REDIRECTED FREE EXPLORATION WITH DISTRACTORS: A LARGE-SCALE, REAL-WALKING LOCOMOTION INTERFACE

Tabitha C. Peck

A dissertation submitted to the faculty of the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Computer Science.

Chapel Hill
2010

Approved by:

Henry Fuchs

Mary C. Whitton

Frederick P. Brooks, Jr.

Victoria Interrante

Dennis R. Proffitt

Anthony Steed

ABSTRACT

TABITHA C. PECK: Redirected Free Exploration with Distractors: A Large-Scale, Real-Walking Locomotion Interface
(Under the direction of Henry Fuchs and Mary C. Whitton)

Immersive Virtual Environments (VEs) enable user controlled interactions within the environment such as head-controlled point-of-view and user-controlled locomotion. In the real world people usually locomote by walking; walking is simple and natural, and enables people not only to move between locations, but also to develop cognitive maps, or mental representations, of environments. People navigate every day in the real world without problem, however users navigating VEs often become disoriented and frustrated, and find it challenging to transfer spatial knowledge acquired in the VE to the real world.

In this dissertation I develop and demonstrate the effectiveness of a new locomotion interface, Redirected Free Exploration with Distractors (RFED) that enables people to freely walk in large scale VEs. RFED is the combination of *distractors*—objects, sounds, or combinations of objects and sounds in the VE that encourage people to turn their heads, and *redirection*—making the user turn herself by interactively and imperceptibly rotating the virtual scene about her while she is turning her head. I demonstrate through user studies that compare RFED to a real-walking locomotion interface that RFED does not diminish user ability to navigate. I further demonstrate that users navigate better in RFED than with joystick and walking-in-place locomotion interfaces. Additionally, RFED does not significantly increase simulator sickness when compared to real walking, walking-in-place, and joystick interfaces.

PREFACE

All we have to believe with is our senses, the tools we use to perceive the world: our sight, our touch, our memory. If they lie to us, then nothing can be trusted. And even if we do not believe, then still we cannot travel in any other way than the road our senses show us; and we must walk that road to the end. (Neil Gaiman, *American Gods*)

Dedicated to my loving parents Martha and Earl Peck.

ACKNOWLEDGMENTS

I would like to thank my advisors, Henry Fuchs and Mary Whitton. I thank Mary for all of the time she spent helping me develop RFED and design user studies, for encouraging me to learn statistics, and for her countless hours editing papers and teaching me how to write. I thank Henry for his continuous enthusiasm, inspiring ideas, and help in finding a dissertation topic that I really interested me. This dissertation would not exist without all of Mary and Henry's patience and hard work. I would also like to thank the rest of my committee, Frederick Brooks, Victoria Interrante, Dennis Proffitt, and Anthony Steed for all of their helpful advice and feedback.

I would like to thank Sharif Razzaque for encouraging me to continue his work on Redirected Walking and taking time out of his busy schedule to answer any questions I had about Redirected Walking.

I wish to thank the EVE research group for all of their feedback on presentations, advice on study designs, and time fixing EVEIL bugs. I especially would like to thank:

- Eric Burns for encouraging me to come to my first EVE meeting
- Luv Kohli and Chris VanderKnyff for the help with EVEIL
- Chris Oates for recording audio and helping with models
- Jeremy Wendt for letting my use his GUD-WIP interface

I would like to thank Eric Knisley for modeling the distractor hummingbird.

I would like to thank The Link Foundation and The Anita Borg Institute for funding parts of this work.

I would like to thank the staff and faculty of The UNC Computer Science Department. This department is a wonderful and collaborative environment and I have truly enjoyed my time here. I would like to personally thank Courtney Ferriter, Kelli Gaskill, Missy Wood, Janet Jones, David Harrison, John Thomas, Bil Hays, David Marshburn, Andrei State, Russ Taylor, and Gary Bishop.

I would also like to thank my many friends and classmates especially Anish Chandak, Brian Clipp, Cory Quammen, Dave Millman, David Feng, Ilknur Kaynar Kabul, Kirsten Morrison, Keith Lee, Srinivas Krishnan, Stephen Olivier, Chad Spensky, Russ Gayle, Jason Sewall, Sarah Kennedy, Xiaoxiao Liu, and Zhimin Ren.

I would like to *The Trio*, Brad Moore, Tim Thirion, and Rick Skarbez for getting my through my first year. Especially Brad Moore for being an amazing friend.

I would like to that Catie Welsh, Laura Kassler, Caroline Green, and Jeanette Olli for the many much needed girls nights.

I thank my parents, Martha and Earl Peck for always encouraging and supporting me to excel throughout my life. This dissertation would not exist without you and for that reason I have dedicated this dissertation to you.

Last, but not least, I would like to thank my husband, David Borland, for always supporting me through this process. Thank you for you patience and understanding while I worked on this dissertation.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

3D	Three-Dimensional
FOE	Field of Expansion
FOR	Field of Regard
FOV	Field of View
MC	Motion Compression
POV	Point Of View
RDW	Redirected Walking
RFED	Redirected Free Exploration with Distractors
ROT	Reorientation Technique
UVN	Unilateral Visual Neglect
VE	Virtual Environment
WIP	Walking-In-Place

CHAPTER 1

Overview

Virtual environments (VEs) have been described as a window into a "mathematical wonderland" of computer generated worlds (Sutherland, 1965). VEs not only enable looking through a window, but walking through a door into a computer generated world. Immersive VEs enable user-controlled interactions within the environment such as head-controlled point-of-view and user-controlled locomotion. In the real world people usually locomote by walking; walking is simple and natural, and enables people not only to move between locations, but also to develop cognitive maps, or mental representations, of environments. People navigate every day in the real world without problem, however users navigating VEs often become disoriented and frustrated, and find it challenging to transfer spatial knowledge acquired in the VE to the real world (Durlach and Mayor, 1995; Psotka, 1995; Darken and Sibert, 1996; Grant and Magee, 1998).

In this dissertation I develop and demonstrate the effectiveness of a new locomotion interface for head-mounted displays, Redirected Free Exploration with Distractors (RFED) that enables people to freely walk in large scale VEs[1]. RFED is the combination of *distractors*—objects, sounds, or combinations of objects and sounds in the VE that encourage people to turn their heads, and *redirection*—making the user turn herself by interactively and imperceptibly rotating the virtual scene about her while she is turning her head (Razzaque, 2005). I demonstrate through user studies that compare RFED to a real-walking locomotion interface that RFED does not diminish user ability to navigate. I further demonstrate that users navigate better in RFED than with joystick and walking-in-place locomotion interfaces. Additionally, RFED does not significantly increase simulator sickness when compared

[1]This dissertation focuses on locomotion interfaces designed for dismounted users and does not discus vehicular transport.

to real walking, walking-in-place, and joystick interfaces.

1.1 Theory Behind RFED Development

Navigation is important for VE training applications where spatial understanding of the VE must transfer to the real world. Navigation is the combination of wayfinding and locomotion and as such is both cognitive and physical. Wayfinding is the cognitive aspect of navigation and does not involve movement. Wayfinding is the building and maintaining of a cognitive map, and is used to determine how to get from one location to another. Locomotion, the physical aspect of navigation, is defined as moving, physically or virtually, between two locations (Darken and Peterson, 2002). Specifically, navigation is essential for large VEs including exploring virtual cities, training ground troops, or visiting virtual models of houses.

Previous research suggests that users navigate best in VEs with real-walking locomotion interfaces (Ruddle and Lessels, 2009). Locomotion interfaces that provide users with vestibular and proprioceptive feedback improve user navigation performance and are less likely to cause simulator sickness than locomotion interfaces that do not stimulate both the proprioceptive and vestibular systems (Chance et al., 1998; Ruddle and Lessels, 2009). For this reason, RFED is designed to enable users to really walk in the VE. This is in contract to other VE locomotion[2] interfaces such as walking-in-place, omni-directional treadmills, or bicycles (Hollerbach, 2002; Darken et al., 1997) that require physical-input from the user, but do not stimulate both the proprioceptive and vestibular systems in the same way as really walking. RFED enables people to really walk, thus supporting user navigation by stimulating both the proprioceptive and vestibular systems.

Real-walking locomotion interfaces are believed to enable better user navigation, are more natural, and produce a higher sense of presence than other locomotion interfaces (Slater et al., 1995b; Usoh et al., 1999). However, because user motion must be tracked,

[2]In virtual environments, the term ‘travel’ instead of ‘locomote’ is often used to describe movement between two virtual locations, since some movement systems do not enable real-world user-physical movement (Bowman et al., 1997). I use ‘travel’ to describe passive motion interfaces like joysticks, and ‘locomote’ for active motion interfaces like walking-in-place or bicycles, to move between two virtual locations.

VEs using a real-walking locomotion interface have typically been restricted in size to the area of the tracked space. Current interfaces that enable real walking in larger-than-tracked-space VEs include redirected walking (RDW) (Razzaque et al., 2001; Razzaque et al., 2002; Razzaque, 2005), scaled-translational-gain (Robinett and Holloway, 1992; Williams et al., 2006a; Williams et al., 2007), seven-league boots (Interrante et al., 2007), and motion compression (MC) (Nitzsche et al., 2004; Su, 2007). Each of these interfaces transforms the VE or user motion by rotating the environment or scaling user motion. While these transformations enable large-scale real-walking in VEs, the effect of the transformations on navigational ability is unknown. I evaluate the effect of rotational transformations on navigational ability through two user studies presented in Chapters 6 and 7.

Additionally, when freely walking in the locomotion interfaces mentioned above, users may find themselves about to walk out of the tracked space, and possibly into a real wall. When a user nears the edge of the tracked space a reorientation technique (ROT) must be used to prevent the user from leaving the tracked space (Peck et al., 2009). ROTs must stop the user and rotate the VE around her current virtual location, returning the immediately predicted path into the tracked space. The user must also reorient herself by physically turning around in the real environment so she can follow her desired path in the newly-rotated VE. ROTs are required for real-walking interfaces enabling free exploration of large VEs and are not required for interfaces using devices that constrain user physical movement such as joysticks, walking-in-place interfaces, treadmills, or bicycles (Brooks, 1987; Christensen et al., 2000; Darken et al., 1997; Iwata, 1999; Slater et al., 1995a). However, many current ROT implementations cause breaks in presence by using disembodied voices, or quickly spinning the VE around the user, which detract from the immersive VE experience.

In this dissertation I develop and evaluate a new ROT, distractors, through three user studies (Chapter 5). Distractors are objects, sounds, or combinations of objects and sounds in the VE that encourage people to turn their heads. Redirection, is best accomplished while the head is turning, is "making the user turn herself by interactively and imperceptibly rotating the virtual scene about her," (Razzaque, 2005) .

I then applied distractors to a new locomotion interface, Redirected Free Exploration with Distractors (RFED) which enables real walking to stimulate the proprioceptive and vestibular systems to support user navigation. RFED enables large-scale real walking by combining *distractors* and *redirection* (Chapter 4). I then evaluated user navigational ability and simulator sickness when using RFED compared to real-walking, walking-in-place, and joystick interfaces (Chapters 6 and 7).

1.2 Thesis Statement

A large-scale, real-walking locomotion interface using distractors and redirection enables people to freely locomote in larger-than-tracked-space virtual environments, navigating no worse than real-walking and better than joystick and walking-in-place interfaces.

To demonstrate the validity of my thesis statement, I

1. Develop a large-scale, real-walking locomotion interface using distractors and redirection, Redirected Free Exploration with Distractors (RFED), that enables people to freely locomote larger-than-tracked-space virtual environments
   - Develop and Evaluate Distractors through three user studies (Chapter 5)
   - Develop and implement RFED (Chapter 4)
   - Demonstrate that RFED users can freely locomote in larger-than-tracked-space VEs in two user studies (Chapters 6 and 7)
2. Compare navigational ability
   - RFED users navigate no worse than users who really walk (Chapter 6)
   - RFED users navigate better than users of walking-in-place and joystick interfaces (Chapter 7)

1.3 Overview of Dissertation

Chapters 2 and 3 introduce background about perception, locomotion interfaces, navigation, and simulator sickness. Chapter 5 introduces, develops, and evaluates distractors, my new

ROT, through three user studies. Chapter 4 discusses the design of RFED in detail and should be a reference for future RFED implementations. Chapters 6 and 7 present two user studies evaluating user navigational ability when using RFED by comparing it to real walking, walking-in-place, and joystick interfaces. Chapter 8 presents an overview of the final results and discusses future research areas that may improve RFED implementations. The questionnaires used in the user studies can be found in Appendix A. The rest of Chapter 1 presents an overview and results of the five user studies.

1.4 Overview of User Studies and Results

1.4.1 Distractors: Chapter 5

Virtual environments that use a real-walking locomotion interface have typically been restricted in size to the area of the tracked lab space. Techniques proposed to lift this size constraint, enabling real walking in VEs that are larger than the tracked lab space, all require reorientation techniques (ROTs) in the worst-case situation–when a user is close to walking out of the tracked space. I propose a new ROT using visual and audial distractors–objects in the VE that the user attends while the VE rotates–and compare my method to current ROTs through three user studies. ROTs using distractors were preferred and ranked more natural by users. Users were also less aware of the rotating VE when ROTs with distractors were used. The findings also suggest that improving visual realism and adding sound to distractors increased a user's feeling of presence. Much of the work reported in this chapter was published in (Peck et al., 2009).

1.4.2 Redirected Free Exploration with Distractors versus Real Walking: Chapter 6

Users in virtual environments often find navigation more difficult and frustrating than in the real world. This effect is compounded by locomotion interfaces that do not enable real walking in the VE, since RW stimulates the proprioceptive and vestibular systems (Chance et al., 1998; Ruddle and Lessels, 2009). I have developed a new locomotion interface, Redirected Free Exploration with Distractors (RFED) to support user navigation in large-scale VEs. The interface is discussed in Chapter 4.

I compared RFED to the current best interface, really-walking, by conducting a user study measuring navigational ability. The results show that RFED users can really-walk through VEs that are larger than the tracked space and can point to targets and complete maps of VEs no worse than when really walking (Peck et al., 2010).

1.4.3 Redirected Free Exploration with Distractors > Walking-In-Place and Joystick: Chapter 7

I performed a between-subjects study comparing navigation ability in Redirected Free Exploration with Distractors, Walking-in-Place, and Joystick locomotion interfaces in VEs that are more than two times the dimensions of the tracked space. I evaluated the interfaces based on navigation and wayfinding metrics and found that participants using RFED were significantly better at navigating and wayfinding through virtual mazes than participants using walking-in-place and joystick interfaces. Participants traveled shorter distances, made fewer wrong turns, pointed to hidden targets more accurately and more quickly, and were able to place and label targets on maps more accurately. Moreover, participants were able to more accurately estimate VE size.

CHAPTER 2

Perception

This chapter introduces the reader to the sensory systems that are stimulated by Redirected Free Exploration with Distractors. After defining the physical systems that are relevant for RFED, I discuss the psychology of visual perception and psychology studies that suggest directions for future development of RFED.

2.1 Sensory Systems

The human sensory systems are part of the nervous system and have sensory receptors that receive stimulation from internal and external sources. The stimuli pass information to the brain via neural pathways, and the brain then processes the information. RFED takes advantage of the human sensory systems by stimulating the senses with rendered images that do not match what should be seen based on the user's physical actions. To understand why people do not "perceive" these inaccuracies and yet physically respond to changes in rendered images, e.g. rotating or scaling the VE, requires an understanding of visual perception and the human sensory systems. An understanding of human perceptual systems provides insight for future RFED developers and guides future RFED development.

2.1.1 The Visual System

The sensing part of the visual system is the eyes (Figure 2.1). The cornea, the transparent layer that covers the iris and pupil, refracts and focuses light that enters the eye. The pupil, the dark center part at the front of the eye, allows light to enter the eye, while the iris, the colorful part surrounding the pupil, contracts and expands to control the amount of light permitted to pass into the eye. The lens, the biconvex structure behind the pupil, along

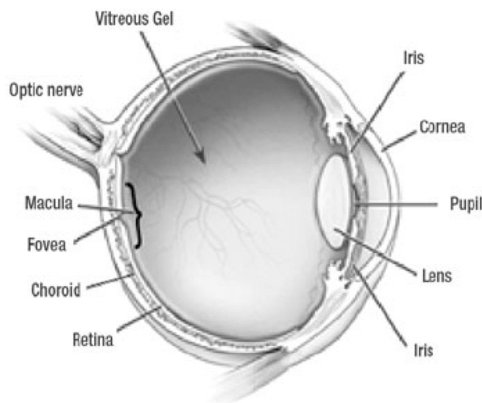

Figure 2.1: The anatomy of the human eye (Connect, 2010).

with the cornea, refracts and focuses the light onto the retina at the back of the eye. The retina is comprised of photoreceptors, called rods and cones, which are most dense on the fovea. The rods and cones are stimulated by the light coming through the cornea and lens. The stimulated photoreceptors send information through the optic nerve along the dorsal and ventral streams in the brain (See Section 2.2.4).

2.1.2 Depth Cues

Having two eyes enables stereopsis, a component of depth perception, which comes from the retinas being horizontally offset acquiring two slightly different images. Each eye sees an image and the brain calculates the difference between corresponding objects in the images, and due to the difference, depth is perceived. Stereopsis can be produced by two two-dimensional images called *stereo pairs*. An example of a stereo pair can be seen in Figure 2.2. To view the image in stereo, one image of the stereo pair is displayed to each eye, such as in a stereopticon. Along with stereopsis, there are other visual cues which convey depth.

- *Motion parallax* is the depth cue that comes from the observer's motion. As people move, closer objects move a greater distance across their field of view.
- *Perspective* is the understanding that parallel lines converge at infinity. This depth cue is often used by artists to convey depth in paintings and drawings such as *The School of Athens* painted by Raphael (1483-1520), (Figure 2.3).

Figure 2.2: A cross-eye stereo photo pair of a kasuga lantern on Wooded Island, Jackson Park, Chicago, Illinois (Scarborough, 2007).

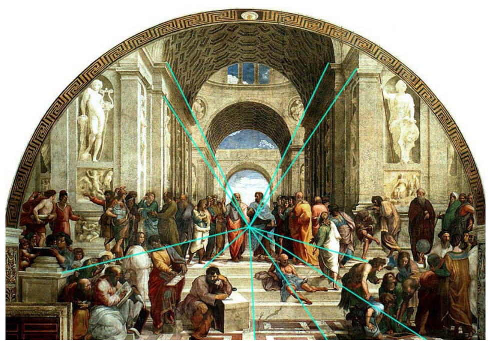

Figure 2.3: *The School of Athens* by Raphael (1483-1520). The superimposed lines are perspective parallel lines meeting at "infinity" in the center of the painting (Raphael, 1520).

- People use *relative size* to determine depth when two objects are known to be the same size. The object that appears larger on the retina is the closer object.

- *Looming* objects are ones that move toward or away from the viewer, yielding information about depth derivative.

- *Accommodation* is an occulomotor cue based on the kinesthetic sensation of the muscles in the eye that focus the lens. The movement of the muscles provide depth information from the kinesthetic sensation that the muscles stretch the lens more for farther away objects.

- *Convergence*, the inward rotation of the eyes, also provides an occulomoter cue about depth from the kinesthetic sensation of the muscles that rotate the eye. The eyes rotate inward greater amounts for nearer objects.

- When an object blocks all or part of another object from view, it is said that the closer object *occludes* the farther object. Although occlusion provides depth information, it only provides relative distance information.

- *Texture gradients* are the patterns of light formed on the retina from light reflected off textured surfaces, and can give an impression as to the shape or direction of a surface. Texture gradients provide depth cues from the granularity of the texture on the surface. As the depth of a surface increases, so does the texture fineness. Examples of texture gradients providing shape or surface information can be seen in Figures 2.10 and 2.11.

- *Haze* lightens distant objects and provides depth information for mostly outdoor scenes at large distances.

2.1.3 The Vestibular System

The vestibular system is the non-auditory part of the inner ear labyrinth and interprets head movement and aids balance. The vestibular system is composed of the semicircular canals and the otolith organs, which are located in the utricle and saccule. See Figures 2.4 and 2.5.

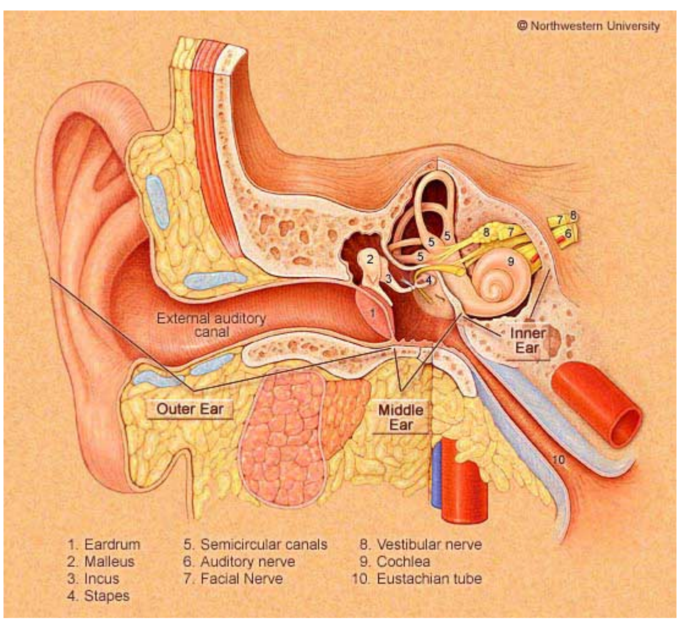

Figure 2.4: The ear (Northwestern, 2001a).

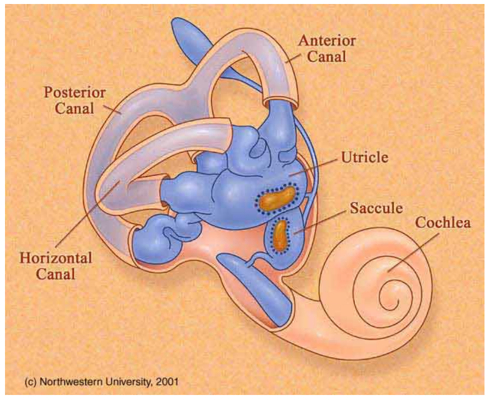

Figure 2.5: The semicircular canals and the utricle and saccule. Notice that three the semicircular canals are orthogonal to each other (Northwestern, 2001b).

Angular acceleration is sensed along the three-axes through the three orthogonal circular structures, called the *semicircular canals* within the inner ear. See Figure 2.4. The semicircular canals are filled with a liquid called *endolymph* and little hairs called *cilia*. The cilia are the sensory receptors in the semicircular canals and are embedded in a gelatinous structure called the *cupula*. As the head turns, the endolymph liquid flows through the three orthogonal canals, and moves the cupula, which stimulates the cilia. The moving cilia send information to the brain. The brain then interprets the person's movement. See Figure 2.6.

The two otolith organs are located in the saccule and utricle in the inner ear, and sense linear acceleration, including gravity, each in two directions. See Figure 2.5. The otolith organs are perpendicular to each other enabling the otoliths to sense acceleration along three orthogonal accelerations. The otolith organs also contain cilia. On top of the cilia is a gelatinous substance and the otolith membrane. *Otoconia crystals* of calcium carbonate sit on top of the otolith membrane adding weight, so that when the head moves, the otolithic membrane tilts in the direction of head motion. The motion of the otolithic membrane stimulates the cilia which then send information along the dorsal and ventral streams in the brain (See Section 2.2.4). See Figure 2.7.

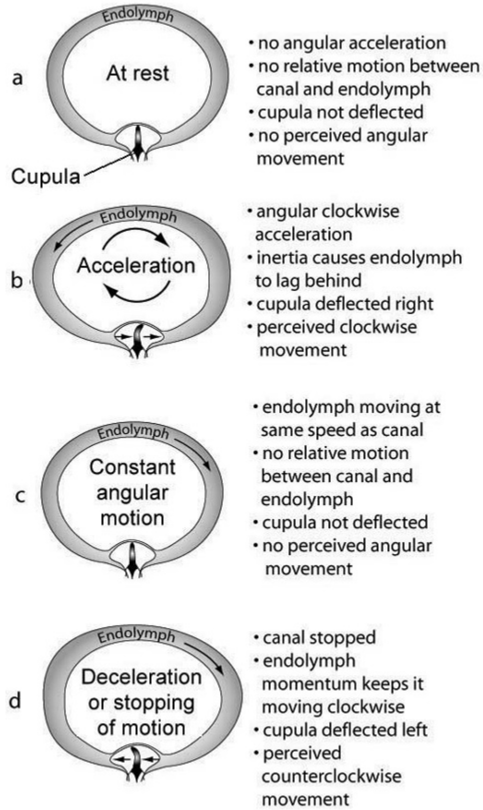

Figure 2.6: Acceleration in a semicircular canal (NASA, 2003).

2.1.4 The Interaction of the Visual and Vestibular Systems

The otolith organs and the semicircular canals detect head motion and convey head motion to the brain as neurological signals. Studies from aircraft simulation (Young, 1967; Young et al., 1969) have led to mechanical systems that can model the vestibular sensors. Figure 2.8 shows, as a function of time, simulation results for the relative response to the visual cues of motion, vestibular cues of motion, and the combination of both visual and vestibular (Borah et al., 1979). For a constant stimulus, the vestibular cues are initially dominant, however over time the visual cues become dominant.

Research (Hosman and van der Vaart, 1981) determined the spectrum of sense sensitivity. The results suggest that vision is most sensitive at low frequencies (e.g. one side-to-side head turn per 100 seconds) of motion followed by the otolith and semicircular canals at higher frequencies (e.g. one side-to-side head turn per 1.25 seconds). See Figure 2.9. These results suggest that when the head is not moving or moving at slow frequencies, that the

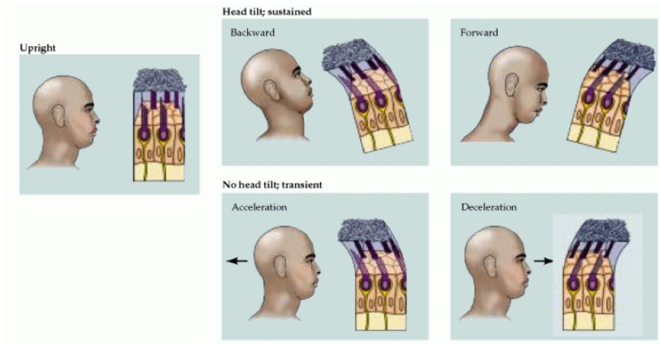

Figure 2.7: The figure shows an otolith membrane with cilia and otoconia crystals. "Forces acting on the head and the resulting displacement of the otolithic membrane of the utricular macula. For each of the positions and accelerations due to translational movements, some set of hair cells will be maximally excited, whereas another set will be maximally inhibited. Note that head tilts produce displacements similar to certain accelerations" (Sinauer Associates, 2001).

visual system is dominant. As head angular velocity increases, vestibular dominance increases.

The implication for this work in redirection is that when people turn their heads, the vestibular system dominates and visual manipulation may go unnoticed (Figure 2.9). Redirection is imperceptible because when people turn their heads at normal angular velocities, the vestibular system dominates the visual system. This enables the VE designer to rotate the VE visuals without the user noticing.

2.1.5 Other Sensory Systems

Other sensory groups that are relevant for RFED are the kinesthetic, somatosensory, and proprioceptive systems, because they are stimulated during locomotion. As stated in Chapter 1, locomotion interfaces that provide users with vestibular and proprioceptive feedback (1.) improve user navigation performance and (2.) are less likely to cause simulator sickness than locomotion interfaces that do not stimulate both the proprioceptive and vestibular systems (Chance et al., 1998; Ruddle and Lessels, 2009). The following definitions are based on definitions from (Stanney, 2002).

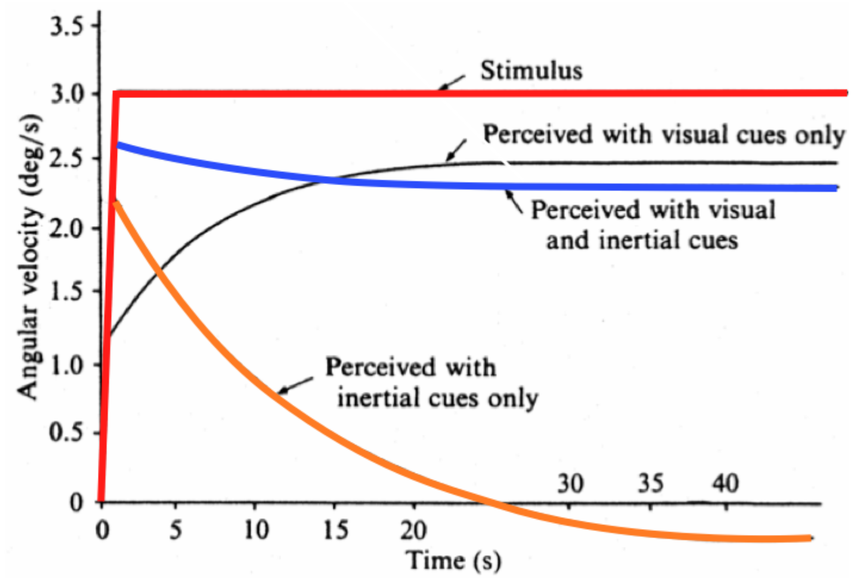

Figure 2.8: The contribution of the visual and vestibular (inertial) senses, in the time domain, to the perception of a step in rotational velocity (about the yaw axis) (adapted from Borah et al., 1979).

Kinesthetic System The kinesthetic system senses the movement of muscles, tendons, and joints.

Somatosensory System The somatosensory system includes all the body senses from the skin, muscles, joints, and internal organs. The somatosensory system includes the kinesthetic system but does not include the vestibular, visual, auditory, or taste and smell. It is comprised of senses from cutaneous, muscle, and joint receptors.

Proprioceptive System Proprioception is the internal sense of body position and movement, and includes the kinesthetic and vestibular senses.

2.2 Visual Perception

Much of the information presented in this section can be found in (Bruce et al., 2003).

Visual perception is more than seeing shapes, objects, and people. It provides understanding and awareness of the surrounding world. Visual perception is achieved through a combination of visual understanding and physical actions through the combination of the

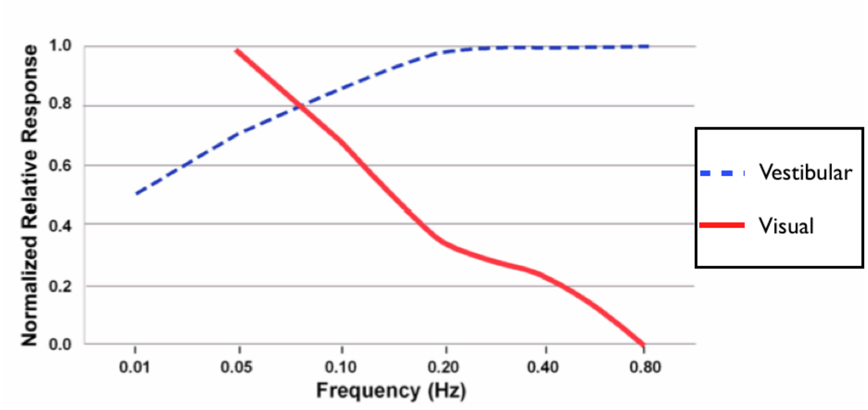

Figure 2.9: The visual-vestibular crossover. This graph shows, in the frequency domain, the relative contributions of visual and linear vestibular cues to postural stability (adapted from Duh et al., 2004) .

visual and kinesthetic systems.

2.2.1 Gibson's Theory of Perception

James Gibson began studying perception in WWII to predict successful and unsuccessful pilots based on the pilot's capability to decipher depth from images. His findings guided his theory that the observer's current evaluation of "depth" and "space" perception should be studied through the perception of textures on surfaces. Gibson claims that the structure of light produced by surfaces provides information for visual perception; visual perception is not just light waves stimulating the retina to enable people to "see", as described in Section 2.1.1.

Gibson's theory that *texture gradients*, the structure formed on the optic array from light reflecting off textured surfaces, can give an impression of the shape or direction of a surface. Examples of structure from texture can be seen in Figures 2.10 and 2.11. The structure from texture theory was further verified by (Beck and Gibson, 1955), who found that changing the texture gradient on a plane changes perceived slope of the plane.

Additionally, texture gradients provide information about distance and size of objects. Gibson's "ground" theory (Gibson, 1950) states that the perception of objects is based on the texture of the ground compared to the object. An example of texture providing information about distance and size of objects can be seen in Figure 2.12. Since texture

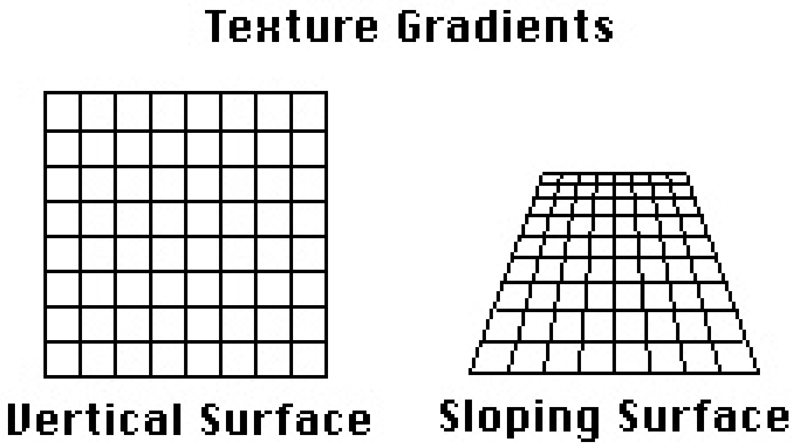

Figure 2.10: Applying different textures makes the surface appear to be vertical or sloping. (Redding, 2000)

gradients are *invariants* on the optic array, implying that properties of texture gradients do not change even if the world conditions change, an object will always occlude the same amount of a texture no matter its distance from an observer. The textured surface in Figure 2.12 informs us that the box on the left is farther away than the box on the right.

Gibson believed that the structure of the light reflecting from textured surfaces is required for perception. An example of this is in a Ganzfeld experiment (Metzger, 1930; Gibson and Dibble, 1952; Gibson and Waddell, 1952), where ping-pong balls are placed over participants' eyes. Although light is reaching the eyes, there is no structure to the light and participants' cannot perceive their surroundings. This is an example of the disparity between the physics of optics and the psychology of perception because texture is needed to perceive objects. Light hitting the eye without texture enables a person to "see," however they will not perceive.

2.2.2 Vision and Locomotion

In addition to Gibson's theory that shape and structure are perceived from texture, Gibson theorized that perception and physical action are connected and cannot be viewed as separate. Gibson believed that, since people actively explore the world, rotations and translations of light on the optic array produced by an observer's active movements are necessary

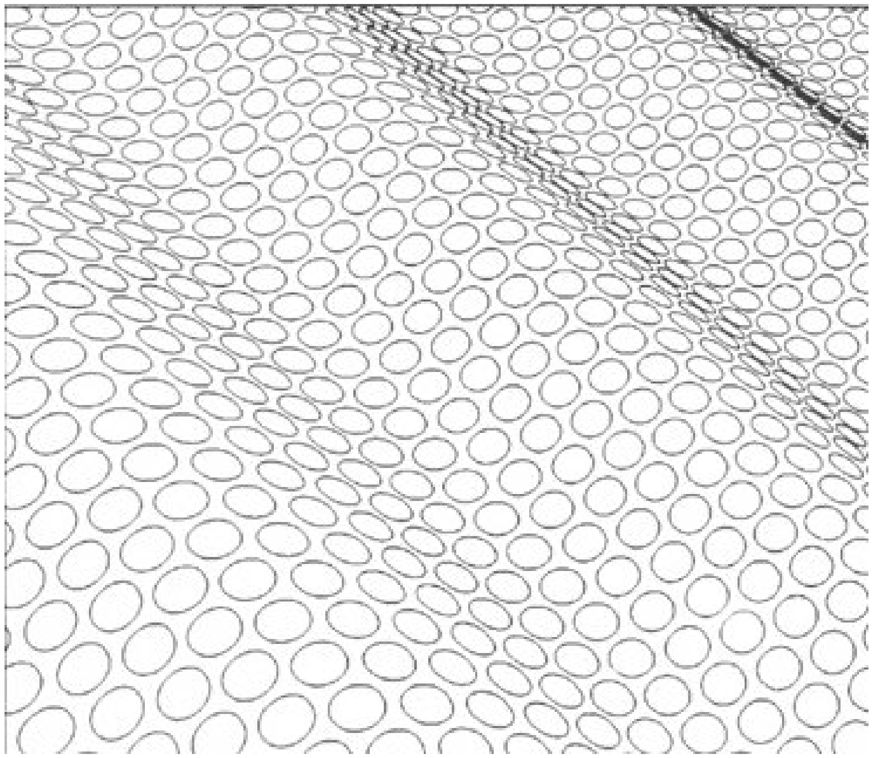

Figure 2.11: The texture causes the image to appear to have a ripple shape (Wiersma, 2007).

for visual perception. As people move, the views of the surroundings change, and information about the layout of the environment and the shape of surfaces, as well as their relative position within the environment, are revealed.

The impression of self-motion, known as *vection*, can be produced by visual stimulation alone. Vection can occur when a person is sitting in a stationary car and the adjacent car starts to move, causing the person to perceive a sensation of backwards motion.

Movement, essential for accurate perception of the environment, causes *optic flow*, the pattern of light on the optic array caused by relative motion between the observer and environment. Optic flow patterns contain information about self-motion, the motion of objects, and the environment's three-dimensional (3D) structure. If an observer is moving forward, the optic flow will radiate outward from the *center of expansion*, the point toward which the person is moving. If a person is riding in a train and looking out the window, the optic flow will move horizontally across the observer's retina. This horizontal optic flow is known as *lamellar flow.* Figure 2.13 shows radial and lamellar flow.

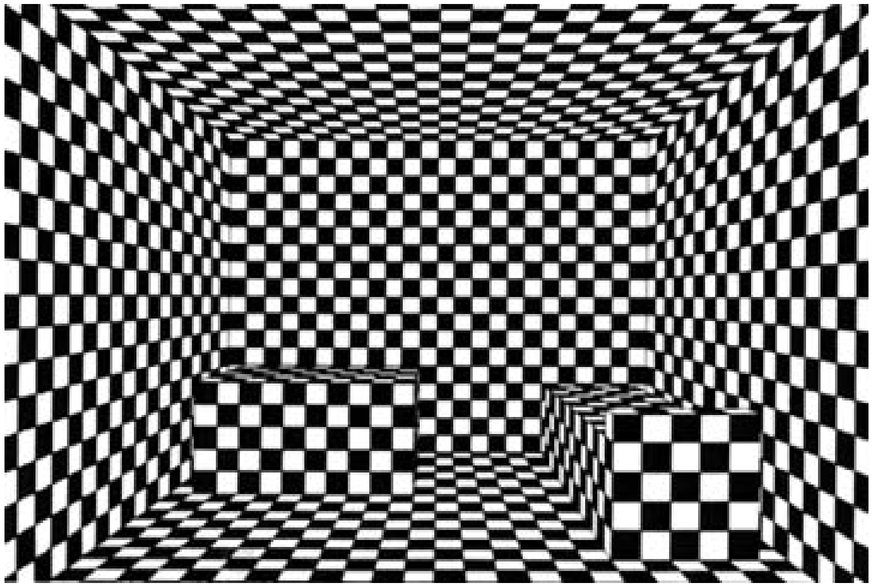

Figure 2.12: The texture provides information about depth, size, and orientation of the boxes (Schouten, 2003).

Gibson (Gibson, 1979), describes the important relationship between locomotion and optic flow as,

1. Locomotion is specified by flow of the *ambient optic array*, structured light that reaches the viewer, while stasis is specified by the absence of flow;

2. The type of locomotion is specified by the optic flow such that outflow is an approach, inflow a retreat;

3. The *center of expansion* of outflow specifies the locomotion direction;

4. A change in the *center of expansion* specifies a change in direction;

As stated by (Warren, 2004, p. 1247), "optic flow is a key example of Gibson's (Gibson, 1979) ecological approach to perception and action." The *ecological approach to perception and action* suggests that vision guides behavior in an environment, and that information provided from optic flow is important for specifying complex properties of an environment's relationship to the observer (Warren, 2004). (Warren, 2004) suggests that the exploitation of optical information could potentially control locomotion, and this is demonstrated through numerous studies (presented in Section 2.2.3) that modify optic flow by moving

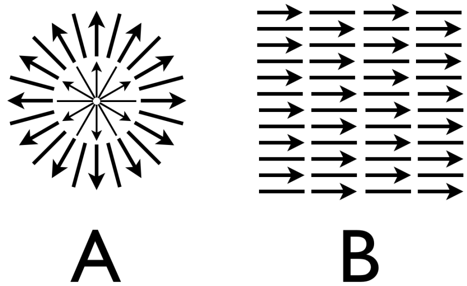

Figure 2.13: A. Outward radial optic flow. B. Lamellar, or horizontal optic flow.

the environment which changes the radial or lamellar optic flow (Figure 2.13) presented to the participant. The studies and results are discussed in Section 2.2.3.

Optic flow is represented as an instantaneous velocity field with vectors corresponding to the optical motion of points in the environment. The translational component of optic flow is produced by an observer "translating" forward or backward in the environment, thus producing a radial flow pattern on the retina (Figure 2.13, A). The rotational component of optic flow occurs when the eye or head rotates, which produces a lamellar flow pattern (Figure 2.13, B). In addition to the flow caused by translation and rotation of the body and head, the eyes rotate (Gibson, 1950) which adds additional retinal optic flow. The decomposition of the flow on a person's eyes into rotational and translational information provide information about the person's locomotion, however the added optic flow from eye rotation makes decomposition non-trivial.

The next section presents results of studies exploring the relationship between optic flow and locomotion.

2.2.3 Perception and VEs: Studies that Guide Large-Scale Real-Walking Interface Design

Gibson described a relationship between optic flow and locomotion, however people can stand and walk even with closed or masked eyes because of feedback provided from other sensory systems (Section 2.1). Nevertheless, vision has an important role in human locomotion because it is constantly used to adjust for errors that may occur in the vestibular or kinesthetic systems.

Numerous studies have explored the relationship between vision and locomotion by manipulating the optic flow or visuals presented to the user. Although these studies are important in understanding the connection between vision and locomotion, they also provide additional evidence supporting the capability of visual manipulation in VEs to redirect users. These studies also suggest guidelines for potentially successful techniques. In this section, it is assumed the reader has some understanding of how Redirection and RFED work (Chapter 4).

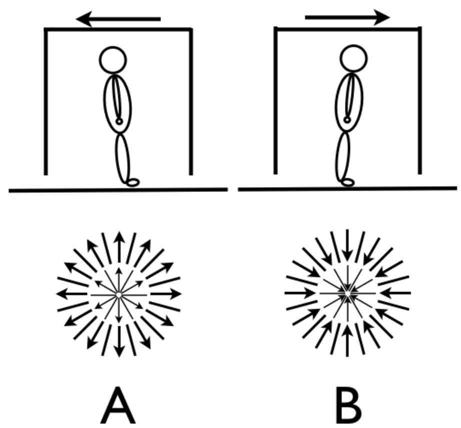

Figure 2.14: Based on an image from (Bruce et al., 2003). A. The room swings toward the subject causing outward optic flow. The participant interprets the optic flow as if he is moving forward and compensates by leaning backward. B. The room swings away from the subject causing inward optic flow. The participant interprets the optic flow as if she is moving backward and compensates by leaning forward.

(Lee and Aronson, 1974; Lee and Lishman, 1975; Lishman and Lee, 1973) studied the importance of vision in maintaining balance using the "swinging room," see Figure 2.14. The swinging room is comprised of textured walls that are suspended around the participant and are moved without participant knowledge. As the walls move toward the participant, an outward optical flow is produced and when the walls move away from the participant, an inward flow is produced. To explore the role of optic flow in balance, (Lee and Aronson, 1974), placed toddlers in the swinging room and found that the toddlers fell toward the wall when it was moving away from them, and away from the wall when it was moving toward them. These results are consistent with what was expected of the children if optic flow is interpreted correctly. If the child were to sway forward, causing an outward radial flow, to compensate the child should lean backward. This result suggests that radial optic flow should not be imperceptibly manipulated for a stationary standing user because the user may unknowingly lean forward or backward.

(Lee and Lishman, 1975) describes participants in the "swinging room" experiments as "visual puppet(s)" whose balance is manipulated by imperceptibly moving the surroundings around the user. Results from "swinging room" experiments suggested that redirection (Razzaque, 2005) and other VE techniques that manipulate optic flow, such as seven-league boots and motion compression, will be able to redirect user motion as you would a puppet.

(Bardy et al., 1996) investigated the role of vision on posture by having participants walk on a treadmill while the experimenters manipulated optic flow and motion parallax on surrounding walls. Participants were found to sway with the optic flow images, however manipulation of motion parallax had a greater effect than manipulation of simple horizontal flow. Their results suggest that people use vision to stay upright, and use motion parallax and optic flow to adjust their bodies. The result that motion parallax has a greater effect suggests that when manipulating optical flow in a VE, a detailed 3D environment should be used, so that motion parallax cues are present.

There are numerous studies evaluating the effect on increasing and decreasing optic flow although the results between some studies are contradictory. An experiment by (Konczak, 1994) explored using a "swinging room" to change the optic flow of participants during walking. They found that as optic flow slowed down, subjects' walking speed slightly increased, however increasing the speed of optic flow appeared to have no effect on participants. Changing the speed of optic flow produces similar results to scaling step size, as in scaled-translational gain, (Robinett and Holloway, 1992; Williams et al., 2006a; Williams et al., 2006b), and seven-league boots, (Interrante et al., 2007). The results from (Konczak, 1994) suggest that increasing user step size will increase optic flow and have little effect on the user. However, caution should be taken in slowing optic flow (decreasing step size) because it may increase the user's walking speed. However, additional studies exploring walking speed as a function of optic flow suggest that walking speed can be inversely controlled by optic flow, (Prokop et al., 1997; Schubert et al., 2005; Varraine et al., 2002) suggesting that increased optic flow will have an effect on users.

In addition to controlling walking speed and posture, vision also guides *heading direction*, the user's forward direction of motion. Gibson's theories suggest that heading is determined from the center of expansion of optic flow. When people walk toward a target, they adjust their movements to align heading direction with the intended goal.

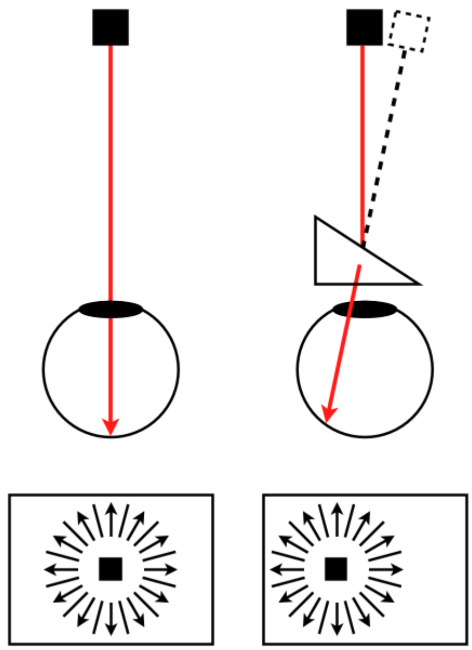

Figure 2.15: **Left.** As the person looks directly along the heading vector, optic flow radiates outward from the center of vision. **Right.** A prism is placed in front of the eye which shifts the visual location of the goal and the location of the radial optic flow. The optic flow on the retina is the same pattern as on the *left* but shifted due to the prism. Based on an image from (Rushton et al., 1998).

An alternate theory to Gibson's optic flow controlling heading, is the *egocentric direction hypothesis*, which states that heading is determined by the anterior-posterior axis of the body. This theory was explored by (Rushton et al., 1998) after observing a subject referred to as WV. WV has unilateral visual neglect (UVN) - damage to one side of the cerebral hemisphere and the inability to respond to stimuli on the side opposite the lesion. UVN is often associated with a misperception of location. (Rushton et al., 1998) observed WV walking in curved paths to reach target objects. To simulate the misperception of the target location for individuals without UVN, (Rushton et al., 1998) had participants wear prisms in front of their eyes. The prism translates not only the target object, but also the optic flow produced when the participant walked toward the target. See Figure 2.15. If the optic flow theory were true, subjects would walk straight to the target because the center of expansion of the optic flow will radiate from the heading direction. However, participants walked along a curved path with the heading direction deviating from the target by approximately the angle of the prism glasses that the participants wore. This suggests that the perceived location of objects guides locomotion direction, and not optic flow.

A study from (Harris and Carre, 2001) suggests that the restricted field of view (FOV) caused by wearing the prism glasses in (Rushton et al., 1998) prevented the use of optic flow in participants' peripheral vision. Optic flow in peripheral vision guides walking direction (Warren and Kurtz, 1992). When participants wear an HMD, the FOV restricts peripheral vision and therefore heading cannot be guided with peripheral vision. For future RFED development with a wide FOV, optic flow may dominate and redirection may not work as well. Further study evaluating the impact of FOV on redirection may show that changes in optic flow in the peripheral vision guide heading.

Another result supporting the theory that redirection may not work as well in wide field-of-view HMDs is work done examining the importance of radial and lamellar flow in determining heading direction. Results from (Crowell and Banks, 1993) suggest that people are more accurate at determining heading from radial flow than from lamellar flow. Additionally, according to (Warren and Kurtz, 1992), lamellar flow is only used to determine heading in the periphery. VE rotation around the user will affect lamellar flow. HMDs without a wide FOV do not provide peripheral vision. Since lamellar flow is only used to determine heading in the periphery, and radial flow in the center of vision guides heading, rotation of the VE around the user should not alter heading direction.

A problem with the prism experiment performed by (Rushton et al., 1998) is that prisms warp the optical flow pattern possibly leading participants to rely more on the egocentric direction hypothesis. A study by (Warren et al., 2001) using virtual environments compared the difference between using prisms to displace optic flow and virtually displacing optic flow in a virtual environment. The results suggest that people wearing prisms generally use the egocentric direction hypothesis to steer to a goal; however, when placed in the same condition with virtual displacement of optic flow, used a combination of egocentric direction and the optic flow theory, with "increasingly dominant behavior" toward the optic flow hypothesis when more optic flow was added to the scene. See Figure 2.16.

(Warren et al., 2001) further investigated whether the egocentric direction hypothesis or the optic flow hypothesis dominates. (Warren et al., 2001) had people walk through virtual environments with different amounts of optic flow (Figure 2.16) to see if the amount of optic flow affected participant heading direction to a target. Their results show that with no optic flow participants followed the egocentric direction hypothesis, however when optic flow was added to the ground plane, participants initially followed the egocentric direction hypothesis, and then after traveling a few meters participants adjusted their heading and used optic flow to aid their guidance.

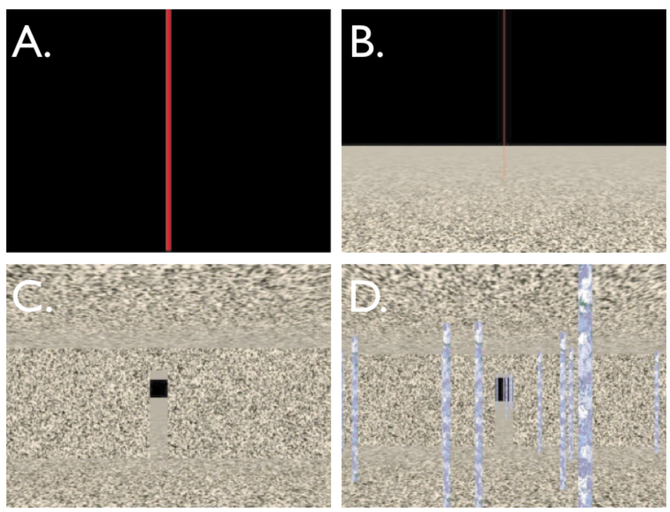

Figure 2.16: The four environments used in (Warren et al., 2001) with different amounts of optic flow. *A.* a target line, *B.* a target line and textured ground plane, *C.* a fully textured environment with a doorway, *D.* a fully textured environment with doorway and posts. Based on an image from (Warren et al., 2001).

The results from (Warren et al., 2001) suggest that humans rely on both optic flow and egocentric direction to guide locomotion. This dual guidance explains how people locomote to a target in low light or at night, where optic flow is not available. (Warren et al., 2001) developed a model of human locomotion guidance toward a goal, combining the egocentric direction theory and the optic flow hypothesis:

$$d\theta/dt = -k(\beta + wv\alpha) \quad (2.1)$$

where the turning rate ($d\theta/dt$) is a sum of egocentric direction and optic flow. In an extrinsic reference frame, θ represents the walking direction and β is the egocentric direction to the goal. α is the visual angle between the "center of expansion" and the goal, w is a measure of the magnitude and angular area of flow presented to the optic array and due to environmental structure, v is the observer's velocity, and k is a turning rate constant. See Figure 2.17.

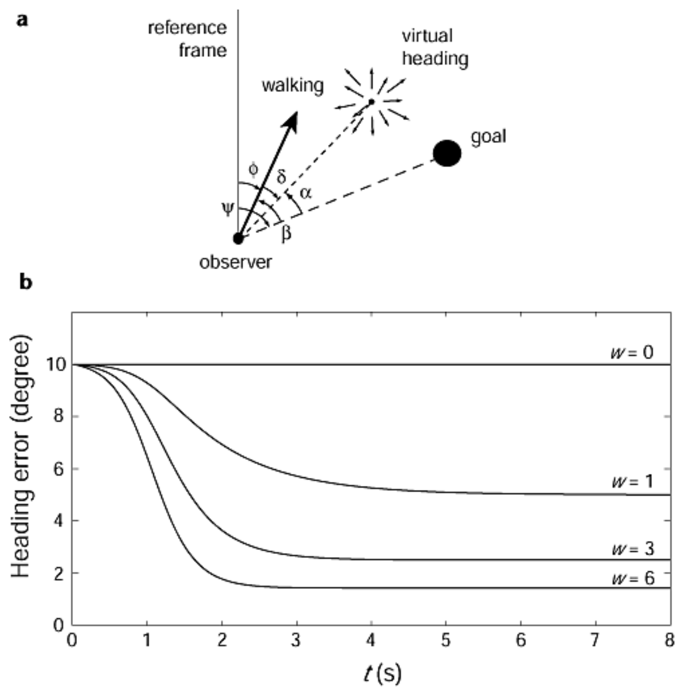

Figure 2.17: Heading error as a function of time and amount of optic flow. An image from (Warren et al., 2001).

Simulations of Equation 2.1 were consistent with results from experimentation (Warren et al., 2001). When optic flow is zero ($w = 0$) then the turning rate is controlled by β, the egocentric direction of the goal. As optic flow increases ($w > 0$), the turning rate is controlled by α, the angle between the "center of expansion" and the goal.

Based on results from (Warren et al., 2001), people use both the egocentric direction hypothesis and optic flow to guide locomotion toward a goal such that as optic flow increases, so does reliance on optic flow for locomotion. All environments used in the RFED studies presented in this dissertation have been fully textured ($w > 0$), providing optic flow throughout the environment. From Equation 2.1, $w > 0$ and therefore α, the visual angle between the "center of expansion" and the goal, controls turning rate. Based on Equation 2.1, when $w = 0$, a non-textured environment, people rely on the egocentric direction hypothesis, and will walk toward *redirected* targets. This suggests that RFED will work with textured and untextured environments. Further study of RFED altering amounts of w may provide insight about different amounts of rotation that can be added to the VE during redirection. Further research should explore the importance of textures on RFED.

2.2.4 The Dorsal and Ventral Streams: Vision for Action and Vision for Perception

Vision is responsible for two independent functions: the control of action and the "construction of conceptual representations" (Goodale and Milner, 2004). The brain has two separate visual systems, one system controls and guides action, and the other is for perception, the two separate functions of vision. (Ungerleider and Mishkin, 1982) studied the visual systems of monkeys, whose visual pathways and visual systems are similar to humans. They found that signals from the eyes first reach the *visual cortex*, a small area at the back of the cerebral cortex. The signal, the *dorsal stream*, is routed along the *dorsal visual pathway* and sent to the posterior partial region at the top of the cerebral hemisphere. The *ventral stream*, sent along the *ventral visual pathway*, is sent to the inferiotemporal region, located on the bottom and sides of the cerebral hemisphere. See Figure 2.18.

Studies that led to the theory that there are the two separate visual systems have involved observation of people with brain damage in either the posterior partial region or the inferiotemporal region. These studies have revealed that the dorsal stream controls vision for action, such as grasping objects or walking over obstacles, whereas the ventral stream controls vision for perception and the recognition of objects. People with damage to their posterior partial cortex have problems reaching and grasping for objects and people with damage to their inferotemporal cortex have problems recognizing the shapes of objects.

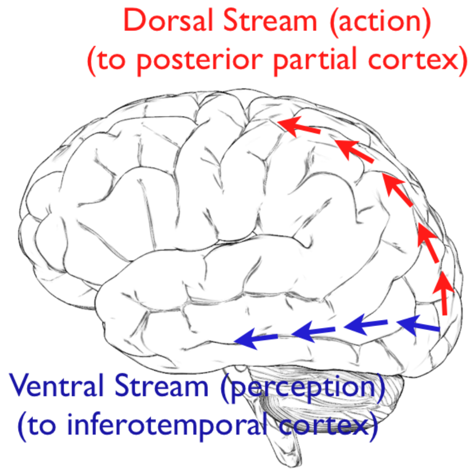

Figure 2.18: The ventral and dorsal streams. Based on an image from (Zuj et al., 2004).

The separation of the two visual streams, combined with the vestibular system's dominance over the visual system at higher frequencies (See Section 2.1.4), enables imperceptible redirection as follows: As the user turns his head the vestibular system dominates the visual system. This high-frequency head motion (greater that 0.07 Hz) enables VE rotation to be visually imperceptible to the user. The dorsal stream, the vision-for-action stream, detects that the VE has rotated causing the user to change heading direction to continue walking along the intended path. Due to the dominance of the vestibular system, the ventral stream, the vision-for-perception stream does not perceive the rotation of the VE. Imperceptible redirection would not be possible without the separation of the two visual systems separating vision for action and vision for perception.

CHAPTER 3

Virtual Locomotion

RFED is first and foremost a virtual locomotion interface for walking through large-scale virtual environments. In this chapter I define locomotion interfaces and discuss commonly used interfaces for walking through virtual environments. I present some of the challenges of using different locomotion interfaces and discuss some of the metrics used to evaluate locomotion interfaces, specifically navigation and simulator sickness.

3.1 Locomotion Interfaces

3.1.1 Introduction

Locomotion, the act of moving from one location to another, is often performed in the real world by walking, running, or crawling, or by transport on a vehicle. To enable locomotion

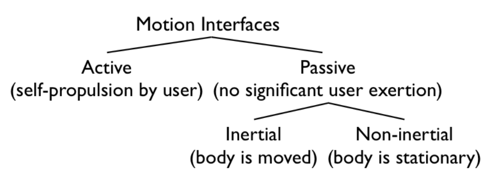

Figure 3.1: A visual representation of motion interfaces for virtual environments.

through virtual environments (VEs), a *motion interface*, the means by which a user's point of view (POV) moves from one location to another in a virtual environment, must be used. Motion interfaces enable user controlled POV within interactive VEs by mapping user actions into VE motion. Motion interfaces are divided into *active interfaces*, also known as *locomotion interfaces* where the user must use self-propulsion or *gait*, repetitive limb motion, to move through the virtual environment, and *passive interfaces*, where the user does not need to provide significant amounts of self-exertion (Durlach and Mayor, 1995). Passive motion interfaces can further be divided into *inertial*, where the user is physically moved as in flight or driving simulators with motion platforms, and *non-inertial*, where the user is not physically moved, as in joystick interfaces (Hollerbach, 2002) and vehicle interfaces without motion platforms. Figure 3.1 visually displays the relationship among different motion interfaces.

I limit the discussion of motion interfaces to those that simulate the act of walking within VEs, because RFED is a locomotion interface for walking between virtual locations. The goal of such motion interfaces is to enable people to perform tasks in the same way in a VE as they would in the real world environment. Examples of these tasks include maneuvering around obstacles, navigating (Section 3.3), or estimating the size of the environment or how far the user has traveled. Research suggests that stimulation of the kinesthetic and proprioceptive systems (Section 2.1.5) aid in representing real walking (Ruddle and Lessels, 2009).

3.1.2 Joystick

Joysticks, a form of passive motion interface, such as the one shown in Figure 3.2, are common interactive devices used to control user POV through VEs. Joysticks are commonly used because they are inexpensive, easy to implement, and can be used in small areas. Joysticks have a base resting position and the joystick is then deflected from that position. The deflected directions on the joystick are mapped to directions in the VE. Some joysticks are isometric, and do not move but sense the user's forces. In some interfaces, the joystick controls forward speed while direction is controlled by the user's body or head direction.

A user's speed is also mapped to the joystick controls. The speed is often set based on the joystick's distance from the resting center position, and the relationship is not necessarily linear. The relationship between joystick deflection and speed is often preset. Joystick controlled speeds can range from crawling, to an average walking speed, running, or even faster. User travel speed using joysticks is often not the same as real-walking travel speed.

This loose coupling between virtual travel speed and real-walking speed may be one reason why people underestimate virtual travel distance (Witmer and Kline, 1998). Additionally, joysticks are limited in the feedback they provide users about their speed or distance traveled in the environment. Motion feedback through joystick interfaces is provided through *vection*, the impression of self-motion produced by visual stimulation alone. Joysticks are also limited in that they do not require user physical exertion, which may be important for training applications. Joystick users traveling a mile in a VE exert almost no energy, unlike users using walking-like interfaces.

Figure 3.2: An xBox 360 game controller. The arrows point to the two joysticks on the controller. This joystick was used in the study described in Chapter 7.

3.1.3 Walking-in-Place

The walking-in-place (WIP) locomotion interface can be thought of as a *gestural interface* where gestures are interpreted and control user speed and heading direction. WIP interfaces use the gesture of moving the feet and knees up and down as if the user is walking in place. The "walking" gesture is detected by the system through tracking of the foot, knee, or head movement and is translated computationally into VE motion. Since the person is walking in one location, little space is needed to implement a WIP system.

User speed, in WIP systems, is controlled by step frequency and stride length, however heading direction is more complicated. Heading direction has been controlled by gaze or torso direction, or with a hand-held joystick. Gaze directed heading, although easy to implement, limits the user's ability to look around the VE while moving. Torso-directed motion requires additional tracking of the torso, but has the advantage of decoupling heading and gaze direction. A hand-held joystick provides easy steering but limits the ability to carry props. An advantage to using gaze or torso directed steering is that it requires users to physically turn, which stimulates the kinesthetic system.

WIP systems have an advantage over joystick interfaces because users of a WIP system have feedback from their kinesthetic system (Section 2.1.5). A disadvantage of WIP interfaces is that users have to wear extra equipment on their feet or legs for tracking of their steps. This encumbrance is however often considered less important than the benefit of kinesthetic stimulation.

WIP systems are promising interfaces in that physical exertion is translated into virtual motion, they stimulate the kinisthetic system, and they can be implemented in small tracking areas. However, current WIP implementations do not accurately map user starting and stopping gestures into starting and stopping virtual motions (Wendt et al., 2010). WIP systems require walking-in-place some fraction of steps before the system registers walking. WIP systems do not immediately register when the user stops walking-in-place, causing additional virtual motion after real motion stops, a disturbing lag. A fundamental reason for the stopping lag is that there are no stopping gestures to map. Also, even though WIP interfaces stimulate the kinesthetic system, there are discrepancies in the proprioceptive system because users see forward motion however they are not physically moving forward. Treadmills and omni-directional treadmills have this similar problem to WIP interfaces, however treadmills enable the correct walking muscle sensations.

3.1.4 Treadmills

Treadmills (unidirectional, circular, and omnidirectional) are used as a locomotion interface to enable users to really walk while staying physically in the same place. Users can walk in a more natural way than when using a WIP interface. The user's physical motion on the treadmill is transferred to motion in the VE. A problem with treadmills is that users can move only in the forward direction and cannot make natural turns (Darken et al., 1997). Some treadmills add turning devices such as handle bars or pressure sensors to enable virtual turning. However, turning devices do not provide kinesthetic feedback to the user that is comparable to turning in the real world.

To enable more accurate physical turning, omni-directional treadmills (ODTs) have been developed. ODTs enable users to walk in any direction while remaining in a confined location, however ODTs are loud, expensive, and diminish friction between foot and ground. Users must wear safety harnesses to stay upright. See Figure 3.3. Although treadmills mimic real walking more accurately than WIP systems, they do not provide the same proprioceptive sensations as real walking because people do not physically move forward. Users often have to re-acclimate to real walking after being on the treadmill for extended periods of time (Darken et al., 1997).

Figure 3.3: A soldier positioned on omni-directional treadmill, inside of CAVE environment. (ARL HRED, 2006)

3.1.5 Real-Walking Interfaces

A locomotion interface that enables users to really walk has to restrict the scope of the walking to the tracked space. Current locomotion interfaces that enable real walking in large-scale VEs apply transformations, such as rotating the VE, or scale user motions. However each interface has limitations.

3.1.5.1 Scaled-Translational Gain

Scaled translational gain (Robinett and Holloway, 1992; Williams et al., 2006a; Williams et al., 2006b) scales user translation. For example, if the user moves one step in the real world, the user motion is scaled to move multiple steps.

An problem with scaled-translational gain is that scaled user motion is not limited to only heading direction. As people walk their heads move side to side (Hicheur et al., 2005). By scaling all motion, the head bobs will also be scaled causing the viewpoint to sway (Interrante et al., 2007).

3.1.5.2 Seven League Boots

Seven league boots improves upon scaled-translational gain by eliminating the viewpoint sway by scaling only user motion in the intended direction of travel (Interrante et al., 2007). This still leaves head bobs in the to and fro direction. The Seven League Boots interface approximates the user's intended direction of travel by using a weighted average of gaze direction and previous displacement over a short period of time, such as two seconds (Interrante et al., 2007).

3.1.5.3 Motion Compression

Motion compression (MC) (Nitzsche et al., 2004; Su, 2007), has a misleading name because it does not compress motion. Instead, MC rotates the VE around the user and remaps areas of the VE that were outside of the tracked-space into the tracked space. The MC algorithm predicts a user's future target location based on points of interest in the VE. The algorithm then maps the straight line of the path from the user to the predicted target location into the largest possible arc that can fit into the tracked space. MC continuously updates the target location and the rotation of the VE relative to the tracked space. However, MC does not make imperceptibility of rotation a primary goal. This may increase the likelihood of simulator sickness (Section 3.2).

3.1.5.4 Redirected Walking

Redirected walking (RDW) (Razzaque et al., 2001; Razzaque et al., 2002; Razzaque, 2005) is a technique that exploits the imprecision of human perception of self-motion, the motion of humans based on sensory cues other than vision. RDW modifies the direction of the user's gaze by imperceptibly rotating the VE around the user and redirecting the user's future path back into the tracked space. Unlike MC, RDW was designed to make rotation imperceptible to the user. RDW achieves imperceptible rotation by exploiting the visual-vestibular crossover (Section 2.1.4). The vestibular system is dominant over the visual system at head frequencies greater than 0.07 Hz causing users to not perceive unmatched VE rotation while turning their heads at frequencies greater than 0.07 Hz. For this reason, an integral part of the design for RDW was to make users turn their heads frequently.

Razzaque added *waypoints* to his environments and tasks, predefined locations that defined the user's virtual route within the VE, for two reasons.

1. A series of waypoints predefined the sequence of the user's future locations. Knowledge of the user's future location enables the system to always know what part of the VE should be rotated in the tracked space.

2. Waypoints are a mechanism designed to make people look around. That is, users were required to turn their heads to find the next waypoint. This enabled RDW to rotate the VE and redirect the user's future path, the path to the next waypoint, into the tracked space.

Although waypoints enable RDW, they limit applications to those that have predefined paths and task related reasons for users to turn their heads.

3.1.6 Reorientation Techniques

Reorientation techniques (ROTs) handle the situation when large-area real-walking techniques fail and the user is close to walking out of the tracked space. Additionally, ROTs should interfere the with the virtual experience as little as possible. Each of the methods described in Section 3.1.5 uses a ROT. When users are close to walking out of the tracked space ROTs must stop the user and rotate the VE around her current virtual location. The rotation of the VE places the immediately expected user path back within the tracked space. The user must also reorient herself by physically turning around in the real environment so she can follow her desired path in the newly-rotated VE.

Redirected walking (Razzaque et al., 2001; Razzaque et al., 2002; Razzaque, 2005), in addition to waypoints, uses a ROT with a loudspeaker within the VE, played through user-worn headphones, that asks the user to stop, turn her head back and forth, and then continue walking in the same direction. Razzaque asked users to turn their heads back and forth because the user is least likely to notice extra rotation while she is turning her head.

The ROT used in motion compression (Nitzsche et al., 2004; Su, 2007) is built into the motion compression algorithm: as the user approaches the edge of the tracked space, the arc of minimum curvature is quite small and the VE rotation is large. Large VE rotation causes the user to feel that the VE is spinning around (Nitzsche et al., 2004).

Williams et al. explored three ”resetting” methods for manipulating the VE when the user nears the edge of the tracked space (Williams et al., 2007). One technique involves turning the HMD off, instructing the user to walk backwards to the middle of the lab, and then turning the HMD back on. The user will then find herself in the same place in the VE but will no longer be near the edge of the tracked space. The second technique turns the HMD off, asks the user to turn in place, and then turns the HMD back on. The user will then find herself facing the same direction in the VE, but she is facing a different direction in the tracked space. Preliminary research (Williams et al., 2007) suggests that the most promising is a third technique which uses an audio request for the user to stop and turn 360°. The VE rotates at twice the speed of the user and stops rotating after a user turn of 180°. The user is supposed to reorient herself by turning only 180° but should think she has turned 360°. This ROT attempts to trick the user into not noticing the extra rotation, however when testing this ROT I noticed that few participants were tricked into thinking they turned 360° after only turning 180°.

In Chapter 5, I introduce a new ROT, *distractors*, objects or sounds to which the user attends while the VE rotates. In Chapter 4, I discuss two distractor implementations: one using distractors as ROTs, and one using distractors to prevent the user from reaching the edge of the tracked space.

3.2 Simulator Sickness

Simulator sickness, also known as *cybersickness*, is the sensation caused by VE systems that produces symptoms similar to motion sickness, e.g. nausea, increased sweating, and dizziness. There are three common theories for the cause of simulator sickness: the *sensory conflict theory*, the *poison theory*, and the *ecological theory*.

The sensory conflict theory is the most widely accepted simulator sickness theory (Stanney, 2002), claiming that sickness results when there is a conflict between the sensory systems (Section 2.1). An example of the sensory conflicts within a VE occurs when vection is used to simulate motion. The visual system is stimulated as if the user is moving, however it may be that neither the kinesthetic nor vestibular systems are stimulated. A common argument against using *redirection* is that there will be conflict between the visual and vestibular sensory systems, since the user's physical rotations are not mapped 1-1 with the VE rotation. An argument against the sensory conflict theory is that simulator sickness is not always induced when there is a cue conflict. This has led to the development of other explanations.

The poison theory suggests that motion sickness was an evolutionary response that would remove poison from the body (Treisman, 1977). When a body ingests poison, the response is to vomit to remove the poison from the body to minimize further poison induced damiage. (Treisman, 1977) suggests that motion is an artificial stimulus that stimulates a response similar to the neurological stimuli produced by digesting poisonous toxins.

The ecological theory suggests that the interactions between an animal and the environment are critical to simulator sickness, and the longer an animal is unstable the greater the simulator sickness (Stanney, 2002). The body must work to maintain balance and if we are not balanced, we compensate to try to regain our balance. For example, in a VE, if a user is virtually moving while physically standing still, she may lean in the direction of motion because she thinks she is physically moving, thus causing herself to become unbalanced. The ecological theory helps explain why some people get sick in VEs while others do not. People who learn to "balance" in VEs by physically responding in a different way to visual stimuli can get their "sea legs" as can sailors on a moving boat.

The measurement of simulator sickness is customarily done through questionnaires. For the work in this dissertation, the Simulator Sickness Questionnaire (SSQ) (Kennedy et al., 1993) was used. The SSQ and an explanation about how to calculate a user's simulator sickness are in Appendix A.3.

3.3 Navigation

Navigation is a common task in the real world; it is what people physically and cognitively do to get from point A to point B. People navigate everyday without problem, yet in virtual worlds people often get frustrated and lost (Durlach and Mayor, 1995; Psotka, 1995; Darken and Sibert, 1996; Grant and Magee, 1998). Navigation is a fundamental task for exploring large environments and enabling people to navigate with equivalent ability in a VE and the real world should be a high priority for an interface designer. Enabling equivalent navigation in VEs and the real world will expand the usefulness and types of applications for which VEs can be used.

Designing a VE interface that enables successful user navigation requires an understanding of how people navigate in the real world. *Navigation* is cognitive and physical and is defined as the combination of *wayfinding* and *locomotion* (Darken and Peterson, 2002). *Wayfinding*, the cognitive part of navigation, is a constantly updating mental process determining how to get from one location to another. *Locomotion* is the physical part of navigation and is the act of moving between two locations. Locomotion interfaces are discussed in Section 3.1.

The cognitive part of navigation, wayfinding, is associated with building and exploiting a mental map (Darken and Peterson, 2002). When people explore environments, real or virtual, they build a mental map of the environment. Mental map building is dependent on many factors and is related to the user's natural ability to navigate. When people locomote around an environment, they begin building a mental map. The fidelity of the mental map improves over time and through multiple exposures to an environment. However, even in the real world, people often develop inaccurate mental maps of the environment even after years of exposure (Waller et al., 1998).

Mental maps are also built using cues besides exploring the environment. Maps are a common tool used everyday in real and virtual worlds for navigational aid. The problem with maps is that people often are unable to build an orientation-independent mental map of the environment when using a "real" map as an aid (Darken and Sibert, 1993). For people to accurately follow maps, especially displayed in VEs, the map must be displayed using the *forward-up* principle (Darken and Sibert, 1993)—the map must be placed such that the forward direction of the user is displayed up when the user is holding the map perpendicular to the floor. Additionally, people who learn the layout of an environment from a map often have a hard time when entering the environment from any direction other than that which was presented to them as the forward-up direction of the map. This implies that people have a hard time developing an orientation-independent mental map of an environment when using only a map (Darken and Sibert, 1993).

Common aids used for VE navigation include maps, gridlines, and landmarks. When displaying a virtual map to a user, the map must continuously follow the forward-up principle and the location of the user should always be displayed (Darken and Sibert, 1993). Although navigational aids are helpful for VE navigation, they do not always accurately simulate the real world. A common design goal for VEs is to emulate the real world as closely as possible, including real world navigational ability without requiring arbitrary landmarks, maps, or gridlines to aid VE navigation.

Evaluation of a locomotion interface's effect on navigational ability is important for VE development. Previous research from (Ruddle and Lessels, 2009) suggests that real walking participants navigate significantly better than participants using joystick interfaces in a tracked-space-size VE. However, to date, I know of no evaluation that studies user navigational ability in any of larger-than-tracked-space, real-walking locomotion interfaces such as the ones presented in Section 3.1.5. Studies to evaluate navigational ability of a larger-than-tracked-space, real-walking interface is the subject of Chapters 6 and 7.

CHAPTER 4

Redirected Free Exploration with Distractors

This chapter describes the Redirected Free Exploration with Distractors (RFED) algorithm that was developed to demonstrate how people can freely walk in VEs that are larger than the tracked space. Any RFED system is composed of two parts, *redirection* and *distractors*. This chapter discusses a generic RFED system, current design challenges, and specific implementations of the complete RFED systems used in Chapters 6 and 7. This chapter should be a reference for future designs and implementations of RFED systems.

4.1 Introduction

For users to walk normally around a VE that is larger than the tracked space, the VE must be remapped to the physical space by translating, rotating, scaling, or skewing. Transforming the VE remaps VE regions that were out of the tracked space into the tracked space, thus enabling users to really walk to new regions in the VE. RFED, based on Razzaque's Redirected Walking system (Razzaque, 2005), uses redirection—imperceptibly rotating the VE model around the user.

Razzaque demonstrated that users can be imperceptibly redirected very little unless the redirection is performed while the head is turning and hence the visual system is desensitized. He found that quite large amounts of redirection could be imperceptibly accomplished during such head turns.

To force head-turns, Razzaque's system used prescribed paths through the VE. The principle aim of this work is to remove this severe limitation and enable users to walk

freely about in a VE.

Free walking in VEs raises a new problem—how to ensure that the user avoids real-space obstacles. A most important special case is the obstacle consisting of a boundary of the tracked space. This dissertation handles the general rectilinear boundary case, however future RFED implementations could handle any boundary shape, or moving obstacles such as people.

The RFED algorithm presented here enables free walking as follows:

1. At each frame, predict the user's real-space future direction.
2. Steer the user's future direction to walk toward the center of the tracked-space.
3. If the user is near the boundary, introduce a distractor to:
   (a) Stop the user.
   (b) Force a head-turn, enabling large amounts of redirection.
   (c) Redirect the user's future direction to the center of the tracked space.

Distractors enable free-exploration, however they are a distraction and therefore inherently impair the VE experience. Hence, one wants to minimize the total number of distractors. To minimize the number of distractors requires redirecting the user away from the boundary of the tracked space. For this implementation of RFED I redirect the user through the center of the tracked space to ensure the longest path between the user and the tracker boundary.

To steer the user to the center of the tracked space requires predicting the user's real future direction and then redirecting the user's future direction toward the center of the tracked space. In this steer-to-center implementation, if there were no limits to the amount of redirection that could occur at any instant, and users are always steered directly toward the center of the tracked space, they will never reach a boundary and a distractor will never be introduced.

However, based on the human visual and vestibular systems there is a limit to the amount of redirection that can occur at any instant. To limit the number of distractions,

Table 4.1: Efficient Redirection Variables

Variables	
θ_{VE}	instantaneous magnitude and direction of VE rotation
v_{future}	user's immediate future direction
s'	a real-world location toward which the user is steered
v'_s	vector from user's real location to s'
θ_{ideal}	ideal magnitude and direction of VE rotation to redirect user toward s'
Note: a transformation consisting of a rotation and translation can convert between the real world (') and the virtual world.	

one wants to redirect as quickly as possible to steer the user away from the boundary. Due to instantaneous redirection being limited, *efficient redirection*, redirecting as quickly as possible, can be thought of as:

- Minimizing total VE rotation used to redirect the user away from the boundary.
- Maximizing instantaneous redirection to quickly redirect the user away from the boundary.

Efficient redirection is designed to steer the user away from the tracker boundary, however when redirection fails a distractor must be used to prevent the user from leaving the tracked space and force user head-turns to enable large amounts of redirect to steer the user back toward the center of the tracked space.

The remainder of this chapter discusses specific RFED implementations.

4.2 Efficient Redirection

Redirection is the two-dimensional (2D) rotation of the virtual model around the user's real 2D-location in the tracked space. At any instance, redirection is a transformation of the VE consisting of a rotation and translation about the user's current location. The redirection transformation is updated **each frame** to redirect the user away from the tracker boundary as quickly as possible, i.e. the user is efficiently redirected. From above, efficient redirection requires minimizing the total VE rotation to steer the user to stay within the tracked space, while maximizing the instantaneous per frame VE rotation. For **each frame**, efficient redirection requires answering two questions:

1. What direction (clockwise (-) or counter-clockwise (+)) should the VE rotate to minimize the total amount of redirection required to steer the user to stay within the tracked space?

2. What is the maximum amount of VE rotation that should be added to this frame?

Answering these two questions determines θ_{VE}, the magnitude and direction of VE rotation to add to the current frame. θ_{VE} is calculated through six steps that are visually illustrated in Figure 4.1. As an aid to the reader, a reference of variables used in following steps is in Table 4.1.

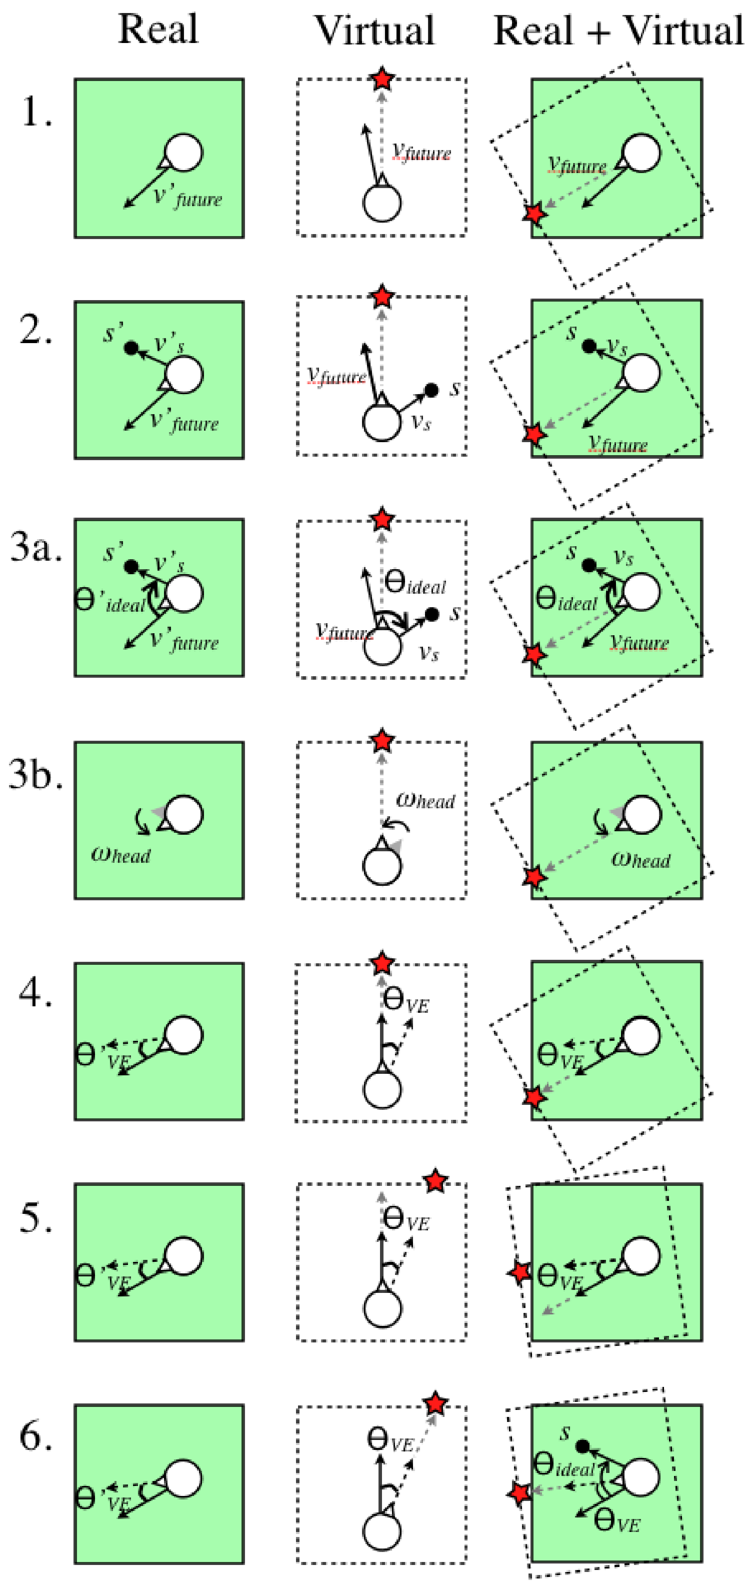

Figure 4.1: The six steps of *efficient redirection*. These steps are discussed in Section 4.2. The star is a virtual reference point. Notice that the star moves with the VE in **Step 5**.

What direction should the VE rotate?

Step 1. Predict the user's immediate future virtual direction v_{future}. In the VE, the user will walk in the direction of v_{future}. Rotating the VE redirects v'_{future}, the user's real-world future direction, enabling steering the user within the real space. An inaccurate prediction of v_{future} inhibits steering thus reducing RFEDs ability to steer the user away from the tracker boundary. I discuss algorithms for predicting v_{future} in Section 4.2.1.

Step 2. Steer v'_{future} toward a real-world steer-to point s'. I defined s' as the center of the tracked space to always direct the user toward the longest path from the user to a boundary. Define a vector v'_s from the user's real location to s'. To steer, or redirect, the user toward s', v_{future} must be rotated to be the same direction as v'_s. Note: the steer-to point s' does not need to be a single location; I discuss my steering algorithm further in Section 4.2.2.

Step 3a. Calculate the minimum angle from v'_{future} to v'_s, the total ideal signed rotation angle of the VE θ'_{ideal}. Rotating the VE by θ'_{ideal}, will rotate the VE such that v'_{future} is in the direction of s'. θ'_{ideal} determines the direction and maximum value of θ'_{VE}.

What is the maximum amount the VE should rotate?

Step 3b. Calculate the angular velocity of the user's head since the previous frame, ω'_{head}. Based on the visual and vestibular systems, (Section 2.1.4), rotation of the VE will be least noticed during head turns.

Step 4. Calculate θ_{VE}, as a function of ω_{head} and θ_{ideal}. My algorithm for calculating θ_{VE} is discussed in Section 4.2.2.

Step 5. Rotate the VE by θ_{VE}.

Step 6. The user physically rotates by θ_{VE}, so as to walk in the direction of s'. Notice that for the current s', θ_{ideal} is now smaller.

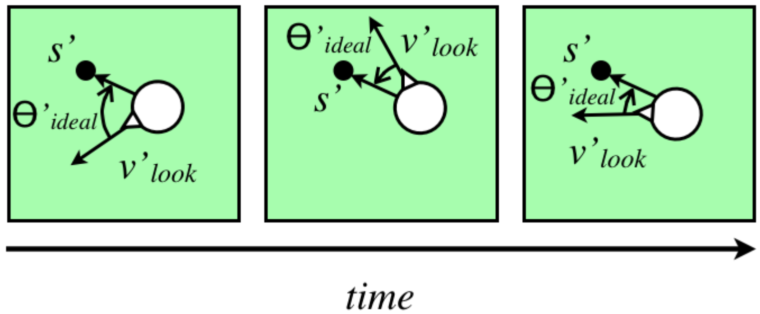

Figure 4.2: Direction prediction is unstable when the user look direction, v_{look} is used for v_{future}. As a user turns her head back and forth, the look direction, v'_{look} quickly changes to the right and left of s', changing the rotation direction of θ_{VE} from positive and negative.

4.2.1 Direction Prediction

As stated above, prediction of the user's immediate future direction, v_{future} is essential in determining θ_{VE} and is one of the hardest parts of redirection. Knowing the user's immediate future direction enables minimization of the VE rotation that rotates the VE to steer v_{future} into the tracked space.

In this section, I describe the development of the direction prediction algorithms used in the studies presented in Chapters 6 and 7. Although the final algorithms enabled every participant to successfully complete the user studies, participants frequently had to be prevented from leaving the tracked space by distractors. Improving the direction prediction algorithm could reduce the number of distractors, and thus improve RFED usability.

Basic Direction Prediction, Version 1. I first implemented basic direction prediction by defining θ_{future} to always be the user's look direction, θ_{look}, reported in the tracker data of the user's head. This implementation of path prediction was based on results from (Hollands et al., 2002) suggesting that gaze direction and heading direction are the same approximately 70% of the time. The problem with this simple path prediction model arose when people turned their heads, quickly changing v_{look}, and thus changing θ_{ideal}. The rapid change in v_{look} changed the direction of θ_{ideal} if v_{look} moved to the right and left of s' (Figure 4.2). The rapid direction change of θ_{ideal} caused θ_{VE} to continually change between clockwise and counter-clockwise VE rotation. If the VE rotates clockwise one frame and then counter-clockwise the next frame (or vise-versa), the redirection cancels itself out, thus hindering redirection of the VE.

Averaging v_{look} over time to predict v_{future} could be an alternative to using instantaneous look direction. However, one wants most redirection to occur during head turns. The average of v_{look} changes the most during head turns because v_{look} is continuously changing, and thus the average of v_{look} during head turns is potentially the least accurate prediction of v_{future}. Since the most redirection occurs when people are turning their heads and the average of v_{look} is the least accurate prediction of v_{future} during head turns, averaging v_{look} is unlikely to provide a usable prediction of v_{future}.

Direction Prediction, Version 2. (Used in Chapter 6). To improve upon approximation of v_{future}, I assume that people continue walking in the same direction in the virtual environment. Measuring direction in the VE, not the real world, is important because redirection causes people to walk curved paths in the real world. Since walk direction must be calculated as the difference between user locations over time, it is not an instantaneous measure like look direction. I calculated the average direction of the user, v_{user} as the average of the difference in user's previous virtual locations between frames. Pilot studies guided the selection of parameters for this method. I sample the user's 2D virtual location l_i every 30 frames (29 differences), and using the most recent samples $\{l_0, l_1, ..., l_{29}\}$ I calculate v_{user} such that:

$$v_{user} = \frac{\sum_{i=0}^{28} l_{i+1} - l_i}{29} \tag{4.1}$$

The user location samples are gathered over approximately 3.5 seconds (approximately 250 frames are computed per second). Further research may suggest a different averaging scheme that may produce better results.

The direction prediction algorithm used in the Chapter 6 study defined $v_{future} = v_{user}$ as calculated from Equation 4.1. This direction prediction implementation enabled all participants to complete the experiment and successfully walk freely in a VE that was larger than the tracked space. However, v_{user} did not provide flawless results. Problems occurred when people reached the edge of the tracked space and were unable to continue moving in the virtual environment since v_{user} was calculated from the average direction the user had moved, sampled over 3.5 second. Participants were able to successfully complete the experiment because the distractor implementation required participants to step backward when they reached the boundary. The backward step required participants to take a step forward to continue walking in the VE. This forward step prevented participants from standing in one place, thus enabling direction prediction.

Direction Prediction Version 3. (Used in Chapter 7). I used information from the VE to improve upon using v_{user} to predict v_{future}. The final direction prediction algorithm, used in the study presented in Chapter 7 is graphically presented in Figure 4.3 and discussed in the following paragraphs.

Step 1. Define a bidirected graph over the VE such that nodes are locations in the environment where people may change direction, and edges are straight paths in the environment. Specifically, for the maze environments used in Chapter 7, I defined the edges of the graph as hallways, and nodes as intersections and dead-ends of hallways. A grid can be used as a generic graph of any environment. Defining a graph over the VE does not restrict user movement because the user is not required to walk to nodes and the user can change directions.

Step 2. Determine the node, p in the virtual space nearest to the user. The user should walk in the direction of one of the nodes connected to p. In figure 4.3, the user will walk to p_a or p_b.

Step 3. Define vectors v_i from the user's virtual location to each node connected to p. In Figure 4.3, define vectors v_a and v_b.

Step 4. Calculate the angles between v_{user}, the average direction of the user as calculated in Equation 4.1, and each v_i. Assume that the user will walk toward the node that has the smallest angle between v_{user} and v_i. In Figure 4.3, the user will walk toward node p_a, and

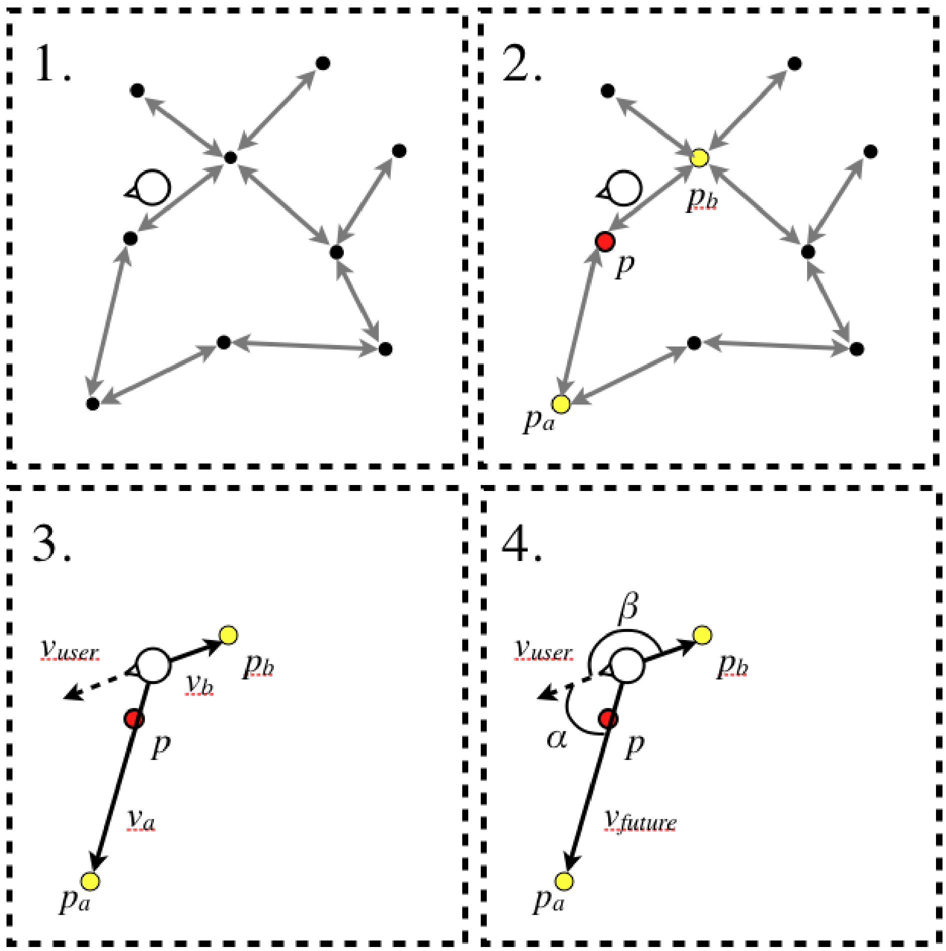

Figure 4.3: **Step 1.** Define a bidirected graph over the VE. **Step 2.** Identify the node closest to the user (p), and the nodes connected to p, (p_a and p_b). **Step 3.** Define vectors v_a and v_b from the user to connected nodes p_a and p_b. **Step 4.** Calculate and compare the angles α and β between the user direction vector, v_{user} (Equation 4.1) and vectors v_a and v_b. Since α is smaller than β, set v_a as v_{future}.

$v_{future} = v_a$

Although this algorithm produced better results than using the average direction v_{user} to predict v_{future}, there were stability problems when p, the node closest to the user changed back and forth between the same two nodes. That is, when the user was located directly between two nodes causing v_{future} to change between different nodes and possibly change the direction of θ_{VE}.

I recommend that future developers focus on improving direction prediction. Improvements to this algorithm may include incorporating robotic path planning algorithms and determining a better way to define the directed graph of the environment. Another improvement to direction prediction may include a better estimation of v_{future} by averaging over different times or incorporating the look direction of the user. Work by (Wendt, 2010) is also promising and may improve path prediction by extracting the user's instantaneous future direction from head-bob information.

4.2.2 Steering

The next step in redirection is "steering" v_{future}, as predicted above, to a "steer-to" direction vector v'_s. A *steering algorithm* is an algorithm designed to steer the user to stay within the tracked space (Razzaque, 2005). Based on the redirection algorithm in Section 4.2, steering determines the direction and maximum angular rotation of the VE to steer v_{future} toward v'_s, a "steer-to" direction vector. The direction of θ_{VE} is calculated by the steering algorithm.

I use a steer-to-center algorithm. In the redirection algorithm described in Section 4.2, I always define the steer-to-point, s', to be the center of the tracked space and calculate the direction and angle of the shortest arc between v'_{future} and v'_s (Figure 4.1, 3a.) In theory, this algorithm should always steer the user through the longest path across the lab. However, evaluation of different steering algorithms such as steer-to-circle, or steer-to-moving target (Razzaque, 2005), may produce better results.

After determining the sign of rotation of v_{future} toward the center of the tracked space, I calculate the maximum magnitude of the rotation based on ω_{head}, the user's angular head speed. The faster the user turns his head, the less aware he will be of VE rotation. Therefore, I rotate the VE by θ_{VE}, where θ_{VE} is a function of how fast the user has turned her head and a predefined rotation constant, c.

$$\theta_{VE} = |\omega_{head}| * c \tag{4.2}$$

That is, for each frame, the faster the user turns her head, the greater the rotation of the VE. Based on pilot experiments and research from (Jerald et al., 2008), I chose c to be 0.10 when the VE is rotating in the same direction as ω_{head} and 0.05 when the VE rotates against ω_{head}. Since the ideal amount of rotation is θ_{ideal}, if $\theta_{VE} > \theta_{ideal}$ then set $\theta_{VE} = \theta_{ideal}$.

I have defined θ_{VE} as a linear function of ω_{head} and capped its value based on θ_{ideal}. Further evaluation of user perception of VE motion during head turns, and determining the maximum amount of acceptable perceptual VE rotation will determine a maximum value for θ_{VE}.

In summary, for each frame the path prediction algorithm predicts the user's future direction v_{future} and the steering algorithm determines θ_{ideal}, the direction and maximum rotation of the VE to steer v_{future} to a steer-to point s'. The instantaneous rotation of the VE, θ_{VE} is a function of θ_{ideal} and ω_{head}, the instantaneous angular head speed of the user. Over several frames the VE is rotated by θ_{ideal} around the user to redirect the user to stay within the tracked space. Although every attempt is made to steer the user away from the edges of the tracked space, redirection does not always work and a reorientation technique (ROT) must be used to steer the user back into the tracked space.

4.3 Distractors

Distractor implementation involves continually answering two questions:

1. **Turn distractor on:** When should the distractor appear?
2. **Turn distractor off:** When should the distractor disappear?

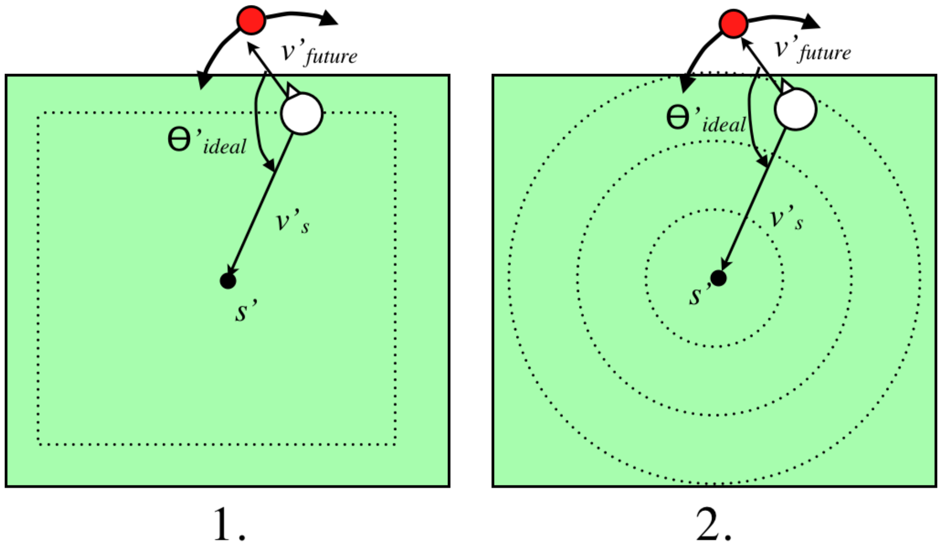

Figure 4.4: **Version 1.** Turn distractor on: The user is stopped by a distractor when he crosses a boundary near the edge of the tracked space. Turn distractor off: When the VE has rotated by θ'_{ideal} around the user. **Version 2.** Turn distractor on: The distractor appears based on a function of the user's distance to the center of the tracked space and the time since the previous distractor. Turn distractor off: When the VE has rotated a fraction of θ_{ideal} based on the distance of the user from the center of the tracked space.

I implemented and ran two distractor algorithms.

Distractor algorithm, Version 1 (Used in Chapter 6). The first distractor implementation (Figure 4.4 Version 1.) was used in the experiment discussed in Chapter 6. For this implementation distractors were only used to stop the user from leaving the tracked space and to redirect the user's path back into the tracked space. For this implementation I defined a border $1m$ from the edge of the tracked space. **Turn distractor on:** When the user crossed the border a distractor would appear to stop the user. Since the user was at the edge of the tracked space I wanted to redirect the user to the center of the tracked space. That is, **Turn distractor off:** when the VE has rotated θ'_{ideal} to within ϵ of v'_s, where $\epsilon = 0.005°$, and v'_s is the vector to the center of the tracked space.

A problem with this implementation was that the **turn distractor off** condition required reorienting the user all of θ_{ideal}. The greater the reorientation amount the longer the reorientation and some reorientations took as long as 30 seconds. Users often reported increased frustration if the distractor did not disappear within 5 seconds after appearing.

Table 4.2: Turn distractor on values for distance and time

Distance from center (d)	Time since previous distractor (t)	percentage of θ_{ideal}
$< 1.5m$	$> 40s$	0.67
$< 2m$	$> 35s$	0.71
$< 2.5m$	$> 15s$	0.77
$< 3m$	$> 5s$	1.0
$< 3.25m$	$> 3s$	1.0

Also, after the distractor disappeared, the user had to reorient her body by θ_{ideal} (Figure 4.1, Step 6.). For large VE rotations, participants were likely to notice that the VE had rotated after the VE stopped rotating and after they stopped turning their heads. For large θ_{ideal} people were more likely to notice that the VE had rotated when they stopped rotating their heads and had to physically reorient themselves within the VE.

A third problem was that **turn distractor on** was triggered by the user being within any location in the boundary region. If people did not step out of the boundary before the distractor disappeared, the distractor would reappear again and again. To address this problem, participants were asked to take one step backwards. Since people were being redirected, the backwards step was not always toward the center of the lab. Occasionally the experimenter had to physically guide the participant toward the center of the lab to make sure the distractors would not keep reappearing.

Distractor algorithm, Version 2 (Used in Chapter 7). I implemented a second distractor algorithm for the study presented in Chapter 7 (Figure 4.4 Version 2.). For this design, distractors appear before participants reach the boundary area of the tracked-space so as to reduce the number of distractions. The **turn distractor on** condition is a function of t, the time since the previous distractor appeared, and d, the distance of the participant from the center of the lab. The inversely related values of d and t turn the distractor on. For example, when the participant is near the edge of the tracked space, a large d, and a distractor has not appeared within t, a small t, a distractor will appear. If the user is near the center of the tracked space, small d, then a distractor will appear only if one has not recently appeared, a large t.

The **turn distractor off** condition is triggered after the VE rotates a percentage of θ_{ideal} based on d, i.e. the closer the participant is to the center of the lab, the smaller the percentage of θ_{ideal} the VE rotates before the distractor disappears. The values for t, d, and the percentage of θ_{ideal} are in Table 4.2. For example, from Table 4.2, if the participant is 2.4 meters from the center of the tracked space and a distractor has not appeared in 16 seconds, the distractor will turn off after the VE rotates by $0.77\theta_{ideal}$.

Overall, this algorithm was promising because participants were kept from the edges of the lab, however distractors still appeared frequently and many participants complained about the over-abundance of distractions. For anyone implementing a distractor algorithm, I would recommend using distractors to steer people away from the edge of the tracked space, but I would look into developing less intrusive distractors that fit seamlessly into the environment. For example, if developing RFED for real estate applications, add a small dog, or children playing within the house. These objects may catch the participant's attention and cause her to turn her head, enabling redirection.

Results from the studies discussed in Chapter 5 suggests that people are less aware of VE rotation, i.e. redirection, when a distractor is used. From this result, I piloted increasing the rotation constant c from Equation 4.2 when distractors are visible. The values of c used in the experiment in Chapter 7 where 0.60 when with head rotation and 0.30 against head rotation. Participants were instructed that the VE would rotate around them during the experiment. Even though the rotation values were extremely high, participants commented that they noticed VE rotation only after they stopped moving their heads and had to reorient their bodies to continue walking in the same direction in the VE. Although some participants noticed that the VE had moved, no participant had problems reorienting to the rotated VE.

4.4 Deterrents

In the version 2 distractor algorithm, a distractor appeared every 3 seconds when the user was near the edge of the tracked space. If the user stayed near the boundary distractors continually reappeared and frustrated users. To guide the participant away from the bound-

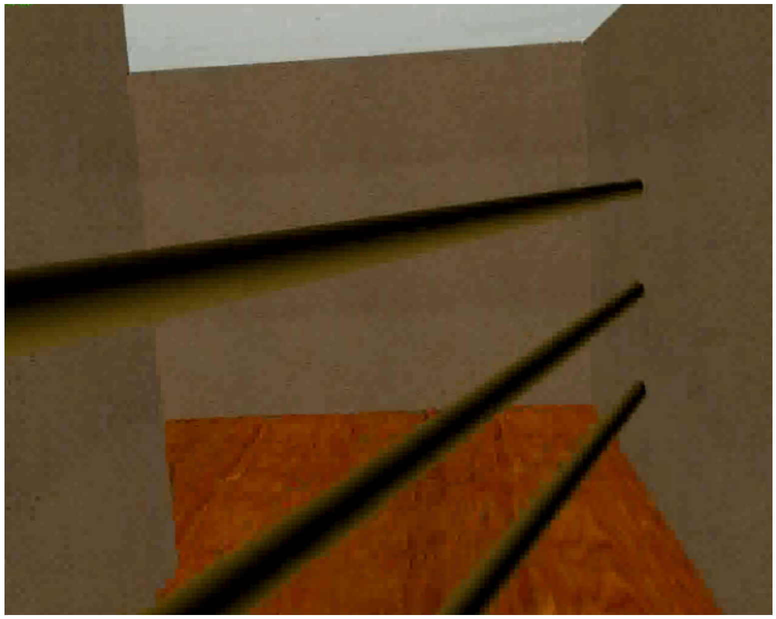

Figure 4.5: A screen shot of the horizontal bars used as deterrents in the version 2 distractor algorithm.

ary, I added deterrents to the environment. *Deterrents* are objects in the environment that people are instructed to stay away from or not to cross. For this implementation, deterrents were virtual horizontal bars that were aligned with the edge of the tracked space. See Figure 4.5. The bars fade in as the user nears the boundary of the tracked space and fade out as the user walks away from the boundary.

The virtual bars provided participants with a visual cue as to which direction to walk to stay in the real space. No participant complained about the bars.

When the bars appear to inform the user where they cannot go in the real space, they also provide the user with a visual cue about the size of the VE and the orientation of the tracked-space in relation to the participant. I originally though this would cause problems with participants, however people often noted that the bars appeared to be the "virtual" part of the virtual environment. Also, since the bars always mark the location of the stationary tracked space they do not move. If the VE rotates while the bars are visible people will quickly notice the VE rotation, even though it appears that the bars are rotating, not the VE. For this reason, I did not reorient the VE while stationary deterrents were in view.

CHAPTER 5

Distractors

This chapter introduces and evaluates *distractors*, a new reorientation technique (ROT) that was developed to enable people to freely explore virtual environments (VEs) without causing breaks in presence. ROTs stop the user and rotate the VE around her current virtual location, placing the predicted user path back within the tracked space. Distractors are objects in the VE that the user attends, causing the user to turn her head, minimizing the observed rotation of a VE during reorientation.

In this chapter I present three formative user studies, where I iteratively implemented, tested, and improved distractors by comparing them to current ROTs. The first two studies were presented in (Peck et al., 2008) and the final user study was presented in (Peck et al., 2009).

5.1 Introduction

Virtual Environments (VEs) using real-walking locomotion interfaces have typically been restricted in size to the area of the tracked lab space. Techniques have been proposed to lift this size constraint, enabling real walking in VEs that are larger than the tracked space (Nitzsche et al., 2004; Razzaque et al., 2001; Razzaque et al., 2002; Razzaque, 2005; Su, 2007; Williams et al., 2006a; Williams et al., 2006b). For free exploration in each of these large-area walking VE methods relies on a reorientation technique (ROT) to handle the case when the technique fails and the user is close to walking out of the tracked space. When such an event happens, ROTs must stop the user and rotate the VE around her current virtual location, placing the immediately expected user path back within the tracked space. The user must also physically reorient herself by turning around in the real environment so

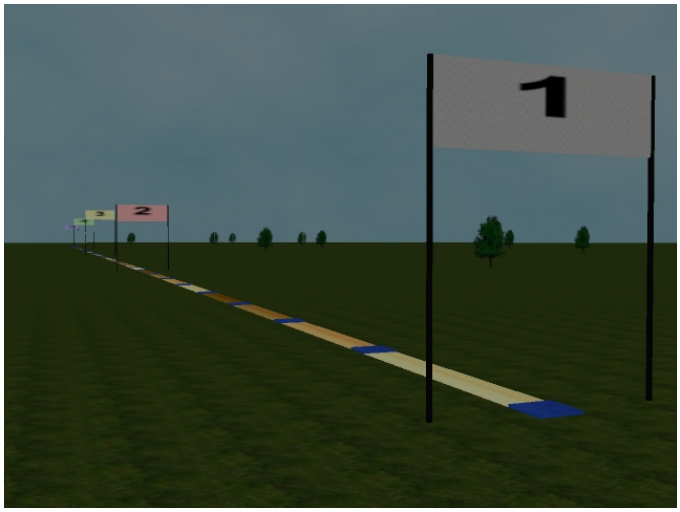

Figure 5.1: Virtual Environment used in Experiments 2 and 3.

she can follow her desired path in the newly-rotated VE.

Unlike staying-in-place interfaces like joysticks, walking-in-place, or treadmills (Brooks, 1987; Christensen et al., 2000; Darken et al., 1997; Iwata, 1999; Slater et al., 1995a), ROTs are required by real-walking locomotion interfaces to enable free exploration of large VEs. I hypothesized that current ROT implementations cause breaks in presence, which detract from the immersive VE experience. In this chapter I introduce a new ROT, distractors, and compare this method to existing ROTs in three user studies. The distractor method was modified between user studies based on participant feedback. Evaluation each ROT is measured by *presence*, user-ranked *preference*, and user-ranked *naturalness*.

My method introduces the concept of a *distractor*, an object, sound, or combination of object and sound in the VE that the user attends while the VE rotates. Distractors reduce perception of VE rotation, and thus reduce the likelihood of a break in presence. In the three studies I compare my new distractor technique to previously reported techniques.

Previously implemented ROTs are discussed in Section 3.1.6.

5.2 Overview of User Studies

The goal of Experiment 1 was to determine the best way to prevent people from leaving the tracked-space when really walking in larger-than-tracked-space VEs. I hypothesized that

Figure 5.2: Laboratory Layout used in Experiments 2 and 3.

current ROTs suggested or implemented by (Nitzsche et al., 2004; Razzaque et al., 2001; Razzaque et al., 2002; Razzaque, 2005; Su, 2007; Williams et al., 2007) caused breaks in presence and therefore developed distractors. I compared distractors to current ROTs based on user subjective sense of *presence*, user-ranked *preference*, and user-ranked *naturalness*. The results suggested that distractors and the ROT suggested by (Razzaque, 2005) produced increased presence, had higher user preference and were more natural to the user.

Experiment 1 user feedback suggested improvements for distractors, including making the distractor more natural to the environment and moving the distractor at a slower speed. I implemented the participant suggested distractor improvements in Experiment 2 and reevaluated the improved distractor by comparing it to the most promising ROTs from Experiment 1, the original distractor technique and the ROT suggested by (Razzaque, 2005). The results from Experiment 2 suggest that ROTs that use distractors reduce the likeness of a users' feeling as if they are turning around while being reoriented, and that participants prefer ROTs with distractors and consider them to be more natural.

Based on user feedback from Experiment 2, I improved the distractor method by using a more realistic model: a hummingbird (Figure 5.16). I also explored adding sound to the visual distractor and using sound alone as a distractor. The results from Experiment 3 suggest that using more realistic distractors can increase a user's feeling of presence and that audio alone can be used as a distractor.

5.2.1 Equipment

Each participant wore a Virtual Research Systems V8 head-mounted display with 60° diagonal FOV (640 x 480 resolution) tracked using a 3rdTech HiBall 3000. Participants were permitted to walk in an 8m x 6m tracked space. The environment used in experiments 2 and 3 is shown in Figure 5.1. A similar environment was used in experiment 1. All environments were rendered in stereo at 60 fps in each eye on a Pentium D dual-core 2.8GHz processor machine with an NVIDIA GeForce 6800 GPU with 2GB of RAM. The cardboard taped to the wooden surface, visible in Figure 5.2, was slightly padded and gave users, who had no self-avatar, haptic confirmation of reaching the markers on the paths. No haptic feedback was given to participants in experiment 1.

5.2.2 Experiment 1

The first study evaluated the ROTs suggested or implemented by (Nitzsche et al., 2004; Razzaque et al., 2001; Razzaque et al., 2002; Razzaque, 2005; Su, 2007; Williams et al., 2007) plus the distractor technique. The measures were *presence*, user-ranked *preference*, and user-ranked *naturalness*.

5.2.2.1 Participants

Twenty-four introductory psychology students (13 men and 11 women) participated in the experiment. Each participant visited the laboratory once for a session lasting approximately 1 hour and received class credit for participation. All participants had normal or corrected to normal vision and were naive to the purpose of the study. Participants were not informed about ROTs and were initially unaware that the VE would rotate.

5.2.2.2 Experimental Design

Experiment 1 consisted of two parts, both taking place in the same VE. The VE was an outdoor space featuring a 200-meter straight wooden path with circular markers placed 5 meters apart along the path, similar to the experiment 2 environment (Figure 5.1). Participants received audio instructions via head phones before the experiment began and received audio trial-specific instructions before each trial. Trial specific instructions included informing participants to physically turn, turn your head back and forth, or watch the distractor. Participants did not have a training session, and no participant had problems performing the experiment.

All VE rotation rates were determined from pilot experiments. For all reorientation techniques, except one, the rotation of the virtual environment is increased only when the head is turning in the same direction that the VE is rotating.

Participants were instructed to really-walk along the virtual path and to stop on each marker. When the participant reached a marker, they were reoriented (the VE rotated) 180° by one of the following ROTs, and then walked back across the lab to the next marker.

Turn without instruction (T). When the user reaches the marker the VE immediately rotates 180° around the user at 120°/second. The rotation relocates the virtual path so it is located within the tracked environment. The user will have reoriented herself by turning only 180° in the real world. This is similar to the technique described in (Nitzsche et al., 2004; Su, 2007).

Turn with audio instruction (TI). Audio instructions in the VE, presented via headphones, ask the user to turn 360° and continue along the path; however, the VE rotates 180°. The rotation of the VE is controlled by the user's head and rotates at twice the speed of the user's head. The user is deceived to think that she has turned 360° in both the virtual and real worlds when she has only turned 180° in the real world. The user needs to reorient herself in the VE by turning only 180°. This is similar to a method described in (Williams et al., 2007).

Head turn with audio instruction (HT). Audio instructions in the VE ask the user to turn her head back and forth and then continue walking along the path. While the user turns her head the rotation applied to the VE is 1.3 times the rotation speed of the participant's head, determined from pilot experiments, the rotation speed of the user's head until the VE has rotated 180°. The participant reorients herself by rotating 180° in the real world. This is similar to a method described in (Razzaque, 2005).

Head turn with visual instruction, distractor (D). A moving sphere appears in front of the user. The user watches the sphere as it moves in a horizontal arc and continues walking along the path once the sphere disappears. The rotation applied to the VE is 1.5 times the rotation speed of the participant's head, determined from pilot experiments, the rotation speed of the user's head until the VE has rotated 180°. The distractor moves along an arc with a 1.75 meter radius, with sinusoidal speed, amplitude = 180° and frequency = 4.5 HZ. The user reorients herself by rotating 180° in the real world. The path and velocity of the distractor are described in Figure 5.3. A general distractor is any object in the VE that distracts the user.

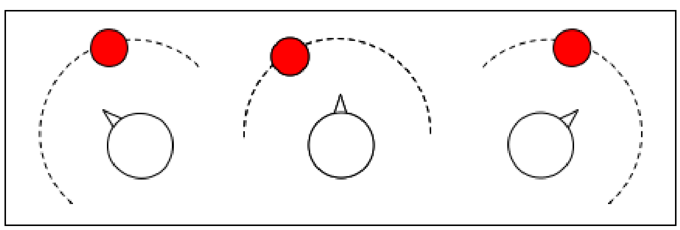

Figure 5.3: The path of all distractors is defined as an arc (dashed line) directly in front of the user. The distractor moves with sinusoidal displacement along the arc causing the participant to turn her head back and forth to keep the distractor in view. The distractor is displayed 1.75 meters away from the user, and the height of all distractors is approximately 1.5 meters. The same path trajectory is used for all distractors in each of the three experiments presented in this chapter.

D - Distractor
HT - Head turn with audio instruction
TI - Turn with audio instruction
T - Turn without instruction

Figure 5.4: Experiment 1–Legend

Table 5.1: Experiment 1 - Mean HIGH scores on SUS Presence Questionnaire

ROT	$\bar{x}$
D	0.47917
HT	0.50000
TI	0.28472
T	0.44444

Part I of the experiment assessed the user's subjective sense of presence in the environment and consisted of four trials, each using one of the four reorientation techniques. The order of the trials was counterbalanced among participants. Each trial was comprised of four sub-trials in which the participant walked along the virtual path and stopped at markers along the path. When the participant reached a marker, an ROT would stop the participant and rotate the VE. Each trial consisted of walking to four markers and experiencing the same reorientation technique four times. Participants then removed the HMD and filled out a modified Slater-Usoh-Steed (SUS) presence questionnaire (Slater and Usoh, 1993; Slater et al., 1994).

Part II consisted of 12 trials, each with two reorientation techniques. Trials were counterbalanced and every ROT was compared to every other ROT twice, with order reversed, to remove the possibility of order effects. Each trial required the participant to walk to a marker, experience an ROT, then walk to the next marker, and experience a different ROT. The participant then made a forced choice regarding which ROT they preferred and which ROT was more natural. At the end of each trial participants were asked by the experimenter to explain why they chose one ROT over another.

At the end of the experiment participants filled out an exit survey and were asked to describe the differences between the four ROTs, explain what they liked or disliked about each of the ROTs, and rank the four ROTs based on naturalness and preference.

We used a modified SUS presence questionnaire (Slater and Usoh, 1993; Slater et al., 1994) to assess the user's participantive sense of presence. Naturalness and preference were each measured in two ways: at the end of the experiment participants ranked the ROTs, and during the experiment participants made a forced-choice ranking between pairs of ROTs.

Table 5.2: Experiment 1 - Results of Logistic Regression of SUS Presence Questionnaire. Statistically significant results are marked with a box.

Contrast	$\chi^2(1)$	$p(\alpha = 0.05)$
D vs. HT	0.15	0.6980
D vs. TI	3.35	0.0672
D vs. T	0.02	0.8912
HT vs. TI	11.97	$\boxed{0.0005}$
HT vs. T	0.46	0.4986
T vs. TI	6.39	$\boxed{0.0115}$

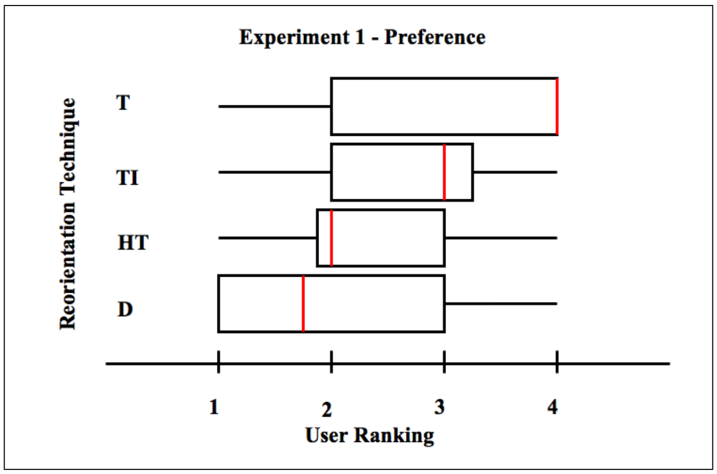

Figure 5.5: Experiment 1–User rated preference scores from 1 (most preferred) to 4 (least preferred). Standard box-and-whisker plots with the median in red.

5.2.2.3 Results

Tables 5.1 and 5.2 and Figures 5.4 through 5.8 show my results from Experiment 1. The SUS presence scores were analyzed using the same binomial logistic regression techniques as applied in previous uses of the questionnaire (Slater et al., 1995b). The response to each question was converted from the 1 to 7 scale to a binary value: responses of 5, 6, or 7 were converted to HIGH (1) and values less than 5 were converted to LOW (0). This conversion avoids treating the subjective ratings as interval data. After this conversion, I further transformed the data to create a new response variable for each participant: the count of their HIGH responses. Tables 5.1 and 5.2 show the average proportion of HIGH responses for each of the four conditions as well as the pairwise contrasts of conditions using logistic regression adjusted for multiple observations for each participant. There is a statistically significant effect between *HT* vs. *TI* ($\chi^2(1) = 11.97, p < 0.05$) and *T* vs. *TI*

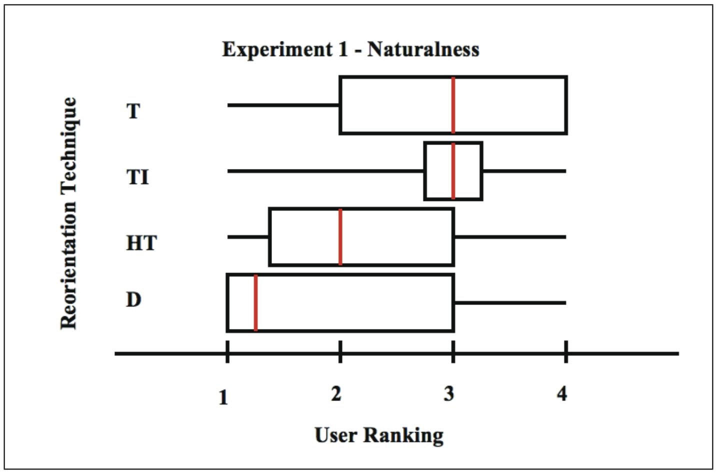

Figure 5.6: Experiment 1–User rated naturalness scores from 1 (most natural) to 4 (least natural). Standard box-and-whisker plots with the median in red.

$(\chi^2(1) = 6.39, p < 0.05)$. I also found a trend between D vs. TI $(\chi^2(1) = 3.35, p = 0.0672)$.

Figures 5.5 and 5.6 show the average user rankings, with 1 being the highest and 4 being the lowest, of preference and naturalness by ROT respectively. The data was analyzed using Friedman's ANOVA. User-ranked naturalness was significantly different between ROTs: $(\chi^2(3) = 9.524, p < 0.05)$, as was user-ranked preference $(\chi^2(3) = 10.958, p < 0.01)$. Wilcoxon tests were used to expand on this finding and a Bonferroni correction was applied. All effects are reported at a 0.0125 level of significance. The Wilcoxon test statistic is *T'* and should not be confused with my condition *T*. Participants significantly found *HT* to be more natural than *TI*, (*T'* = 220.00, $r = 0.38$) and significantly preferred *D* and *HT* to *T*, *T'* = 237.50, $r = 0.37$ and *T'* = 235.50, $r = 0.36$ respectively.

Figures 5.7 and 5.8 show user preference and user-ranked naturalness of paired ROTs. The frequency at which a participant preferred one ROT over another was compared to random choice, a frequency of 0.50, using Wilcoxon tests. I found participants significantly preferred D over TI $(T' = 184.00, p < 0.05, r = 0.31)$, HT over TI $(T' = 176.00, p < 0.05, r = 0.35)$, and HT over T $(T' = 165, p < 0.5, r = 0.28)$ and participants significantly considered *HT* to be more natural than *TI*, (*T'* = 170.00, $p < 0.01, r = 0.50$).

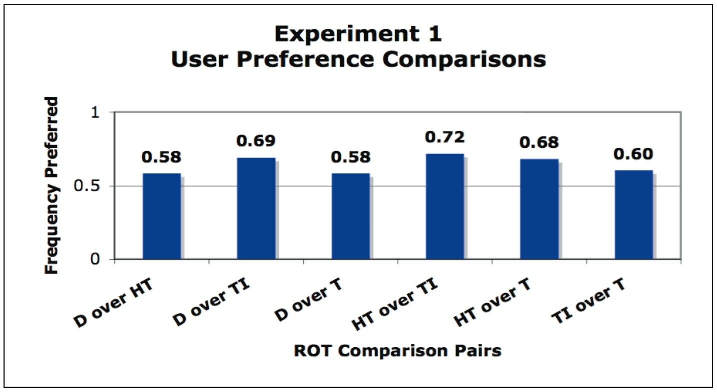

Figure 5.7: Experiment 1–User forced-choice comparisons of preference across ROTs.

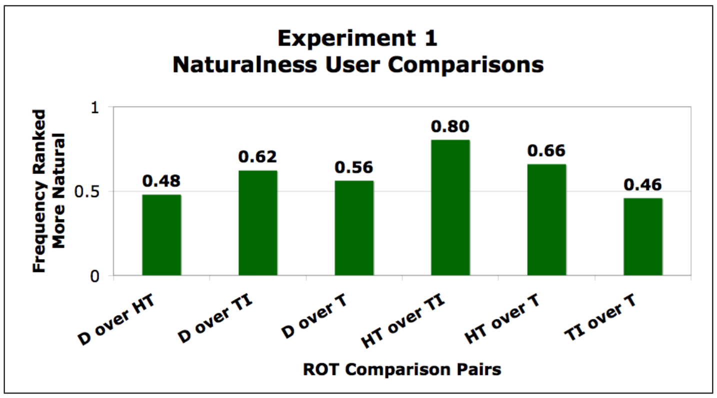

Figure 5.8: Experiment 1–User forced-choice comparisons of naturalness across ROTs.

5.2.2.4 Discussion

Participants' exit surveys and responses during the experiment provided useful information about each ROT. Participants' reasons for favorably rating ROTs included: the method provided instruction (either audio or visual), they did not notice rotation, and the method was realistic or natural. I believe that D and HT were rated higher by participants than T and TI because both rotate the VE while the participant is stimulating the vestibular system by turning her head and is less likely to notice the rotation of the VE.

Participants were confused during the first few sub-trials of T and often needed extra instruction from the experimenter to determine which direction to walk in the lab. After the first sub-trial of T one participant exclaimed, "Where am I?" and had to be stopped before walking out of the lab space. This occurred with several participants, however after three sub-trials participants often no longer needed extra instruction to determine the correct direction to walk in the lab. Participants described T as dizzying, and complained about disorientation in the VE after the world "spun." Some participants found T to be "fun" and simple because the participant just waited for "the flip" and then the virtual would moved as they expected.

Participants were occasionally confused by the audio instructions in TI asking for the participant to turn 360° but seeing the VE stop rotating after the participant only turned 180°. Participants would occasionally follow the audio instructions and turn 360° in the real world and then turn an additional 180° to walk the correct direction along the path. Participants also noticed the VE spinning at a much faster rate than they were turning. One participant complained about the disembodied voice that did not fit into the environment. Participants praised this technique for giving them some control over the VE by spinning when the participant turned and participants also found audio instructions helpful for determining how to turn around in the VE.

When using HT, participants complained about noticing the path in the VE not being in the right place once they started turning their heads but also commented on not seeing the rotation as much as other ROTs. Some participants would occasionally stop turning their heads before the VE had rotated 180° and would stand and wait until given more instruction to continue turning their heads. These participants would no longer need extra instruction after three sub-trials. Participants liked having control over the rotation of the VE that was offered by turning their heads.

Participants commented that the distractor was dizzying because it moved too fast, or that they would not be able to turn their heads fast enough to keep it in view. Participants also complained that a "big red ball is not normal." Some participants also complained about the ball's sudden appearance and disappearance. Other participants found D entertaining and engaging and found that when looking at the ball they were not paying attention to the moving VE.

The results revealed that D and HT were better ROTs than TI and T by producing increased presence, having higher user preference and being more natural to the user. However, user feedback suggested further improvements; these were explored in Experiment

2.

5.2.3 Experiment 2

Based on the results and user feedback from Experiment 1, the distractor method was improved by using a butterfly with flapping wings instead of a sphere because it is more natural for the VE being used in the experiment. The butterfly model is shown in Figure 5.9. The butterfly also flew in and out of the VE instead of suddenly appearing and disappearing, a common user complaint about the distractor from Experiment 1. I compared the improved distractor to the most promising ROTs from Experiment 1: the original red sphere distractor and head turn with audio instruction (Razzaque, 2005).

Figure 5.9: Butterfly distractor used in Experiment 2.

To have the butterfly appear more realistic and to respond to the complaints from Experiment 1 that the distractor was "dizzying," the speed of the butterfly was slowed down as it flew along an arc in front of the participant. To compare the difference in natural versus unnatural distractors the speed of the sphere was changed to match that of the butterfly.

5.2.3.1 Participants

Twelve participants (6 men and 6 women), mostly computer science graduate students in their twenties, participated in the experiment. Each participant visited the laboratory once for a session lasting approximately 1 hour and received \$7.50 for participation during the week and \$10.00 for weekend participation. All participants had normal or corrected-to-normal vision and were naive to the purpose of the study. Participants were not informed about ROTs and were initially unaware that the VE would rotate.

5.2.3.2 Experimental Design

Experiment 2 consisted of two parts, both taking place in the same VE. The VE was an outdoor space similar to Experiment 1, with a 180-meter straight wooden path and square markers placed 5 meters apart along the path. The environment is shown in Figure 5.1. Participants were instructed to walk along the designated path in the environment and to stop at each marker along the path. Once a participant had reached a marker, the participant experienced one of three reorientation techniques:

Head turn with audio instruction (HT). Audio instructions in the VE, presented via headphones, ask the user to turn her head back and forth and then continue walking along the path. While the user turns her head the rotation applied to the VE is 1.3 times the rotation speed of the user's head until the VE has rotated 180°. The participant reorients herself by rotating 180° in the real world. This is similar to a method described in (Razzaque, 2005).

Head turn with visual instruction, distractor (D). A moving sphere appears in front of the user. The user watches the sphere as it moves in a horizontal arc and continues walking along the path once the sphere disappears. The rotation applied to the VE is 1.5 times the rotation speed of the user's head until the VE has rotated 180°. The distractor moves along the arc with sinusoidal speed, amplitude = 180° and frequency = 1.125 Hz. The user reorients herself by rotating 180° in the real world. The path and velocity of the distractor are described in Figure 5.3.

Head turn with visual instruction, improved distractor (ID). A butterfly flies into the scene towards the participant, and then flies in a horizontal arc in front of the participant. The participant continues walking along the path once the butterfly flies away. Subjects are instructed before the trial to turn their heads to watch the butterfly. While the user is watching the butterfly the rotation applied to the VE is 1.5 times the rotation speed of the user's head until the VE has rotated 180°. The distractor moves along the arc with sinusoidal speed, amplitude = 180° and frequency = 1.125 Hz The user reorients herself by rotating 180° in the real world.

Part I of the experiment assessed the user's subjective sense of presence, how aware the user was of turning around, and how aware the user was of the VE rotation. Part I consisted of three trials, each using one ROT. The order of the trials was counterbalanced among participants. Each trial was comprised of eight sub-trials requiring the participant to walk along the virtual path to the next marker along the path. Once the participant reached a marker a ROT would stop the participant and rotate the VE. Each trial consisted of walking to eight markers, experiencing the same ROT eight times. Participants then removed the HMD and filled out the SUS presence questionnaire. In addition to the presence questionnaire, questions of interest about the VE rotating and the participant turning were embedded in the following list of questions:

Did you notice anything unnatural or odd during your virtual experience? Please rate the following on a scale from 0 to 7. Where 0 = did not notice or happen, 7 = very obvious and took away from my virtual experience.

___I felt like I was turning around

___I saw the virtual world get smaller or larger

___I saw the virtual world flicker

___I saw the virtual world rotating

___I felt like I was getting bigger or smaller

___I saw the virtual world get brighter or dimmer

Table 5.3: Experiment 2 - Mean percentage of HIGH scores on SUS Presence Questionnaire

ROT	$\bar{x}$
ID	0.52778
D	0.45833
HT	0.41667

Table 5.4: Experiment 2 - Results of Logistic Regression of SUS Presence Questionnaire

Contrast	$\chi^2(1)$	$p(\alpha = 0.05)$
ID vs. D	1.09	0.2974
ID vs. HT	1.72	0.1895
D vs. HT	0.63	0.4291

Part II consisted of 6 trials, each with two ROTs. Trials were counterbalanced and each ROT was compared to every other ROT twice with order reversed to remove the possibility of order effects. Each trial required the participant to walk to a marker, experience an ROT, and then walk to the next marker and experience a different ROT. The participant then made a forced-choice decision as to which ROT they preferred and which ROT was more natural. Participants were also asked to explain why they chose one ROT over another.

At the end of the experiment, participants filled out an exit survey and ranked the three ROTs based on naturalness and preference.

ID - Improved distractor
D - Distractor
HT - Head turn with audio instruction

Figure 5.10: Experiment 2–Legend

5.2.3.3 Results

Tables 5.3 and 5.4 and Figures 5.10 through 5.15 show my results from Experiment 2. The analysis of the SUS presence scores was done in the same manner as reported in Section 3.2.3. Tables 5.3 and 5.4 show the proportion of HIGH responses for each of the three conditions and the results of the pairwise contrasts of conditions. I found no statistical significance in user reported presence scores between ROTs.

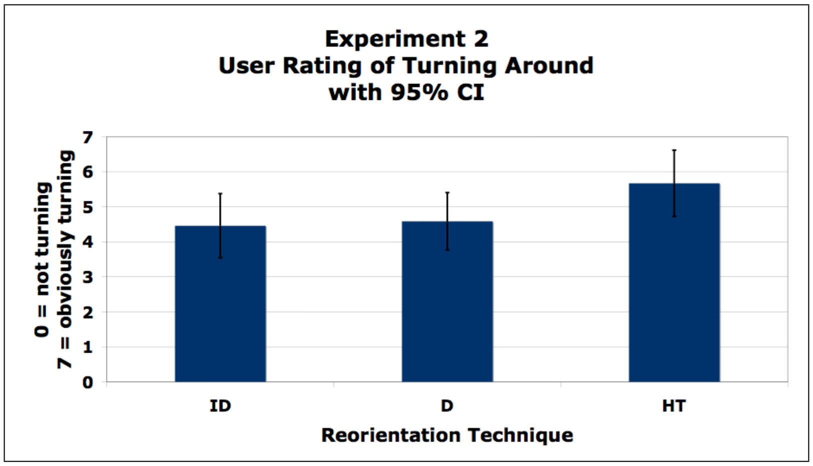

Figure 5.11: Experiment 2–User rating - "I felt like I was turning around" with ± 1 standard deviation.

Figure 5.11 shows, by ROT, the average user scores of response to the question about feeling as if they were turning around. I analyzed the data using Friedman's ANOVA and found significant differences between ROTs: $\chi^2(2) = 7.550, p < 0.05$. Wilcoxon tests were used to follow-up this finding. A Bonferroni correction was applied and all effects are reported at a 0.025 level of significance. Participants significantly rated the question "I felt like I was turning around," higher in *HT* than *D* (*T'* = 51.50, $r = 0.74$), and a trend was found that participants rated the feeling of turning around higher in *HT* than *ID* (*T'* = 46.50, $r = 0.56$).

Figure 5.12 shows, by ROT, the average user scores of response to the question about participants noticing that the VE was rotating. Using Friedman's ANOVA I found no significant difference between ROTs: $\chi^2(2) = 3.630, p = 0.187$.

Figures 5.13 and 5.14 show results from user ranked preference and naturalness by ROT, with 1 being the highest preference and 3 being the lowest. Trends were found between participant rankings of preference ($\chi^2(2) = 4.667, p = 0.108$) and participant rankings of naturalness ($\chi^2(2) = 5.167, p = 0.080$).

Figure 5.15 shows user preference and user-ranked naturalness of paired ROTs. The frequency at which a participant preferred one ROT over another was compared to random choice, a frequency of 0.50, using Wilcoxon tests. Participants preferred both *ID* and *D* to *HT* (*T'* = 65.00, $r = 0.47$, and *T'* = 77.00, $r = 0.51$ respectively), and ranked *ID* and *D*

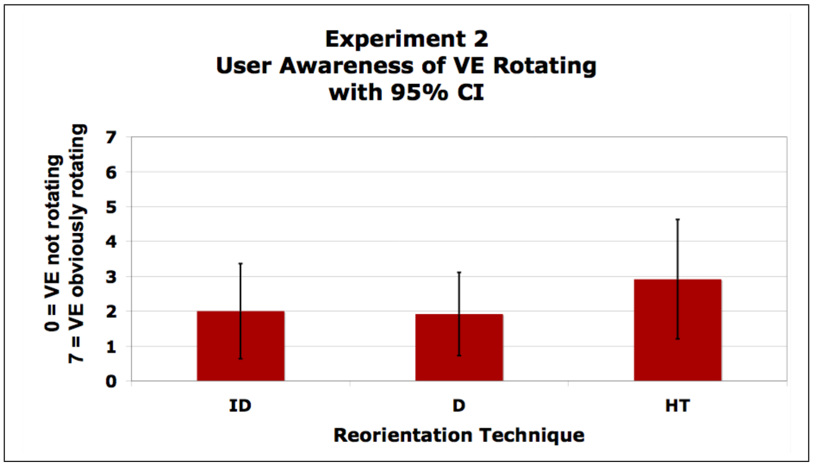

Figure 5.12: Experiment 2–User rating - "I saw the virtual world rotating" with ± 1 standard deviation.

to be more natural than *HT* (*T'* = 82.50, $r = 0.44$, and *T'* = 65.00, $r = 0.47$ respectively). A trend suggests that *ID* is more natural than *D* (*T'* = 63.00, $r = 0.28, p = 0.11$).

5.2.3.4 Discussion

The results from Experiment 2 suggest ROTs that use distractors reduce the liklihood of a users' feeling as if they are turning around while being reoriented. The results also suggest that participants prefer ROTs with distractors and consider them to be more natural than ROTs that do not use distractors. I account for the difference between *D* and *HT* in Experiment 2 compared to Experiment 1 by the reduced peak angular velocity of the sphere from 80°/sec to 20°/sec.

The VE rotates 1.3 times the rotation speed of the user's head in *HT* and 1.5 times the rotation speed of the user's head in *D* and *ID*. This difference in rotation speeds was an invertent design flaw to distractors disadvantage. That is, the VE rotated faster in the *D* and *ID* conditions and therefore the rotation should be more noticeable. However, no significant difference was found between ROTs and user awareness of head rotation speed. Further studies comparing different rotation speeds of the VE relative to head-turn speeds may reveal further differences between ROTs with and without distractors.

Exit surveys and responses during Experiment 2 provided useful information about each ROT. In the *HT* condition participants reported that turning their heads back and forth

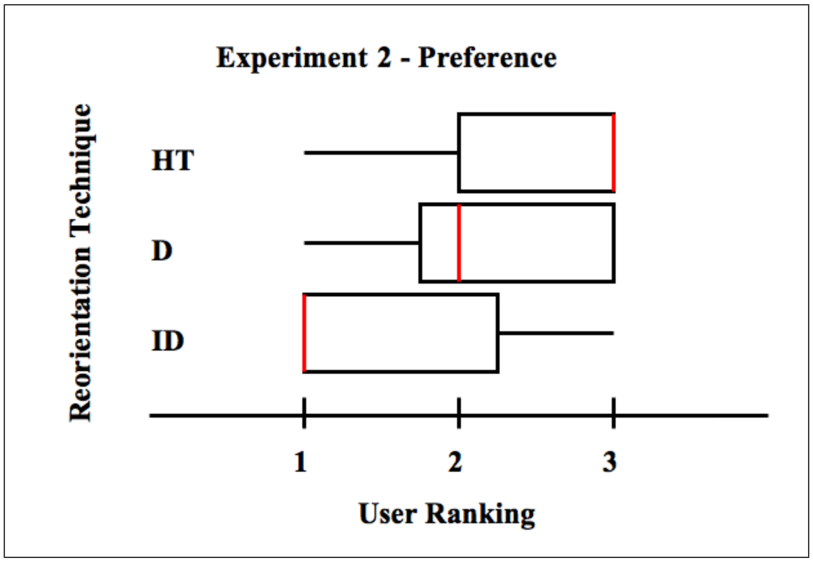

Figure 5.13: Experiment 2–User rated preference scores from 1 (most preferred) to 3 (least preferred). Standard box-and-whisker plots with the median in red.

for no reason to be annoying and "silly." One participant noted, "The voice destroys being there." Participants in HT were aware that the path had moved when they rotated their heads and complained of being more lost than with visual instruction. Two participants reported HT to provide more freedom and the ability to look around the environment during reorientation.

Participants found D to be easy to follow and some participants found D less distracting than the flapping butterfly wings of ID. Participants continued to complain about the sphere in HT not being natural to the environment and noted that it "defies the laws of physics." Participants commented on the naturalness of the butterfly, but some found the flapping of the butterfly wings to be "annoying." Participants enjoyed watching the butterfly fly in and out of the scene but, in Experiment 2, no negative comments were made about the sudden appearance and disappearance of the sphere. However, based on the numerous complaints about the sudden appearance and disappearance of the sphere in Experiment 1, I believe the distractor should engage the user in a manner natural to the scene in appearance and motion.

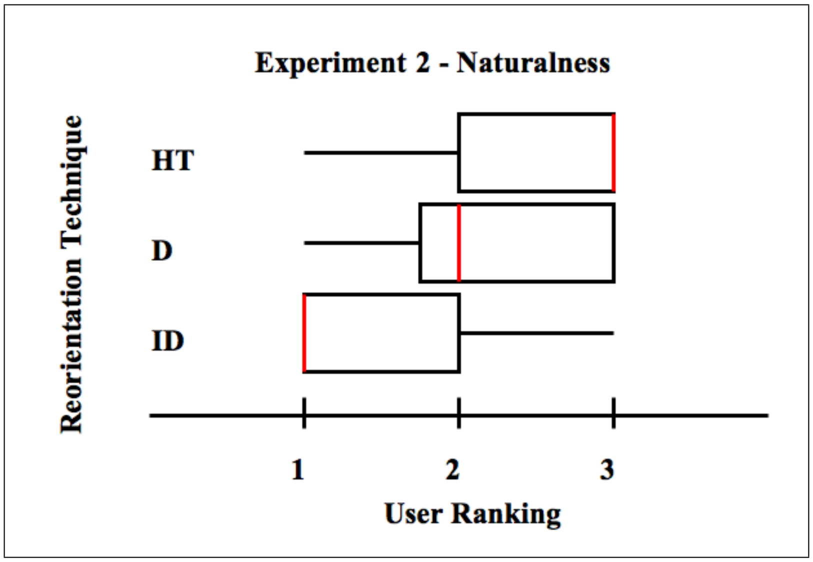

Figure 5.14: Experiment 2–User rated naturalness scores from 1 (most natural) to 3 (least natural). Standard box-and-whisker plots with the median in red.

5.2.4 Experiment 3

Based on user feedback from Experiment 2, I further improved the distractor method by using a more realistic model: a hummingbird (Figure 5.16)with flapping wings. I hypothesized that adding sound and improving the quality of the model would increase presence, be ranked more natural, and be preferred by users. In addition to using a more realistic model created using a realistic texture map and modeled by an artist, I explored adding sound to the visual distractor and using sound alone as a distractor. All distractors in this experiment had the same motion path and speed as the butterfly from Experiment 2.

5.2.4.1 Participants

Twelve participants, mostly graduate students and researchers (7 men and 5 women) participated in the experiment. The age range was 23 to 50, with an average age of 32. Each participant visited the laboratory once for a session lasting approximately 1 hour and received \$7.50 for participation during the week and \$10.00 for weekend participation. All participants had normal or corrected-to-normal vision and were naive to the purpose of the study. Participants were not informed about ROTs and were initially unaware that the VE would rotate.

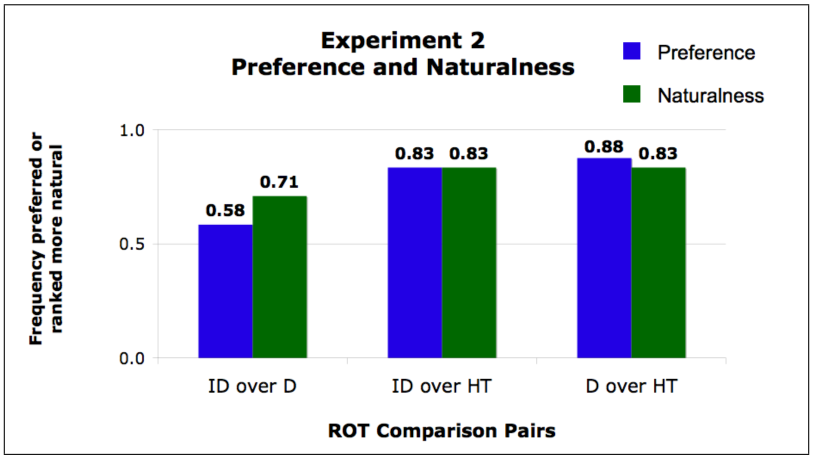

Figure 5.15: Experiment 2–User forced-choice comparisons of preference and naturalness across ROTs.

5.2.4.2 Experimental Design

Experiment 3 consisted of two parts, both taking place in the same VE. The VE was the same as Experiment 2 and consisted of a 180-meter straight wooden path with square markers placed five meters apart along the path. Participants were instructed to walk along the path in the environment and to stop at each marker along the path. Upon reaching each marker, the participant experienced one of three ROTs:

Distractor, visual (DV). A hummingbird flies into the scene towards the participant, and then flies in a horizontal arc in front of the participant. Before the trial, participants were instructed to turn theirs heads to watch the bird. The participant continues walking along the path once the hummingbird flies away. While the user is watching the hummingbird the rotation applied to the VE is 1.5 times the rotation speed of the user's head until the VE has rotated 180°. The distractor moves along an arc with 1.75m diameter, with sinusoidal speed, amplitude = 180° and frequency = 1.125 Hz. The user reorients herself by rotating 180° in the real world.

Distractor, visual and audio (DVA). A hummingbird flies into the scene towards the participant, and then flies in a horizontal arc in front of the participant. The hummingbird is accompanied by spatialized 3D audio of hummingbird wings flapping, presented via head-

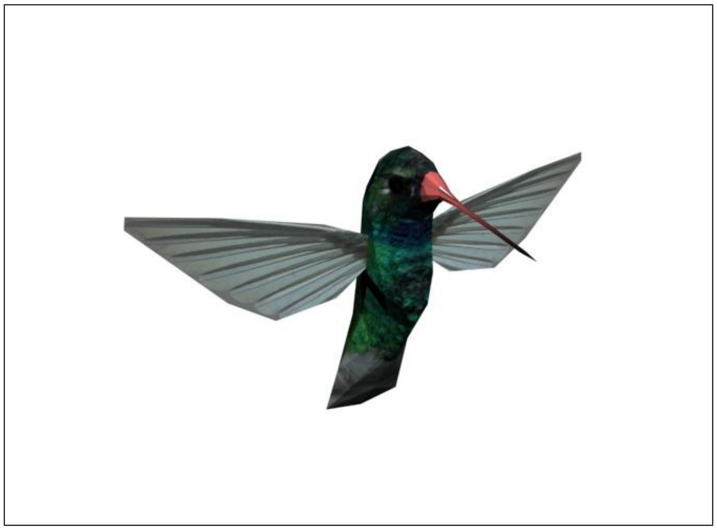

Figure 5.16: Hummingbird distractor used in Experiment 3.

phones. The hummingbird moves along an arc with sinusoidal speed, amplitude = 180° and frequency = 1.125 Hz. Before the trial, participants are given instructions to watch the bird. The participant continues walking along the path once the hummingbird flies away. While the user is watching the hummingbird the rotation applied to the VE is 1.5 times the rotation speed of the user's head until the VE has rotated 180°. The user reorients herself by physically rotating 180° in the real world.

Distractor, audio (DA). A sound of hummingbird wings flapping flies into the scene towards the participant, and then spatially moves in a horizontal arc in front of the participant. The sound has sinusoidal speed along the arc, amplitude = 180° and frequency = 1.125 Hz. Before the trial, participants are given instructions to watch the bird. There is no visual hummingbird to accompany the sound. The participant continues walking along the path once the sound of the hummingbird flies away. While the user is listening to the hummingbird the rotation applied to the VE is 1.5 times the rotation speed of the user's head until the VE has rotated 180°. The user reorients herself by rotating 180° in the real world.

Experiment 3 had the same experimental design as Experiment 2. Part I of the exper-

Table 5.5: Experiment 3 - Mean percentage of HIGH scores on SUS Presence Questionnaire

ROT	$\bar{x}$
DV	0.77780
DA	0.62500
DVA	0.69444

iment assessed the user's subjective sense of presence, how aware the user was of turning around, and how aware the user was of the VE rotation. Part I consisted of three trials, each using one ROT. The order of the trials was counterbalanced among participants. Each trial was comprised of eight sub-trials requiring the participant to walk along the virtual path to the next marker along the path. Once the participant reached a marker a ROT would stop the participant and rotate the VE. Each trial consisted of walking to eight markers, experiencing the same ROT eight times. Participants then filled out the SUS presence questionnaire. In addition to the presence questionnaire, participants also answered the embedded questions about the VE rotating and the user turning around that were presented in Experiment 2 (Section 5.2.3).

Part II consisted of 6 trials, each with two ROTs. Trials were counterbalanced and every ROT was compared to every other ROT twice with order reversed to remove possible order effects. Each trial required the participant to walk to a marker, experience a ROT, and then walk to the next marker and experience a different ROT. The participant then made a forced-choice decision as to which ROT they preferred and which ROT was more natural. Participants were also asked to explain why they chose one ROT over another.

At the end of the experiment, participants filled out an exit survey and ranked the three ROTs based on naturalness and preference.

DV - Distractor, visual
DVA - Distractor, visual and audio
DA - Distractor, audio

Figure 5.17: Experiment 3–Legend

Table 5.6: Experiment 3 - Results of Logistic Regression of SUS Presence Questionnaire. Statistically significant results are marked with a box.

Contrast	$\chi^2(1)$	$p(\alpha = 0.05)$
DV vs. DA	6.23	$\boxed{0.0126}$
DV vs. DVA	1.60	0.2060
DVA vs. DA	1.99	0.1581

Table 5.7: Experiments 2 and 3 - Results of Logistic Regression of SUS Presence Questionnaire comparing data from Experiment 3 to data from Experiment 2. Statistically significant results are marked with a box.

Contrast	$\chi^2(1)$	$p(\alpha = 0.05)$
Exp 3 vs. Exp 2		
DV vs. ID	6.18	$\boxed{0.0129}$
DV vs. D	10.73	$\boxed{0.0011}$
DV vs. HT	10.44	$\boxed{0.0012}$
DVA vs. ID	3.29	0.0699
DVA vs. D	7.76	$\boxed{0.0053}$
DVA vs. HT	9.06	$\boxed{0.0026}$
DA vs. ID	1.09	0.2969
DA vs. D	3.84	$\boxed{0.0500}$
DA vs. HT	6.60	$\boxed{0.0102}$

5.2.4.3 Results

Tables 5.5 and 5.6 and Figures 5.17 through 5.23 show the results from Experiment 3. Note that Figure 5.18 shows results from both Experiments 2 and 3. The analysis of the SUS presence scores was performed in the same manner as reported in Section 3.2.3. Tables 5.5 and 5.6 show the proportion of HIGH responses for each of the three conditions and the results of the pairwise contrasts of conditions. I found users felt significantly more present in *DV* than *DA* ($\chi^2(1) = 6.23, p < 0.05$).

Experiments 2 and 3 used an identical experimental design: participants perform the same number of trials and used the same environment. Differences in presence scores between experiments may occur because of differences in ROTs displayed to participants, however, participants for both experiments came from the same pool. I compared presence scores between Experiment 2 and Experiment 3. Figure 5.18 and Table 5.7 show the percentage of HIGH responses for each of the three conditions and the results of the pairwise contrasts of conditions. I found users felt significantly more present in *DV* than *ID*, *D*,

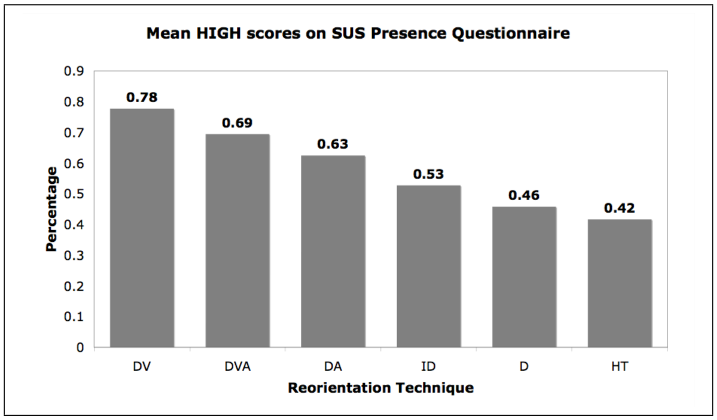

Figure 5.18: Experiments 2 and 3–User rating - Mean percentage of HIGH scores on SUS Presence Questionnaire.

and HT ($\chi^2(1) = 6.18, p < 0.05$, $\chi^2(1) = 10.73, p < 0.01$, $\chi^2(1) = 10.44, p < 0.01$ respectively). Users statistically felt more present in DVA than D and HT ($\chi^2(1) = 7.76, p < 0.01$, $\chi^2(1) = 9.06, p < 0.01$, respectively), and a trend suggests that users feel more present in DVA than ID ($\chi^2(1) = 3.29, p = 0.07$). Users also felt significantly more present in DA than D and HT ($\chi^2(1) = 3.84, p = 0.05$, $\chi^2(1) = 6.60, p < 0.05$, respectively).

Figure 5.19 shows average scores of response to the question about feeling like they were turning around for each ROT. I analyzed the data using Friedman's ANOVA and found no significant differences between ROTs: $\chi^2(2) = 0.712, p = 0.514$.

Using Friedman's ANOVA I found no significant difference between ROTs and participants noticing that the VE (Figure 5.20) was rotating $\chi^2(2) = 1.372, p = 0.298$.

Figures 5.22 and 5.21 show participants' ranked preference and naturalness of ROTs with 1 being the highest rank and 3 being the lowest. I found significant differences between ROTs of participant ranked preference ($\chi^2(2) = 16.875, p < 0.05$) and participant ranked naturalness ($\chi^2(2) = 102.308, p < 0.001$). Wilcoxon tests were used to follow-up this finding. A Bonferroni correction was applied and all effects are reported at a 0.025 level of significance. Participants significantly preferring DVA to DV and DA ($T' = 66.00$, $r = 0.352$, and $T' = 75.50$, $r = 0.433$ respectively), and a trend was found with participants preferred DV to DA ($T' = 62.00$, $r = 0.306$). Participants ranked DVA to be more natural than DV and DA, ($T' = 66.00$, $r = 0.387$, and $T' = 72.00$, $r = 0.342$ respectively).

Figure 5.23 shows user preference and user-ranked naturalness of paired ROTs. The

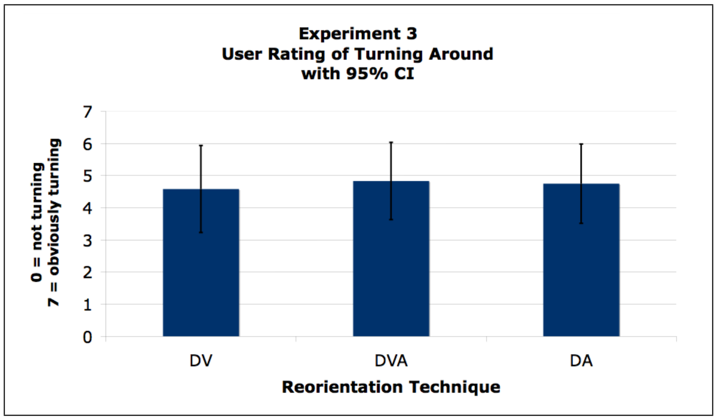

Figure 5.19: Experiment 3–User rating - "I felt like I was turning around" with ± 1 standard deviation.

frequency at which a participant preferred one ROT over another was compared to random choice, a frequency of 0.50, using Wilcoxon tests. Participants preferred *DVA* to both *DV* and *DA* ($T' = 55.00$, $r = 0.575$, and $T' = 55.00$, $r = 0.575$ respectively). Participants also preferred *DV* to *DA* ($T' = 60.00$, $r = 0.45$). Participants ranked *DVA* to to be more natural than both *DV* and *DA* ($T' = 55.00$, $r = 0.575$, and $T' = 54.00$, $r = 0.352$ respectively).

5.2.4.4 Discussion

The results from Experiment 3 suggest that users felt an increased subjective sense of presence with a realistic visual distractor without audio than with only an audio distractor. Although group differences may effect results, I performed contrasts between Experiment 2 and 3 and found that improving the visual quality of the distractor from an unrealistic butterfly to a more realistic hummingbird produced a higher feeling of presence among users. Note that the motion path and animation of the distractors was not modified between Experiments 2 and 3. The results suggest that using more realistic distractors can increase a user's feeling of presence.

Adding natural audio sounds to a visual distractor resulted in no significant increase of user-reported presence when compared to a visual distractor without audio. However, users prefer the addition of audio cues to the visual distractor and find the audio plus visual stimuli to be more natural than visual or audio alone. Many users claimed that the hummingbird with the sound of wings flapping stimulated more senses and was therefore

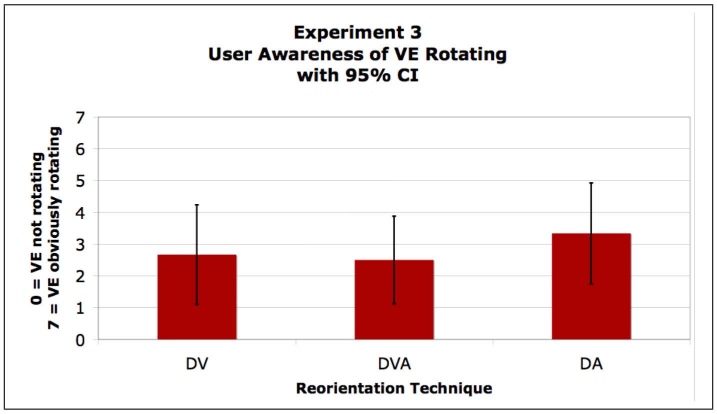

Figure 5.20: Experiment 3–User rating - "I saw the virtual world rotating" with ± 1 standard deviation.

more natural. No significant change in user-reported presence was found between having visual cues and when the visual cue of the hummingbird were removed and only the three-dimensional audio cues were presented to the user.

When comparing presence data from Experiment 2 and Experiment 3, I found that natural audio as a distractor without visual cues produces a higher sense of presence than using the unnatural red sphere distractor from Experiment 2. The ability to use only audio as a distractor extends the range of VEs in which distractors are applicable. Possible applications for audio distractors include military applications where environment-appropriate moving visual objects in front of the user would be distracting. Military training applications may have loud noises or explosions that naturally suit the environment and can be used as distractors. However, further studies need to be conducted to determine if distractors cause mis-training in military applications. Audio distractors may be especially useful for VEs because they do not require model changes and modeling and animation expertise.

One user commented that the audio distractor was hard to track and while he was searching to find the (audio) hummingbird he was much less aware of the VE rotating. Other users found the audio frustrating because they had a hard time determining the location of the sound source. This may be the reason that users ranked the audio distractor lower than the distractors with a visual hummingbird. Users may prefer natural distractors with audio to audio distractors alone, but audio distractors may still be effective.

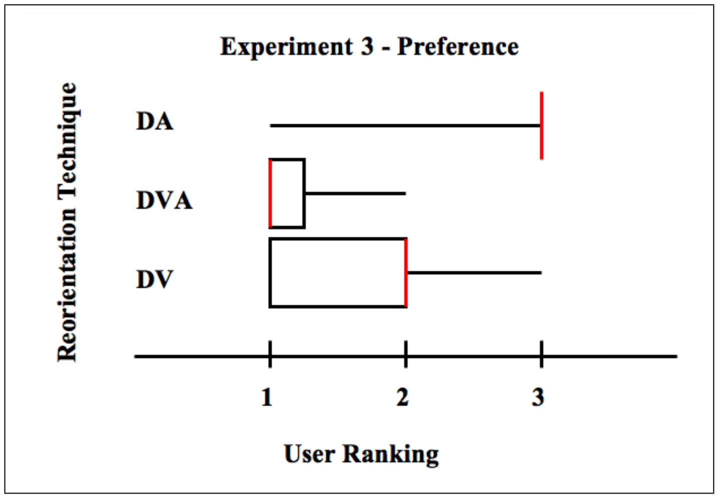

Figure 5.21: Experiment 3–User rated preference scores from 1 (most preferred) to 3 (least preferred). Standard box-and-whisker plots with the median in red.

5.3 Conclusion

I successfully implemented and tested eight ROTs to handle the worst-case scenario in large-walking VEs–when the user is about to walk out of the tracked space. Five of these ROTs use a novel technique, distractors–objects in the VE that the user focuses on while the VE rotates–to minimize the observed rotation of a VE during reorientation. In addition to reducing observed rotation of the VE, ROTs using distractors were preferred and ranked more natural by users than currently available ROTs that do not use distractors. I also found participants were less aware of physically turning around in the VE when reorienting using distractors.

User feedback suggests that ROTs should be realistic and users should not notice the rotation of the VE. Unlike non-distractor ROTs, distractors can be realistic and the results suggest distractors reduce the likelihood of perceiving VE rotation during reorientation. Distractors exhibiting smooth movements that are easy and interesting to watch received positive feedback from users. Improving the realism of the distractor increases a user's feeling of presence, and adding natural audio to a visual distractor is preferred and considered more natural to users than using a visual or audio distractor alone.

An audio alone distractor doesn't produce as high a feeling of presence as a natural audio plus visual distractor, however it does produce a higher feeling of presence than an

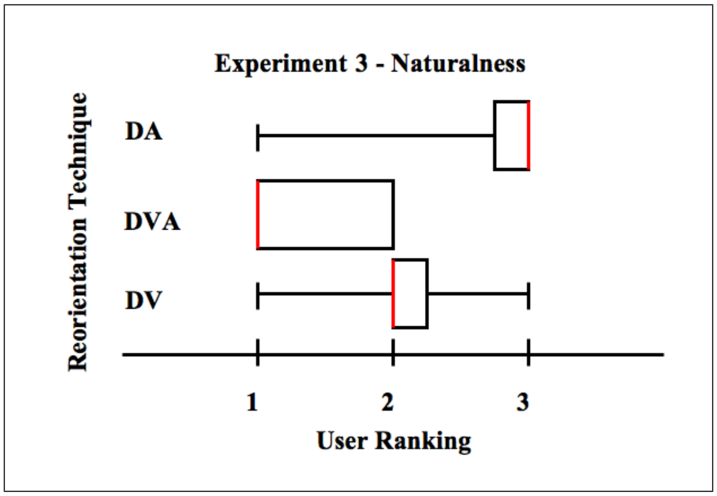

Figure 5.22: Experiment 3–User rated naturalness scores from 1 (most natural) to 3 (least natural). Standard box-and-whisker plots with the median in red.

unnatural distractor without audio. Audio distractors are easier to implement than visual distractors as they require no model changes. Audio distractors may also be useful for VEs in which the addition of visual distractors may be unnatural or detract from the VE experience.

I believe that optimal distractors are VE-dependent and should be designed to be as natural as possible to the VE. Possible implementations of distractors include: exploring a virtual house and having a dog run by, walking through a virtual art museum and having a docent point you in a new direction, and training dismounted infantry to successfully navigate enemy territory while snipers are heard in the distance.

This chapter introduced, evaluated, and incrementally improved distractors. The results from this chapter suggest that using distractors as ROTs are the current best solution and should be included in the design of large-scale real-walking interfaces to prevent users from leaving the tracked space.

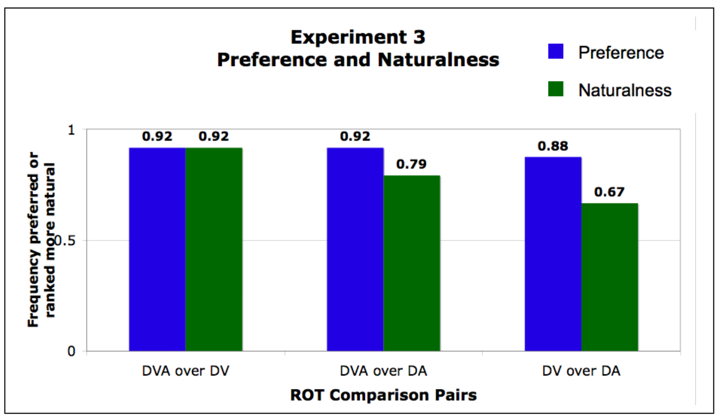

Figure 5.23: Experiment 3–User forced-choice comparisons of preference and naturalness across ROTs.

CHAPTER 6

An Evaluation of Navigational Ability Comparing Redirected Free Exploration with Distractors to Real Walking

Users in virtual environments often find navigation more difficult and frustrating than in the real world. In this chapter I compare Redirected Free Exploration with Distractors (RFED) to the current best virtual locomotion interface, really-walking. I compare the two locomotion interfaces by measuring navigational ability with standard wayfinding and locomotion metrics. The study results show that RFED users can really-walk through VEs that are somewhat larger than the tracked space and can navigate and wayfind no worse than when really walking.

6.1 Experiment

I performed a between-participants study that required participants to navigate virtual mazes and find targets. The virtual mazes used in this experiment are shown in Figure 6.1.

The University of North Carolina's Effective Virtual Environment's tracked space is $9m \times 9m$. Since I was comparing RFED to RW, the mazes were restricted to $8m \times 8m$ so they could fit completely within the tracked space while still having good tracking around the edges. Participants in the RW condition really-walked through the mazes. Participants in the RFED condition used the interface described in Chapter 4 and were restricted to walking in a space that was $6.5m \times 6.5m$ in the center of the $9m \times 9m$ tracked space. Restricting the walking space to $6.5m \times 6.5m$ to show that participants in RFED could

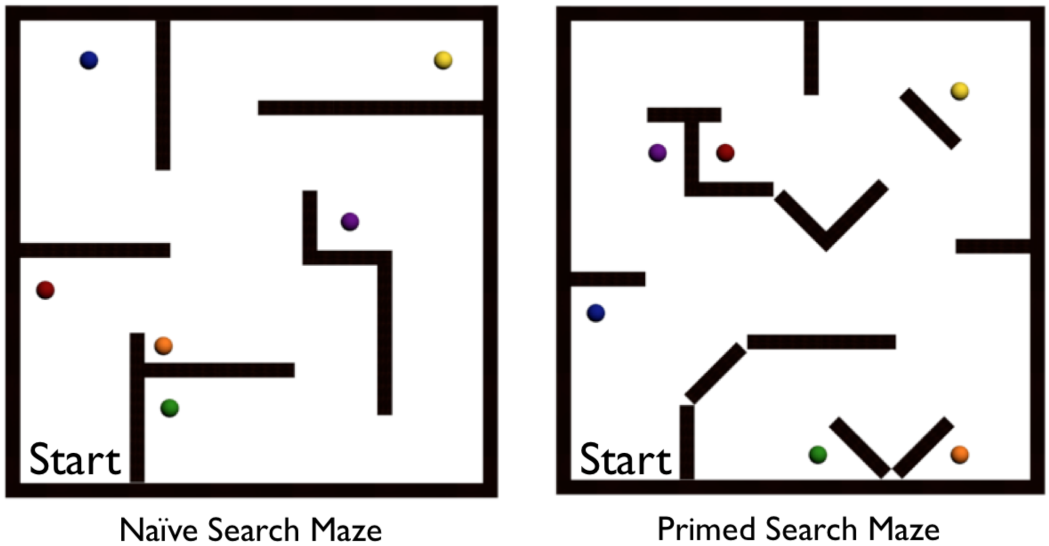

Figure 6.1: An overhead view of the mazes and target locations used in the naïve and primed searches. Participants started each maze in the bottom left corner.

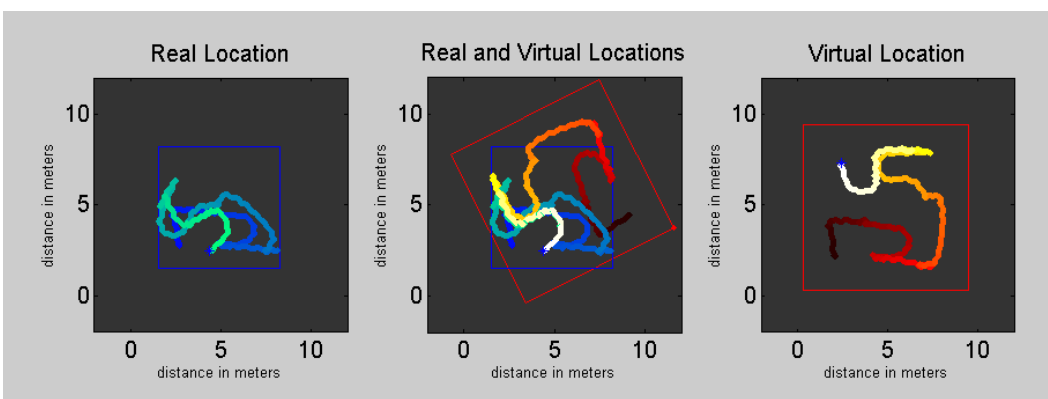

Figure 6.2: A participant walking through the maze in Part 1, a naïve search. The blue box is the boundary of the tracked space and the red box is the boundary of the VE. The left most image is the participant's real path over time. The right most image is the participant's virtual path over time. The start of the participant's real and virtual paths are dark and the ends are light. The center image is the final composite of the left and right images with the final transformation applied to the VE and user's virtual path.

walk in larger than tracked-space VEs.

Based on pre-test observation of the current implementation of RFED, using a tracked space smaller than $6.5m \times 6.5m$ causes RFED users to stop too frequently and increases user frustration. Therefore, a tracked space of at least $6.5m \times 6.5m$ is recommended when implementing RFED. This is discussed again in Chapter 8.

6.1.1 Hypotheses and Measures

Evaluation of RFED is based on being "no worse" than RW using 95% confidence interval equivalence testing. For each measure, if the 95% confidence intervals of the *mean difference* between RFED and RW falls within $\pm$ our predefined acceptable value of the difference, then RFED is no worse than RW (Wellek, 2002).

I compared RFED to RW through three common navigation and wayfinding tasks: *search* for specified targets within the VE, *point-to-targets* that are not visible, and *map completion*.

Navigation. *Search tasks*, which are common VE locomotion tasks (Bowman, 2002), are used to evaluate navigational ability and VE training-transfer of spatial knowledge (Waller et al., 1998; Witmer et al., 1996) for locomotion interfaces. The search task evaluation includes a *naïve search*, in which targets have not yet been seen, and a *primed search*, in which targets have previously been seen. For each locomotion interface I measure the total distance participants traveled using each interface, and the number of times participants revisit already seen areas of the virtual mazes. I claim that participants who travel shorter distances have a better spatial understanding of the environment and of previously visited locations within the environment. Thus, participants who travel shorter distances are less likely to retrace previous steps. That is, participants with better mental maps of the environment should walk shorter distances and not revisit areas of the maze.

Wayfinding. *Point-to-target* techniques require participants to point to targets that they have previously seen, but are currently out of view. Pointing-to-targets measure a user's ability to wayfind within VEs (Chance et al., 1998) by requiring a mental model of the relationship of target locations to the user's current location. A small sum of absolute angular pointing errors suggests that participants have a good understanding of the location of targets.

Map completion requires users to place and label targets at their corresponding VE locations on a paper map of the VE. Map completion is often used as a wayfinding metric because maps are a familiar navigation metaphor (Darken and Peterson, 2002). Participants with a better mental model of the VE should be able to more accurately place targets in correct locations and correctly label targets on the map.

6.1.2 Participants

Twenty-two participants, 18 men and 4 women, with average age 26, participated in the IRB-approved experiment. One participant's data was not used because the experimenter believed the participant not to be trying because the participant completed the experiment in half the time of all other participants. Eleven participated in the RW condition, (9 men and 2 women) and ten participated in the RFED condition, (8 men and 2 women). Not all participants were naïve to RFED, therefore all participants in the RFED condition were informed about the locomotion interface.

6.1.3 Equipment

Each participant wore a stereo nVis nVisor SX head-mounted display with 1280x1024 resolution in each eye and a diagonal FOV of 60°. The tracked-space was $9m$ x $9m$ and tracked using a 3rdTech HiBall 3000. The textured maze environment was rendered on a Pentium D dual-core 2.8GHz processor with an NVIDIA GeForce GTX 280 GPU and 4GB of RAM. The interface was implemented in a locally developed EVEIL intermediate level library that communicates with the Gamebryo® software game engine from Emergent Technologies. The Virtual Reality Peripheral Network (VRPN) was used for tracker and button communication.

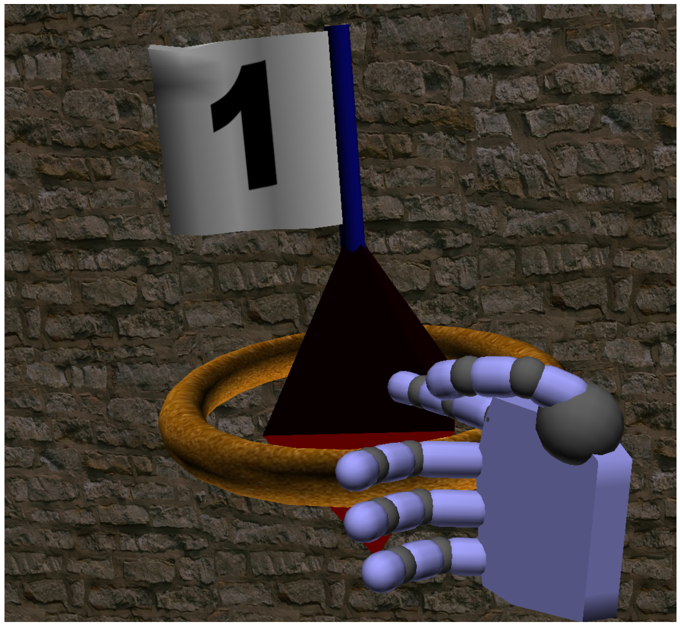

Figure 6.3: A screen shot of the virtual avatar hand selecting Target 1, the red target.

6.1.4 Experimental Design

The experiment used three virtual environments, a training environment and the two testing environments shown in Figure 6.1. The environments are $8m \times 8m$ mazes with uniquely colored and numbered targets placed at predefined locations. All mazes had the same textures, and the same coloring and numbering of targets. The location and total number of targets changed with each environment. All participants completed the same trials in the same order to control for training effects. Participants were randomly assigned the RW or RFED condition, and completed all parts of the experiment, including training, within the assigned condition.

6.1.4.1 Training

Training. The first environment, the training environment (Figure 6.4), was a directed maze with all walls placed at 90° angles. Participants read written instructions before beginning each section of the experiment and were advised to ask questions if they were unclear about tasks. Participants walked through the training environment and used a hand-held tracked device, with trigger button, to select each of the four targets located

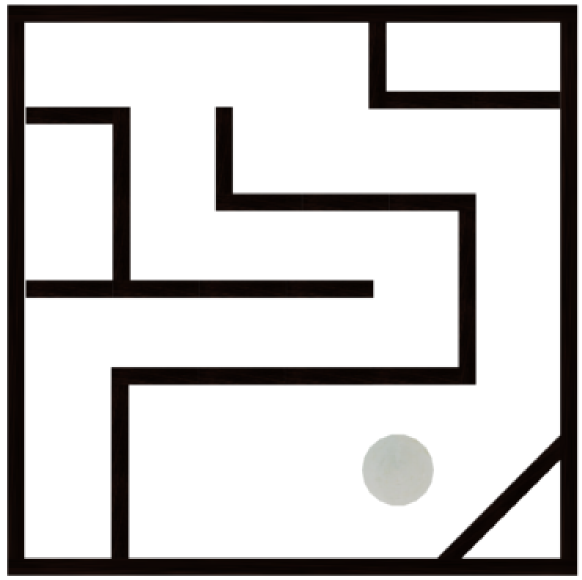

Figure 6.4: An overhead view of the training maze.

along the path. When a target was selected, a ring appeared around it and audio feedback was played to signify that the target had been selected (Figure 6.3). Once participants had found the four targets, they were asked to stand inside a circle on the floor of the VE and practice pointing and clicking at a target that was 1.5m in front of them. The training session ended when the participant successfully pointed within 6cm of the center of the target. No success feedback was given.

After participants completed the training maze, the head-mounted display was removed and participants were asked to complete a 8.5" $\times$ 11" paper map of an overhead view of the maze with targets missing. Maps were presented such that the initial starting direction was up and away from the user, and the starting location was given to the participant. By hand, participants placed a dot at the corresponding location to each target, and labeled each target with its number or color. Participants were not given performance feedback during any part of the experiment after the training session.

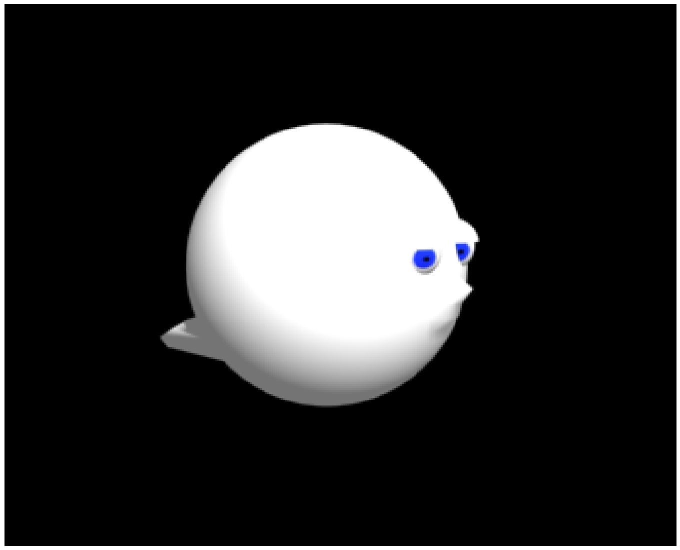

Figure 6.5: A screen shot of the ghost distractor.

6.1.4.2 Part 1: Naïve Search

Naïve search. After training, participants read instructions for Part 1, the naïve search. The non-directed maze and target locations for Part 1 can be seen in Figure 6.1. Participants were instructed to find, select, and remember the location of the six targets within the maze and were reminded they would have to complete a map, just as in the training session. Participants in the RFED condition were reminded that a ghost might appear within the environment, and if the ghost appeared, they were to take one step backward (so as to step out of the **turn distractor on** area) and turn their heads to follow the distractor ghost (see Figure 6.5) as it moved in an arc in front of them. Participants were allowed to continue walking once the ghost disappeared. As soon as participants selected all targets, the virtual environment faded to white and participants were instructed to remove the head-mounted display. Participants then completed a map as in the training.

6.1.4.3 Part 2: Primed Search

Primed search. Participants were given written instructions for Part 2, the primed search, and participants in the RFED condition were again reminded about the distractor ghost. The maze and target locations for Part 2 can be seen in Figure 6.1. The VE is similar to the maze from Part 1 except that not all walls were axis-aligned. This was done to make the experiment more challenging by removing feedback that enables users to determine cardinal directions from the walls. Participants first followed a arrow directed *priming path* that led to each of the six targets. After participants reached the end of the priming path, marked by the screen fading to white, they removed the HMD and moved in the real world to the starting point in the VE.

Participants put the HMD back on and were asked to walk, as directly as possible, to one of the targets in the maze. Once the participant reached the specified target, they were instructed, via audio instruction, to point, in turn, to each of the other targets. In the audio instructions, targets were referenced by both color and number. After the participants had pointed in the direction of each target and clicked the hand held wand, they were audio instructed to walk to another target and to repeat a similar pointing task. All participants locomoted between the six targets in the same order. If a participant could not find a target within three minutes, arrows appeared to direct the participant to the target. Once the participant reached the target, the experiment continued as before, with the participant then pointing to all other targets.

Participants walked to each of the six targets in the order 4-2-1-3-6-5 and, from each, pointed to each of the other targets in numerical order. At the end of Part 2, participants had pointed to each target five times, for a total of 30 pointing tasks per participant.

After completing the search and pointing tasks, participants removed the HMD and completed a map just as in the previous parts of the experiment.

After the experiment, participants completed a modified Slater-Usoh-Steed Presence Questionnaire (Slater and Steed, 2000) and a Simulator Sickness Questionnaire (Kennedy et al., 1993).

6.2 Results and Discussion

6.2.1 Part 1: Naïve Search

Navigation. The total distance each participant traveled was calculated from head-tracked data. A trend suggests that participants walked greater distances using RFED

Naïve Search	Predefined Equivalence Rate	Explanation	95% Confidence Interval of mean difference	Mean	Comparison
Distance traveled	2.5 meters	10% of shortest possible distance	(-35.05m, **-9.38**m)	-22.12m	9.38 ≮ 2.5
Revisited routes	1 route	smallest measurable unit	(-0.72 routes, **0.88** routes)	0.08 routes	0.88 < 1
Map					
% correctly placed targets	16.7%	within 1 target, the smallest measurable unit	(**-4.6%**, 27.7%)	11.50%	16.7% > 4.6%
% correctly placed and labeled targets	16.7%	within 1 target, the smallest measurable unit	(**-21.4%**, 38.9%)	8.63%	16.7% ≯ 21.4%
Primed Search					
Distance traveled	0.55 meters	10% of average of shortest distances to each target	(-0.547m, **1.172**m)	0.312m	1.172 ≮ 0.55
Revisited routes	1 route	smallest measurable unit	(-0.22 routes, **0.21** routes)	0.004 routes	0.21 < 1
Absolute pointing angular error	15°	results from (Grant and Magee, 1998)	(**-8.074°**, **8.165°**)	0.045°	15°>8.07°, 8.17°<15°
Pointing time	1 second	audio instruction time was greater than equivalence rate	(-1.15s, **0.81**s)	-0.16s	0.81 < 1
Map					
% correctly placed targets	16.7%	within 1 target, the smallest measurable unit	(**-35.2%**, 36.4%)	0.60%	16.7% ≯ 35.2%
% correctly placed and labeled targets	16.7%	within 1 target, the smallest measurable unit	(**-7.0%** and 21.0%)	7.00%	16.7% > 7.0%
Post Tests					
Presence	1 "high" score	smallest measurable unit	(**-1.97**, 0.44)	-0.76	1 ≮ 1.97

Figure 6.6: A summary of the predefined equivalence values and results from the equivalence tests performed in this study. Bold faced 95% CI values were compared to the predefined equivalence values to evaluate RFED being "no worse" that RW. "No worse" than results are highlighted with dashed lines.

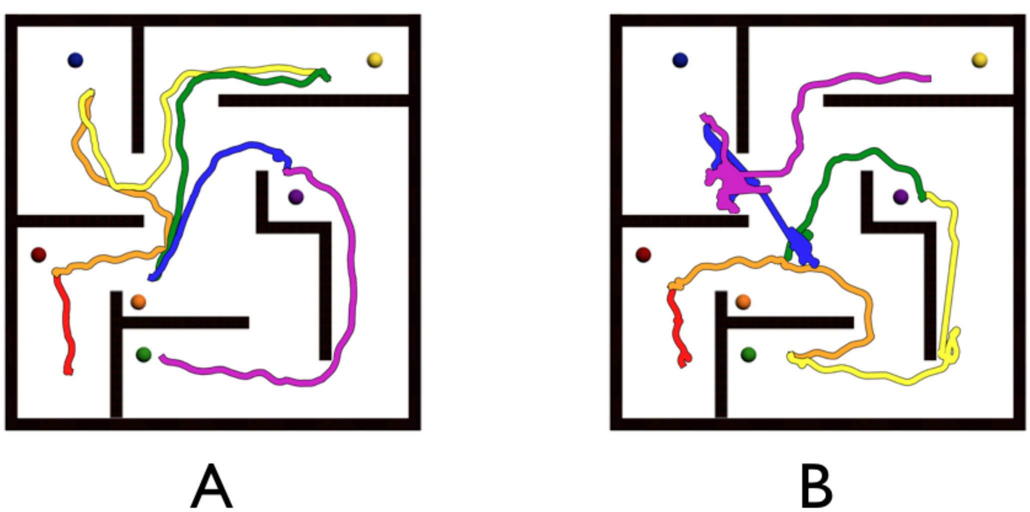

Figure 6.7: The virtual-space routes taken by the median performing participant from each locomotion interface during the naïve search. **A:** real walking (virtual space = real space); **B:** Redirected Free Exploration with Distractors. Notice the areas in **B.** where the participant continues walking in the same area. This is caused by distractor appearances. Participants started in the bottom left corner; path segment color follows the colors of the rainbow, ROYGBV, to mark the participant's finding and selecting a new target.

than RW, $t(19) = 4.08$, $p = 0.058$. Walking longer distances suggests that participants in the RFED condition were more lost than participants in the RW condition. However, when a distractor appeared in the RFED condition participants had to take one step backwards which added to their total traveled distance. Figure 6.7 shows virtual routes from median participants in the RFED and RW conditions.

To account for participants in the RFED condition having to take an extra step every time a distractor appeared, I performed two different analyses on the route data: **1.** filtering the routes and **2.** transforming the data. I present the analysis of **2.** transforming the data and both analyses produced similar results. To transform the data I subtracted 0.874 meters, an average step size, for every distractor that appeared. Although participants took one step backwards and then had to retrace that step, accounting for two steps, I subtracted only one step as a conservative estimate. See Figure 6.8. Since the naïve search could be completed by walking a minimum of 25 meters (the actual average RW distance was 39 meters), I predefined that RFED would be no worse than RW if RFED participants traveled no more than 10% of the shortest distance, or 2.5 meters more than RW participants. The 95% confidence interval of the mean difference between interface conditions, with modified RFED data, was -35.05 meters to -9.38 meters, $\bar{x} = -22.21$, $SE = 6.133$. **The results suggest that with 95% confidence, real-walking participants will travel shorter distances than RFED participants**.

For further route evaluation, the number of revisited routes was counted. I defined a *revisited route* as an area in the maze that a participant revisits. I predefined that RFED would be no worse than RW if participants revisited no more than one route more, the smallest measurable unit, when compared to RW. The 95% CI of the mean difference between the number of revisited routes for the two conditions was -0.72 routes to 0.88 routes, $\bar{x} = 0.08$, $SE = 0.38\%$. That is, with 95% confidence, participants using RFED will revisit no more than 0.88 routes more than and no fewer than 0.72 routes less than real walking. Since 0.88, is less than one, the predefined equivalence value, **I conclude that, when performing a naïve search, participants using RFED do not revisit more routes than real-walking participants**.

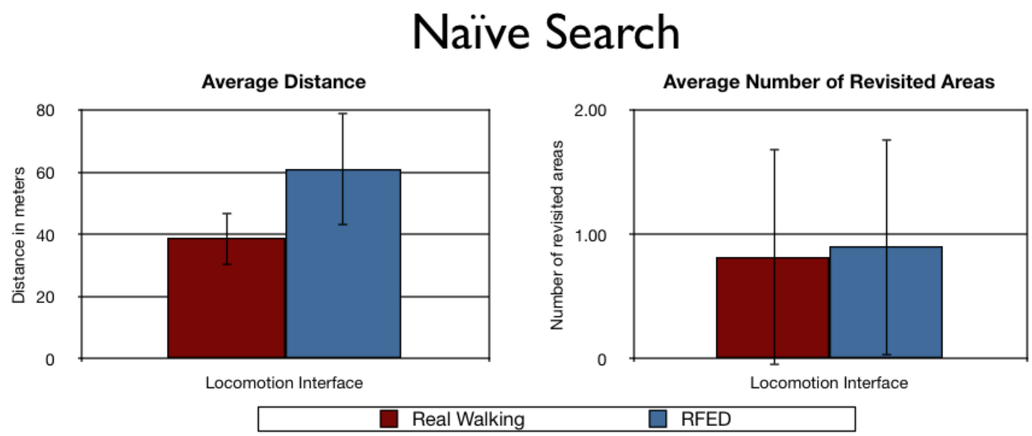

Figure 6.8: The average total distance traveled, and the average number of revisited areas between RW and RFED during Part 1, a naïve search, with ± one standard deviation error bars.

Wayfinding. After participants found all targets within the maze, they completed a map by placing dots at the corresponding locations to each target, and labeled each target with its corresponding color or number. Each participant was given a score of the percentage of correctly placed targets. A target was scored as correctly placed if a dot was within $2cm$ on the map (1 meter in the environment) of an actual target. Additionally, the dot, corresponding to the location of a target, had to be placed on the correct sides of walls. A correctly placed-and-labeled target had to be correctly placed, based on the rules above, and had to be labeled with the correct number or color. The results can be seen in Figure 6.10.

For map placing and labeling of targets, I predefined an acceptable range of differences between RFED and RW to be answering correctly within one question, the smallest measurable unit, or 16.7%. See Figure 6.9. The 95% confidence interval of the mean difference between the percentage of correctly placed targets was -4.6% to 27.7%, $\bar{x} = 11.5\%$, $SE = 7.7\%$. With 95% confidence, we found that when performing a naïve search, participants in the RFED condition will incorrectly place no more than 4.6% more of the targets compared to participants in RW. Since -4.6% is greater than -16.7%, **I claim that, when performing a naïve search, participants in the RFED condition are no worse at placing targets on a map than participants in the RW condition**.

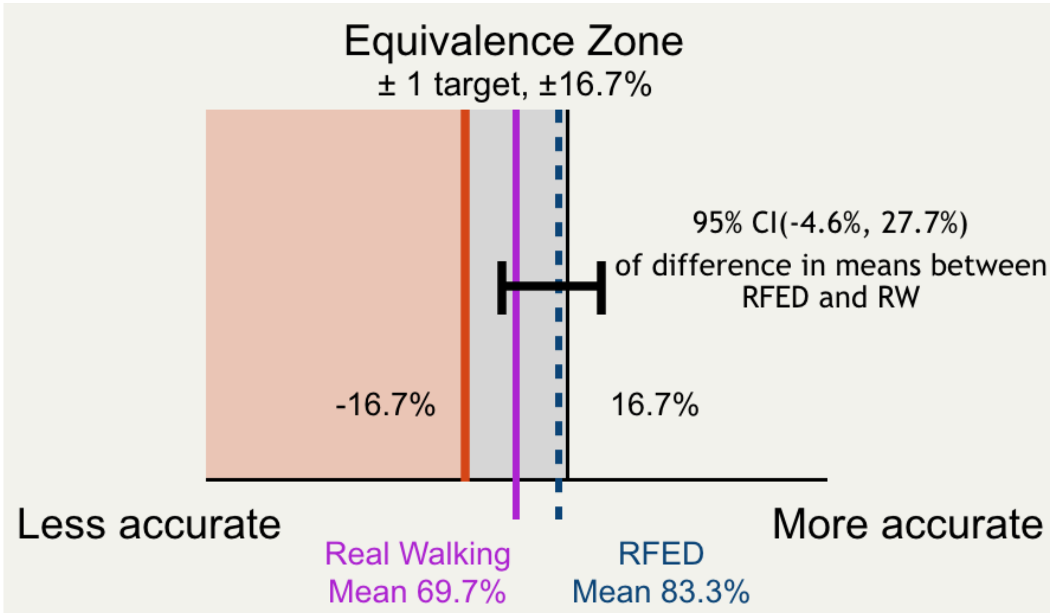

Figure 6.9: Map placing of targets for a naïve search. The 95% CI of the difference of the means of RFED and real walking is the horizontal bar and the "less accurate" zone is on the left. Since the 95% CI is greater than the "less accurate" zone, RFED is "no worse" than real walking.

The 95% confidence interval of the mean difference between interface conditions of the percentage of correctly placing-and-labeling targets was -21.4% and 38.9%, $\bar{x} = 8.63\%$, $SE = 14.3\%$. Since 21.4% is greater than 16.7% **I make no claim about user ability of placing-and-labeling targets between interface conditions when performing a naïve search**.

Conclusion To summarize, when performing a naïve search, participants in the RFED condition traveled longer distances to find targets, but do not revisit more routes than RW participants. The longer travel distance is most likely due to the extra steps required when distractors appear in the RFED condition. An improvement to the distractor implementation may reduce the difference in travel distance, as well as improve the overall RFED interface usability. Additional results suggest that RFED participants are no worse at correctly placing targets on a map compared to RW participants. Based on my hypotheses, (Section 6.1.1) RFED participants' mental models of the VE were no worse than RW participants' mental models.

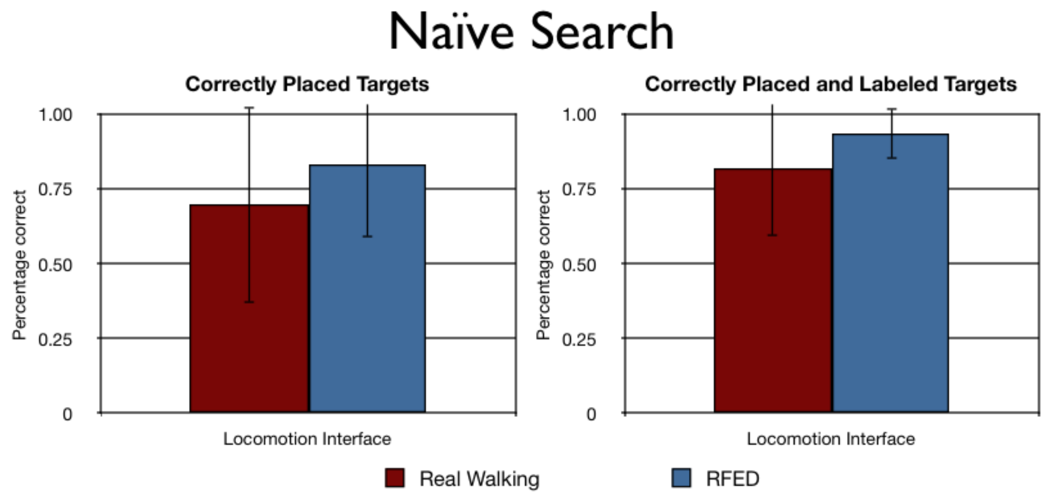

Figure 6.10: The average percentage of correctly placed, and correctly placed-and-labeled targets on maps for RW and RFED during Part 1, a naïve search. With ± one standard deviation error bars.

6.2.2 Part 2: Primed Search

Navigation. The priming path was the same for all participants, so I compared priming path distance between RFED and RW. I found participants traveled significantly greater distances, approximately 20% longer, in RFED when traveling the same virtual path, $t(19) = 6.07$, $p = 0.023$. That is, the RFED algorithm increases the total distance participants travel when following identical routes. This is likely due to participants taking extra steps when distractors appear. The routes that participants locomoted in both the RFED and RW conditions was directed and identical, therefore participants in the RFED condition are not making wrong turns or revisiting routes, and are not more lost than RW participants. I believe an improved steering algorithm will reduce the number of distractors and thus reduce the difference in distance traveled between RFED and RW.

The virtual and corresponding real routes of a participant in the RFED condition is shown in Figure 6.11. I evaluated the difference between locomotion conditions RFED and RW in distance traveled during the primed search. See Figure 6.12. Based on results from the naïve search, 0.874 meters, an average step size, was subtracted from the trial path length for each distractor appearance during a trial to calculate each adjusted trial distance. This was done to account for participants in the RFED condition having to take an extra

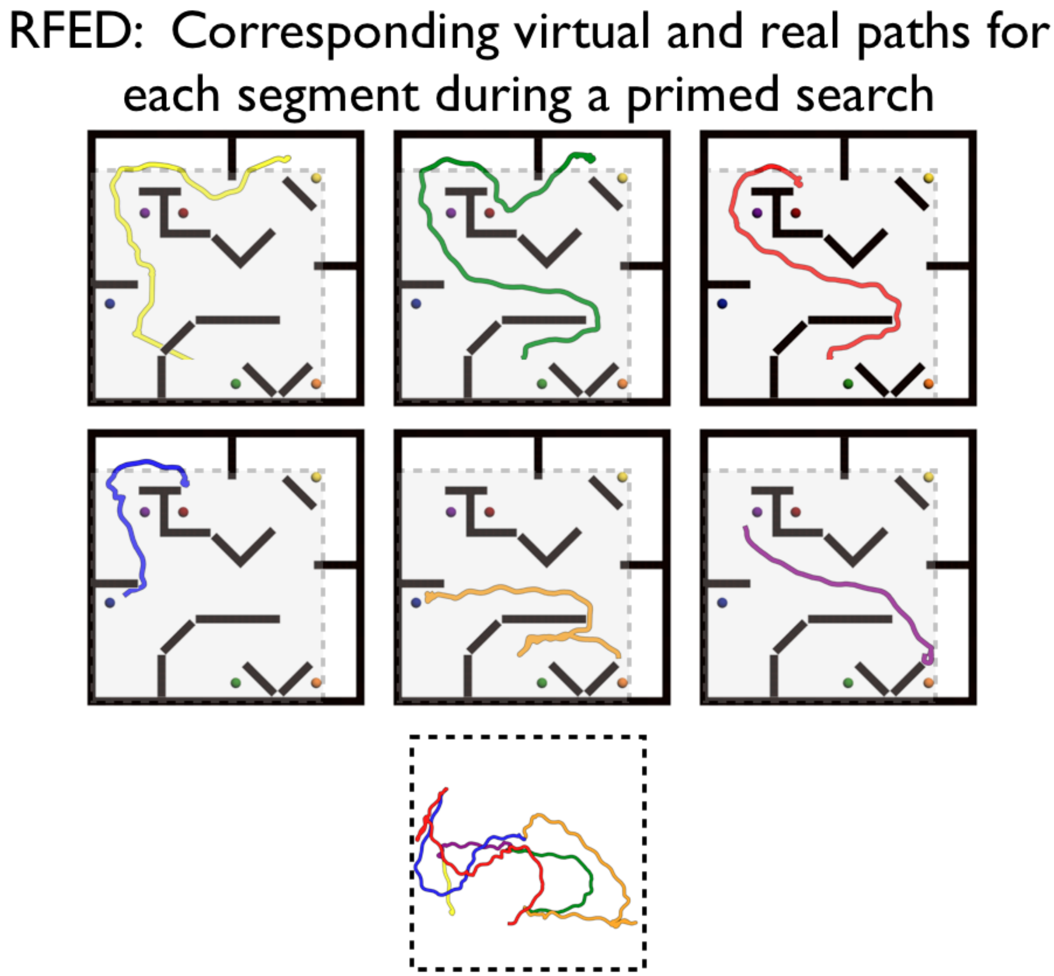

Figure 6.11: The virtual route of an RFED participant to each target during the primed search. The dashed boxes represent the size of the real area participants walked in. The corresponding real routes are displayed in the dashed box at the bottom. Participants started in the bottom left corner of the maze and walked to the yellow target. Participants then walked to targets in the following order: green, red, blue, orange, and purple.

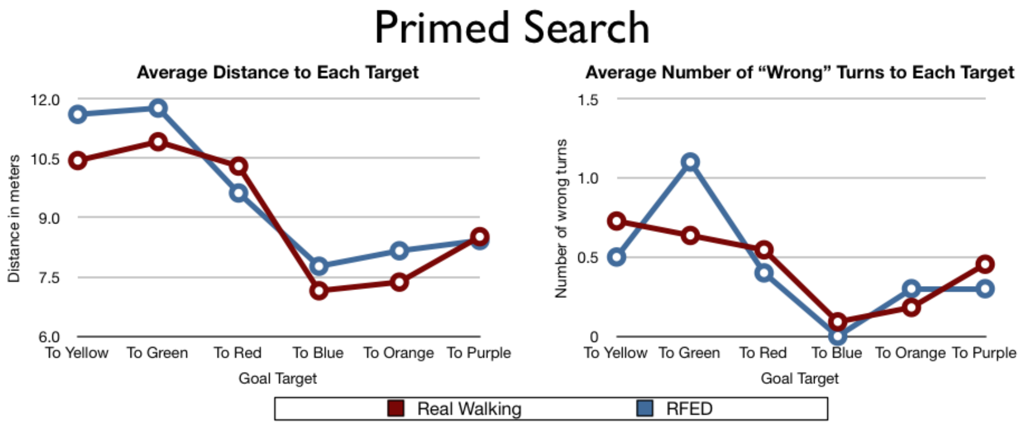

Figure 6.12: The average distance traveled between successive targets, and the average number of wrong turns to each target for RW and RFED during the primed search.

step every time a distractor appeared. I predefined an expectable equivalence region of the difference of the mean distances as 10% of 5.5m, since 5.5m was the average of the shortest distance to each of the six targets. The average actual distance to each of the targets was 9.3 meters. Using a Mixed Model ANOVA with locomotion interface as the between-subject factor, and the adjusted trial distance as a repeated measure, I found the 95% confidence interval of the mean difference to be -0.547 to 1.172 meters, $\bar{x} = 0.312$, $SE = 0.411$. That is, with 95% confidence, participants using RFED will travel no more than 1.172 meters more to each target than with RW. Since $1.172 > 0.55$, **I make no claim that participants using RFED will not travel greater distances than real-walking participants**.

For further route evaluation, I counted the number of wrong turns. I defined a *wrong turns* as when at an intersection not taking the shortest route to the goal target. I assumed that RFED would be no worse than RW if participants made no more than one wrong turn more, the smallest measurable unit, when compared to RW. Using a Mixed Model ANOVA with locomotion interface as the between-subjects factor, and the number wrong turns locomoting to each target as the repeated measure, I found the 95% confidence interval of the mean difference to be -0.221 to 0.213, $\bar{x} = -0.004$, $SE = 0.104$. Since, with 95% confidence, participants using RFED will make no more than 0.213 wrong turns, which is less than the predefined equivalent interval, **I claim that participants using RFED do not make more wrong turns than participants who really walk**.

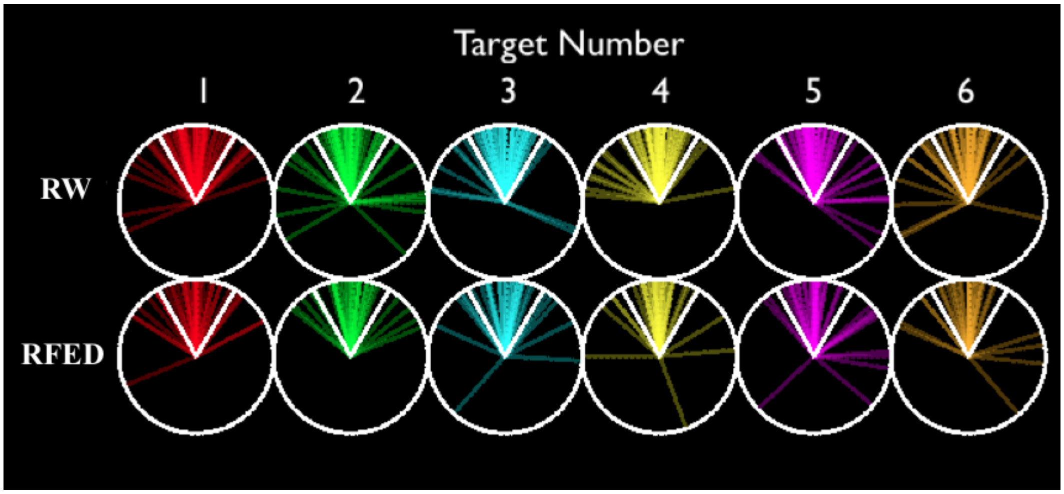

Figure 6.13: Pointing data from all participants, with RW and RFED as separate rows. Each circle is the composite data of the differences in degrees between pointing directly to a specific target (denoted as 12 o'clock) and where the participant actually pointed. Each circle contains all pointing data to the specific numbered and colored target, from all participants in the corresponding condition. The white angle lines provide a reference at $\pm 30°$.

Wayfinding. The composite results of all pointing data are shown in Figure 6.13. Results from (Grant and Magee, 1998) suggest that in the real world people point within $\pm 33°$ of a target, and $\pm 66°$ when in a VE. Based on these results, I predefined as a conservative estimate that RFED would be no worse than RW if participants were able to point within $\pm 15°$ of those in RW. I evaluated the value of the absolute angular difference between pointing direction and direction to target location using a Mixed Model ANOVA with locomotion interfaces as the between-subjects factor and pointing data as the repeated measure. Results show that the 95% CI of the mean difference between participant pointing is $(-8.074°, 8.165°)$, $\bar{x} = 0.045$, $SE = 4.133$. With 95% confidence, participants will point no less than $-8.074°$ less and no more than $8.165°$ more when using RFED compared to RW. Since $-8.074° > -15°$ and $8.165° < 15°$, **I conclude that pointing ability is equivalent between RFED and RW**.

The total time participants took to point to targets can be seen in Figure 6.14. Since audio instruction time was $3s$ and the number of the target was given at $1s$, I predetermined that if participants could point within $1s$, RFED was no worse than RW. I used a Mixed Model ANOVA with locomotion interface as the between-subjects factor and time-to-point

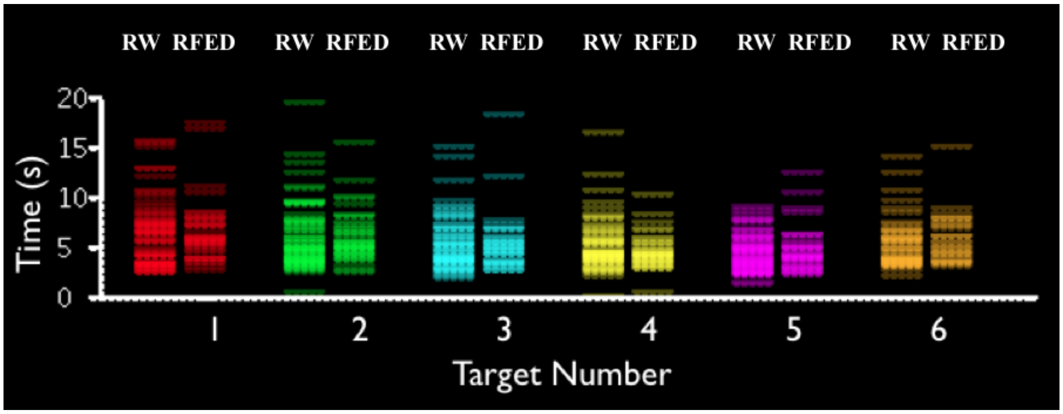

Figure 6.14: Total pointing time from all participants, for RW and RFED conditions. Each column is the composite data, from all participants in RW and RFED, of the total time taken to point to the target of the specified color.

as the repeated measure, and found the 95% CI for the mean difference between condition pointing time was (-1.153s, 0.811s), $\bar{x} = -0.156$, $SE = 0.492$. With 95% confidence, participants using RFED will take no more than 0.811s more to point to a target than RW. Since 0.811s is less than 1s, **I claim that RFED is no worse than RW for time taken to point to targets**.

Map data was calculated in the same way as in Part 1 with acceptable difference of one target, or 16.7%. The mean difference between percentage of correctly placed targets has 95% CI (-35.2%, 36.4%), $\bar{x} = 0.6\%$, $SE = 017.1\%$. Since 35.2% is greater than 16.7%, **I make no claim about target placement**. The mean difference between conditions of the percentage of correct placing-and-labeling of targets was, 95% CI (-7.0% and 21.0%), $\bar{x} = 7.0$, $SE = 6.9$, see Figure 6.15. With 95% confidence, participants in the RFED score no more than 7.0% lower on the placement of targets than RW participants. Since 7.0% is less than the predetermined acceptable range of 16.7%, **I claim that participants performed no worse in RFED than RW when placing and labeling targets on maps**.

6.2.3 Post Tests

After the experiment, participants completed a modified Slater-Usoh-Steed presence questionnaire (Slater et al., 1994) and a simulator-sickness questionnaire (Kennedy et al., 1993).

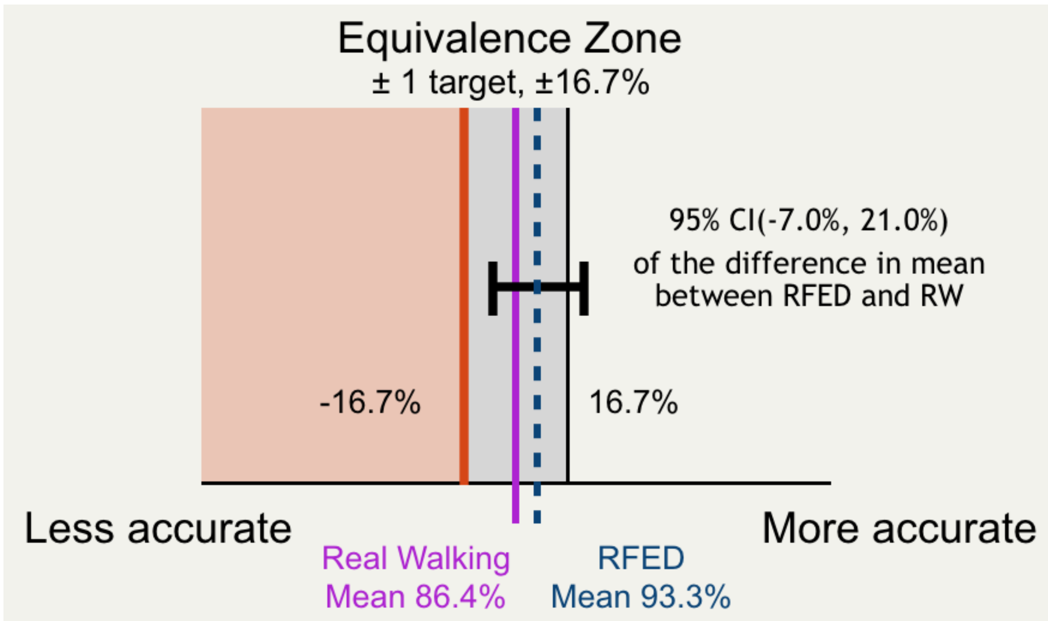

Figure 6.15: Map placing-and-labeling of targets for the primed search. The 95% CI of the difference of the means of RFED and real walking is the horizontal bar and the "less accurate" zone is on the left. Since the 95% CI is greater than the "less accurate" zone, RFED is "no worse" than real walking.

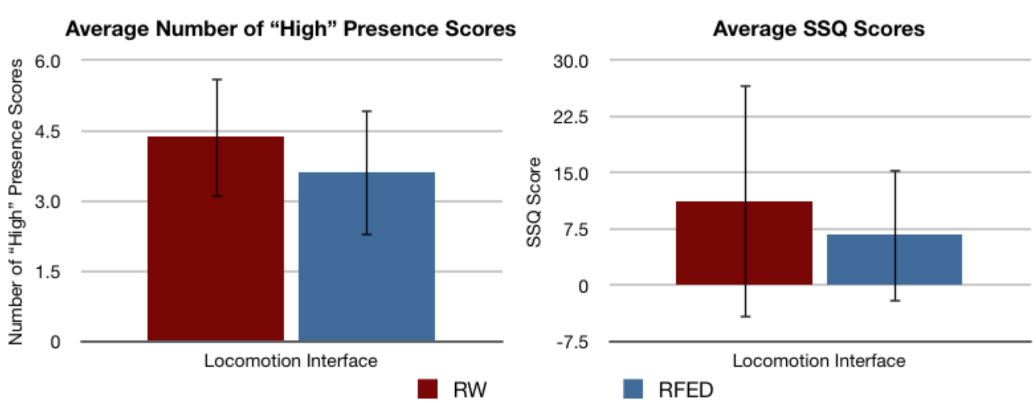

Figure 6.16: The average number of "high" presence scores and the average simulator sickness score for RFED and RW with ± one standard deviation error bars.

See Figure 6.16. On the SUS, the number of "high" scores, scores five or greater, were calculated for each participant. As a conservative estimate, I predetermined that if the number of "high" scores was within one, the smallest measurable unit, that RFED was no worse than RW. The 95% confidence interval of the mean difference between the total number of "high" scores in RFED and RW was -1.968 to 0.441, $\bar{x} = -0.764$, $SE = 0.575$. That is, with 95% confidence, participants in RFED will have no fewer than 1.968 and no greater than 0.441 "high" scores, compared to RW. Since participants could have greater than one fewer "high" score, **I make no claim about RFED being no worse than RW for presence scores**.

Based on my hypothesis that RFED does not significantly increase simulator sickness, I performed a t-test on the simulator-sickness scores calculated from Kennedy's simulator-sickness questionnaire. I found no significant difference simulator-sickness scores between locomotion interfaces, t = 0.91, p = 0.51. Because the variances were large I can make no equivalence claim.

6.3 Conclusion

In this study, I evaluated RFED by comparing it to real-walking, measuring user navigational ability.

For map completion, the results suggest that users are no worse using RFED than RW when placing targets on a map after a naïve search and when placing and labeling targets on a map after a primed search. I also found that participants in RFED can accurately point to previously seen targets equivalently to participants pointing to the same targets using RW. Also, participants using RFED do not take any longer to point to previously seen targets. This suggests that, even with the VE continuously rotating around users, users can wayfind no worse in RFED than RW.

A problem with RFED is that when walking the same path compared to RW, participants walked significantly farther. I believe this is due to the current implementation of the distractor algorithm in RFED, requiring users to take a step backwards (Chapter 4). I believe that improving the distractor algorithm to eliminate the extra step, which caused people take take multiple steps, or improving the steer-to-center algorithm to reduce the total number of distractor appearances, could reduce the distance participants walk using RFED.

I have shown that users can wayfind no worse when using our interface compared to the current best technique, real walking. RFED is designed to enable people to really walk in VEs that are much larger than the tracked space.

CHAPTER 7

An Evaluation of Navigational Ability Comparing Redirected Free Exploration with Distractors to Walking-in-Place and Joystick

In this chapter I present a user study evaluating Redirected Free Exploration with Distractors (RFED) by comparing it to Walking-in-Place (WIP) and Joystick (JS), two common locomotion interfaces used to locomote large-scale VEs. The three interfaces were compared based on navigation, especially including wayfinding metrics. The results from this study support my thesis statement that people navigate better using RFED than WIP and JS interfaces.

The evaluation of RFED compared to WIP and JS was performed in a between-subjects study requiring participants to locomote through virtual mazes. The two experimental mazes used in the study are shown in Figure 7.1. The mazes were $15.85m \times 15.85m$ and designed to be more than twice the dimension (four times the area) of the tracked space. Participants in the RFED condition were restricted to walking in a space that was $6.5m \times 6.5m$, while participants in the WIP and JS conditions where users stay in one place were confined to $1.5m \times 1.5m$ area. See Figure 7.2.

Turning, which stimulates the kinesthetic system, is believed to aid navigation, (Chance et al., 1998; Ruddle and Lessels, 2009). I eliminated turning as a confounding factor by controlling heading direction by physical heading direction, thus requiring participants to physically turn in each locomotion interface, RFED, WIP, and JS.

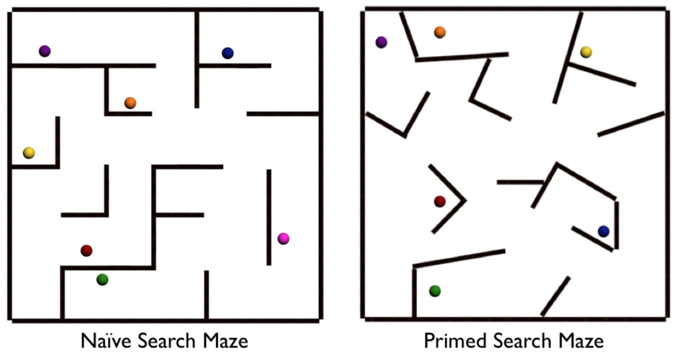

Figure 7.1: The 15.85m x 15.85m mazes used in this study. Left: the maze used during the naive search with seven targets. Right: the maze used during the primed search with six targets. Participants started each maze in the bottom left corner.

7.1 Conditions

7.1.1 Redirected Free Exploration with Distractors (RFED)

A complete description of RFED can be found in Chapter 4.

7.1.2 Walking-In-Place (WIP)

Subjects in a WIP system condition locomoted by stepping in place. Advantages of WIP interfaces include: participants receive kinesthetic feedback from the in-place steps, which moves the viewpoint, and WIP interfaces can be implemented in small spaces. I used the GUD-WIP locomotion interface, because it is the WIP interface that models and most closely simulates real-walking (Wendt et al., 2010). Subjects wore shin-guards equipped with Phase Space beacons for direction of shin motion (Figure 7.2). Phase Space cameras encircled the subject. This setup enabled forward direction to be determined by the participant's average shin direction.

7.1.3 Joystick (JS)

Participants in the JS condition controlled forward speed with a hand-held X-Box 360 controller that was spring loaded. Deflection controlled speed. See Figure 3.2. The maximum speed of the participant in the joystick condition was chosen to be an average walking speed of 3 $miles/hour$. Subjects in the JS condition also wore shin guards equipped with Phase

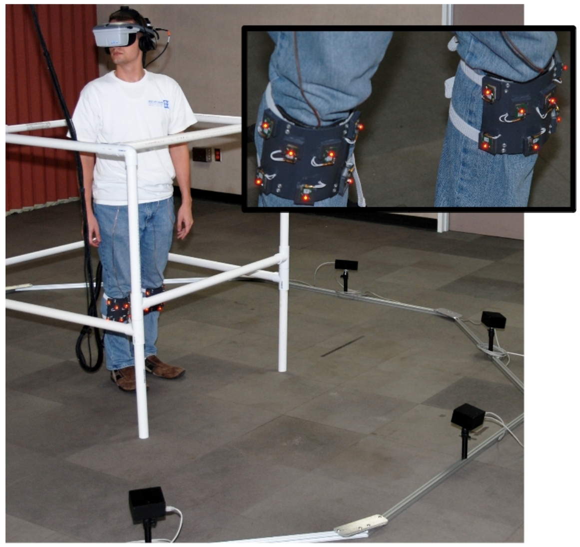

Figure 7.2: The GUD WIP locomotion interface set-up.

Space beacons, just like the participants in the WIP condition. Pushing forward on the joystick translated the participant's viewpoint in the average direction of the participant's shins.

7.2 Hypotheses and Measures

Navigation I used the same hypotheses and measures that I used in Chapter 6, Section 6.1.1.

7.3 Participants

Thirty-six participants, 25 men and 11 women, with average age 26, participated in the IRB-approved experiment. Twelve participants were in each condition (8 men and 4 women in both RDW and WIP, and 9 men and 3 women in JS).

7.4 Equipment

Each participant wore a stereo nVisor SX head-mounted display with 1280x1024 resolution in each eye and a diagonal FOV of 60°. The tracked-space was $9m$ x $9m$ and tracked using a 3rdTech HiBall 3000. RFED participants were restricted to walking in a $6.5m \times 6.5m$ area. The environment was rendered on a Pentium D dual-core 2.8GHz processor machine with an NVIDIA GeForce GTX 280 GPU with 4GB of RAM. The interface was implemented in our locally developed EVEIL intermediate level library that communicates with the Gamebryo® software game engine from Emergent Technologies. The Virtual Reality Peripheral Network (VRPN) was used for tracker communication.

The Walking-in-Place and Joystick systems used an eight-camera PhaseSpace Impulse optical motion capture system with the cameras placed in a circle around the user. The user wore shin guards with seven beacons attached to each shin. PhaseSpace tracked the forward-direction and stepping motion of each leg. The GUD-WIP interface and Joystick direction detection code ran on a PC with an Intel Core2 2.4GHz CPU, NVIDIA GeForce 8600 GTS GPU, and 3 GB RAM.

7.5 Experimental Design

Participants locomoted through three virtual mazes: a training environment and two testing environments (Figure 7.1). The virtual environments were $15.85m \times 15.85m$ mazes with uniquely colored and numbered targets placed at specified locations, see Figure 6.3. All environments used the same textures on the walls and floors, and the same coloring and numbering of targets. The naive search included seven targets and the primed search included six targets. The location of the targets changed between the naive and primed searches. All subjects completed the same trials in the same order to control for training effects. Subjects were randomly assigned to the RFED, WIP or JS condition, and completed all parts of the experiment, including training, in the assigned condition. The experimental design is similar to the design presented in Section 6.1.4.

7.5.1 Training

Subjects received oral instructions before beginning each section of the experiment and were advised to ask questions if they were unclear of tasks. The first environment, the training environment, was a directed maze with all walls placed at 90° angles. Subjects walked through the training environment and used a hand-held tracked device to select each of the seven targets, placed at eye-height and located along the path. When a target was selected, a ring appeared around it and audio feedback was played to signify that the target had been found (Figure 6.3). Subjects were not given performance feedback during any part of the experiment.

After subjects completed the training maze, the head-mounted display was removed and participants were asked to complete a 8.5" $\times$ 11" paper map of the environment. The map representation of the environment was a $16cm \times 16cm$ overhead view of the maze with the targets missing. Participants were given their starting location and maps were presented such that the initial starting direction was away from the user. By hand, subjects placed a dot at the location corresponding to each target and labeled each target with its corresponding number or color.

7.5.2 Part 1: Naïve Search

After training, participants were given oral instructions for Part 1, the naïve search. The maze and target locations for Part 1 can be seen in Figure 7.1. Participants were instructed to find in any order, and remember the location of the seven targets within the maze. Participants were also reminded they would have to complete a map, just as in the training session. As soon as subjects found and selected all targets, the virtual environment faded to white and subjects were instructed to remove the head-mounted display. Subjects then completed a map in the same manner as in the training part of the experiment.

7.5.3 Part 2: Primed Search

After completing the naïve search, subjects were given oral instructions for Part 2, the primed search. The maze and target locations for the primed search can be seen in Figure 7.1. The VE is similar to the maze from Part 1 except that the walls are not all placed at 90° angles. This was done to make the experiment more challenging by removing feedback that enables users to determine cardinal directions from axis-aligned walls. Participants first followed a directed *priming path* that led to each of the six targets in a pre-specified order.

After participants reached the end of the priming path the HMD faded to white, and the participants were placed at the starting point in the VE. Participants in the RFED condition had to remove the HMD and physically walk to the starting location in the tracked-space. Participants using WIP or JS were asked if they wanted to remove the HMD, none did, and then turned in place so they would be facing the starting forward direction in the virtual maze.

Participants were then asked to walk, as directly as possible, to one of the targets in the maze. Participants had to be within an arm's length to select a target. Once the participant reached and selected the specified target, they were instructed, via audio instruction, to point, in turn, to each of the other targets. The audio instructions referenced targets by both color and number. After participants pointed to each other target, they were instructed, via audio instructions, to walk to another target where they repeated the pointing task. If a participant could not find a target within three minutes, arrows appeared on the floor directing the participant to the target. Once the participant reached the target, the experiment continued as before, with the participant pointing to all other targets.

Participants walked to the six targets in the order 3-5-4-1-2-6 and, from each, pointed to each of the other targets in numerical order. At the end of Part 2, subjects had pointed to each target five times, for a total of 30 pointing tasks per subject.

After completing the search and pointing tasks, subjects removed the HMD and completed a map just as in the previous parts of the experiment.

After the experiment, subjects completed a modified Slater-Usoh-Steed Presence Questionnaire (Slater and Steed, 2000) and a Simulator Sickness Questionnaire (Kennedy et al., 1993).

7.6 Results and Discussion

7.6.1 Part 1: Naïve Search

Navigation People who really walk in VEs often walk more slowly than normal and increasing head bobs. This caused head bob signal to appear in the head pose data. See

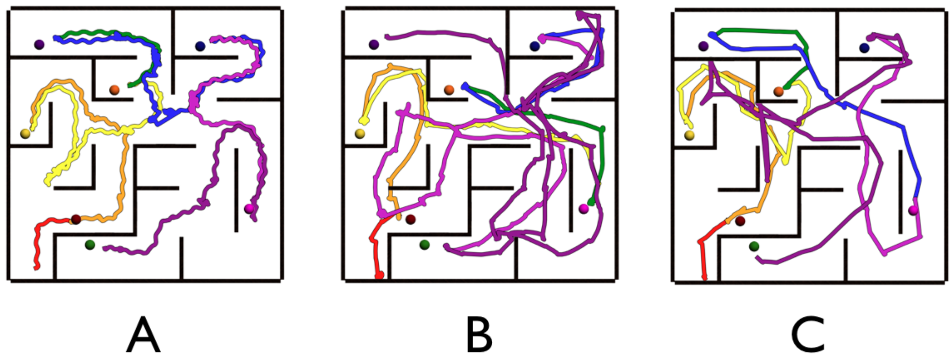

Figure 7.3: The virtual routes of three participants, one using each of the three locomotion interfaces, when performing the naive search. The routes of the median performing participants in each locomotion interface is displayed. **A.** The virtual route a participant took using RFED. Note the side-to-side head bob characteristic of real walking. **B.** The virtual route a participant took using WIP. **C.** The virtual route a participant took using JS.

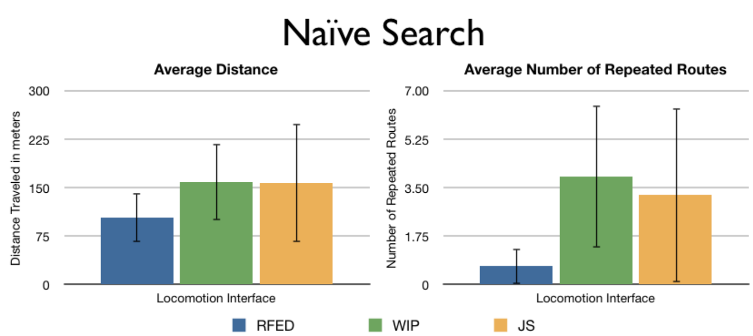

Figure 7.4: The total average distance traveled and the average number of repeated routes, by locomotion interface, when performing the naive search to find seven targets within the maze, with ±1 standard deviation.

Figure 7.3. For this reason all head pose data, including WIP and JS data, were filtered with a box filter over approximately 3 seconds, to remove the head bobbing signals. Participant travel distance was calculated from the filtered head pose data. I assume participants who travel shorter distances have a better spatial understanding of the environment and of previously visited locations. From this assumption, I evaluated the null hypothesis that there was no difference in locomoted distances among locomotion interfaces, Figure 7.4. I used a Mixed Model ANOVA with locomotion interface as the between-subjects variable and distance traveled as the dependent variable and found a significant difference among locomotion interfaces, $F(2,35)=4.688$, $p=0.016$, $r=0.353$.

I performed Tukey pair-wise, post-hoc tests on the distance traveled data, and applied a Bonferroni correction. Participants using RFED traveled significantly shorter distances than participants using either WIP and JS, $p=0.028$ and $p=0.037$ respectively. No significant difference was found in locomoted distance between WIP and JS, $p=0.992$. These results suggest that participants using RFED had a better spatial understanding of the environment.

The number of times participants revisited routes were counted. See Figure 7.4. I interpret revisiting routes of the maze to indicate that participants were more lost, or were having a harder time building a mental model of the environment. I performed a Kruskal-Wallis test on the number of repeated routes and found a significant difference among locomotion interfaces for the number of times participants revisited areas of the maze when performing a naive search, $H(2)=7.869$, $p=0.02$. Pair-wise comparison post-hoc tests were performed and a Bonferroni correction was applied. I found that participants using RFED revisited significantly fewer routes of the maze than participants using WIP, $H(1)=-11.000$, $p=0.026$. This suggests that, participants using RFED were not as lost, or were having an easier time building a mental modal of the environment than participants using WIP. No significant difference was found comparing RFED to JS, or WIP to JS.

Wayfinding I evaluated participants' ability to place and label each virtual target onto a map of the VE. Targets were counted as correctly placed if they were within one meter scaled of the actual target and on the correct side of walls. Targets were counted as correctly

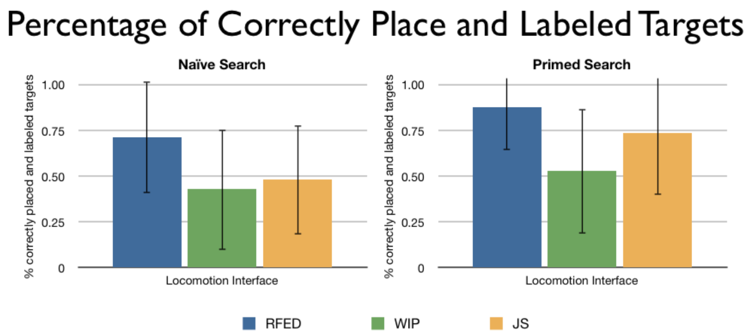

Figure 7.5: The average percentage of correctly placed and correctly labeled targets on paper maps after completing the naive and primed searches. ±1 standard deviation.

labeled if they were both correctly placed and were labeled with either the correct number or color. I performed two Mixed Model ANOVAs with locomotion interface as the between-subjects variable and percentage of correctly placed, and correctly placed and labeled targets as the dependent variables. No significant difference was found among locomotion interfaces in user ability to place targets on maps. However, a trend was found among the three locomotion interfaces comparing participant ability to correctly place and label targets after the naive search, F(2,30)=2.591, p=0.092, $\omega = -0.683$, see Figure 7.5.

Conclusion Summarizing, based on the results from the naive search, RFED participants traveled significantly shorter distances than both WIP and JS participants, and revisited significantly fewer routes in the maze than participants using WIP. These results suggest that, when performing a naive search, participants using RFED had a better understanding of where they had already been within the VE, and had a better spatial understanding of the VE than participants using either WIP or JS.

7.6.2 Part 2: Primed Search

Navigation The real and virtual routes from an RFED participant can be seen in Figure 7.6. The head-pose log files were filtered with a box filter. From the filtered files, I calculated each participant's total travel-distance to find each of the targets for the primed search. I assert that participants who traveled shorter distances to each target were able to build a better mental model of the environment while locomoting the directed training path and while traveling through the environment to each of the targets. I performed a MANOVA with locomotion interface as a between-subjects variable and distance traveled to each of the six targets as a within-subjects repeated measure. See Figure 7.7. I found a significant difference between locomotion interfaces on distance traveled, $F(2,32)=7.150$, $p=0.003$, $r=0.427$. Tukey post-hoc tests show that participants using RFED traveled significantly shorter distances than participants using WIP, $p=0.002$. No other significant results were found. This implies that participants using RFED were better at navigating the VE than participants using WIP.

An additional path data evaluation was performed by using a Kruskal-Wallis test on the total number of wrong turns taken by each participant during the primed search. A *wrong turn* occurs when at an intersection, the participant does not take the shortest route to the current target goal. A significant difference was found between locomotion interfaces, $H(2)=11.251$, $p=0.004$. Pairwise comparisons, reported with a Bonferroni correction, show that participants using RFED made significantly fewer wrong turns than those using either WIP, $H(1)=-13.667$, $p=0.004$, or JS, $H(1)=-10.708$, $p=0.038$. No significant difference was found between JS and WIP users, $H(1)=2.958$, $p=1.00$. These results suggest that participants in RFED had a better understanding of where they were going within the virtual maze, and had a better mental model of the environment, after receiving the same amount of training as participants in WIP and JS interfaces.

Analysis of the routes taken to each individual target among locomotion interfaces show significant difference between walking to the red (No. 1), and green (No. 2) targets, $H(2)=6.505$, $p=0.039$, and $H(2)=8.881$, $p=0.012$ respectively. Post-hoc tests reveals that participants using WIP made significantly more wrong turns when navigating to these two targets than participants using RFED, $H(1)=-9.352$, $p=0.034$, and $H(1)=-11.727$, $p=0.01$ respectively. It is interesting to note that during the directed route portion of the task, participants visited the red (No. 1) target first, and visited the green (No. 2) target last. This may suggest that participants using WIP have problems in the beginning and end of the VE experience. Note: subjects regularly stopped and started their routes many times as they walked the directed path, and participants had to "walk" to get to the red (No.

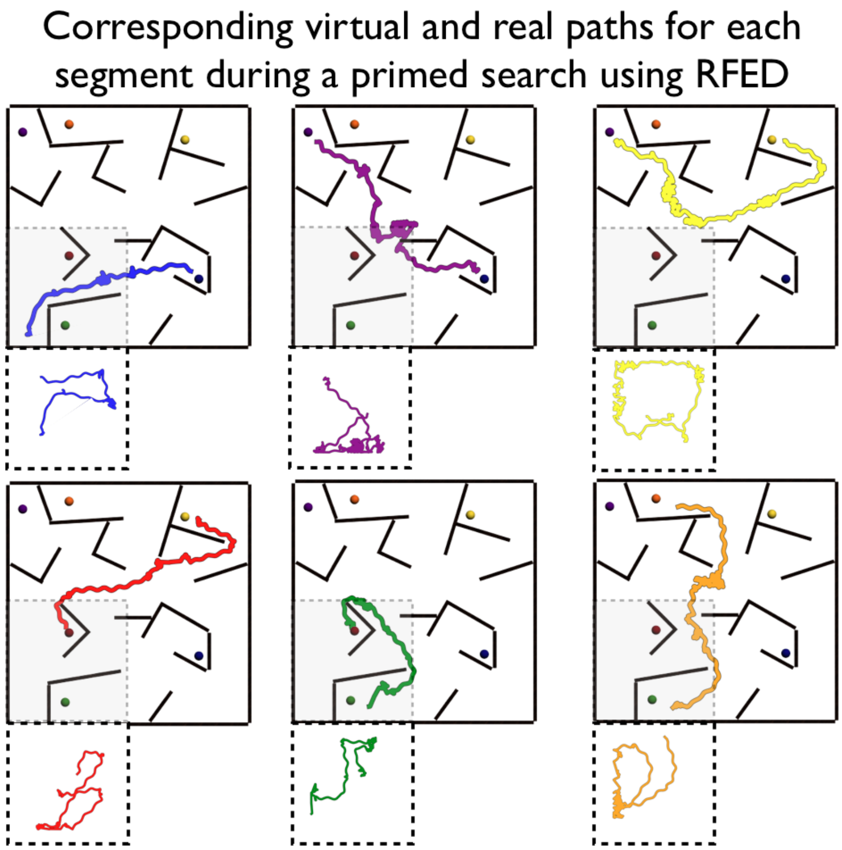

Figure 7.6: The virtual paths with corresponding real path taken by a participant in the RFED condition during the primed search part of the experiment. Participants were really walking in one-quarter of the area of the VE. The large boxes are the virtual routes, and the small dashed line boxes are the corresponding real routes. Routes are displayed to scale.

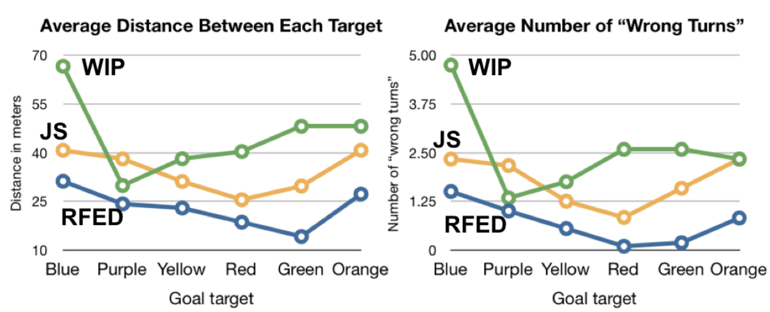

Figure 7.7: The total average distance traveled and the average number of "wrong turns", by locomotion interface when performing the primed search to each of the six targets within the maze.

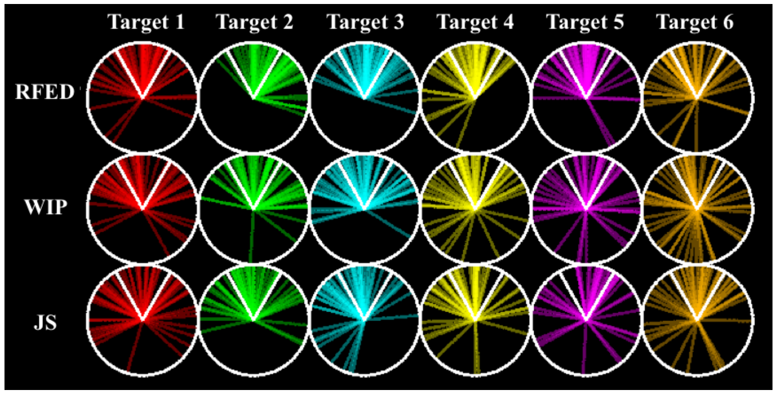

Figure 7.8: The pointing data for all participants to each of the six targets(columns) by each locomotion interface (row). The white lines denote ±30°.

1) target. Further evaluation of WIP interfaces should be explored, specifically looking at cognitive load at the beginning and end of a virtual experience. There was no significant difference for any of the individual routes between JS and RFED or JS and WIP.

Wayfinding During the primed search, when subjects reached a target they then had to point to each other target. See Figure 7.8. Small absolute angular pointing error would suggest that participants have a better understanding of the location of targets. I ran a Mixed Model ANOVA with locomotion interface as the between-condition variable and absolute pointing error to each target as the repeated measure. There was a significant difference among locomotion interfaces for the absolute angular error when pointing to targets, $F(2,28)=5.314$, $p=0.011$, $r= 0.399$. Tukey pair-wise post-hoc tests reveal that participants using RFED had significantly smaller absolute pointing errors than either WIP and JS, $p=0.021$ and $p=0.024$ respectively. There was no significant difference between in absolute pointing error between WIP and JS, $p=0.993$. That is, participants using RFED had significantly better understanding of the location of targets in relation to their current location.

In addition to evaluating pointing ability, I also analyzed how long participants took to point to each target. See Figure 7.9. I hypothesize that participants with a clearer mental model would be able to point more quickly to targets. The first pointing trial was also the first time participants pointed, thus I considered this as a training trial and removed

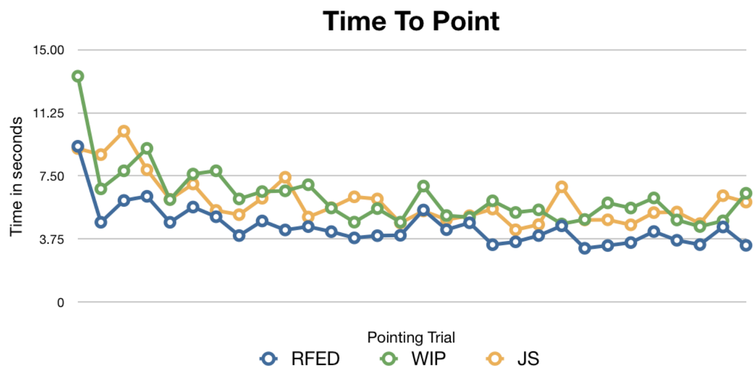

Figure 7.9: The average pointing time for each pointing trial by locomotion interface.

it from the data. I ran a Mixed Model ANOVA with locomotion interface as the between-condition variable, and time to point to each target as the repeated measure and found a trend suggesting a difference in pointing time among locomotion interfaces, $F(1,19)=2.992$, $p=0.074$, $r=0.369$.

Further analysis of the first 14 trials, with the first trial removed, shows a significant difference between locomotion interfaces, $F(2,23)=4.636$, $p=0.02$, $r=0.410$. Tukey post-hoc tests show a significant difference between RFED and either WIP and JS, $p=0.031$ and $p=0.050$ respectively. This suggests that participants using RFED were more confident in pointing ability when compared to participants in WIP and JS, during the first half of the primed search. This result may imply that participants using RFED train faster than participants in either WIP or JS conditions. Further studies should evaluate training time among RFED, WIP, and JS.

I compared the difference in map completion ability among locomotion interfaces, see Figure 7.5, and found a significant difference among interfaces in participant ability to correctly place and label targets, $F(2,30)=3.534$, $p=0.042$, $\omega = -0.603$. Tukey pair-wise post-hoc tests revealed a significant difference between RFED and WIP in correctly placing and labeling targets on maps after completing the primed search part of the experiment, $p=0.034$. No other significant differences were found.

Conclusion The primed search results suggest that participants using RFED navigate and wayfind significantly better than participants using WIP or JS. RFED participants travel shorter distances than participants using WIP, suggesting that RFED participants have a better spatial understanding of the environment and consequently walk more directly to targets. Additionally, participants using RFED make fewer wrong turns than either WIP and JS participants, implying that RFED participants walk more directly to the goal targets, and hence are better at navigating the environment.

Participants in RFED were significantly better at wayfinding than participants in WIP or JS. RFED participants had significantly smaller absolute pointing errors than those using either WIP and JS. In addition to pointing to targets more accurately, participants using RFED are also better at placing and labeling the targets on maps than participants using WIP. This further suggests that participants in RFED develop a better mental model than WIP participants.

Finally, RFED participants point more quickly to targets in the beginning of the experiment than participants in both WIP and JS, suggesting that participants using RFED build mental models faster. Overall, participants using RFED point to targets more accurately, complete maps with fewer mistakes, and are quicker at pointing to targets in the first half of the experiment.

7.6.3 Post Tests

After completing the final map, participants were asked to estimate the size of the VEs compared to the size of the tracker space they were currently in. See Table 7.1. Subjects were told that all three environments were the same size and were given the dimensions of the tracked space. I found a significant difference between VE size predictions based on locomotion condition, $F(2,31)=6.7165$, $p=0.006$, $r=0.742$. Tukey pair-wise post-hoc tests reveal differences between RFED and both WIP ($p=0.033$) and JS ($p=0.007$). The results suggest that people have a better understanding of VE size when using RFED than with either WIP or JS.

One additional factor was that participants in the RFED condition saw virtual bars in the environment that represented the location of the bounds of the real lab. This "real world"-sized reference gave people in the RFED condition an advantage in estimating the

Table 7.1: The average VE size estimate and area underestimate by locomotion interface.

Locomotion Interface	Dimension Estimate	Area Underestimate (%)
RFED	15.0 m x 15.0 m	10%
WIP	10.5 m x 10.5 m	56%
JS	9.1 m x 9.1 m	67%
Actual	**15.85** m x **15.85** m	**0%**

size of the VE. However, two participants in the RFED condition asked to walk around the room before making a guess as to the dimensions of the VE. No participants in JS or WIP asked to walk around the room. This suggests that two participants in the RFED condition realized that their physical walking steps could help measure the size of the VE. The two participants who asked to walk around the room were permitted to walk.

Presence was evaluated using a modified Slater-Usoh-Steed presence questionnaire (Slater and Steed, 2000). The number of "high" presence scores were counted, scores with a 5 or higher, and a Pearson's chi-square test was performed on the transformed data. No significant difference was found among locomotion interfaces and the number of "high" presence scores, $\chi^2(12) = 14.143$, p=0.292.

Participant simulator sickness scores were calculated using Kennedy's simulator sickness questionnaire (Kennedy et al., 1993). A Pearson's chi-square test was performed on the results. No significant difference was found between locomotion interfaces and simulator sickness scores, $\chi^2(40) = 42.800$, p=0.337. These results are discussed further in the next chapter.

CHAPTER 8

Conclusion and Future Work

8.1 Discussion of Results

In this dissertation I have developed and evaluated Redirected Free Exploration with Distractors and have results to support my thesis statement:

A large-scale, real-walking locomotion interface using distractors and redirection enables people to freely locomote in larger-than-tracked-space virtual environments, navigating no worse than real-walking and better than joystick and walking-in-place interfaces.

My results supporting my thesis are in three parts.

1. Develop a large-scale, real-walking locomotion interface using distractors and redirection, Redirected Free Exploration with Distractors (RFED).
   - I developed and evaluated the RFED locomotion interface. RFED incorporates *redirection*—rotation of the VE around the user to redirect the user's path back into the tracked space—and *distractors*—objects or sounds in the VE that stop the user from leaving the tracked space, encourage the user to turn her head, and redirect the user's predicted future path back into the tracked space. A description of the RFED algorithm can be found in Chapter 4.
2. Demonstrate RFED enables people to freely locomote larger-than-tracked-space virtual environments
   - I evaluated RFED's ability to enable people to freely walk in larger than tracked space VEs in two user studies. All participants in the user studies (n=34), pre-

sented in Chapters 6 and 7, freely locomoted in VEs that were 1.5 to 3.9 times the area of the tracked space. One pilot participant was unable to successfully locomote a trial version of RFED.

3. Compare navigational ability among RFED, real walking, walking-in-place, and joystick virtual locomotion interfaces

   - I first compared RFED to the locomotion interface that best emulates real-world locomotion, namely real walking (see Chapter 6). Participants walked through three virtual mazes that were the size of the tracked space ($8m \times 8m$) since real walking, by definition, restricts the size of the VE to the size of the tracked space. I reduced the size of the tracked space that RFED participants were permitted to walk in to $6.5m \times 6.5m$, the minimum tracked space size as determined from pilot experiments.

     Comparison between RFED and real walking on navigation tasks suggests that RFED is "no worse" than real walking. Participants were "no worse" at completing maps of the VE, did not revisit more areas of the VE, did not make more wrong turns, were able to point-to-targets that were out of view, and did not take longer to point to targets than real walking participants.

   - Although the results comparing RFED to real walking were promising, the intent for RFED is to enable people to walk in VEs that are much larger than the tracked space. Since the available tracked-space size was limited to $8m \times 8m$, I could not compare RFED to real walking on this larger scale. Therefore, I compared RFED to locomotion interfaces that are commonly used for locomoting large scale VEs. Specifically, I compared RFED to walking-in-place and joystick interfaces (Chapter 7).

     I compared RFED to walking-in-place (WIP) and joystick (JS) interfaces in a VE, that was 3.9 times the area of the tracked space, in a between-subjects user study. I evaluated the three locomotion interfaces on navigational metrics. The results suggest that in RFED participants traveled significantly shorter distance to either WIP and JS, revisited fewer areas, were better able to complete maps, pointed to hidden targets with more accuracy, and learned the location of targets more quickly.

Two positive results emerged from this work. First, this work presents the first locomotion interface proven to approach the gold standard locomotion interface, real walking and is noticeably as good as real walking on many navigation metrics. Second, the studies in this dissertation show that real walking is significantly better than walking-in-place based on navigation metrics. Researchers have shown that walking interfaces are significantly superior to joystick interfaces on many kinds of measures (Witmer et al., 1996; Usoh et al., 1999). In the studies in UNC's Effective Virtual Environments lab, trends have been seen suggesting that real walking is superior to walking-in-place. But, no previous results have been able to show a statistically significant superiority. The studies presented in this dissertation do that. I used a WIP system, GUD-WIP that has been demonstrated to be state-of-the-art (Wendt et al., 2010). I developed a real-walking system, RFED, that enabled free exploration of larger-than-tracked-space virtual environments. Pairwise comparisons showed that RFED was significantly superior to GUD-WIP on several navigation measures.

8.2 Future Work

RFED is a promising locomotion interface in that it is the first large-scale interface to be proven to approach the gold standards of real walking. However, limitations, namely the user reaching the boundary of the tracked space, currently exist. The ideal view for RFED and RFED-like interfaces is to eliminate the need for user training and instruction and interruption. My inspiration for this work was to enable a user to freely walk around a virtual model of a French cathedral, experiencing the cathedral's beauty as if he were really there.

Further development of RFED will improve usability and approach eliminating the requirement of user training and instruction, and help reach the intended goal of free exploration without user awareness. In this section I discuss future research areas for improving RFED and aid future designers of RFED-like interfaces.

8.2.1 RFED Algorithm

I presented a generic redirection algorithm and the metrics used in the RFED implementation used in this dissertation in Chapter 4. Participants were able to successfully locomote and navigate VEs that were larger than the tracked space, however many participants complained about the distractors in the current RFED implementation. Specifically, participants complained that distractors appeared too frequently, a result of participants reaching the edge of the tracked space too often. Improving the current redirection design and implementation, as well as to determine how to encourage users to quickly turn their heads, will reduce the number of times participants reach the edge of the tracked space, thus reducing the number of distractor appearances and duration.

8.2.1.1 Redirection

Redirection can be thought of as determining the instantaneous rotation, θ_{VE}, of the VE around the user such that the total VE rotation over time is minimized, while the instantaneous per frame VE rotation is maximized. One area of future study is to determine the maximum per frame θ_{VE} to maximize efficient redirection. Research by (Jerald et al., 2008) suggests lower bounds for imperceptible rotation that can be added to the VE during head turns, however the rotation amounts used in Chapter 7 were larger than those presented in (Jerald et al., 2008). Participants in Chapter 7 did not complain about VE rotation, nor did the VE rotation significantly increase simulator sickness compared to participants using WIP or JS interfaces, although the null-hypothesis cannot be proved. This suggests that rotation in RFED does not need to be imperceptible, however an upper bound for redirection is currently not known. Determining the maximum amount of rotation will enable maximum redirection.

Direction prediction One of the hardest parts of redirection is predicting the user's future direction. Having accurate direction prediction enables redirection to minimize VE rotation. Based on user feedback and experimenter observation, I would recommend enabling the user to interact with the system to define her own future direction. This would remove the guess work and inaccuracies from the path prediction algorithm as well as reduce the time required to redirect the user. This technique would be specifically useful for non-naïve users, while providing the benefits obtained of really walking.

Designing path prediction without user input to the system is more challenging. Creating a statistical model of the environment to determine the most likely future user path will drastically improve the current direction prediction algorithm. I also recommend using motion planning algorithms, or previous user path information of specific environments to determine common user paths within environments.

Steering The current implementation of steering always directs the user's predicted future path to the center of the tracked space. Steer-to-farthest-corner, steer-to-circle, or steer-to-moving-targets may steer the user to stay within the tracked space better than steer-to-center. Simulations of different steering algorithms on different virtual paths may provide insight into the best steering algorithm for RFED.

8.2.1.2 Distractors

Distractors enable users to move by really walking in VEs that are larger than the tracked lab space; however, further investigation is needed to determine the potential effects of using distractors. The goals for distractor implementation are, (1.) minimizing distractor appearance frequency and duration, (2.) minimizing user extra work, and (3.) minimizing distractor-related instruction. Future research areas for distractors include:

Minimize awareness. Results from (Bailey et al., 2007) provide promising results for encouraging head turns through image modulations in peripheral vision. These results could remove the requirement of instructing users to watch distractors.

Naïve users. Current distractor implementations require initial instruction causing users to be non-naïve to distractors. Requiring users to be non-naïve may increase cognitive load or have unknown negative effects on usability. Developing distractors for a completely naïve user without requiring initial instruction will enable visually and cognitively imperceptible reorientation. For example, research should determine types of objects in an environment that occasion head turns without user instruction.

Appearance. Results from Chapter 5 suggest that the appearance of distractors has an effect on user preference. Additional study of distractor appearance includes studying animated versus rigid-body distractors, realistic versus non-realistic distractors, and looking at different colors, shapes, sounds, or sizes of distractors.

Minimize cognitive load. Future evaluation of distractors should focus on determining if distractors increase cognitive load. Additional studies may therefore also focus on designing distractors to minimize cognitive load.

Motion paths. Current distractor implementations only move distractors in arcs located directly in front of the user. Further evaluation of different distractor motion paths may reduce user frustration or may result in more effectively encouraging user head turns. Additional research should focus on the motion and appearance of distractors in different parts of the FOV, specifically objects in the user's periphery when using a wide FOV HMD.

Algorithm. Further evaluation of the distractor algorithms discussed in Chapter 4 may determine more efficient algorithms for distracting people to encourage user head turns.

8.2.1.3 Deterrents

Deterrents were implemented as horizontal bars that marked the edge of the physical tracked space (Figure 4.5). Currently deterrents have been implemented as stationary virtual objects, however implementing deterrents as dynamic virtual objects will provide additional ways to "steer" the user away from the edge of the tracked space. One implementation of deterrents could be virtual avatars walking around the environment, such as visitors at a museum or shoppers in a store, to deter the user from the boundary of the tracked space.

8.2.2 Implementation and Evaluation

8.2.2.1 Wide Field-Of-View Head-Mounted-Displays

Wide-field-of-view (150°) HMDs are desired. I used an HMD with a 60° diagonal field of view (FOV). Psychology studies evoke concern that redirection may not work as well with a wide-field-of-view HMD due to optic flow in peripheral vision which guides locomotion (Warren and Kurtz, 1992).

RFED rotates the VE around the user, which effects lamellar optic flow. However, lamellar flow is used to determine heading direction in peripheral vision (Warren and Kurtz, 1992) which may hinder user reorientation. Additional caution concerning simulator sickness or user instability may occur from the discrepancy between user physical motion and inaccuracies between lamellar flow in the user's peripheral vision.

Work by (Razzaque et al., 2002) using redirection in a CAVE with a 205° FOV provides evidence that redirection will turn people in wide-field-of-view HMDs. However, the effect of changes in lamellar flow on locomotion for redirection is currently unknown.

Wide FOV HMDs may allow earlier or more subtle distractors to stimulate head turns.

8.2.2.2 Training

The results presented in this dissertation suggest that people navigate better when using RFED than WIP or JS, however the effect of RFED on training-transfer is currently unknown. A promising future opportunity involves evaluation of the effects of RFED on training-transfer as well as developing distractors that integrate into the environment or aid training.

8.2.3 Projector systems

Work by (Razzaque et al., 2001) provides evidence that redirection works in CAVEs. This suggests that, in addition to HMDs, RFED will work in large projected rooms.

8.2.3.1 Multiuser interfaces

RFED can be adapted to enable multiple users, a common goal for many VEs. Multiple users will be able to walk in the same physical and virtual spaces, however people will be at a different virtual and real proximities to each other. The system would need to steer users away from moving boundaries, the other users, as well as the physical boundary. It would be easiest to implement multiuser interfaces with wireless HMDs with shared audio presence.

8.2.3.2 Size of the tracked space

The tracked space used in the experiments in this dissertation was $6.5m \times 6.5m$. Pilot experiments suggest that, based on the current RFED implementation, using a smaller tracked space will increase user frustration. Determining the smallest usable tracked space, as well as the ideal area of the tracked space is an area for future research. I believe that increasing the size of the tracked space will produce better results for RFED. A larger tracked space will enable more redirection to occur away from the tracked space boundaries. This will enable fewer distractors and thus improve the overall system.

Future work developing RFED-like systems, specifically determining the correct tracked-space size, accurately predicting the future user path, and creating imperceptible distractors that encourage people to rapidly turn their heads while not requiring user instruction, will help achieve the vision of being able to freely explore a virtual model of a French cathedral as if you were really there.

APPENDIX A

Questionnaires

A.1 Modified Slater-Usoh-Steed Presence Questionnaire

1. Please rate your sense of walking in the environment, on a scale of 1 to 7, where 7 represents your natural experience of walking in a real environments.
2. To what extent was the experience within the virtual environment reality for you (7), or not reality for you (1)?
3. When you think back to the experience, do you think of the virtual environment more as images you saw (1) or more as somewhere you visited (7)?
4. Consider your memory of being in the virtual environment. How similar in terms of the structure of the memory is this to the structure of the memory of other places you have been today, where 1 is not similar and 7 is very similar?

   By 'structure of the memory' consider things like the extent to which you have a visual memory of the virtual environment, whether that memory is in color, the extent to which the memory seems vivid or realistic, its size, locations in your imagination, the extent to which it is panoramic in your imagination, and other such structural elements.
5. During the time of your experience, did you often think to yourself that you were actually in the virtual environment (7) or that you were not in the virtual environment (1)?
6. During the time of the experience, which was the strongest on the whole, your sense of being in the virtual environment (7) or of being elsewhere not in the virtual environment(1)?

A.2 Embedded Questions

Did you notice anything unnatural or odd during your virtual experience? Please rate the following:					
	I did not notice or it did not happen	I may have noticed it but I am not sure	I noticed it sometimes but it didn't bother me	I noticed it	Very obvious and took away from my experience
I felt like I was getting bigger or smaller	○	○	○	○	○
I saw the virtual world flicker	○	○	○	○	○
I saw the virtual world getting brighter or dimmer	○	○	○	○	○
I felt like the virtual world was turning me around	○	○	○	○	○
I saw the virtual world rotating	○	○	○	○	○
I saw the virtual world get smaller or larger	○	○	○	○	○

A.3 Simulator Sickness Questionnaire

For each of the following conditions, please indicate how you are feeling right now, on the scale of "none" through "severe". Here are definitions for some of the conditions:

Fatigue Weariness or exhaustion of the body

Eye Strain Weariness or soreness of the eyes

Nausea Stomach Distress

Vertigo Surrounding seem to swirl

Stomach Awareness Just a short feeling of nausea

Fullness of head Sinus pressure

	none	slight	moderate	severe
General Discomfort	○	○	○	○
Fatigue	○	○	○	○
Headache	○	○	○	○
Eye Strain	○	○	○	○
Difficulty Focusing	○	○	○	○
Increased Salivation	○	○	○	○
Sweating	○	○	○	○
Nausea	○	○	○	○
Difficulty Concentrating	○	○	○	○
Fullness of Head	○	○	○	○
Blurred Vision	○	○	○	○
Dizzy (with eyes open)	○	○	○	○
Dizzy (with eyes closed)	○	○	○	○
Vertigo	○	○	○	○
Stomach Awareness	○	○	○	○
Burping	○	○	○	○

A.4 Embedded Questions

Did you notice anything unnatural or odd during your virtual experience? Please rate the following on a scale from 0 to 7. Where 0 = did not notice or happen, 7 = very obvious and took away from my virtual experience.

___I felt like I was turning around

___I saw the virtual world get smaller or larger

___I saw the virtual world flicker

___I saw the virtual world rotating

___I felt like I was getting bigger or smaller

___I saw the virtual world get brighter or dimmer

Bibliography

Bailey, R., McNamara, A., Sudarsanam, N., and Grimm, C. (2007). Subtle gaze direction. page 44.

Bardy, B. G., Warren, W. H., and Kay, B. A. (1996). Motion parallax is used to comntrol postural sway during walking. *Experimental Brain Research*, 111:271–282.

Beck, J. and Gibson, J. J. (1955). The relation of apparent shape to ajpparent slant in the perception of objects. *Journal of Experimental Psychology*, 50:125–133.

Borah, J., Young, L. R., and Curry, R. E. (1979). Sensory mechanism modelling. *Human Resources Laboratory*, pages 78–83.

Bowman, D. A. (2002). *Principles for the Design of Performance-oriented Interaction Techniques*, chapter 13, pages 277–300. The Handbook of Virtual Environments. Lawrence Erlbaum Associates, Mahwah, NJ.

Bowman, D. A., Koller, D., and Hodges, L. F. (1997). Travel in immersive virtual environments: An evaluation of viewpoint motion control techniques. *Virtual Reality Annual International Symposium*, 0:45.

Brooks, F. P. J. (1987). Walkthrough—a dynamic graphics system for simulating virtual buildings. *SI3D '86: Proceedings of the 1986 workshop on Interactive 3D graphics*, pages 9–21.

Bruce, V., Green, P. R., and Georgeson, M. A. (2003). *Visual perception: physiology, psychology and ecology*. Psychology Press, New York, NY.

Chance, S., Gaunet, F., Beall, A., and Loomis, J. (April 1998). Locomotion mode affects the updating of objects encountered during travel: The contribution of vestibular and proprioceptive inputs to path integration. *Presence*, 7(2):168–178.

Christensen, R. R., Hollerbach, J. M., Xu, Y., and Meek, S. G. (2000). Inertial-force feedback for the treadport locomotion interface. *Presence: Teleoperators and Virtual Environments*, 9(1):1–14.

Connect, F. (2010). The human eye [image]. Retreived 2010, http://www.familyconnect.org.

Crowell, J. A. and Banks, M. S. (1993). Perceiving heading with different retinal regions and types of optic flow. *Perception & Psychophysics*, 53(3):325–337.

Darken, R. P., Cockayne, W. R., and Carmein, D. (1997). The omni-directional treadmill: a locomotion device for virtual worlds. *UIST '97: Proceedings of the 10th annual ACM symposium on User interface software and technology*, pages 213–221.

Darken, R. P. and Peterson, B. (2002). *Spatial Orientation, Wayfinding, and Representation*, chapter 24, pages 493–518. The Handbook of Virtual Environments. Lawrence Erlbaum Associates, Mahwah, NJ.

Darken, R. P. and Sibert, J. L. (1993). A toolset for navigation in virtual environments. *UIST '93: Proceedings of the 6th annual ACM Symposium on User interface software and technology*, pages 157–165.

Darken, R. P. and Sibert, J. L. (1996). Wayfinding strategies and behaviors in large virtual worlds. *CHI '96: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 142–149.

Duh, H. B.-L., Parker, D. E., Phillips, J., and Furness, T. A. (2004). "conflicting" motion cues at the frequency of crossover between the visual and vestibular self-motion systems evoke simulator sickness. *Human Factors*, 46:142–153.

Durlach, N. I. and Mayor, A. S. (1995). *Virtual reality: scientific and technological challenges*. Washington, DC: National Academy Press.

Gibson, J. and Dibble, F. (1952). Exploratory experiments on the stimulus conditions for the perception of a visual surface. *Journal of Experimental Psychology*, 43:414–419.

Gibson, J. and Waddell, D. (1952). Homogeneous retinal stimulation and visual perception. *American Journal of Psychology*, 65:263–270.

Gibson, J. J. (1950). *The Perception of the Visual World*. Boston: Houghton Mifflin.

Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.

Goodale, M. and Milner, D. (2004). *Sight Unseen*. Oxford University Press.

Grant, S. C. and Magee, L. E. (Sept. 1998). Contributions of proprioception to navigation in virtual environments. *Human Factors*, 40.3:489.

Harris, M. G. and Carre, G. (2001). Is optic flow used to guide walking while wearing a displacing prism? *Presence*, 30:811–818.

Hicheur, H., Vieilledent, S., and Berthoz, A. (2005). Head motion in humans alternating between straight and curved walking path: Combination of stabilizing and anticipatory orienting mechanisms. *Neuroscience Letters*, 383(1-2):87 – 92.

Hollands, M. A., Patla, A. E., and Vickers, J. N. (2002). "Look where you're going!": Gaze behaviour associated with maintaining and changing the direction of locomotion. *Experimental Brain Research*, 143:221–230.

Hollerbach, J. M. (2002). *Locomotion Interfaces*, chapter 11, pages 493–518. The Handbook of Virtual Environments. Lawrence Erlbaum Associates, Mahwah, NJ.

Hosman, R. and van der Vaart, J. (1981). Effects of vestibular and visual motion perception on task performance. *Acta Psychologica*, 48(1-3):271 – 287.

Interrante, V., Ries, B., and Anderson, L. (2007). Seven league boots: A new metaphor for augmented locomotion through moderately large scale immersive virtual environments. *3D User Interfaces*, pages 167–170.

Iwata, H. (1999). Walking about virtual environments on an infinite floor. *VR '99: Proceedings of the IEEE Virtual Reality*, page 286.

Jerald, J., Peck, T., Steinicke, F., and Whitton, M. (2008). Sensitivity to scene motion for phases of head yaws. *APGV '08: Proceedings of the 5th symposium on Applied perception in graphics and visualization*, pages 155–162.

Kennedy, R., Lane, N., Berbaum, K., and Lilenthal, M. (1993). Simulator sickness questionnaire: An enhanced method for quantifying simulator sickness. *The International Journal of Aviation Psychology*, pages 203–220.

Konczak, J. (1994). Effects of optic flow on the kinematics of human gait - a comparison of young and older adults. *Journal of Motor Behavious*, 26:225–236.

Lee, D. N. and Aronson, E. (1974). Visual proprioceptive control of stance. *Perception and Psychophysics*, 15:529–532.

Lee, D. N. and Lishman, J. R. (1975). Visual proprioceptive control of stance. *Journal of Human Movement Studies*, 1:87–95.

Lishman, J. R. and Lee, D. N. (1973). The autonomy of visual kinaesthesis. *Perception*, 2:287–294.

Metzger, W. (1930). Optische untersuchungen in ganzfeld ii. *Psychologische Forschung*, 13:6–29.

NASA (2003). Semicircular canal [image]. Retreived 2010, from: http://weboflife.nasa.gov/images/figure5RevA.jpg.

Nitzsche, N., Hanebeck, U. D., and Schmidt, G. (2004). Motion compression for telepresent walking in large target environments. *Presence: Teleoper. Virtual Environ.*, 13(1):44–60.

Northwestern (2001a). The ear [image]. Retreived 2010, from: http://www.dizziness-and-balance.com/disorders/bppv/otoliths.html.

Northwestern (2001b). The inner ear labyrinth [image]. Retrived 2010, from: http://www.dizziness-and-balance.com/disorders/bppv/otoliths.html.

Peck, T., Whitton, M., and Fuchs, H. (2008). Evaluation of reorientation techniques for walking in large virtual environments. *Virtual Reality Conference, 2008. VR '08. IEEE*, pages 121–127.

Peck, T. C., Fuchs, H., and Whitton, M. C. (2009). Evaluation of reorientation techniques and distractors for walking in large virtual environments. *IEEE Transactions on Visualization and Computer Graphics*, 15(3):383–394.

Peck, T. C., Fuchs, H., and Whitton, M. C. (2010). Improved redirection with distractors: A large-scale-real-walking locomotion interface and its effect on navigation in virtual environments. *IEEE Conference on Virtual Reality*, pages 35–38.

Prokop, T., Schubert, M., and Berger, W. (1997). Visual influence on human locomotion. modulation to changes in optic flow. *Experimental Brain Research*, 114:63–70.

Psotka, J. (1995). Immersive training systems: Virtual reality and education and training. *Instructional Science*, 23, Numbers 5-6 / November:405–431.

Raphael (1483-1520). The school of athens [painting]. Retreived 2010, from: http://www.ski.org.

Razzaque, S. (2005). *Redirected Walking*. PhD dissertation, University of North Carolina at Chapel Hill, Department of Computer Science.

Razzaque, S., Kohn, Z., and Whitton, M. C. (2001). Redirected walking. EUROGRAPHICS 2001 / Jonathan C. Roberts Short Presentation The Eurographics Association 2001.

Razzaque, S., Swapp, D., Slater, M., Whitton, M. C., and Steed, A. (2002). Redirected walking in place. pages 123–130, Aire-la-Ville, Switzerland, Switzerland. Eurographics Association.

Redding, G. (2000). Shape from texture [image]. Retreived 2010, from: www.mind.ilstu.edu.

Robinett, W. and Holloway, R. (1992). Implementations of flying, scaling and grabbing in virtual worlds. *Interactive 3D graphics*, pages 189–192.

Ruddle, R. A. and Lessels, S. (2009). The benefits of using a walking interface to navigate virtual environments. *ACM Trans. Comput.-Hum. Interact.*, 16(1):1–18.

Rushton, S. K., Harris, J. M., Lloyd, M. R., and Wann, J. P. (October 1998). Guidance of locomotion on foot uses perceived target location rather than optic flow. *Current Biology*, 8:1191–1194.

Scarborough, K. (2007). A cross-eye stereo pair [image]. Retreived 2010, from: http://commons.wikimedia.org/wiki/.

Schouten, T. (2003). Texture, distance, and size [image]. Retreived 2010, from: http://www.cs.ru.nl/ ths/ rt2/col/h9/9gebiedENG.html.

Schubert, M., Prokop, T., Brocke, F., and Berger, W. (2005). Visual kinesthesia and locomotion in parkinson's disease. *Movement Disorders*, 20:141–150.

Sinauer Associates, I. (2001). Otolithic membrane and of the utricular macula [image]. Retreived 2010, from: http://www.ncbi.nlm.nih.gov.

Slater, M. and Steed, A. (2000). A virtual presence counter. *Presence*, 9(5):413–434.

Slater, M., Steed, A., and Usoh, M. (1995a). The virtual treadmill: a naturalistic metaphor for navigation in immersive virtual environments. *VE '95: Selected papers of the Eurographics workshops on Virtual environments '95*, pages 135–148.

Slater, M. and Usoh, M. (1993). Presence in immersive virtual environments. *Virtual Reality Annual International Symposium, 1993., 1993 IEEE*, pages 90–96.

Slater, M., Usoh, M., and Steed, A. (1994). Depth of Presence in Virtual Environments. *Presence: Teleoperators and Virtual Environments*, 3.2 MIT Press:130–144.

Slater, M., Usoh, M., and Steed, A. (1995b). Taking steps: the influence of a walking technique on presence in virtual reality. *ACM Trans. Comput.-Hum. Interact.*, 2(3):201–219.

Stanney, K. M., editor (2002). *Handbook of Virtual Environments-Design, Implementation, and Applications*. Lawrence Erlbaum Associates, Mahwah, NJ.

Su, J. (2007). Motion compression for telepresence locomotion. *Presence*, 16(4):385–398.

Sutherland, I. E. (1965). The ultimate display. *IFIP Congress*, 2:506–508.

Treisman, M. (1977). An evolutionary hypothesis. *Science*, 197:493–495.

Ungerleider, L. G. and Mishkin, M. (1982). Two cortical visual systems. pages 549–586.

Usoh, M., Arthur, K., Whitton, M. C., Bastos, R., Steed, A., Slater, M., and Frederick P. Brooks, J. (1999). Walking > walking-in-place > flying, in virtual environments. *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 359–364.

Varraine, E., Bonnard, M., and Pailhous, J. (2002). Interaction between different sensory cues in the control of human gait. *Experimental Brain Research*, 142:374–384.

Waller, D., Hunt, E., and Knapp, D. (April 1998). The transfer of spatial knowledge in virtual environment training. *Presence*, 7, No. 2:129–143.

Warren, W. H. (2004). *Optic Flow*, chapter 84, pages 1247–1259. The MIT Press.

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

Warren, W. H. J., Kay, B. A., Zosh, W. D., Duchon, A. P., and Sahuc, S. (February 2001). Optic flow is used to control human walking. *Nature*, 4 Number 2:213–216.

Wellek, S. (2002). *Testing Statistical Hypotheses of Equivalence*. Chapman and Hall.

Wendt, J. (2010). *Real-Walking Models Improve Walking-In-Place Systems*. PhD dissertation, University of North Carolina at Chapel Hill, Department of Computer Science.

Wendt, J., Whitton, M. C., and Brooks, F. (2010). GUD WIP: Gait-understanding-driven walking-in-place. *IEEE Conference on Virtual Reality*, pages 51–58.

Wiersma, O. B. (2007). Shape from texture [image]. Retreived 2010, from: http://www.ottobw.dds.nl/ filosofie/perspect.htm.

Williams, B., Narasimham, G., McNamara, T. P., Carr, T. H., Rieser, J. J., and Bodenheimer, B. (2006a). Updating orientation in large virtual environments using scaled translational gain. *APGV '06: Proceedings of the 3rd symposium on Applied perception in graphics and visualization*, pages 21–28.

Williams, B., Narasimham, G., Rump, B., McNamara, T. P., Carr, T. H., Rieser, J., and Bodenheimer, B. (2006b). Exploring large virtual environments with an hmd on foot. *APGV '06: Proceedings of the 3rd symposium on Applied perception in graphics and visualization*, pages 148–148.

Williams, B., Narasimham, G., Rump, B., McNamara, T. P., Carr, T. H., Rieser, J., and Bodenheimer, B. (2007). Exploring large virtual environments with an hmd when physical space is limited. *APGV '07: Proceedings of the 4th symposium on Applied perception in graphics and visualization*, pages 41–48.

Witmer, B. G., Bailey, J. H., Knerr, B. W., and Parsons, K. C. (1996). Virtual spaces and real world places: transfer of route knowledge. *Int. J. Hum.-Comput. Stud.*, 45(4):413–428.

Witmer, B. G. and Kline, P. B. (1998). Judging perceived and traversed distance in virtual environments. *Presence: Teleoperators and Virtual Environments*, 7(2):144–167.

Young, L. R. (1967). Some effects of motion cues on manual tracking. *Journal of Spacecraft*, 4(10):1300–3.

Young, L. R., Meiry, J. L., Newman, J. S., and Feather, J. E. (1969). Research in design and development of a functional model of the human nonauditory labyrinths. *Air Force Aerospace Medical Research Laboratory*, pages 68–102.

Zuj, K., Koonstra, K., Siemiginowski, M., Hui, T., and Crawford, J. (2004). The visual brain [image]. Retreived 2010, from: http://ahsmail.uwaterloo.ca/ kazuj/index.htm.