectively). We also performed a 3 (navigation method) x 3 (block) repeated measures ANOVA on orientation time and target acquisition time for the final 3 blocks and again found main effects of block ($F(3,45) = 4.424$, $p = 0.021$, $\eta_p^2 = 0.23$, and $F(3,45) = 3.900$, $p = 0.031$, $\eta_p^2 = 0.21$, respectively). Thus, both components improve significantly by block.

4.3.6 Conclusions

The framework of spatial microgenesis [48, 15, 17] provides a structured methodology to frame the acquisition of spatial knowledge in humans. We use this framework as the primary measure of the usability for the three methods of navigating large virtual environments presented in this chapter. These three methods were resetting, WiP using the Samsung Gear VR, and WiP using the

Microsoft Kinect. Three metrics measured the acquisition of spatial knowledge in each of these three methods. The first metric was the initial angular error, which assessed the directional component of survey knowledge. We found no significant difference between the walking methods, and a Bayes Factor analysis found high odds that the Gear only and Resetting methods had identical performance. We have no theoretical reasons for why the Kinect might differ from the other conditions, and more work would be needed to ascertain if it indeed does.

The second metric was the path length of the novel shortcuts, which measured how well subjects were able to encode distance information into their cognitive map of the environment. Confirming our first hypothesis, subjects performed significantly better in the resetting condition than the WiP conditions, where they overshot the true distance by 68% in the Gear only condition and 81% in the Kinect condition. WiP does not seem to permit the same level of acquisition of metric interpoint distance survey knowledge as does resetting. Subject's overwalking in the WiP conditions might be explained by the complexity of the maze. This complexity may make the maze seem larger than it actually is, and it produced a pattern of overshooting in the absence of developed survey knowledge. It is possible, however, that this method of locomotion biases any attempt to measure the encoded metric information. Subjects were tested with only optic flow as feedback in the WiP conditions whereas subjects in the resetting condition have full idiothetic feedback.

The final aspect of spatial knowledge we examined was recall time. Across all conditions, subjects consistently improved in both measurements of map recall time. Over blocks, they were more rapidly able to localize themselves in the maze and to remember the relative direction of a target object. The steady decrease in recall time could be attributed to a strengthening of subjects' cognitive maps. Thus, while the maps did not get any more accurate, subjects recalled them faster. Since the directional component of survey knowledge is representative of direction between objects, subjects may be building a stronger but incorrect map from repeated attempts to recall their survey knowledge.

4.4 Experiment 4

Our third experiment demonstrated that there was a difference between the two WiP methods and the resetting method. Subjects could not accurately reproduce distance in our WiP methods. Our fourth experiment seeks to determine why WiP prevented subjects from being able to reproduce distance. We separated the mechanism by which subjects learned and tested their acquisition of spatial knowledge. This separation gave us four conditions and allowed us to determine if the walking method used for testing or learning prevented the accurate reproduction of distance. The four conditions are described in complete detail in Section 4.4.5.

Moreover, individual ability in navigation performance varies wildly, a trend that persisted during our task. We also wish to understand what causes the origins of these differences so we also include standard measures of individual difference. First, by reducing the individual variance inherent in our task we can begin to understand the full effect on navigation performance of our methods of locomotion. Second, measuring individual differences allows us to begin to select a locomotion method appropriate for a subject based on skill, strategy, and experience.

4.4.1 Hypotheses

This experiment modifies the previous experiment and, depending on condition (Table 4.1), subjects may use a different locomotion method during the learning and testing phases. The resulting four conditions allows us to separate the effects of the locomotion method used to learn from the one used to test for the acquisition of knowledge. Based on the work of Hanson et al. [35] and others [53, 7], we know that optic flow may not be sufficient for subjects to fully integrate the distance they have traveled. Additionally since the environments provide a large number of cues for distance traveled we expect that the difficultly will lie in reproduction of distance traveled, not in the acquisition of the scale of the environment.

Our second hypothesis is that while we expect to see a difference in distance reproduction for either the mechanism of learning or testing, we do not expect this to effect the angular error measured in our task. By testing for individual differences we expect that we will be able to reduce variance and strengthen the conclusions reached in the previous experiment. This will require some correlation between performance and our measures of individual difference, so our third hypothesis is that these correlations will be positive as in prior work on navigation in similar tasks [39, 17].

4.4.2 ID Tests

Participants completed short pretests to assess their individual levels of skill in various measures which have been shown to relate to performance in similar tasks. The first was the extended range vocabulary test (EVRT), which served to measure the general verbal abilities of subjects. Subjects next completed a mental rotation test (MRT) that has shown high correlation with navigation performance in VR Hegarty et al. [39]. The MRT serves to measure small scale spatial abilities that should be interesting in the case of resetting performance, which requires mental rotation. The final test is the Santa Barbara Sense of Direction test (SBSOD) [39]. Various papers have shown this self report measure to correlate with general navigation performance, but this correlation is lessened in VR. Questions ask subjects to report their own abilities in wayfinding.

4.4.3 Participants

This experiment had the same subject pool of college age students from our institution (mean = 20.4, median = 20) and all 114 participating subjects (52 male, 62 female gave written consent and were compensated $15 for 90 minutes of their time. Instructions were identical to Experiment 3 with the exception that some subjects would be told about both methods of locomotion. Fourteen subjects withdrew from the study and 2 subjects were excluded due to system errors.

4.4.4 Equipment

The equipment was the same as in Experiment 3 with the exception of the tracking systems. The Kinect was eliminated from consideration and the Vicon system was replaced with the World-

Viz Precise Position Tracking (PPT) system, which used four cameras to track the position of a single optical marker. Orientation was tracked using the Gear VR's internal IMU.

4.4.5 Procedure

At the start of the experiment, subjects took several tests to measure their innate spatial cognition. First, subjects received a standard vocabulary test. Afterward, subjects matched identical three-dimensional shapes that had been rotated to different perspectives [137]. Lastly, subjects personally rated themselves on the Santa Barbara Sense-of-Direction Scale [39]. All these tests gave us a baseline understanding of the subjects initial spatial aptitude before commencing the virtual reality portion of the experiment. The virtual reality test consisted of two five-minute training sections, one ten-minute learning section, and an un-timed assessment section, in that order. The first training section and the learning section employed the same technique of walking, and so did the second training section and assessment section. Thus, there were four conditions in this experiment, one for each permutation of walking-in-place and resetting. Note that for homogeneous pairs (e.g., two resetting), the two five-minute training sections were combined into a single ten-minute section. See Table 4.1 for more detail.

Condition	Training 1 (5 minutes)	Training 2 (5 minutes)	Learning (10 minutes)	Testing (40 trials)
1	Resetting	Resetting	Resetting	Resetting
2	Resetting	WiP	Resetting	WiP
3	WiP	Resetting	WiP	Resetting
4	WiP	WiP	WiP	WiP

Table 4.1: Virtual reality based phases of Experiment 4 for the four between subjects conditions. Twenty-four subjects completed each of the above conditions.

Subjects were first given instructions on how to move in their walking technique. The first phase of the experiment was to place the subject in the training maze and give them five minutes to freely explore the maze to learn how to walk around in their condition. Subjects were required to complete the full allotted time of this phase to ensure they were confident and competent at navigating. If subjects were in a heterogeneous pair (i.e., first resetting then walking-in-place), then subjects would repeat this phase with a different walking technique. After training was over, subjects were taken out of the headset and given instructions on what to do in the next maze. Subjects were also told to try and remember relative locations of objects as they would be tested later on them. Next they were placed in the learning maze and given 10 minutes to freely explore and learn the layout. This walking technique matched the technique employed in the first training section. They were not able to walk through the walls.

At the end of 10 minutes subjects began the assessment phase and were placed in the Voronoi textured environment where they were given their next set of instructions. In this phase subjects experienced a series of trials to find objects from various locations within the maze. To begin, subjects pressed the touchpad on the Gear VR and were placed back in the learning maze directly in front of an object, which allowed subjects to orient themselves. They were instructed not to walk around to prevent seeing any more of the maze. When oriented, subjects pressed the touchpad again placing them back in the Voronoi environment. The time taken to orient themselves in the environment was recorded as orientation time. Upon being placed in the Voronoi environment subjects were given another object in the maze to walk to via a heads up display. Subjects were instructed to walk directly to the target object in a straight line. A second measurement of time (target acquisition time) was recorded here to denote the time taken to recall the direction of the target object. This straight line condition ensured their walked path was a novel shortcut. The position in the maze, time, acceleration (for the Gear only condition), and orientation of the headset were recorded at every frame for potential reconstruction. Subjects indicated the conclusion of a trial by pressing the touchpad a final time. Each subject completed forty trials in total. Five repetitions of eight paths as in Chrastil and Warren [16].

4.4.6 Results

Similar to Experiment 3 we reduced the amount of data and variability by collapsing across trials to reach a single average for each metric that will be discussed in this section. Data from Experiment 3 showed that there was no effect of time between the conditions, either as a main effect or an interaction. This trend persisted in Experiment 4 and for that reason we do not include it in our analysis. We present the remaining two measures of survey knowledge as well as our results regarding individual differences. All ANOVAs were performed using SPSS and the tests for normality and homogeneity of variances were met. Error bars in all figures denote standard errors of the mean.

4.4.6.1 Path Length

The main purpose of this experiment was to determine the cause of the interpoint distance errors revealed in Experiment 3. We ran a 2x2 (testing x learning) ANOVA on path length error with MRT and SBSOD as covariates to determine the root cause. We found a main effect of testing ($F(1,90) = 60.047, p < .001$). Post hoc comparisons showed that the conditions in which subjects tested in WiP resulted in subjects walking 112% of the true distance further than those who tested in resetting. Learning showed no effect on path length. Figure 4.9 shows the effect of testing on distance recreation.

4.4.6.2 Initial Angular Error

To detect differences in configural knowledge we ran a second 2x2 (testing x learning) ANOVA on initial angular error with MRT and SBSOD as covariates. The ANOVA revealed no significant effect of learning or testing but showed that the MRT was a significant predictor of performance in this metric ($F(1,90) = 28.694, p < .001$). Furthermore, MRT was significantly correlated with initial angular error ($r = -.510, p < .001$).

4.4.6.3 Individual Differences

Analysis of the correlations of SBSOD with the initial angular error performance of those who learned in resetting and those who learned in WiP revealed large differences (r = .064 for resetting vs r = -.39 for WiP). We computed a Fisher r-to-z transformation to compare these and found a significant difference ($Z = 2.27, p = .0232$).

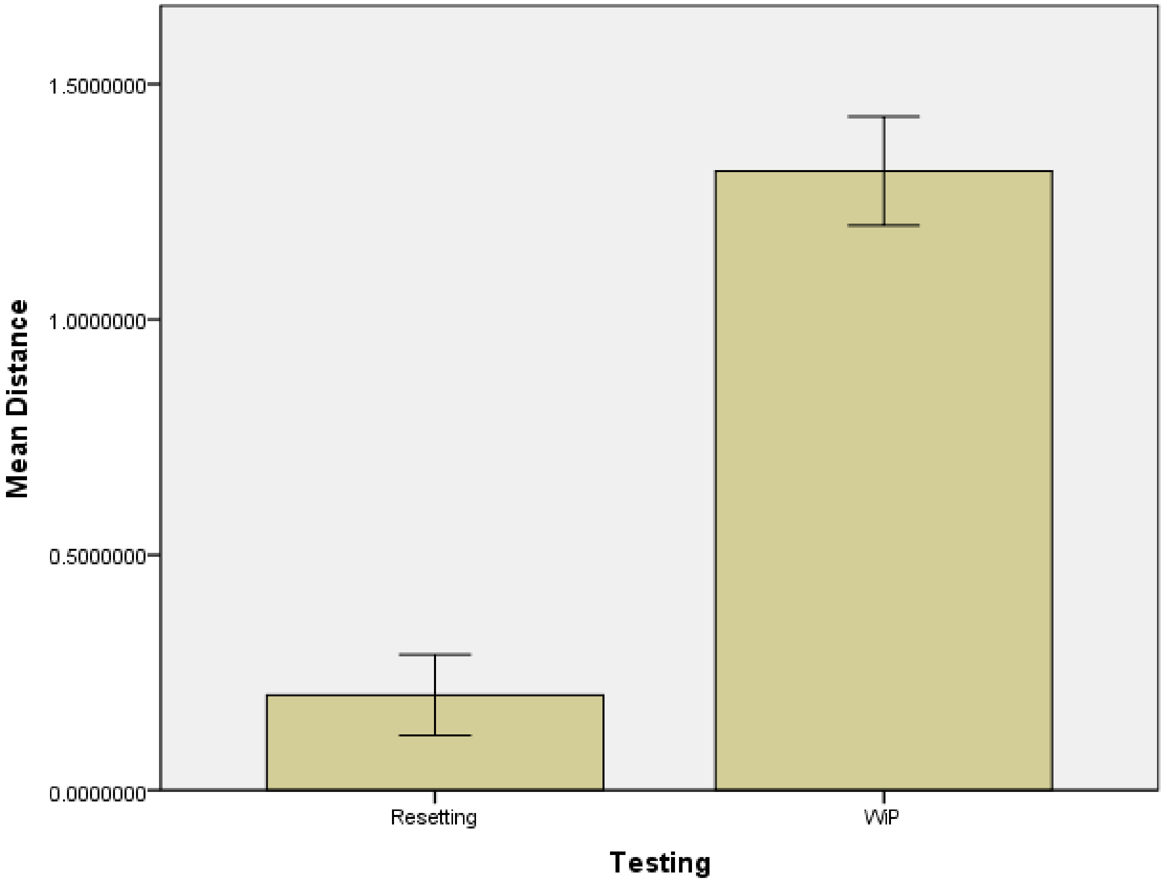

Figure 4.9: This figure shows the distance overwalked when testing in each condition.

	Mean	Median
SBSOD	4.158	4.130
MRT	29.47	29.50

Table 4.2: Mean and median SBSOD and MRT scores among all participants

To investigate these differences we coded those with high SBSOD as having higher than the median SBSOD and similarly coded MRT scores. The median scores are reported along with mean scores in Table 4.2 We computed a 2x2x2x2 (learning x testing x MRT x SBSOD) ANOVA and found a interaction of learning x MRT x SBSOD ($F(1,80) = 5.006, p = .028$). Post hoc t-tests showed that having stronger mental rotation abilities resulted in lower angular error, this trend was broken for those subjects with low SBSOD who learned the maze with the WiP locomotion method. This can be seen in Table 4.3 which shows pairwise t-test comparisons among four groups.

Learning Method	SBSOD Score	High MRT	Low MRT	p < .01
WiP	High	36.209	63.887	**
WiP	Low	57.652	64.607	
Resetting	High	46.216	70.934	**
Resetting	Low	37.188	66.190	**

Table 4.3: Comparison of improvement in angular error between those with weak and strong mental rotation abilities.

Additional post hoc t-tests revealed that resetting outperformed WiP only when subjects had low SBSOD and high MRT scores ($\delta = 27.419, p = .009$). Additional correlations between our dependent measures and individual difference metrics are shown in Table 4.4

	Direction	Distance	SBSOD	MRT
Direction	1.000	-0.032	-0.172	0.250
Distance	-0.032	1.000	0.004	0.036
SBSOD	-0.172	0.004	1.000	-0.510
MRT	0.250	0.036	-0.510	1.000

Table 4.4: Correlations between dependent measures and the individual differences measured in this experiment.

4.4.7 Discussion

4.4.7.1 Testing Method Effects

The results of the fourth experiment showed that the distance expansion revealed from the third experiment is caused by a testing effect when subjects test using WiP. This effect likely indicates that subjects do not have a sense of how far they have traveled under WiP and are relying on the visual information provided by the environment rather than their own body-based cues. This finding also tells us that any future testing should be done using resetting or some real walking technique to prevent a testing effect on expression. There were no testing effects on direction which makes sense because subjects effectively choose their direction before any locomotion method is used.

4.4.7.2 Learning Method Effects

In our analysis we have four groups of people based on an individual's mental rotation abilities and sense of direction. Subjects can be strong in both, neither, or only one of the two. Subjects were able to learn an environment in one of two ways: either by resetting or WiP. WiP outperformed resetting, but, as our 3-way interaction of learning x MRT x SBSOD revealed, only for subjects with high MRT scores and low SBSOD scores. In the literature, e.g., Hegarty et al. [39], navigation performance is generally correlated with high MRT scores and this work confirms those results. SBSOD scores are also correlated with navigation performance, however, in this work this result only held when subjects learned by WiP. Resetting differs from typical walking in that it is incongruent to directional body-based cues. This correlation may explain why those who have a better sense of direction show lower performance in resetting. They may either decide to ignore the directional information or may have a more imprecise map created by bad information. We believe the reason this finding only shows up in those who have strong MRT scores may be that they naturally make use of directional body-based cues while the other group does not.

4.5 General Discussion

These experiments show that resetting is a good general purpose method for navigation. Survey knowledge can be acquired regardless of the characteristics of the environment. None of our measures of individual differences indicate that there are people who have difficulty using resetting for locomotion. This finding is perhaps a surprising result given the "2-1" turn that is fundamental to resetting.

In contrast, WiP results in some difficulty in assessing the distance component of survey knowledge. This difficulty only exists in the reproduction of distance knowledge as subjects could accurately reproduce distances learned with WiP but tested with resetting. Looking at other work [35, 91] the question then arises, why was optic flow not enough in this experiment to give subjects accurate knowledge of their distance traveled? In this study, optic flow was tied to the step rate of the subject and may have been mismatched from what they would normally expect. Various strategies can be employed to reproduce the distance and subjects who tested in walking in place may have used a strategy based on time spent walking rather than optic flow. This strategy could lead to a mismatch if participants acquired an accurate sense of scale, but not of the speed at which they would walk in our implementation of WiP.

It is important to note that subjects were able to acquire distance knowledge using WiP, which means that with sufficient landmarks reproduction does not seem to be an issue. In the work of Hanson et al. [35] subjects did not get optic flow as feedback during distance testing but were accurate at estimating distances with WiP. However, in that work they trained and likely became calibrated to WiP. While there are a number of differences between that work and the present one, it may come down to the second training session. The subjects in condition 3 (see Table 4.1) were allowed to actually walk around a similar maze and may have begun to gain an understanding of the scale of the maze. Still, the two tasks differed in the time between learning and testing. In Hanson et al. subjects immediately attempted to reproduce the distance whereas in our task they needed to commit and then recall the distance. This difference may have caused a shift in strategy and follow up work may be warranted to understand how subjects approached each task.

Perhaps the most interesting result is that for WiP individual differences seem to play a key role. Users with low SBSOD scores may require additional information or more time training and learning to perform adequately. Designers of virtual environments will need to consider these results and adjust accordingly.

4.5.1 Future Work

This study leads to many interesting questions on how different navigators make use of different information. Future work could be done to determine if subjects are making imprecise maps or ignoring the information available to them. The trend, only occurring in some navigators, makes testing the precision of one's cognitive map difficult to test for, but should inspire the use of other individual difference tests when looking at these walking methods.

Chapter 5

Specific Aim 3

5.1 Introduction

Virtual reality (VR) can provide us with experiences that would normally be prohibitively expensive or infeasible in the real world. It even allows for impossible and non-existent environments (e.g., Mordor or ancient Rome). With VR, something as amazing and breathtaking as taking a walk around Paris is not only possible, but almost trivial. However, in this case walking is a misnomer. Many VR systems allow “walking” through the use of a controller to move or by teleporting. Teleportation and the use of a controller are more aptly described as locomotion methods, or methods of self induced motion that permit the exploration of an environment. Those two are examples of unfulfilling or unnatural locomotion methods. Walking, as we do it in real life, is a locomotion mode we want to get to in VR systems.

Locomotion in VR is a heavily researched topic with many ways of evaluation [114]. Some studies take the approach of developing ways to move or walk through a virtual environment [93]. These studies examine whether the method is feasible or enjoyable. Other studies are concerned with evaluating methods based on the mathematics behind locomotion, in particular, redirected walking [127, 94]. These studies look at minimizing the breaks in presence. Finally, a number of studies look at how users perform when using different locomotion methods [149]. These studies take a more objective measurement to determine which method is better. This approach is the one we will take in this work, with the ability to navigate as our metric.

Navigation is a critical part of experiencing and exploring the world. Learning the layout or distances and directions between objects is necessary in understanding one’s environment. Locomotion is important because studies[105, 111] have shown that the method of walking, and hence the information provided to our senses, is important in understanding or navigating the environment we are in. The ability to navigate naturally occurs when exploring a real environment, making it a good measurement of the fidelity of the locomotion method.

The proliferation of tetherless VR has marked a major change in what we think of in terms of locomotion. There is now more freedom in how we are able to move around in VR without cables or the restrictions they cause. Still, newer tetherless HMDs do not necessarily have innate tracking systems. The means that a large gap in the research on room size and tracking has proven to be critical moving forward. The majority of research on room size has related to large spaces [4, 99]. These studies don't fit with consumer level VR technology as it exists out of the box today (e.g., HTC Vive or Oculus Rift).

Consumer level technology is limited by the available space a consumer has. In many instances that is no more than a single room. We address this challenge in this work. How do we pick a locomotion method for a given, limited consumer-size space? We consider the maximal available space to be 4x4 m^2, the default maximum space of the Vive. To understand the effect of smaller spaces we also look at 3x3 m^2, 2x2 m^2, and standing/personal space (1x1 m^2).

We will choose a reorientation locomotion method [149] as one of the methods to employ in our evaluation. Reorientation is a classic method used in a number of locomotion methods for last second collision avoidance. Hence, we choose to evaluate its efficacy in some of our selected spaces. Piloting showed reorientation to be incredibly uncomfortable and undesirable in standing space, however, so we choose to use walking in place (WiP) as an alternative. Including WiP allows us to evaluate reorientaion against a technique that lacks some translational body-based information but has more accurate rotational body-based information (see Table 2.1).

Thus, this work evaluates reorientation, in three differently sized spaces, and WiP, in standing space, based on how well each affords the acquisition of spatial knowledge. To measure this affordance we turn to the theory of "Spatial Cognitive Microgenesis", a structured framework which has describes three stages of spatial knowledge [118]. We are interested in survey knowledge, which represents one's structural knowledge of an environment. Survey knowledge has two components, interpoint distance and direction, which we test individually following the framework of Chrastil and Warren [17].

Several studies [90, 17, 48, 39] have shown large variations in performance on navigation tasks. Some navigators have stronger abilities than others [39] and these individual differences in ability make testing the acquisition of spatial knowledge difficult. In this study we test each subject for a number of individual differences to reduce the variation and strengthen our conclusions.

5.2 Research Design

5.2.1 Locomotion Methods

This work implements two locomotion methods. The first locomotion method is the WiP algorithm taken from Paris et al. [90] and Hanson et al. [35]. We use the basic WiP method for the Samsung Go. WiP works by inducing forward motion when a step is detected. It is a form of simulated walking and does not require its user to physically translate. The second method was the reorientation technique called resetting. It was taken directly from Paris et al. [90] and adjusted only to fit into one of three spaces. These spaces will be introduced further in 5.2.4. Resetting involves adjusting the rotational gain when a boundary is encountered causing the world to rotate around its user when a reset is required. Resetting was designed so that the user does not notice that the rotational gain is different from 1.0, when in fact it is closer to 2.0.

5.2.2 Sample Size

In this experiment we are interested in both accounting for and detecting individual differences in navigation ability and performance. The work completed in Chapter IV showed a correlation coefficient between SBSOD and navigation performance for the resetting only condition of 0.38. A power analysis revealed that to achieve a power of 0.8 for this correlation coefficient we needed 52 subjects in the resetting conditions.

Another goal of this experiment was to detect any useful differences in angular error caused by the locomotion method or by the available tracked space. We completed a power analysis for a between group study with four groups. We assumed the only effect worth finding would be a medium effect size (Cohen's d = 0.3) and determined that one hundred four subjects (26 per condition) would be sufficient. The latter analysis yielded a higher sample size than the former so for this study we used a design of one hundred four subjects.

5.2.3 Metrics

In this study we want to measure how well different room sizes allow for the acquisition of survey knowledge. As with our prior studies we choose to use the framework presented by Chrastil and Warren[16]. This gives us two independent metrics: direction and distance. This experiment uses those measures to determine how strongly one has acquired survey knowledge.

5.2.4 Tracked Spaces

The four conditions in this experiment corresponded to four differently sized spaces, they were 4x4 m^2, 3x3 m^2, 2x2 m^2, and 1x1 m^2 (standing space). Piloting showed that resetting in a 1x1 space was too difficult and so we used only walking in place in that sized space. In the other three spaces we used the resetting locomotion method. Table 5.1 has complete details on the conditions, tracked space, and locomotion method used in each part of the experiment.

Condition	Training	Learning	Testing
WiP	WiP	WiP	Resetting (4x4 m^2)
4x4	Resetting (4x4 m^2)	Resetting (4x4 m^2)	Resetting (4x4 m^2)
3x3	Resetting (3x3 m^2)	Resetting (3x3 m^2)	Resetting (4x4 m^2)
2x2	Resetting (2x2 m^2)	Resetting (2x2 m^2)	Resetting (4x4 m^2)

Table 5.1: The four conditions in this study as well as the locomotion method (and tracked space if applicable) in each of the three VR phases of the experiment.

5.3 Methods

Experiment 5 had five parts which are presented in this section. Subjects first completed a number of tests of individual difference (Section 5.3.3). The following three phases were in VR and are described in detail in Section 5.3.5. Finally subjects completed a post-test debriefing.

5.3.1 Participants

We recruited college age subjects from our city between the ages of 18 and 25. One hundred forty (58 men, 82 women) subjects between four conditions participated and were compensated $15 for their time. All subjects who entered VR were included in the simulator sickness portion of the experiment. Otherwise, 25 subjects dropped out before completion of the experiment and were excluded from part of the analysis. Eleven subjects had to be excluded from part of the analysis due to computer errors. Conditions were balanced so that 26 subjects completed each of the four conditions and subjects were assigned their condition randomly.

5.3.2 Equipment

The environment was developed in Unity and based on the maze developed by Chrastil and Warren [17]. A Samsung Go head-mounted display (HMD) provided visual information to subjects. The resolution in each eye is 1280x1440 with refresh rate of 60hz. The field of view of the Samsung GO was at least 110°. We tracked position in two ways. In all conditions we tracked subjects using a WorldVIZ Precise Position Tracking system, which allowed us to provide 6DOF tracking. The physical space was roughly 6x5 meters and the tracked space was 5x5 meters. We placed foam interlocking mats on the floor to mark off the 5x5 meter space, which ensured subjects could not walk into a wall. For one of the four conditions, WiP, we used the IMU of the Samsung Go to detect vertical linear acceleration. To allow subjects to interact with the experiment, they were given a Samsung Go controller.

5.3.3 Individual Difference Measures

Prior to the experiment subjects completed a number of tests of individual differences. The first was the expanded vocabulary range test (EVRT) used as a measure of general intelligence to ensure all effects were ability based and not intelligence based [26]. Second was the Vandenburg mental rotation test (MRT) which has subjects determine objects which are rotations of the given object [137]. To test the capacity of a subject's working memory we administered the CORSI block tapping test [22]. Finally subjects completed the Santa Barbara sense of direction test (SBSOD) which measures a subject's confidence in navigating [39].

5.3.4 Environment

Figure 5.1: Top-down view of the environment used in the practice phase of Experiment 1. There are four objects for the subject to find and four landmarks (paintings).

In this experiment each subject was presented with three distinct environments. The first was a training maze, shown in Figure 5.1. The training maze was roughly 6x6 m^2. Subjects were placed in a practice maze in order to train in their assigned room size. The second was the learning

Figure 5.2: Top-down view of the environment used in the learning phase of Experiment 1. There are eight objects for the subject to find and four landmarks (paintings) to facilitate learning.

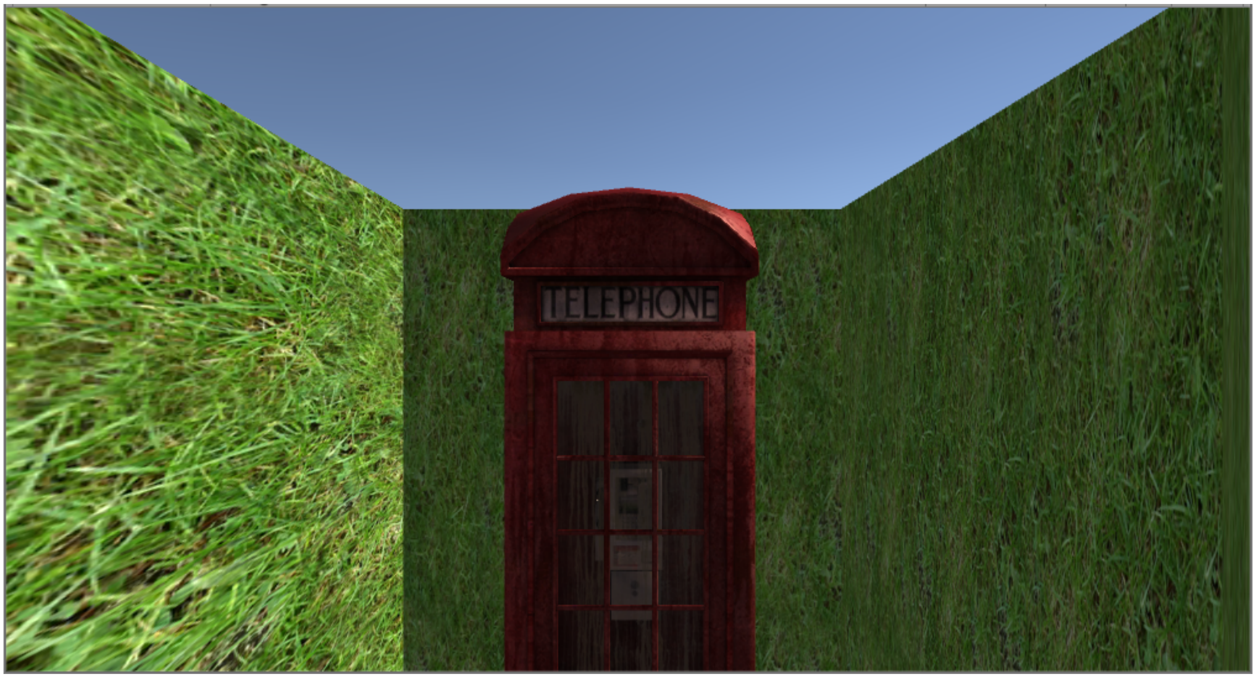

Figure 5.3: This figure shows the telephone booth as would be seen by subjects during the orientation phases of the experiment.

Please go to Snowman

Figure 5.4: At the beginning of the testing phase subjects are informed of the target object via a heads up display. This disappears shortly so as not to distract the subject during walking.

maze shown in Figure 5.2; subjects were instructed to learn the spatial relations among the eight objects contained in this maze. Four landmarks were present to aid in learning the overall layout. A first person view of the learning maze is shown in Figure 5.3, with one of the eight objects. The learning maze was roughly 10x10 m^2. Due to the geometry of the maze, a subject could not see any two objects at the same time. The final environment is presented during the actual testing phase. Figure 5.4 shows a first person view while the subject is in a sparse environment with a Voronoi textured ground plane to give the subject some ocular flow for feedback on distance traveled.

The study presented in Chapter IV demonstrated the importance of keeping the locomotion method used for testing one's survey knowledge consistent across conditions. In this study we leveraged that knowledge and for every subject the testing phase was completed in the 4x4 m^2 tracked space (see Section 5.2.4) using the resetting locomotion method.

5.3.5 Navigation Task

Subjects were first given instructions on how to move in their technique of walking. The first phase of the experiment was to place the subject in the training maze (Figure 5.1) and give them five minutes to freely explore the maze to learn how to walk around in their condition. Subjects were required to complete the full five minutes of this phase to ensure they were confident and competent at navigating and using their assigned locomotion method. After this phase subjects were taken out of the headset and given instructions on what to do in the next maze. Subjects were also told to try and remember relative locations of objects as they would be tested later on them.

Next they were placed in the learning maze (Figure 5.2) and given 10 minutes to freely explore and learn the layout. They were not able to walk through the walls. At the end of 10 minutes subjects began the assessment phase and were placed in the Voronoi textured environment where they were given their next set of instructions. Those subjects who learned in the WiP condition were given instruction on how resetting works.

In the testing phase subjects experienced a series of trials to find objects from various locations within the maze. To begin, subjects pressed the a button on the Go and were placed back in the learning maze directly in front of an object, which allowed subjects to orient themselves. They were instructed not to walk around to prevent seeing any more of the maze. When oriented, subjects pressed the button again placing them back in the Voronoi environment. Upon being placed in the Voronoi environment subjects were given another object in the maze to walk to via a heads up display (see Figure 5.4). Subjects were instructed to walk directly to the target object in a straight line. This straight line condition ensured that their walked path was a novel shortcut. Subjects indicated the conclusion of a trial by pressing the button a final time. In order to reduce potential variance the real world location of each subject was the center of the tracked space. Each subject completed forty trials in total consisting of five repetitions of eight pairs of objects.

5.3.6 Metrics

The objective metrics analyzed in this study were the same as those in Experiment 4. For our first metric, angular error, we took the difference between the walked angle and the correct angle to the target object after 1 meter. Our second metric, interpoint distance, we measured the ratio of the distance traveled in the virtual environment divided by the true distance from the starting object to the target object.

5.3.7 Simulator Sickness

To assess the undue simulator sickness caused by each of the four conditions (see 5.2.4) we measured the discomfort induced as in Rebenitsch and Owen [100]. During the learning portion of the experiment (see Section 5.3.5), which lasted 10 minutes, every minute subjects reported their current level of simulator sickness on a scale from 1-10. A baseline measurement was taken at the beginning of the learning phase immediately following the subject donning the helmet.

5.3.8 Post Test

After completing the experiment subjects completed a post test questionnaire designed to determine if they were able to notice any induced rotation from the system. Each subject was then interviewed and asked questions regarding the rotation of the environment and strategies for exploring, learning, and recalling the environment. We were interested in seeing if subjects could detect the rotation induced from resetting and asked various masking questions to ensure they did not know our intent. The questions of the questionnaire can be seen in Table 5.2 and were presented as a Likert scale from 1 to 5. The interview was semi-structured and questions were asked based on the responses to the questions in Table 5.3.

1	I felt like the virtual world was turning
2	I saw the virtual world get smaller or larger
3	I saw the virtual world flicker
4	I was getting bigger or smaller
5	I saw the virtual world get brighter or dimmer
6	I felt like I was turning when I wasn't

Table 5.2: Post test questionnaire presented to each subject. Questions 3-6 were masking questions and all 6 questions were presented in a random order.

1	Did you notice anything unusual about the environment?
2	How did you go about exploring the environment?
3	What was your strategy to learn the objects?
4	What was your strategy to recall the locations of the objects?
5	How did you decide how far to walk?
6	Did you use the resetting intervention to measure distance?

Table 5.3: Post test interview presented to each subject. Followup questions were asked based on responses to each question in this table.

5.4 Hypotheses

Of the four conditions only WiP provides fully accurate rotational body-based information. It is also the only condition that does not have an associated cognitive cost[149]. Our first hypothesis is that WiP will result in the lowest angular error of the four conditions and that as the size of the space available for resetting shrinks, the angular error will go up.

Mismatches in body-based and visual information can lead to increased simulator sickness. So our second hypothesis is that the increased number of resets will lead to increased simulator sickness.

While we are testing and accounting for many individual differences in this study, only the mental rotation test (MRT) and the Santa Barbara sense of direction test (SBSOD) have already shown to be predictive of navigation performance in this dissertation. Additionally, the task in this study and the CORSI block tapping task both utilize spatial memory. This leads to our second hypothesis that all three of these tests will be predictive of navigation performance in this study.

5.5 Results

5.5.1 Angular Error

Confirming our first hypothesis, angular error was lowest during the WiP condition. Prior work in Chapter IV showed that subjects could acquire accurate configural information through both resetting and WiP. We conducted a 4x2 repeated measures ANOVA (condition by gender) with covariates for MRT, CORSI, SBSOD, and simulator sickness and found a main effect of condition ($F(3,92) = 2.742, p = .048$) and SBSOD ($F(1,92) = 10.31, p = .002$). Post test comparisons revealed significant differences between the 2x2 and WiP conditions, as well as the 3x3 and WiP conditions. The WiP condition resulted in significantly better configural knowledge than those (3x3 and 2x2) conditions. These results can be seen in Figure 5.5.

	Mean	Median
SBSOD	3.953	4.000
MRT	28.09	27.00
Corsi	6.722	6.333

Table 5.4: Mean and median SBSOD and MRT scores among all participants

We performed a correlation analysis on our dependent measures and covariates and found sev-

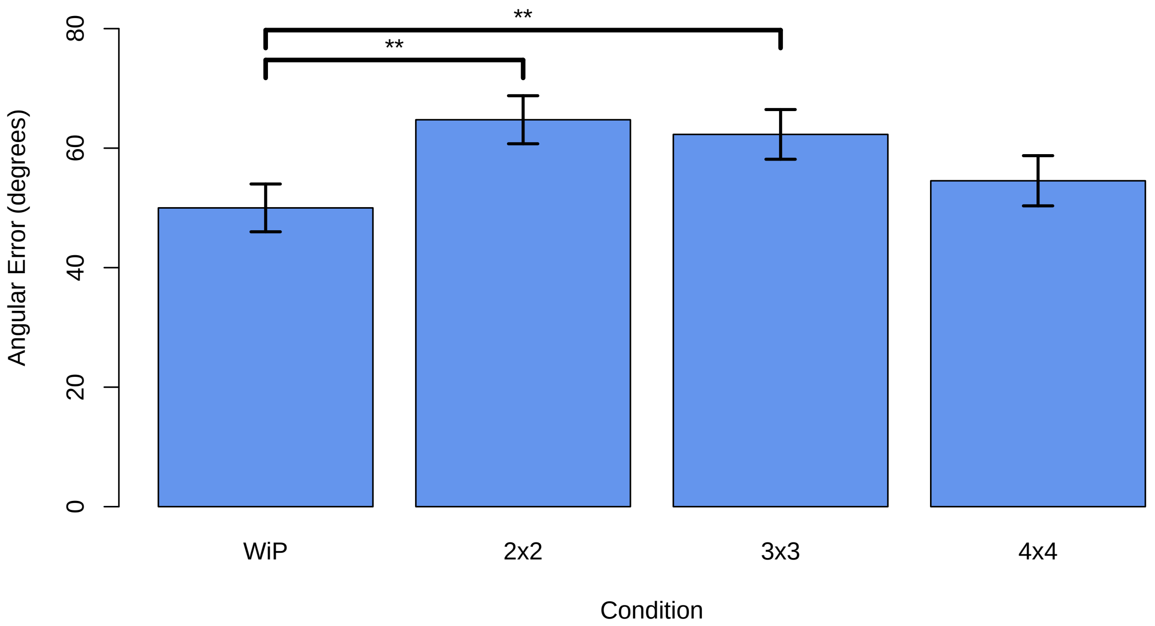

Figure 5.5: This chart shows the estimated marginal mean angular error in each of the four conditions and the standard error of the mean. These means take into account the individual differences measured in this experiment.

eral significant correlations. Angular error was significantly correlated with both MRT scores ($r = -.211, p = .032$) and SBSOD scores ($r = -.386, p < .001$). MRT scores were also correlated with SBSOD scores ($r = .290, p = .003$) and CORSI scores ($r = .306, p = .002$). Refer to Table 5.5 for the full correlation analysis and to Table 5.4 for mean and median individual difference scores among subjects in this experiment.

5.5.2 Simulator Sickness

To analyze the simulator sickness scores given, we performed a repeated measures ANOVA with two factors (gender and condition) with three covariates (MRT, CORSI, and SBSOD). This analysis revealed no main effects of either factor, nor did it reveal an effect of time. Subjects did not in general show an increase in simulator sickness in any of the conditions. There were,

	Direction	Distance	SBSOD	MRT	CORSI	SSQ
Direction	1.000	0.010	-0.386	-0.211	-0.189	0.030
Distance	0.010	1.000	0.023	0.141	0.082	-0.150
SBSOD	-0.386	0.023	1.000	0.289	0.109	0.044
MRT	-0.211	0.141	0.289	1.000	0.306	-0.001
CORSI	-0.189	0.082	0.109	0.306	1.000	0.030
SSQ	0.030	-0.150	0.044	-0.001	0.030	1.000

Table 5.5: Correlations between dependent measures and the individual differences measured in this experiment.

however, a number of subjects who did get sicker and even dropout. The dropout rate was 18% but the condition did not seem to affect the number of dropouts as shown by Table 5.6. A dropout rate of 18% is higher than in other locomotion studies, and some aspect of our experiment may have led to a higher than normal dropout rate.

Condition	Dropouts
WiP	6
4x4	5
3x3	7
2x2	7

Table 5.6: Number of dropouts due to simulator sickness occurring in each of the four conditions.

5.5.3 Questionnaire Responses

Nearly every participant (92%) noticed that something was occurring during the resetting locomotion method. Many of those subjects reported that they noticed something because they walked much further than was reasonable in the physical environment. A few (15%) remarked that they used something external to the virtual world to determine what was going on. For example, some employed a strategy such as turning exactly 90° twice to realize they were only turning halfway around. There did not seem to be any differences in how quickly subjects realized an intervention was occurring based on resetting condition.

Most participants attempted to explore the maze in one of two ways. Many tried a gridlike approach by trying to explore the length of the maze, turn, and then quickly turn again to walk a parallel path. Others tried to explore the entire perimeter and then explore the inner corridors of the maze. When reporting how subjects memorized object locations, reports were split between egocentric and allocentric representations, with egocentric being more common overall. However, there was no difference in the split between conditions. Many subjects had to be given examples of strategies in order to explain what strategy they used. Interestingly, some subjects used the feather not as an indicator of where an object was located, but to aid in knowing the distance they had walked. Subjects indicated that for long distances they would expect multiple interventions, and tried to stay consistent in the number of interventions among similar paths.

5.6 Discussion

This chapter looked at how room size affected spatial cognition and the acquisition of spatial knowledge using the resetting locomotion method. We compared resetting in three differently sized spaces to WiP using survey knowledge as our principal metric. Subjects showed improved directional survey knowledge absent any resets, and in general improved in performance as the space grew. These results also showed that navigational performance can also be predicted based on individual differences in ability. Subjects with better mental rotation abilities or stronger sense of direction performed markedly better at this task. The findings of this experiment grant us a better understanding of how resetting affects one's ability to navigate, particularly showing that the cognitive cost [149] extends to directional survey knowledge. It is important to consider these findings when deciding on a locomotion method. If one's space is sufficiently large resetting is the more enticing method as it involves real walking. Alternatively WiP is the preferable locomotion method if learning the environment is the key goal.

5.6.1 Future Work

Regardless of the locomotion method chosen, these findings confirm [39] the importance of considering individual differences when employing a navigation task. Future research into navigation must continue to consider these individual differences and explore new measures to reduce the variation in performance.

Chapter 6

Conclusion

This dissertation compares and contrasts standing space locomotion (WiP) to room scale locomotion (resetting) in terms of how each method affects spatial cognition and is affected by individual differences in navigation ability. Navigation in large virtual environments is difficult in part due to a lack of available body-based cues of many standard locomotion methods [31, 95]. Teleporting [31] and joystick locomotion remove all translational cues and leave only rotational cues. The techniques used in this dissertation enable as many of these body-based cues as possible in their given space.

Locomotion is important in virtual reality systems to allow for the exploration of large (beyond visible space) virtual environments. When moving through large virtual environments, it is easy to become disoriented. Disorientation may defeat the learning or training purposes of many application using large immersive virtual environments. To remedy this problem, we identified some issues associated with locomotion methods. Our criteria for selecting locomotion methods was the reliability and presence of the body-based cues and information they provide.

A number of commodity level virtual reality systems have been developed recently, and these systems can largely be divided into two classes. One class, which consists of HMDs such as the HTC Vive and the Oculus Rift, offers room scale tracking on the order of 4m by 4m. Tracking spaces of this size allow for the use of overt reorienting methods [130] such as resetting [149] or distractors [93]. Other systems, such as the Oculus Go and Samsung Gear VR, use inertial tracking systems and therefore have only standing scale tracking space available. These systems having only rotational tracking means that a technique that doesn't require that positional tracking is needed. There are a number of controller or joystick based techniques that could be used, however, as shown by Riecke et al. [105], the lack of body-based cues (e.g., translational cues in complex environments) can be detrimental to spatial cognition and navigation ability. For this reason we chose to develop and investigate a WiP locomotion method as well.

In order to allow unbounded exploration there are three classes of manipulations or locomotion methods. One can manipulate curvature, rotational information, or positional information. Chapter 2 examined the literature on how these manipulations were manipulated and presents the various evaluations and refinements of each of these methods. We focus on these three classes as they roughly correspond to the three space size requirements, though there is some overlap. This chapter presents the importance of body-based information in reinforcing spatial cognition and discussed how to provide that information.

The latter half of Chapter 2 contains a primer on spatial cognition as defined by the "Spatial Cognitive Microgenesis" framework. We then discussed individual differences in strategy, skill, and experience and how those differences affect general performance in navigation. This dissertation looked at these individual differences to try and predict navigation performance of the locomotion methods presented here to tailor locomotion to an individual.

Chapter 3 of this dissertation examined different ways in which subjects could locomote freely through a large virtual environment that is larger than the typical 4x4 m^2 space available in commodity VR. As room scale VR continues its proliferation, understanding the effects of the various room scale locomotion methods on navigation becomes increasingly important. This importance stems from the need to learn virtual environment layouts. Applications such as architecture and military can benefit from a user having a greater understanding of the virtual space. In this research we select only those techniques that allow for self-generated locomotion that has full body-based information. Self-generated locomotion is very effective in supporting the development of a cognitive map and learning the locations of landmarks within an environment.

The work reported in Chapter 3 showed that the room scale techniques of resetting and distractors are equivalent in their ability to support spatial cognition in terms of path integration and route knowledge. In Experiment 1 we implemented a path integration task that was designed to load one's cognitive resources. This task allowed us to compare resetting, distractors, and redirected walking, all of which have been shown to increase cognitive demand [149]. Experiment 1 showed the cognitive cost of a reorientation in either of the reorienting methods was equivalent.

Experiment 2 continued to compare resetting and distractors using route knowledge as a measurement tool. Subjects were tasked with learning a simple city-like environment as they visited various storefronts. This work showed that while redirected walking allowed for the same level of knowledge in our path integration task, it was not nearly as comfortable. The small curvature radius in the implementation mandated by our lab space caused sickness and was quickly noticed. This conclusion is supported by prior work indicating that our lab is too small to reliably allow unbounded redirected walking. We also looked at time requirements of resetting and distractors and found that using distractors has a much higher time cost. We then concluded that for our purposes, resetting was a suitable selection for a room scale locomotion method.

In Chapter 4 we developed and evaluated a WiP technique which could be used in standing scale spaces. To first develop a WiP technique we employed two different methods of tracking and compared them to that of resetting. The first used an external device, the Kinect, which used skeletal tracking for step detection. The second detected steps using an IMU and, within the accelerometer signal, identified a successive peak and valley, indicating a step had likely occurred. We then evaluated all three locomotion methods.

Our evaluation criteria was that of the "Spatial Cognitive Microgenesis" framework for testing the acquisition of spatial knowledge. This framework had two metrics available for testing: distance and direction. In Experiment 3, direction showed very little difference between conditions but high variation between subjects. Distance showed large differences between conditions; subjects indicated almost 3 times the true distance in the WiP conditions.

Experiment 4 was our follow-up work intended to determine the cause of the vast over-walking from Experiment 3. The key idea here involved splitting the locomotion methods used for learning and testing. This split gave us four conditions which allowed us to make direct comparisons between learning methods and testing methods. From this study we determined that there is difficulty expressing an intended distance with WiP. The exact cause of that difficulty is unknown and should be investigated in future work. The most interesting result we found is that most subjects could acquire spatial knowledge just as accurately with WiP as they could resetting.

We also designed this experiment as a first step in explaining the large variation between subjects' individual performances. As suggested by the literature, we investigated individual differences in three areas: general navigation ability, small scale spatial ability, and general intelligence.

Chapter 5 then looked at comparing WiP to resetting when the tracked space was around 5x5 m^2. Since WiP requires much less room we completed the work in Chapter 5 to determine the space required for resetting to continue to provide information on par with WiP. Resetting comes at a cognitive cost [149] meaning an increase in the number of resets should lead to a decrease in the amount of survey information one can acquire. As we shrink the available space, the number of resets will necessarily increase, and this work showed that this shrinking led to degraded performance in our navigation task. This result could mean that there is a room size large enough so that resetting will outperform WiP and make the cognitive cost negligible. This chapter also showed us that SBSOD is an important factor in explaining large variations in individuals' navigation performance.

Every living room or gaming room is different and every person has a different amount of available space. With the proliferation of VR technology, particularly mobile technology, the effects that one's available space has on user experience has become an important consideration and must be researched. This work used two measures of user experience: spatial awareness and simulator sickness. We presented and evaluated five locomotion methods: resetting, distractors, redirected walking, WiP using an IMU, and WiP using the Kinect. Those locomotion methods were evaluated against those two criteria. We selected methods that leave the hands free for other potential interaction and provide in part both vestibular and proprioceptive cues.

Specifically in this work we assessed three real walking methods of navigating in large scale virtual environments using a medium sized physical space such as those available in a typical living room or gaming room. This assessment gave us a method with which we can compare standing scale locomotion methods. We developed a standing scale locomotion method, WiP, which allows spatial knowledge to be acquired. This WiP method was compared to resetting and differences in spatial knowledge acquired were found. We have taken the first steps in determining how individual differences in ability effect one's ability to acquire spatial knowledge with resetting and/or WiP. Finally we have examined the effect room size has on resetting performance. These results allowed us to make a number of suggestions on which locomotion method to employ based on an individual's ability, the virtual task, or physical space available.

6.1 Future Directions

This work is just the beginning stages of determining an appropriate method of walking based on the factors such as individual differences and room size. Secondly, this dissertation looked only at methods that involve real walking or simulate real walking. Techniques that don't require real walking (e.g., teleporting and controller-based) have been developed and should be studied within the Chrastil and Warren [17] framework used in this dissertation. Lastly, the WiP technique used here is simple in nature and makes a number of assumptions that limit the individualization possible.

6.1.1 Individual Differences

This work is a first analysis of the relationship between individual differences (e.g., skill and strategy) and navigation performance in resetting and WiP locomotion methods. Small scale spatial abilities, as measured by the MRT, were significantly correlated with performance in both methods of walking. Subjects with strong spatial abilities, then, could be given a locomotion method that is more cognitively taxing. Future work would then be required to determine if subjects could then still acquire sufficient spatial knowledge.

In this work we found that for those subjects using WiP, small scale spatial abilities were important only when subjects also had high SBSOD scores as well. Future work should investigate the cause and effects of this trend and may be important in designing a VR system individualized to a user and their space. This work also confirmed several correlations between measures of individual differences (e.g., CORSI and MRT). Even with a number of tests considered in this dissertation large variations still exist and further investigation into which other measures can be used is important.

6.1.2 Larger and Irregular Spaces

We know from this work as well as prior work that there is a cognitive cost associated with resetting. The number of resets that occur in a given path or during a time period is largely dependent on the size and shape of the tracked space. While this work has begun to show trends of how resetting performance depends on the number of resets and the size of the tracked space, it will be important in future work to extend this for larger and irregular spaces.

HTC intends to create tracking spaces as large as 10x10 meters[138], much larger than those investigated in this work. VR cafes and those with the available space will require locomotion methods which are appropriate for larger spaces. Work from Simeone et al. [119] and Dong et al. [25] has dealt with how to keep people away from obstacles in the space. This line of work stems from the need to plan around irregular spaces and allow those with obstacles in their tracked space to still walk through a virtual environment. Follow up work should be done to determine if irregular spaces can still be used with resetting locomotion methods and how to best implement that resetting to minimize the number of interventions required.

6.1.3 Other Locomotion Affordances

This dissertation did not consider locomotion affordances other than walking interfaces. There are a number of other ways to move around a virtual environment as discussed in Chapter 2. However, those methods do not involve or simulate real walking. As newer locomotion methods are developed, it becomes vital to understand the relationship between navigation performance and the physical involvement required. While walking has been shown superior to methods such as joystick locomotion, using methods which involve arm motion or stimulate the vestibular system may prove to be powerful replacements of real walking.

6.1.4 Walking in Place Improvements

Individual differences in skill, strategy, and experience are one of the key aspects of this dissertation. However, there are a number of other differences that can be considered. When walking in our simple WiP system subjects are expected to adapt to the system, not the other way around. This can lead to the need for increased training to adapt one's own walking motion to the system as well as large lag when attempting to start or stop. Future work should examine the different ways subjects walk and tailor the method to their individual acceleration curves and tendencies.

Current inertial measurement units (IMUs) have six streams of data (three acceleration and three orientation streams) and can report their data as quickly as 90 times per second. These streams give us a tremendous amount of data which makes it difficult to understand how starting and stopping appear differently than normal turning or walking. Deep learning is a typical approach to analyzing a large amount of data and is applicable in this area. Precedent comes from Usoh et al. [136] who used this approach previously when tracking the up axis of head movement of a subject to identify walking.

WiP in general and the machine learning approach mentioned above specifically can benefit from the understanding of the mechanics of natural walking. Feasel et al. [28] use an approach that takes into account gait analysis of real walking to reduce the latency associated with starting and stopping in WiP. Their WiP methods use logs of leg tracking, but may be extendable to head acceleration tracking with the inclusion of motion transfer mechanics. The naive approach to deep learning for WiP will require the recognition of an entire step. By understanding the difference in signals between a step occurring and a stop in stepping occurring we can further reduce stopping latency in WiP. Taking all of these improvements together could make WiP a strong competitor of resetting.

Glossary

Acquisition of spatial knowledge The process by which people acquire and encode spatial knowledge for later recall.

Allocentric Allocentric spatial knowledge is spatial knowledge between two external objects or landmarks.

Body-based cues Salient events generated by the body that allow for the acquisition of body-based information.

Body-based information Any information, such as directional information or positional information, gained from body-based cues.

Direction cues Salient body-based or environment-based events which allow for the acquisition of directional information.

Direction information Information indicating absolute direction from yourself (egocentric) of with respect to another object (allocentric).

Egocentric Placing oneself at the center. Egocentric spatial knowledge is spatial knowledge with respect to oneself.

Graph Knowledge An additional stage of knowledge proposed by Chrastil and Warren [18]. The ability to combine multiple routes to create a novel path without survey knowledge.

Interpoint distance cues Salient body-based or environment-based events which allow for the acquisition of interpoint distance information.

Interpoint distance information Information indicating how far away an object is from yourself (egocentric) or two objects are from one another (allocentric).

Joystick navigation A locomotion method that allows virtual translation and/or rotation using a handheld device without changing real world translation and/or rotation.

Landmark Knowledge The first stage of Spatial Cognitive Microgenesis [118]. Recognition and recall of landmarks in the environment.

Locomotion methods A control system that allows users of virtual reality to control their movement through a virtual environment.

Navigation Knowledge of where one is, where one is going, and how one will get there.

Orientation Tracking A computer system which measures the orientation of the user's head. This type of tracking typically uses gyroscopes or magnetic compasses.

Positional Tracking A computer system which measures the change in position of the users head. This type of tracking typically uses cameras to detect infrared markers.

Proprioceptive system An internal system which allows a person to know the configuration and exertion of their body at all times.

Redirected walking A locomotion method that continuously manipulates the rotational and translational gain to steer users away from obstacles/boundaries. The constant rotational gain causes a difference in curvature between the real and virtual worlds.

Reorientation (ROT) A class of locomotion methods which allow users to recover from colliding with an obstacle or the physical limits. At the conclusion of an ROT, users are no longer facing an obstacle in the real world, but have retained their virtual orientation.

Resetting A locomotion method that allows users to recover from colliding with an obstacle or the physical limits of the tracking system. The rotational gain is manipulated and users are instructed to turn around, thus reorienting users away from the physical obstacle while retaining their virtual orientation.

Route Knowledge The second stage of Spatial Cognitive Microgenesis [118]. The ability to navigate along a route and know when and which direction to turn.

Spatial Cognitive Microgenesis (SCM) The dominant framework on how spatial knowledge is acquired. The framework consists of three stages of knowledge: landmark, route, and survey. It is hierarchical, so each stage builds upon the last.

Spatial working memory The subcomponent of working memory responsible for using spatial information such as configuration.

Suppression task Related to working memory, a suppression task tries to tax one subcomponent of memory to force all acquisition of spatial knowledge to be done by the remaining systems.

Survey Knowledge The final stage of Spatial Cognitive Microgenesis [118]. An understanding of the orientation and distance between objects in the environment.

Teleportation A locomotion method that instantaneously moves a user to a new desired virtual location without changing their orientation or real world position.

Tracking Any computer system that transmits real world position and/or orientation to the virtual environment control system.

Translational gain A locomotion method that increases the distance traveled each step by some scaling factor, thus for much more virtual space to be covered in the same physical space.

Verbal working memory The subcomponent of working memory responsible for using verbal information such as descriptors or instructions.

Vestibular system An internal system which measures the acceleration of the inner ear to determine body-based motion.

Visual working memory The subcomponent of working memory responsible for using visual information such as color or shape.

Walking in Place (WiP) A locomotion method that allows users to simulate walking and move through a virtual environment without physically translating. Motion is determined by detecting in-place steps.

Way-finding The process by which one finds their way around an environment.

Working memory A measure of how much information can be acquired and processed by the brain at any one time. There are three subcomponents: verbal, visual, and spatial.

BIBLIOGRAPHY

[1] R. C. Allen and M. J. Singer. Landmark direction and distance estimation in large scale. In *Proceedings of the Human Factors and Ergonomics Society 41st Annual Meeting*, pages 1213–1217, 1997.

[2] Aiden EGF Arnold, Ford Burles, Taisya Krivoruchko, Irene Liu, Colin D Rey, Richard M Levy, and Giuseppe Iaria. Cognitive mapping in humans and its relationship to other orientation skills. *Experimental Brain Research*, 224(3):359–372, 2013.

[3] Daniel H Ashmead, Everett W Hill, and Charles R Talor. Obstacle perception by congenitally blind children. *Perception & Psychophysics*, 46(5):425–433, 1989.

[4] Mahdi Azmandian, Timofey Grechkin, and Evan Suma Rosenberg. An evaluation of strategies for two-user redirected walking in shared physical spaces. In *Virtual Reality (VR), 2017 IEEE*, pages 91–98. IEEE, 2017.

[5] Olesya Blazhenkova and Maria Kozhevnikov. The new object-spatial-verbal cognitive style model: Theory and measurement. *Applied Cognitive Psychology*, 23(5):638–663, 2009.

[6] Doug A Bowman, Elizabeth T Davis, Larry F Hodges, and Albert N Badre. Maintaining spatial orientation during travel in an immersive virtual environment. *Presence*, 8(6):618–631, 1999.

[7] F Bremmer and M Lappe. The use of optical velocities for distance discrimination and reproduction during visually simulated self motion. *Experimental Brain Research*, 127(1): 33–42, 1999.

[8] Gerd Bruder, Frank Steinicke, and Klaus H. Hinrichs. Arch-explore: A natural user interface for immersive architectural walkthroughs. *3D User Interfaces*, 0:75–82, 2009. doi: http://doi.ieeecomputersociety.org/10.1109/3DUI.2009.4811208.

[9] Gerd Bruder, Paul Lubas, and Frank Steinicke. Cognitive resource demands of redirected walking. *IEEE Transactions on Visualization and Computer Graphics*, 21(4):539–544, 2015.

[10] Axel Buchner and Petra Jansen-Osmann. Is route learning more than serial learning? *Spatial Cognition & Computation*, 8(4):289–305, 2008.

[11] Heather Burte and Daniel R Montello. How sense-of-direction and learning intentionality relate to spatial knowledge acquisition in the environment. *Cognitive Research: Principles and Implications*, 2(1):18, 2017.

[12] Xiaoqian J Chai and Lucia F Jacobs. Sex differences in directional cue use in a virtual landscape. *Behavioral Neuroscience*, 123(2):276, 2009.

[13] Sarah S Chance, Florence Gaunet, Andrew C Beall, and Jack M Loomis. Locomotion mode affects updating of objects encountered during travel: The contribution of vestibular and proprioceptive inputs to path integration. *Presence*, 7(2):168–178, 1998.

[14] Elizabeth Chrastil and William Warren. Contributions of attention and decision-making to spatial learning. *Journal of Vision*, 12(9):203–203, 2012.

[15] Elizabeth R Chrastil. Neural evidence supports a novel framework for spatial navigation. *Psychonomic Bulletin & Review*, 20(2):208–227, 2013.

[16] Elizabeth R Chrastil and William H Warren. Active and passive contributions to spatial learning. *Psychonomic Bulletin & Review*, 19(1):1–23, 2012.

[17] Elizabeth R Chrastil and William H Warren. Active and passive spatial learning in human navigation: Acquisition of survey knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(5):1520, 2013.

[18] Elizabeth R Chrastil and William H Warren. Active and passive spatial learning in human

navigation: acquisition of graph knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 41(4):1162, 2015.

[19] ER Chrastil and WH Warren. Spatial navigation: Why is active exploration better than passive exploration? In *Poster, Vision Sciences Society Annual Meeting*, 2011.

[20] Jacob Cohen. Statistical power analysis. *Current directions in psychological Science*, 1(3): 98–101, 1992.

[21] Noah Coomer, Sadler Bullard, William Clinton, and Betsy Williams-Sanders. Evaluating the effects of four vr locomotion methods: Joystick, arm-cycling, point-tugging, and teleporting. In *Proceedings of the 15th ACM Symposium on Applied Perception*, SAP '18, pages 7:1–7:8, New York, NY, USA, 2018. ACM. ISBN 978-1-4503-5894-1. doi: 10.1145/3225153.3225175. URL http://doi.acm.org/10.1145/3225153.3225175.

[22] P Corsi. Memory and the medial temporal region of the brain. *Unpublished Doctoral Dissertation), McGill University, Montreal, QB*, 1972.

[23] Sarah H. Creem-Regehr, Jeanine K. Stefanucci, William B. Thompson, Nathan Nash, and Michael McCardell. Egocentric distance perception in the oculus rift (dk2). In *Proceedings of the ACM SIGGRAPH Symposium on Applied Perception*, SAP '15, pages 47–50, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3812-7. doi: 10.1145/2804408.2804422. URL http://doi.acm.org/10.1145/2804408.2804422.

[24] Guy Deutscher. Does your language shape how you think? https://www.nytimes.com/2010/08/29/magazine/29language-t.html?pagewanted=all&_r=1&, 2010. Accessed: 2018-5-29.

[25] Zhi-Chao Dong, Xiao-Ming Fu, Chi Zhang, Kang Wu, and Ligang Liu. Smooth assembled mappings for large-scale real walking. *ACM Transactions on Graphics (TOG)*, 36(6):211, 2017.

[26] Ruth B Ekstrom, Diran Dermen, and Harry Horace Harman. *Manual for kit of factor-referenced cognitive tests*, volume 102. Educational testing service Princeton, NJ, 1976.

[27] David Engel, Cristóbal Curio, Lili Tcheang, Betty Mohler, and Heinrich H. Bülthoff. A psychophysically calibrated controller for navigating through large environments in a limited free-walking space. In *VRST '08: Proceedings of the 2008 ACM symposium on Virtual reality software and technology*, pages 157–164, New York, NY, USA, 2008. ACM. ISBN 978-1-59593-951-7. doi: http://doi.acm.org/10.1145/1450579.1450612.

[28] Jeff Feasel, Mary C Whitton, and Jeremy D Wendt. Llcm-wip: Low-latency, continuous-motion walking-in-place. In *3D User Interfaces, 2008. 3DUI 2008. IEEE Symposium on*, pages 97–104. IEEE, 2008.

[29] Alexa W Fields and Amy L Shelton. Individual skill differences and large-scale environmental learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32(3):506, 2006.

[30] Julia Frankenstein, Betty J Mohler, Heinrich H Bülthoff, and Tobias Meilinger. Is the map in our head oriented north? *Psychological Science*, 23(2):120–125, 2012.

[31] S. Freitag, D. Rausch, and T. Kuhlen. Reorientation in virtual environments using interactive portals. In *2014 IEEE Symposium on 3D User Interfaces (3DUI)*, pages 119–122, March 2014. doi: 10.1109/3DUI.2014.6798852.

[32] Michael N. Geuss, Jeanine K. Stefanucci, Sarah H. Creem-Regehr, William B. Thompson, and Betty J. Mohler. Effect of display technology on perceived scale of space. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 57(7):1235–1247, 2015.

[33] Google. Google cardboard: Official vr headsets. https://store.google.com/product/google_cardboard, 2017. Accessed: 2017-5-25.

[34] Timofey Grechkin, Jerald Thomas, Mahdi Azmandian, Mark Bolas, and Evan Suma. Revisiting detection thresholds for redirected walking: Combining translation and curvature gains. In *Proceedings of the ACM Symposium on Applied Perception*, SAP '16, pages 113–120, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-4383-1. doi: 10.1145/2931002.2931018. URL http://doi.acm.org/10.1145/2931002.2931018.

[35] S. Hanson, R. A. Paris, H. A. Adams, and B. Bodenheimer. Improving walking in place methods with individualization and deep networks. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pages 367–376, March 2019. doi: 10.1109/VR.2019.8797751.

[36] Abraham M Hashemian and Bernhard E Riecke. Leaning-based 360 interfaces: investigating virtual reality navigation interfaces with leaning-based-translation and full-rotation. In *International Conference on Virtual, Augmented and Mixed Reality*, pages 15–32. Springer, 2017.

[37] Denise Head and Marlisa Isom. Age effects on wayfinding and route learning skills. *Behavioural Brain Research*, 209(1):49–58, 2010.

[38] Mary Hegarty, Anthony E Richardson, Daniel R Montello, Kristin Lovelace, and Ilavanil Subbiah. Development of a self-report measure of environmental spatial ability. *Intelligence*, 30(5):425–447, 2002.

[39] Mary Hegarty, Daniel R Montello, Anthony E Richardson, Toru Ishikawa, and Kristin Lovelace. Spatial abilities at different scales: Individual differences in aptitude-test performance and spatial-layout learning. *Intelligence*, 34(2):151–176, 2006.

[40] Mary Hegarty, Raymond D Crookes, Drew Dara-Abrams, and Thomas F Shipley. Do all science disciplines rely on spatial abilities? preliminary evidence from self-report questionnaires. In *International Conference on Spatial Cognition*, pages 85–94. Springer, 2010.

[41] E. Hodgson, E. Bachmann, and T. Thrash. Performance of redirected walking algorithms in a constrained virtual world. *Visualization and Computer Graphics, IEEE Transactions on*, 20(4):579–587, 2014.

[42] Eric Hodgson and Eric Bachmann. Comparing four approaches to generalized redirected walking: Simulation and live user data. *Visualization and Computer Graphics, IEEE Transactions on*, 19(4):634–643, 2013.

[43] Eric Hodgson, Eric Bachmann, and David Waller. Redirected walking to explore virtual environments: Assessing the potential for spatial interference. *ACM Transactions on Applied Perception (TAP)*, 8(4):22, 2011.

[44] John M. Hollerbach, Damaso Checcacci, Haruo Noma, Yasuyuki Yanagida, and Nobuji Tetsutani. Simulating side slopes on locomotion interfaces using torso forces. *International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 0:91, 2003. doi: http://doi.ieeecomputersociety.org/10.1109/HAPTIC.2003.1191241.

[45] Corinne A Holmes, Daniele Nardi, Nora S Newcombe, and Steven M Weisberg. Children's use of slope to guide navigation: Sex differences relate to spontaneous slope perception. *Spatial Cognition & Computation*, 15(3):170–185, 2015.

[46] Giuseppe Iaria, Liana Palermo, Giorgia Committeri, and Jason JS Barton. Age differences in the formation and use of cognitive maps. *Behavioural Brain Research*, 196(2):187–191, 2009.

[47] Victoria Interrante, Brian Ries, and Lee Anderson. Seven league boots: A new metaphor for augmented locomotion through moderately large scale immersive virtual environments. In *IEEE Symposium on 3D User Interfaces*, pages 167–170, 2007.

[48] Toru Ishikawa and Daniel R Montello. Spatial knowledge acquisition from direct experience in the environment: Individual differences in the development of metric knowledge and the integration of separately learned places. *Cognitive Psychology*, 52(2):93–129, 2006.

[49] Hiroo Iwata. Walking about virtual environments on an infinite floor. In *Virtual Reality, 1999. Proceedings., IEEE*, pages 286–293. IEEE, 1999.

[50] Harold Jeffreys. *Theory of Probability*. Oxford Classic Texts in the Physical Sciences. Oxford University Press, 3rd edition, November 1998.

[51] Robert E. Kass and Adrian E. Raftery. Bayes factors. *Journal of the American Statistical Association*, 90(430):773–795, 1995. doi: 10.1080/01621459.1995.10476572. URL http://amstat.tandfonline.com/doi/abs/10.1080/01621459.1995.10476572.

[52] Yoshinobu Kato and Yoshiaki Takeuchi. Individual differences in wayfinding strategies. *Journal of Environmental Psychology*, 23(2):171–188, 2003.

[53] M. Kearns, W. Warren, A. Duchon, and M. Tarr. Path integration from optic flow and body senses in a homing task. *Perception*, 31:349–374, 2002.

[54] KinectVR. Kinect vr: Room scale tracking for mobile vr. http://kinectvr.com, 2017. Accessed: 2017-5-18.

[55] Alexandra Kitson, Bernhard E Riecke, Abraham M Hashemian, and Carman Neustaedter. Navichair: Evaluating an embodied interface using a pointing task to navigate virtual reality. In *Proceedings of the 3rd ACM Symposium on Spatial User Interaction*, pages 123–126. ACM, 2015.

[56] Alexandra Kitson, Abraham M Hashemian, Ekaterina R Stepanova, Ernst Kruijff, and Bernhard E Riecke. Comparing leaning-based motion cueing interfaces for virtual reality locomotion. In *3D User Interfaces (3DUI), 2017 IEEE Symposium on*, pages 73–82. IEEE, 2017.

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

[58] Alexander Klippel, Thora Tenbrink, and Daniel R Montello. The role of structure and function in the conceptualization of directions. *Motion Encoding in Language and Space*, 2010.

[59] David JM Kraemer, Victor R Schinazi, Philip B Cawkwell, Anand Tekriwal, Russell A Epstein, and Sharon L Thompson-Schill. Verbalizing, visualizing, and navigating: The effect of strategies on encoding a large-scale virtual environment. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 43(4):611, 2017.

[60] Ernst Kruijff, Alexander Marquardt, Christina Trepkowski, Robert W Lindeman, Andre Hinkenjann, Jens Maiero, and Bernhard E Riecke. On your feet!: Enhancing vection in leaning-based interfaces through multisensory stimuli. In *Proceedings of the 2016 Symposium on Spatial User Interaction*, pages 149–158. ACM, 2016.

[61] Scott A. Kuhl. Recalibration of rotational locomotion in immersive virtual environments. In *APGV 2004*, pages 23–26, August 2004.

[62] Scott A. Kuhl, Sarah H. Creem-Regehr, and William B. Thompson. Recalibration of rotational locomotion in immersive virtual environments. *ACM Transactions on Applied Perception*, 5(3):1–11, 2008. ISSN 1544-3558. doi: http://doi.acm.org/10.1145/1402236.1402241.

[63] Enia Labate, Francesca Pazzaglia, and Mary Hegarty. What working memory subcomponents are needed in the acquisition of survey knowledge? evidence from direction estimation and shortcut tasks. *Journal of Environmental Psychology*, 37:73–79, 2014.

[64] Eike Langbehn, Paul Lubos, Gerd Bruder, and Frank Steinicke. Application of redirected walking in room-scale vr. In *Virtual Reality (VR), 2017 IEEE*, pages 449–450. IEEE, 2017.

[65] Carol A Lawton. Gender differences in way-finding strategies: Relationship to spatial ability and spatial anxiety. *Sex Roles*, 30(11-12):765–779, 1994.

[66] Carol A Lawton. Strategies for indoor wayfinding: The role of orientation. *Journal of environmental psychology*, 16(2):137–145, 1996.

[67] Carol A Lawton. Gender and regional differences in spatial referents used in direction giving. *Sex Roles*, 44(5-6):321–337, 2001.

[68] Carol A Lawton and Kevin A Morrin. Gender differences in pointing accuracy in computer-simulated 3d mazes. *Sex Roles*, 40(1-2):73–92, 1999.

[69] Robert Leeb, Claudia Keinrath, Doron Friedman, Christoph Guger, Reinhold Scherer, Christa Neuper, Maia Garau, Angus Antley, Anthony Steed, Mel Slater, and Gert Pfurtscheller. Walking by thinking: The brainwaves are crucial, not the muscles! *Presence: Teleoperators and Virtual Environments*, 15(5):500–514, 2006. doi: 10.1162/pres.15.5.500. URL http://dx.doi.org/10.1162/pres.15.5.500.

[70] Jack M Loomis, Roberta L Klatzky, Reginald G Golledge, and John W Philbeck. Human navigation by path integration. *Wayfinding behavior: Cognitive mapping and other spatial processes*, pages 125–151, 1999.

[71] Jeremy R Manning, Timothy F Lew, Ningcheng Li, Robert Sekuler, and Michael J Kahana. Magellan: A cognitive map–based model of human wayfinding. *Journal of Experimental Psychology: General*, 143(3):1314–1330, June 2014.

[72] T. P. McNamara. How are the locations of objects in the environment represented in memory? In C. Freksa, W. Brauer, C. Habel, and K. F. Wender, editors, *Spatial cognition III: Routes and navigation, human memory and learning, spatial representation and spatial learning*, pages 174–191. Springer Berlin / Heidelberg, 2003.

[73] Frank Meijer, Branko L Geudeke, and Egon L Van den Broek. Navigating through virtual environments: Visual realism improves spatial cognition. *CyberPsychology & Behavior*, 12 (5):517–521, 2009.

[74] Tobias Meilinger, Markus Knauff, and Heinrich H Bülthoff. Working memory in wayfindinga dual task experiment in a virtual city. *Cognitive Science*, 32(4):755–770, 2008.

[75] Tobias Meilinger, Julia Frankenstein, and Heinrich H Bülthoff. Learning to navigate: experience versus maps. *Cognition*, 129(1):24–30, 2013.

[76] Scott D Moffat. Aging and spatial navigation: what do we know and where do we go? *Neuropsychology Review*, 19(4):478, 2009.

[77] Daniel R Montello. A new framework for understanding the acquisition of spatial knowledge in large-scale environments. *Spatial and Remporal Reasoning in Geographic Information Systems*, pages 143–154, 1998.

[78] W Mou and T P McNamara. Intrinsic frames of reference in spatial memory. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 28(1):162–170, 2002.

[79] W Mou, T P McNamara, C M Valiquette, and B Rump. Allocentric and egocentric updating of spatial memories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30(1):142–157, 2004.

[80] Weimin Mou, Timothy P McNamara, Björn Rump, and Chengli Xiao. Roles of egocentric and allocentric spatial representations in locomotion and reorientation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32(6):1274, 2006.

[81] Weimin Mou, Yanli Fan, Timothy P McNamara, and Charles B Owen. Intrinsic frames of reference and egocentric viewpoints in scene recognition. *Cognition*, 106(2):750–769, 2008.

[82] Stefan Münzer and Christoph Stahl. Learning routes from visualizations for indoor wayfinding: Presentation modes and individual differences. *Spatial Cognition & Computation*, 11 (4):281–312, 2011.

[83] Stefan Münzer, Hubert D Zimmer, Maximilian Schwalm, Jörg Baus, and Ilhan Aslan. Computer-assisted navigation and the acquisition of route and survey knowledge. *Journal of Environmental Psychology*, 26(4):300–308, 2006.

[84] Daniele Nardi, Corinne A Holmes, Nora S Newcombe, and Steven M Weisberg. Sex differences and errors in the use of terrain slope for navigation. *Cognitive Processing*, 16(1): 323–326, 2015.

[85] C.T. Neth, J.L. Souman, D. Engel, U. Kloos, H.H. Bu? andlthoff, and B.J. Moh̃ler. Velocity-dependent dynamic curvature gain for redirected walking. In *Virtual Reality Conference (VR), 2011 IEEE*, pages 151 –158, March 2011. doi: 10.1109/VR.2011.5759454.

[86] Nora Newcombe and Lynn S Liben. Barrier effects in the cognitive maps of children and adults. *Journal of Experimental Child Psychology*, 34(1):46–58, 1982.

[87] Ehren L Newman, Jeremy B Caplan, Matthew P Kirschen, Igor O Korolev, Robert Sekuler, and Michael J Kahana. Learning your way around town: How virtual taxicab drivers learn to use both layout and landmark information. *Cognition*, 104(2):231–253, 2007.

[88] Norbert Nitzsche, Uwe Hanebeck, and Gunther Schmidt. Motion compression for telepresent walking in large target environments. *Presence*, 13(1):44–60, 2004.

[89] Lace M Padilla, Sarah H Creem-Regehr, Jeanine K Stefanucci, and Elizabeth A Cashdan. Sex differences in virtual navigation influenced by scale and navigation experience. *Psychonomic Bulletin & Review*, pages 1–9, 2016.

[90] Richard Paris, Miti V. Joshi, Qiliang He, Gayathri Narasimham, Timothy P. McNamara, and Bobby Bodenheimer. Acquisition of survey knowledge using walking in place and resetting methods in immersive virtual environments. In *Proceedings of the ACM Symposium on Applied Perception*, SAP '17, New York, NY, USA, 2017. ACM.

[91] Richard Paris, Joshua Klag, Priya Rajan, Lauren Buck, Timothy P. McNamara, and Bobby Bodenheimer. How video game locomotion methods affect navigation in virtual environments. In *ACM Symposium on Applied Perception 2019*, SAP '19, pages 12:1–12:7, New York, NY, USA, 2019. ACM. ISBN 978-1-4503-6890-2. doi: 10.1145/3343036.3343131. URL http://doi.acm.org/10.1145/3343036.3343131.

[92] T. C. Peck, H. Fuchs, and M. C. Whitton. Improved redirection with distractors: A large-scale-real-walking locomotion interface and its effect on navigation in virtual environments. In *IEEE Virtual Reality Conference*, March 2010.

[93] Tabitha C. Peck, Henry Fuchs, and Mary C. Whitton. Evaluation of reorientation techniques and distractors for walking in large virtual environments. *IEEE Transactions on Visualization and Computer Graphics*, 15:383–394, 2009. ISSN 1077-2626. doi: http://doi.ieeecomputersociety.org/10.1109/TVCG.2008.191.

[94] T.C. Peck, H. Fuchs, and M.C. Whitton. The design and evaluation of a large-scale real-walking locomotion interface. *Visualization and Computer Graphics, IEEE Transactions on*, 18(7):1053–1067, 2012. ISSN 1077-2626. doi: 10.1109/TVCG.2011.289.

[95] Peng Peng, Bernhard E. Riecke, Betsy Williams, Timothy P. McNamara, and Bobby Bodenheimer. Navigation modes in virtual environments: walking vs. joystick. Poster at the Symposium on Applied Perception in Graphics and Visualization, August 2008.

[96] P. Péruch, L. Belingard, and C. Thinus-Blanc. Transfer of spatial knowledge from virtual to real environments. In C. Freska, editor, *Spatial Cognition II*. Springer, Berlin, 2000.

[97] H L Pick, J J Rieser, D Wagner, and A E Garing. The recalibration of rotational locomotion. *Journal of Experimental Psychology: Human Perception and Performance*, 25(5):1179–1188, 1999.

[98] Shairf Razzaque. *Redirected Walking*. PhD thesis, University of North Carolina, Chapel Hill, 2005.

[99] Sharif Razzaque, Zachariah Kohn, and Mary C Whitton. Redirected walking. *Eurographics Short Presentation*, 2001.

[100] Lisa Rebenitsch and Charles Owen. Individual variation in susceptibility to cybersickness. In *Proceedings of the 27th annual ACM symposium on User interface software and technology*, pages 309–317. ACM, 2014.

[101] Jan D Restat, Sibylle D Steck, Horst F Mochnatzki, and Hanspeter A Mallot. Geographical slant facilitates navigation and orientation in virtual environments. *Perception*, 33(6):667–687, 2004.

[102] B. E. Riecke, D. W. Cunningham, and H. H. Bülthoff. Spatial updating in virtual reality: the sufficiency of visual information. *Psychological Research*, 71(3):298–313, 2006.

[103] Bernhard E. Riecke, Markus Von Der Heyde, and Heinrich H. Bülthoff. Visual cues can be sufficient for triggering automatic, reflexlike spatial updating. *ACM Transactions on Applied Perception*, 2(3):183–215, 2005. ISSN 1544-3558. doi: http://doi.acm.org/10.1145/1077399.1077401.

[104] Bernhard E Riecke, Aleksander Väljamäe, and Jörg Schulte-Pelkum. Moving sounds enhance the visually-induced self-motion illusion (circular vection) in virtual reality. *ACM Transactions on Applied Perception (TAP)*, 6(2):7, 2009.

[105] Bernhard E. Riecke, Bobby Bodenheimer, Timothy P. McNamara, Betsy Williams, Peng Peng, and Daniel Feuereissen. Do we need to walk for effective virtual reality navigation? physical rotations alone may suffice. In *Spatial Cognition*, pages 234–247, 2010.

[106] J J Rieser, H L Pick, D A Ashmead, and A E Garing. The calibration of human locomotion and models of perceptual-motor organization. *Journal of Experimental Psychology: Human Perception and Performance*, 21:480–497, 1995.

[107] John J. Rieser. Access to knowledge of spatial structure at novel points of observation. *Journal of Experimental Psychology*, 15(6):1157–1165, 1989.

[108] Jeffrey N. Rouder, Paul L. Speckman, Dongchu Sun, Richard D. Morey, and Geoffrey Iverson. Bayesian t-tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, 16:225–237, 2009.

[109] Roy A. Ruddle. Navigation: Am i really lost or virtually there? *Engineering Psychology and Cognitive Ergonomics*, 6:135–142, 2001.

[110] Roy A Ruddle. The effect of translational and rotational body-based information on navigation. In *Human Walking in Virtual Environments*, pages 99–112. Springer, 2013.

[111] Roy A Ruddle and Simon Lessels. For efficient navigation search, humans require full physical movement but not a rich visual scene. *Psychological Science*, 17(6):460–465, 2006.

[112] Roy A. Ruddle and Simon Lessels. The benefits of using a walking interface to navigate virtual environments. *ACM Transactions on Computer-Human Interaction*, 16(1):1–18, 2009. ISSN 1073-0516. doi: http://doi.acm.org/10.1145/1502800.1502805.

[113] Roy A Ruddle, Ekaterina Volkova, and Heinrich H Bülthoff. Walking improves your cognitive map in environments that are large-scale and large in extent. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 18(2):10, 2011.

[114] Roy A Ruddle, Ekaterina Volkova, Betty Mohler, and Heinrich H Bülthoff. The effect of landmark and body-based sensory information on route knowledge. *Memory & Cognition*, 39(4):686–699, 2011.

[115] Roy A Ruddle, Ekaterina Volkova, and Heinrich H Bülthoff. Learning to walk in virtual reality. *ACM Transactions on Applied Perception (TAP)*, 10(2):11, 2013.

[116] Silvia Serino and Giuseppe Riva. How different spatial representations interact in virtual environments: the role of mental frame syncing. *Cognitive processing*, 16(2):191–201, 2015.

[117] M Jeanne Sholl. From visual information to cognitive maps. *The construction of cognitive maps*, pages 157–186, 1996.

[118] Alexander W Siegel and Sheldon H White. The development of spatial representations of large-scale environments. *Advances in child development and behavior*, 10:9–55, 1975.

[119] Adalberto L Simeone, Ifigeneia Mavridou, and Wendy Powell. Altering user movement behaviour in virtual environments. *IEEE Transactions on Visualization and Computer Graphics*, 23(4):1312–1321, 2017.

[120] M. Slater and M. Usoh. Presence in immersive virtual environments. In *Proceedings of IEEE Virtual Reality Annual International Symposium*, pages 90–96, Sep 1993. doi: 10.1109/VRAIS.1993.380793.

[121] M. Slater, M. Usoh, and A. Steed. Depth of presence in virtual environments. *Presence: Teleoperators and Virtual Environments*, 3(2):15, 1994.

[122] M Slater, M Usoh, and A Steed. Taking steps: The influence of a walking technique on presence in virtual reality. *ACM Transactions on Human Computer Interaction*, 2(3):201–219, 1995.

[123] J. L. Souman, P. Robuffo Giordano, M. Schwaiger, I. Frissen, T. Thü mmel, H. Ulbrich, A. De Luca, H. H. Bülthoff, and M. O. Ernst. Cyberwalk: Enabling unconstrained omnidirectional walking through virtual environmen ts. *ACM Transactions on Applied Perception*, 8(4):25:1–25:22, December 2008. ISSN 1544-3558. doi: 10.1145/2043603.2043607. URL http://doi.acm.org/10.1145/2043603.2043607.

[124] Jan L. Souman, Ilja Frissen, Manish N. Sreenivasa, and Marc O. Ernst. Walking straight into circles. *Current Biology*, 19(18):1538 – 1542, 2009. ISSN 0960-9822. doi: https://doi.org/10.1016/j.cub.2009.07.053. URL http://www.sciencedirect.com/science/article/pii/S0960982209014791.

[125] Frank Steinicke, Gerd Bruder, Jason Jerald, Harald Frenz, and Markus Lappe. Analyses of human sensitivity to redirected walking. In *Proceedings of the 2008 ACM symposium on Virtual Reality Software and Technology*, pages 149–156, New York, NY, USA, 2008. ACM. ISBN 978-1-59593-951-7. doi: http://doi.acm.org/10.1145/1450579.1450611.

[126] Frank Steinicke, Gerd Bruder, Luv Kohli, Jason Jerald, and Klaus Hinrichs. Taxonomy and implementation of redirection techniques for ubiquitous passive haptic feedback. *Cyberworlds, International Conference on*, 0:217–223, 2008. doi: http://doi.ieeecomputersociety.org/10.1109/CW.2008.53.

[127] Frank Steinicke, Gerd Bruder, Jason Jerald, Harald Frenz, and Markus Lappe. Estimation of detection thresholds for redirected walking techniques. *IEEE Transactions on Visualization and Computer Graphics*, 16:17–27, 2010. ISSN 1077-2626. doi: http://doi.ieeecomputersociety.org/10.1109/TVCG.2009.62.

[128] Frank Steinicke, Yon Visell, Jennifer Campos, and Anatole Lécuyer. *Human Walking in Virtual Environments*. Springer, 2013.

[129] Evan A Suma, David M Krum, Samantha Finkelstein, and Mark Bolas. Effects of redirection on spatial orientation in real and virtual environments. In *3D User Interfaces (3DUI), 2011 IEEE Symposium on*, pages 35–38. IEEE, 2011.

[130] Evan A Suma, Gerd Bruder, Frank Steinicke, David M Krum, and Mark Bolas. A taxonomy for deploying redirection techniques in immersive virtual environments. In *Virtual Reality Short Papers and Posters (VRW), 2012 IEEE*, pages 43–46. IEEE, 2012.

[131] Evan A Suma, Gerd Bruder, Frank Steinicke, David M Krum, and Mark Bolas. A taxonomy for deploying redirection techniques in immersive virtual environments. In *2012 IEEE Virtual Reality Workshops (VRW)*, pages 43–46. IEEE, 2012.

[132] Jennifer E Sutton, Melanie Buset, and Mikayla Keller. Navigation experience and mental representations of the environment: do pilots build better cognitive maps? *PloS one*, 9(3): e90058, 2014.

[133] David Swapp, Julian Williams, and Anthony Steed. The implementation of a novel walking interface within an immersive display. In *3D User Interfaces (3DUI), 2010 IEEE Symposium on*, pages 71–74. IEEE, 2010.

[134] J N Templeman, P S Denbrook, and L E Sibert. Virtual locomotion: Walking in place through virtual environments. *Presence*, 8(6):598–617, 1999.

[135] Sam Tregillus and Eelke Folmer. Vr-step: Walking-in-place using inertial sensing for hands free navigation in mobile vr environments. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 1250–1255. ACM, 2016.

[136] Martin Usoh, Kevin Arthur, Mary C. Whitton, Rui Bastos, Anthony Steed, Mel Slater, and Frederick P. Brooks. Walking $>$ walking-in-place $>$ flying, in virtual environments. In *SIGGRAPH 99*, pages 359–364, August 1999.

[137] Steven G Vandenberg and Allan R Kuse. Mental rotations, a group test of three-dimensional spatial visualization. *Perceptual and Motor Skills*, 47(2):599–604, 1978.

[138] VIVE. Vive: Discover virtual reality beyond imagination. https://www.vive.com/us/, 2019. Accessed: 2019-9-30.

[139] D Waller, E Hunt, and D Knapp. The transfer of spatial knowledge in virtual environment training. *Presence*, 7(2):129–143, 1998.

[140] William H Warren, Bruce A Kay, Wendy D Zosh, Andrew P Duchon, and Stephanie Sahuc. Optic flow is used to control human walking. *Nature neuroscience*, 4(2):213–216, 2001.

[141] Steven M Weisberg and Nora S Newcombe. A slippery directional slope: Individual differences in using slope as a directional cue. *Memory & cognition*, 42(4):648–661, 2014.

[142] Steven M Weisberg and Nora S Newcombe. How do (some) people make a cognitive map? routes, places, and working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 42(5):768, 2016.

[143] Steven M Weisberg, Daniele Nardi, Nora S Newcombe, and Thomas F Shipley. Up by upwest: Is slope like north? *The Quarterly Journal of Experimental Psychology*, 67(10): 1959–1976, 2014.

[144] Wen Wen, Toru Ishikawa, and Takao Sato. Individual differences in the encoding processes of egocentric and allocentric survey knowledge. *Cognitive Science*, 37(1):176–192, 2013.

[145] Wen Wen, Toru Ishikawa, and Takao Sato. Instruction of verbal and spatial strategies for the learning about large-scale spaces. *Learning and Individual Differences*, 35:15–21, 2014.

[146] Jeremy D Wendt, Mary C Whitton, and Frederick P Brooks. Gud wip: Gait-understanding-driven walking-in-place. In *Virtual Reality Conference (VR), 2010 IEEE*, pages 51–58. IEEE, 2010.

[147] Mary C. Whitton, Joseph V. Cohn, Jeff Feasel, Paul Zimmons, Sharif Razzaque, Sarah J. Poulton, Brandi McLeod, and Jr. Frederick P. Brooks. Comparing ve locomotion interfaces. *Virtual Reality Conference, IEEE*, 0:123–130, 2005. ISSN 1087-8270. doi: http://doi.ieeecomputersociety.org/10.1109/VR.2005.25.

[148] Betsy Williams, Gayathri Narasimham, Tim P. McNamara, Thomas H. Carr, John J. Rieser, and Bobby Bodenheimer. Updating orientation in large virtual environments using scaled

translational gain. In *Symposium on Applied Perception in Graphics and Visualization*, pages 21–28, Boston, MA, July 2006.

[149] Betsy Williams, Gayathri Narasimham, Bjoern Rump, Timothy P. McNamara, Thomas H. Carr, John J. Rieser, and Bobby Bodenheimer. Exploring large virtual environments with an hmd when physical space is limited. In *Symposium on Applied Perception in Graphics and Visualization*, pages 41–48, Tübingen, Germany, July 2007.

[150] Betsy Williams, Stephen Bailey, Gayathri Narasimham, Muqun Li, and Bobby Bodenheimer. Evaluation of walking in place on a wii balance board to explore a virtual environment. *ACM Transactions on Applied Perception*, 8(3):19:1–19:14, August 2011. ISSN 1544-3558. doi: 10.1145/2010325.2010329. URL http://doi.acm.org/10.1145/2010325.2010329.

[151] Preston Tunnell Wilson, Kevin Nguyen, Alyssa Harris, and Betsy Williams. Walking in place using the microsoft kinect to explore a large ve. In *Proceedings of the 13th ACM SIGGRAPH International Conference on Virtual-Reality Continuum and its Applications in Industry*, pages 27–33. ACM, 2014.

[152] Preston Tunnell Wilson, William Kalescky, Ansel MacLaughlin, and Betsy Williams. Vr locomotion: walking¿ walking in place¿ arm swinging. In *Proceedings of the 15th ACM SIGGRAPH Conference on Virtual-Reality Continuum and Its Applications in Industry-Volume 1*, pages 243–249. ACM, 2016.

[153] Xianshi Xie. *An Examination of Navigation Methods for Large Immersive Virtual Environments with Application to the Study of Human-Robot Teams*. PhD thesis, Vanderbilt University, 2015.

[154] Xianshi Xie, Qiufeng Lin, Haojie Wu, Gayathri Narasimham, Timothy P. McNamara, John Rieser, and Bobby Bodenheimer. A system for exploring large virtual environments that

combines scaled translational gain and interventions. In *Proceedings of the 7th Symposium on Applied Perception in Graphics and Visualization*, APGV '10, pages 65–72, New York, NY, USA, 2010. ACM. ISBN 978-1-4503-0248-7. doi: http://doi.acm.org/10.1145/1836248.1836260. URL http://doi.acm.org/10.1145/1836248.1836260.

[155] Sen Xu, Alexander Klippel, Alan M MacEachren, and Prasenjit Mitra. Exploring regional variation in spatial language using spatially stratified web-sampled route direction documents. *Spatial Cognition & Computation*, 14(4):255–283, 2014.

[156] Run Yu, Wallace S Lages, Mahdi Nabiyouni, Brandon Ray, Navyaram Kondur, Vikram Chandrashekar, and Doug A Bowman. Bookshelf and bird: Enabling real walking in large vr spaces. In *3D User Interfaces (3DUI), 2017 IEEE Symposium on*, pages 116–119. IEEE, 2017.

[157] Ye Zheng, Matthew McCaleb, Courtney Strachan, and Betsy Williams. Exploring a virtual environment by walking in place using the microsoft kinect. In *Proceedings of the ACM Symposium on Applied Perception*, pages 131–131. ACM, 2012.