Towards Understanding and Expanding Locomotion in Physical and Virtual Realities

Yashas Joshi

A Thesis

in

The Department

of

Engineering and Computer Science

Presented in Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy (Computer Science) at

Concordia University

Montréal, Québec, Canada

April 2023

CONCORDIA UNIVERSITY

School of Graduate Studies

This is to certify that the thesis prepared

By: Mr. Yashas Joshi

Entitled: Towards Understanding and Expanding Locomotion in Physical and Virtual Realities

and submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy (Computer Science)

complies with the regulations of this University and meets the accepted standards with respect to originality and quality.

Signed by the Final Examining Committee:

_____ Chair
Dr. Lata Narayanan

_____ External Examiner
Dr. Ali Arya

_____ Examiner
Dr. Marta Kersten

_____ Examiner
Dr. Tse-Hsun (Peter) Chen

_____ Examiner
Dr. Anjali Agarwal

_____ Supervisor
Dr. Charalambos Poullis

Approved by _____
Dr. Leila Kosseim, Graduate Program Director
Department of Computer Science and Software Engineering

_____ 2023

_____
Mourad Debbabi, Dean
Faculty of Engineering and Computer Science

Abstract

Towards Understanding and Expanding Locomotion in Physical and Virtual Realities

Yashas Joshi, Ph.D.

Concordia University, 2023

Among many virtual reality interactions, the locomotion dilemma remains a significant impediment to achieving an ideal immersive experience. The physical limitations of tracked space make it impossible to naturally explore theoretically boundless virtual environments with a one-to-one mapping. Synthetic techniques like teleportation and flying often induce simulator sickness and break the sense of presence. Therefore, natural walking is the most favored form of locomotion. Redirected walking offers a more natural and intuitive way for users to navigate vast virtual spaces efficiently. However, existing techniques either lead to simulator sickness due to visual and vestibular mismatch or detract users from the immersive experience that virtual reality aims to provide.

This research presents innovative techniques and applications to enhance the user experience by expanding walkable, physical space in Virtual Reality. The thesis includes three main contributions. The first contribution proposes a mobile application that uses markerless Augmented Reality to allow users to explore a life-sized virtual library through a divide-and-rule approach. The second contribution presents a subtle redirected walking technique based on inattentional blindness, using dynamic foveated rendering and natural visual suppressions like blinks and saccades. Finally, the third contribution introduces a novel redirected walking solution that leverages a deep neural network, to predict saccades in real-time and eliminate the hardware requirements for eye-tracking.

Overall, this thesis offers valuable contributions to human-computer interaction, investigating novel approaches to solving the locomotion dilemma. The proposed solutions were evaluated through extensive user studies, demonstrating their effectiveness and applicability in real-world scenarios like training simulations and entertainment.

Acknowledgments

First and foremost, I would like to express my deepest gratitude to my supervisor, Dr. Charalambos Poullis, for his unwavering support, guidance, and encouragement throughout the course of my research. His invaluable insights, expertise, and mentorship have been instrumental in shaping this thesis and my overall growth as a researcher. I am grateful to the members of my thesis committee, Dr. Marta Kersten, Dr. Tse-Hsun (Peter) Chen, and Dr. Anjali Agarwal, for their valuable feedback, constructive criticisms, and the time they invested in reviewing my work.

I would also like to thank my colleagues and friends at the Immersive and Creative Technologies Lab (ICT Lab) for the stimulating discussions, collaborative atmosphere, and camaraderie that have enriched my research experience.

My heartfelt appreciation goes to my parents, Navnit and Aaradhana Joshi, for their constant love, encouragement, and understanding throughout my academic journey. Their unwavering belief in me has been a source of strength and motivation, and I am forever grateful for their support. A special thanks to my beautiful fiance Khyati Raval, lovely sister Urvi Joshi, cousin Nikunj Joshi, sister-in-law Urmi Joshi and my dear friend, Dhaval Chavada, for their unwavering support throughout the research process.

This research is based upon work supported by the Natural Sciences and Engineering Research Council of Canada Grants No. N01670 (Discovery Grant).

Last but not least, I would like to extend my gratitude to all the participants who took part in the user studies, without whom this research would not have been possible.

This thesis is dedicated to all those who have inspired, supported and believed in me throughout this journey.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Immersive technologies, such as augmented reality (AR) and virtual reality (VR), have come a long way since Ivan Sutherland's pioneering work on the first head-mounted display in 1968 I. Sutherland (1968). Today, AR and VR systems are becoming increasingly prevalent in various aspects of our lives, from entertainment and education to professional training and healthcare. These advances are made possible by the rapid development of graphics processing units, computer vision algorithms, and display technologies, allowing for the seamless integration of computer-generated content with our real-world surroundings Chatzopoulos, Bermejo, Huang, and Hui (2017); A. Dash, Behera, Dogra, and Roy (2018).

There are numerous platforms for delivering immersive experiences, including mobile devices, smart glasses, and head-mounted displays (HMDs). The extent to which virtual content is blended with real-world elements is determined by the Reality-Virtuality Continuum proposed by Milgram et al. Milgram, Takemura, Utsumi, and Kishino (1995). This continuum ranges from fully real environments to entirely virtual ones, with various degrees of mixed reality in between Chatzopoulos et al. (2017).

While AR and augmented virtuality (AV) focus on incorporating virtual elements into the real world or the virtual environment, respectively, VR aims to provide a fully immersive experience, replacing the user's field of view with computer-generated content Michael et al. (2019). However, despite the potential of VR to revolutionize various aspects of our lives, it also comes with several

inherent challenges that need to be addressed to ensure seamless, engaging, and comfortable user experiences.

One of the most critical aspects of VR is the sense of presence – the feeling of truly being inside the virtual environment. Achieving a high degree of presence requires careful consideration of various factors, including the user's perception of space, body, and interactions with the virtual world. To further enhance the sense of presence, developers must design natural and intuitive interaction methods, such as gestural interfaces or natural locomotion, as alternatives to traditional controller-based interactions.

This thesis investigates one of the most critical aspects of VR interactions – locomotion. Specifically, it addresses the challenges associated with navigating virtual environments that exceed the available physical tracked space (PTS) and proposes novel solutions to enhance the user's sense of presence, immersion, and comfort.

By delving into the locomotion dilemma in VR, this research aims to provide insights into the development of natural and intuitive navigation techniques that can be employed in a wide range of VR applications, ultimately enhancing the overall user experience and unlocking new possibilities for immersive technologies in various domains.

1.1 The Locomotion Dilemma in Virtual Reality

In our initial work with life-sized virtual libraries, we found that the immersive technologies faced a terrible dilemma that needed immediate attention Joshi and Poullis (2020c). The ability to move and navigate within a virtual environment is a crucial aspect of the VR experience. However, the design and implementation of locomotion techniques in VR present a unique set of challenges, particularly when it comes to preserving the user's sense of presence and immersion while also ensuring their comfort and safety.

One of the primary concerns in VR locomotion is the mismatch between the size of the virtual environment (VE) and the physical tracked space (PTS). In many cases, the VE is significantly larger than the PTS, meaning that users must find a way to traverse extensive virtual spaces within a confined real-world area. Traditional locomotion methods, such as joystick-based movement or

teleportation, may allow users to navigate the VE but often result in a reduced sense of presence and immersion due to their unnatural feel and potential to induce simulator sickness.

Natural walking, on the other hand, provides a more intuitive and realistic mode of locomotion. By using position tracking through external sensors or inside-out tracking, high-end VR systems enable users to navigate virtual environments simply by walking. However, the limited size of the PTS poses a significant challenge when it comes to providing the illusion of exploring boundless virtual spaces.

Redirected walking (RDW) has emerged as a promising solution to the locomotion dilemma in VR. RDW is a set of techniques that enables users to walk in the real world on paths that differ from those they perceive in the virtual environment, creating the illusion of exploring vast virtual spaces within a confined PTS. RDW offers several advantages over other locomotion techniques, including an increased sense of presence, reduced motion sickness risk, and more natural navigation without the need for external hardware interactions.

Despite the potential benefits of RDW, there are still many challenges to overcome in its implementation. Users can become disoriented or concerned about collisions with real-world obstacles, which can hinder immersion. Additionally, the redirection must be carefully designed to avoid detection by the user and adapt to the user's walking pace, which can vary significantly.

This thesis aims to address these challenges by exploring novel RDW techniques and examining their effectiveness in providing natural, intuitive, and immersive locomotion solutions in VR. By tackling the locomotion dilemma head-on, we hope to contribute valuable insights to the development of future VR systems and applications that can deliver seamless, engaging, and comfortable experiences for users.

1.2 Motivation

The growing popularity and advancements in VR technology have led to a wide range of applications, including entertainment, education, training, and healthcare. To fully realize the potential of these applications, it is essential to provide users with a truly immersive and engaging experience. One of the critical components of achieving this goal is the design and implementation of effective

locomotion techniques that allow users to explore and interact with virtual environments naturally and intuitively.

The motivation behind this thesis lies in addressing the limitations and challenges associated with existing VR locomotion methods. Traditional techniques, such as joystick-based movement or teleportation, can hinder immersion and presence due to their unnatural feel, potential for inducing simulator sickness, and reliance on external hardware. On the other hand, natural walking provides a more realistic mode of locomotion but is constrained by the limited size of the PTS.

Redirected walking has shown promise as a solution to the locomotion dilemma in VR. However, current RDW techniques often require specialized equipment, visual augmentation, or expensive eye-tracking hardware. This thesis is driven by the desire to develop novel, accessible RDW techniques that can be applied without reliance on additional equipment or visual alterations, while still providing an immersive, comfortable, and natural locomotion experience.

Our article, "Portal to Knowledge: a Virtual Library using Marker-less Augmented Reality System for Mobile Devices" Joshi and Poullis (2020c), published in SPIE AR/VR/AR, in 2020, unknowingly explored the divide-and-rule method to tackle the locomotion dilemma in mixed reality. By dividing virtual environments into smaller chunks and mapping them repeatedly onto the available physical space, the application allowed users to navigate each chunk one at a time Joshi and Poullis (2020c) while furnishing an illusion of exploring a far bigger virtual space. Therefore, the application served as a foundation for understanding the locomotion dilemma and offered a practical solution for exploring virtual spaces, paving the way for further developments presented in this thesis. This article is included in the appendix A of this thesis.

By exploring innovative solutions to the challenges of VR locomotion, this thesis aims to contribute to the development of more immersive, engaging, and user-friendly VR systems and applications. In turn, these advancements have the potential to greatly enhance the experiences of users across various domains, fostering a more experience-driven global community where the boundaries between the real and virtual worlds continue to blur.

1.3 Challenges

Designing and implementing effective redirected walking techniques for virtual reality environments involve overcoming several challenges, which are discussed in this section.

(1) Limited Physical Tracked Space (PTS): One of the primary challenges in VR locomotion is the limited PTS available for users to navigate. Most users have confined spaces for VR setups, making it difficult to provide a truly immersive experience that allows users to explore vast virtual environments.

(2) Intrusive Hardware and Visual Augmentation: Many existing RDW techniques rely on specialized equipment or visual alterations, which can be expensive, intrusive, and deviate from the content creator's original vision. There is a need to develop RDW techniques that do not rely on such methods.

(3) Real-time Adaptation: To effectively redirect users during locomotion, it is necessary to have real-time, accurate predictions of their actions. Developing systems that can make these predictions quickly and accurately is a significant challenge.

(4) Maintaining Presence: Ensuring that users remain immersed and engaged in the virtual environment is crucial for an effective VR experience. Locomotion techniques should not compromise the sense of presence or cause discomfort, as this can negatively impact the overall experience.

(5) Simulator Sickness: Traditional VR locomotion methods, such as joystick-based movement, can induce simulator sickness due to the visual-vestibular mismatch. It is essential to design locomotion techniques that minimize the risk of motion sickness for users.

These challenges motivate the need to explore novel approaches to VR locomotion and redirected walking, with the goal of improving user experience, accessibility, and overall immersion in virtual environments.

1.4 Contributions

This thesis aims to address the challenges discussed in the previous section and makes several key contributions to the field of virtual reality locomotion and redirected walking techniques. The major contributions are as follows:

(1) Foveated Rendering-Based Redirected Walking: We develop a new redirected walking technique that leverages foveated rendering and the psychological phenomenon of inattentional blindness Joshi and Poullis (2020a, 2020b, 2023a). This approach does not rely on expensive eye-tracking hardware, making it more accessible and cost-effective for a broader range of users.

(2) Deep Learning-Based Saccade Prediction for Redirected Walking: We investigate the potential of deep neural networks for predicting saccades during apparent head rotations, with the goal of developing a more effective and low-budget alternative to eye tracking for redirected walking. Our neural network-based approach leverages the phenomenon of change blindness to redirect users without additional eye-tracking hardware Joshi and Poullis (2023b).

(3) Comprehensive Evaluation and User Studies: To validate the effectiveness of our proposed methods, we conduct a series of user studies that evaluates the efficacy and performance of our RDW techniques Joshi and Poullis (2020b, 2020c, 2023b). These studies provide valuable insights into the effectiveness of our methods in enhancing user experience and immersion in virtual environments.

(4) Scalability and Accessibility: Our contributions focus on creating redirected walking techniques that are not only effective but also scalable and accessible to a wider audience. By leveraging the psychological phenomenon of change and intentional blindness, foveated rendering Joshi and Poullis (2020a, 2020b, 2023a), and deep learning Joshi and Poullis (2023b), we aim to incrementally develop RDW techniques that do not rely on specialized hardware or intrusive visual augmentation, making them more suitable for widespread consumer use.

These contributions represent significant advancements in the field of virtual reality locomotion, addressing the challenges and limitations of existing redirected walking techniques to provide a more immersive, accessible, and enjoyable user experience.

1.5 Organization

The remainder of this thesis is organized into several chapters, each focusing on a specific aspect of our research and contributions. The structure of the thesis is as follows:

- Chapter 1 - Introduction: This chapter presents a comprehensive introduction to the thesis, encompassing the locomotion dilemma in virtual reality, the motivation behind our research, the challenges we address, and our contributions to the field.

- Chapter 2 - Background and Related Work: This chapter provides an overview of the existing literature on virtual reality, locomotion techniques, and redirected walking. It discusses the current state of the art and highlights the challenges and limitations faced by researchers in the field.

- Chapter 3 - Foveated Rendering-Based Redirected Walking: In this chapter, we present our foveated rendering-based redirected walking technique. We discuss the underlying psychological principles, the implementation details, and the results of our user studies, demonstrating the technique's effectiveness in enhancing user immersion without the need for intrusive visual augmentation of the user's field of view. It is based on our publications Joshi and Poullis (2020a, 2020b, 2023a).

- Chapter 4 - Deep Learning-Based Saccade Prediction for Redirected Walking: This chapter delves into our deep learning-based approach for saccade prediction during apparent head rotations. We present the neural network design, the training process, and the results of our user studies, highlighting the potential of our approach as a low-budget alternative to eye tracking for redirected walking. It is based on our publication Joshi and Poullis (2023b)

- Chapter 5 - Conclusion: The thesis concludes with a summary of our contributions and their

impact on the field of virtual reality locomotion. We also reflect on the broader implications of our work and its potential to shape the future of immersive experiences in VR.

The organization of the thesis enables a logical progression through our contributions, facilitating a clear understanding of the challenges we address and the solutions we propose for enhancing locomotion in virtual reality. This structure ensures a comprehensive and coherent presentation of our research, ultimately leading to a greater understanding of our work and its implications within the field.

Chapter 2

Related Work

This chapter provides an overview of influential works and state-of-the-art research most relevant to our research goals. The literature review is organized into sections discussing virtual reality and the locomotion dilemma, redirected walking, and the assessment methodologies.

2.1 Virtual Reality and The Locomotion Dilemma

Virtual Reality (VR) technology enhances our visual perception by entirely replacing real-world images with contextual, computer-generated information. The first example of VR can be traced back to 1968 I. Sutherland (1968), when Ivan Sutherland developed the immersive headset for the Sword of Damocles, as shown in Figure 2.1. One could argue that this was the first example of an Augmented Reality (AR) system since its displays were partially transparent R. Azuma et al. (2001a).

Although virtual reality has been a research topic for several decades, the term gained popularity when Jaron Lanier developed the Dataglove, as shown in Figure 2.2. By incorporating ultrasonic and magnetic hand position tracking technologies, the Dataglove successfully operated as a hand gesture interface device for NASA in its Virtual Interface Environment Workstation (VIEW) system Fisher, Wenzel, Coler, and Mcgreevy (1988). Lanier also developed the first-of-its-kind virtual reality head-mounted display, the EyePhone Darken et al. (1986), in 1986, making it the first commercially available VR solution. Today, consumer-centric solutions range from mobile-based VR experiences

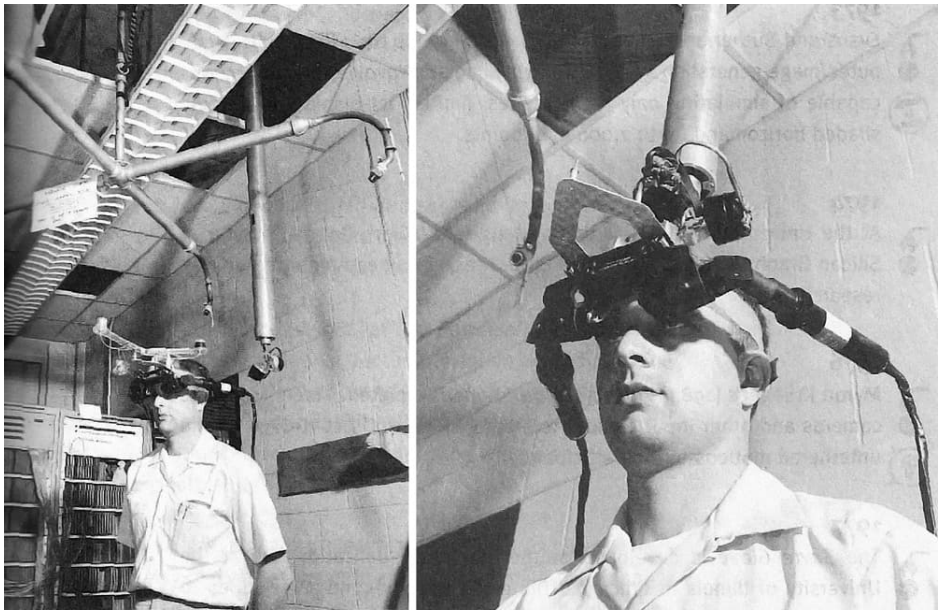

Figure 2.1: Ivan Sutherland wearing the first ever AR headset developed in 1968 I. Sutherland (1968)

to additional accessories like haptic gloves Xiong et al. (2022) and full-body haptic suits Weber et al. (2022).

2.1.1 The Locomotion Dilemma

Although locomotion is not the primary goal of virtual reality, it is considered one of the most critical components of interaction since it is crucial for navigating virtual worlds Cherni, Métayer, and Souliman (2020). If physical and virtual environments match each other, users can easily navigate the virtual space by simply walking. However, exploring a Virtual Environment (VE) larger than the available Physical Tracked Space (PTS) requires alternative methods.

Several locomotion techniques rely on pointing devices or walking in place and have become customary in immersive applications Christou, Tzanavari, Herakleous, and Poullis (2016a). Some of the most common but unnatural means of locomotion are teleportation Langbehn, Lubos, and Steinicke (2018) and flying Sikström, de Götzen, and Serafin (2015). Nevertheless, users find these methods cumbersome and unnatural to work with. These methods also tend to induce disorientation among users while experiencing VR Lathrop and Kaiser (2002).

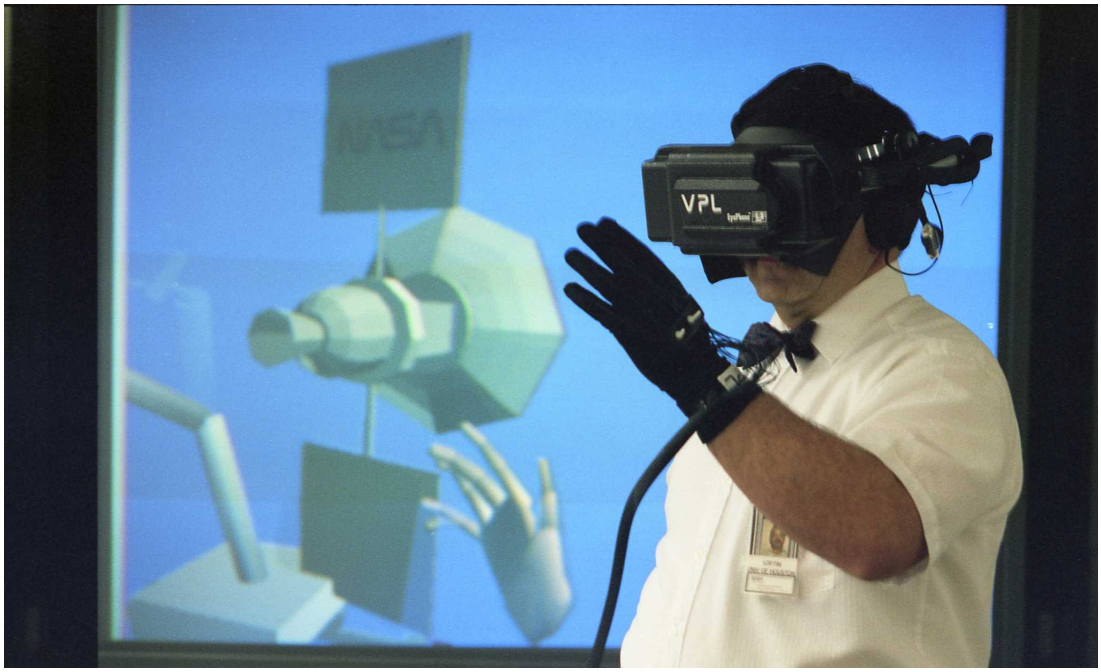

Figure 2.2: DataGlove and EyePhone HMD being used at NASA Darken et al. (1986)

2.1.2 Hardware-Based Solutions

Popular hardware-based solutions include low-friction surfaces Iwata and Fujii (1996), Virtusphere Medina et al. (2008), omnidirectional treadmill Infinadeck - an omni-directional treadmill (n.d.), and suspended walking Walther-Franks et al. (2013). For instance, Souman et al. Souman, Giordano, Frissen, Luca, and Ernst (2010) proposed a solution using an omnidirectional treadmill algorithm to improve the perception of lifelike walking in VR. In this work, a position-based controller was used to determine the user's VR movement speed to rotate the treadmill accordingly. However, it usually overshoots when the user suddenly stops and rocks the user back and forth. Several similar attempts were made to create a perfect omnidirectional treadmill, such as Iwata (1999) and Darken, Cockayen, and Carmein (1997).

In another example, Fernandes et al. Fernandes, Raja, and Eyre (2003) proposed a solution consisting of a life-sized hamster-ball-like structure. Following this, Medina et al. Medina et al. (2008) proposed a similar solution, the Virtusphere. Unfortunately, while each of these techniques addresses the locomotion dilemma in one way or another, they are cumbersome, complicated, limit overall freedom of movement, are extremely expensive, and, most importantly, always pose some

Figure 2.3: Treadmill Infinadeck - an omni-directional treadmill (n.d.)

Figure 2.4: Suspended Walking Walther-Franks et al. (2013)

Figure 2.5: Virtusphere Medina et al. (2008)

form of accidental threat to the users. Additionally, they need to provide vestibular cues equivalent to natural walking Razzaque et al. (2001). Figure 2.3, 2.4, and 2.5 illustrate these solutions.

2.1.3 Software-Based Solutions

Software-based techniques are relatively cost-effective and independent of external hardware, such as the omnidirectional treadmill or low-friction surfaces, which leads to a lower risk of accidental injury. Using natural walking as the primary form of locomotion in PTS Christensen, Hollerbach, Xu, and Meek (2000), software-based techniques provide proper vestibular cues and inertial force feedback, reducing VR sickness LaViola Jr. (2000). Furthermore, with an increased sense of presence in the VE Usoh et al. (1999) and improved spatial understanding Peck, Fuchs, and Whitton (2011); R. A. Ruddle, Volkova, and Bulthoff (2011), natural walking is considered the most preferred locomotion technique in virtual reality. However, the main challenge here is the locomotion dilemma, especially for simulations involving large-scale virtual environments Poullis and You (2009).

Figure 2.6: Applying subtle rotations to the VE to redirect the user Rothacher et al. (2018)

2.2 Redirected Walking

An elegant solution to the locomotion dilemma is redirected walking. As shown in Figure 2.6, the concept of redirected walking is relatively straightforward. Studies have shown that visual senses often dominate over any other human senses. Therefore, the target object in VE is slightly but consistently shifted toward any particular direction. The conflicting visual and vestibular senses lead the users into subconscious correction through motor commands. Hence, when the solutions are carefully and strategically designed, the users unknowingly end up taking curved paths in the PTS without detecting any visual artifacts.

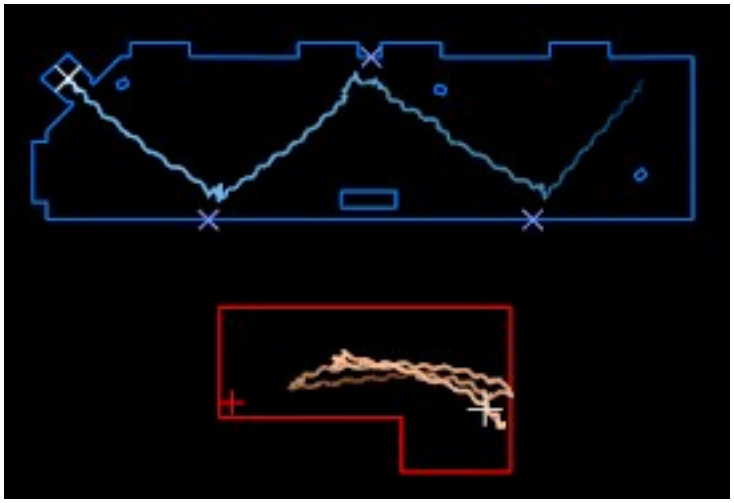

Figure 2.7: Virtual (blue) and physical (red) paths for a user Razzaque et al. (2001)

Razzaque et al. Razzaque et al. (2001) proposed the concept of 'Redirected Walking' for the first time, circa 2001. The authors designed a VR application simulating a mock fire drill with predefined tasks for the users. Ambient sounds and prerecorded instructions were played in the VE to strengthen the immersion. Users followed predefined waypoints inside the application. At every waypoint, the authors applied scaled head rotations from the users' HDM to the entire VE when the users were looking left and right for the next waypoint. Large rotations applied to the VE at every waypoint trick the users into walking back and forth in the PTS while maintaining a zigzag path in the VE. These subtle rotations applied to the VE were enough to convince the users that they had explored a comparatively larger virtual area than the actual available PTS.

Furthermore, the VE rotation algorithm employed three separate components; a constant base-line rotation, a rotation component related to the user's walking speed, and finally, a rotation proportional to the user's angular velocity at each waypoint. The maximum of these three rotations rotates the VE at any given time. Finally, a fail-safe mechanism is implemented that warns users about hitting the wall. A prerecorded message through the virtual radio instructs the users to stop and look left and right before proceeding with the task. When the user rotates his head to look left and right, the system will automatically scale the VE rotations such that the users will now have an ample amount of space in front of them to walk. Figure 2.7 and 2.8 shows virtual (in blue) and physical (in red) paths taken by a user during the experiments. We can see that the users walked

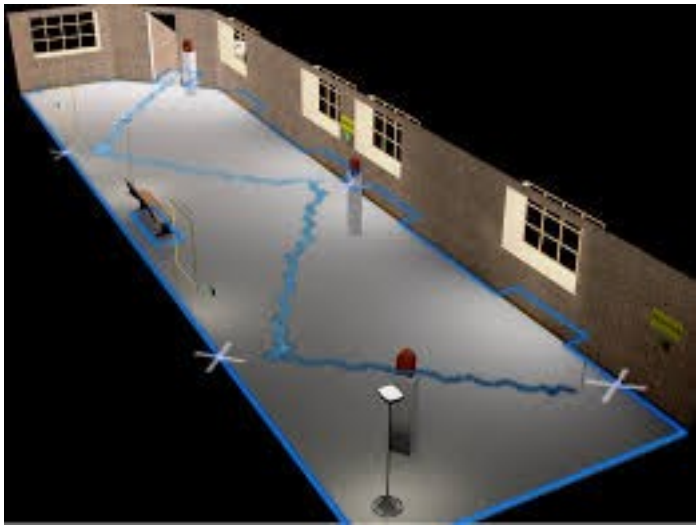

Figure 2.8: Zig Zag path in the virtual environment Razzaque et al. (2001)

along a zigzag path in the VE (blue) while back-and-forth in the PTS (red).

Since the original concept of redirected walking, several attempts have been made to achieve a similar effect. Some researchers even tried to achieve this by incorporating physical props Cheng et al. (2015) or manipulating the entire physical environment E. A. Suma et al. (2011); E. A. Suma, Lipps, Finkelstein, Krum, and Bolas (2012). However, these solutions failed to gain mainstream attention due to their dependencies on factors other than the user's actions. The following successful approaches can mainly be divided into two groups: (a) techniques that use head rotations and translations to dynamically scale the VE based on the requirements as in Azmandian, Grechkin, and Rosenberg (2017); Razzaque et al. (2001); Sakono, Matsumoto, Narumi, and Kuzuoka (2021), and (b) techniques that partially or fully warp the virtual environment Dong, Fu, Zhang, Wu, and Liu (2017); Sun, Wei, and Kaufman (2016).

2.2.1 Steering Algorithms

In order to calculate the amount of redirection, two parameters are required: the target direction of the user (a) in the VE, and (b) in the PTS. There are many methods for predicting the target direction in VE ranging from using the user's past walking direction Zank and Kun (2015), to head rotations Steinicke, Bruder, Kohli, Jerald, and Hinrichs (2008), and gaze direction Zank and Kun (2016). As there are spatial constraints in terms of the available PTS, the user must be steered away

from the boundaries of the PTS i.e. typically walls. To locate the user's target direction in PTS, Razzaque proposed a number of related agorithms in Razzaque, Kohn, and Whitton (2005), namely steer-to-center, steer-to-orbit, and steer-to-multiple-targets: (i) steer-to-center steers the users towards the center of the PTS; (ii) steer-to-orbit steers them towards an orbit around the center of the PTS; and (iii) steer-to-multiple-targets steers the users towards several assigned waypoints in the PTS.

In a recent example, Huiyu et al. proposed a Voronoi-based redirection algorithm to generate a skeleton graph consisting of paths that guide users through certain way-points H. Li and Fan (2020). They also proposed a static graph mapping method that adopts relocation and curvature adjustment to map the skeleton graph of the virtual space to our PTS. In another recent example, Williams et al. proposed a visibility polygons-based redirection algorithm that uses techniques from robot motion planning to compute the redirection gains that steer the user on collision-free paths in the physical space N. L. Williams, Bera, and Manocha (2021b). The algorithm computes the entire physical space visible from the user's perspective and is, hence, walkable. Another ingenious approach by Williams et al. uses an alignment-based redirection controller to guide users in PTS, such that their proximity to any object in the physical space closely matches with the objects in VR N. L. Williams, Bera, and Manocha (2021a). In all of the above cases, the user's direction of movement is constantly changing, which implies that the amount of redirection must also be constantly computed.

2.2.2 Resetting Phase

The resetting phase is one of the most critical aspects of every RDW technique because users may always cross the PTS boundaries. In such situations, a full reset generates more walkable physical space in front of the users.

Several solutions were proposed to address this problem, with the most common being B. Williams et al. (2007): (i) Freeze-Turn is the method where the field-of-view (FoV) remains unchanged while the users turn towards walkable space in the PTS, (ii) Freeze-Back-up is the method where FoV remains unchanged while the user backs-up from the boundary making enough room to continue walking, and (iii) 2:1 Turn is the method where twice the rotational gain is applied to the user's head rotations. For example, if the user turns by 180°, a rotation of 360° is applied in the opposite

direction.

Furthermore, most of the redirected walking techniques follow a stop-and-go paradigm where if the user crosses a boundary, he/she will have to perform a full reset to generate more walkable space in front Azmandian et al. (2017). We follow the same paradigm for our work.

2.2.3 Natural Visual Suppressions

The human visual system is not perfect. Due to several involuntary actions, humans face temporary blindness frequently O'Regan, Deubel, Clark, and Rensink (2000); Rensink (2002); Rensink, O'Regan, and Clark (1997). These involuntary actions are called visual suppressions. Saccades and blinks are two of the most common involuntary visual suppressions Volkmann (1986). Saccades are incredibly unpredictable, rapid, and ballistic eye movements that occur when we abruptly shift our fixation point from one object to another Bahill, Clark, and Stark (1975). These are frequent and fast eye movements. While blinks are much less frequent and slower. In each scenario, our visual perception is suppressed before, during, and after performing the act Volkmann (1986); Volkmann, Riggs, and Moore (1980).

2.2.4 Change Blindness

Change blindness is a fascinating psychological phenomenon where individuals fail to notice significant changes in their visual environment, despite being convinced that they are paying close attention Rensink (2002). This cognitive illusion reveals the limitations of our perceptual and attentional systems and has significant implications for understanding how our brains process the visual world.

Change blindness occurs due to the selective nature of attention and the limitations of our visual short-term memory. Our brains are constantly bombarded with a vast amount of visual information, and as a result, we can only focus on a small fraction of it at any given time. Several studies have shown that users fail to perceive life-sized changes, like the changing positions of doors and walls, while immersed in a VE Steinicke, Bruder, Hinrichs, and Willemsen (2011); E. A. Suma, Clark, Finkelstein, and Wartell (2010). When a change occurs outside our focus, it often goes unnoticed. Furthermore, our visual short-term memory can only hold a limited amount of information, and it

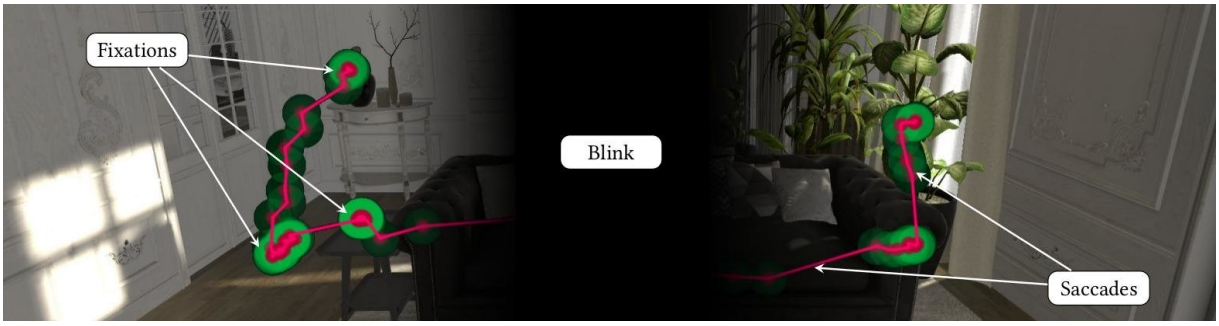

Figure 2.9: Rotating the VE during the change blindness induced from eye blinks Langbehn and Steinicke (2018a)

is prone to interference, making it difficult to detect changes that occur between brief interruptions or changes in viewpoint Rensink et al. (1997).

2.2.5 Blink-Induced Change Blindness for Redirection

Researchers exploit this limitation in a human's ability to perceive a change in the visual scene, especially during visual suppressions like blinks and saccades, for redirection. Due to the dominance of the visual sense over the other human senses, the redirection techniques that partially or fully wrap the VE focus mainly on reducing the effect of subtle visual distractions resulting from repeated redirections by hiding them behind the visual suppressions Volkmann (1986). One such technique is proposed by Langbehn et al. Langbehn and Steinicke (2018a) that leverages the change blindness induced due to naturally occurring blinks to update the VE and redirect the user, Figure 2.9.

The change blindness induced from blinks lasts much longer than the ones due to saccades, making it easier to detect even with readily available off-the-shelf eye trackers. This temporary blindness typically lasts about 100 to 400 ms Ramot, Daniel (n.d.), which is enough to make rotational or translational adjustments to the VE. Authors in Langbehn and Steinicke (2018a) performed extensive experiments to determine the thresholds of rotational and translational gains that can be applied to the VE during a blink without any noticeable visual artifacts. According to O'Regan et al. O'Regan et al. (2000), the longer the suppression, the higher the thresholds, i.e., more redirection per blink.

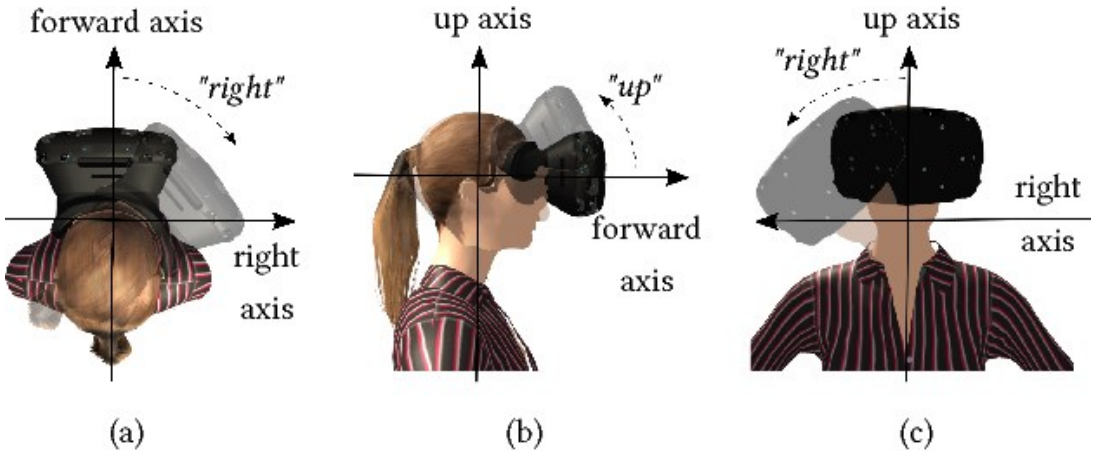

Figure 2.10: Positive rotation offsets for experiment 1 Langbehn and Steinicke (2018a)

Langbehn et al. designed two experiments to determine thresholds for each rotational and translational gain during an eye blink. Sixteen participants with normal or corrected to normal vision performed each of these experiments. In the first experiment, rotational offsets of □ {0, ±3, ±6, ±9, ±12, ±15} degrees on each axis (”up”, ”right” and ”forward”) were applied to the VE during a blink. At the end of each trial, participants answered a two-alternative forced-choice question about the direction of VE rotation. The answer could be either ”left” or ”right” for the ”up” axis (yaw). Similarly, participants also answered questions for rotation on other axes, i.e., ”right” axis (pitch) and ”forward” axis (roll). Every participant repeated each trial six times. Figure 2.10 shows all the possible positive offset rotation directions i.e., {3, 6, 9, 12, 15}. Likewise, all the possible negative offset rotation directions are just the opposite of the ones shown in the figure. Subsequently, all the first experiment procedures were repeated for translation instead of rotation in the second experiment. I.e., for each trial, a translation offset of □ {0, ±3, ±6, ±9, ±12, ±15} cm was applied to the VE. Similar to figure 2.10, figure 2.11 shows all the possible positive translation offset directions.

Upon examining the data collected from the first experiment, the authors found that since humans are used to identifying rotations on the ”up” axis, the yaw’s detection threshold was maximum at 5.3 degrees. While for the ”right” (pitch) and ”forward” (roll) axes, the detection thresholds were 2.1 degrees and 3.5 degrees, respectively. Similarly, from the second experiment, the detection threshold for translation along the forward axis was maximum at 8.7 cm since humans are used to seeing translations on this axis. Furthermore, the translation thresholds for the ”right” and ”up’ axis

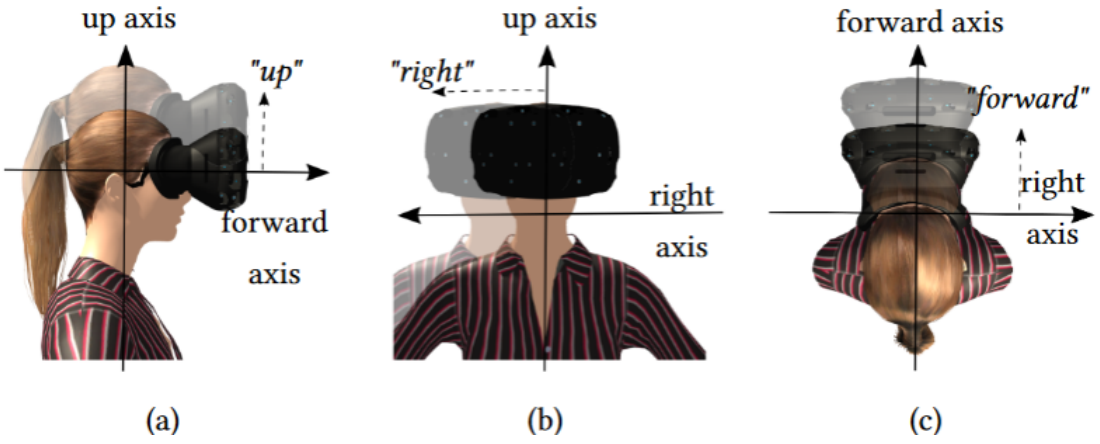

Figure 2.11: Positive translation offsets for experiment 2 Langbehn and Steinicke (2018a)

were 4.5 cm and 4 cm, respectively.

These results are significant since the authors also performed a final confirmatory study with 5 participants to prove the hypothesis of an application in redirected walking. Participants performed a trial ten times of walking along a curved path. They blinked at every ”beep” sound played via headphones. During each trial, the system randomly played the beep twice and applied a rotation gain of 5 degrees during one of the two blinks. Figure 2.12 shows the real and virtual paths during this confirmatory study. To the authors' satisfaction, participants could not detect these gains and confirmed the hypothesis. However, since the blinks do not occur as frequently as saccades, there are still many future research opportunities. A study performed by Bentivoglio et al. Bentivoglio et al. (1997), circa 1997, showed that an average person blinks at an average rate of 17 times per minute, whereas the rate of saccades typically goes up to 1 per second. Hence, using saccades instead of blinks will provide more opportunities for redirection and, therefore, an even better performance.

2.2.6 Saccade-Induced Change Blindness for Redirection

Naturally occurring saccades are ballistic eye movements that a person makes to change the focus from one object to another. For a short period before, during, and after these rapid eye movements, even healthy humans a functionally blind. This temporary blindness induced from saccades is much quicker as compared to blinks. They typically last 100-200 ms, with the eyes' angular speeds reaching up to 900°/s Bahill et al. (1975). However, according to Langbehn et al.

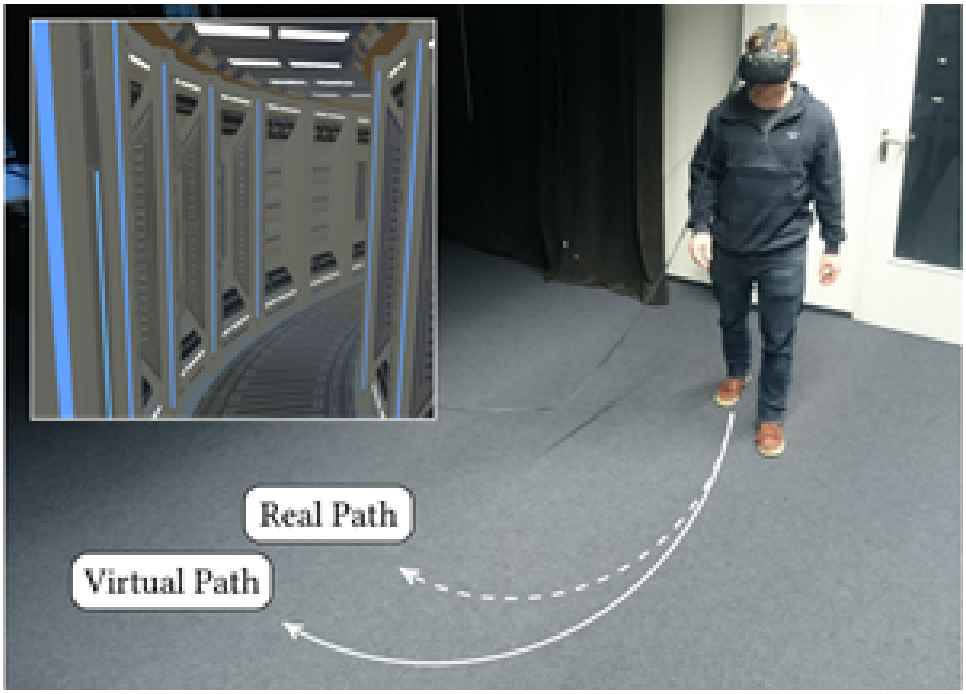

Figure 2.12: Real and virtual paths during the confirmatory study Langbehn and Steinicke (2018a)

Langbehn and Steinicke (2018a), and Bentivoglio et al. Bentivoglio et al. (1997), saccades have lower detection thresholds and are much more frequent than blinks. Therefore, more redirections per unit time are possible when using change blindness induced due to saccades instead of blinks as the visual masks to hide redirection artifacts. Sun et al. in Sun et al. (2018) leveraged the perceptual blindness occurring before, during, and after the saccades to update the VE over several frames. In this section, we present an overview of the redirected walking technique implemented in Sun et al. (2018), followed by a discussion on its importance and limitations, Figure 2.13.

Sun et al. in Sun et al. (2018) conducted several experiments to determine the rotational threshold for saccade and evaluate their proposed redirected walking technique. To estimate the rotational thresholds for saccades, participants immerse themselves in state-of-the-art VR experiences like "Van Gogh Room," "NVIDIA VR Funhouse," and "The Brookhaven Experiment" for 10 minutes straight in a pilot study. During their experience, subtle gaze rotations are applied to the VE at every saccade detected. Using subjective feedback about the visual artifacts like "Yes, I noticed something in the camera orientation", or "No, I do not. They are all the same", the authors adjusted rotations to determine the optimum threshold with no observable visual artifacts. Results indicated

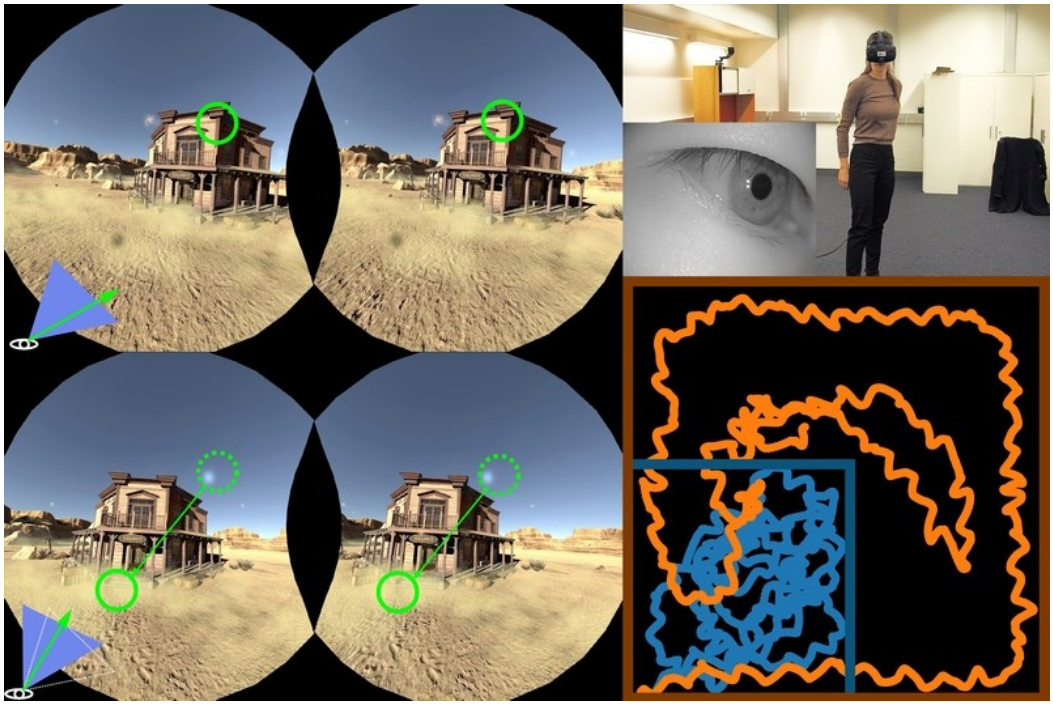

Figure 2.13: Leveraging saccades to apply subtle rotations to the VE. (a) Left: visual distractions to stimulate saccades, and (b) Right: paths in physical (blue) and corresponding virtual (orange) space Sun et al. (2018)

that participants were unable to detect rotations less than 12.6° for saccades with an angular velocity of more than 180°/s. This rotation threshold is also linearly scaled with longer saccades Bolte and Lappe (2015) and used for the redirected walking process.

The redirected walking process in Sun et al. (2018) is performed in three main steps: saccade detection, dynamic path planning, and subtle gaze directions. A saccade is first detected using high-speed eye trackers fitted inside the HMD. Next, the current angular velocity of the users' gaze is estimated using gaze positions from the previous two frames. The system then assumes a saccade if this velocity goes above 180°. Once a saccade is detected, the dynamic path planning algorithm computes the required redirection considering static boundaries and moving obstacles. This algorithm uses an importance-based dynamic sampling mechanism focusing on objects closer to the users and in front of them. Since this process runs on GPU, with multiple threads running in parallel to compute the local importance sampling, it only takes one frame to determine the amount of redirection. Finally, image-space as-well-as object-space subtle gaze directions, as shown in figure

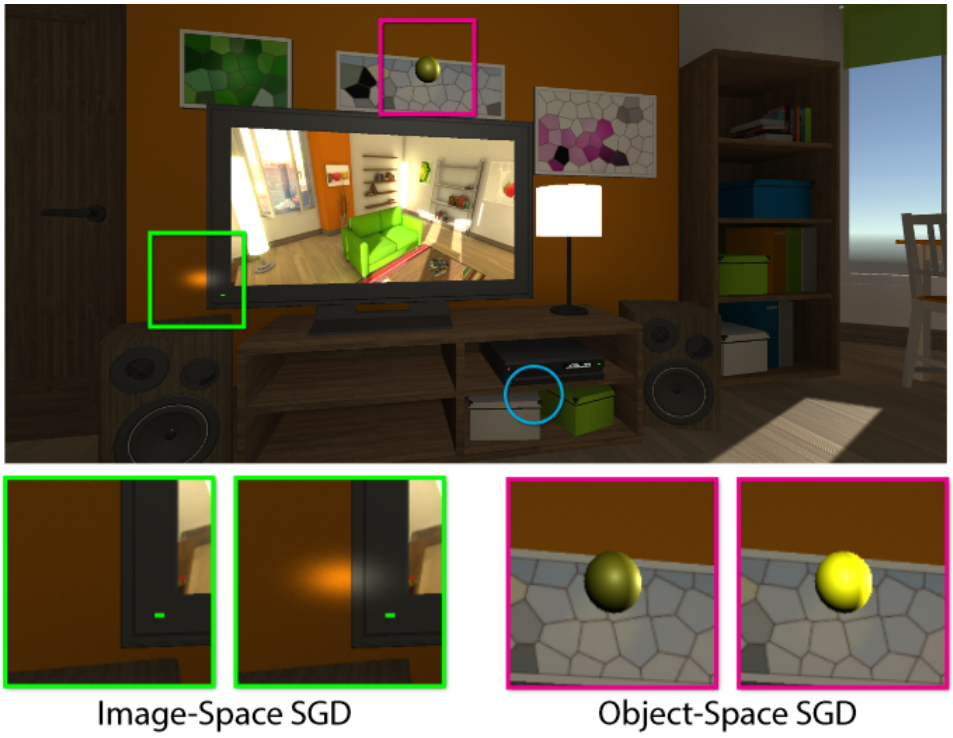

Figure 2.14: Image-space and object-space subtle gaze directions Sun et al. (2018)

2.14, help increase the frequency of redirection by stimulating artificial saccades. Image-space distractors stimulate saccades by temporally modulating pixels around the high-contrast regions in a frame far from the user's current gaze direction. On the other hand, object-space subtle gaze directions modulate the luminance of textures on particular objects in the scene. Upon completing these three steps, the system applies a rotation to the V E over several frames.

Sun et al. performed two experiments to evaluate their proposed technique. The first experiment compared the redirections using head-rotations, saccades, and saccades with image-space subtle gaze directions. The experiment showed a significant effect of using saccades over head-rotations, with an average of 68.7% reduced errors. However, contrary to the literature suggesting the benefits of visual distractors Peck et al. (2011), results from this experiment showed no significant effect of using image-space distractions on redirection. Furthermore, another experiment was devised to compare the effects of image-space and object-space subtle gaze directions on redirection. Results show that the difference between image-space and object-space subtle gaze direction techniques is

not statistically significant. Nevertheless, using both in conjunction showed a significant improvement in redirection than using either image-space or object-space gaze directions.

Although introducing subtle gaze directions can increase redirection opportunities, it also acts as a common disadvantage. Introducing artifacts in the rendered image distracts the user's attention and forces them to look away. Figure 2.13 (left) clearly shows the artifacts (small contrasting flashes in the rendered image) for saccade stimulation. The image-space subtle gaze directions, however subtle, are continuously flashed inside the users' field of view and are disruptive to the cognitive task at hand. These continuous distractions can easily break the immersion.

2.2.7 Dynamic Foveated Rendering

The concept of foveated rendering was first introduced by Reder in 1973 Reder (1973). It is a technique that drastically reduces GPU workload while providing undistinguishable VR experiences. Small changes in our peripheral vision are imperceptible to humans. Thus, the area of the image corresponding to the peripheral vision is rendered at a much lower resolution. In contrast, the area corresponding to the foveal vision is rendered at full resolution Patney et al. (2016).

Although in the past decade, researchers have proposed many software-based techniques to simulate perceptually-guided foveated rendering He, Gu, and Fatahalian (2014); Patney et al. (2016); Stengel, Grogorick, Eisemann, and Magnor (2016); Sun et al. (2018). However, with the recent introduction of Nvidia's Variable Rate Shading NVIDIA (2018), foveated rendering is supported at a hardware level and integrated into the rendering pipeline. This integration has reduced latency and aliasing effects close to zero, as demonstrated by the real-time visualization platform ZeroLight O'Connor, Chris (2018). As a result, we implement dynamic foveated rendering in our projects to reduce GPU workload and get real-time performance.

2.3 Assessment Methodologies

User studies with specific research questions and hypotheses must be designed and conducted to determine the efficacy of any redirected walking technique. Valuable data is gathered from the tasks performed and questionnaires filled by the participants during these experiments, which can

be used to evaluate any proposed technique Langbehn and Steinicke (2018a). The data collected is either subjective or objective. Subjective data is collected from the questionnaires filled before and after performing any predefined task. For example, Sun et al. Sun et al. (2018), and Langbehn et al. Langbehn and Steinicke (2018a) both used subjective feedback as one of their evaluation factors. On the other hand, objective data is collected indirectly by saving the values of predefined variables during the experiments. Both subjective as-well-as objective data is then analyzed to confirm whether or not it supports the proposed hypothesis. Techniques like Analysis of Variance (ANOVA) identify the disparities observed between different variables in the collected data Bolte and Lappe (2015). This chapter presents critical measures for evaluating any redirected walking technique used in the reviewed literature.

2.3.1 Experiment Design and Evaluation

User studies are designed using either within-subjects or between-subjects approaches. In the within-subjects approach, data is collected from all the participants performing every task. While in the between-subjects approach, each participant performs a single task so that only one user interface is exposed to the user. For example, in the first experiment performed by Langbehn et al. Langbehn and Steinicke (2018a), each participant had to perform trials for all the predefined rotational and translational offsets. Hence, the experiment followed a within-subjects approach. Hypothetically, in the case of a between-subjects setting, each participant would have performed trials with only a single offset. However, since the primary goal of this experiment was to compare rotational and translational offsets and determine a threshold for the maximum tolerable rotation and translation gain, using a between-subjects approach did not make any sense. Therefore, either of these approaches could be selected to design a user study depending on the scenario.

The within-subjects approach significantly reduces the number of participants to gather the required data points. Subsequently, it will also reduce the cost of the user study. Additionally, since within-subjects design minimizes the random noise, it is also less likely that an existing real disparity between the predefined tasks will remain obscured Maxwell and Delaney (2018). The works of both Langbehn et al. Langbehn and Steinicke (2018a), and Sun et al. Sun et al. (2018), use the within-subjects approach in their experimental setups. Langbehn et al. use this approach to

determine the rotational thresholds for blinks, whereas Sun et al. use it to evaluate their proposed redirected walking techniques.

Langbehn et al. Langbehn and Steinicke (2018a) and Steinicke et al. Steinicke et al. (2011) also use Two-Alternative Forced-Choices tasks or 2-AFCT for their experiments. In this experimental setup, the participants only have two possible answers choices, out of which only one is correct. For example, in the first experiment performed by Langbehn et al. Langbehn and Steinicke (2018a), participants only had two choices for every question, i.e., ”Left” and ”right” for rotation on up and forward axes, and ”up” and ”down” for rotation on the right axis.

2.3.2 Participants for the User Study

Determining the sample size for the study is a crucial step in the evaluation process. Calculating the number of required participants to support the hypotheses beforehand helps avoid swaying towards the early trends. By doing this, we can also determine the demographics of the group based on our requirements. ”G*Power,” a tool proposed by Faul et al. Faul, Erdfelder, Lang, and Buchner (2007), can calculate the sample size depending on various dependent or independent statistical variables involved in the study.

2.3.3 Evaluation

Once all the user studies have been performed and the necessary data is collected, an in-depth analysis identifies disparities among different variables. Many statistical tests like Z-test, t-test, ANOVA test, and chi-square test help determine subtle disparities in the data. Both Langbehn et al. Langbehn and Steinicke (2018a) and Sun et al. Sun et al. (2018) use ANOVA as their evaluation method. These tests compare means across one or more variables to identify the statistical significance in the obtained results. These variables are dependent on repeated observations. For example, Sun et al. Sun et al. (2018) use ANOVA to determine the statistical significance of the differences between independent variables obtained for the evaluation factors like saving ratios and saccadic angular gains observed during the experiments.

2.3.4 Evaluation Factors

Regarding the works reviewed in chapter two, we discuss some of the main evaluation factors used to distinguish the proposed techniques from other alternatives.

$$\xi = \frac{1 - \frac{\epsilon(P_r)}{\epsilon(P_v)}}{\int_R h(t)dt} \tag{1}$$

Saving Ratio: For the redirected walking technique presented by Sun et al. Sun et al. (2018), authors compared the error areas (i.e., the combined area outside the PTS and inside the obstacles as shown in Figure 2.15) for both physical (P_r - redirected walking enabled) and virtual (P_v - equivalent physical path with redirected walking disabled) paths shown in Figure 2.13 (Right). The total head rotations then normalize this ratio of the error areas to produce the Saving Ratio. The saving ratio ξ is computed using Equation 1. $h(t)$ is the user's head rotations angle at time frame t, and $\epsilon(P_r)$ and $\epsilon(P_v)$ are the error areas for physical and virtual paths respectively. Higher saving ratios indicate effective redirection.

Saving ratios observed by Sun et al. Sun et al. (2018) in their first experiment were as follows, the average saving ratio for saccadic redirection was 3.39e-3 (SD = 1.98e-3), and the average saving ratio for only head rotations was 2.01e-3 (SD = 1.95e-3). The authors obtained a statistically significant difference between these saving ratios and their effect on redirection. It was also evident that this effect was positive since the saving ratio for saccadic redirection was more significant than the saving ratio for non-saccadic (only head rotations) redirection.

The saving ratio is an excellent measure to evaluate any redirected walking technique; however, it has one limitation. Since it depends on the error area calculated for virtual as-well-as physical paths, it will fail when the PTS is confined by physical bounds, like walls, as the users will not be able to walk outside.

Saccadic Angular Gain: Saccadic angular gain is the average rotation applied to any VE during the change blindness induced by one saccade. Typically a system that can handle massive saccadic angular gains is considered a sound redirected walking system since larger rotations applied to the VE per saccade lead to more redirection. Sun et al. Sun et al. (2018) used this evaluation factor to distinguish the effects of both image-space and object-space, subtle gaze directions on redirection.

Figure 2.15: Total error area for the orange path is shown stripped Sun et al. (2018)

They noticed that using both types of gaze directions in conjunction leads to higher saccadic angular gains and more redirection.

Subjective Feedback: Subjective feedback refers to asking subjective questions about the immersive experience and the perception of the visual artifacts due to repeated redirection. These questions are often asked in a post-test questionnaire at the end of the experiments. This subjective feedback provided by the users contains the most valuable data as it explores their perception of redirection during the experiments. For example, both Sun et al. Sun et al. (2018) and Langbehn et al. Langbehn and Steinicke (2018a) relied on this evaluation factor to determine the efficacy of their proposed redirected walking techniques.

Simulator Sickness: The feeling of motion sickness (nausea, dizziness, and vertigo) induced by the tasks at hand while immersed in VR is called Simulator Sickness. Contradictory visual and vestibular cues induced by the immersive experience lead to simulator sickness. A method proposed by Kennedy et al. R. S. Kennedy et al. (1993) quantifies the measure of a participants' simulator sickness using a Simulator Sickness Questionnaire. Both state-of-the-art works presented in the report, Langbehn and Steinicke (2018a) and Sun et al. (2018), have chosen the simulator sickness as an evaluation factor to quantify their participants' comfort level during the experiments. Langbehn et al. Langbehn and Steinicke (2018a) reported an increase in the simulator sickness scores among their participants after each experiment, while Sun et al. Sun et al. (2018) reported a decrease in their simulator sickness scores compared to the warping-based approach Dong et al. (2017).

Chapter 3

Inattentional Blindness for Redirected Walking Using Dynamic Foveated Rendering

In this chapter, we introduce a rotation-based redirection technique that leverages dynamic-foveated rendering and the psychological phenomenon of inattentional blindness induced by a cognitive task. This work is based on our paper "Inattentional Blindness for Redirected Walking Using Dynamic Foveated Rendering" Joshi and Poullis (2020b), published in IEEE Access. This method supports open-world environments and enables users to walk freely while maintaining high levels of immersion. Key contributions of this work include the development of the rotation-based redirection technique, the integration of dynamic-foveated rendering, the use of inattentional blindness and the evaluation of the method's effectiveness through extensive user studies. Although this work relies on the eye trackers embedded within the VR headset for foveated rendering, it further expands the scope of VR navigation solutions, addressing the challenges of physical space limitations and enhancing the user's sense of presence in large-scale open-world virtual environments. The work has also been published as a full US and Canadian patent Joshi and Poullis (2023a), and as a poster Joshi and Poullis (2020a) at ACM SIGGRAPH 2020, included as appendix B and C, respectively.

3.1 Abstract

Redirected walking is a Virtual Reality(VR) locomotion technique which enables users to navigate virtual environments (VEs) that are spatially larger than the available physical tracked space.

In this work we present a novel technique for redirected walking in VR based on the psychological phenomenon of inattentional blindness. Based on the user's visual fixation points we divide the user's view into zones. Spatially-varying rotations are applied according to the zone's importance and are rendered using foveated rendering. Our technique is real-time and applicable to small and large physical spaces. Furthermore, the proposed technique does not require the use of stimulated saccades but rather takes advantage of naturally occurring saccades and blinks for a complete refresh of the framebuffer. We performed extensive testing and present the analysis of the results of three user studies conducted for the evaluation.

3.2 Introduction

Since the early days of virtual reality researchers have been investigating ways of navigating virtual environments that are spatially larger than the available Physical Tracked Space (PTS). A number of locomotion techniques relying on pointing devices or walking in-place were proposed which have since become customary in VR applications Christou et al. (2016a); but as could be expected, users find these methods rather cumbersome and unnatural. The concept of redirected walking was then introduced circa 2000 Moshell (1999) as a more natural way of navigating VEs, albeit with many restrictions on the shape and size of the physical and virtual spaces.

A number of approaches have since been proposed for implementing redirected walking. Hardware-based techniques e.g. omni-directional treadmills Souman et al. (2010), virtusphere Medina et al. (2008), are not only an expensive solution to this problem but also fail to provide inertial force feedback equivalent to natural walking Christensen et al. (2000).

In contrast, software-based techniques are more cost effective and typically involve applying perceptually subtle rotations to the VE causing the users' to unknowingly change their walking direction. Applying these rotations to the VE, however subtle, can negatively impact the sense of immersion of the user. The reason for this is because these techniques either employ warping which

introduces visual artifacts and distortions in the VE or even simulation sickness Dong et al. (2017), or rely on forcing the user to look away by stimulating saccades in order to update the environment during a rapid eye movement as in Sun et al. (2018).

In this work, we address the problem of redirected walking and present a novel technique based on the psychological phenomenon of inattentional blindness. Inattentional blindness refers to the inability of an individual to see a salient object in plain sight, due to lack of attention. The phenomenon was popularized in a world-famous awareness test described in D. J. Simons and Chabris (1999) [1] in which participants were asked to watch a video and count the number of times a basketball is passed between the 6 people shown in the video wearing a white shirt. During the video a person dressed as a gorilla walks from right side of the screen to the left, passing through the people playing basketball, and at one point stopping to beat its chest. At the end of the experiment participants were asked to state the number of passes, and whether they noticed anything different in the video; the majority of which reported that they hadn't noticed anything out of the ordinary i.e. the gorilla walking and beating its chest.

We exploit this phenomenon and further strengthen its effect with foveated rendering. Foveated rendering is a rendering technique whose primary objective is to reduce the rendering workload. Using eye tracking the user's eyes are tracked within the VR headset in real-time. The zone in the image corresponding to the foveal vision, i.e. the zone gazed by the fovea which provides sharp and detailed vision, is then rendered at very high quality. On the other hand, the zone in the image corresponding to the peripheral vision is rendered at a much lower quality since peripheral vision lacks acuity albeit it has a wider field of view. This process is performed without causing any perceptual change to the user. Nowadays, foveated rendering is supported by hardware such as NVIDIA's RTX graphics card series which allows for real-time ray-tracing and hence, real-time performance.

Our proposed technique applies spatially-varying rotations to the VE according to the zone's importance using foveated rendering to strengthen the effect of inattentional blindness. We present the results of three user-studies. The first user study was conducted in order to determine the maximum rotation angle and field-of-view for which participants do not perceive a change. The objective

[1] Video of the test can be found here: https://www.youtube.com/watch?v=vJG698U2Mvo

of the second user study was to confirm the results of the first and also verify using only in-situ gaze redirection that a rotation of the VE results in an equivalent rotation of the participant in the PTS. Lastly, the third user study was conducted to evaluate redirected walking using dynamic foveated rendering during inattentional blindness in a small PTS.

The chapter is organized as follows: Section 3.3 summarizes the state-of-the-art in the areas of redirected walking and foveated rendering. In Section 3.4 we provide a technical overview of the proposed technique. The user-study for determining the maximum rotation angle and the field-of-view of the foveal zone for which change is imperceptible to a user is presented in Section 3.5. In Section 3.7 we present the results and analysis of our third and final user-study for evaluating the technique for redirected walking in VR.

3.3 Related Work

Over the recent years many techniques have been proposed relating to interaction in VR and in particular, navigation. Below we present a brief overview of the state-of-the-art most relevant to our proposed technique. The related work is categorized in terms of (a) redirection and VR, (b) steering algorithms, resetting phase and natural visual suppressions, and (c) dynamic foveated rendering.

3.3.1 Redirection and VR

One of the most important forms of interaction in VR is locomotion. Natural walking is the most preferred (and natural) technique primarily because it provides an increased sense of presence in the VE Usoh et al. (1999) and improved spatial understanding Peck et al. (2011); R. Ruddle and Lessels (2009); R. A. Ruddle et al. (2011) while reducing the signs of VR sickness LaViola Jr. (2000). However, the main difficulty of using natural walking as a locomotion technique in VR is the requirement that the size of the PTS is comparable in size with the VE, which is often not the case; especially for simulations involving large-scale environments Poullis and You (2009). To this day it remains a very active area of research with a particular focus on locomotion techniques which do not carry, in any degree, the spatial constraints imposed by the physical space over to the theoretically boundless virtual space.

In 2001, Razzaque et al. Razzaque et al. (2001) proposed the concept of 'Redirected Walking' for the first time. By making subtle rotations to the VE the users were tricked into walking on a curved path in the PTS while maintaining a straight path in the VE. These subtle rotations applied to the VE were enough to convince the users that they had explored a comparatively larger virtual area than the actual available play space.

Since the original concept of redirected walking, a number of attempts were made to achieve the same effect based on software and/or hardware. Some researchers even tried to achieve this by incorporating physical props Cheng et al. (2015) or manipulating the entire physical environment E. A. Suma et al. (2011, 2012). However, these type of solutions failed to gain the mainstream attention due to their many dependencies on factors other than the actions of the user himself.

Hardware-based approaches were also explored to resolve this problem such as the omnidirectional treadmill Souman et al. (2010), suspended walking Walther-Franks et al. (2013), low-friction surfaces Iwata and Fujii (1996), walking in a giant Hamster Ball Medina et al. (2008), etc. Souman et al Souman et al. (2010) proposed a technique using an omnidirectional treadmill algorithm to improve the perception of life-like walking in the VE. A position-based controller was used to determine the speed of the user's movements in the VE and to rotate the treadmill accordingly. Several other attempts were also made to create a perfect omnidirectional treadmill Darken et al. (1997); Huang (2003); Iwata (1999); Nagamori, Wakabayashi, and Ito (2005) but they are not considered cost-effective solutions primarily due to the heavy and expensive equipment involved. Following a similar hardware-based approach, a giant hamster-ball-like structure was proposed by Fernandes et al in Fernandes et al. (2003) where the user was placed in a human-sized hollow sphere which could spin in any direction. Another similar prototype was proposed by Medina et al in Medina et al. (2008). These prototypes, although they are possible solutions to the problem of infinite walking in VR, they all lack inertial force feedback. For this reason natural walking is considered to be the most preferred and natural way for locomotion in VR Christensen et al. (2000). Moreover, the multimodal nature of natural walking allows free user movements such as jumping or crouching.

Software-based techniques have also been proposed for achieving the same effect by solely manipulating the VE. These can be divided into two groups: (a) techniques that use the user's head rotations and translations for scaling the VE dynamically based on the scenario as in Azmandian

et al. (2017); Razzaque et al. (2001); Razzaque, Swapp, Slater, Whitton, and Steed (2002), and (b) techniques that partially or fully warp the virtual environment Dong et al. (2017); Sun et al. (2016). Due to the dominance of the visual sense over the other human senses, these techniques focus mainly on reducing the effect of subtle visual distractions resulting from repeated redirection. These visual distractions are mainly created during the naturally occurring or stimulated visual suppressions such as a blink or a saccade. Langbehn et al in one of their recent studies Langbehn and Steinicke (2018a) proposed a redirection technique which leverages the naturally occurring blinks to redirect the user. The authors performed extensive experiments to determine the thresholds of rotational and translational gains that can be introduced during the blink. Concurrently, Sun et al in Sun et al. (2018) leveraged the perceptual blindness occurring before, during, and after the saccades to update the environment quickly over several frames. A common disadvantage of these techniques is the fact that they are disruptive to the cognitive task at hand, since they rely on stimulating saccades by introducing artifacts in the rendered image to distract the user's attention.

3.3.2 Steering Algorithms, Resetting Phase and Natural Visual Suppressions

Steering Algorithms: In order to calculate the amount of redirection, two parameters are required: the target direction of the user (a) in the VE, and (b) in the PTS. There are many methods for predicting the target direction in VE ranging from using the user's past walking direction Zank and Kun (2015), to head rotations Steinicke et al. (2008), and gaze direction Zank and Kun (2016). As there are spatial constraints in terms of the available PTS, the user must be steered away from the boundaries of the PTS i.e. typically walls. To locate the user's target direction in PTS, Razzaque proposed a number of related agorithms in Razzaque et al. (2005), namely steer-to-center, steer-to-orbit, and steer-to-multiple-targets: (i) steer-to-center steers the users towards the center of the PTS; (ii) steer-to-orbit steers them towards an orbit around the center of the PTS; and (iii) steer-to-multiple-targets steers the users towards several assigned waypoints in the PTS. In Hodgson and Bachmann (2013), Hodgson et al proposed the steer-to-multiple-centers algorithm which is essentially an extension of steer-to-center algorithm. In this case, the steer-to-center algorithm proved to be working better than this extension. Another experiment performed by Hodgson et al in Hodgson, Bachmann, and Thrash (2014) showed that the steer-to-orbit algorithm worked better

in the setting where the directions of walking were limited by virtual objects in the VE. In all of the above cases, the user's direction of movement is constantly changing which implies that the amount of redirection must also be constantly computed.

Recently, Langbehn et al Eike, Paul, Gerd, and Frank (2017) proposed a redirected walking technique using predefined curved paths in the VE. In this experiment, users were instructed to follow the predefined curves inside the VE and were allowed to change their direction of walking only when they reached the intersection of their current curve with another predefined curve. Since the curves were mapped within the PTS, this eliminated the possibility of crossing over a boundary as long as the user followed the predefined path. In our work, we employ the steer-to-center algorithm for redirecting the user towards the center of the PTS when a collision is predicted to occur.

Resetting Phase: The resetting phase is one of the most important aspects of all redirected walking techniques because there is always a small possibility of the user crossing over the boundary of the PTS. If this occurs, the user has to be stopped and their position has to be reset before starting to walk again.

Several solutions were proposed to address this problem with the most common being in B. Williams et al. (2007): (i) Freeze-Turn is the method where the field-of-view (FoV) of the user remains unchanged while she turns towards the empty space in the PTS, (ii) Freeze-Back-up is the method where the FoV remains unchanged while the user backs-up from the boundary making enough room for walking, and (iii) 2:1 Turn is the method where twice the rotational gain is applied to the user's head rotations, i.e. if the user turns by 180°, a rotation of 360° is applied so that the user is again facing the same direction that she was facing before. Visual effects which may result from the resetting can be addressed with the use of visual distractors as proposed by Peck et al in Peck et al. (2011).

Furthermore, most of the redirected walking techniques (e.g. Redirected Walking Toolkit Azmandian et al. (2017)) follow the stop-and-go paradigm where if the user crosses over a boundary, and before she starts walking again, she will have to perform an in-situ rotation towards a direction where there is no obstacle. We follow the same paradigm in our work.

Natural Visual Suppressions: The human visual system is not perfect. Due to several very frequent involuntary actions, humans face temporary blindness for short periods of time O'Regan

et al. (2000); Rensink (2002); Rensink et al. (1997). These involuntary actions are called visual suppressions. Saccades, eye-blinks, the phase of nystagmus, and vergence movements are some of the involuntary visual suppressions Volkmann (1986). Saccades are the brief rapid eye movements that occur when we quickly glance from one object to another; eye-blink is a rapid opening and closing of the eyelids - these eye movement can occur either voluntarily, involuntarily or as a reflex; the phase of nystagmus is a condition where uncontrolled rapid eye movements occur from side-to-side, top-to-bottom or in circular motion; vergence movement occurs to focus on the objects with different depths, one after the other Volkmann (1986). In the following, we review only techniques employing saccades and eye-blinks since our proposed technique only utilizes these.

Saccades are extremely unpredictable, rapid, and ballistic eye movements that occur when we abruptly shift our fixation point from one object to the other Bahill et al. (1975). The visual suppression occurs before, during, and after the saccadic movement Burr, Morrone, and Ross (1994) and could last for 20-200ms while the speed of these saccadic movements can reach up to 900°/s Bahill et al. (1975). Researchers are still trying to identify an exact reasoning behind this involuntary human behavior Burr et al. (1994); Ibbotson and Cloherty (2009). A saccade occurs very frequently and can last for several frames Langbehn and Steinicke (2018a) which makes it possible to render the updated VE without the user noticing.

In contrast to the saccades, blinks occur much less frequently and are much slower which gives more flexibility for reorientation due to the longer induced change blindness O'Regan et al. (2000). Depending upon the scenario, one blink can give the users a temporary blindness of 100-400 ms which is much longer than a saccade Ramot, Daniel (n.d.). This makes them easier to detect even with readily available off-the-shelf eye trackers. Similar to saccades, our visual perception is also suppressed before, during, and after the opening and closing movements of the eyelids Volkmann (1986); Volkmann et al. (1980). A study performed by Bentivoglio in Bentivoglio et al. (1997) showed that an average person blinks at an average rate of 17 times per minute.

In our work, we leverage this physiological phenomenon to refresh the foveal zone render and therefore redirect the user multiple times per minute during blinks. Furthermore, we leverage information reported in recent studies for determining the maximum rotational and translational thresholds for VR during blinks and saccades Bolte and Lappe (2015); Langbehn, Bruder, and Steinicke

(2016); Langbehn and Steinicke (2018a); Sun et al. (2018) to update the VE and refresh the render without the user perceiving anything.

3.3.3 Dynamic Foveated Rendering

The concept of foveated rendering was first introduced by Reder in 1973 Reder (1973). This is a technique which drastically reduces the overall GPU workload while providing the same VR experiences as before. Foveated rendering leverages the fact that small changes occurring in our peripheral vision are imperceptible to humans. Thus, the area of the image corresponding to the peripheral vision can be rendered at a much lower resolution while the area of the image corresponding to the foveal vision is rendered at full resolution Patney et al. (2016).

Although in recent years researchers have proposed many software-based techniques for simulating perceptually-guided foveated rendering e.g. He et al. (2014); Patney et al. (2016); Stengel et al. (2016); Sun et al. (2018), it was until the very recent introduction of Nvidia's Variable Rate Shading NVIDIA (2018) that foveated rendering was supported in hardware and integrated into the rendering pipeline. This integration has reduced latency and aliasing effects close to zero as it was demonstrated by the real-time visualization platform called ZeroLight O'Connor, Chris (2018).

In this chapter, we employ Nvidia's VRS not only for reducing the overall GPU workload but also for blending foveal and non-foveal (peripheral) zones rendered from two co-located cameras, respectively.

3.4 Technical Overview

Figure 3.1 shows the pipeline of the proposed technique. Two co-located cameras Cam_{foveal}, $Cam_{non-foveal}$ render the VE. Based on the results of the first user study we have determined that the field-of-view for Cam_{foveal} is $\delta = 60^{\circ}$, and the rotation angle applied to the VE and rendered from $Cam_{non-foveal}$ is $13.5^{\circ} > \theta > 0^{\circ}$. A mask corresponding to $\delta = 60^{\circ}$ is applied on the rendered image of Cam_{foveal}, and the inverse of the same mask is applied on the rendered image of $Cam_{non-foveal}$. The resulting masked renders are then composited into the final render displayed to the user. As demonstrated by the results of the user studies #2 and #3, the users fail to perceive

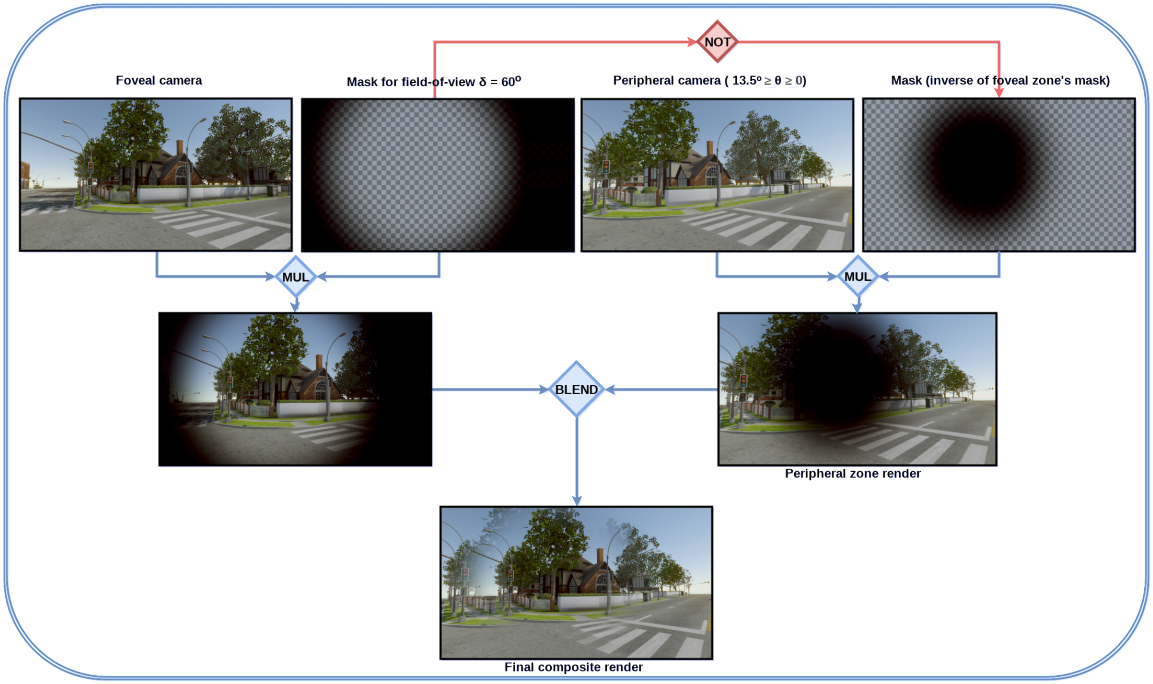

Figure 3.1: The pipeline of the proposed technique involves rendering the VE from two co-located cameras Cam_{foveal} with field-of-view $\delta = 60^{\circ}$ and $Cam_{peripheral}$ with rotation angle $13.5^{\circ} > \theta > 0^{\circ}$. As demonstrated by the results of user study #2 the users fail to perceive any visual distractions or artifacts in the final composite render while they are preoccupied with a cognitive task; which is almost always the case with VR applications.

any visual distractions or artifacts in the final composite render while they are preoccupied with a cognitive task; which is almost always the case in VR applications.

3.5 User Study #1: Determining Maximum Rotation Angle and Field-of-View of Foveal Zone

The efficacy of redirected walking is tightly coupled with the user's perception of the redirection taking place. In our first user-study we determine the maximum angle for which the rotation of the non-foveal zone (i.e. the area in the rendered image corresponding to the peripheral vision) remains imperceptible by the user.

3.5.1 Application

We designed an immersive VR application using the HTC Vive Pro Eye HMD with an integrated Tobii Eye Tracker. The application depicted a serene urban city environment in which red spheres were hidden at random locations. The environment was developed in Unity3D and foveated rendering was supported using an NVIDIA RTX 2080Ti graphics card. Three zones were specified for the foveated rendering:

(1) Foveal zone: The foveal zone is the circular area in the rendered image centered at the fixation point captured by the eye tracker. For rendering, this zone must have the highest quality since this is where the user's foveal vision is focused. Thus, pixels in the foveal zone are rendered with a 1:1 sampling.

(2) Non-foveal zone: The non-foveal zone is the area in the rendered image (complementary to the foveal zone) which corresponds to the peripheral vision. This zone is of lower importance than the foveal zone since it is not the main focus of the user's vision. Hence, pixels in the non-foveal zone are rendered with a 16:1 sampling.

(3) Transition zone: The transition zone is the overlapping area of the foveal and non-foveal zones. This zone was introduced after initial experiments had shown that having only the foveal and non-foveal zones results in sharp boundary edges on the circular area separating

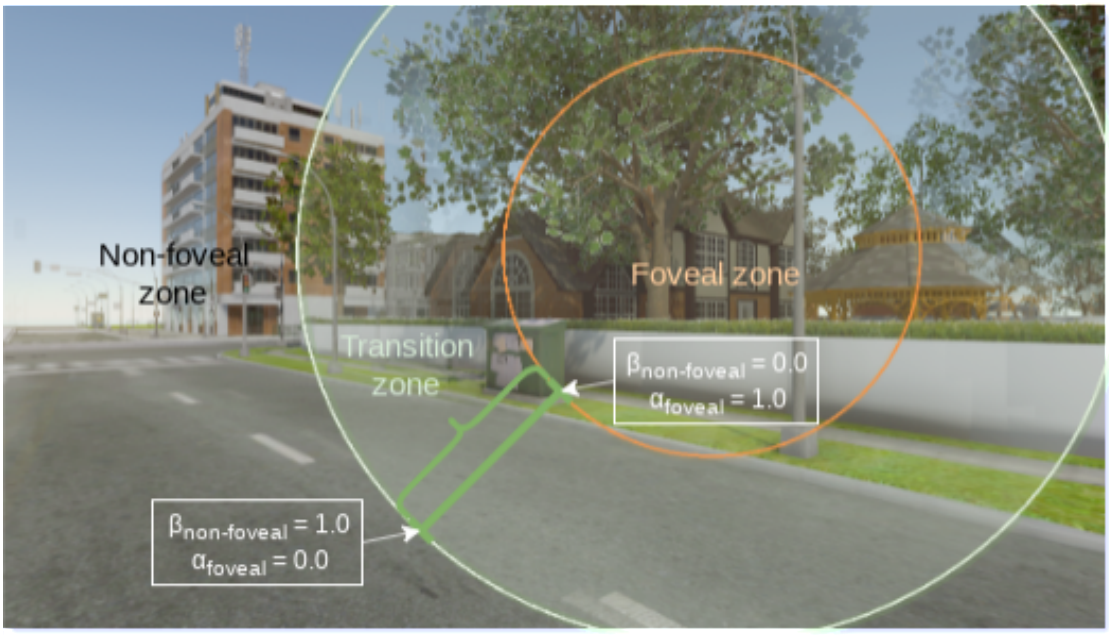

Figure 3.2: User's perspective during the user study #1. The foveal zone marked with the orange circle is rendered at the high quality (1:1 sampling). The non-foveal zone is the complement of the foveal zone and is rendered at a lower resolution (16:1 sampling)

them which are immediately perceived by the user. Pixels in the transition zone are rendered with a 4:1 sampling.

Figure 3.2 shows a frame [2] from the application with the three zones annotated. The foveal zone corresponding to a field-of-view of $\delta_{foveal} = 60^{\circ}$ is marked with an orange circle and is rendered at the highest quality with 1:1 sampling and no rotation. The non-foveal zone is rendered at a lower resolution with 16:1 sampling and a rotation of $\theta_{non-foveal} = 6^{\circ}$.

The transition zone is also shown as the green ring around the foveal zone, rendered with 4:1 sampling. This indicates the overlapping area between the foveal and non-foveal zones for which alpha-blending is performed at each pixel with $C_{blended} = \alpha_{foveal} \times C_{foveal} + \beta_{non-foveal} \times C_{non-foveal}$ where $\beta_{non-foveal} = 1.0 - \alpha_{foveal}$, $C_{blended}$ is the resulting blended color, and C_{foveal}, $C_{non-foveal}$ are the foveal and non-foveal colors at a pixel, respectively. This requires two co-located cameras (with different rotations) in order to render the foveal and non-foveal frames from which the color values C_{foveal} and $C_{non-foveal}$ are retrieved. The boundary values [0.0, 1.0] for α_{foveal} and $\beta_{non-foveal}$ are also shown in this figure. These coincide with the boundaries of the transition zone. The field-of-view $\delta_{transition}$ corresponding to the transition zone is defined empirically as an offset to the field-of-view δ_{foveal} of the foveal zone, given by $\delta_{transition} = \delta_{foveal} + 40^{\circ}$.

3.5.2 Procedure

The premise of our hypothesis is inattentional blindness which implies that the user's attention must be directed towards another cognitive task. Thus, we instructed the participants to perform a target-retrieval task. More specifically, the participants were asked to search and count red spheres hidden in the VE. At each iteration of the experiment, the red spheres were hidden at random locations. This was done in order to eliminate the possible bias that may be introduced by memorizing the locations between iterations.

The first user study involved 11 participants (2 females, 18.2%). Average age was 24.64 with a SD of 2.8. Median of their reported experiences with using VR devices was 3, and the median of their experiences with using an eye tracking device was also 3 on a 5-point Likert scale, with 1

[2] The visible frame in VR is panoramic and considerably larger than the one shown here. We are showing only the part of the frame relevant to the discussion.

being least familiar, and 5 being most familiar. The participants performed this task from a seated position and were only allowed to rotate their chair in-place. The participants were instructed to press the trigger button on the Vive controller if and when they noticed a visual distortion or felt nausea due to simulator sickness. During the experiment, the rotation angle θ of the non-foveal zone was gradually increased and grouped by increasing order of the field-of-view δ of the foveal zone i.e. $([(\delta_1, \theta_1), (\delta_1, \theta_2), \ldots, (\delta_1, \theta_n), (\delta_2, \theta_1), \ldots, (\delta_2, \theta_n), \ldots])$. Each time the trigger was pressed, the pair of (δ_i, θ_i) was recorded, and then the experiment continued with the now increased field-of-view δ_{i+1} and a reinitialized rotation angle θ_1.

The range of values for the field-of-view was 20° to 60°. The step of each increment was 10° after the completion of one cycle of the rotation angle, or until triggered by the user. During a cycle, the rotation angle ranged from 0° to 15° and the step of each increment was 1° per second.

Preliminary experiments during the design of the application had shown that repeated increments of the field-of-view of the foveal zone can lead to nausea and severe dizziness. For this reason, the participants were instructed to take a short break after each cycle of increments of the field-of-view. Furthermore, the sequence of the cycles i.e. the field-of-view values, was randomized for each participant in order to eliminate any bias. Figure 3.2 shows the view from the user's perspective during the experiment.

3.5.3 Analysis of results

The results are shown in Figure 3.3. A cycle of the rotation angle $\theta \in [0^\circ, 15^\circ]$ was performed for each field-of-view $\delta \in [20^\circ, 60^\circ]$. The results show that as the field-of-view δ increases the tolerance for higher rotation angle θ also increases, which can also be confirmed by the exponential trendlines shown for each participant. For the reasons mentioned above, we select the smallest rotation angle for which users did not perceive a change associated with the largest field-of-view for which the majority of the users did not perceive a change (i.e. 9 out of 11). Thus, the ideal pair values for (δ, θ) is determined to be $(60^\circ, 13.5^\circ)$; where 13.5° is the maximum allowed rotation angle.

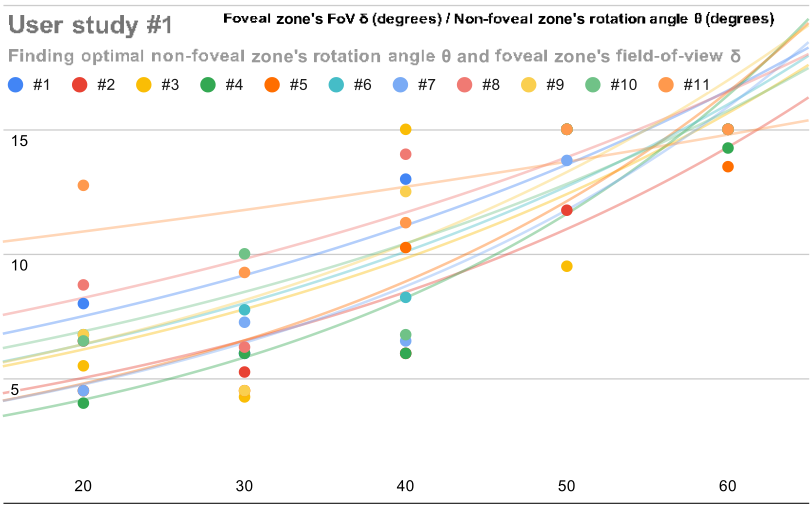

Figure 3.3: Participants' responses for pairs of (δ, θ). We select the smallest rotation angle for which users did not perceive a change associated with the largest field-of-view for which the majority of the users did not perceive a change (i.e. 9 out of 11). Thus, the optimal pair values for (δ, θ) is determined to be (60°, 13.5°). The exponential trendlines are also shown which confirm that as the field-of-view δ increases, the tolerance for higher rotation angle θ also increases.

Simulator Sickness Questionnaire (SSQ)

Upon completing the experiment all participants were asked to complete Kennedy Lane's Simulation Sickness Questionnaire (SSQ). The Total Severity (TS) and the sub-scales Nausea, Oculomotor, and Disorientation were calculated using the formulas from R. S. Kennedy et al. (1993). Based on the SSQ categorization provided by Kennedy et al. in R. Kennedy et al. (2003), 55% of the participants reported no signs (TS=0) or minimal signs (TS< 10) of simulator sickness. All the participants completed the test, with 0 dropouts. Upon further analysis, the disorientation sub-scale had the highest average score of 29.11 with a maximum score of 139.2. This was expected, considering the fact that the rotation angle was constantly increasing and thus the VE rendered in the HMD induces conflicts between the vestibular and visual signals, leading to vestibular disturbances such as vertigo or dizziness. The results from SSQ responses are summarized in Table 3.1 and Figure 3.4.

Table 3.1: Results from the responses of SSQ. The Total Severity (TS) and the corresponding sub-scales such as Nausea, Oculomotor, and Disorientation were calculated using the formulas from R. S. Kennedy et al. (1993). The majority (55%) of the participants reported no signs (TS=0) or minimal signs (TS< 10) of simulator sickness. Highest average score for disorientation as expected.

Scores	Mean	Median	SD	Min	Max
Nausea (N)	11.27	9.54	14.67	0	38.16
Oculomotor (O)	19.29	15.16	24.06	0	83.38
Disorientation (D)	29.11	0	45.09	0	139.2
Total Score (TS)	21.76	7.48	28.47	0	93.5

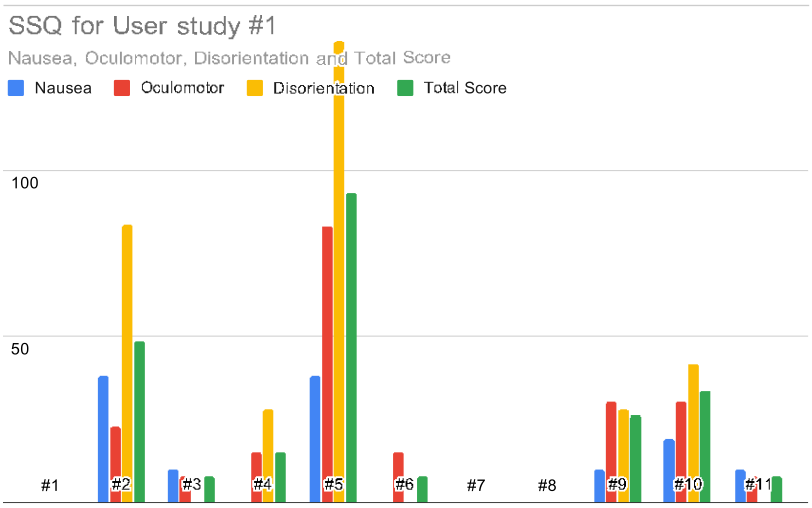

Figure 3.4: Results from the SSQ scores (Left to right: Nausea, Oculomotor, Disorientation and Total Severity). The Total Severity and sub-scales such as Nausea, Oculomotor, and Disorientation were calculated based on the formulas in R. S. Kennedy et al. (1993)

3.5.4 Discussion

The design of the application of user study #1 involved making a decision on the range of values for (a) the field-of-view δ for the foveal zone, and (b) the rotation angle θ of the non-foveal zone:

- Range of δ: Humans have a maximum horizontal field-of-view of about 210°. This is further reduced to 110° by the hardware i.e. maximum field-of-view for HTC Vive Pro.

  The foveal and non-foveal zones are by definition complementary to each other. Thus, there is a tradeoff between the sizes of the foveal and non-foveal zones. If δ is large, then the foveal zone is large, and the non-foveal zone is small. Having a small non-foveal zone means that only a small area of the final composite image will be rendered from the non-foveal camera $Cam_{non-foveal}$, as shown in Figure 3.1, leading to smaller possible redirections. When $\delta = 60^\circ$ the foveal zone occupies 54.55% of the final composite render, and the non-foveal zone occupies 45.45% (including a transition zone of 35.45%). Similarly, if $\delta = 90^\circ$ the foveal zone occupies 90.91% of the final composite render which does not leave much for the non-foveal zone. In contrast, if δ is small, then the foveal zone is small and the non-foveal zone is large. Although this allows for larger redirections, we have found in our preliminary tests that when $\delta < 20^\circ$ it can cause severe nausea and simulation sickness.

  For these reasons we have selected the range of values $\delta \in [20^\circ, 60^\circ]$ which balances the sizes of the foveal and non-foveal zones, and is large enough that it does not cause user discomfort.

- Range of θ: Recent experiments reported in Sun et al. (2018) have shown that users cannot tolerate a rotational angle of more than 12.6° in their field-of-view during a saccade having a velocity of 180°/sec. Based on this, we have selected the range $\theta \in [0^\circ, 15^\circ]$.

3.6 User Study #2: In-situ Gaze Redirection using Dynamic Foveated Rendering

The objective of the second user study is twofold:

(1) Firstly, to determine whether a rotation of the VE by an angle below the maximum (i.e. $\theta < 13.5^\circ$) is indeed imperceptible and does not cause simulation sickness.

(2) Secondly, to experimentally prove with quantitative measures that using the proposed redirection technique with gaze-only (and without walking) the rotation of the VE results in the equivalent redirection of the participant in the PTS.

3.6.1 Application and Procedure

An experiment was devised similar to the one in Section 3.5. In contrast to the user study #1, the participants were only allowed to spin in-situ from a standing position. This change from a seated to a standing position eliminates the possibility of the participants perceiving any orientation and navigation cues coming from the orientation of the chair. The participants were given a target-retrieval task and instructed to retrieve, by directing their gaze to, as many virtual targets (i.e. orange pumpkins) as possible. The virtual targets disappeared as soon as the gaze direction intersected their bounding box. The positions of the targets were randomized for each participant.

The duration of the experiment was 60 seconds. Unbeknownst to the participants, in the first 30 seconds there was no redirection applied i.e. $\theta_{Cam_{non-foveal}} = 0$. This served as a baseline for participants who had little to no experience in using HMDs. Once accustomed to the VE, during the following 30 seconds the rotation angle of the VE was increased at a rate of $6^{\circ}/s$.

Hence, the hypothesis is that after 30 seconds of a constant smooth rotation of the VE at a moderate rate of $6^{\circ}/s$ the participant should face 180° away from their initial orientation i.e. the opposite direction. To prove this, the initial (i.e. at $time = 0s$) and the final (i.e. at $time = 60s$) gaze directions of each participant were recorded. Additionally, before each participant removed the HMD at the end of the experiment they were asked to face towards what they believed to be their initial directions in the PTS using visual landmarks from within the VE to orient themselves.

3.6.2 Analysis of Results

The study involved 11 participants (2 females: 18.2%, average age of 26.27 ± 3.13). Based on a 5-point Likert scale the medians of their experience with using VR or any other eye tracking devices were 3. Five of the participants had not taken part in user study #1 (Participants #2, #3, #6, #9, #11). After the experiment, the participants completed the SSQ. The angle between the initial and final gaze directions was calculated for each participant. The average deviation was 171.26°

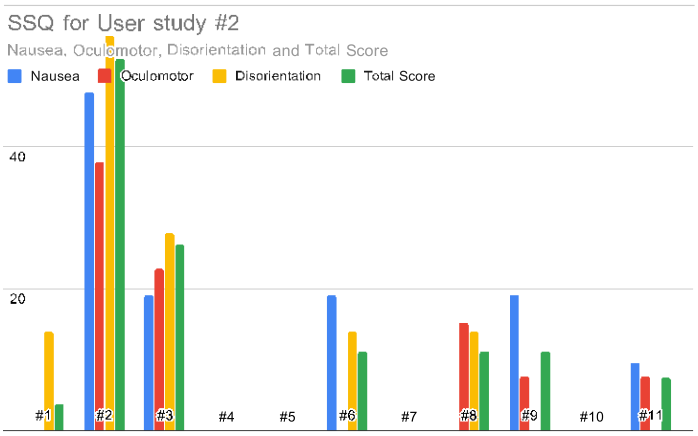

Figure 3.5: Results from the SSQ scores (Left to right: Nausea, Oculomotor, Disorientation and Total Severity). The Total Severity and sub-scales such as Nausea, Oculomotor, and Disorientation were calculated based on the formulas in R. S. Kennedy et al. (1993)

(4.77 SD) which means that the participants thought that their initial orientation was towards the opposite direction. In fact, all participants reported that they did not perceive the redirection and were surprised by how off their 'sensed' orientations were.

Simulator Sickness Questionnaire (SSQ)

Based on the scores reported by the participants in the post-test SSQ, the majority of the participants (55%) showed no signs (TS=0) or minimal signs (TS< 10) of simulator sickness. The highest score and average was reported for the sub-scale disorientation although reduced by a factor of 2 from user study #1. This was anticipated since the rotation angle was less than the maximum determined from user study #1. As it can be seen from Figure 3.5, one of the participants (#2) had no previous experience with VR and reported perceptual anomalies including difficulty concentrating, fullness of head and difficulty focusing. The results for the SSQ are summarized in Table 3.2.

Table 3.2: Results from the responses of SSQ for user study #2.

Scores	Mean	Median	SD	Min	Max
Nausea (N)	10.41	0	15.06	0	47.7
Oculomotor (O)	8.27	0	12.43	0	37.9
Disorientation (D)	11.39	0	17.41	0	55.68
Total Score (TS)	11.22	7.48	15.78	0	52.36

3.7 User Study #3: Redirected Walking using Dynamic Foveated Rendering

The objective of the third user study is to evaluate the efficacy of redirected walking during inattentional blindness using dynamic foveated rendering.

3.7.1 Application

As previously explained, inattentional blindness refers to the inability of an individual to see a salient object in plain sight due to lack of attention. This is true for the majority of the VR applications where the user is actively engaged and preoccupied with a cognitive task e.g. games, training simulations, etc. Thus, for the purposes of this user study, we designed a first-person VR game where the objective is to stop an alien invasion. To achieve this the user has to walk through a deserted urban city to a predefined location indicated by a large purple orb and destroy the alien-mothership (appearing in the form of a giant brain) while zapping green aliens along the way. Zapping one alien will award one score point to the player. The green alien enemies are randomly spawned (and are therefore independent of the orientation of the current redirection) only within the field-of-view of the user while also making a sound effect. An example of in-game gameplay is shown in Figure 3.6.

The shortest distance the participants had to travel in the VE was 42m while the available PTS has a size of $4 \times 4\text{m}^2$. The PTS is shown as the cyan-colored box in Figure 3.7 and the position of the user w.r.t. the PTS is indicated with the camera icon. For safety reasons, a resetting mechanism of 2:1 was implemented. In cases where the game predicts that the user is about to cross over a boundary of the PTS, it would pause and prompt the user to rotate in-situ by 180°. During the user's rotation, the VE was also rotated by the same angle but in the opposite direction. The user was then allowed to proceed with playing the game.

Redirection was primarily performed by blending in real-time the foveal and non-foveal renders. Furthermore, redirection was also performed during the tracked naturally occurring blinks and saccades. In contrast to the state-of-the-art Sun et al. (2018), our approach does not stimulate saccades nor blinks since these are disruptive to the cognitive task at hand.

Figure 3.6: The game designed for user study #3. The objective is to stop an alien invasion by walking to a predefined location in the VE and destroying the alien-mothership (appearing in the form of a giant brain) while shooting aliens along the way.

Figure 3.7: Examples from two participants' tests of the experimental condition of user study #3. The path walked in PTS up to that point is shown in orange color. The corresponding path in the VE is shown in blue color. The cyan-colored box indicates the $4 \times 4\text{m}^2$ available PTS and the camera icon inside the box indicates the location of the user w.r.t. the PTS. Statistics are shown at the top left corner about the current redirection (measured in degrees), distance traveled in PTS (measured in meters), number of resets required, and the score at that point in time. For safety reasons the resetting mechanism of 2:1 was implemented. Steer-to-center algorithm was used for redirecting.

3.7.2 Procedure

In order to evaluate the efficacy of our technique, a controlled experiment is conducted where the independent variable being tested is the proposed redirection technique. The participants were instructed to complete the objective of the game twice: the first time with the experimental condition i.e. with redirected walking, and after a short break a second time with the control condition i.e. without redirected walking. For the experiment with the control condition, participants had to navigate themselves to the destination by relying solely on the resetting mechanism every time they went out-of-bounds from the available $4 \times 4m^2$ PTS.

A sample size estimation with an effect size of 0.25 showed that a total of 25 participants were required for the experiment. All the participants were randomly chosen (12% female, average age of 25.88 years with a SD of 3.06). Based on a 5-point Likert Scale, the median of their experiences using VR headsets or any other eye tracking devices was 3.

Before the experiment, participants were briefed on their objective. Instructions were also given on how the resetting mechanism works in case they are prompted with an out-of-bounds warning and are required to reset their orientation. Moreover, they were instructed to walk at a normal pace which will allow them to complete the task along the way. Once both the objectives were completed, participants were also asked to complete the SSQ R. S. Kennedy et al. (1993). Furthermore, for the experimental condition, at the end of the first experiment the participants were asked "Did you feel the redirection or any other scene or camera modulation during the experience?".

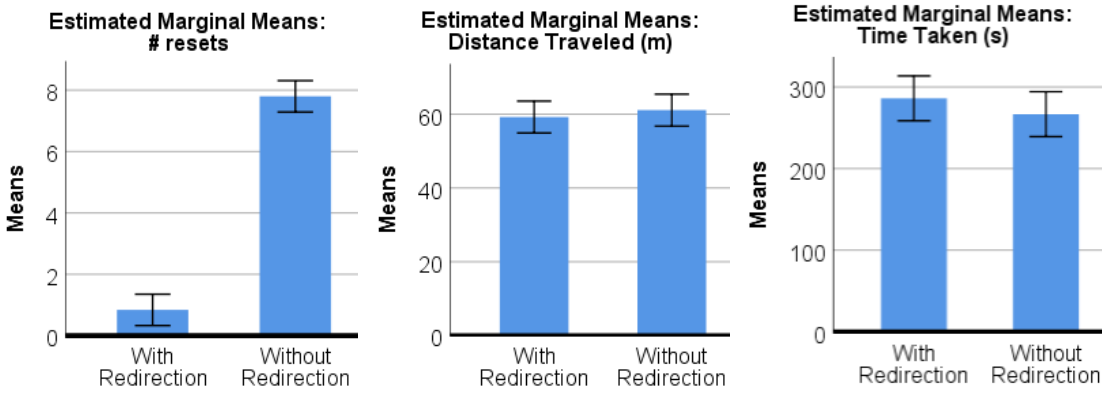

Figure 3.8: Left to right: ANOVA results for (a) Number of Resets, (b) Distance traveled in PTS, and (c) Total Time Taken. Confidence Interval = 95%

3.7.3 Analysis of Results

A one-way between groups ANOVA ($\alpha = 0.05$) with repeated measures was performed to compare the effects of with- and without-using the redirection on the dependent variables: (a) number of resets, (b) distance traveled in PTS, (c) total time taken, and (d) scores. We used partial eta squared (η_p^2) to report the obtained effect sizes for each variable.

Based on the results of Levene's test Levene (1960), it was found that the outcomes for the number of resets ($F(2.14) = 375.710, p > 0.05$), distance traveled ($F(1.176) = 0.348, p > 0.05$) and total time taken ($F(0.971) = 1.001, p > 0.05$) were normally distributed and hence equal variances were assumed. However, the outcome for scores ($F(4.103) = 0.054, p < 0.05$) showed the opposite. As scores violated the homogeneity of variances assumption, the variable was omitted during the ANOVA analysis.

The results from ANOVA showed a statistically significant difference between the number of resets when the game was played with- and without- using the proposed redirection technique ($F(1, 48) = 375.710, p < 0.001$) with $\eta_p^2 = 0.887$. Nonetheless, these results also showed a statistically insignificant effect of redirection on distance traveled ($F(1, 48) = 0.384, p > 0.05$; $\eta_p^2 = 0.008$), and total time taken ($F(1, 48) = 1.001, p > 0.05$; $\eta_p^2 = 0.020$). The η_p^2 values shows that 88.7% of the variability in the required number of resets is accounted for by our independent variable i.e.- redirection. However, the effects on distance traveled and total time taken remains negligible. The results from this test can be seen in Figure 3.8. The error bars in the graphs shows a confidence interval of 95%.

Besides this, results of the first experiment also showed that the system applied an average of 1547.55° (SD = 152.26°) of absolute angular gain to each participant's orientation during the entire test. An average of 3.15° of absolute angular gain was applied per redirection with an average of 1.73 redirections/s. As the participants were cognitively preoccupied with the task of zapping aliens, they were unaware of this angular gain. Furthermore, since this is a real-time technique and thus the frames were rendered without additional lagging other than what is typically imposed by the hardware, none of the participants reported perceiving any scene nor camera manipulation. In the post-study questionnaire, one of the participant stated that "I felt like I was walking straight. I was completely unaware of my actual movements.".

Simulator Sickness Questionnaire (SSQ)

After completing the experiment, participants were asked to fill out an SSQ. Based on the scores reported, the majority of the participants (80%) showed no signs (TS = 0) or minimal signs (TS < 10) of simulator sickness, and only 8% of the participants reported TS > 12. The highest score and mean were reported for the disorientation sub-scale even though the rotation angle was at all times well within the limits of tolerance as determined by user study #1. This can be attributed to the fact that the cognitive workload of the task involved in this user study was more demanding than in the previous user studies. Although this caused an increase in the highest score for disorientation, the mean decreased when compared to that of user study #2. The median values for all sub-scales, as before, were reported as 0. Table 3.3 summarizes the results from this SSQ. As it is evident, the mean scores were dropped significantly from user study #1 and #2.

Table 3.3: Results from the SSQ responses for user study #3. 80% of the participants reported no (TS = 0) to minimal (TS < 10) signs of simulator sickness

Scores	Mean	Median	SD	Min	Max
Nausea (N)	2.29	0	5.7	0	19.08
Oculomotor (O)	6.67	7.58	9.1	0	37.9
Disorientation (D)	8.35	0	17.05	0	69.6
Total Score (TS)	6.43	3.74	10.04	0	44.88

3.8 Discussion

The results of the user study #3 are indicative of the efficacy of the proposed technique in VR applications where the cognitive workload on the user is moderate. Examples of such applications are immersive games, training simulations, cinematic VR, etc.

Further to exploiting inattentional blindness, and unlike other state-of-the-art techinques, our technique relies only on naturally occurring saccades and blinks, and not stimulated. This is distinct from other reported work in the literature and an important advantage since stimulating saccades is both, disruptive to the cognitive task at hand and increases the effects of VR sickness. For example, in Sun et al. (2018) saccades are stimulated by introducing orbs of light in image- and object-space to forcefully divert the user's attention in order to perform the redirection. In contrast to

this, and based on the psychological effect of inattentional blindness, rather than divert the user's attention we exploit the fact that the user is fixated on a particular task leading to "tunnel-vision"-like focus. This allows us to constantly update in real-time the non-foveal (peripheral) zone without the user perceiving a change. Metaphorically speaking, the foveal vision/zone acts as an update-brush of the framebuffer: whenever it moves, based on the tracking of the user's eyes, everything within the foveal zone is rendered from the Cam_{foveal} without any rotations being applied, and everything outside i.e. the non-foveal zone, is rendered from the $Cam_{non-foveal}$ with a rotation $0 < \theta_{Cam_{non-foveal}} < 13.5^{\circ}$ applied to the VE calculated in real-time based on the required redirection.

The experiment in user-study #3 used a PTS of $4 \times 4m^2$. The results show that even with room-scale PTS such as the one used, the users were able to walk distances in the VE which are up to almost 18× orders of magnitude higher than the longest distance in the PTS (i.e. diagonal of $\sqrt{32}$). The maximum recorded distance in our experiments was 103.9m with no resets. Furthermore, the traveled distance can include long straight walks as shown in Figure 3.7.

3.9 Conclusion and Future Work

In this work, we presented a rotation-based redirection technique using dynamic-foveated rendering which leverages the effect of inattentional blindness induced by a cognitive task. The technique uses natural visual suppressions such as eye blinks and saccades (without any artificial stimuli) to make subtle rotations to the VE without the user's knowledge. Furthermore, we conducted extensive tests and presented the results of three user studies. The results confirmed that the technique is indeed effective and can handle long-straight walks. This allows the users to freely explore open world VEs.

Currently, our technique only uses rotational gains for redirection. In the future work, we plan to incorporate translational gains while maintaining the current real-time performance.

Lastly, there is an ample amount of research opportunities in enhancing the redirected walking systems which include saccade prediction algorithms and using other forms of visual suppression e.g. phase of nystagmus, etc.

Chapter 4

SaccadeNet: Towards real-time saccade prediction for virtual reality infinite walking

This chapter presents an innovative event-based redirection technique that employs a CNN-based deep neural network to predict visual suppression resulting from saccades during apparent head rotations and redirect the users. By capitalizing on these predicted suppressions, the technique induces subtle rotations to the virtual environment, redirecting the user's physical walking direction without compromising their sense of presence in VR. It is based on our paper, "Enabling Saccadic Redirection Through Real-time Saccade Prediction", invited for publication in the Journal of Computer Animation and Virtual Worlds Joshi and Poullis (2023b), and accepted for presentation at the 36th International Conference on Computer Animation and Social Agents (CASA) 2023. Additionally, due to the page limitations of the publication, some additional images and contents that provide further context and insights into the research have been included in the Appendix D.

The main contributions of this manuscript include developing a deep-learning-based approach to solving the locomotion dilemma, and a thorough evaluation of the technique's effectiveness through extensive user studies. As compared to our previous work, this research eliminates the need for high-end VR headsets with embedded eye trackers and brings us closer to our ultimate goal by providing

a cost-effective and efficient redirection solution that is compatible with widely available consumer-grade VR hardware, thus increasing its accessibility and applicability in various VR applications.

4.1 Abstract

Our proposed Redirected Walking technique improves on current solutions by leveraging temporary visual disruption caused by saccades for redirection without requiring additional eye-tracking hardware and enables saccadic redirection on widely used consumer-grade hardware, like mobile VR. Using a deep neural network trained on head rotation data, we predict saccades during an apparent head rotation and apply rigid transformations to the virtual environment for redirection. This allows users to infinitely walk in a perceived straight direction from within a physical tracked space of 3.5 x 3.5m^2. The three user studies illustrated in this chapter validate our approach and demonstrate its efficacy.

4.2 Introduction

In recent years, virtual reality (VR) devices have gained immense popularity, thanks to the recent advancements in graphics processing units (GPUs) and low-cost displays. These devices provide an immersive experience and allow users to explore vast virtual environments (VE) that often exceed the size of the available physical-tracked space (PTS). However, navigating these virtual spaces while maintaining a sense of presence still poses a significant challenge. Traditional approaches, such as joystick-based locomotion or teleportation, can accomplish the task successfully but breaks the sense of presence, while physical walking is often limited by the available PTS.

Redirected walking (RDW) offers a more natural approach to locomotion, allowing users to walk in a straight line in a limited PTS while virtually exploring larger VE. However, existing RDW techniques require specialized equipment, such as high-frequency eye-trackers embedded within the VR headset, to detect temporary blindness and add subtle VE rotations to redirect the users. Additionally, some RDW techniques augment the visual content, which can deviate from the content creator's intention, negatively impacting the user's experience. Thus, there is a need for RDW techniques that are more accessible and do not rely on additional equipment or visual

augmentation.

Software-based techniques, such as those introduced by Sun et al. Sun et al. (2018) and Joshi et al. Joshi and Poullis (2020b), have attempted to augment the visual content presented to the user. These techniques leverage the natural phenomenon of change blindness and inattentional blindness to introduce observable artifacts, such as light flashes and subtle peripheral rotations, to redirect the users. However, they distract the users' attention from the primary task, and rely on expensive, integrated eye trackers that are not feasible for widespread consumer use.

To address these limitations, we propose a new RDW technique that leverages the temporary visual disruption caused by saccades during apparent head rotations for redirection without additional eye-tracking hardware. Our approach involves using a neural network to predict saccades during head rotations and apply rigid transformations to the VE during these saccades. Due to the lack of specialized hardware and distractions, it provides a cost-effective, immersive solution that maintains the intended VE's fidelity.

In this chapter, we propose and present a comprehensive evaluation of a deep-learning-based novel RDW technique. Firstly, we explore the relationship between head and eye rotations and provide evidence that users predominantly fixate around the center of the field-of-view (FoV) while primarily using head rotations to change their focus. This finding motivated our neural network's development, which trains on head rotation data to predict saccades in real-time during an apparent head rotation. To validate our approach, we conducted three user studies, including one to confirm the hypothesis, one to collect training data for our neural network, and one to evaluate the performance of our RDW technique. We demonstrate that our approach allows users to walk more than 38 meters in a perceived straight direction from within a physical tracked space of $3.5 \times 3.5 m^2$, while our VE updates are imperceptible to the users.

4.3 Background & Related Work

Locomotion is essential to achieving a truly immersive VR experience, and researchers have been investigating various techniques to make it more natural and intuitive. Current locomotion techniques, such as teleportation and flying, require external controllers and gamepads, which can

break the sense of presence Usoh et al. (1999) and induce simulator sickness LaViola Jr. (2000). In contrast, natural walking is the most favored form of locomotion due to its intuitive nature, which also increases spatial understanding Peck et al. (2011); R. Ruddle and Lessels (2009); R. A. Ruddle et al. (2011). This section provides a detailed review of the most relevant research and state-of-the-art for RDW.

4.3.1 Redirected walking

Despite the advantages of natural walking, the limited availability of the PTS constrains users within a finite boundary while the virtual space can be boundless Poullis and You (2009). Redirected walking Razzaque et al. (2001), leverages the dominance of our visual sense over vestibular or proprioceptive senses Dichgans and Brandt (1978) to manipulate the users' FoV, resulting in discrepancies between the physical and virtual paths Langbehn and Steinicke (2018b); Nilsson et al. (2018); E. Suma, Bruder, Steinicke, Krum, and Bolas (2012). To achieve this, researchers use various methods to predict target directions in the VE, including users' past walking direction Zank and Kun (2015), head rotations Steinicke et al. (2008), and gaze direction Zank and Kun (2016). While, algorithms like steer-to-center, -orbit, and -multiple-targets Razzaque et al. (2005) determine the target direction in the PTS.

Since the inception of RDW, many similar solutions have been developed. Some scaled the user's head rotations, and translations Azmandian, Grechkin, Bolas, and Suma (2016); Azmandian et al. (2017); Eike et al. (2017); Razzaque et al. (2001), while others relied on producing self-overlapping virtual spaces by partially or fully wrapping the entire VE Dong et al. (2017); Sun et al. (2016). More recently, techniques like Langbehn and Steinicke (2018a) and Sun et al. (2018) rely on subtle VE rotations during visual suppressions. Humans face temporary blindness from time to time due to the actions known as visual suppressions Rensink (2002); Rensink et al. (1997). Two of the most frequent visual suppressions are blinks and saccades. Blinks are the rapid opening and closing of the eyes, while Saccades are the ballistic eye movements that change focus from one object to another Volkmann (1986).

4.3.2 Natural Visual Suppression

With speeds reaching up to 900°/s Bahill et al. (1975), the temporary blindness during a saccade can last for 20 to 200 ms Burr et al. (1994). On the contrary, blinks are scarce and more gradual in comparison. The temporary blindness induced due to a blink can typically last for about 100 to 400 ms Ramot, Daniel (n.d.). Due to change blindness, humans fail to notice subtle changes introduced during these visual suppressions O'Regan et al. (2000). In addition, humans also make joint head-eye movements, like smooth pursuits and reflexive, compensatory eye movements such as vestibulo-ocular reflexive (VOR) eye movements Kothari et al. (2020). These oculomotor behaviors do not have similar dynamics, i.e., smooth pursuit and VOR movements have lower velocities and accelerations than saccades. Since the purpose of VOR and smooth pursuit is to keep an object of interest stable on the retina as the head moves, it becomes challenging to differentiate between saccades and these other voluntary behaviors using just head rotations; hence a machine-learning approach to predict saccade during an apparent head rotation is justified.

A method proposed by Langbehn et al. Langbehn and Steinicke (2018a) leverages blinks to inject subtle VE rotations, while Sun et al. Sun et al. (2018) leverages saccades. Both approaches successfully manage to redirect the users, but the latter relies on stimulating artificial saccades by flashing orbs in both image and object space, distracting the users from the task at hand. Following these approaches, Joshi et al. Joshi and Poullis (2020b) proposed a technique that combines the effects of change and inattentional blindness. The FoV was divided into peripheral, foveal, and transitional zones and rendered using dynamic foveated rendering. Based on their importance, these zones are updated one at a time, replacing the entire frame buffer. Finally, they update the foveal zone using the temporary blindness caused due to naturally occurring saccades.

4.3.3 Reset Mechanism

Although each technique is imperceptible and thus preferable in most cases, overt techniques are sometimes favored due to safety or practical limitations. Techniques like freeze-back-up, freeze-turn, and 2:1 turn are some standard approaches proposed by Williams et al. in B. Williams et al. (2007). Freeze-back-up allows users to step back upon hitting the PTS boundary with a frozen FoV,

freeze-turn allows users to turn by 180° with a frozen FoV, and 2:1 turn allows users to make half a turn (180°) in any direction while concurrently making a complete turn (360°) of the FoV in the opposite direction. These are mainly used in hybrid systems as a safety reset mechanism when subtle redirection techniques fail Joshi and Poullis (2020b).

In this chapter, we propose a novel event-based redirection technique that utilizes deep learning to predict visual suppressions caused by saccades during head rotations. We aim to leverage this change blindness effect, where subtle changes in the visual field can go unnoticed by the user due to saccades, and utilize it to perform redirection without the need for specialized eye-tracking hardware.

4.4 Head-Eye Relationship

The first study investigates the relationship between head and eye rotations during VR locomotion while users perform a cognitive task that induces repeated head rotations.

In this user study, participants were presented with an immersive VR environment featuring an open sky, eliminating directional cues. They were instructed to locate and tag small stationary targets using a firing wand, which were spawned at a distance of 10 meters and separated by at least 20° from each other to ensure purely eye-driven saccades Fang, Nakashima, Matsumiya, Kuriki, and Shioiri (2015); Stahl (1999). Upon elimination, the targets would respawn in another random direction. A score was kept to encourage more target hits. This elicited frequent gaze shifts, resulting in simultaneous head and eye rotations. Participants performed the task five times with one minute of the task, followed by a one-minute break. The gaze data from each experiment were recorded to generate a heatmap to analyze average gaze duration within the FoV.

The study involved 12 participants with an average age of 24.54 years and a standard deviation of 4.38. For each participant, we conducted a pre-test questionnaire to measure the demographics, and a post-test questionnaire to obtain simulator sickness levels R. S. Kennedy et al. (1993) after the VR experience. The participants reported a median of four on a 5-point Likert scale for VR device experience.

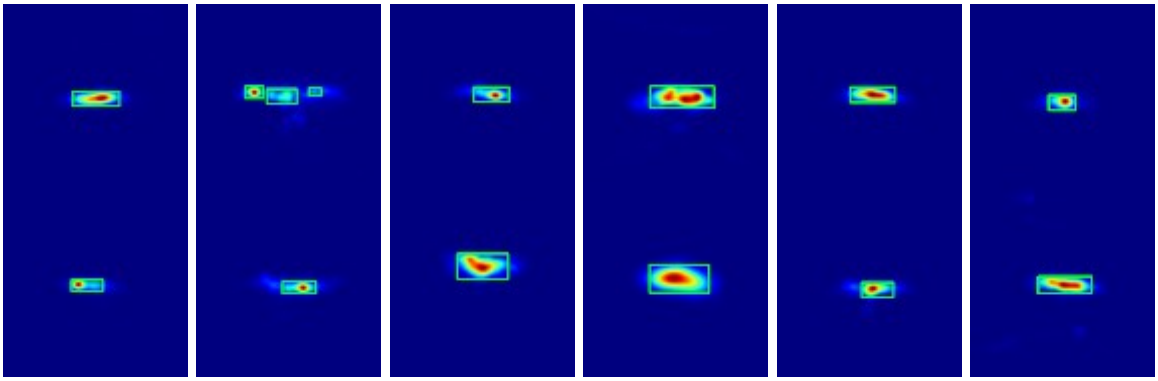

Figure 4.1: Heatmaps and bounding boxes. Time spent fixating on the viewport increases from navy blue to dark red.

4.4.1 Equipment and Safety

The experiments were conducted on a workstation with an Intel(R) Core(TM) i9 - 9900K CPU @3.60GHz and an NVIDIA RTX 2080 Ti GPU. An HTC Vive Pro Eye with an integrated Tobii Eye Tracker was used as the primary VR headset, operating at a frame rate of 120 Hz with eye-tracking accuracy of $0.5^{\circ} - 1.1^{\circ}$. Eye calibration for each participant was performed before each experiment in every user study to ensure accurate gaze estimation. The experiments were approved by the institution's Environmental Health and Safety Board (EHS) and the Ethics Research (ER) board and conducted in accordance with relevant guidelines and regulations, with informed consent obtained from each participant.

4.4.2 Analysis - Average Foveal Region

Figure 4.1 shows heatmaps for all the participants, indicating the average fixation distribution in the FoV. The following equation shows the smallest ellipse that encloses all the bounding boxes from each participant's heatmap Gärtner and Schönherr (1997).

$$\frac{(x - 0.491)^2}{0.28^2} + \frac{(y - 0.451)^2}{0.144^2} = 1$$

It marks the average gaze region during a chaotic VR experience where the users must repeatedly change their focus and rotate their heads from side to side. It shows that the users mainly focused around the center of the viewport, indicating saccade-like actions during rapid head rotations. Since

the human's horizontal FoV spans roughly 150° Glaholt (2016), we hypothesize that when users rotate their heads in VR at a velocity of more than 150°/sec (for multiple frames), it is safe to assume that the FoV is distorted enough to cause temporary blindness.

4.4.3 Simulator Sickness Questionnaire (SSQ)

We calculated the Total Severity (TS) and its sub-scales (Nausea (N), Oculomotor (O), and Disorientation (D)) using the formulas from R. S. Kennedy et al. (1993). Most participants (83.34%) reported no significant signs of simulator sickness, with the highest expected average for disorientation (16.24) due to vestibular disturbances from repeated head rotations. Table 4.1 summarizes the post-test SSQ responses.

Table 4.1: SSQ responses - Preliminary study

Scores	Mean	Median	SD	Min	Max
N	7.95	0	16.19	0	57.24
O	9.475	7.58	11.71	0	37.9
D	16.24	0	26.43	0	83.52
TS	12.16	5.61	18.19	0	63.58

4.5 Technical overview

To validate our hypothesis and minimize the impact of vestibulo-ocular reflex (VOR) and smooth pursuit eye movements, we devised a real-time deep neural network that employed head rotation data to predict saccades and execute redirection. Figure 4.2 shows the system overview of our technique whereupon predicting a saccade, the VE is dynamically adjusted in accordance with the Redirection algorithm, effectively guiding users during the onset period of the anticipated saccades. This technique was designed and optimized to achieve processing speeds comparable to high-end eye trackers. In our conclusive user study, we confirmed the effectiveness of our approach in VR, with users experiencing seamless and uninterrupted redirection without any visual artifacts. Our results thus demonstrate the immense potential of this approach in providing a highly immersive and engaging VR experience.

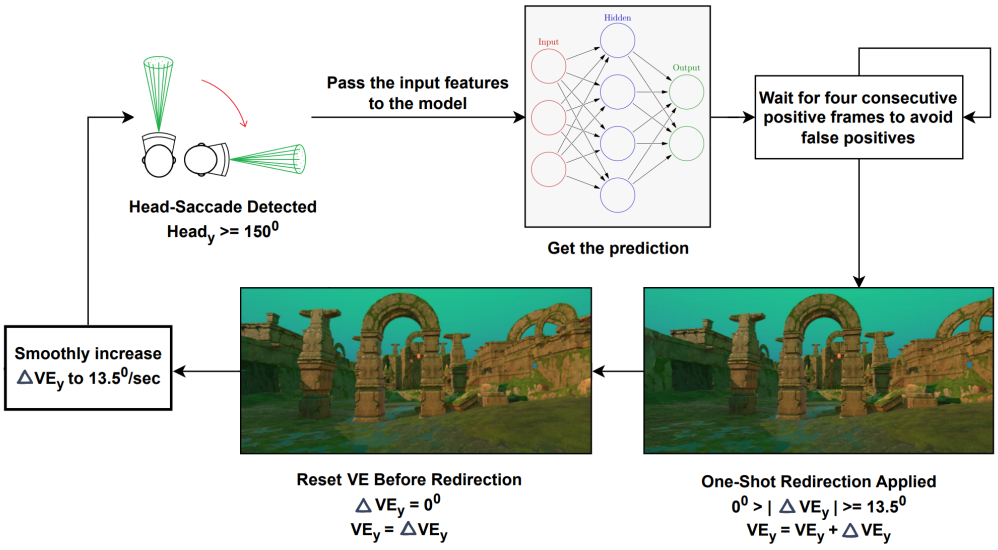

Figure 4.2: System Overview

4.6 Data Acquisition & Learning

The present study proposes a real-time saccade prediction model that is specifically designed to operate during tracked head rotations. Figure 4.3 illustrates the model architecture, which was developed using the Pytorch framework. The model consists of four 1D convolutional layers, each specified with a kernel size of 3, a padding and stride of 1, and a constant width of 10. The convolutional layers are followed by a layer to flatten features and four fully connected layers. Given that our data is a time series, an input window size of nine was chosen, with consecutive samples selected from the dataset. The first convolutional layer comprised nine input channels and 16 output channels, while the second, third, and fourth layers had (16, 32), (32, 64), and (64, 128) input and output channels, respectively. The output of the convolutional layers is then flattened to produce a 128 x 10 input for the first fully connected layer, which has 1024 output channels. The subsequent fully connected layers have (1024, 512), (512, 256), (256, 128), and (128, 1) input and output channels, respectively.

To address potential vanishing gradients problems, the convolutional layers and the first three fully connected layers in our model employed a Leaky Rectified Linear Unit (Leaky ReLU) activation function in the forward pass. On the other hand, the fourth fully connected layer employed

a Rectified Linear Unit (ReLU) activation function to produce entirely positive values for binary classification in the final layer. A Sigmoid activation function was then applied to the output of the final layer to obtain probabilities between 0 and 1. Finally, the Adam optimizer was used for backpropagation with a binary cross-entropy (BCE) loss function and a learning rate of 0.001, and the model was trained over ten epochs with a batch size of 128.

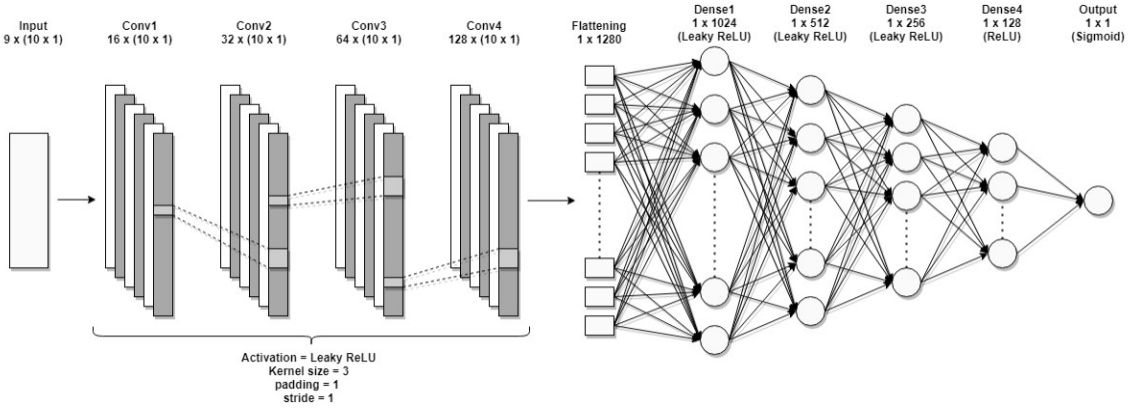

Figure 4.3: Model architecture

4.6.1 Data Acquisition

To gather data for the training, we designed an application similar to that used in our preliminary study. The primary task remained the same, and the tests were divided into three timed-trials of five, ten, and fifteen minutes each to minimize the risk of simulator sickness. Following each trial, participants were allowed a short break as needed. User interactions in this study were consistent with those of the first study, and participants were subject to the same experimental constraints. Finally, Kennedy's sickness simulator questionnaire R. S. Kennedy et al. (1993) was administered after the immersive experience to quantify the participants' comfort.

A cohort of 14 participants, with an average age of 25.86 years and a standard deviation of 4.37, was recruited for the data collection phase. To mitigate any redundancy in data, none of the participants were involved in the first confirmatory study. The participants had normal or corrected-to-normal vision, and their self-reported familiarity with VR devices and eye-tracking devices had a median rating of four and three, respectively, on a 5-point Likert scale, with one being the least familiar.

4.6.2 Learning & Inference

We used time series data consisting of historical head rotations and a binary value indicating the onset duration of saccades as the ground truth to train our model. To identify the most effective features for saccade detection, we initially computed several hand-crafted features from head rotations. After careful examination of the impact of each feature on saccade detection, we determined that the head's historical fixation direction, angular velocity, and acceleration between three successive frames were the most pertinent features for training our model. Specifically, assuming f_{t-2}, f_{t-1} and f_t denote the last three frames, we filtered the following features:

- h_{t-2}: head rotation at f_{t-2}
- h_{t-1}: head rotation at f_{t-1}
- h_t: head rotation at f_t
- ΔD_y: Change in direction (f_{t-1} to f_t)
- V_2: Angular Velocity (f_{t-2} to f_{t-1})
- V_1: Angular Velocity (f_{t-1} to f_t)
- ΔV : Change in Angular Velocity
- A_2: Angular Acceleration (f_{t-2} to f_{t-1})
- A_1: Angular Acceleration (f_{t-1} to f_t)
- ΔA: Change in Angular Acceleration

Given that humans are less sensitive to horizontal changes, we opted to only rotate the virtual environment horizontally for redirection purposes. Consequently, each of the ten features was measured along the yaw/y/UP axis. Furthermore, to ensure independence and accuracy, each feature was recorded separately in world space, i.e., eye-in-world and head-in-world velocities.

During the data collection process, a total of 1,009,667 data points were saved, with each data point consisting of ten features and the corresponding ground truth labels for saccadic events. In line with the definition established by Sun et al. Sun et al. (2018), saccadic events were defined as ballistic eye movements with an angular velocity greater than $180^{\circ}/s$. Consequently, for the duration of each saccadic event, the ground truth was set to 1, and for the remainder of the frames,

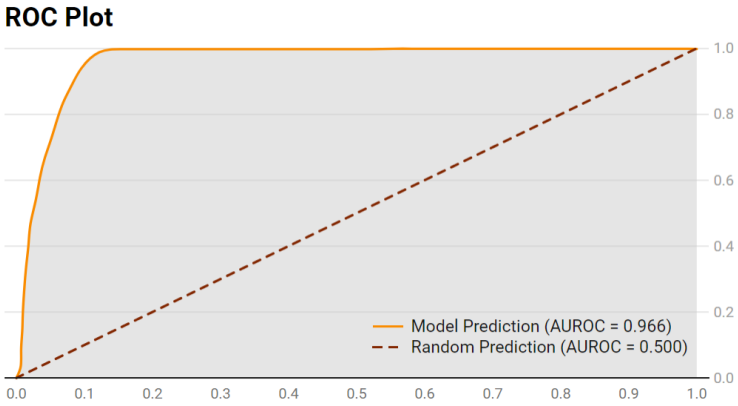

Figure 4.4: ROC is plotted against FPR on the x-axis and TRP on the y-axis.

it was set to 0. In addition, as our technique only predicts saccades during apparent head rotations, any saccadic event that occurs when the head rotation velocity is less than 150°/sec was excluded from the dataset. The ground truth was artificially set to 0 for those frames. We established this threshold in our first user study that confirms the head-eye relationship. The threshold accounts for smooth pursuit and VOR eye movements. Finally, we trained our binary classification model using an input window size of nine.

4.6.3 Training, validation, and testing

The dataset was partitioned into three distinct subsets: a training set, a validation set, and a test set, in an 80:10:10 ratio. Following each epoch of the training process, the model's predictive ability was assessed using the validation set, while the final performance evaluation was conducted using the test set. The validation phase resulted in an average precision of 89.91% and a mean accuracy of 93.41%, whereas the precision achieved on the test set was 88.72% with an accuracy of 93.51%.

In addition, the model was validated by computing the Area-Under-the-Curve (AUC) of the Receiver Operating Characteristic (ROC) curve. The ROC graph was plotted for different thresholds against the True Positive Rate (TPR) and False Positive Rate (FPR), as shown in Figure 4.4. The dashed line represents the worst-case scenario of completely random predictions (AUC = 0.5), while the orange curve represents the actual ROC curve. The computed AUC for the ROC curve was higher than 0.5 and close to 1 (AUC = 0.966), indicating good performance.

4.6.4 Simulator Sickness Questionnaire (SSQ)

Table 4.2 presents the results of the SSQ responses obtained after the data acquisition phase. Expectedly, the data shows that a significant proportion (71.43%) of the participants did not report any signs of simulator sickness. Among those who reported experiencing symptoms, disorientation had the highest expected average score.

Table 4.2: SSQ responses - Data Acquisition.

Scores	Mean	Median	SD	Min	Max
N	5.45	0	7.21	0	19.08
O	10.83	3.79	17.77	0	60.64
D	14.91	6.96	17.66	0	41.76
TS	11.49	7.48	14.99	0	48.62

4.7 Evaluation of RDW

In the final experiment, we developed an application to assess the overall effectiveness of our event-based redirected walking system that incorporates saccade predictions to adjust the VE in real-time. To ensure an uninterrupted and seamless immersive experience, we leverage the principle of change blindness, which allows us to mask the redirection events.

4.7.1 Method

Change blindness refers to the inability of individuals to perceive significant changes within their FoV D. Simons and Rensink (2005), which is often associated with attentional deficits. This perceptual phenomenon is frequently observed in VR, where users are typically engaged in various cognitive tasks such as training simulations and games. In light of this, we developed a first-person treasure hunt game set on a mysterious island inhabited by dragons and swamp crawlers to evaluate the effectiveness of our redirected walking system.

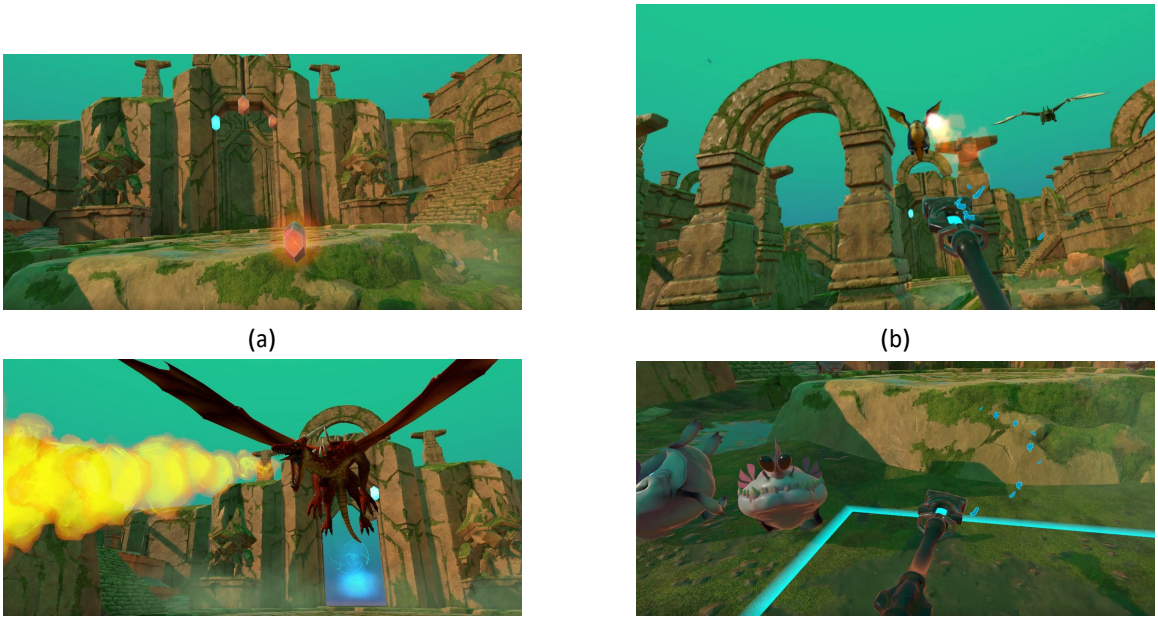

(a) (b) (c) (d)

Figure 4.5: (a) Virtual environment used during the evaluation of the proposed RDW technique, and the directional indicators (glowing crystals) that reveal the next destination, one at a time. (b) User's perspective of dragon enemies and the mystical staff. (c) The Final Dragon (d) Swamp crawlers

4.7.2 Application & Task:

The primary aim of this immersive game entails collecting three crystals in the ruins of an abandoned ancient arena. To achieve this objective, participants were asked to navigate various pre-defined virtual locations identifiable by glowing crystals from their initial spawn position. Prior to the start of the game, a brief tutorial elucidating diverse interactions and the ultimate goal is provided to the participants. After completing the tutorial, the location of the first crystal and directions to the second crystal is disclosed. Each crystal obtained unlocks a unique magical power. The first crystal confers the power to throw lightning bolts.

The effect of change blindness was further intensified in our study by introducing a secondary target retrieval task in the form of enemies Peck, Whitton, and Fuchs (2008) during the traversal towards the second destination. These targets, in the form of tiny dragons, spawned at a distance, flew in a random orbital pattern around the user with at least 20° of separation between them and were accompanied by an audio cue. Figures 4.5b, 4.5a, 4.5c and 4.5d show the screen captures from the user's perspective for the final user study, and Figure 4.5c shows the participant's perspective

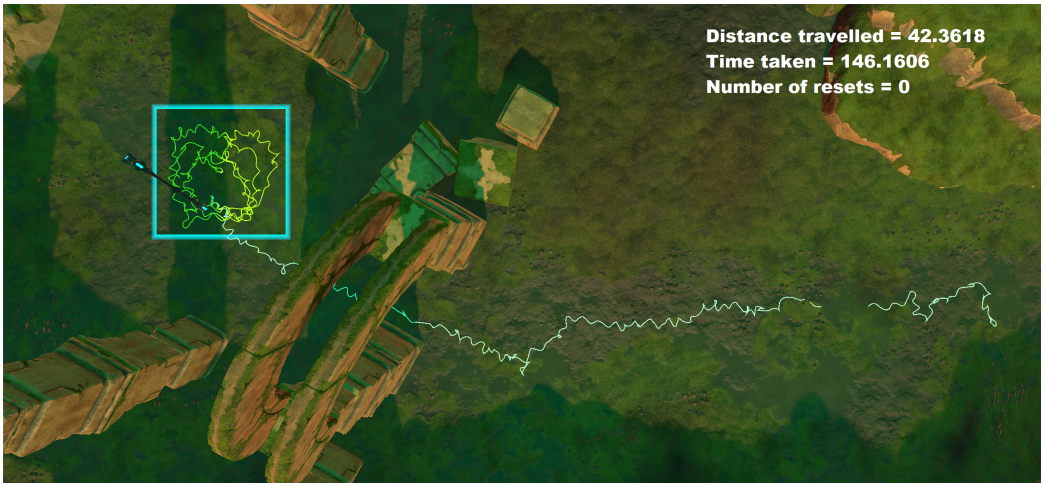

Figure 4.6: Physical and virtual paths taken by a user during evaluation. The current statistics are shown on the top right. Cyan-coloured box indicates the 3.5 × 3.5m^2 PTS, while physical and virtual paths are marked with gray and yellow, respectively.

during the task. Participants were instructed to use their newly acquired lightning bolt power, which could be activated by pressing a button on the Vive controller, to zap these dragons.

The experiment comprised a cognitive task that induced recurrent head rotations to augment the likelihood of predicting saccades, thus enhancing redirection frequency. The minimum linear distance between the initial and the second destinations was substantially greater (38m) than the maximum feasible straight distance covered in the physical tracking space (4.95m). The PTS is identified by the cyan-colored box in Figure 4.6, while the user's paths during the final trials are depicted by the yellow and grey lines, representing the virtual and physical paths, respectively.

Redirection algorithm:

We implemented a steer-to-center redirection algorithm and employed a 2:1 turn reset mechanism for the safety of our users.

Procedure:

The study involved two distinct experiments that examined the effectiveness of the proposed redirected walking technique as an independent variable. The first experiment tested the technique in an experimental condition where both redirection and reset mechanisms were enabled, while the

second experiment repeated the same course of actions in a controlled condition where only the reset mechanism was enabled. Data collection at the end of each experiment included relevant variables such as the number of resets, total distance covered, and total time taken, which were used for later analysis. Furthermore, the experiments were separated by a brief intermission to reduce simulator sickness.

4.7.3 Participants

While designing the experiment, we conducted a power analysis with an effect size of 0.7 to determine a sample size of 20, resulting in a power of 0.956. We recruited 32 participants for the study, with an average age of 27.38 years and a standard deviation of 3.74. Twelve participants were selected for the fine-tuning experiments, while the remaining 20 participated in the final evaluation. To gather information about their experience with VR headsets and eye-tracking devices, we used a 5-point Likert scale where one represented the least familiarity. The participants reported a median score of three for their experience with VR headsets and two for their experience with eye-tracking devices, with normal or corrected-to-normal vision.

Prior to commencing the experiment, the participants were briefed about the emergency reset feature and the objectives of the game. They were advised to maintain an average pace during walking and remain fully engaged with the task throughout. Additionally, at the experiment's conclusion, the participants were requested to provide subjective feedback, with a particular query on their observations of any visual disparities or shifts in the virtual environment.

4.7.4 Results

This section provides a comprehensive analysis of the outcomes obtained from the final study. In the initial segment, we present quantitative results that assess the efficacy of our prediction model. Subsequently, we exhibit quantitative outcomes that determine the overall performance of our redirected walking technique.

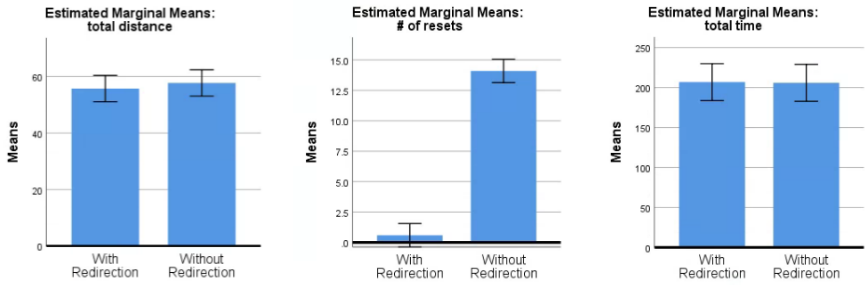

Figure 4.7: Marginal means reported by one-way ANOVA analysis of (a) Distance Travelled, (b) Number of Resets in PTS, and (c) Total Time Taken. Confidence Interval = 95%

Quantitative Performance Evaluation:

Our prediction model exhibited satisfactory performance within real-time constraints, with an average accuracy of 94.75%, a recall of 99.99%, and a sensitivity of 94.68%. The model's F1-score for real-time data was 0.72. Our analysis of the results revealed that the training dataset was heavily skewed towards negative training examples, resulting in a greater number of false positives due to the relatively low number of true positives. Consequently, while the model can correctly predict a saccade, it often mispredicts its duration by a few frames before and after the actual saccade, resulting in more false positives. Therefore, the question arises about when one should apply redirection to ensure its imperceptibility. We conducted four fine-tuning experiments involving 12 participants to address this question, compensating for the 56.52% precision obtained in the final study.

Fine-tuning:

In the initial fine-tuning experiment, redirection was performed at the first positive saccade prediction. As a result, two participants noticed the redirection as the VE was rotated just before the actual saccade occurred. In subsequent fine-tuning experiments, the VE was only rotated if two or three consecutive frames produced positive saccade predictions. As a result, all participants in the second experiment and two in the third experiment noticed an angular shift, while one participant in the third experiment experienced no distractions. Finally, in the fourth experiment, redirection was only applied if positive predictions were made over four consecutive frames. All participants reported no distractions, and none could perceive any angular shifts in the VE.

Based on the results obtained from these fine-tuning experiments, we have determined that a

window size of four consecutive positive predictions is the optimal choice for applying the redirection in the final user study for evaluation.

Performance Evaluation of RDW:

In order to thoroughly evaluate the effectiveness of our proposed redirected walking (RDW) technique, we carried out a comprehensive statistical assessment. For this purpose, we employed a one-way within-group analysis of variance (ANOVA) with repeated measures and a significance level of $\alpha = 0.05$. This method of analysis allowed us to distinguish the distinct impacts our RDW technique had on an array of dependent variables, including the number of resets, the total distance covered, and the overall time required to complete tasks. To further analyze our findings, we calculated the effect sizes for all the dependent variables using partial eta squared (η_p^2) values obtained from the ANOVA results.

The results of our analysis indicated that toggling redirection accounted for almost 86.5% of the observed variance in the number of resets while contributing only 0.6% and 0.0% to the variances observed in total distance traveled and total time taken, respectively. Furthermore, we observed a statistically significant difference between the number of resets ($F(1, 62) = 396.094, p < 0.001$) with and without using our proposed redirection technique. However, there was no statistically significant effect of toggling redirection on the total distance traveled ($F(1, 62) = 0.366, p > 0.05$) or time taken ($F(1, 62) = 0.003, p > 0.05$).

We plotted the means for the dependent variables with respect to the independent variable in Figure 4.7, with error bars at a confidence interval of 95%. These findings demonstrate the effectiveness of our proposed RDW technique in reducing the number of resets while having no significant impact on the total distance traveled and time taken.

Moreover, as reported by Sun et al. Sun et al. (2018), the proposed technique resulted in an average absolute gain of 12.59° per redirection to the virtual environment (VE), with a frequency of approximately 0.55 redirections per second. Consequently, on average, each participant's field of view (FoV) gained 1375.1° of absolute gain, with a standard deviation of 432.131°. Furthermore, during the final user study, all participants walked at least 38 meters straight in the VE while walking within a $3.5 \times 3.5 m^2$ physical tracked space. Figure 4.6 illustrates the path taken by a participant

Table 4.3: SSQ responses - Evaluation

Scores	Mean	Median	SD	Min	Max
N	10.14	9.54	12.1	0	38.16
O	9.95	3.79	15.21	0	60.64
D	13.05	13.92	16.18	0	55.68
TS	12.39	9.35	14.55	0	56.1

during the final study. Additionally, by redirecting the user only after four consecutive positive predictions, compensating for the lower precision, none of the participants reported any perceptible visual disparity or distraction in the VE.

The majority of our participants (71.87%) did not report any significant signs of simulator sickness. Consistent with our previous studies, the mean score for disorientation was higher than the other subscales. However, we observed a decrease in the disorientation score from our first study (16.24) to the Data Collection phase (14.91) and the final user study (13.05) due to our continuous efforts to mitigate this effect. An overview of the post-test SSQ responses for the final user study is presented in Table 4.3. Overall, every participant had a seamless experience, with one participant remarking, "I felt like I was walking straight inside the game, but I was actually walking in circles to my surprise."

4.8 Discussion

Our technique aims to predict saccades and perform saccadic redirection on commodity hardware, eliminating the need for specialized eye-tracking equipment. However, the technique is best suited for use in applications with a moderate cognitive workload, such as battlefield training simulations or game-like scenarios. It proves to be particularly effective when users engage in tasks that necessitate repeated head rotations. It is known that saccades tend to be purely eye-driven when the object of interest is within 18° of the fixation point Stahl (1999). As such, our technique performs optimally when the targets that illicit head rotations are separated by 20°. The versatility of our saccade prediction technique presents a range of cost-effective potential solutions to applications, like redirected walking and foveated rendering for consumer-grade hardware.

4.8.1 Evaluation

In our final user study, we thoroughly evaluated the effectiveness of our proposed system for redirected walking through both qualitative and quantitative means. The results demonstrate that our system eliminates the need for high-end eye trackers and allows users to explore relatively straight virtual distances of at least 38 meters within a confined room-scale physical tracking space of 3.5x3.5m^2. Figure 4.6 shows a long, straight virtual walk with a completely circular physical path. Despite the inherent challenges of working with an imbalanced dataset resulting from artificially setting the ground truth to 0 for frames with head rotation velocity less than 150°/sec, we were able to fine-tune the hyper-parameters to apply redirection only when four consecutive positive predictions were made. As a result, each of the 20 participants completed the task without any visual disparities or distractions, despite the model's precision rate of 56.52%. When asked about any noticeable visual disparities or shifts in the virtual environment during the post-test questionnaire, one participant responded, "I did not notice any visual disparity or shift. The experience was smooth."

4.8.2 Limitations

Seemingly Forced Head-Rotations:

Our saccadic RDW technique only applies to scenarios with repeated left-right head rotations. As this work primarily focuses on training simulations and hyperactive game-like experiences, we have yet to incorporate natural head rotations and expand the scope of our technique to a broader range of use cases.

Tunnel/focused vision:

It is important to note that the effectiveness of any saccadic redirected walking (RDW) technique is inherently limited when users refrain from making head or eye movements. In such cases, it is difficult to achieve the intended redirection effect. However, in our final user study, the task in our final user study elicits repeated head rotations and serves as a cognitive workload to introduce an attention deficit, reinforcing the effect of change blindness.

Comparison:

Rather than directly comparing our technique to other redirected walking (RDW) methods requiring embedded eye-trackers, we used a quantitative analysis using a resets-only baseline. Adopting this approach gave us a fair and accurate measure of our technique's effectiveness. This methodology allowed us to evaluate the performance of our technique in a controlled setting, without the confounding effects of eye-tracking hardware, while also enabling its comparison with future RDW methods.

4.8.3 Merits

Our technique does not require additional eye-tracking hardware but relies only on positional tracking. This approach gives us essential advantages, including lower computing requirements, improved accessibility, and reduced hardware cost, distinguishing it from state-of-the-art Eike et al. (2017); Joshi and Poullis (2020b); Langbehn and Steinicke (2018a); Sun et al. (2018). In addition, recent developments in positional tracking for smartphones make it possible to optimize the model for cell phone processors, potentially enabling saccade prediction on widely available mobile VR.

4.9 Conclusion and Future Work

This work presents a novel event-based redirection technique that employs a CNN-based deep neural network to predict visual suppression resulting from saccades during apparent head rotations. By exploiting these predicted suppressions, the technique applies subtle rotations to the VE, which induces a change in the user's physical walking direction while perceiving a straight motion in VR. The effectiveness of the proposed technique was evaluated in three user studies, which confirmed its ability to handle long straight walks. The first two studies aimed to establish the relationship between head and eye directions and acquire training data for the prediction model. Finally, the third user study evaluated RDW with this predictor and demonstrated its effectiveness in VR applications. It enabled users to explore large virtual spaces by walking in applications like hyper-active immersive games or training simulations.

Our technique is only effective during specific tasks that elicit repeated head rotations; we acknowledge this limitation. However, it offers several advantages over existing RDW techniques, including accessibility to a broader audience and maintaining the intended VE's fidelity. Our approach unlocks saccadic redirection on widely used hardware without eye-tracking, making it an attractive solution for VR applications, including foveated rendering.

Future Work

There are many potential avenues for improvement in our saccadic redirection technique. We plan to predict saccades in normal viewing conditions, explore the correlation of saccades with different stimuli, and incorporate saliency and depth maps of the FoV into our dataset. Additionally, based on the results obtained by Sidenmark et al. Sidenmark and Gellersen (2019), we will exploit the correlation between user torso movements and gaze and explore the possibility of adding features from hand, foot, and torso movements to our model. Lastly, we hope to expand the scope of redirection to include translational gains.

Chapter 5

Conclusion and Future Work

Over the past years, virtual and augmented reality (VR/AR) technologies have witnessed remarkable advancements, resulting in improved user experiences and expanding the range of applications in various domains. However, a significant challenge in this field is the locomotion dilemma. This thesis addresses the locomotion dilemma in virtual reality, presenting a comprehensive investigation of various techniques to enable natural navigation in virtual environments within confined physical spaces. The primary focus of this thesis is to investigate and propose innovative solutions that contribute to the greater understanding and expansion of locomotion techniques in VR while enhancing the user's sense of presence and immersion. The research focused on three primary approaches, each with a unique contribution to the field, supported by experimental evidence and user studies.

The first approach, "Portal to Knowledge," introduced a divide-and-rule strategy to provide users with the illusion of exploring a substantially larger virtual environment while being in a limited physical space. This method was applied to develop an alternative to traditional physical libraries and online book retrieval systems. Through a user study with 45 participants, we demonstrated the efficiency of this approach in reducing mental, physical, and overall workload while providing a more convenient experience. Furthermore, the study highlighted the application's potential to revamp library exploration, as most participants preferred the proposed solution over traditional libraries and online retrieval methods. This work was presented at SPIE AR/VR/MR Joshi and Poullis (2020c), one of the highly reputed venues in the field of immersive technologies.

In the second approach, we presented a rotation-based redirection technique that leveraged the psychological phenomenon of inattentional blindness induced by cognitive tasks and strengthened its effect using dynamic foveated rendering. Users' entire field of view is divided into zones, and spatially-varying rotations are applied according to their importance. The final frame buffer is rendered by combining the foveal and peripheral zones from two cameras. Each view comes from the same scene but is slightly rotated according to the redirection algorithm. Finally, by exploiting naturally occurring visual suppressions, such as eye blinks and saccades, the technique refreshes the entire framebuffer to apply subtle rotations to the virtual environment without the user's knowledge. The effectiveness of this method was validated through three user studies, which confirmed its ability to handle long straight walks and enable users to explore open-world virtual environments. This approach significantly reduced hardware costs and made the technique accessible to a broader audience by eliminating the need for intrusive visual augmentations in the user's field of view. This work was published by IEEE in their open-access journal, IEEE Access Joshi and Poullis (2020b), and was also presented at SIGGRAPH 2020 Joshi and Poullis (2020a). Furthermore, after two years of market research, we recently obtained a complete US and Canadian patent for this work Joshi and Poullis (2023a).

Finally, our last approach aimed to eliminate the requirement of using eye trackers for an efficient redirected walking solution. This work proposes an event-based redirection technique that employs a CNN-based deep neural network to predict saccades during apparent head rotations. This method allowed for subtle rotations to the virtual environment during predicted visual suppressions, enabling users to walk freely in large-scale virtual spaces. The effectiveness of the proposed technique was evaluated in three user studies, which confirmed its potential to improve navigation in various VR applications, such as hyper-active immersive games and training simulations. This approach unlocked saccade prediction on widely used consumer-grade hardware without eye-tracking, making it an attractive solution for a wide range of VR applications, including foveated rendering and redirected walking. This work is invited for publication in the Journal of Computer Animation and Virtual Worlds Joshi and Poullis (2023b) and is accepted for presentation at the 36th International Conference on Computer Animation and Social Agents (CASA 2023). Additionally, a

complete US and Canadian patent application for this work is under review.

These research outcomes collectively contribute to the development of more immersive, natural, and cost-effective locomotion techniques that maintain the fidelity of virtual content and minimize the need for specialized hardware. Furthermore, by addressing the limitations of current state-of-the-art techniques, the proposed solutions in this thesis offer valuable insights and pave the way for future advancements in VR navigation systems.

5.1 Future Work

As we reflect on the contributions of this thesis, numerous opportunities for future research emerge, aiming to advance VR locomotion techniques and user experiences further. The following directions present a cohesive and inclusive approach to guide future research endeavors in this domain:

(1) Investigating Novel Forms of Visual Suppression: Researching other forms of visual suppression, such as phases of nystagmus, and their potential application in redirected walking techniques. This can lead to the development of more efficient and effective locomotion systems that take advantage of natural visual phenomena to enhance user experiences.

(2) Context-aware Locomotion Techniques: Developing context-aware locomotion solutions that dynamically adapt to the user's current environment and task requirements. This could incorporate environmental cues, user preferences, and task constraints into the locomotion algorithms to provide a more seamless and intuitive navigation experience.

(3) Collaborative and Multi-user Experiences: Investigating how the proposed locomotion techniques can be adapted to support collaborative and multi-user virtual experiences. This could involve exploring methods for synchronizing users' movements and interactions within shared virtual environments, enabling more engaging and immersive social VR experiences.

(4) Combining Locomotion Techniques: Investigating how the proposed locomotion techniques can be combined and integrated with other existing methods to create more comprehensive

and versatile navigation solutions. This could involve developing hybrid systems that utilize multiple techniques in parallel or sequentially, depending on the user's needs and preferences.

(5) Incorporating Multimodal Feedback: Expanding the proposed techniques to include multimodal feedback, such as haptic or auditory cues, to enhance the sense of presence and immersion in virtual environments. By integrating additional sensory information, we can create a more convincing and engaging VR/AR experience for users.

Pursuing these future research directions will enhance the techniques presented in this thesis and contribute to advancing virtual reality navigation solutions. Furthermore, by refining and expanding the techniques proposed in this research, we can further push the boundaries of VR applications and improve their impact on various domains, including education, entertainment, and training.

Moreover, collaborations with industry partners and researchers from other fields can lead to the development of new applications that incorporate these navigation techniques, enabling better user experiences and opening up new possibilities for virtual and augmented reality. As we continue to explore and innovate, the ultimate goal is to make immersive technologies more accessible, user-friendly, and efficient, transforming how we interact with digital environments in our daily lives.

Appendix A

Portal to Knowledge: a divide-and-rule approach to the locomotion dilemma

This chapter is based on our paper "Portal to Knowledge: a Virtual Library using Marker-less Augmented Reality system for Mobile Devices" Joshi and Poullis (2020c), published in SPIE AR/VR/AR, 2020. The manuscript evaluates an application, Portal to Knowledge, that implements the divide-and-rule approach to tackle the locomotion dilemma in mixed reality. By dividing virtual environments into smaller chunks and mapping them repeatedly onto the available physical space, the application allows users to explore a vast virtual library while walking in a confined physical area. The related contribution of this manuscript includes developing and implementing the divide-and-rule locomotion method in the context of a life-sized virtual library exploration application. We evaluate the efficacy of this application through a user study and compare the user experience with traditional library navigation and online book retrieval methods. Furthermore, as the divide-and-rule approach repeatedly remaps different virtual environments on top of the same physical space, it only supports closed virtual environments to maintain immersion. However, this research serves as the foundation for understanding the locomotion dilemma and offers a practical solution for exploring virtual spaces, paving the way for further developments presented in this thesis.

A.1 Abstract

Since exceedingly efficient hand-held devices became readily available to the world, while not being a relatively recent topic, Augmented Reality (AR) has rapidly become one of the most prominent research subjects. These robust devices could compute copious amounts of data in a mere blink of an eye. Making it feasible to overlap computer-generated, interactive, graphics over real-world images in real-time to enhance the comprehensive immersive experience of the user.

In this chapter, we present a novel mobile application that allows users to explore and interact with a virtual library in their physical space using markerless AR. Digital versions of books are represented by 3D book objects on bookcases similar to an actual library. Using an in-app gaze controller, the user's gaze is tracked and mapped into the virtual library. This allows the users to select (via gaze) a digital version of any book and download it for their perusal. To complement the immersive user experience, continuity is maintained using the concept of Portals while making any transition from AR to immersive VR or vice-versa, corresponding to transitioning from a "physical" to a virtual space. The use of portals makes these transitions simple and seamless for the user. The presented application was implemented using Google AR Core SDK and Unity 3D, and will serve as a handy tool to spawn a virtual library anytime and anywhere, giving the user an imminent mixed sense of being in an actual traditional library while having the digital version of any book on the go.

A.2 Introduction

In a recent study by Sansone et al Sansone and Sansone (2013) it was shown that the modern society, especially the young generation, is reclining progressively towards the digital realms for almost any kind of solution. With changing tides, it has become imperative for the organizations to serve their purpose through the medium that their users are most familiar with. One of the common example of such organization is a library. For some, due to the tedious taxonomy and specific jargon, it is really hard and often frustrating to navigate their way through the library without any special aid. While some prefer not to go to the library at all because of the unavailability of any in their vicinity. In this chapter, we present a novel mobile application, named "Portal to Knowledge", which serves the exact purpose of providing users with an experience of going to an actual library

without physically going there. Using our application, users can drop a portal to their personal library anywhere and anytime they want. The application employs marker-less AR technology which enables users to roam freely inside their virtual library, browsing books, without worrying about losing the track of AR markers.

As AR is not a very recent technology, due to the paucity of devices with good processing power and their unavailability amongst public before a couple of decades, applications of AR were only limited to Head Mounted Displays (HMDs) and hence, expensive research. After the introduction of hand-held devices with lightening fast processing speeds and graphics rendering powers, AR has gained it's now known fame. Augmented Reality simply means to present the real actuality augmented by an additional layer of contextual computer generated content to enhance the overall immersive experience of the user. As per Azuma et al in R. T. Azuma (1997), this additional content is superimposed on the user's field of view, complementing his/her visual perception, and not completely replacing it. In Craig (2013), Craig et al also suggested that this additional information should be spatially and temporally registered with the physical world and should be interactive in real-time Biocca (1997),Biocca, Harms, and Gregg (2001),Loomis, Blascovich, and Beall (1999). This makes it a little different from Virtual Reality (VR). In VR, the user's field of view is completely replaced by the virtual content which gives him/her the sensation of actually being in a whole different virtual world.

A.3 Background and Related Work

Below we present a brief history and an overview of the state-of-the-art. We have categorized the related work in terms of (a) a brief overview of VR, (b) a brief overview of AR and various tracking methods, and (c) enhancing library experiences.

A.3.1 A brief overview of VR

Virtual Reality and Augmented reality systems can both be used to display digital content to the user and enhance or completely replace the user's perception of the real world at that moment. Commonly, VR is defined as an immersive system in which an HMD completely blocks the user's

field of view and replaces it with two small high resolution stereoscopic screens. These two screens then display the virtual computer-generated environment to the user which can be explored and interacted-with in real-time Siegrist et al. (2019). Because of the technological and hardware limitations, there weren't many HMDs available to public before the period of 90's. Though Virtual reality is quite an old topic of research, the term was popularised by Jason Lanier when he first launched the Dataglove, during 1987 which was then used by NASA in it's VIEW system, and the EyePhone, making it the first commercially available VR solution to the market. Nowadays, consumers have a wide range of alternatives for their VR needs to choose from ranging from PC-based VR to haptic gloves and full body haptic suits.

Few of the best examples of recent publicly available PC-based VR HMDs are the Oculus Rift and HTC Vive. Both the HMDs consists of 1080 × 1200 resolution displays for both eyes and a field of view of 110 degrees making them highly immersive to the user. These HMDs have ruled the market since 2016, until recently when HTC launched an upgrade of it's vastly famous HMD, Vive. In April of 2018, HTC launched Vive Pro, with wireless functionality and with a 78% increase in the resolution of it's predecessor making it even more immersive to the user. While these are some of the best PC-driven VR solutions, there are also various mobile VR devices. Treating two halves of the phone screens as individual stereoscopic screens on the user's eye, we can get the similar immersive effect at a reasonably lower cost and better overall performance Steed and Julier (2013). Few of the good recent examples using this technique are Google Cardboard, Samsung Gear VR and Google Daydream Powell, Powell, Brown, Cook, and Uddin (2016).

Another variation of immersive VR is using a CAVE. A CAVE is a virtual reality system which enables users to experience virtual environments in a fully immersive way. It consists of four projection screens (three back-projections, one front projection) where stereoscopic images are being projected. The user can literally walk-in virtual environments and move freely in the area. Further exploration and interaction can be achieved through the use of other specialised equipment, such as data gloves, etc as in Christou, Tzanavari, Herakleous, and Poullis (2016b).

The user can perceive the stereoscopic images using a pair of active stereo glasses in 3D, therefore increasing the sense of immersion and reducing the gap between physical and virtual reality. Moreover, the user's head movements are tracked in real-time and are "translated" in movements

in the virtual environment. Thus, by moving in real space the system automatically adjusts the projections according to the user's current position and point-of-view, giving a more realistic feel. Although a number of successful applications have been reported Sutcliffe et al. (2019), a CAVE is not considered cost-effective in most scenarios.

A.3.2 A brief overview of AR and various tracking methods

Augmented Reality is a technology with which the vision sense of any human being is enhanced by superimposing some additional computer generated, context related, information onto their field of view. We can trace the first ever example of AR technology all the way back in 1968 I. E. Sutherland (1968), when Ivan Sutherland made the headset for the Sword of Damocles. Although the headset is taken to be the prime example for a VR system, one can fairly argue that it was the first example of an AR system as the headset proposed by Sutherland was partially transparent R. Azuma et al. (2001b). Mainly, AR is used to enhance the user's visual perception of the real world but is not just limited to that. Researchers have also implemented the concept using high definition spatial audio to help blind people navigate in their day to day lives Loomis, Golledge, and Klatzky (1998). Few of the most recent, famous and commercially available AR headsets are Microsoft Hololens, Magic Leap One and Meta 2. But the refined circuitry and many sensitive systems like spatial mapping, voice recognition and gesture recognition makes it too costly for the consumers to buy it for their day to day lives. Because of this costliness, Mobile AR (its equivalent for mobile devices) is becoming more and more popular amongst the consumers.

In the case of mobile AR, techniques used for tracking the environment were dramatically improved after the introduction of better hardware which could compute thousands and thousands of bits of data in real-time. From around 2011, since the hand-held AR supporting devices became readily available to the public, one can see a dramatic surge in the research conducted on AR which, according to Akayr et al in Akcayır and Akcayır (2017), was only going to escalate by 2012. We have also observed advancements in the process of registering the additional 3D content with the real world and maintaining its spatial continuity from the user's perspective R. Azuma et al. (2001b). One such advancement in the field of environment tracking and 3D content registration was the introduction of vision-based tracking systems Pang, Yuan, Nee, Ong, and Youcef-Toumi (2006). In

these types of systems, spatial data of the user with respect to his/her surroundings is extracted using computer vision algorithms Barandiaran, Paloc, and Gra na (2009), which takes a RAW RGB camera feed as it's input and detects and tracks the feature points such as contrasting colors in textures of the scene, edges, corners etc.

Easily trackable images, such as predefined markers, were predominantly used in AR systems to directly register and track the 3D content before the introduction of Marker-less tracking algorithms. Few of the good marker-based AR application examples are geun Kim and jung Kim (2014)Cieza and Lujan (2018)A. K. Dash, Behera, Dogra, and Roy (2018)M. Sagayam, Ho, Henesey, and Bestak (2018)Shatte, Holdsworth, and Lee (2014). However, the main limitation in such systems is that in most of the marker-based AR applications, the user is constrained in his/her movements with the AR device as he/she has to face the camera towards the marker at all times. This limitation is manageable for tiny play-spaces with multiple markers and hybrid approaches but it requires an efficient marker-less approach for larger play-spaces. One such hybrid approach was presented by Li et al in J. Li, Slembrouck, and Deboeverie (2015) where the authors proposed a hybrid approach to extract the user's spatial information using markers and the built-in gyroscope of the AR device. This system was better than any other simple marker-based AR application, but it still required further improvement to get rid of the markers completely.

When researchers realized the need for marker-less systems, many different techniques for tracking the environment without any markers emerged. Some researchers used Scale-Invariant Feature Transform (SIFT) algorithm for the registration of their 3D contents Kao and Shih (2013), while some used Speed-ed Up Robust Features (SURF) as their base feature extraction and tracking algorithm Yang and Cheng (2012). Perhaps the most famous algorithm was Simultaneous Localization And Mapping (SLAM). This is a technique to derive a three dimensional geometry from a systematic range of two dimensional images alongside with the rotational data from the gyroscope of the AR device. A bigger scale variant of this algorithm is known as Structure from Motion (SfM). SLAM determines the depth of any feature point in a scene by triangulating the 2D feature points in two different images al. (2006) and creates a three dimensional mesh after scanning all the detected feature points. Most popular examples of marker-less mobile AR SDKs are ARToolkit (by ARToolWorks) and Google AR Core. The application presented in this chapter uses the latest version

of Google AR Core to track the environment and maintain the position of virtual objects in their relative physical space.

A.3.3 Enhancing Library Experiences

Since the widespread availability and use of AR supporting mobile devices, mid 2010s, people have noticed some significant applications of this technology in Educational Industry. Amongst people predicting the digital age being our future and complex taxonomy systems used by the traditional libraries, an ample amount of opportunities became evident in the library sector of the educational industry that it became imperative for the librarians to move towards a hybrid model and introduce their services on a more familiar and widely used platform. By now, researchers have also realized the potential of this technology in libraries and started working on different applications and conducting surveys to help the librarians make a fully informed decision. Current state of AR was examined Massis (2015) and studies on the expanding use of AR and VR in educational industry (libraries) were conducted Hahn (2012). In this section we explore several articles that encouraged libraries to use AR technology to somehow aid the patron's overall experience. A better path-finding for the library users, the ability to highlight exclusive bookshelves and also the aided shelf reading are some of the examples of several library services that were improved Meredith (2015), Hahn (2012).

In Armstrong, Hodgson, Manista, and Ramirez (2012), Armstrong et al presented an initial version of an application under their SCARLET project which enhanced student's learning about primary as well as secondary sources in special collections using AR. It was found that in special collections environment, AR allows the students to experience both the sense of holding the actual material while simultaneously enhancing the learning ability by augmenting the contextual information around the object itself. Also similar to this, Arnhem and Spiller in Arnhem and Spiller (2014) present a mobile AR system which enhances the patron's experience by providing additional information about the local art in the library of the College of Charleston. Apart from these, some scholars also suggested applications to facilitate new users ease their way into the library using various digital orientation methods. Few examples of such applications including other library related AR applications are,

- Audio-Visual tours using QR codes Whitchurch (2012)
- Tours using tiny beacons and videos Bradley et al. (2016)
- ShelvAR: Mobile AR application for browsing and book shelf reading Hahn (2012)
- AR-Game-based approach to learn the taxonomy of a Chinese library Chen and Tsai (2012)
- Aurasma: Videos as AR overlays for self-guided library tours Mulch (2014)
- Improving library navigation for children using AR Meredith (2015)
- The SCARLET Project Armstrong et al. (2012)
- Training library employees using QR Codes and AR Rodríguez and Rivero (2016)

Against the endless improvements in AR technology and its contextual applications, one can find only anecdotal evidence in the literature of AR technology assisting the use of library services without utilizing any actual, physical, library. In this chapter, we present the novel marker-less mobile AR application: ”Portal to Knowledge” which digitally replaces the physical library with a virtual while preserving the user experience of the traditional library.

A.4 Application Overview: Portal to Knowledge

In this section, we describe the processes involved in the design and creation of the application which are divided into the following four subsections: (a) Tracking using Google AR Core, (b) Application Run-time Flow, (c) Creation of Book and Library Models, and (d) Creating the Portal.

A.4.1 Tracking using Google AR Core

The proposed AR application was developed with Unity 3D Game Engine and AR Core Software Development Kit (SDK) for Android devices. The latter is an SDK developed by Google which is mainly used for tracking the user’s current positional and rotational information in physical space, and in real-time. In order to give the users an immersive, seamless and realistic experience

of the AR virtual library, we utilized several built-in features of this SDK in our application. To determine the user's local position in the real world, Google's AR Core uses Concurrent Odometry and Mapping (COM) Langlotz et al. (2011). Although the user's location in physical space is output from COM, to get more accurate positional and rotational data, this information is combined with the inertial measurements made by the Inertial Measurement Unit (IMU) of the mobile device. IMU is a small electronic device which consists of a bundle of sensors and meters such as accelerometers, gyroscopes and magnetometers. This small device (IMU) is used to measure specific natural or human induced forces and angular rates of the mobile device and report it back to the main processing unit. Using the information output from IMU, AR Core finally determine the orientation of the user in physical space Won, Melek, and Golnaraghi (2009).

To make the tracking marker-less and even more accurate, AR Core detects and tracks feature points in the sequence of images captured by the device's primary camera in run-time. Once feature points are detected they are constantly tracked. If the feature points move outside the user's field of view, anchors are used Andujar, Mejías, and Márquez (2010) to relate the out-of-screen feature points with the screen's current content. Anchors are similar to landmarks we use to determine the specific address of any building. Once the additional digital content is dropped into the real world, some feature point is assigned to it. Hence this additional content is then tracked in physical space with respect to the feature point it has been assigned to. We call these type of feature points, anchors. Thus in terms of real world analogy, if you want to find some additional content in the physical space, you can look for the anchor point it has been assigned to and you will get to the object you were looking for Onime and Abiona (2016).

A.4.2 Application Run-time Flow

When the application is started, the user is prompted with a list of sections to choose from, having books related to different topics. Once a section of library is selected, AR Core initializes and removes any pre-stored information about the environment or the user's location in it from any previous instances of the application. Once the initialization concludes, the application displays the camera feed onto the device's screen as the AR background and starts searching for any surfaces in real-world to track. If a surface is found the application places a grid mesh on top for visualization

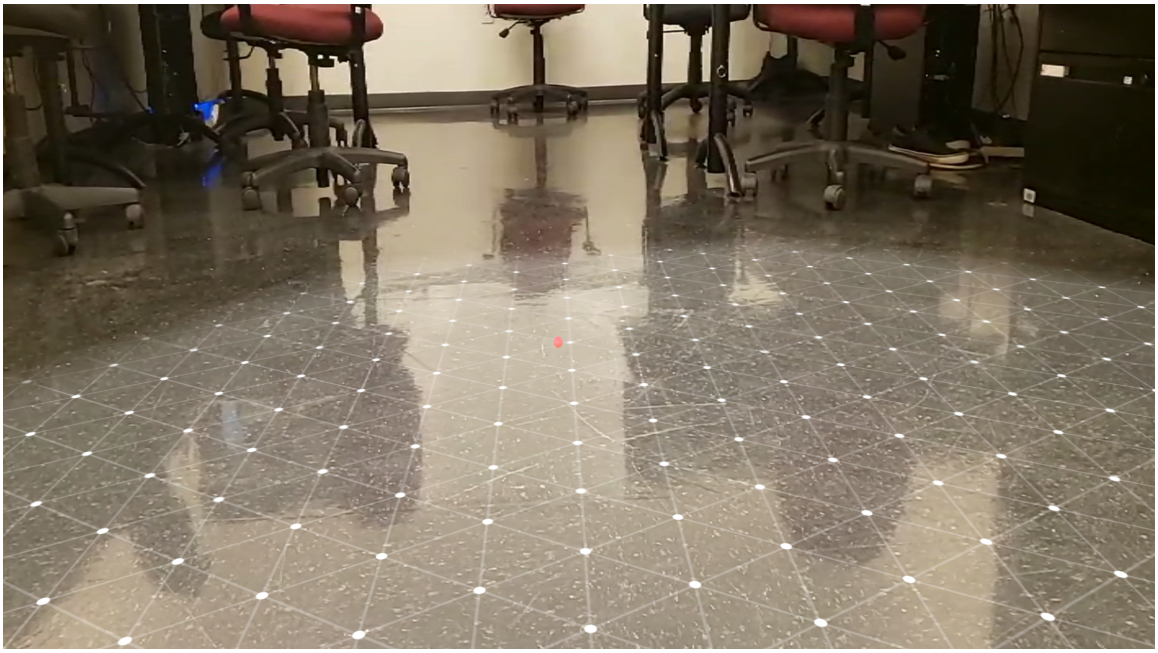

Figure A.1: Scanned surface visualized using a white colored grid mesh

purposes. A unique grid mesh with a unique color is assigned to every detected surface. Fig. A.1 shows a white grid being superimposed on the detected surface i.e. floor. While users can control the visibility of these grid meshes, they can also place additional content on these surface as per their convenience.

Once the application assigns a grid to at least one surface, it awaits for the touch input from the user which will indicate where the door of the virtual library will be placed. When a touch input occurs, a ray is cast into the virtual space shown on the screen representing the tracked physical space from the user's point of view. If the ray intersects a detected surface, a feature point on that surface is created as an anchor and the virtual library is attached to it. In contrast, if the ray does not hit any surface the application continues to await for more touch inputs. Once the portal to the virtual library becomes visible, the user can enter through it (by walking straight into it) and explore the virtual space while actually moving in their physical space. Fig. A.2 shows an example of what the screen displays once the user walks 'into the portal' transitioning from AR (as in Fig. A.1) to VR by completely blocking the view of the real world.

Once the user is inside the virtual library the application awaits for the user to select a book. The selection is done using the gaze input method. Every book in the virtual library has a white,

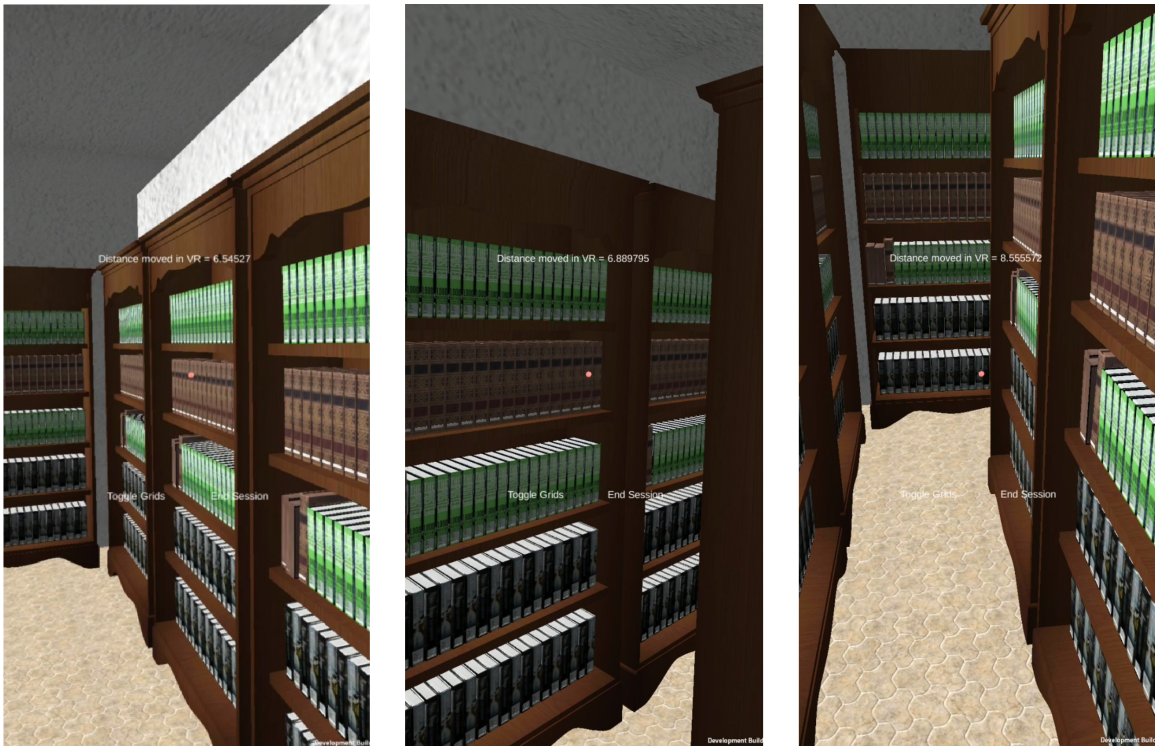

Figure A.2: Inside the virtual library from user's point of view

rectangular, gaze tag at the bottom of its visible side on the shelf. Fig. A.3 shows a number of books with the tags on each book highlighted by a red circle. To use this input method, users have to focus the gaze pointer (a pink dot in the center of the screen) onto the tag of the book that they want to open for 1 second. This will trigger an event which in turn will, if it's not already downloaded, download the PDF version of the selected book, minimize the current application, and open this PDF on the user's device.

Once the user has finished perusing the PDF version of the book, he/she can return to the virtual library just by pressing the back button or reopening the application. The application will switch back to AR and continue to track the environment, load the updated positional data from the device, and continue from where it left off. Thus, the user can peruse as many books as he/she wants to read within a single session. The flow chart showing the overall run-time flow of our application is illustrated in Fig. A.4.

Figure A.3: Gaze tags highlighted by the red circles

A.4.3 Creation of Book and Library Models

One of the final objectives of this work is its application to a large-scale context and in particular with the library of Concordia University. Through this collaboration we will incorporate thousands of books present in the library that are available in both, hard-copy version and the digital online version, into a compact virtual library AR application which can run on any pocket-sized device. For the purposes of the evaluation of the design of this application a generic 3D book model was designed using Blender. As previously mentioned, in the deployed application, each book will have its unique 3D model and texture. To create the generic asset, we used a Blender plugin called ”BookGen” which allows you to create a stack of books on one go. In order to create books of different thickness, we created the stack with just one book and adjusted the dimensions of that asset approximately similar to the dimensions of the actual book. Once we had a texture-less book asset and the original book with us, we scanned the texture of the actual book and mapped it onto our 3D asset. All the sides were perfectly mapped on the exterior of the texture-less 3D book asset that we created, but the middle portion of the book, where we can see the sides of all the pages in

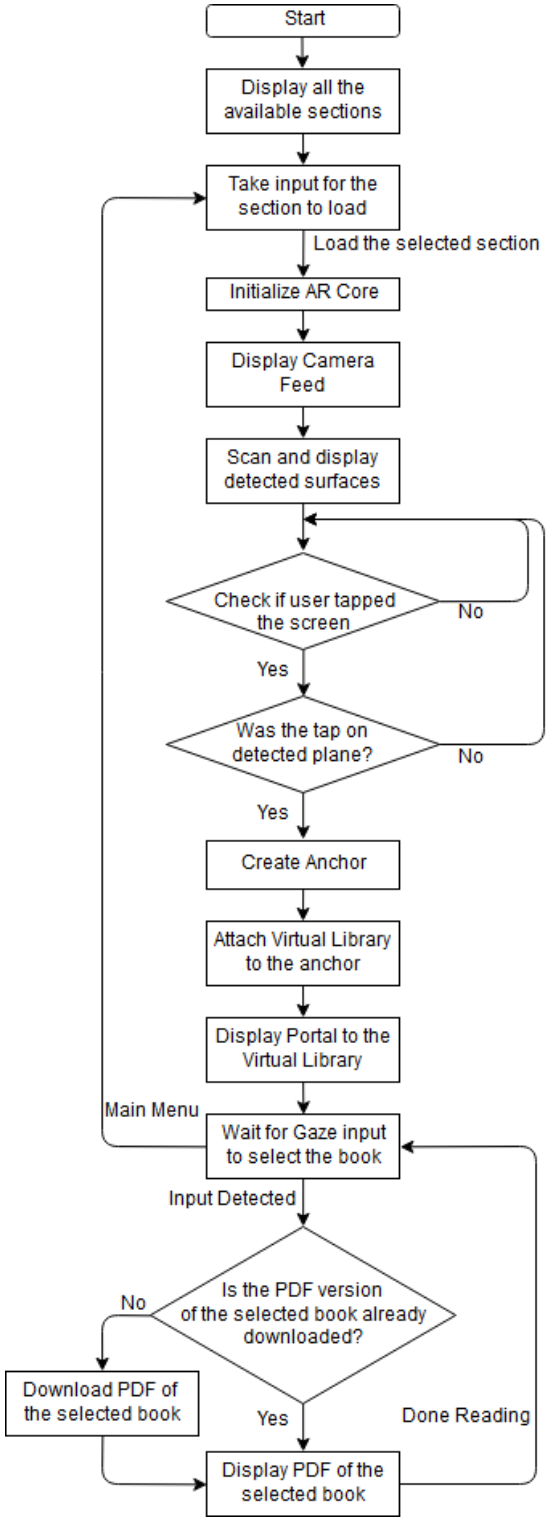

Figure A.4: Flow chart showing overall workflow of the application

the book, was still texture-less. An additional image was added with parallel black and white lines of different thickness to show the density of pages, in the texture mapping to cover that texture-less surface. This made our 3D book asset recognizable, and even more realistic.

The final 3D asset for the book is then imported into the unity project to be used as one book in our virtual library. The scanned texture of the actual book with the additional page density image and the final mapped 3D asset can be seen in Fig. A.5a and Fig. A.5b, respectively.

A.4.4 Creating the Virtual Portal

One of the main goals of the application is to give the users an immersive experience of enjoying the realistic, life-sized, virtual library while roaming in their physical space. To achieve this we employed the AR technology as the users in our application needed to move in their own physical space. The dilemma was that our library was completely virtual i.e. while exploring the virtual library, users will not be able to see anything from the real-world. Therefore, at one point or another in the application, we had to make a transition from being in a completely real-world to a completely virtual world i.e. our virtual library.

As we wanted to make the application as immersive as possible, transition from the real-world to the virtual world had to be smooth and seamless. There were several ways to implement this; one of which was by making the transition from real to virtual world (or the other way around) in a single cut i.e. very short time. Although this is too simple to implement and does not involve any heavy computations, it lacks the sense of continuity during the transition since the abrupt change in all the surrounding environment of the user was not at all smooth and seamless. Instead, another method was employed: that of fading the virtual environment into the scene while fading out the real world background. This was smooth, not perfectly seamless but better than the first method although it too lacked the sense of continuity that we were looking for.

After trying out several other methods, we finally decided to implement and use the Portal Effect for these transitions. In this method, the user would just have to place the portal in their physical space anywhere they want and walk right into the virtual library. When the user places the portal, it will fade in a simple door in his/her physical space through which he/she can simply walk and enter our virtual library. This method was more GPU intensive but the final effect was very smooth and

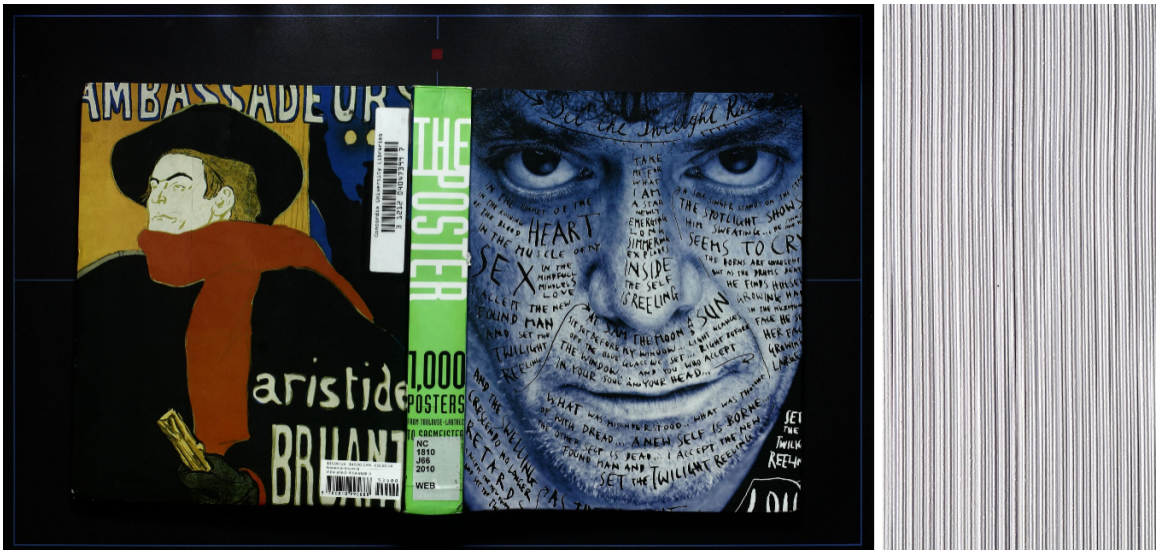

(a) Scanned texture of the actual book

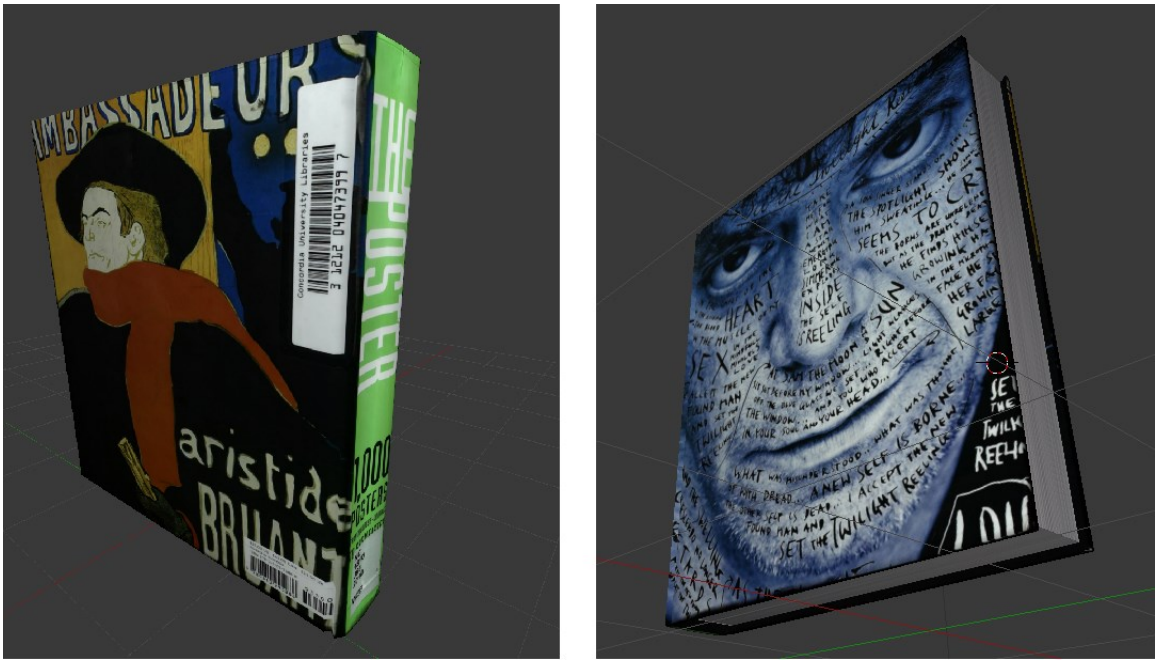

(b) Final 3D asset

Figure A.5: First image (a) shows the scanned texture of the actual book and the second image (b) shows the final 3D asset of the book used in our virtual library

Figure A.6: Portal effect from outside of the portal

seamless, and as we were using the state-of-the-art hardware, it didn't really effect the performance of the application. Also, as the user had to walk into the virtual library through the life-sized portal door, just like any other door in the real world, this method added a sense of realism and continuity to the transition that we were looking for.

To make the illusion of a portal, we had to display only a certain portion of the virtual library. This would be the only portion of the library which is visible from the user's point of view if he/she is looking inside the portal door. This had to be implemented such that the inside of the library is visible if the user is looking through the door but the normal real-world background should be visible around the portal door. Furthermore, the user should only be able to go inside the library if he/she walks through the portal door and not around it. This effect in our application was achieved using Shaders. For the area through which the user can see the inside of the library, we made the mask by a plane covering the whole portal door. A shader was applied to this plane which instructs the GPU to fill the stencil buffer for the area covered by this mask (plane) by the value of 1 instead of drawing on the screen. Given the mask, we applied another shader to all the materials of our

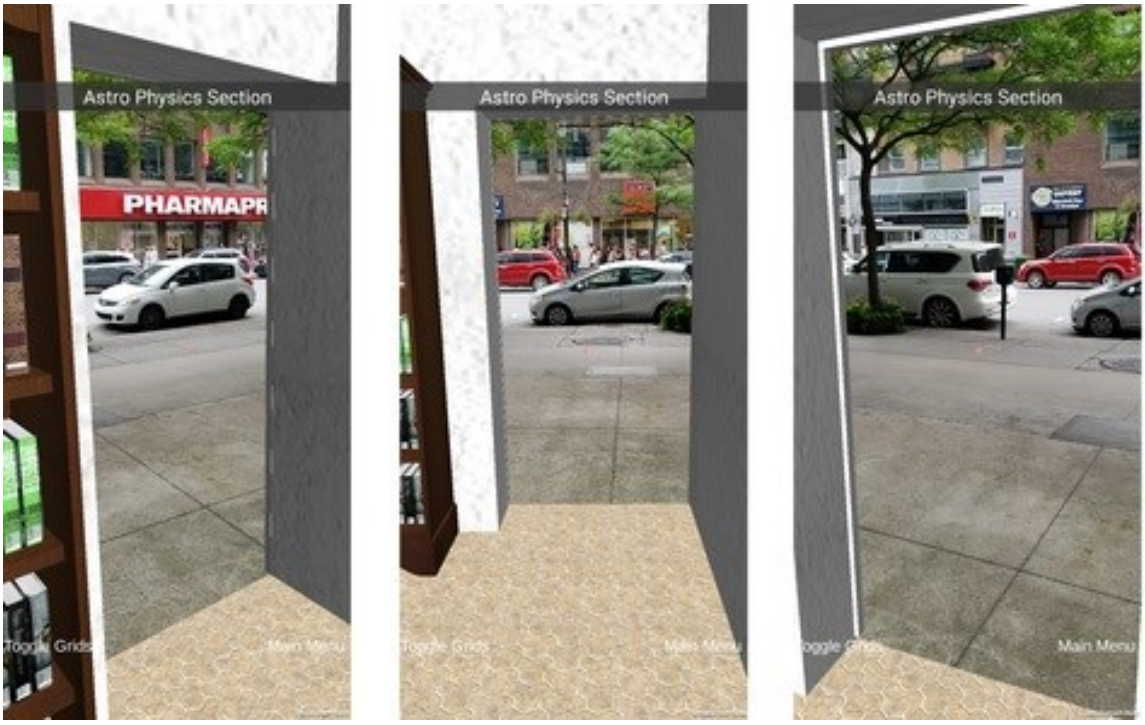

Figure A.7: Portal effect from inside of the portal

library which gave commands to the GPU to only draw the fragments of the library where the value of the stencil buffer is equal to the value 1. Hence, if the user is outside, the application will only draw the portions of the library behind the portal door which are only visible through the mask. Once the user enters the portal, we remove the condition in all the virtual library materials to only draw where the value of stencil buffer is equal to 1. The application will always draw the virtual library irrespective of where the user is looking and when the user comes out of the portal, the condition in all the materials is again invoked to only draw where the value of stencil buffer is equal to 1. Fig. A.6 shows screenshots of the portal effect, through various angles, from the outside of the virtual library while Fig. A.7, shows screenshots of the same portal effect, through various angles, from the inside of the virtual library.

A.5 Methodology

For the evaluation of our application, we conducted a user study with a number of participants. The user study compared the user experiences of three different methods to retrieve an actual book

or its PDF version. As previously mentioned one of our main goals was to give our users a realistic experience exploring the library at their own convenience and in the platform they are most familiar with. Since there is barely any other research literature available prior to ours exploring this concept, or any other similar approach to give users this kind of experience, we compared our technique of retrieving the book with that of a traditional library, and a simple online type and search method. We have divided this section in three subsections, one for each step in our study: (a) pre-test questionnaire, (b) performing book retrieval tasks using different methods, and (c) post-test questionnaire.

A.5.1 Pre-Test Questionnaire

The main focus of the pre-test questionnaire was to derive the basic information about the demographics of the users, prior experience with AR, their current preference to the method of getting any book etc. At the beginning of the trial of the application, for every user, we requested that they fill out this questionnaire and provide us with some of this basic information about them. This information was vital for our study as it gave us a decent understanding of the crowd that we were dealing with, especially their familiarity with the AR technology and how comfortable were they using it.

A.5.2 Performing Book Retrieval Tasks using Different Methods

We asked the participants to complete three simple tasks for this study: (i) downloading a book's PDF version and opening it on a desktop workstation, (ii) finding a book in a traditional library, and (iii) finding a book in the virtual library using our application.

First Task: Desktop method

In the first task, we gave users a standard desktop and a secure internet connection to browse on. We gave them the name of a book and asked them to find that particular book on any of the online libraries, download its PDF version into the hard drive of the system, and finally open that PDF on the desktop. While most of the users succeeded, several users who were not familiar with the online libraries or with using desktop computers, struggled a bit to finish this task. For most

users, this task was too mentally demanding since they had to think a lot while making each and every click decision. As the task seemed too simple, users were trying to minimize the time and complete the task as soon as possible. Most of them succeeded, but the ones who couldn't find the book after looking into several google searches, they got too stressed out which in turn made their performance worse.

This task provided the users with an experience of using a Simple Type and Search method to retrieve the book.

Second Task: Search in Traditional Library

In the second task, we took users to our University's Library and gave them the name of a book. Users were then asked to find the book inside this seven-story library building using nothing more than just the title of the book. They had to use the library's computer-based search systems, and signs on the walls to reach a particular section of the library in which the book was shelved. Once the users reached their desired section they then had to go through all the books in that section and locate one book that was required in those shelves. As the library consisted of seven floors and the required book was placed on the fifth floor, this was a lot more physically as well as mentally intensive task when compared to the first task where the users just had to stay at one place and not move at all.

This task provided the users with an experience of using a traditional library method to retrieve the book.

Third Task: Portal to Knowledge

In the third and final task, users finally got to test the AR application and play around with it. Similar to the first two tasks, the primary job of the user was to find a book based on a visual description in the virtual library and open the PDF of that book on the device. Instructions on how to operate the application and interact with the books inside the virtual library were provided to the users before hand. Almost all of our users found the application to be more productive and convenient.

This third task was performed in an open airy space of around $7x10m^2$ which is equivalent to an

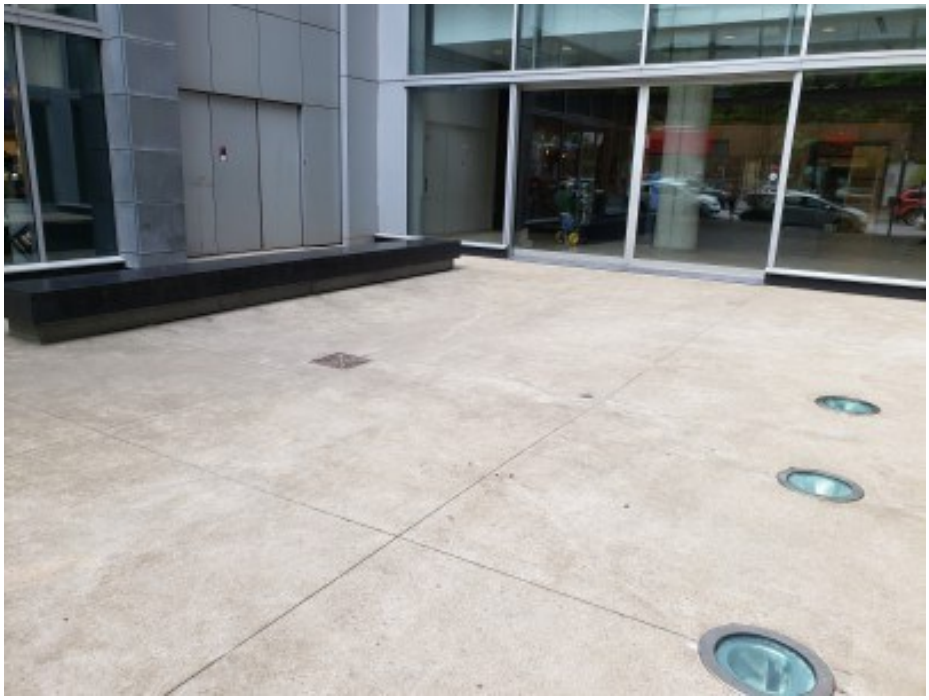

Figure A.8: Play area used for the user study

average large sized room. We used this large space to give our users more space to roam around, but our experiments have shown that one can easily perform the same task using our application with a space as small as 4x4m^2. This task provided the users with an experience of using our application to retrieve the book. Fig. A.8 shows the available physical space (outdoors) that we used to complete this last step of the user study. This is the area utilized by the users to explore the virtual library and accomplish this third task.

A.5.3 Post-Test Questionnaire

In the post-test questionnaire, after the users completed all three tasks mentioned above, we recorded their feedback on locating a book either digitally or physically. In this questionnaire we included a number of questions for understanding the experience of our users and compare them amongst each method of retrieving a book. One of the important tests that we included in this questionnaire was NASA TLX, a scale to measure the overall workload of anything based on 6 different parameters. Using the NASA TLX, based on the mental demand, physical demand, temporal demand, effort needed, performance and frustration level, an overall workload can be calculated for

the three task. On the scale of 1 to 5, we took the feedback for all the above mentioned parameters for all three tasks from the user. This data helped us calculate the overall workload for all three tasks and hence compare it with each other.

Apart from these, the questionnaire took feedback on whether the users preferred (or not) our application over any other method and also if there is anything that can be improved. We also asked users about the convenience offered by the simple structure and portability of the application. The questionnaire took some additional information in the form of open questions about where the users generally faced issues while operating the application. Finally, we asked feedback from all the users about the overall rate that they would give to our application on a scale of 1 to 5.

A.6 Analysis of Evaluation Results

Filling up the post-test questionnaire was an end to the user study for one user. Once we collected all the data from each and every user, we performed a detailed analysis of the evaluation results. In this section, we preset an analysis based on various factors. It is mainly categorized into four different subsections: (a) participants overview, (b) analysis, (c) task load and, (d) user preferences.

A.6.1 Participants Overview

The user study was conducted amongst 45 participants out of which 20 were female and the age groups included are from 14 years to 44 years. The majority (64.4%) were from the age group of 18-34 years. Due to a vast diversity in the group of people that can be the potential users of this application, all of our participants were selected randomly and had nothing to do with their past experiences using AR. The participants were evenly spread among avid traditional and online library users, almost everyone was comfortable with using e-reading material on their mobile devices.

A.6.2 Analysis

ANOVA ($\alpha = 0.05$) with repeated measures was conducted for different methods of locating a book which includes our application, traditional library, and type and search. We had mental

workload, physical workload and the convenience offered by the methods as independent variables for the test. Using Mauchly's test, we confirmed that the Sphericity assumption was preserved during the entire test. We used partial eta squared (η_p^2) to report the obtained Effect sizes.

Mental Demand

A repeated measured analysis of variance (ANOVA) was conducted to evaluate the amount of mental workload required in accomplishing a single task using our application, the traditional library approach and the simple type and search approach with a diverse group of participants ($N = 45$). The results of ANOVA showed a statistically significant effect of using different methods on the mental demand, $Wilks' Lambda = 0.67, F(2, 43) = 10.587, p < 0.05, \eta_p^2 = 0.33$. This gives us a significant evidence of mental workloads being different in each case.

As expected, the mental workload required to accomplish the tasks in the case of our application was the least. Followed by the traditional library method and topped by the simple type and search method. As the task seemed to be really simple when first given to the participants, most of them made a mental picture of how quick they would be in completing it before even they got on the internet. While failing to locate the book after making a few searches, a significant number of participants got frustrated and started to make even more blunders. This compromised their performance and increased their stress levels which in turn increased the amount of mental work done to accomplish the task. Although the difference between type and search, and traditional library is not very evident, the mental workload for our application is significantly less.

Physical Demand

We performed an ANOVA with repeated measurements test to evaluate the amount of physical workload required for the same tasks with a diverse group of participants ($N = 45$). The results of ANOVA showed a statistically significant effect of using different methods on the physical demand, $Wilks' Lambda = 0.15, F(2, 43) = 124.518, p < 0.05, \eta_p^2 = 0.853$. This gives us a significant evidence of physical workloads being different in each case.

Physical workload in case of the traditional library was found to be the highest as compared to any other technique. This can be attributed to the fact that the library used for this task was seven

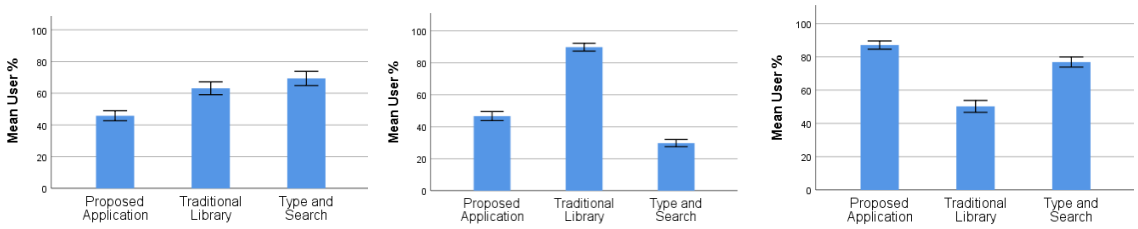

Figure A.9: Left to right: ANOVA results for (a) Mental workload experienced by the participants. (b) Physical workload experienced by the participants. (c) Convenience offered by different methods. Error bars in all three figures is +/- 1*Standard Error.

stories tall and participants had no clue on the where-about of the book. They had to roam a lot inside the library in order to get to the book. While in the case of Type and search, users had almost no physical workload as they just had to sit on a desktop and browse internet. Hence the physical workload observed in the case of type and search method was the least. Since users only had to walk in a small virtual room, in case of our application, the physical workload is not significantly higher than the type and search.

Convenience

Upon reapplying the ANOVA test with repeated measurements on the convenience offered by different methods, it showed us that there is an effect of using different methods on the convenience offered by those respective methods, and it is statistically significant. With 45 people in our survey ($N = 45$), the following results were calculated from ANOVA, Wilks' Lambda = $0.38, F(2, 43) = 34.69, p < 0.05, \eta_p^2 = 0.617$.

Upon taking further detailed feedback, we got some valid reasons that support the above hypothesis. According to majority of the users, as the application was portable and gave them a realistic experience, it was more convenient for them to just drop a portal and get the book instead of going to an actual library. While the rest of the users preferred the experience of an actual library and using the online libraries as they were more accustomed to these socially accepted methods, and also were quite hesitant towards the change. In Fig. A.9c, we can clearly see that most of the users found our proposed application and typical type and search methods way more convenient than the traditional library. Moreover, it can also be seen that the proposed application is almost 10% more convenient than the type and search.

A.6.3 Task Load

Evaluations from NASA TLX showed several differences in the perceptions of various task loads between the three methods. It also showed the overall mean for those task loads. We performed the repeated measures ANOVA upon all the six factors of NASA TLX with the following results:

- Mental Demand: $Wilks' Lambda = 0.67, F(2, 43) = 10.587, p < 0.05, \eta_p^2 = 0.33$
- Physical Demand: $Wilks' Lambda = 0.15, F(2, 43) = 124.518, p < 0.05, \eta_p^2 = 0.853$
- Temporal Demand: $Wilks' Lambda = 0.645, F(2, 43) = 11.850, p < 0.05, \eta_p^2 = 0.355$
- Performance: $Wilks' Lambda = 0.546, F(2, 43) = 17.898, p < 0.05, \eta_p^2 = 0.454$
- Effort: $Wilks' Lambda = 0.414, F(2, 43) = 30.381, p < 0.05, \eta_p^2 = 0.0.586$
- Frustration: $Wilks' Lambda = 0.546, F(2, 43) = 17.898, p < 0.05, \eta_p^2 = 0.454$

Upon examining the task loads of individual methods in Fig. A.10, using a traditional library revealed a much higher task load in almost all the attributes of NASA TLX (except mental) than any of the other two methods. It is indeed peak to peak with the type and search method in the case of Mental demand. On comparing it with the proposed application, the overall work load for the application has revealed itself to be much lower (25.5%) than that of the traditional library.

After the first look at the tasks, one would expect for the overall workload of type and search method to be lower than the proposed application as there was no movement involved and also comparatively the task was really simple. However, it came out to be much higher than what we expected. The overall workload for type and search method is 47.6%, which makes it 10.3% higher than that of the proposed application. Hence, it makes the overall workload for our proposed application the least as compared to any of the other two methods.

A.6.4 User Preferences

Fig. A.11 shows that according to most of our users (75.6%), the proposed application is much better for the navigation inside the library as compared to the traditional library. Due to the smaller

Figure A.10: Average ratings out of hundred for six different workloads of NASA TLX questionnaire and the overall workload for each of the three methods

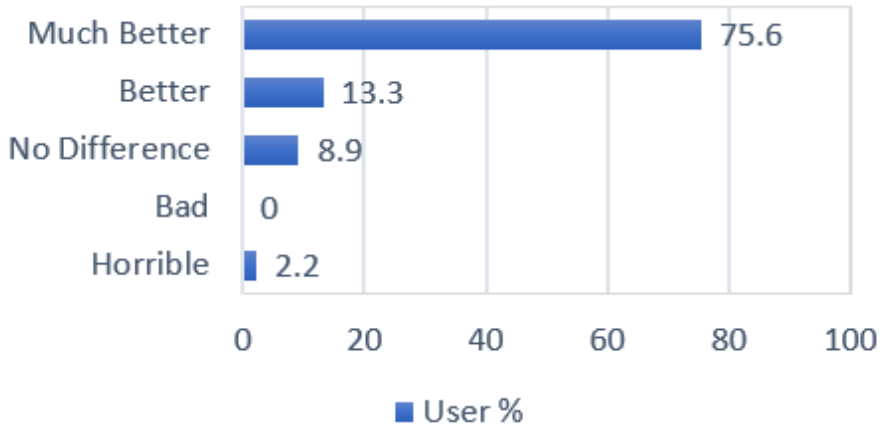

Figure A.11: Navigation in the proposed application as compared to the Traditional Library

walking space and interchangeable sections, the navigation in the virtual library is extremely easy e.g. (p23: "As there was only small space to search for the book, it was really easy for me to find that book. And also I was able to switch between the sections quickly"). Only a couple of shelves in one virtual section gives the application it's easier and faster navigation (p16: "Easy and faster navigation"). However, some participants still found it difficult to locate the book suggesting, e.g. p6: "Need direction arrow for finding book location". Participants also reported that the multiple interchangeable sections made it highly likeable and helpful in the case of large libraries e.g. (p2: "Its easy to use and can be very helpful for very large libraries").

Although the majority of our users did not encounter any problems while operating the application, some faced the following common issues during the survey: difficulty scanning the surfaces (reported by 17.8% of users), difficulty keeping the AR images in view (reported by 17.8% of users) and most importantly, difficulty using gaze as an input (reported by 22.2% of users).

Due to the portability of our proposed application e.g. (p7: ”Can be opened anywhere”, p8: ”Library in my pocket”), better navigation e.g. (p15: ”Not much walking but I get the same experience”, p38: ”Big time saver with similar experience”) and least overall workload e.g. (p41: ”Similar experience with less effort”), the majority of the participants preferred the proposed application over the actual traditional library. The user distribution based on their preference of the proposed application over the traditional library can be seen in Fig. A.12.

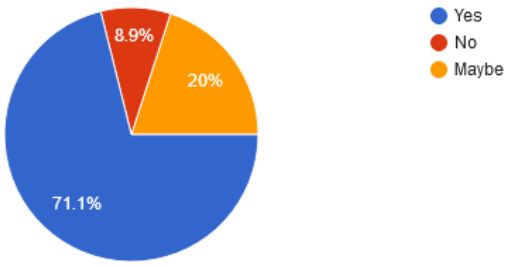

Figure A.12: User preference of the proposed application over the traditional library

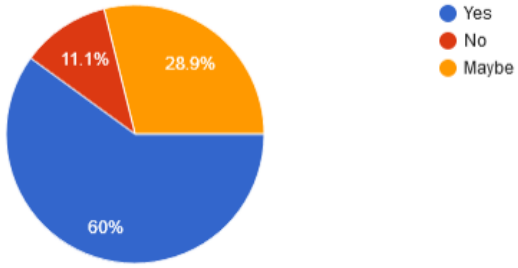

Figure A.13: User preference of the proposed application over the type and search

Despite the fact that the type and search is a less physically demanding method, the majority of our users preferred to have an experience of actually being in the library e.g. (p15: ”I prefer the traditional library experience”). Additionally, the interactive aspect e.g. (p37: ”Yes, it is interactive, easy and something different. I would love to see more of it”) of the application appealed the users and preferred the proposed application over the type and search method. At the same time, participants not familiar with using online libraries also preferred the proposed application over type and search. In Fig. A.13, we can see the distribution of users based on their preferences.

Finally, in Fig. A.14, we show the distribution of users based on their overall rating responses

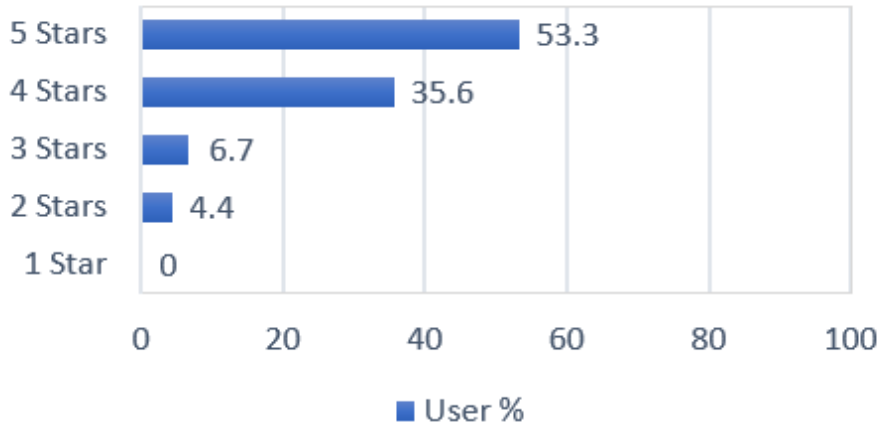

Figure A.14: Overall user rating for the proposed application

on a scale of one to five stars; where one being the least favourite and five being the most.

A.7 Conclusion

To sum it all up, in this chapter, we proposed a novel application to provide users with an experience of visiting an entire library while actually being in a tiny play space using Marker-less AR. We also investigated various other methods for the same task to draw a comparison between the proposed method and two other socially accepted methods. A user study was conducted with 45 participants which revealed various flaws in existing techniques and showed how the proposed application can bring positive advancements in the ways we explore libraries. NASA TLX showed that using a traditional library had the highest workload. Upon further evaluation, we found that the proposed application had the least overall workload and was the most preferred method by the participants. The survey also further established that the navigation in the proposed application is much better than that in the traditional library. Traditional libraries use cumbersome taxonomy system which is implicitly used for the navigation. The proposed application uses smaller and contained walking area making the physical workload experienced by the users far less, hence more than 71% of the participants opted to definitely not prefer the traditional library and switch to using the proposed application in the future. In the case of the type and search, as our application provided a realistic library experience at their fingertips, 60% of the participants preferred the application.

Overall, the participants gave our application a combined rating of 4.74 out of 5.

For future works, this application will be imported to an immersive AR platform e.g. Microsoft's Hololens. Various additional interaction techniques will be explored in order to improve the interaction. Currently, this application consists of 7800 books, however we are in the process of expanding the database to our University's library's database which houses thousands of books.

Appendix B

Patent:

Managing real world and virtual motion

US011557105B2

(12) **United States Patent**
Joshi et al.

(10) **Patent No.: US 11,557,105 B2**
(45) **Date of Patent: Jan. 17, 2023**

(54) **MANAGING REAL WORLD AND VIRTUAL MOTION**

(71) Applicant: **VALORBEC SOCIETE EN COMMANDITE**, Quebec (CA)

(72) Inventors: **Yashas Joshi**, Montreal (CA); **Charalambos Poullis**, Beaconsfield (CA)

(73) Assignee: **Concordia University**, Montreal (CA)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **17/467,844**

(22) Filed: **Sep. 7, 2021**

(65) **Prior Publication Data**

US 2022/0084301 A1 Mar. 17, 2022

Related U.S. Application Data

(60) Provisional application No. 63/079,089, filed on Sep. 16, 2020.

(51) **Int. Cl.**

G06T 19/20	(2011.01)
G06T 19/00	(2011.01)
G06T 15/20	(2011.01)
G06F 3/01	(2006.01)

(52) **U.S. Cl.**

CPC ***G06T 19/20*** (2013.01); ***G06F 3/013*** (2013.01); ***G06T 15/20*** (2013.01); ***G06T 19/006*** (2013.01)

(58) **Field of Classification Search**

CPC G06T 19/20; G06T 15/20; G06T 19/006; G06F 3/013; G06F 2203/012; G06F 3/011

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

2018/0165881	A1*	6/2018	Pohl	G06F 3/013
2019/0012824	A1*	1/2019	Sun	G06F 3/011
2021/0373657	A1*	12/2021	Connor	G02B 27/017

OTHER PUBLICATIONS

Bolte et al. "Subliminal reorientation and repositioning in immersive virtual environments using saccadic suppression." IEEE Transactions of Visualization and Computer Graphics, 21(4):545-552, 2015.

Langbehn et al. "Subliminal re-orientation and re-positioning in virtual reality during eye blinks." SUI '16, pp. 213-213, 2016.

(Continued)

Primary Examiner — Chong Wu

(74) *Attorney, Agent, or Firm* — Rosenberg, Klein & Lee

(57) **ABSTRACT**

Navigation of a virtual environment (VE) can mean navigating a VE that is spatially larger than the available Physical Tracked Space (PTS). Accordingly, the concept of redirected walking was introduced in order to provide a more natural way of navigating a VE, albeit with many restrictions on the shape and size of the physical and virtual spaces. However, prior art techniques have limitations such as negatively impacting the sense of immersion of the user, motion sickness, or forcing the user to look away by stimulating major saccades. Accordingly, the inventors have established a novel technique which overcomes these limitations. The technique is based on the psychological phenomenon of inattentional blindness allowing for re-directed walking without requiring the triggering major saccades in the users, complex expensive systems, etc.

14 Claims, 11 Drawing Sheets

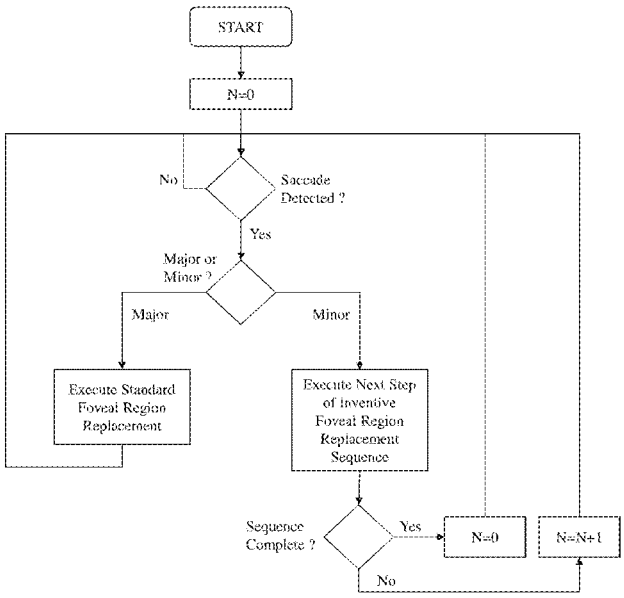

(56) **References Cited**

OTHER PUBLICATIONS

Langbehn et al. "In the blink of an eye—leveraging blink induced suppression for imperceptible position and orientation redirectionin virtual reality." ACM Transactions on Graphics, 37:1-11, 2018.

Patney et al. "Perceptually-based foveated virtual reality." ACM SIGGRAPH 2016 Emerging Technologies, SIGGRAPH '16, p. 17:1-17:2. ACM, New York, NY, USA, 2016.

Razzaque et al. "Redirected walking." Proceedings of Eurographics, 9:105-106, 2001.

Stengel et al. "Adaptive Image-Space Sampling for Gaze-Contingent Real-time Rendering." Computer Graphics Forum, 2016. doi: 10.1111/cgf.12956.

Suma et al. "Leveraging change blindness for redirection in virtual environments." 2011 IEEE Virtual Reality Conference, pp. 159-166. IEEE, 2011.

Sun et al. "Towards virtual reality infinite walking: dynamic saccadic redirection." ACM Transactions on Graphics (TOG), 37(4):67, 2018.

* cited by examiner

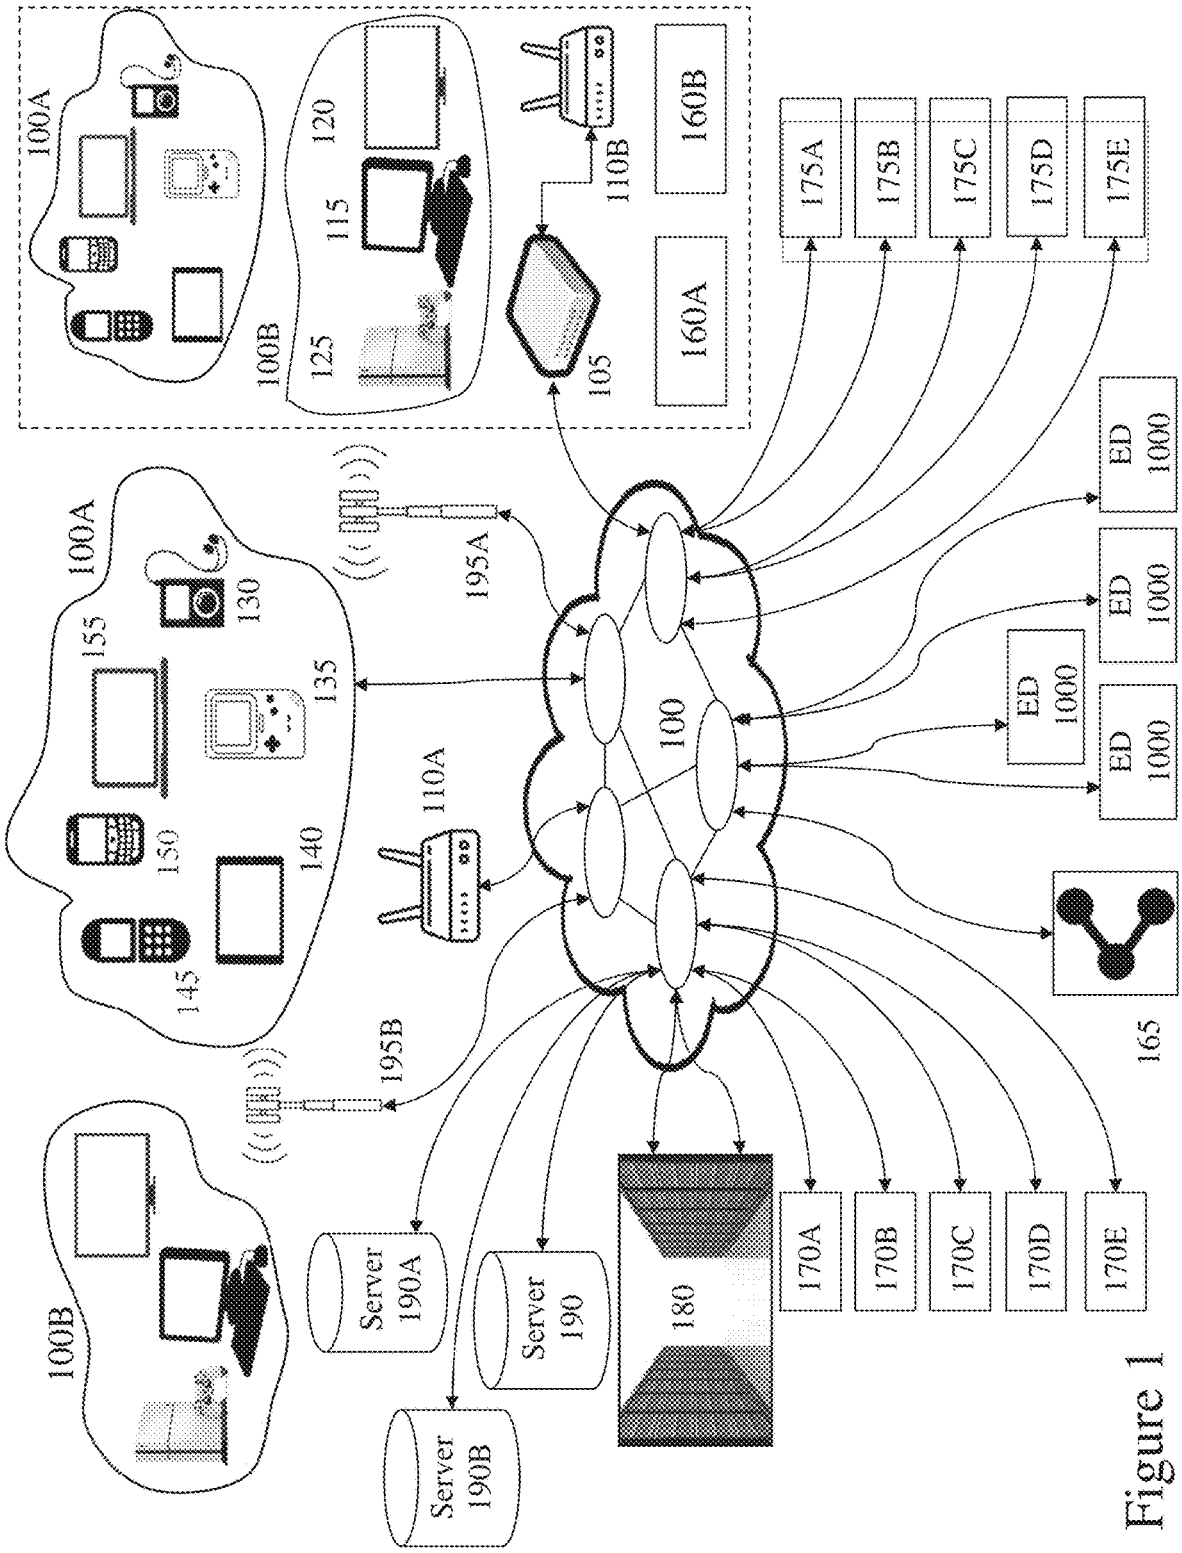
100A
100B
120
115
125
105
110B
160B
160A
175A
175B
175C
175D
175E
ED 1000
ED 1000
ED 1000
ED 1000
100
195A
195B
110A
155
130
135
150
140
145
Server 190A
Server 190
Server 190B
180
165
170A
170B
170C
170D
170E

Figure 1

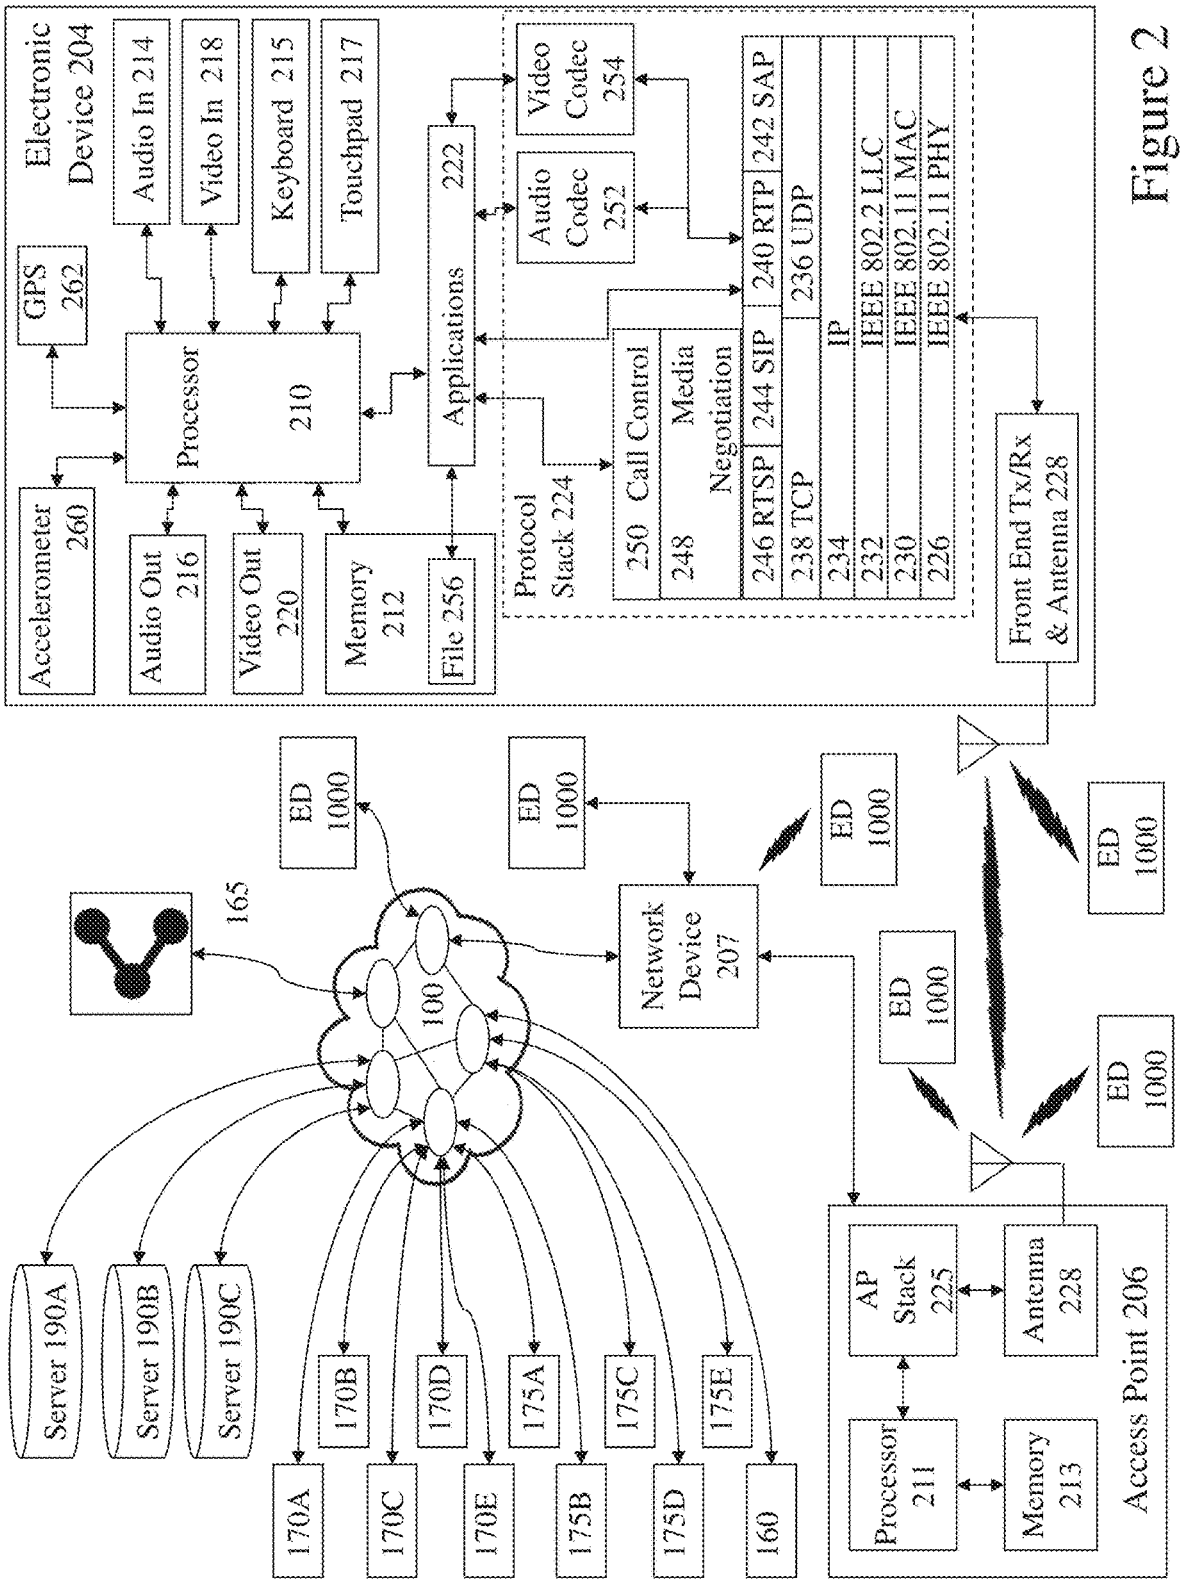
Electronic Device 204
Accelerometer 260
GPS 262
Audio In 214
Video In 218
Keyboard 215
Touchpad 217
Processor 210
Audio Out 216
Video Out 220
Memory 212
File 256
Applications 222
Protocol Stack 224
Audio Codec 252
Video Codec 254
250 Call Control
248 Media Negotiation
246 RTSP
244 SIP
240 RTP
242 SAP
238 TCP
236 UDP
234 IP
232 IEEE 802.2 LLC
230 IEEE 802.11 MAC
226 IEEE 802.11 PHY
Front End Tx/Rx & Antenna 228
ED 1000
165
100
Network Device 207
Server 190A
Server 190B
Server 190C
170A
170B
170C
170D
170E
175A
175B
175C
175D
175E
160
Access Point 206
Processor 211
Memory 213
AP Stack 225
Antenna 228

Figure 2

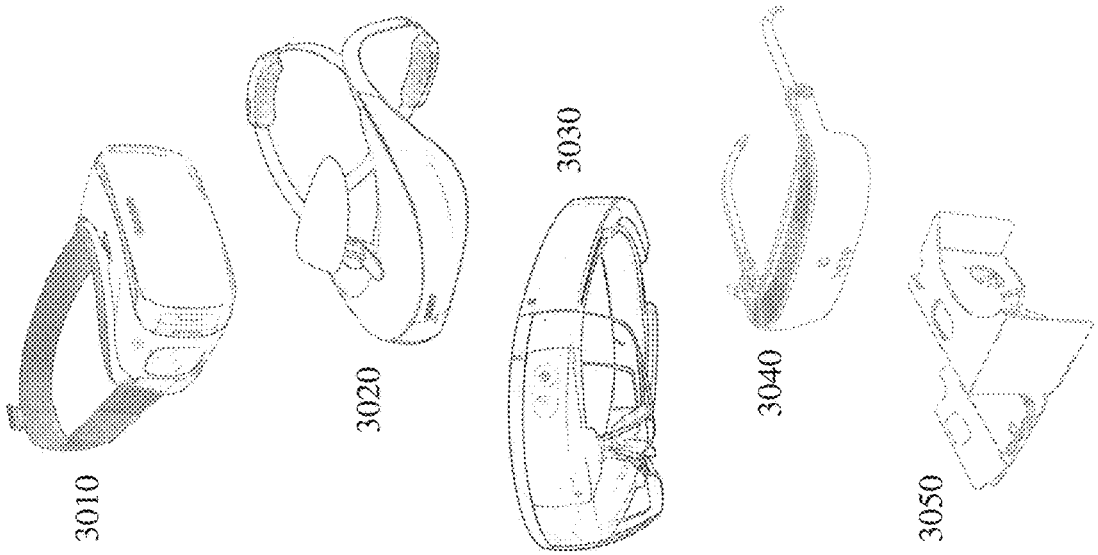

Figure 3B

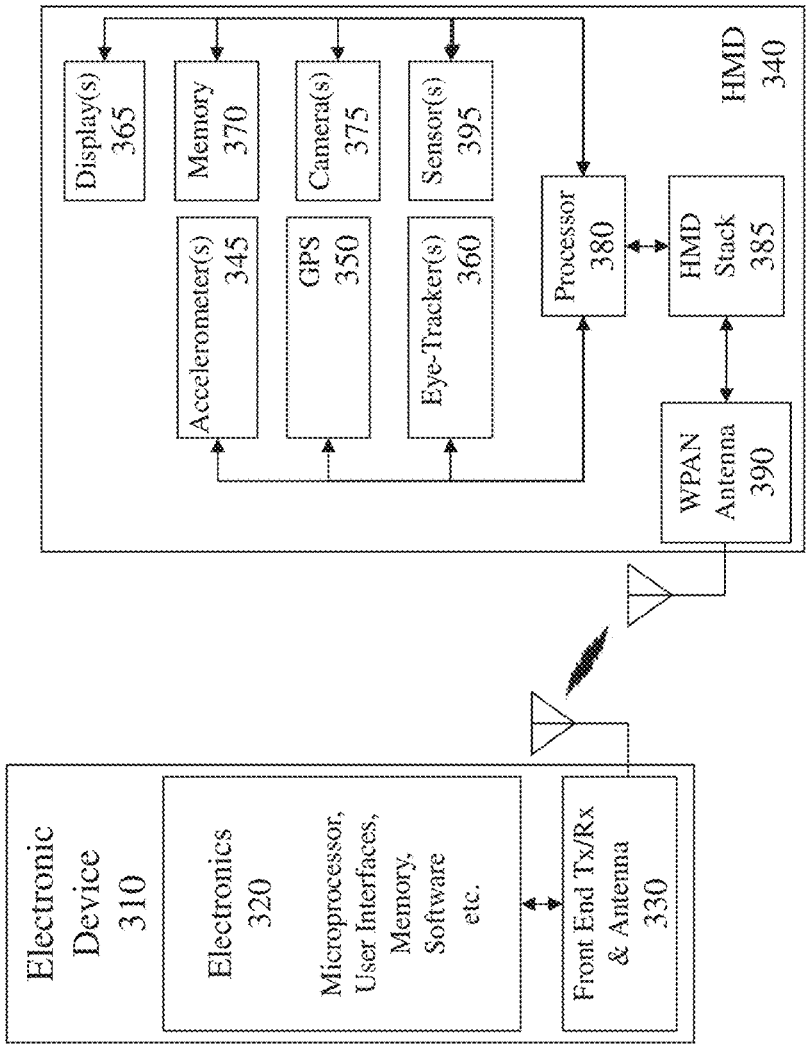

Figure 3A

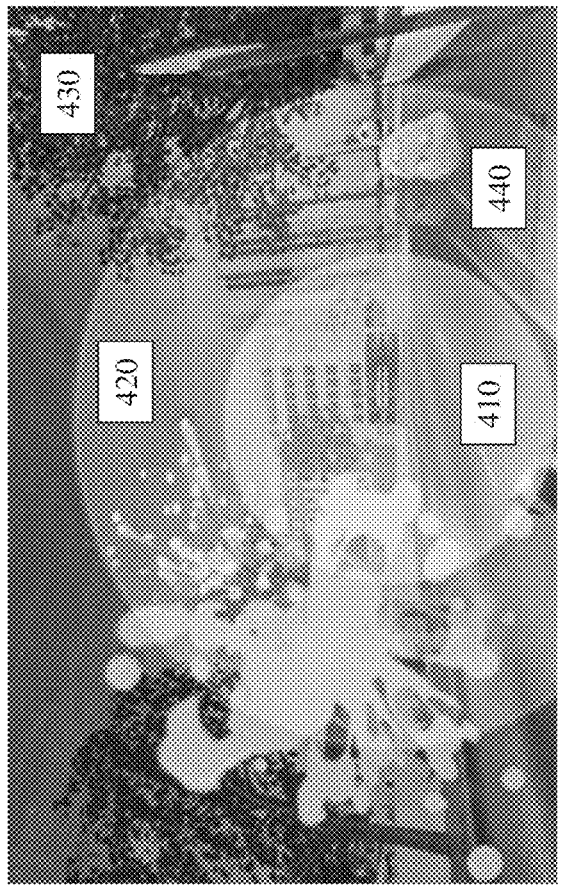

Figure 4A

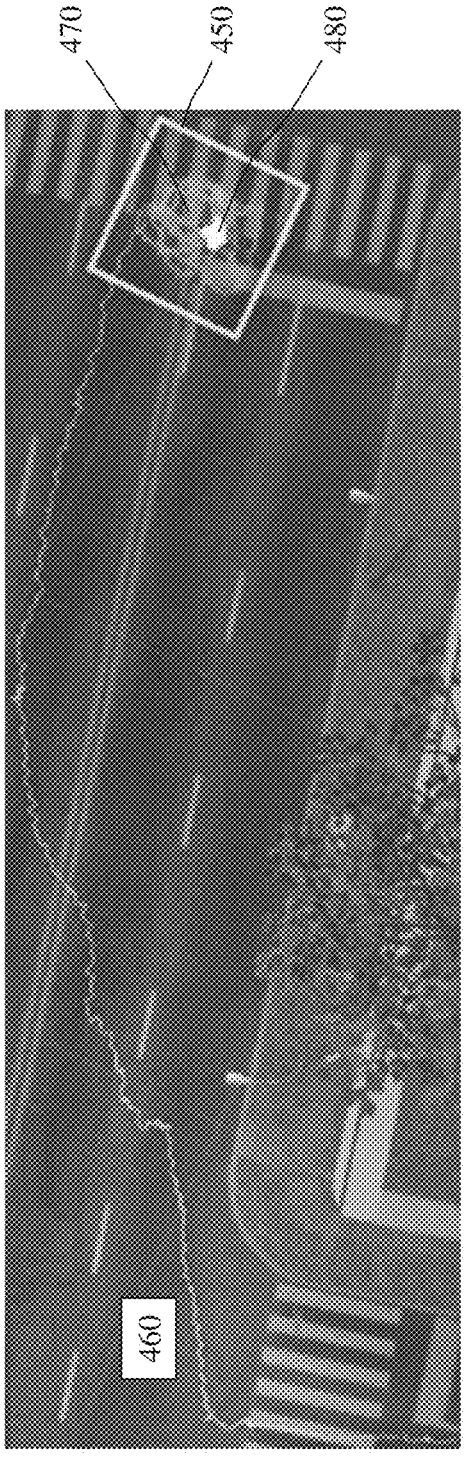

Figure 4B

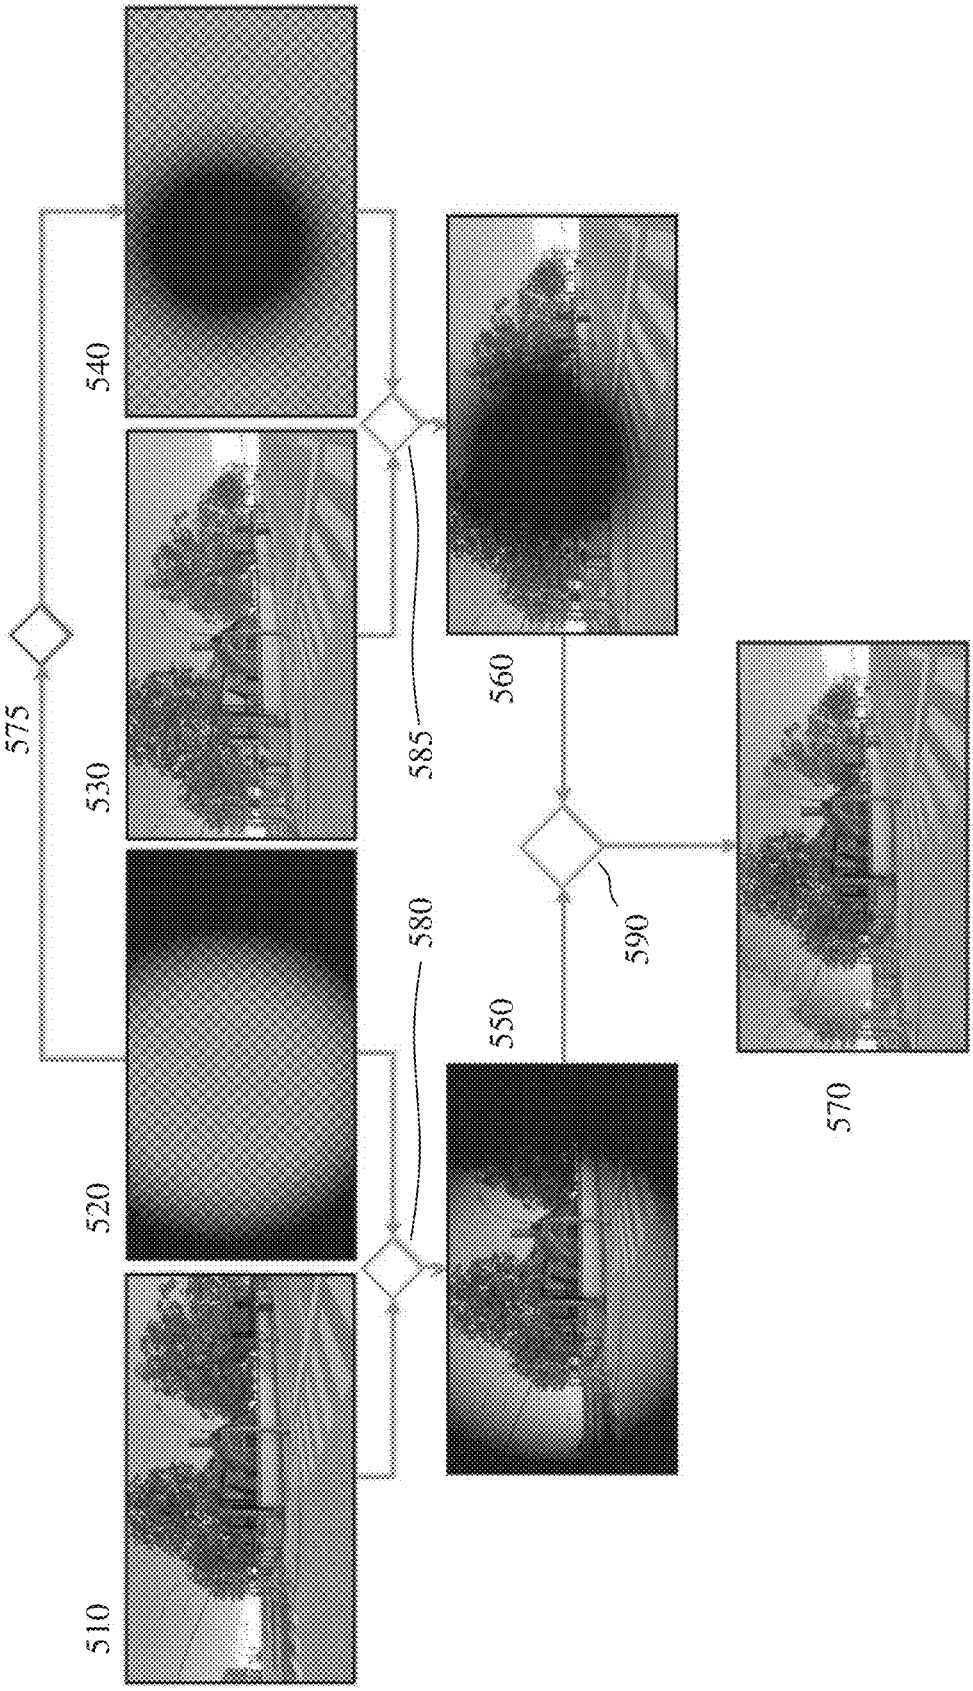

Figure 5

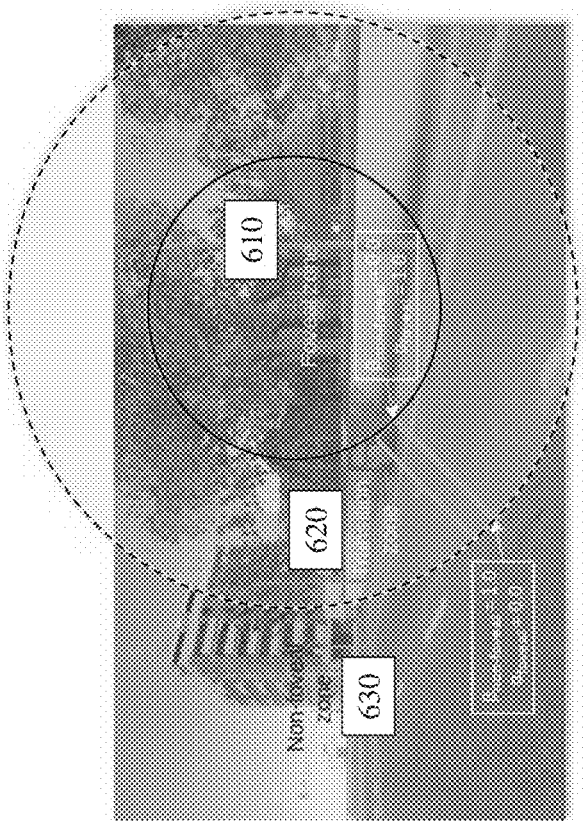

Figure 6

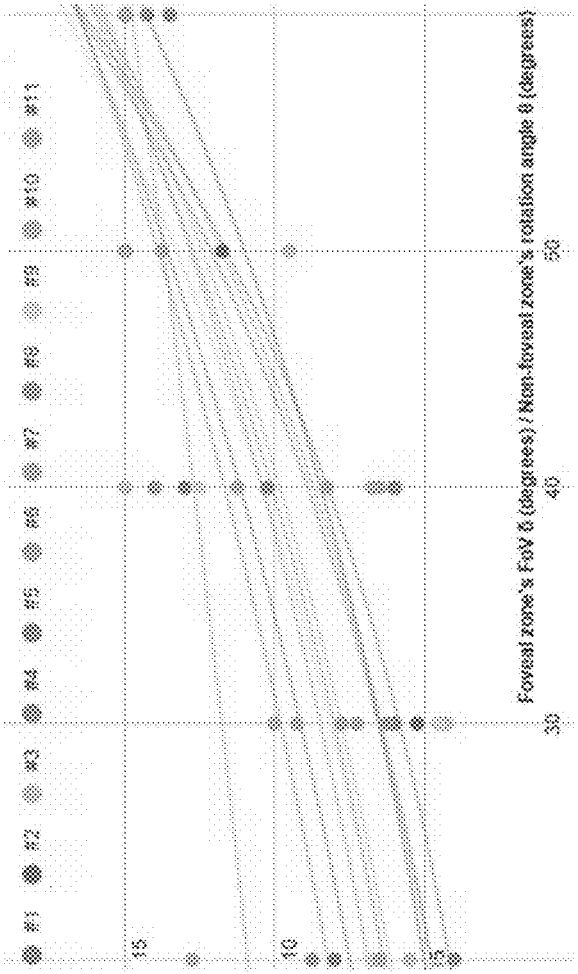

Figure 7

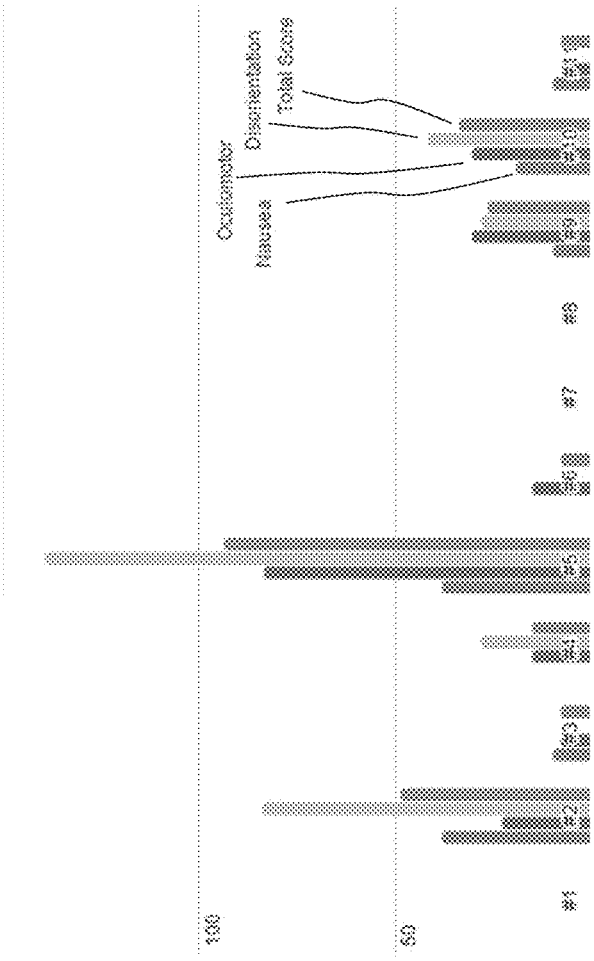

Figure 8

Figure 9

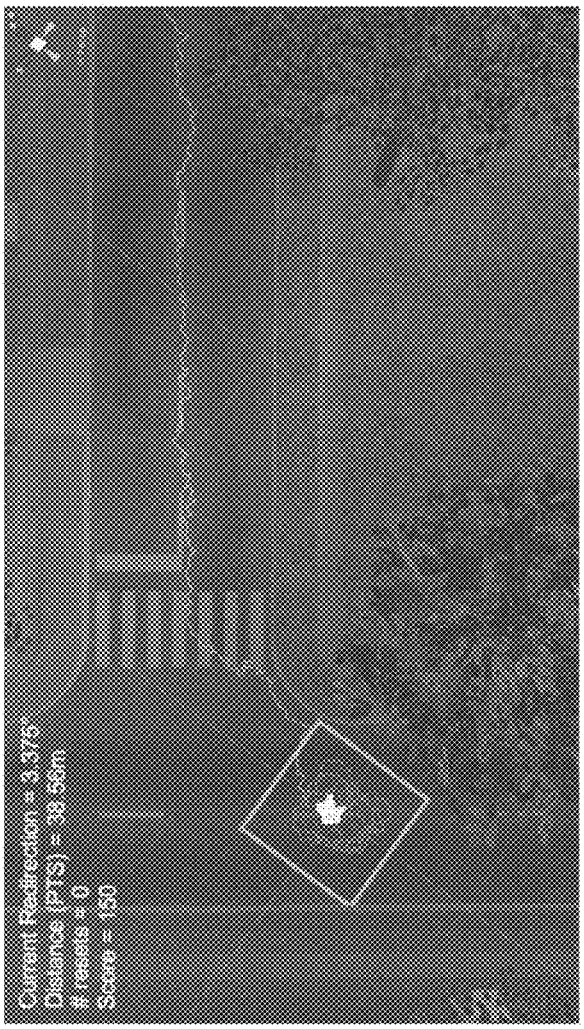

Figure 10A

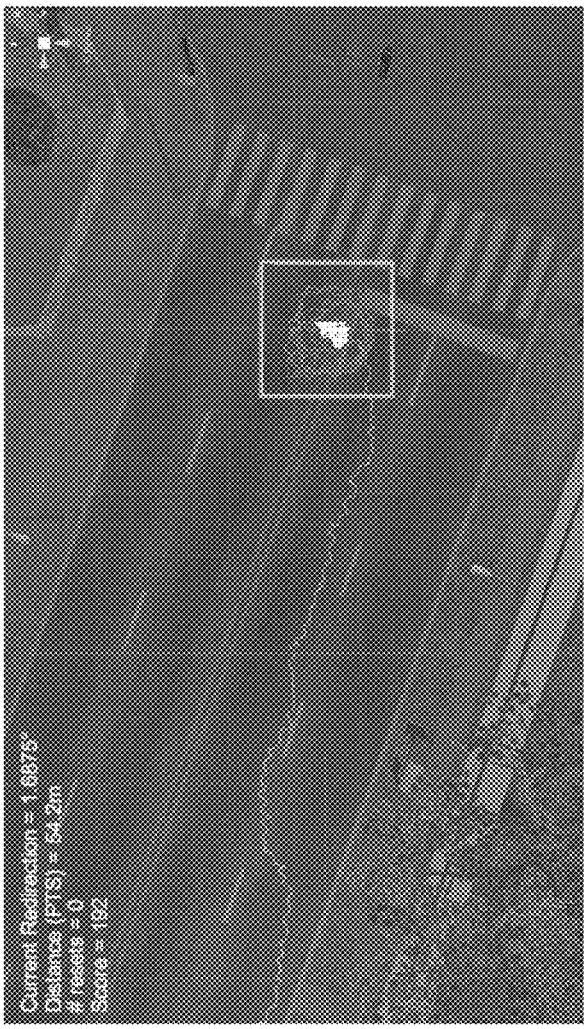

Figure 10B

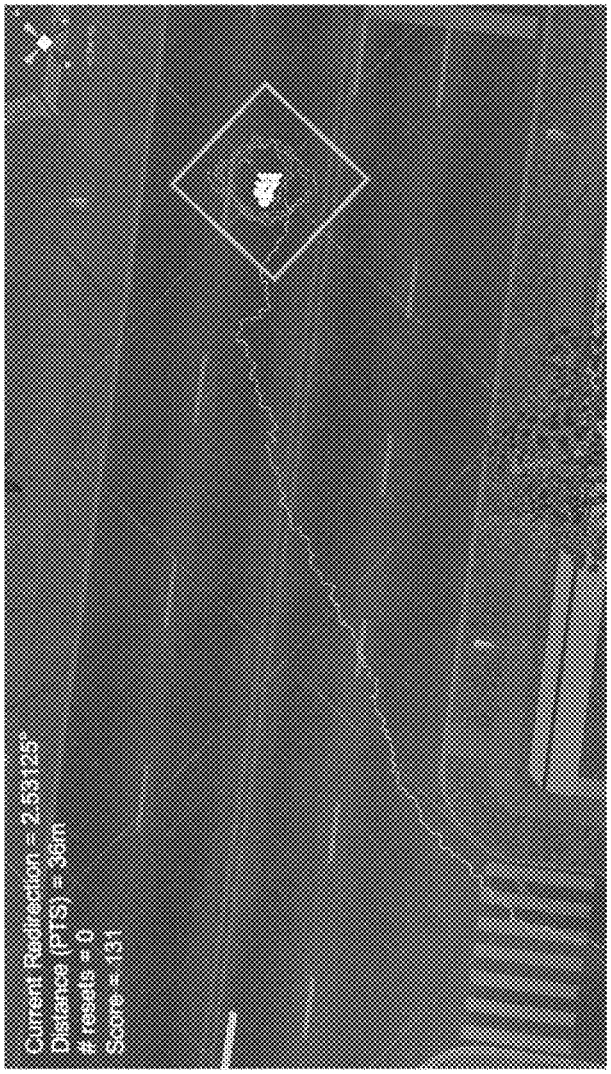

Figure 10C

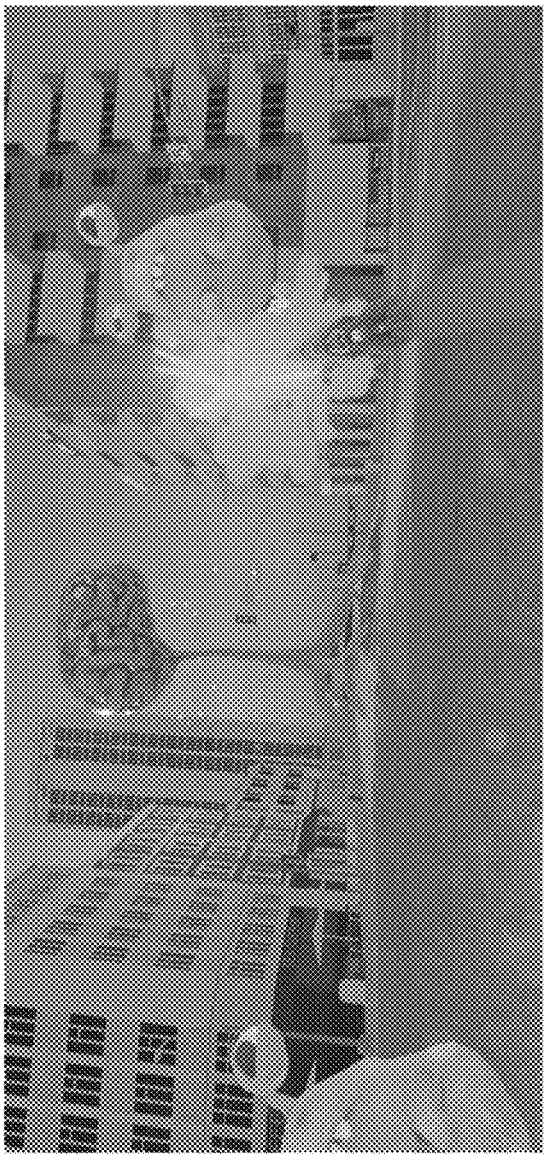

Figure 11

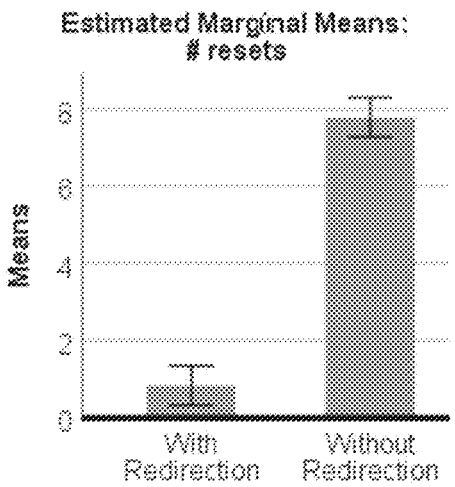

Figure 12A

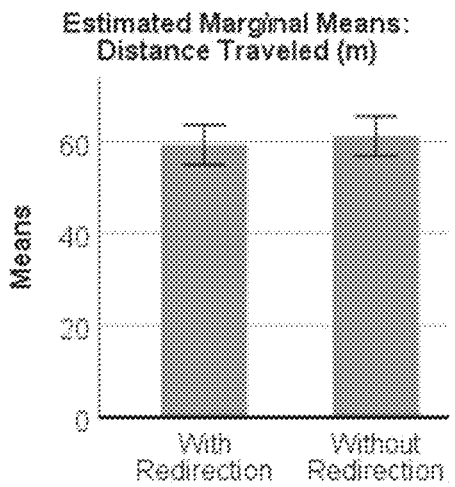

Figure 12B

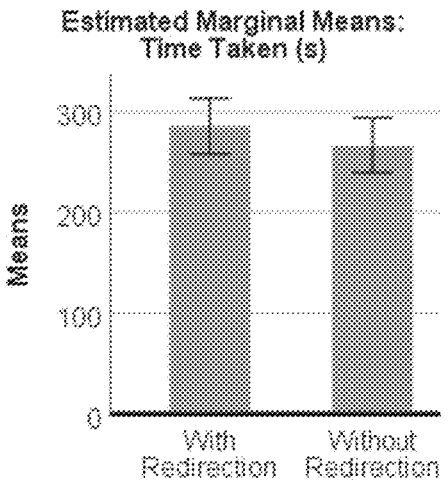

Figure 12C

Figure 13

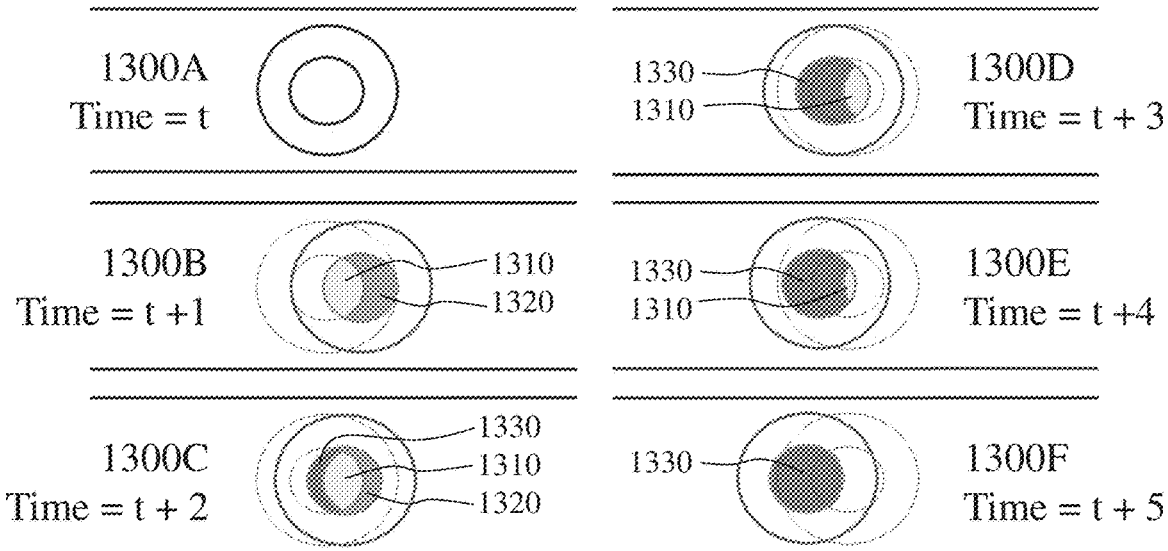

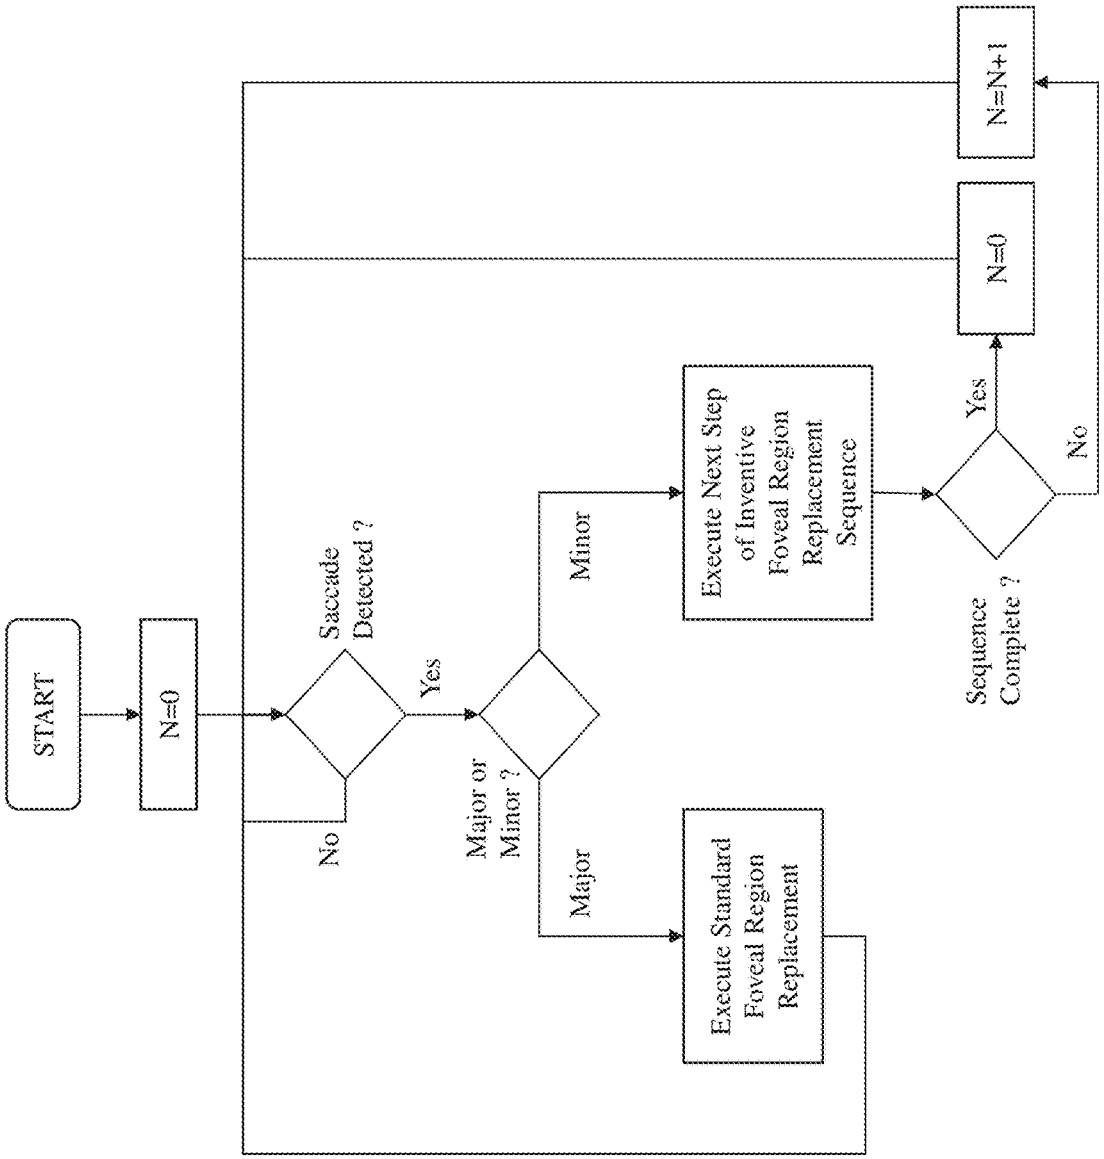
START
N=0
Saccade Detected ?
No
Yes
Major or Minor ?
Major
Minor
Execute Standard Foveal Region Replacement
Execute Next Step of Inventive Foveal Region Replacement Sequence
Sequence Complete ?
Yes
No
N=0
N=N+1

Figure 14

MANAGING REAL WORLD AND VIRTUAL MOTION

CROSS-REFERENCE TO RELATED APPLICATIONS

This patent application claims the benefit of priority from U.S. Provisional Patent Application 63/079,089 filed Sep. 16, 2020; the entire contents of which are incorporated herein by reference.

FIELD OF THE INVENTION

This patent application relates to virtual reality and more particularly to locomotion within virtual reality environments where redirected walking allows a user to navigate a virtual environment spatially larger than the available physical space.

BACKGROUND OF THE INVENTION

Since the early days of virtual reality researchers have investigated ways of users navigating virtual environments (VEs) that are spatially larger than the available Physical Tracked Space (PTS). A number of locomotion techniques relying on pointing devices or walking in-place were proposed which have since become customary in VE applications. However, users find these methods cumbersome and unnatural. The concept of redirected walking was introduced about 20 years ago in order to provide a more natural way of navigating VEs, albeit with many restrictions on the shape and size of the physical and virtual spaces.

A number of approaches have since been proposed for implementing redirected walking based upon hardware or software techniques. Hardware-based techniques such as omni-directional treadmills, the VirtuSphere, etc. are not only expensive solutions to this problem but also fail to provide inertial force feedback equivalent to natural walking. In contrast, software-based techniques are more cost effective and typically involve applying perceptually subtle rotations to the VE causing the user to unknowingly change their walking direction. Applying these rotations to the VE, however subtle, can negatively impact the sense of immersion of the user. This arises as these techniques either employ warping which introduces visual artifacts and distortions in the VE or even simulation sickness, or rely on forcing the user to look away by stimulating major saccades in order to update the environment during the subsequent rapid eye movement resulting from the stimulated major saccades.

Accordingly, it would be beneficial to provide a method of redirected walking for users exploiting virtual reality (VR) or a VE allowing virtual distances and spaces to be traversed which are larger than the available physical space to the user. The inventors have therefore established a novel technique based on the psychological phenomenon of inattentional blindness without requiring the triggering major saccades in the users, complex expensive systems, etc.

Other aspects and features of the present invention will become apparent to those ordinarily skilled in the art upon review of the following description of specific embodiments of the invention in conjunction with the accompanying figures.

SUMMARY OF THE INVENTION

It is an object of the present invention to mitigate limitations within the prior art relating to virtual reality and more particularly to locomotion within virtual reality environments where redirected walking allows a user to navigate a virtual environment spatially larger than the available physical space.

In accordance with an embodiment of the invention there is provided a method comprising rendering to the user a sequence of images relating wherein each image comprises:

a foveal region representing a predetermined portion of a first field of view of the user at a current location of the user with a first direction having a first predetermined geometry and second predetermined dimensions;

a non-foveal region representing a predetermined portion of a second field of view of the user at the current location with a second direction having a second predetermined geometry and second predetermined dimensions; and

a transition region having a predetermined third predetermined geometry and third predetermined dimensions; wherein

in a first image of the sequence of images the first direction and second direction are the same;

in a second image of the sequence of images the second direction is offset relative to the first direction by a predetermined angular offset;

in a third image of the sequence of images the first direction is set to the second direction in the second image of the sequence of images such that the first direction and second direction are the same but are now offset from the original first direction in the first image of the sequence of images by the predetermined angular offset;

the transition region blends from the foveal region to the non-foveal region according to a predetermined function; and

transitioning from the first image of the sequence of images to the second image of the sequence of images and the second image of the sequence of images to the third image of the sequence of images is established in dependence upon determining a predetermined natural event with respect to an eye of the user.

In accordance with an embodiment of the invention there is provided a system comprising:

a display for rendering content to a user;

an eye tracker for tracking movement or motion with respect to an eye of the user;

a microprocessor for generating the content to be rendered by the display;

a non-volatile, non-transitory storage medium comprising executable instructions for execution by the microprocessor, where the executable instructions when executed by the microprocessor configure the microprocessor to perform a process comprising the steps of: receive data from the eye tracker;

determine a predetermined natural event with respect to the eye of the user in dependence upon the data received; and

render to the user a next image of a sequence of images; wherein

each image of the sequence of images comprises:

a foveal region representing a predetermined portion of a first field of view of the user at a current location of the user with a first direction having a first predetermined geometry and second predetermined dimensions;

a non-foveal region representing a predetermined portion of a second field of view of the user at the current location with a second direction having a second predetermined geometry and second predetermined dimensions; and

a transition region having a predetermined third predetermined geometry and third predetermined dimensions; wherein

in a first image of the sequence of images the first direction and second direction are the same;

in a second image of the sequence of images the second direction is offset relative to the first direction by a predetermined angular offset;

in a third image of the sequence of images the first direction is set to the second direction in the second image of the sequence of images such that the first direction and second direction are the same but are now offset from the original first direction in the first image of the sequence of images by the predetermined angular offset; and

the transition region blends from the foveal region to the non-foveal region according to a predetermined function.

In accordance with an embodiment of the invention there is provided a method comprising: rendering to the user a sequence of images relating wherein each image comprises:

a foveal region having a first predetermined geometry and second predetermined dimensions;

a non-foveal region having a second predetermined geometry and second predetermined dimensions; and

a transition region having a predetermined third predetermined geometry and third predetermined dimensions; wherein

in a first image of the sequence of images the foveal region is a first predetermined portion of a first field of view of the user at a current location of the user with a first direction;

in a last image of the sequence of images the foveal region is a first predetermined portion of a second field of view of the user at a current location of the user with a second direction;

the non-foveal region in each image of the sequence of images is a predetermined portion of a second field of view of the user at the current location with the second direction:

the second direction is offset from the first direction by a predetermined amount;

in each subsequent image of the sequence of images between the first image of the sequence of images and the second image of the sequence of images the foveal region comprises:

a first predetermined foveal portion comprising a first predetermined portion of the first field of view of the user at the current location of the user with the first direction; and

a second predetermined foveal portion comprising a first predetermined portion of the second field of view of the user at the current location of the user with the second direction.

In accordance with an embodiment of the invention there is provided a system comprising: a display for rendering content to a user;

an eye tracker for tracking movement or motion with respect to an eye of the user;

a microprocessor for generating the content to be rendered by the display;

a non-volatile, non-transitory storage medium comprising executable instructions for execution by the microprocessor, where the executable instructions when executed by the microprocessor configure the microprocessor to perform a process comprising the steps of: receive data from the eye tracker;

determine a minor saccade with respect to the eye of the user in dependence upon the data received; and

render to the user a next image of a sequence of images; wherein

each image of the sequence of images comprises:

a foveal region having a first predetermined geometry and second predetermined dimensions;

a non-foveal region having a second predetermined geometry and second predetermined dimensions; and

a transition region having a predetermined third predetermined geometry and third predetermined dimensions;

in a first image of the sequence of images the foveal region is a first predetermined portion of a first field of view of the user at a current location of the user with a first direction:

in a last image of the sequence of images the foveal region is a first predetermined portion of a second field of view of the user at a current location of the user with a second direction;

the non-foveal region in each image of the sequence of images is a predetermined portion of a second field of view of the user at the current location with the second direction;

the second direction is offset from the first direction by a predetermined amount;

in each subsequent image of the sequence of images between the first image of the sequence of images and the second image of the sequence of images the foveal region comprises:

a first predetermined foveal portion comprising a first predetermined portion of the first field of view of the user at the current location of the user with the first direction; and

a second predetermined foveal portion comprising a first predetermined portion of the second field of view of the user at the current location of the user with the second direction.

In accordance with an embodiment of the invention there is provided a method redirecting motion of a user comprising:

a) determining a requirement to redirect the motion of the user from a current direction to a target direction;

b) determining an occurrence of an event with respect to an eye of a user;

c) determining whether the saccade is a minor saccade, a major saccade, or a blink;

d) upon determination of a minor saccade executing a next step of a first process which redirects the motion of the user from the current direction to the target direction in a number of steps;

e) upon determination of a major saccade or a blink executing a second process which redirects the motion of the user from the current direction to the target direction in a single step;

f) repeating steps (b) to (f) until motion of the user has been redirected to the target direction.

In accordance with an embodiment of the invention there is provided a system comprising:

a display for rendering content to a user;

an eye tracker for tracking movement or motion with respect to an eye of the user;

a microprocessor for generating the content to be rendered by the display;

a non-volatile, non-transitory storage medium comprising executable instructions for execution by the microprocessor, where the executable instructions when executed by the microprocessor configure the microprocessor to perform a process comprising the steps of:

a) determine a requirement to redirect the motion of the user from a current direction to a target direction;
b) determine an occurrence of an event with respect to an eye of a user;
c) determine whether the saccade is a minor saccade, a major saccade, or a blink;
d) upon determination of a minor saccade executing a next step of a first process which redirects the motion of the user from the current direction to the target direction in a number of steps;
e) upon determination of a major saccade or a blink executing a second process which redirects the motion of the user from the current direction to the target direction in a single step; and
f) repeat steps (b) to (f) until motion of the user has been redirected to the target direction.

Other aspects and features of the present invention will become apparent to those ordinarily skilled in the art upon review of the following description of specific embodiments of the invention in conjunction with the accompanying figures.

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments of the present invention will now be described, by way of example only, with reference to the attached Figures, wherein:

FIG. **1** depicts an exemplary network environment within which configurable electrical devices according to and supporting embodiments of the invention may be deployed and operate; and

FIG. **2** depicts an exemplary wireless portable electronic device supporting communications to a network such as depicted in FIG. **1** and supporting embodiments of the invention;

FIG. **3**A depicts an exemplary head mounted display (HMD) supporting communications to a network such as depicted in FIG. **1**, an electronic device such as described in FIG. **2**, and supporting embodiments of the invention;

FIG. **3**B depicts exemplary commercial HMDs supporting embodiments of the invention;

FIG. **4**A depicts orientation-based redirection using dynamic-foveated rendering which leverages the effect of inattentional blindness induced by a cognitive task according to an embodiment of the invention;

FIG. **4**B depicts an image of a path walked in physical tracked space (PTS) with the corresponding path in the virtual environment (VE) achieved by a user exploiting an HMD supporting embodiments of the invention;

FIG. **5** depicts an exemplary pipeline for implemented redirecting walking according to an embodiment of the invention exploiting rendering the VE from two co-located cameras;

FIG. **6** depicts a user's perspective during a first user study within a VE as rendered according to an embodiment of the invention;

FIG. **7** depicts participant responses from the first user study employing VE rendered user perspectives such as depicted in FIG. **6**;

FIG. **8** depicts results of Simulation Sickness Questionnaire (SSQ) from the first user study employing VE rendered user perspectives such as depicted in FIG. **6** for determining maximum rotation angle and field-of-view of the foveal zone;

FIG. **9** depicts results of Simulation Sickness Questionnaire (SSQ) from a second user study with dynamic foveated rendering and parameters established from the first user study;

FIGS. **10**A to **10**C depict exemplary results of participants from a third user study showing the real world PTS and virtual environment path using inattentional blindness parameters as established by the inventors with dynamic foveated rendering;

FIG. **11** depicts an exemplary image from the game designed for the third user study;

FIGS. **12**A to **12**C depict analysis of variable (ANOVA) results for number of resets, distance travelled in PTS, and total time taken;

FIG. **13** depicts schematically exploiting minor saccades within embodiments of the invention; and

FIG. **14** depicts a flowchart illustrating a process based upon an embodiment of the invention.

DETAILED DESCRIPTION

The present invention is directed to virtual reality and more particularly to locomotion within virtual reality environments where redirected walking allows a user to navigate a virtual environment spatially larger than the available physical space.

The ensuing description provides representative embodiment(s) only, and is not intended to limit the scope, applicability or configuration of the disclosure. Rather, the ensuing description of the embodiment(s) will provide those skilled in the art with an enabling description for implementing an embodiment or embodiments of the invention. It being understood that various changes can be made in the function and arrangement of elements without departing from the spirit and scope as set forth in the appended claims. Accordingly, an embodiment is an example or implementation of the inventions and not the sole implementation. Various appearances of "one embodiment," "an embodiment" or "some embodiments" do not necessarily all refer to the same embodiments. Although various features of the invention may be described in the context of a single embodiment, the features may also be provided separately or in any suitable combination. Conversely, although the invention may be described herein in the context of separate embodiments for clarity, the invention can also be implemented in a single embodiment or any combination of embodiments.

Reference in the specification to "one embodiment", "an embodiment", "some embodiments" or "other embodiments" means that a particular feature, structure, or characteristic described in connection with the embodiments is included in at least one embodiment, but not necessarily all embodiments, of the inventions. The phraseology and terminology employed herein is not to be construed as limiting but is for descriptive purpose only. It is to be understood that where the claims or specification refer to "a" or "an" element, such reference is not to be construed as there being only one of that element. It is to be understood that where the specification states that a component feature, structure, or characteristic "may", "might", "can" or "could" be included, that particular component, feature, structure, or characteristic is not required to be included.

Reference to terms such as "left", "right", "top", "bottom", "front" and "back" are intended for use in respect to the orientation of the particular feature, structure, or element within the figures depicting embodiments of the invention. It would be evident that such directional terminology with

respect to the actual use of a device has no specific meaning as the device can be employed in a multiplicity of orientations by the user or users.

Reference to terms "including", "comprising", "consisting" and grammatical variants thereof do not preclude the addition of one or more components, features, steps, integers or groups thereof and that the terms are not to be construed as specifying components, features, steps or integers. Likewise, the phrase "consisting essentially of", and grammatical variants thereof, when used herein is not to be construed as excluding additional components, steps, features integers or groups thereof but rather that the additional features, integers, steps, components or groups thereof do not materially alter the basic and novel characteristics of the claimed composition, device or method. If the specification or claims refer to "an additional" element, that does not preclude there being more than one of the additional element.

A "wireless standard" as used herein and throughout this disclosure, refer to, but is not limited to, a standard for transmitting signals and/or data through electromagnetic radiation which may be optical, radio-frequency (RF) or microwave although typically RF wireless systems and techniques dominate. A wireless standard may be defined globally, nationally, or specific to an equipment manufacturer or set of equipment manufacturers. Dominant wireless standards at present include, but are not limited to IEEE 802.11, IEEE 802.15, IEEE 802.16, IEEE 802.20, UMTS, GSM 850, GSM 900, GSM 1800, GSM 1900, GPRS, ITU-R 5.138, ITU-R 5.150, ITU-R 5.280, IMT-1000, Bluetooth, Wi-Fi, Ultra-Wideband and WiMAX. Some standards may be a conglomeration of sub-standards such as IEEE 802.11 which may refer to, but is not limited to, IEEE 802.1a, IEEE 802.11b, IEEE 802.11g, or IEEE 802.11n as well as others under the IEEE 802.11 umbrella.

A "wired standard" as used herein and throughout this disclosure, generally refer to, but is not limited to, a standard for transmitting signals and/or data through an electrical cable discretely or in combination with another signal. Such wired standards may include, but are not limited to, digital subscriber loop (DSL), Dial-Up (exploiting the public switched telephone network (PSTN) to establish a connection to an Internet service provider (ISP)), Data Over Cable Service Interface Specification (DOCSIS), Ethernet, Gigabit home networking (G.hn), Integrated Services Digital Network (ISDN), Multimedia over Coax Alliance (MoCA), and Power Line Communication (PLC, wherein data is overlaid to AC/DC power supply). In some embodiments a "wired standard" may refer to, but is not limited to, exploiting an optical cable and optical interfaces such as within Passive Optical Networks (PONs) for example.

A "sensor" as used herein may refer to, but is not limited to, a transducer providing an electrical output generated in dependence upon a magnitude of a measure and selected from the group comprising, but is not limited to, environmental sensors, medical sensors, biological sensors, chemical sensors, ambient environment sensors, position sensors, motion sensors, thermal sensors, infrared sensors, visible sensors, RFID sensors, and medical testing and diagnosis devices.

A "portable electronic device" (PED) as used herein and throughout this disclosure, refers to a wireless device used for communications and other applications that requires a battery or other independent form of energy for power. This includes devices, but is not limited to, such as a cellular telephone, smartphone, personal digital assistant (PDA), portable computer, pager, portable multimedia player, portable gaming console, laptop computer, tablet computer, a wearable device and an electronic reader.

A "fixed electronic device" (FED) as used herein and throughout this disclosure, refers to a wireless and/or wired device used for communications and other applications that requires connection to a fixed interface to obtain power. This includes, but is not limited to, a laptop computer, a personal computer, a computer server, a kiosk, a gaming console, a digital set-top box, an analog set-top box, an Internet enabled appliance, an Internet enabled television, and a multimedia player.

A "server" as used herein, and throughout this disclosure, refers to one or more physical computers co-located and/or geographically distributed running one or more services as a host to users of other computers, PEDs, FEDs, etc. to serve the client needs of these other users. This includes, but is not limited to, a database server, file server, mail server, print server, web server, gaming server, or virtual environment server.

An "application" (commonly referred to as an "app") as used herein may refer to, but is not limited to, a "software application", an element of a "software suite", a computer program designed to allow an individual to perform an activity, a computer program designed to allow an electronic device to perform an activity, and a computer program designed to communicate with local and/or remote electronic devices. An application thus differs from an operating system (which runs a computer), a utility (which performs maintenance or general-purpose chores), and a programming tools (with which computer programs are created). Generally, within the following description with respect to embodiments of the invention an application is generally presented in respect of software permanently and/or temporarily installed upon a PED and/or FED.

An "enterprise" as used herein may refer to, but is not limited to, a provider of a service and/or a product to a user, customer, or consumer. This includes, but is not limited to, a retail outlet, a store, a market, an online marketplace, a manufacturer, an online retailer, a charity, a utility, and a service provider. Such enterprises may be directly owned and controlled by a company or may be owned and operated by a franchisee under the direction and management of a franchiser.

A "service provider" as used herein may refer to, but is not limited to, a third party provider of a service and/or a product to an enterprise and/or individual and/or group of individuals and/or a device comprising a microprocessor. This includes, but is not limited to, a retail outlet, a store, a market, an online marketplace, a manufacturer, an online retailer, a utility, an own brand provider, and a service provider wherein the service and/or product is at least one of marketed, sold, offered, and distributed by the enterprise solely or in addition to the service provider.

A "third party" or "third party provider" as used herein may refer to, but is not limited to, a so-called "arm's length" provider of a service and/or a product to an enterprise and/or individual and/or group of individuals and/or a device comprising a microprocessor wherein the consumer and/or customer engages the third party but the actual service and/or product that they are interested in and/or purchase and/or receive is provided through an enterprise and/or service provider.

A "user" as used herein may refer to, but is not limited to, an individual or group of individuals. This includes, but is not limited to, private individuals, employees of organizations and/or enterprises, members of community organizations, members of charity organizations, men and women. In

its broadest sense the user may further include, but not be limited to, software systems, mechanical systems, robotic systems, android systems, etc. that may be characterised by an ability to exploit one or more embodiments of the invention. A user may also be associated through one or more accounts and/or profiles with one or more of a service provider, third party provider, enterprise, social network, social media etc. via a dashboard, web service, website, software plug-in, software application, and graphical user interface.

"Biometric" information as used herein may refer to, but is not limited to, data relating to a user characterised by data relating to a subset of conditions including, but not limited to, their environment, medical condition, biological condition, physiological condition, chemical condition, ambient environment condition, position condition, neurological condition, drug condition, and one or more specific aspects of one or more of these said conditions. Accordingly, such biometric information may include, but not be limited, blood oxygenation, blood pressure, blood flow rate, heart rate, temperate, fluidic pH, viscosity, particulate content, solids content, altitude, vibration, motion, perspiration, EEG, ECG, energy level, etc. In addition, biometric information may include data relating to physiological characteristics related to the shape and/or condition of the body wherein examples may include, but are not limited to, fingerprint, facial geometry, baldness, DNA, hand geometry, odour, and scent. Biometric information may also include data relating to behavioral characteristics, including but not limited to, typing rhythm, gait, and voice.

"User information" as used herein may refer to, but is not limited to, user behavior information and/or user profile information. It may also include a user's biometric information, an estimation of the user's biometric information, or a projection/prediction of a user's biometric information derived from current and/or historical biometric information.

A "wearable device" or "wearable sensor" relates to miniature electronic devices that are worn by the user including those under, within, with or on top of clothing and are part of a broader general class of wearable technology which includes "wearable computers" which in contrast are directed to general or special purpose information technologies and media development. Such wearable devices and/or wearable sensors may include, but not be limited to, smartphones, smart watches, e-textiles, smart shirts, activity trackers, smart glasses, environmental sensors, medical sensors, biological sensors, physiological sensors, chemical sensors, ambient environment sensors, position sensors, neurological sensors, drug delivery systems, medical testing and diagnosis devices, motion sensors, and head mounted displays (HMDs).

"Electronic content" (also referred to as "content" or "digital content") as used herein may refer to, but is not limited to, any type of content that exists in the form of digital data as stored, transmitted, received and/or converted wherein one or more of these steps may be analog although generally these steps will be digital. Forms of digital content include, but are not limited to, information that is digitally broadcast, streamed or contained in discrete files. Viewed narrowly, types of digital content include popular media types such as MP3, JPG, AVI, TIFF, AAC, TXT, RTF, HTML, XHTML, PDF, XLS, SVG, WMA, MP4, FLV, and PPT, for example, as well as others, see for example http://en.wikipedia.org/wiki/List_of_file_formats. Within a broader approach digital content mat include any type of digital information, e.g. digitally updated weather forecast, a GPS map, an eBook, a photograph, a video, a Vine™, a blog posting, a Facebook™ posting, a Twitter™ tweet, online TV, etc. The digital content may be any digital data that is at least one of generated, selected, created, modified, and transmitted in response to a user request, said request may be a query, a search, a trigger, an alarm, and a message for example.

A "profile" as used herein, and throughout this disclosure, refers to a computer and/or microprocessor readable data file comprising data relating to settings and/or limits of an adult device. Such profiles may be established by a manufacturer/supplier/provider of a device, service, etc. or they may be established by a user through a user interface for a device, a service or a PED/FED in communication with a device, another device, a server or a service provider etc.

A "computer file" (commonly known as a file) as used herein, and throughout this disclosure, refers to a computer resource for recording data discretely in a computer storage device, this data being electronic content. A file may be defined by one of different types of computer files, designed for different purposes. A file may be designed to store electronic content such as a written message, a video, a computer program, or a wide variety of other kinds of data. Some types of files can store several types of information at once. A file can be opened, read, modified, copied, and closed with one or more software applications an arbitrary number of times. Typically, files are organized in a file system which can be used on numerous different types of storage device exploiting different kinds of media which keeps track of where the files are located on the storage device(s) and enables user access. The format of a file is defined by its content since a file is solely a container for data, although, on some platforms the format is usually indicated by its filename extension, specifying the rules for how the bytes must be organized and interpreted meaningfully. For example, the bytes of a plain text file are associated with either ASCII or UTF-8 characters, while the bytes of image, video, and audio files are interpreted otherwise. Some file types also allocate a few bytes for metadata, which allows a file to carry some basic information about itself.

A "major saccade" as used herein, and throughout disclosure, refers to a quick, simultaneous movement of both eyes between two or more phases of fixation in the same direction. A major saccade is a rapid, ballistic movement of the eyes that abruptly changes the point of fixation. A major saccade is typically defined by movement of the eyes with a velocity greater than 180°.

A "minor saccade" or "microsaccade" as used herein, and throughout disclosure, refers to a fixational eye movements which are small, jerk-like, involuntary eye movements, similar to miniature versions of major saccades but typically occurring during visual fixation. A minor saccade is typically defined by movement of the eyes with a velocity less than 180°/second.

A "Head Mounted Display" (HMD) as used herein, and throughout this disclosure, refers to is a display device, worn on the head or as part of a helmet which has a small display device in front of one (monocular HMD) of a user or each eye of the user (binocular HMD). HMDs differ in whether they can display only computer-generated imagery (CGI) in what are referred to commonly as virtual reality (VR) or a virtual environment (VE), only live imagery from the physical world, or a combination thereof in what is referred to as augmented reality (AR) or mixed reality (MR).

An "eye-tracker" as used herein, and throughout this disclosure, refers to a device which measures rotations of an eye. Eye-trackers principally exploit measurement of the

movement of an object (e.g. a special contact lens) attached to the eye; optical tracking without direct contact to the eye; and measurement of electric potentials using electrodes placed around the eyes. Optical methods are widely used for gaze-tracking and are favored for being non-invasive and inexpensive where light, for example infrared light, is reflected from the eye and sensed by a camera or some other specially designed optical sensor. The information is then analyzed to extract eye rotation from changes in reflections.

1. Re-Directed Walking

Since the early days of virtual reality researchers have investigated ways of users navigating virtual environments (VEs) that are spatially larger than the available Physical Tracked Space (PTS). A number of locomotion techniques relying on pointing devices or walking in-place were proposed which have since become customary in VE applications. However, users find these methods cumbersome and unnatural. The concept of redirected walking was introduced about 20 years ago in order to provide a more natural way of navigating VEs, albeit with many restrictions on the shape and size of the physical and virtual spaces. Re-directed walking is important as it has wide range of applications outside the dominant use of VEs for gaming. For example, an architect can "walk-through" an environment to assess aspects of its design; emergency responders can assess fire exits, emergency procedures etc.; users can wander virtual museums; users can physically move during participation in VEs for any application.

As noted above whilst a number of approaches have been proposed both in hardware and software these suffer limitations such as complexity, cost, introducing artifacts and distortions into the VE etc. or forcing the user to look away by inducing major saccades in order to allow the VE to be updates during a rapid eye movement. However, the inventors have established an alternate solution based upon the psychological phenomenon of inattentional blindness in conjunction with managing/partitioning the image content rendered to the user.

Inattentional blindness refers to the inability of an individual to see a salient object in plain sight, due to their lack of attention. The inventor's methodology exploits this phenomenon and further strengthen its effect with foveated rendering. Foveated rendering is a rendering technique whose primary objective is to reduce the workload on the microprocessor and/or graphics processor generating the images for rendering. Using eye tracking the user's eyes are tracked within the VE headset in real-time. The zone in the image corresponding to the foveal vision, i.e. the zone gazed by the fovea which provides sharp and detailed vision, is rendered at high quality. On the other hand, the zone in the image corresponding to the peripheral vision is rendered at a lower quality since the user's peripheral vision lacks acuity albeit it has a wider field of view. This process is performed without causing any perceptual change to the user.

Foveated rendering is supported today within hardware such as NVIDIA's RTX graphics card series which allows for real-time ray-tracing and software, such as hence, real-time performance as well as by techniques from eSight Corp. for managing content to be rendered through multiple files which are generated for an image to be rendered to the user and transmitted to the display for rendering.

Accordingly, the inventor's inventive methodology applies spatially varying rotations to different zones of the image rendered to the user according to the zone's importance using foveated rendering to strengthen the effect of inattentional blindness. Accordingly, a user's real world motion can be re-directed such that their motion within the real world is within a small physical space whilst their apparent VE motion is over a larger physical space. For example, a user may walk through a building navigating, for example, 300 meters (985 feet) where the actual physical space they walked within was a 4 m (13 foot) square room. User studies were performed to determine the maximum rotation angle and field-of-view for which participants do not perceive a change. These then were employed with foveated rendering to verify the in-situ gaze redirection in the VE environment results in equivalent motion of the user within the PTS. Beneficially, the inventor's methodology provides for increased VE travel within the PTS with reduced resets due to the user impacting a limit in the PTS, e.g. a wall. Further, as will become evident the techniques implemented by the inventors can be extended to provide not only redirected walking within a PTS which is devoid of obstacles, e.g. an empty room or space, but also within a PTS with obstacles, either fixed or mobile, including other users. Accordingly, through tracking multiple users a VE simulating, for example, an emergency evacuation can be performed with multiple users concurrently.

2. Exemplary Head-Mounted Display, Associated Electronic Devices, and Network Environment Supporting Embodiments of the Invention

Referring to FIG. **1** there is depicted a Network **100** within which embodiments of the invention may be employed supporting Redirect Walking (RW) Systems, Applications and Platforms (RW-SAPs) according to embodiments of the invention. Such RW-SAPs, for example, supporting multiple communication channels, dynamic filtering, etc. As shown first and second user groups **100**A and **100**B respectively interface to a telecommunications Network **100**. Within the representative telecommunication architecture, a remote central exchange **180** communicates with the remainder of a telecommunication service providers network via the Network **100** which may include for example long-haul OC-48/OC-192 backbone elements, an OC-48 wide area network (WAN), a Passive Optical Network, and a Wireless Link. The central exchange **180** is connected via the Network **100** to local, regional, and international exchanges (not shown for clarity) and therein through Network **100** to first and second cellular APs **195**A and **195**B respectively which provide Wi-Fi cells for first and second user groups **100**A and **100**B respectively. Also connected to the Network **100** are first and second Wi-Fi nodes **110**A and **1101**B, the latter of which being coupled to Network **100** via router **105**. Second Wi-Fi node **1101**B is associated with commercial service provider **160** and comprises other first and second user groups **100**A and **100**B. Second user group **100**B may also be connected to the Network **100** via wired interfaces including, but not limited to, DSL, Dial-Up, DOCSIS, Ethernet, G.hn, ISDN, MoCA, PON, and Power line communication (PLC) which may or may not be routed through a router such as router **105**. As will become evident in respect of FIGS. **2** and **3** PEDs and/or FEDs within first and second user groups **100**A and **100**B may provide the role of an electronic device, e.g. Electronic Device **204** or Electronic Device **310**, to which a Head Mounted Display (HMD) may be interfaced in order to provide communications to/from Network **100** and therein other devices, systems, servers, etc.

Within the cell associated with first AP **110**A the first group of users **100**A may employ a variety of PEDs including for example, laptop computer **155**, portable gaming console **135**, tablet computer **140**, smartphone **150**, cellular telephone **145** as well as portable multimedia player **130**. Within the cell associated with second AP **110**B are the

second group of users **100**B which may employ a variety of FEDs including for example gaming console **125**, personal computer **115** and wireless/Internet enabled television **120** as well as cable modem **105**. First and second cellular APs **195**A and **195**B respectively provide, for example, cellular GSM (Global System for Mobile Communications) telephony services as well as 3G and 4G evolved services with enhanced data transport support. Second cellular AP **195**B provides coverage in the exemplary embodiment to first and second user groups **100**A and **100**B. Alternatively the first and second user groups **100**A and **100**B may be geographically disparate and access the Network **100** through multiple APs, not shown for clarity, distributed geographically by the network operator or operators. First cellular AP **195**A as show provides coverage to first user group **100**A and environment **170**, which comprises second user group **100**B as well as first user group **100**A. Accordingly, the first and second user groups **100**A and **100**B may according to their particular communications interfaces communicate to the Network **100** through one or more wireless communications standards such as, for example, IEEE 802.11, IEEE 802.15, IEEE 802.16, IEEE 802.20, UMTS, GSM 850, GSM 900, GSM 1800, GSM 1900, GPRS, ITU-R 5.138, ITU-R 5.150, ITU-R 5.280, and IMT-1000. It would be evident to one skilled in the art that many portable and fixed electronic devices may support multiple wireless protocols simultaneously, such that for example a user may employ GSM services such as telephony and SMS and Wi-Fi/WiMAX data transmission, VOIP and Internet access. Accordingly, portable electronic devices within first user group **100**A may form associations either through standards such as IEEE 802.15 or Bluetooth as well in an ad-hoc manner.

Also connected to the Network **100** are Social Networks (SOCNETS) **165**, first and second service providers **170**A and **170**B respectively, first and second third party service providers **170**C and **170**D respectively, and a user **170**E who may receive data from one or more RW-SAPs and/or HMD(s). Also connected to the Network **100** are first and second enterprises **175**A and **175**B respectively, first and second organizations **175**C and **175**D respectively, and a government entity **175**E who may receive data from one or more RW-SAPs and/or HMD(s). Also depicted are first and second servers **190**A and **190**B may host according to embodiments of the inventions multiple services associated with a provider of Redirect Walking (RW) Systems, Applications and Platforms (RW-SAPs); a provider of a SOCNET or Social Media (SOME) exploiting RW-SAP features; a provider of a SOCNET and/or SOME not exploiting RW-SAP features; a provider of services to PEDS and/or FEDS; a provider of one or more aspects of wired and/or wireless communications; an Enterprise **160** exploiting RW-SAP features; license databases; content databases; image databases; content libraries; customer databases; websites; and software applications for download to or access by FEDs and/or PEDs exploiting and/or hosting RW-SAP features. First and second primary content servers **190**A and **190**B may also host for example other Internet services such as a search engine, financial services, third party applications and other Internet based services.

Also depicted in FIG. **1** are Electronic Devices (EDs) **1000** according to embodiments of the invention such as described and depicted below in respect of FIGS. **3**A and **3**B which support RW-SAPs functionality and features as described and depicted in respect of FIGS. **4**A to **12** respectively. As depicted in FIG. **1** an ED **1000** may communicate directly to the Network **100** through one or more wireless or wired interfaces included those, for example, selected from the group comprising IEEE 802.11, IEEE 802.15, IEEE 802.16, IEEE 802.20, UMTS, GSM 850, GSM 900, GSM 1800, GSM 1900, GPRS, ITU-R 5.138, ITU-R 5.150, ITU-R 5.280, IMT-1000, DSL, Dial-Up, DOCSIS, Ethernet, G.hn, ISDN, MoCA, PON, and Power line communication (PLC).

Accordingly, a consumer and/or customer (CONCUS) may exploit a PED and/or FED within an Enterprise **160**, for example, and access one of the first or second primary content servers **190**A and **190**B respectively to perform an operation such as accessing/downloading an application which provides RW-SAP features according to embodiments of the invention; execute an application already installed providing RW-SAP features; execute a web based application providing RW-SAP features; or access content. Similarly, a CONCUS may undertake such actions or others exploiting embodiments of the invention exploiting a PED or FED within first and second user groups **100**A and **100**B respectively via one of first and second cellular APs **195**A and **195**B respectively and first Wi-Fi nodes **110**A. It would also be evident that a CONCUS may, via exploiting Network **100** communicate via telephone, fax, email, SMS, social media, etc.

Now referring to FIG. **2** there is depicted an Electronic Device **204** and network access point **207** supporting RW-SAP features according to embodiments of the invention. Electronic Device **204** may, for example, be a PED and/or FED and may include additional elements above and beyond those described and depicted. Also depicted within the Electronic Device **204** is the protocol architecture as part of a simplified functional diagram of a system that includes an Electronic Device **204**, such as a smartphone **155**, an Access Point **206**, such as first AP **110**, and one or more network devices **207**, such as communication servers, streaming media servers, and routers for example such as first and second servers **190**A and **190**B respectively. Network Devices **207** may be coupled to Access Point **206** via any combination of networks, wired, wireless and/or optical communication links such as discussed above in respect of FIG. **1** as well as directly as indicated. Network devices **207** are coupled to Network **100** and therein Social Networks (SOCNETS) **165**, first and second service providers **170**A and **170**B respectively, first and second third party service providers **170**C and **170**D respectively, a user **170**E, first and second enterprises **175**A and **175**B respectively, first and second organizations **175**C and **175**D respectively, and a government entity **175**E.

The Electronic Device **204** includes one or more Processors **210** and a Memory **212** coupled to Processor(s) **210**. Access Point **206** also includes one or more Processors **211** and a Memory **213** coupled to processor(s) **211**. A non-exhaustive list of examples for any of processors **210** and **211** includes a central processing unit (CPU), a digital signal processor (DSP), a reduced instruction set computer (RISC), a complex instruction set computer (CISC) and the like. Furthermore, any of Processors **210** and **211** may be part of application specific integrated circuits (ASICs) or may be a part of application specific standard products (ASSPs). A non-exhaustive list of examples for Memories **212** and **213** includes any combination of the following semiconductor devices such as registers, latches, ROM, EEPROM, flash memory devices, non-volatile random access memory devices (NVRAM), SDRAM, DRAM, double data rate (DDR) memory devices, SRAM, universal serial bus (USB) removable memory, and the like.

Electronic Device **204** may include an audio input element **214**, for example a microphone, and an audio output

element **216**, for example, a speaker, coupled to any of Processors **210**. Electronic Device **204** may include a video input element **218**, for example, a video camera or camera, and a video output element **220**, for example an LCD display, coupled to any of Processors **210**. Electronic Device **204** also includes a Keyboard **215** and Touchpad **217** which may for example be a physical keyboard and touchpad allowing the user to enter content or select functions within one of more Applications **222**. Alternatively, the Keyboard **215** and Touchpad **217** may be predetermined regions of a touch sensitive element forming part of the display within the Electronic Device **204**. The one or more Applications **222** that are typically stored in Memory **212** and are executable by any combination of Processors **210**. Electronic Device **204** also includes Accelerometer **260** providing three-dimensional motion input to the Processor **210** and GPS **262** which provides geographical location information to Processor **210**.

Electronic Device **204** includes a Protocol Stack **224** and AP **206** includes an Access Point Stack **225**. Protocol Stack **224** is shown as an IEEE 802.11 protocol stack but alternatively may exploit other protocol stacks such as an Internet Engineering Task Force (IETF) multimedia protocol stack for example. Likewise, AP Stack **225** exploits a protocol stack but is not expanded for clarity. Elements of Protocol Stack **224** and AP Stack **225** may be implemented in any combination of software, firmware and/or hardware. Protocol Stack **224** accordingly, when providing an IEEE 802.11 protocol stack includes an IEEE 802.11-compatible PHY module that is coupled to one or more Front-End Tx/Rx & Antenna **228**, an IEEE 802.11-compatible MAC module coupled to an IEEE 802.2-compatible LLC module. Protocol Stack **224** includes a network layer IP module, a transport layer User Datagram Protocol (UDP) module and a transport layer Transmission Control Protocol (TCP) module.

Protocol Stack **224** also includes a session layer Real Time Transport Protocol (RTP) module, a Session Announcement Protocol (SAP) module, a Session Initiation Protocol (SIP) module and a Real Time Streaming Protocol (RTSP) module. As depicted Protocol Stack **224** also includes a presentation layer Call Control and Media Negotiation module **250**, one or more Audio Codecs **252** and one or more Video Codecs **254**. Applications **222** may be able to create maintain and/or terminate communication sessions with any of Network Devices **207** by way of AP **206**. Typically, Applications **222** may activate any of the SAP, SIP, RTSP, media negotiation and call control modules for that purpose. Typically, information may propagate from the SAP, SIP, RTSP, media negotiation and call control modules to PHY module through TCP module, IP module, LLC module and MAC module.

It would be apparent to one skilled in the art that elements of the Electronic Device **204** may also be implemented within the Access Point **206** including but not limited to one or more elements of the Protocol Stack **224**, including for example an IEEE 802.11-compatible PHY module, an IEEE 802.11-compatible MAC module, and an IEEE 802.2-compatible LLC module. The Access Point **206** may additionally include a network layer IP module, a transport layer User Datagram Protocol (UDP) module and a transport layer Transmission Control Protocol (TCP) module as well as a session layer Real Time Transport Protocol (RTP) module, a Session Announcement Protocol (SAP) module, a Session Initiation Protocol (SIP) module and a Real Time Streaming Protocol (RTSP) module, media negotiation module, and a call control module. Portable and fixed electronic devices represented by Electronic Device **204** may include one or more additional wireless or wired interfaces in addition to the depicted IEEE 802.11 interface which may be selected from the group comprising IEEE 802.15, IEEE 802.16, IEEE 802.20, UMTS, GSM 850, GSM 900, GSM 1800, GSM 1900, GPRS, ITU-R 5.138, ITU-R 5.150, ITU-R 5.280, IMT-1000, DSL, Dial-Up, DOCSIS, Ethernet, G.hn, ISDN, MoCA, PON, and Power line communication (PLC).

Also depicted in FIG. **2** are Electronic Devices (EDs) **100** according to embodiments of the invention such as described and depicted below in respect of FIGS. **3**A and **3**B. As depicted in FIG. **2** an ED **1000** may communicate directly to the Network **100**. Other EDs **1000** may communicate to the Network Device **207**, Access Point **206**, and Electronic Device **204**. Some EDs **1000** may communicate to other EDs **1000** directly. Within FIG. **2** the EDs **1000** coupled to the Network **100** and Network Device **207** communicate via wired interfaces but these may alternatively be wireless interfaces. The EDs **1000** coupled to the Access Point **206** and Electronic Device **204** communicate via wireless interfaces and/or wired interfaces. Each ED **1000** may communicate to another electronic device, e.g. Access Point **206**, Electronic Device **204** and Network Device **207**, or a network, e.g. Network **100**. Each ED **1000** may support one or more wireless or wired interfaces including those, for example, selected from the group comprising IEEE 802.11, IEEE 802.15, IEEE 802.16, IEEE 802.20, UMTS, GSM 850, GSM 900, GSM 1800, GSM 1900, GPRS, ITU-R 5.138, ITU-R 5.150, ITU-R 5.280, IMT-1000, DSL, Dial-Up, DOCSIS, Ethernet, G.hn, ISDN, MoCA, PON, and Power line communication (PLC).

Accordingly, FIG. **2** depicts an Electronic Device **204**, e.g. a PED, wherein one or more parties including, but not limited to, a user, users, an enterprise, enterprises, third party provider, third party providers, wares provider, wares providers, financial registry, financial registries, financial provider, and financial providers may engage in one or more activities and/or transactions relating to an activity including, but not limited to, e-business, P2P, C2B, B2B, C2C, B2G, C2G, P2D, D2D, gaming, regulatory compliance, architectural design, emergency services, etc. via the Network **100** using the Electronic Device **204** or within either the Access Point **206** or Network Device **207** wherein details of the transaction are then coupled to the Network **100** and stored within remote servers. Optionally, rather than wired and/or wireless communication interfaces devices may exploit other communication interfaces such as optical communication interfaces and/or satellite communications interfaces. Optical communications interfaces may support Ethernet, Gigabit Ethernet, SONET, Synchronous Digital Hierarchy (SDH) etc.

Within embodiments of the invention the Electronic Device **204** may itself be an HMD or as described and depicted in respect of FIG. **3**A an Electronic Device **310**, such as Electronic Device **204**, for example may interface to an HMD **340**. Accordingly, as depicted in FIG. **3**A the HMD **340** is coupled to an Electronic Device **310** through a Wireless Personal Area Network (WPAN) interface between Front End Tx/Rx & Antenna (Antenna) **330** and WPAN Tx/Rx & Antenna **390**. Antenna **330** is connected to Electronics **320** which comprises the microprocessor, user interfaces, memory, software etc. as described with respect to Electronic Device **204** in FIG. **2**.

Within HMD **340** the WPAN Antenna **390** is connected to HMD Stack **385** and therein to Processor **380**. As depicted the Processor **340** is coupled to elements of the HMD **340** which include, but are not limited to, Accelerometer(s) **345**, GPS Receiver **350**, Eye-Tracker(s) **360**, Display(s) **365**,

Memory **370**, Camera(s) **375** and Sensor(s) **395**. Within other embodiments of the invention the HMD **340** may be designed solely for immersive applications and not include a Camera(s) **375**. Similarly, the HMD **340** may not include a GPS Receiver **350** to provide location data in conjunction with motion tracking provided by the Accelerometer(s) **345**. Optionally, the HMD **340** may not include Sensor(s) **395**. However, within embodiments of the invention the Sensor(s) **395** may include orientation sensors, such as tilt sensors for example, and distance measurement sensors. For example, a distance measurement sensor may include, but not be limited to, a Light Detection and Ranging sensor (lidar), sound navigation ranging (sonar), a structured light source or structured light 3D scanner, such that the HMD **340** can determine a distance or distances to objects within the environment of a user of the HMD **340** which supports embodiments of the invention. Accordingly, within the embodiments of the invention described and depicted in respect of FIGS. **4**A to **12** for the redirection of a user's motion a HMD **340** may establish an initial measurements of environment the user of the HMD **340** is within together with the user's location within that environment and employ this data to define aspects of the redirection process according to embodiments of the invention. Further, the redirection process according to embodiments of the invention may receive continuous and/or periodic environment updates to provide additional data for the redirection process according to embodiments of the invention or initiate a reset.

Further, whilst the redirection process according to embodiments of the invention is described and depicted in FIGS. **4**A to **12** with respect to a user within a defined space devoid of obstacles it would be evident that the exploitation of such Sensor(s) **395** to provide environment information can also be employed within the redirection process according to embodiments of the invention to allow the user's motion to be redirected such that they avoid real world physical objects whilst interacting with the virtual environment being presented to them. For example, a user employing a HMD **340** with visual defects or degradations may be presented with a virtual environment wherein the redirection process according to embodiments of the invention employed allow the user to navigate their environment where their motion would otherwise lead them to hit an element of their environment, e.g. wall, door, chair, etc. Accordingly, whilst the redirection process according to embodiments of the invention are described and depicted with respect to a user achieving a large virtual travel distance within a space smaller than the distance travelled the redirection process according to embodiments of the invention may also be employed where the user's space is not as restricted or limited.

Optionally, within other embodiments of the invention the Sensor(s) **395** may be associated with other electronic devices worn and/or carried by the user for example wherein the data from these Sensor(s) **395** is communicated to the HMD **340** where the redirection process according to embodiments of the invention are performed by the Processor **380** within the HMD **340**. Accordingly, the Sensor(s) **395** may be associated with wearable devices as well as PEDs and/or FEDs. Alternatively, where the redirection process according to embodiments of the invention are performed by another electronic device, e.g. Electronic Device **310** in FIG. **3**A, then the data from the Sensor(s) **395** within the HMD **340** or other electronic devices worn and/or carried by the user would be communicated to the processor within the Electronic Device **310**. Alternatively, processing may be provided by a remote service upon a remote server wherein the data from the Sensor(s) **395** within the HMD **340** or other electronic devices worn and/or carried by the user would be communicated to the remote server or service.

Optionally, the HMD **340** may include one or more haptic interfaces which provide information to the user by means other than through the Display(s) **365**. For example, the HMD **340** may include a vibratory motor, a loudspeaker, a buzzer etc. which is triggered by a redirection process according to an embodiment of the invention, such as for example when a reset is to be performed, when the redirection process is exceeding a predetermined limit to seek redirection without a reset, etc. Optionally, the one or more haptic interfaces may be within one or more other electronic devices worn and/or carried by the user and receive data from one or more of the HMD **340**, Electronic Device **310**, or remote server etc. Accordingly, the haptic interface(s) may be associated with wearable devices.

Accordingly, HMD **340** may, for example, utilize the processor within Electronic Device **310**, e.g. Processor **210** within Electronic Device **204** in FIG. **2**, processing functionality such that a lower power Processor **380** may be employed within HMD **340** controlling, for example, acquisition of image data from Camera(s) **375**, Sensor(s) **395**, Eye-Tracker(s) **360**, Accelerometer(s) **345** **2076** and presentation of image data to the user via Display(s) **365** with information including one or more of instruction sets, algorithm(s), user profile(s), user biometric data etc. being stored within a memory of the HMD **340**, not depicted for clarity or an electronic device, such as Electronic Device **310**, which is exchanging data/information with HMD **340**, e.g. the Electronic Device **310** and HMD **340** may be paired as known in wireless interfacing.

As noted above a HMD **340** may be employed to simply present visual content to the user, e.g. within an immersive virtual reality environment, or it may be employed to provide visual content to the user which has been processed to address and/or compensate for visual defects and/or vision degradations etc. either in an immersive virtual reality environment, an immersive real world environment, or an augmented reality environment. Accordingly, the HMD **340** and/or an electronic device associated with the user, e.g. Electronic Device **310**, may store data relating to a particular individual's visual defects within a memory, such as Memory **212** of Electronic Device **204** in FIG. **2**A or Memory **370** of HMD **340**. This information may be remotely transferred to the Electronic Device **204** and/or HMD **340** from a remote system via Network Device **207** and Access Point **206**, for example. Whilst FIG. **3**A depicts a wireless interface between Electronic Device **310** and HMD **340** it would be evident that within other embodiments of the invention a wired connection may be employed discretely or in combination with a wireless interface. For example, an eSight Generation 3 HMD, as depicted in fourth image **3040**, supports a wired USB connection to a PED/FED as well as a Bluetooth connection. Within HMD **340** the Processor **380** may execute embodiments of the invention discretely, in combination with the processor of Electronic Device **310**, or through communications directly from HMD **340** or via Electronic Device **310** to one or more remote devices and/or services (e.g. a cloud based service). Similarly, the processing of image data acquired from the Camera(s) **375** may be solely within the HMD **340**, solely within the PED **310**, distributed between them, capable of executed independently upon both, or dynamically allocated according to constraints such as processor loading, battery status etc.

Accordingly, the image acquired from a Camera **375** associated with the HMD **340** may be processed by the HMD **340** directly but image data to be displayed may also be acquired from an external source directly or acquired and processed by the Electronic Device **310** for combination with that provided by the HMD **340** itself, in replacement of a predetermined portion of the image acquired by the HMD **340** or as the image(s) to be displayed to the user of the HMD **340**. Whilst, within the following description with respect to FIGS. **4**A to **12** the HMD **340** is employed in an immersive mode, e.g. no external environment is presented to the user, such as a game, virtual reality environment etc. However, it would be evident that the embodiments of the invention may also be employed within non-immersive modes of an HMD **340** such as those provide augmented reality or assistance to a user.

Within embodiments of the invention the HMD **340** and any user interface it provides directly or upon a PED, such as Electronic Device **310**, may be context aware such that the user is provided with different interfaces, software options, and configurations for example based upon factors including but not limited to cellular tower accessed, Wi-Fi/WiMAX transceiver connection, GPS location, local associated devices, dimensions of measured environment etc. Accordingly, the HMD **340** may be reconfigured upon the determined context of the user based upon the PED determined context. Optionally, the HMD **340** may determine the context itself based upon any of the preceding techniques where such features are part of the HMD **340** configuration as well as, for example, based upon processing the received image from the Camera(s) **375** and/or content being rendered upon the Display(s) **365**. For example, the HMD **340** configuration and processing according to embodiments of the invention may change according to whether the user is walking, running, crawling, riding a bicycle, driving a vehicle, etc.

Referring to FIG. **3**B there are depicted exemplary first to fifth HMDs **3010** to **3050** of HMDs which can exploit a redirection process according to an embodiment of the invention. Accordingly, these are:

First HMD **3010** Samsung Gear VR HMD;

Second HMD **3020** Sony HMZ-T3 W HMD;

Third HMD **3030** Microsoft™ Hololens HMD;

Fourth HMD **3040**, eSight Generation 3 providing immersive and non-immersive environments through a bioptic tilt of the displays relative to the frame; and

Fifth HMD **3050**, Google Cardboard which supports insertion of a smartphone for example and in conjunction with software splits the display of the smartphone into two regions each presented to only one eye of the user.

3. Re-Directed Walking

One of the most important forms of interaction in VE is locomotion. Natural walking (or even jogging, running etc.) is the most preferred (and natural) technique primarily because it provides an increased sense of presence in the VE with improved spatial understanding whilst reducing signs of VR sickness. However, the main difficulty of using natural walking as a locomotion technique in VE is the requirement that the size of the PTS be comparable in size with the VE, which is often not the case; especially for simulations involving large-scale environments. Today this is still an active area of research with a particular focus on locomotion techniques which do not carry, in any degree, the spatial constraints imposed by the physical space over to the theoretically boundless virtual space of the VE.

Accordingly, over recent years many techniques have been proposed for a user's interaction in a VE and in particular, their navigation. This related work is broadly categorized in terms of being (a) redirection, (b) steering algorithms, resetting phase and natural visual suppressions, and (c) dynamic foveated rendering.

3A. Redirection and VE

"Redirected Walking" has been an active research topic for nearly 20 years and exploits subtle rotations in the VE which are presented to the VE users so that they were tricked into walking on a curved path in the PTS while maintaining a straight path in the VE. These subtle rotations applied to the VE were enough to convince the users that they had explored a comparatively larger virtual area than the actual available play space. Since the original concept of redirected walking was proposed a number of attempts were made to achieve the same effect based on software and/or hardware. Some researchers have even tried incorporating physical props or manipulating the entire physical environment. However, these type of solutions failed to gain the mainstream attention due to their many dependencies on factors other than the actions of the user himself.

Hardware-based approaches were also explored to resolve this problem such as the omnidirectional treadmill, suspended walking, low-friction surfaces, walking in a giant hamster ball (e.g. VirtuSphere), etc. However, whilst these prototypes, although bulky, expensive and not easily established in any PTS, offer potentially infinite walking in VE within a defined PTS they all lack inertial force feedback. For this reason, natural walking is considered to be the most preferred and natural way for locomotion in VE. Moreover, the multimodal nature of natural walking allows free user movements such as jumping or crouching as well as variable speeds from walking through to running.

Software-based techniques have also been proposed for achieving the same effect by solely manipulating the VE and are generally divided into two groups, namely those techniques that use the user's head rotations and translations for scaling the VE dynamically based on the scenario and those techniques that partially or fully warp the virtual environment. Due to the dominance of the visual sense over the other human senses, these techniques focus mainly on reducing the effect of subtle visual distractions resulting from repeated redirection. These visual distractions are mainly created during the naturally occurring or stimulated visual suppressions such as a blink or a saccade. However, a disadvantage of these techniques is the fact that they are disruptive to the cognitive task at hand, since they rely on stimulating saccades by introducing artifacts in the rendered image to distract the user's attention.

3B. Steering Algorithms, Resetting Phase and Natural Visual Suppressions

Considering initially steering algorithms then in order to calculate the amount of redirection, two parameters are required, namely, the target direction of the user (a) in the VE, and (b) in the PTS. There are many methods for predicting the target direction in VE ranging from using the user's past walking direction, to head rotations, and gaze direction. As there are spatial constraints in terms of the available PTS, the user must be steered away from the boundaries of the PTS, typically a wall for example. To locate the user's target direction a variety of techniques and related algorithms have been proposed and assessed including, but not limited to, steer-to-center, steer-to-orbit, and steer-to-multiple-targets. The first steers the user towards the center of the PTS whilst steer-to-orbit steers them towards an orbit around the center of the PTS and steer-to-multiple-

targets steers the users towards several assigned waypoints in the PTS. Extensions of these have included, the steer-to-multiple-centers algorithm, limiting the directions of walking by virtual objects in the VE to enhance the steer-to-orbit algorithm, and predetermined curves inside the VE where a user was allowed to change their direction of walking only when they reached the intersection of their current curve with another predetermined curve.

Within the following descriptions with respect to embodiments of the invention and the experimental results of user studies etc. the inventors have employed the steer-to-center algorithm for redirecting the user towards center of the PTS when a collision is predicted to occur. However, it would be evident that the methodologies presented, and embodiments of the invention may be employed with other steering algorithms, etc. without departing from the scope of the invention. It would also be evident to one of skill in the art that the steering algorithm etc. employed may be defined for a VE system, for a VE session, based upon the PTS, or be dynamically established based upon aspects of one or more of the VE session, PTS and the user. For example, the steering mechanism may vary during a VE session, for example through the addition of a second user. Accordingly, the scope of the invention is not limited by the embodiments of the invention presented below but by the scope of the claims.

With respect to the resetting phase then this is an important aspect of all redirected walking techniques because there is always a small possibility of the user crossing over the boundary of the PTS, e.g. reaching a physical boundary such as a wall or walking into a physical object within the PTS. If this occurs, the user has to be stopped and their position has to be reset before starting to walk again, hence the term re-set phase. Several solutions have been proposed to address this problem including, but not limited to, "Freeze-Turn" where the field-of-view (FoV) of the user remains unchanged whilst they turn in the PTS, "Freeze-Back-up" where the FoV remains unchanged whilst the user backs-up from the boundary to a point motion can resume, and "Turn" where a multiplier of the rotational gain, e.g. x2, is applied to the user's head rotations, i.e. if the user turns by 180°, a rotation of 360° is applied so that the user is again facing the same direction that they were facing before. Visual effects which may result from the resetting can be addressed with the use of visual distractors.

Within the following descriptions with respect to embodiments of the invention and the experimental results of user studies etc. the inventors have employed the "stop-and-go" paradigm where if the user crosses over a boundary, then before they start walking again, they have to perform an in-situ rotation towards a direction where there is no obstacle. However, it would be evident that that the methodologies presented, and embodiments of the invention may be employed with other resetting phase processes, etc. without departing from the scope of the invention. It would also be evident to one of skill in the art that the reset employed may be defined for a VE system, for a VE session, based upon the PTS, or be dynamically established based upon aspects of one or more of the VE session, PTS and the user. For example, the reset mechanism may vary during a VE session, for example through the addition of a second user. Accordingly, the scope of the invention is not limited by the embodiments of the invention presented below but by the scope of the claims.

The human visual system of even a healthy individual with 20/20 vision (a term used to express normal visual acuity (the clarity or sharpness of vision) measured at a distance of 20 feet) is not perfect. Due to several very frequent involuntary actions, humans face temporary blindness for short periods of time, called visual suppressions. Saccades, eyeblinks, the phase of nystagmus, and vergence movements are some of the involuntary visual suppressions. Saccades are the brief rapid eye movements that occur when we quickly glance from one object to another whilst an eye-blink is a rapid opening and closing of the eyelids, where these eye movement can occur either voluntarily, involuntarily or as a reflex. The phase of nystagmus is a condition where uncontrolled rapid eye movements occur from side-to-side, top-to-bottom or in circular motion whilst vergence movement occurs to focus on the objects with different depths, one after the other.

Within the following only techniques employing saccades and eye-blinks are reviewed as embodiments of the invention may also exploit these. Saccades are extremely unpredictable, rapid, and ballistic eye movements that occur when we abruptly shift our fixation point from one object to another. The visual suppression occurs before, during, and after the saccadic movement and can last for 20-200 ms whilst the speed of these saccadic movements can reach up to 900°/s. As a saccade occurs very frequently and can last for several frames (considering typical frame rates of standard and high definition systems) which makes it possible to render an updated VE without the user noticing. In contrast to the saccades, blinks occur much less frequently and are much slower thereby giving more flexibility for reorientation due to the longer induced change blindness. Depending upon the scenario, one blink can give the users a temporary blindness of 100-400 ms which is much longer than a saccade. This also makes them easier to detect even with readily available commercial eye trackers. Similar to saccades, our visual perception is also suppressed before, during, and after the opening and closing movements of the eyelids. Studies have shown that the average person blinks at an average rate of 17 times per minute.

Within the following descriptions with respect to embodiments of the invention and the experimental results of user studies etc. the inventors have exploited this physiological phenomenon to refresh the foveal zone render and therefore redirect the user multiple times per minute during blinks and/or minor saccades. Additionally, embodiments of the invention exploit reported data for the maximum rotational and translational thresholds for VE during blinks and saccades to update the VE and refresh the render without the user perceiving anything. However, it would be evident that that the methodologies presented, and embodiments of the invention may be employed with other natural vision suppressions according to the capabilities of the HMD and/or associated software, firmware etc. For example, improved eye-tracking may allow detection of minor and/or major saccades allowing partial or full updating during these. For example, partial updating may update a portion of the foveal region and/or peripheral region. It would also be evident to one of skill in the art that the natural vision suppression(s) employed may be defined for a VE system, for a VE session, based upon the PTS, or be dynamically established based upon aspects of one or more of the VE session, PTS and the user. For example, the natural vision suppression may vary during a VE session, for example through the addition of a second user, detection of a change in a physiological aspect of the user, or based upon characterization of one or more vision degradations and/or defects of the user. Accordingly, the scope of the invention is not limited by the embodiments of the invention presented below but by the scope of the claims.

3C. Dynamic Foveated Rendering

This technique can significantly reduce the overall workload on a microprocessor, e.g. a central processor unit (CPU) or dedicated graphics processor unit (GPU), while providing the same VE experiences to the user. Foveated rendering leverages the fact that small changes occurring in our peripheral vision are imperceptible to us. Thus, the area of the image corresponding to the peripheral vision can be rendered at a much lower resolution while the area of the image corresponding to the foveal vision is rendered at full resolution. Solutions to this have been established with different software-based techniques for simulating perceptually guided foveated rendering as well as hardware solutions

Within the following descriptions with respect to embodiments of the invention and the experimental results of user studies etc. the inventors have exploited a hardware based solution, Nvidia's Variable Rate Shading (VRS) such that foveated rendering was supported in hardware and integrated into the rendering pipeline. However, it would be evident that that the methodologies presented, and embodiments of the invention may be employed with foveated rendering techniques exploiting other hardware and/or software based solutions according to the capabilities of the HMD and/or associated software, firmware etc. For example, embodiments of the invention may exploit hardware based foveated rendering, software based foveated rendering and/or a combination of hardware and software based foveated rendering. For example, software based foveated rendering may exploit approaches to process the image data in multiple files to parallel pipelines and process foveal content differently to peripheral content, or an effective frame-rate of presenting peripheral content relative to foveal content may be lower. It would also be evident to one of skill in the art that the foveated rendering employed may be defined for a VE system, for a VE session, based upon the PTS, or be dynamically established based upon aspects of one or more of the VE session, PTS and the user. For example, the foveated rendering may vary during a VE session, for example through the addition of a second user, detection of a change in a physiological aspect of the user, or based upon characterization of one or more vision degradations and/or defects of the user. Accordingly, the scope of the invention is not limited by the embodiments of the invention presented below but by the scope of the claims.

It would also be evident from the following descriptions with respect to embodiments of the invention that the inventors have exploited the Nvidia's VRS not only for reducing the overall GPU workload but also for blending foveal and non-foveal (peripheral) zones rendered from two co-located cameras, respectively. However, the inventors note that within the prior art software based solutions for the merging/blending/overlay of content from multiple sources exist within other HMD applications such as for augmented reality etc. Accordingly, the scope of the invention is not limited by the embodiments of the invention presented below but by the scope of the claims.

4. Technical Overview

Referring to FIG. **4**A there is depicted an exemplary image depicting orientation-based redirection according to an embodiment of the invention in conjunction with dynamic foveated rendering which leverages the effect of inattentional blindness induced by a cognitive task. The blending (e.g. parallax of green Alien **440**) in the Transition Zone **420** between the rotated non-foveal virtual environment (Non-Foveal VE **430**) and the non-rotated foveal (Foveal VE **410**) is imperceptible by the users due to inattentional blindness. The angular gain for the frame as shown in FIG. **4**A was at the maximum employed, i.e. 13.5°.

Accordingly, a user exploiting an embodiment of the invention may traverse as depicted in FIG. **4**B a path within the VE (VE Path **460**) which is significantly larger than the Physical Tracked Space (PTS **450**) where through re-direction they execute a PTS Path **470**. The PTS **450** is a 4 meter×4 meter (approximately 13 feet by 13 feet). The Camera **1480** icon inside the box indicates the location of the user with respect to the PTS **450**. The PTS **450** represents a room-scale PTS wherein users were able to walk distances in the VE which were substantially larger. The VE Path **460** depicted is 58.51 meters (approximately 192 feet) without a reset which is 15 times the length or width of the PTS **450**. The longest distance recorded in experiments for this PTS **450** with prototype embodiments of the invention was 103.9 meters (approximately 341 feet) or 26 times the length or width of the PTS **450**.

Now referring to FIG. **5** there is depicted a pipeline of an embodiment of the invention. Two cameras, Cam_{Foveal} and $Cam_{Non\text{-}Foveal}$, are employed within this exemplary pipeline to render the VE to the user. Based upon results of the first user study outlined below the inventors determined that the appropriate field-of-view, S, for the foveal region was 60° within the image acquired by Cam_{Foveal} and that the rotation angle applied to the non-foveal image. $Cam_{Non\text{-}Foveal}$, had a maximum value of 13.5°, i.e. $0°<\theta<13.5°$. Accordingly, a first Mask **520** corresponding to $\delta=60°$ was applied to the Foveal Image **510**, i.e. the image of Cam_{Foveal}, in First Process **580** to generate Foveal Masked Render **550**. In Second Process **585** an inverse of the first Mask **520**, second Mask **540** generated by an inversion process **575**, is applied to the Non-Foveal Image **530**, i.e. the image of $Cam_{Non\text{-}Foveal}$ which is rotated relative to Cam_{Foveal} by θ, to generate Non-Foveal Masked Render **560**. Then in Third Process **590** the Foveal Masked Render **550** and Non-Foveal Masked Render **560** are combined to generate the Rendered Image **570** presented to the user upon their HMD. Accordingly, the rotation applied between the Foveal Image **510** and Non-Foveal Image **530** is evident in FIG. **5**. As will be evident from the results of the second and third user studies the users fail to perceive any visual distractions or artifacts in the final composite render, Rendered Image **570**, while they are preoccupied with a cognitive task; which is almost always the case in VE applications.

5. User Study 1 Determining Maximum Rotation Angle and Field of View of the Foveal Image Rendered

The efficacy of redirected walking is coupled with the user's perception of the redirection taking place. In the inventor's first user-study they determined the maximum angle for which the rotation of the non-foveal zone (i.e. the area in the rendered image corresponding to the peripheral vision) remained imperceptible to the user.

5A. Application

The inventor's designed an immersive VE application using the HTC Vive Pro Eye HMD with an integrated Tobii Eye Tracker. The application depicted a serene urban city environment in which red spheres were hidden at random locations. The environment was developed in Unity3D and foveated rendering was supported using an NVIDIA RTX 2080Ti graphics card. Three zones were specified for the foveated rendering, foveal zone, non-foveal zone and transition zone.

5A.1. Foveal Zone

The foveal zone was defined as the circular area in the rendered image centered at the fixation point captured by the eye tracker. For rendering, this zone should have the highest

quality since this is where the user's foveal vision is focused. Within the following description with respect to an embodiment of the invention the pixels in the foveal zone were rendered with a 1:1 sampling was employed although other samplings may be employed without departing from the scope of the invention. It would be evident that within other embodiments of the invention the foveal zone may be defined by another geometry, e.g. elliptical, or that the foveal zone may be defined in dependence upon characterisation of the user's foveal zone where the user has visual defects or vision degradations. This assessment may also determine that the user has a visual axis of vision offset relative to that of the fovea-retina in order to overcome a visual defect or vision degradation which may be factored into the calibration and settings of the eye-tracker.

5A.2. Non-Foveal Zone

The non-foveal zone is the area in the rendered image which corresponds to the peripheral vision. This being complementary to the foveal zone in the exemplary pipeline of FIG. **5** although it would be evident that this represents one specific embodiment of the invention and that other non-foveal zones may be defined with or without specific reference to the foveal zone. This zone is of lower importance than the foveal zone since it is not the main focus of the user's vision. Hence, pixels in the non-foveal zone are rendered with a different sampling to the foveal zone. Within the following description with respect to an embodiment of the invention a 16:1 sampling was employed although other samplings may be employed without departing from the scope of the invention.

5A.3. Transition Zone

The transition zone is an overlapping or non-overlapping area as the rendered image transitions from the foveal zone to the non-foveal zone. This zone was introduced by the inventors in order to avoid a sharp boundary edge between the foveal and non-foveal zones which may be perceived by the user. Within the following description with respect to an embodiment of the invention the transition zone was rendered at a constant sampling of 4:1 between that of the foveal zone and the non-foveal zone. However, it would be evident that within other embodiments of the invention the transition zone may employ graded sampling between that of the foveal zone and the non-foveal zone where the transition function may be, for example, linear or Gaussian although other functions may be employed either mathematically or non-mathematically defined.

FIG. **6** depicts a frame from the application with the three zones annotated. The visible frame within the VE application was panoramic and significantly larger than the image actually shown which depicts only part of the frame relevant to the discussion, The Foveal Zone **610** corresponding to a field-of-view is marked with a solid circle and was rendered at the highest quality with 1:1 sampling and no rotation. The Non-Foveal Zone **630** was rendered at a lower resolution with 16:1 sampling with a rotation of $\theta_{NON\text{-}FOVEAL}$. The Transition Zone **620** is depicted as a dashed circle around the Foveal Zone **610** as was rendered at a 4:1 sampling. This indicates the overlapping area between the foveal and non-foveal zones for which alpha-blending was performed at each pixel using Equation (1), where $C_{BLENDED}$ is the final colour of any pixel in the transition zone, α_{FOVEAL} is the alpha value for the foveal image which changes from 1 to 0 from the foveal to the non-foveal zones, and C_{FOVEAL} and $C_{NON\text{-}FOVEAL}$ are the foveal and non-foveal colours at a pixel respectively.

$$C_{BLENDED}=\alpha_{FOVEAL}\times C_{FOVEAL}+(1-\alpha_{FOVEAL}\times C_{FOVEAL} \quad (1)$$

This requires two co-located cameras (or images representing images acquired by two co-located cameras) which have a rotational offset between them. Within the following embodiments of the invention this rotational offset was not fixed for all user studies. Within the first user study this rotational offset was increased until the users either saw artifacts in their field of view of felt that there were artifacts in their field of view. This was to measure the minimum field of view for the foveal zone for which the rotational offset between the two collocated cameras is maximized. Thereafter, in all user studies we increased the rotational offset at a constant rate of 6°/s from 0° to 13.5°. However, within other embodiments of the invention the rotational offset may be determined by one or more factors including, but not limited to, the PTS, the VE, a VE-PTS scaling, and the user. Optionally, a calibration routine may be executed by the HMD upon initial use by a user to assess the maximum rotation angle etc. perceived or imperceptible to the user. This calibration routine may also establish other parameters of the rendering including, for example, the shape and/or size of the Foveal Zone, the shape and/or size of the Transition Zone, and the rendering within the Transition Zone (e.g. linear, constant, non-linear etc.).

Accordingly, for the two co-located cameras employed to render the foveal and non-foveal frames the color values are retrieved. The boundary values [0.0, 1.0] are also shown in FIG. **6**. These coincide with the boundaries of the transition zone. The field-of-view $\delta_{TRANSITION}$ corresponding to the transition zone was defined empirically as an offset to the field-of-view of the foveal zone of +40°.

5B. Procedure

The premise of the inventor's invention is inattentional blindness which implies that the user's attention should be directed towards a cognitive task. Accordingly, the inventors instructed the participants to perform a target-retrieval task. More specifically, the participants were asked to search and count for red spheres hidden in the VE. At each iteration of the experiment, the red spheres were hidden at random locations. This was done in order to eliminate the possible bias that may be introduced by memorizing the locations between iterations.

The first user study involved 11 participants (2 females, 18.2%). Average age was 24.64 with a SD of 2.8. Median of their reported experiences with using VE devices was 3, and the median of their experiences with using an eye tracking device was also 3 on a 5-point Likert scale, with 1 being least familiar, and 5 being most familiar. The participants performed this task from a seated position and were only allowed to rotate their chair in-place. The participants were instructed to press the trigger button on the HMD controller if and when they noticed a visual distortion or felt nausea due to simulator sickness. During the experiment, the rotation angle of the non-foveal zone (θ) was gradually increased and grouped by increasing order of the field-of-view δ of the foveal zone i.e. ([(δ_1, θ_1), (δ_1, θ_2), . . . , (δ_1, θ_n), (δ_2, θ_1), . . . , (δ_2, θ_n), . . .]). Each time the trigger was pressed the (δ_i, θ_i) was recorded, and then the experiment continued with the now increased field of view δ_{i+1} and a reinitialized rotation angle θ_1.

The range of values for the field-of-view assessed was 20° to 60°. The step of each increment was 10° after the completion of one cycle of the rotation angle, or until triggered by the user. During a cycle, the rotation angle ranged from 0° to 15° and the step of each increment was 1° per second.

Preliminary experiments during the design of the application had shown that repeated increments of the field-of-

view of the foveal zone can lead to nausea and severe dizziness. For this reason, the participants were instructed to take a short break after each cycle of increments of the field-of-view. Furthermore, the sequence of the cycles i.e. the field-of-view values, was randomized for each participant in order to eliminate any bias. FIG. **6** shows the view from the user's perspective during the experiment.

5C. Analysis of Results

The results for the first study are shown in FIG. **7**. A cycle of the rotation angle as performed for each field of view. The results show that as the field-of-view δ increases the tolerance for higher rotation angle also increases, which can also be confirmed by the exponential trendlines shown for each participant. For the reasons mentioned above, we select the smallest rotation angle for which users did not perceive a change associated with the largest field-of-view for which the majority of the users did not perceive a change (i.e. 9 out of 11). Thus, the ideal pair values for (δ, θ) was determined to be (60°, 13.5°); where 13:5° is the maximum allowed rotation angle.

5C.1. Simulator Sickness Questionnaire (SSQ)

Upon completing the experiment all participants were asked to complete Kennedy Lane's Simulation Sickness Questionnaire (SSQ). The Total Severity (TS) and the subscales Nausea, Oculomotor, and Disorientation were calculated using the formulas from Kennedy et al. in "Simulator Sickness Questionnaire: An Enhanced Method for Quantifying Simulator Sickness" (Int. J. Aviation Psychol. Vol 3, No. 2, pp 203-220; hereinafter Kennedy1). Based on the SSQ categorization provided by Kennedy et al. in "Configural Scoring of Simulator Sickness, Cybersickness and Space Adaptation Syndrome: Similarities and Differences" (Virtual and Adaptive Environments: Applications, Implications, and Human Performance Issues, Taylor and Francis, 2003, p. 247; hereinafter Kennedy2), 55% of the participants reported no signs (TS=0) or minimal signs (TS<10) of simulator sickness. All the participants completed the test, with 0 dropouts. Upon further analysis, the disorientation sub-scale had the highest average score of 29.11 with a maximum score of 139.2. This was expected, considering the fact that the rotation angle was constantly increasing and thus the VE rendered in the HMD induces conflicts between the vestibular and visual signals, leading to vestibular disturbances such as vertigo or dizziness. The results from SSQ responses are summarized in Table 1 and FIG. **8**.

TABLE 1

Results from the responses of SSQ for User Study #1. The Total Severity (TS) and the corresponding sub-scales such as Nausea, Oculomotor, and Disorientation were calculated using the formulas of Kennedy1.

Score	Mean	Median	Standard Deviation	Minimum	Maximum
Nausea (N)	11.27	9.54	14.67	0	38.16
Oculomotor (O)	19.29	15.16	24.06	0	83.38
Disorientation (D)	29.11	0	45.09	0	139.2
Total Score (TS)	21.76	7.48	28.47	0	93.5

5D. Discussion

The design of the application of user study #1 involved making a decision on the range of values for (a) the field-of-view δ, and (b) the rotation angle θ of the non-foveal zone.

Range of δ

Humans have a maximum horizontal field-of-view of about 210°. This is further reduced to 110° by the HMD hardware i.e. maximum field-of-view for HTC VivePro.

Based upon the default parameters established by the inventors the foveal and non-foveal zones are complementary to each other. Thus, there is a tradeoff between the sizes of the foveal and non-foveal zones. If δ is large, then the foveal zone is large, and the non-foveal zone is small. Having a small non-foveal zone means that only a small area of the final composite image will be rendered from the non-foveal camera as shown in FIG. **5**, leading to smaller possible redirections. When δ=60° the foveal zone occupies 54.55% of the final composite render, based upon the defined parameters, HMD etc., and the non-foveal zone occupies 45.45% (including a transition zone of 35.45%). Similarly, if δ=90° the foveal zone occupies 90.91% of the final composite render which does not leave much for the non-foveal zone. In contrast, if δ is small, then the foveal zone is small, and the non-foveal zone is large. Although this allows for larger redirections, the inventors found in their preliminary tests that when δ<20° it can cause severe nausea and simulation sickness. For these reasons the inventors selected the range of values $\delta \in [20°, 60°]$ which balances the sizes of the foveal and non-foveal zones, and is large enough that it does not cause user discomfort.

Range of θ

Recent upon experiments reported in the prior art it has been shown that users cannot tolerate a rotational angle of more than 12.6° in their field-of-view during a saccade having a velocity of 180°/sec. Based upon this the inventors selected the range $\theta \in [0°, 15°]$.

6. User Study 2 In-Situ Gaze Redirection Using Dynamic Foveated Rendering

The objectives of the second user study were twofold. Firstly, to determine whether a rotation of the VE by an angle below the maximum (i.e. θ<13.5°) is indeed imperceptible and does not cause simulation sickness. Secondly, to experimentally prove with quantitative measures that using the proposed redirection technique with gaze-only (and without walking) the rotation of the VE results in the equivalent redirection of the participant in the PTS.

6A. Application and Procedure.

An experiment was devised similar to the one in Section 4. In contrast to the user study #1, the participants were only allowed to spin in-situ from a standing position. This change from a seated to a standing position eliminates the possibility of the participants perceiving any orientation and navigation cues coming from the orientation of the chair. The participants were given a target-retrieval task and instructed to retrieve, by directing their gaze to, as many virtual targets (i.e. orange pumpkins) as possible. The virtual targets disappeared as soon as the gaze direction intersected their bounding box. The positions of the targets were randomized for each participant.

The duration of the experiment was 60 seconds. Unbeknownst to the participants, in the first 30 seconds there was no redirection applied. This served as a baseline for participants who had little to no experience in using HMDs. Once accustomed to the VE, during the following 30 seconds the rotation angle of the VE was increased at a rate of 6°/s. Hence, the hypothesis is that after 30 seconds of a constant smooth rotation of the VE at a moderate rate of 6°/s the participant should face 180° away from their initial orientation, i.e. the opposite direction. To prove this, the initial (i.e. at time=0 s) and the final (i.e. at time=60 s) gaze directions of each participant were recorded. Additionally,

before each participant removed the HMD at the end of the experiment they were asked to face towards what they believed to be their initial directions in the PTS using visual landmarks from within the VE to orient themselves.

6B. Analysis of Results

The study involved 11 participants with average age of 26.27±3.13. Based on a 5-point Likert scale the medians of their experience with using VE or any other eye tracking devices were 3. Five of the participants had not taken part in user study #1 (Participants #2, #3, #6, #9, #11). After the experiment, the participants completed the SSQ. The angle between the initial and final gaze directions was calculated for each participant. The average deviation was 171.26° (4.77 standard deviation) which means that the participants thought that their initial orientation was towards the opposite direction. In fact, all participants reported that they did not perceive the redirection and were surprised by how off their "sensed" orientations were.

6B.1. SSQ

Based on the scores reported by the participants in the post-test SSQ, the majority of the participants (55%) showed no signs (TS=0) or minimal signs (TS<10) of simulator sickness. The highest score and average were reported for the sub-scale disorientation although reduced by a factor of 2 from user study #1. This was anticipated since the rotation angle was less than the maximum determined from user study #1. As it can be seen from FIG. **9**, one of the participants (#2) had no previous experience with VE and reported perceptual anomalies including difficulty concentrating, fullness of head and difficulty focusing. The results for the SSQ are summarized in Table 2.

TABLE 2

Results from the responses of SSQ for User Study #2.

Score	Mean	Median	Standard Deviation	Minimum	Maximum
Nausea (N)	10.41	0	15.06	0	47.7
Oculomotor (O)	8.27	0	12.43	0	37.9
Disorientation (D)	11.39	0	17.41	0	55.68
Total Score (TS)	11.22	7.48	15.78	0	52.36

7. User Study #3 Re-Directed Walking Using Dynamic Foveated Rendering

The objective of the third user study is to evaluate the efficacy of redirected walking during inattentional blindness using dynamic foveated rendering.

7A. Application

As discussed above, inattentional blindness refers to the inability of an individual to see a salient object in plain sight due to lack of attention. This is true for the majority of the VE applications where the user is actively engaged and preoccupied with a cognitive task e.g. games, training simulations, etc. Thus, for the purposes of this user study, the inventors designed a first-person VE game where the objective is to stop an alien invasion. To achieve this the user has to walk through a deserted urban city to a predetermined location indicated by a large purple orb and destroy the alien-mothership (appearing in the form of a giant brain) while zapping green aliens along the way. Zapping one alien will award one score point to the player. The green alien enemies are randomly spawned (and are therefore independent of the orientation of the current redirection) only within the field-of-view of the user while also making a sound effect. An example of in-game gameplay is shown in FIG. **11**.

The shortest distance the participants had to travel in the VE was 42 meters (approximately feet) while the available PTS had a size of 4×4 meters (approximately 13 feet by 13 feet). In the first to third user results depicted in FIGS. **10**A to **10**C respectively the PTS is depicted as a white colored square box and the position of the user with respect to the PTS is indicated with the camera icon in each instance. For safety reasons, a resetting mechanism of 2:1 was implemented. In cases where the game predicts that the user is about to cross over a boundary of the PTS, it would pause and prompt the user to rotate in-situ by 180°. During the user's rotation, the VE was also rotated by the same angle but in the opposite direction. The user was then allowed to proceed with playing the game.

Redirection was primarily performed by blending in real-time the foveal and non-foveal renders. Furthermore, redirection was also performed during the tracked naturally occurring blinks and saccades. In contrast to the prior art the methodology does not stimulate saccades nor blinks since these are disruptive to the cognitive task at hand.

7B. Procedure

In order to evaluate the efficacy of our technique, a controlled experiment was conducted where the independent variable being tested is the proposed redirection technique. The participants were instructed to complete the objective of the game twice; the first time with the experimental condition i.e. with redirected walking, and after a short break a second time with the control condition i.e. without redirected walking. For the experiment with the control condition, participants had to navigate themselves to the destination by relying solely on the resetting mechanism every time they went out-of-bounds from the available PTS.

A sample size estimation with an effect size of 0.25 showed that a total of 25 participants were required for the experiment. All the participants were randomly chosen (12% female, average age of 25.88 years with a standard deviation of 3.06). Based on a 5-point Likert Scale, the median of their experiences using VE headsets or any other eye tracking devices was 3.

Before the experiment, participants were briefed on their objective. Instructions were also given on how the resetting mechanism works in case they are prompted with an out-of-bounds warning and are required to reset their orientation. Moreover, they were instructed to walk at a normal pace which will allow them to complete the task along the way. Once both the objectives were completed, participants were also asked to complete the SSQ. Furthermore, for the experimental condition, at the end of the first experiment the participants were asked "Did you feel the redirection or any other scene or camera modulation during the experience?".

7C. Analysis of Results

A one-way between groups analysis of variable (ANOVA) (α=0.05) with repeated measures was performed to compare the effects of with- and without-using the redirection on the dependent variables; (a) number of resets, (b) distance traveled in PTS, (c) total time taken, and (d) scores. The inventors used partial eta squared (η_p^2) to report the obtained effect sizes for each variable.

Based on the results of Levene's test, it was found that the outcomes for the number of resets ($F(2:14)=375.710$; $p>0.05$), distance traveled ($F(1:176)=0:348$; $p>0.05$) and total time taken ($F(0:971)=1:001$; $p>0.05$) were normally distributed and hence equal variances were assumed. However, the outcome for scores ($F(4:103)=0.054$; $p<0.05$) showed the opposite. As scores violated the homogeneity of variances assumption, the variable was omitted during the ANOVA analysis.

The results from ANOVA showed a statistically significant difference between the number of resets when the game was played with- and without-using the proposed redirection technique ($F(1; 48)=375.710$; $p<0.001$) with $\eta_p^2=0.887$. Nonetheless, these results also showed

a statistically insignificant effect of redirection on distance traveled ($F(1; 48)=0.384$; $p>0.05$; $\eta_p^2=0.008$), and total time taken ($F(1; 48)=1.001$; $p>0.05$; $\eta_p^2=0.020$). The η_p^2 values shows that 88.7% of the variability in the required number of resets is accounted for by our independent variable i.e.—redirection. However, the effects on distance traveled and total time taken remains negligible. The results from this test can be seen in FIGS. **12**A to **12**C respectively for resets, distance travelled, and time taken respectively. The error bars in the graphs of Figurers **12**A to **12**C respectively show a confidence interval of 95%.

Referring to FIGS. **10**A to **10**C there are depicted VE and PTS paths for three users during the third study with:

FIG. **10**A shows a distance of 38.56 meters (approximately 126.5 feet) without any resets representing a VE distance approximately 10 times the PTS dimension;

FIG. **10**B shows a distance of 54.2 meters (approximately 178 feet) without any resets representing a VE distance approximately 14 times the PTS dimension; and

FIG. **10**C shows a distance of 36 meters (approximately 118 feet) without any reset representing a VE distance approximately 9 times the PTS dimension.

Besides this, results of the first experiment also showed that the system applied an average of 1547.55° (standard deviation 152.26°) of absolute angular gain to each participant's orientation during the entire test. An average of 3.15° of absolute angular gain was applied per redirection with an average of 1.73 redirections/s. As the participants were cognitively preoccupied with the task of zapping aliens, they were unaware of this angular gain.

Furthermore, since this is a real-time technique and thus the frames were rendered without additional lagging other than what is typically imposed by the hardware, none of the participants reported perceiving any scene nor camera manipulation. In the post-study questionnaire, one of the participant stated that "I felt like I was walking straight. I was completely unaware of my actual movements."

7C.1. SSQ

After completing the experiment, participants were asked to fill out an SSQ. Based on the scores reported, the majority of the participants (80%) showed no signs (TS=0) or minimal signs (TS<10) of simulator sickness, and only 8% of the participants reported TS>12. The highest score and mean were reported for the disorientation sub-scale even though the rotation angle was at all times well within the limits of tolerance as determined by user study #1. This can be attributed to the fact that the cognitive workload of the task involved in this user study was more demanding than in the previous user studies. Although this caused an increase in the highest score for disorientation, the mean decreased when compared to that of user study #2. The median values for all sub-scales, as before, were reported as 0. Table 3 summarizes the results from this SSQ. As it is evident, the mean scores were dropped significantly from user study #1 and #2.

TABLE 3

Results from the responses of SSQ for User Study #2.

Score	Mean	Median	Standard Deviation	Minimum	Maximum
Nausea (N)	2.29	0	5.7	0	19.08
Oculomotor (O)	6.67	7.58	9.1	0	37.9
Disorientation (D)	8.35	0	17.05	0	69.6
Total Score (TS)	6.43	3.74	10.04	0	44.88

8. Review

The results of the user study #3 are indicative of the efficacy of the proposed technique in VE applications where the cognitive workload on the user is moderate. Examples of such applications are immersive games, training simulations, cinematic VE, etc. Further to exploiting inattentional blindness, and unlike other state-of-the-art techniques, the inventor's innovative methodology relies only on naturally occurring saccades and blinks, and not stimulated saccades and blinks.

The inventive methodology is distinct from other reported work in the literature and an important advantage since stimulating saccades is both, disruptive to the cognitive task at hand and increases the effects of VE sickness. For example, within the prior art saccades have been stimulated by introducing orbs of light in image- and object-space to forcefully divert the user's attention in order to perform the redirection. In contrast to this, and based on the psychological effect of inattentional blindness, rather than divert the user's attention we exploit the fact that the user is fixated on a particular task leading to a "tunnel-vision""-like focus. This allows the inventors to constantly update in real-time the non-foveal (peripheral) zone without the user perceiving a change. Metaphorically speaking, the foveal vision/zone acts as an update-brush of the framebuffer: whenever it moves, based on the tracking of the user's eyes, everything within the foveal zone is rendered without any rotations being applied, and everything outside i.e. the non-foveal zone, is rendered applied to the VE calculated in real-time based on the required redirection.

The experiment in user-study #3 used a PTS of 4×4 meters (approximately 13 feet square). The results show that even with a typical room-scale PTS such as the one used, the users were able to walk distances in the VE which were significantly larger. The longest distance recorded in experiments for this PTS **450** with prototype embodiments of the invention was 103.9 meters (approximately 341 feet) or 26 times the length or width of the PTS **450** or 18 times the diagonal of the PTS without a reset. Furthermore, the traveled distance can include long straight walks as shown in FIGS. **4**A and **10**A respectively.

Accordingly, the inventors have presented a rotation-based redirection technique using dynamic-foveated rendering which leverages the effect of inattentional blindness induced by a cognitive task. The technique uses natural visual suppressions such as eye blinks and saccades (without any artificial stimuli) to make subtle rotations to the VE without the user's knowledge. Furthermore, we conducted extensive tests and presented the results of three user studies. The results confirmed that the technique is indeed effective and can handle long-straight walks. This allows the users to freely explore open world VEs.

The embodiments of the invention presented above only use rotational gains for redirection. However, within other embodiments of the invention translational gains may be incorporated whilst maintaining the real-time performance.

Within other embodiments of the invention redirected walking systems may employ a saccade prediction algorithm(s).

Within other embodiments of the invention redirected walking systems may employ other forms of visual suppression e.g. phase of nystagmus, etc.

Within the embodiments of the invention the embodiments of the invention have been presented from the viewpoint of allowing a user to achieve a larger VE distance of travel than that of the PTS they are within. However, it would be evident that the redirection of motion according to embodiments of the invention may also be employed to redirect motion of a user within a PTS where the user is employing a HMD wherein the images rendered to the user are now not of a VE but their physical environment. These physical environment images may be processed to address visual defect(s) and/or vision degradation(s) of the user. Accordingly, a user with visual impairments wearing an HMD may thereby have their physical motion redirected to avoid their impacting elements within their physical environment. Accordingly, the HMD may through structured optical systems, lidar, sonar, radar, etc. establish a map of objects/surfaces etc. For example, such a redirected walking enhanced HMD may be used by the elderly with degraded vision to keep them mobile with reduced risk of injury/incident etc.

9. Minor Saccade Vision Suppression and Re-Directed Walking

Referring to FIG. **13** there are depicted first to sixth images **1300**A to **1300**F respectively to a process of redirected walking based upon executing a series of renderings that exploit minor saccades to sequentially shift the foveal region presented to a user such that they redirect their walking as a result. Referring to these then first image **1300**A depicts concentric circles referring to the foveal and transitional zones respectively at time "t" whereas the shaded concentric circles in second to sixth image **1300**B to **1300**F respectively depict the foveal and transition zones at times "'t+n" (where n=n+1 after the detection of each minor saccade). The First Region **1310** depicts in each sequential minor saccades the foveal region which is to be replaced, Second Region **1320** the transition region which is to be replaced and the Third Region **1330** shows the replaced foveal region to new straight. Accordingly, as depicted by the 6th step (time=t+5), the entire initial foveal region (time=t) is replaced and is now aligned with the new straight. Accordingly, the process is repeated for each detected minor saccade so that every 5^{th} minor saccade the initial foveal region is replaced. If a major saccade is detected in the middle of this process then a regular redirection algorithm is followed and the entire foveal region is directly replaced with the new straight. Whilst the process described and depicted with respect to FIG. **13** is depicted as being performed in 6 steps, i.e. N=6, or 5 minor saccades it would be evident that other number of steps, N, or N−1 minor saccades may be employed, such as for example N=3, 4 or any positive integer. Repainting the foveal region may take any number of steps. Basically, when a minor saccade is detected the area of the new foveal region which was not part of the original foveal region is replaced with the new image from the new foveal region. The intersection of the new and old foveal region remains the same, i.e. the image from the old foveal region. This step is repeated for each following minor saccade until the entire old foveal region is replaced with the foveal region of the new straight (redirected).

Accordingly, based upon an embodiment of the invention a process such as that depicted in FIG. **14** may be executed. This executes such that:

1. User is oriented to original straight ahead.
2. Periphery moves to new straight ahead.
3. Minor saccade is detected and the eye moves to location x (Foveal region X). (Current foveal region can be a blend of any of these 3 regions: original foveal region, original transition region and the new straight periphery)
4. Minor saccade is detected and the eye moves to location y (Foveal region Y). (In foveal region Y, Y−X is now replaced with new straight foveal region and Y∩X remains as it is)

Step 4 is repeated until the entire original foveal region is painted with the new straight.

Accordingly, after this process the user's body is reoriented to new straight ahead so that the user's body orients from the original straight ahead to the new physical orientation by the redirection. If, during this process a major saccade is detected then a regular redirection algorithm is employed such that the entire foveal region is directly replaced with the new straight. Based upon the measurements performed by the inventors on test users it takes an average of 2.2 second to complete this cycle which is rarely noticed by the user due to the inattentional blindness induced by the cognitive task at hand within the VE. The onset of each minor saccade or major saccade being determined from the signal(s) acquired from an eye-tracker forming part of the HMD worn by the user.

Within the embodiments of the invention described and depicted above the process has been described and depicted with respect to images from two cameras although it would be evident that these are "virtual" cameras, the first representing the VE field of view as viewed by the user in their foveal region and the other representing the VE field of view as viewed by the user if they had turned by a predetermined rotational angle. This second VE field of view is rendered in the peripheral vision of the user whilst a transition region is generated to "blend" between the first and second field of views so that user is not aware of the boundary between them.

Specific details are given in the above description to provide a thorough understanding of the embodiments. However, it is understood that the embodiments may be practiced without these specific details. For example, circuits may be shown in block diagrams in order not to obscure the embodiments in unnecessary detail. In other instances, well-known circuits, processes, algorithms, structures, and techniques may be shown without unnecessary detail in order to avoid obscuring the embodiments.

Implementation of the techniques, blocks, steps and means described above may be done in various ways. For example, these techniques, blocks, steps and means may be implemented in hardware, software, or a combination thereof. For a hardware implementation, the processing units may be implemented within one or more application specific integrated circuits (ASICs), digital signal processors (DSPs), digital signal processing devices (DSPDs), programmable logic devices (PLDs), field programmable gate arrays (FPGAs), processors, controllers, micro-controllers, microprocessors, other electronic units designed to perform the functions described above and/or a combination thereof.

Also, it is noted that the embodiments may be described as a process which is depicted as a flowchart, a flow diagram, a data flow diagram, a structure diagram, or a block diagram. Although a flowchart may describe the operations

as a sequential process, many of the operations can be performed in parallel or concurrently. In addition, the order of the operations may be rearranged. A process is terminated when its operations are completed, but could have additional steps not included in the figure. A process may correspond to a method, a function, a procedure, a subroutine, a subprogram, etc. When a process corresponds to a function, its termination corresponds to a return of the function to the calling function or the main function.

Furthermore, embodiments may be implemented by hardware, software, scripting languages, firmware, middleware, microcode, hardware description languages and/or any combination thereof. When implemented in software, firmware, middleware, scripting language and/or microcode, the program code or code segments to perform the necessary tasks may be stored in a machine readable medium, such as a storage medium. A code segment or machine-executable instruction may represent a procedure, a function, a subprogram, a program, a routine, a subroutine, a module, a software package, a script, a class, or any combination of instructions, data structures and/or program statements. A code segment may be coupled to another code segment or a hardware circuit by passing and/or receiving information, data, arguments, parameters and/or memory content. Information, arguments, parameters, data, etc. may be passed, forwarded, or transmitted via any suitable means including memory sharing, message passing, token passing, network transmission, etc.

For a firmware and/or software implementation, the methodologies may be implemented with modules (e.g., procedures, functions, and so on) that perform the functions described herein. Any machine-readable medium tangibly embodying instructions may be used in implementing the methodologies described herein. For example, software codes may be stored in a memory. Memory may be implemented within the processor or external to the processor and may vary in implementation where the memory is employed in storing software codes for subsequent execution to that when the memory is employed in executing the software codes. As used herein the term "memory" refers to any type of long term, short term, volatile, nonvolatile, or other storage medium and is not to be limited to any particular type of memory or number of memories, or type of media upon which memory is stored.

Moreover, as disclosed herein, the term "storage medium" may represent one or more devices for storing data, including read only memory (ROM), random access memory (RAM), magnetic RAM, core memory, magnetic disk storage mediums, optical storage mediums, flash memory devices and/or other machine readable mediums for storing information. The term "machine-readable medium" includes, but is not limited to portable or fixed storage devices, optical storage devices, wireless channels and/or various other mediums capable of storing, containing or carrying instruction(s) and/or data.

The methodologies described herein are, in one or more embodiments, performable by a machine which includes one or more processors that accept code segments containing instructions. For any of the methods described herein, when the instructions are executed by the machine, the machine performs the method. Any machine capable of executing a set of instructions (sequential or otherwise) that specify actions to be taken by that machine are included. Thus, a typical machine may be exemplified by a typical processing system that includes one or more processors. Each processor may include one or more of a CPU, a graphics-processing unit, and a programmable DSP unit. The processing system further may include a memory subsystem including main RAM and/or a static RAM, and/or ROM. A bus subsystem may be included for communicating between the components. If the processing system requires a display, such a display may be included, e.g., a liquid crystal display (LCD). If manual data entry is required, the processing system also includes an input device such as one or more of an alphanumeric input unit such as a keyboard, a pointing control device such as a mouse, and so forth.

The memory includes machine-readable code segments (e.g. software or software code) including instructions for performing, when executed by the processing system, one of more of the methods described herein. The software may reside entirely in the memory, or may also reside, completely or at least partially, within the RAM and/or within the processor during execution thereof by the computer system. Thus, the memory and the processor also constitute a system comprising machine-readable code.

In alternative embodiments, the machine operates as a standalone device or may be connected, e.g., networked to other machines, in a networked deployment, the machine may operate in the capacity of a server or a client machine in server-client network environment, or as a peer machine in a peer-to-peer or distributed network environment. The machine may be, for example, a computer, a server, a cluster of servers, a cluster of computers, a web appliance, a distributed computing environment, a cloud computing environment, or any machine capable of executing a set of instructions (sequential or otherwise) that specify actions to be taken by that machine. The term "machine" may also be taken to include any collection of machines that individually or jointly execute a set (or multiple sets) of instructions to perform any one or more of the methodologies discussed herein.

The foregoing disclosure of the exemplary embodiments of the present invention has been presented for purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise forms disclosed. Many variations and modifications of the embodiments described herein will be apparent to one of ordinary skill in the art in light of the above disclosure. The scope of the invention is to be defined only by the claims appended hereto, and by their equivalents.

Further, in describing representative embodiments of the present invention, the specification may have presented the method and/or process of the present invention as a particular sequence of steps. However, to the extent that the method or process does not rely on the particular order of steps set forth herein, the method or process should not be limited to the particular sequence of steps described. As one of ordinary skill in the art would appreciate, other sequences of steps may be possible. Therefore, the particular order of the steps set forth in the specification should not be construed as limitations on the claims. In addition, the claims directed to the method and/or process of the present invention should not be limited to the performance of their steps in the order written, and one skilled in the art can readily appreciate that the sequences may be varied and still remain within the spirit and scope of the present invention.

What is claimed is:

1. A method of redirecting motion of a user comprising:

rendering to the user a sequence of images relating wherein each image comprises:

a foveal region representing a predetermined portion of a first field of view of the user at a current location

of the user with a first direction having a first predetermined geometry and first predetermined dimensions;

a non-foveal region representing a predetermined portion of a second field of view of the user at the current location with a second direction having a second predetermined geometry and second predetermined dimensions; and

a transition region having a predetermined third predetermined geometry and third predetermined dimensions; wherein

in a first image of the sequence of images the first direction and second direction are the same;

in a second image of the sequence of images the second direction is offset relative to the first direction by a predetermined angular offset;

in a third image of the sequence of images the first direction is set to the second direction in the second image of the sequence of images such that the first direction and second direction are the same but are now offset from the original first direction in the first image of the sequence of images by the predetermined angular offset;

the transition region blends from the foveal region to the non-foveal region according to a predetermined function; and

transitioning from the first image of the sequence of images to the second image of the sequence of images and the second image of the sequence of images to the third image of the sequence of images is established in dependence upon determining a predetermined natural event with respect to an eye of the user.

2. The method according to claim **1**, wherein

the predetermined natural event is a minor saccade, a major saccade or a blink.

3. The method according to claim **1**, wherein

the method is performed based upon natural events with respect to the eye of the user and independent of any induced saccade.

4. The method according to claim **1**, wherein

the determination of the predetermined event with respect to the eye of the user is determined in dependence upon signals from an eye tracker associated with the eye of the user.

5. The method according to claim **1**, wherein

the sequence of images result in physical motion of the user being redirected from the first direction to the second direction.

6. The method according to claim **1**, wherein

the field of view is one of a virtual environment, a physical real world environment of the user, the physical real world environment of the user augmented with additional digital content, and the physical real world environment of the user augmented with another virtual environment.

7. A method of redirecting motion of a user comprising:

rendering to the user a sequence of images relating wherein each image comprises:

a foveal region having a first predetermined geometry and first predetermined dimensions;

a non-foveal region having a second predetermined geometry and second predetermined dimensions; and

a transition region having a predetermined third predetermined geometry and third predetermined dimensions; wherein

in a first image of the sequence of images the foveal region is a first predetermined portion of a first field of view of the user at a current location of the user with a first direction;

in a last image of the sequence of images the foveal region is a first predetermined portion of a second field of view of the user at a current location of the user with a second direction;

the non-foveal region in each image of the sequence of images is a predetermined portion of a second field of view of the user at the current location with the second direction;

the second direction is offset from the first direction by a predetermined amount;

in each subsequent image of the sequence of images between the first image of the sequence of images and the second image of the sequence of images the foveal region comprises:

a first predetermined foveal portion comprising a first predetermined portion of the first field of view of the user at the current location of the user with the first direction; and

a second predetermined foveal portion comprising a first predetermined portion of the second field of view of the user at the current location of the user with the second direction.

8. The method according to claim **7**, further comprising

determining a natural minor saccade with respect to an eye of the user;

rendering the next image of the sequence of images to the user.

9. The method according to claim **7**, wherein

the rendering each image of the sequence of images is triggered by determining a natural minor saccade with respect of an eye of the user and independent of any induced minor saccade or major saccade.

10. The method according to claim **7**, wherein

physical motion of the user is redirected from the first direction to the second direction in dependence upon their viewing the sequence of images.

11. The method according to claim **7**, wherein

the field of view is one of a virtual environment, a physical real world environment of the user, the physical real world environment of the user augmented with additional digital content, and the physical real world environment of the user augmented with another virtual environment.

12. A method of redirecting motion of a user comprising:

a) determining a requirement to redirect the motion of the user from a current direction to a target direction;

b) determining an occurrence of an event with respect to an eye of the user;

c) determining whether the event is a minor saccade, a major saccade, or a blink;

d) upon determination of a minor saccade executing a next step of a first process which redirects the motion of the user from the current direction to the target direction in a number of steps;

e) upon determination of a major saccade or a blink executing a second process which redirects the motion of the user from the current direction to the target direction in a single step;

f) repeating steps (b) to (f) until motion of the user has been redirected to the target direction: wherein

the first process comprises rendering to the user a sequence of images relating wherein each image comprises:

a foveal region having a first predetermined geometry and first predetermined dimensions;

a non-foveal region having a second predetermined geometry and second predetermined dimensions; and

a transition region having a predetermined third predetermined geometry and third predetermined dimensions; wherein

in a first image of the sequence of images the foveal region is a first predetermined portion of a first field of view of the user at a current location of the user with a current direction;

in a last image of the sequence of images the foveal region is a first predetermined portion of a second field of view of the user at a current location of the user with a target direction;

the non-foveal region in each image of the sequence of images is a predetermined portion of a second field of view of the user at the current location with the target direction;

the second direction is offset from the first direction by a predetermined amount;

in each subsequent image of the sequence of images between the first image of the sequence of images and the second image of the sequence of images the foveal region comprises:

a first predetermined foveal portion comprising a first predetermined portion of the first field of view of the user at the current location of the user with the current direction; and

a second predetermined foveal portion comprising a first predetermined portion of the second field of view of the user at the current location of the user with the target direction.

13. The method according to claim **12**, wherein

in the second process a current image rendered to the user relating to a field of view from a current location in the current direction is replaced with a new image which is rendered to the user relating to a new field of view from the current location in the target direction.

14. A method of redirecting motion of a user comprising:

a) determining a requirement to redirect the motion of the user from a current direction to a target direction;

b) determining an occurrence of an event with respect to an eye of the user;

c) determining whether the event is a minor saccade, a major saccade, or a blink;

d) upon determination of a minor saccade executing a next step of a first process which redirects the motion of the user from the current direction to the target direction in a number of steps;

e) upon determination of a major saccade or a blink executing a second process which redirects the motion of the user from the current direction to the target direction in a single step;

f) repeating steps (b) to (f) until motion of the user has been redirected to the target direction; wherein

the first process comprises rendering to the user a sequence of images relating wherein each image comprises:

a foveal region having a first predetermined geometry and second predetermined dimensions;

a non-foveal region having a second predetermined geometry and second predetermined dimensions; and

a transition region having a predetermined third predetermined geometry and third predetermined dimensions; wherein

in a first image of the sequence of images the foveal region is a first predetermined portion of a first field of view of the user at a current location of the user with a first direction;

in a last image of the sequence of images the foveal region is a first predetermined portion of a second field of view of the user at a current location of the user with a second direction;

the non-foveal region in each image of the sequence of images is a predetermined portion of a second field of view of the user at the current location with the second direction;

the second direction is offset from the first direction by a predetermined amount;

in each subsequent image of the sequence of images between the first image of the sequence of images and the second image of the sequence of images the foveal region comprises:

a first predetermined foveal portion comprising a first predetermined portion of the first field of view of the user at the current location of the user with the first direction; and

a second predetermined foveal portion comprising a first predetermined portion of the second field of view of the user at the current location of the user with the second direction.

* * * * *

Appendix C

ACM SIGGRAPH 2020 Poster:

Dynamic Foveated Rendering for Redirected Walking in Virtual Reality

Dynamic Foveated Rendering for Redirected Walking in Virtual Reality

Yashas Joshi
Immersive and Creative Technologies (ICT) Lab, Department of Computer Science and Software Engineering, Concordia University, Montreal, Quebec, Canada
yashasjoshi1996@gmail.com

Charalambos Poullis
Immersive and Creative Technologies (ICT) Lab, Department of Computer Science and Software Engineering, Concordia University, Montreal, Quebec, Canada
charalambos@poullis.org

Figure 1: Left: Orientation-based redirection in Virtual Reality (VR) using dynamic-foveated rendering which leverages the effect of inattentional blindness induced by a cognitive task. The blending (e.g. parallax of green alien) in the transition zone between the rotated virtual environment (VE) (non-foveal) and the non-rotated VE (foveal) is imperceptible by the users due to inattentional blindness. Angular gain of the non-foveal zone for the frame shown is at its maximum i.e. 13.5°. Right: The path walked in physical tracked space (PTS) up to that point is shown in orange color. The corresponding path in the VE is shown in blue color. The cyan-colored box indicates the $4 \times 4m^2$ of available PTS and the camera icon inside the box indicates the location of the user w.r.t. the PTS. Even with room-scale PTS such as the one used, the users were able to walk distances in the virtual environment (VE) which are up to almost 18 times higher than the longest distance in the PTS; maximum recorded in experiments was 103.9m.

ACM Reference Format:
Yashas Joshi and Charalambos Poullis. 2020. Dynamic Foveated Rendering for Redirected Walking in Virtual Reality. In Special Interest Group on Computer Graphics and Interactive Techniques Conference Posters (SIGGRAPH '20 Posters), August 17, 2020, Virtual Event, USA. ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3388770.3407443

1 INTRODUCTION

In this work we present a novel technique for redirected walking in VR based on the psychological phenomenon of inattentional blindness. Based on the user's visual fixation points we divide the user's field of view (FoV) into zones. Spatially-varying rotations are then applied according to the zone's importance and are rendered using foveated rendering. Our technique is real-time and applicable to small and large physical spaces. Furthermore, the proposed technique does not require the use of stimulated saccades [Sun et al. 2018] but rather takes advantage of naturally occurring major and minor saccades and blinks to perform a complete refresh of the framebuffer. We performed extensive testing with the analysis of the results presented from three user studies conducted for the evaluation. Results show that the proposed technique is indeed viable and users were able to walk straight for more than 100m in VE within the confines of $4x4m^2$ of PTS.

SIGGRAPH '20 Posters, August 17, 2020, Virtual Event, USA

ACM ISBN 978-1-4503-7973-1/20/08.
https://doi.org/10.1145/3388770.3407443

2 TECHNICAL OVERVIEW

A blended render of the VE is shown to the user with the help of two co-located cameras, Camfoveal and Camnon−foveal. Based on the results from our first user study, we have determined that the FoV for Camfoveal is $\delta = 60°$, and the rotation angle applied to the VE and rendered from Camnon−foveal is $13.5° > \theta > 0°$. For making a smooth transition from foveal to non-foveal zones, a circular alpha mask with smooth boundaries corresponding to $\delta = 60°$ is applied on the rendered image of Camfoveal, and the inverse of the same mask is applied on the rendered image of Camnon−foveal. Smoothing the boundaries results in a transition zone between foveal and non-foveal zones. The resulting masked renders are then composited into a final render displayed to the user (Fig.1 Left). A safety reset mechanism is also added for ensuring the user does not collide with objects in the PTS. In such cases, this prompts the user

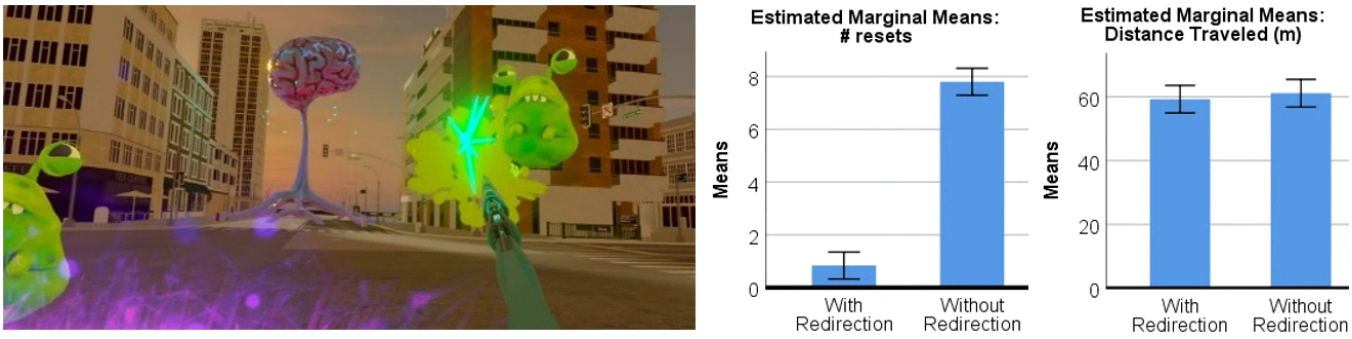

Figure 2: Left: Screen capture of the game for final user study. Steer-to-center algorithm was used for redirection. Right: ANOVA results for number of resets and distance traveled. Confidence interval = 95%

to turn in-situ by 180° while the VE is rotated by the same amount in the opposite direction. As demonstrated by the results from the subsequent two user studies, #2 and #3, the users fail to perceive any visual distractions or artifacts in the final composite render as they are preoccupied with a significant cognitive task on hand; which is almost always the case in VR applications.

The proposed system was developed using Unity3D game engine and the hardware used for its evaluation was HTC Vive Pro Eye HMD with integrated Tobii Eye Tracker. Using dynamic foveated rendering supported by NVIDIA RTX 2080Ti graphics card enables real-time performance. The pixels in the foveal zone are rendered at a higher resolution (1:1 sampling), in the transition zone at a medium resolution (4:1 sampling) and in the non-foveal (Peripheral) zone at a lower resolution (16:1 sampling).

3 EVALUATION

We performed 3 user studies for the evaluation. In user study #1 we determined the maximum rotation angle for the peripheral zone (13.5°) and FoV of the foveal zone (60°); in user study #2 we confirmed the results of the first #1; in user study #3 we assessed the eficacy of redirected walking using the proposed technique in a room-scale PTS of 4x4m^2 within the context of a custom designed first person shooter game (Fig. 2 left) which contained long straight walks. Fig. 1 show examples of the physical (orange) and virtual (blue) paths for a user during the final study. The cyan colored box indicates the PTS and the camera icon in it corresponds to the user's position. As shown the user covered a distance of 54.2m almost in a straight line to reach the final goal. Furthermore, we compared the number of resets required for each user to reach their goal with and without using redirection as shown in Fig. 2. We also performed Kennedy's simulator sickness test after each study which showed an insignificant impact on the sickness experienced amongst the participants.

Furthermore, a one-way between groups ANOVA (α =0.05) with repeated measures was performed post user study #3 to compare the effects of with and without redirection on the dependent variables: (a) number of resets, (b) distance traveled in PTS, and (c) total time taken. These variables passed Levene's homogeneity test. Partial eta squared (η^2_p) is used to report the obtained effect sizes for each variable.

ANOVA showed a statistically significant difference between the number of resets when the game was played with and without using redirection ($F(1, 48) = 375.710$, $p<0.001$) with $\eta^2_p = 0.887$. Nonetheless, these results also showed a statistically insignificant effect of redirection on distance travelled ($F(1, 48) = 0.384$, $p>0.05$; $\eta^2_p = 0.008$) and total time taken ($F(1, 48) = 1.001$, $p>0.05$; $\eta^2_p = 0.020$). The η^2_p values shows that 88.7% of the variability in the required number of resets is accounted for by our independent variable i.e redirection. However, the effects on distance travelled and total time taken remain negligible. The results from this test can be seen in Fig. 2.

Moreover, the analysis of the results from the final study showed that the system applied an average of 1547.55° (SD = 152.26°) of absolute angular gain to each participant's orientation during the entire test. An average of 3.15° of absolute angular gain was applied per redirection with an average of 1.73 redirections/s. As the participants were cognitively preoccupied with the task in the game, they were unaware of the angular gain occurring. In a post-study ques-tionnaire, one participant stated, "I felt like walking straight and was completely unaware of my actual movements". All evaluation details can be found in [Joshi and Poullis 2020].

4 CONCLUSION

In this work we presented a rotation-based redirection technique using dynamic foveated rendering which uses natural visual suppressions such as eye blinks and saccades (without any artificial stimuli) to make subtle rotations to the VE and leverages the effect of inattentional blindness induced by a cognitive task at hand. We conducted three user studies which confirmed the eficacy of the proposed technique in VR applications in which the cognitive workload on the user is moderate/high. It also confirmed that the technique can handle long straight walks which can be useful for applications such as immersive games, training simulations, etc.

REFERENCES

Qi Sun, Anjul Patney, Li-Yi Wei, Omer Shapira, Jingwan Lu, Paul Asente, Suwen Zhu, Morgan McGuire, David Luebke, and Arie Kaufman. 2018. Towards virtual reality infinite walking: dynamic saccadic redirection. ACM Transactions on Graphics (TOG). ACM 37, 4 (August 2018), 1-13. DOI:https://doi.org/10.1145/3197517.3201294

Yashas Joshi and Charalambos Poullis. 2020. Inattentional Blindness for Redirected Walking Using Dynamic Foveated Rendering. IEEE Access. IEEE 8, 19 (February 2020), 39013-39024. DOI:https://doi.org/10.1109/ACCESS.2020.2975032

Appendix D

Supplementary Material:

Enabling Saccadic Redirection Through Real-time Saccade Prediction

Figure D.1 shows the system overview for our redirected walking system.

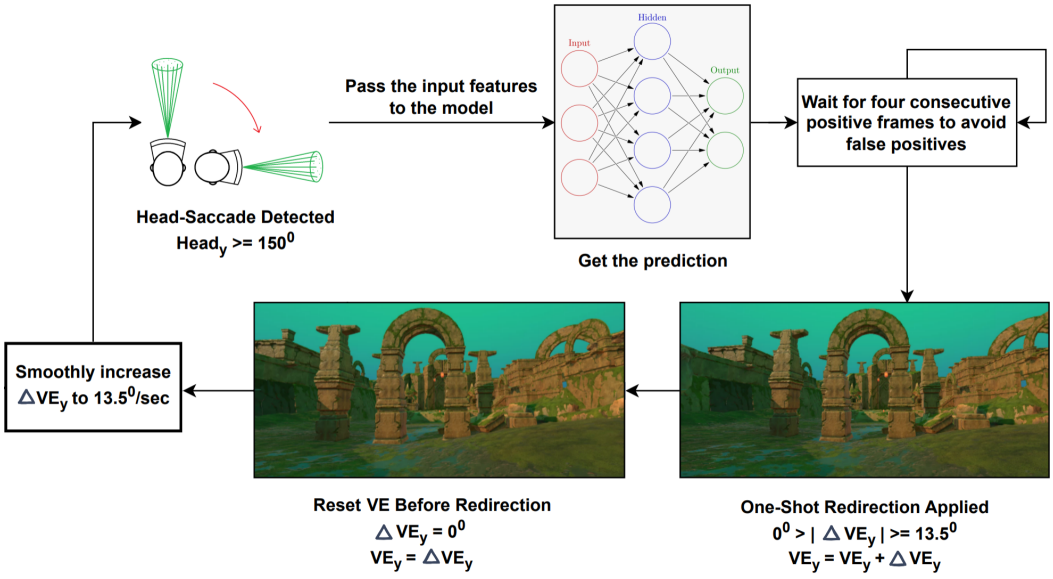

Figure D.1: Technical Overview. SaccadeNet performs in real-time and predicts saccades. During a predicted saccade, we adjust the VE according to where the user must be redirected.

Figure D.2 shows the experimental setup used for the user studies.

Figure D.2: Experimental setup for the final study, 3.5x3.5m^2 of PTS

Application for User study #1: Correlation of head and eye directions, and User study #2: Data collection

We developed an immersive VR experience to examine the correlation between head and eye rotations. The application portrays an open sky environment to eliminate any directional cues. Figure D.3 shows a screenshot from the application used in this first user study, showing the participant's view with two targets appearing in the open environment.

Figure D.4 shows the heatmaps for the 12 participants generated from the gaze data collected during the user study for a thorough analysis of the users' distribution of fixation points. A Gaussian filter was applied to smooth the heatmaps. The different colours indicate the average time spent by the users fixating around any particular region of the viewport. This time increases as the colours shift from navy blue to dark red. Furthermore, we marked the bounding boxes on each heatmap to quantify this region of average gaze. The corner points for all bounding boxes were plotted on a single graph to find the smallest enclosing ellipse. Gartner and Schoenherr's smallest enclosing

Figure D.3: User's perspective during the user study #1. Floating enemy targets in random directions are circled red, and the shooting weapon can be seen in the bottom right corner.

ellipse algorithm was then applied to determine the smallest ellipse that encloses all these points.

Gaze data for 12 participants is plotted as heatmaps for a thorough analysis of the users' distribution of fixation points. A Gaussian filter was applied to smooth the heatmaps and bounding boxes were generated. The corner points for all bounding boxes were plotted on a single graph to find the smallest enclosing ellipse. Gartner and Schoenherr's smallest enclosing ellipse algorithm was then applied to determine the smallest ellipse that encloses all these points.

Application for User study #3: Performance Evaluation of the Proposed Redirected Walking Technique

The main objective in this immersive game is for the user to collect three crystals from the ruins of the ancient abandoned arena. To achieve this, the participants have to walk from their initial spawn position to several predefined locations in the virtual environment (VE) marked by glowing crystals. Figure D.5b shows one of these crystals. Participants were given a quick tutorial at the beginning of the experience on the various interactions with the magical staff attached to their handheld controller. Using the staff, one can cast several magical spells such as a lightning bolt or a magic missile and make power strikes with melee attacks. Upon completing the tutorial, the first crystal and the location of the second crystal are revealed. Each of these crystals unlocks a new

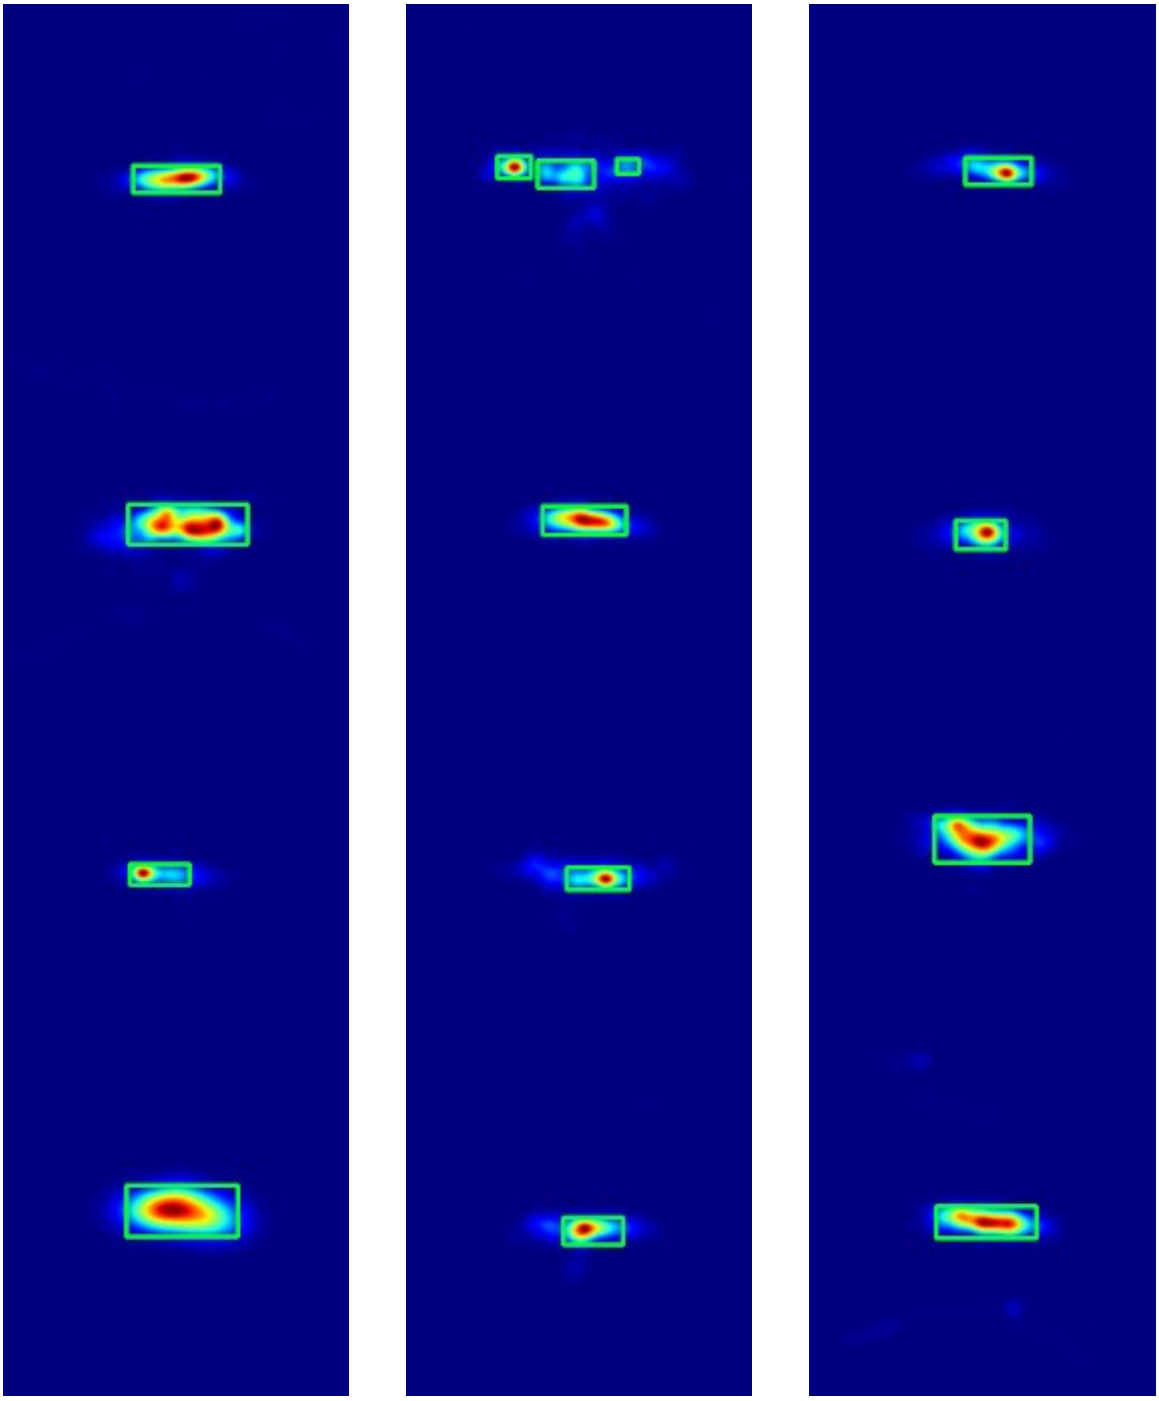

Figure D.4: Heatmap for each participant in the first user study

magical power. For example, the power for throwing lightning bolts from the staff is unlocked with the first crystal.

When walking towards the second destination, tiny dragons spawned far away flew to a random

position in a predefined orbital pattern around the character. They also make sound effects while flying as an additional audio cue. Figures D.5a, D.5b, D.5c, D.5d show screen captures of a participant's perspective during the experiments. Participants were directed to zap these dragons using their newly gained lightning bolt power triggered by a button on the Vive controller.

The immersive VR experience was designed to evaluate the redirected walking system. It involves a moderate cognitive task that causes repeated head rotations to eliminate the dragons, which means a higher probability of saccades occurring, therefore more redirections. Furthermore, the shortest straight distance to the second destination in the VE (38m) is multiple magnitudes larger than the longest possible straight distance in the PTS. The cyan-coloured box in Figure D.6 indicates the physical tracked space, while the grey and yellow lines indicate the physical and virtual paths taken by the user during the final experiments, respectively.

Upon reaching the destination and collecting the second crystal, redirection is toggled off, and the power strike is unlocked with increased melee damage to defeat swamp crawlers. (Figure D.5d). Defeating the herd of crawlers revealed the third and final crystal, unlocking an ability to shoot magical missiles. Grabbing the final crystal awakens the airborne mother dragon that spits fire, Figure D.7. The magical missiles were then used to destroy her before teleporting out of the forbidden island. Figure D.5a shows the abandoned arena depicted by the application, and Figure D.1 shows the technical overview of the proposed redirected walking system.

Figure D.5: (a) Virtual environment used during the evaluation of the proposed RDW technique. (b) Directional indicators (glowing crystals) reveal the next destination, one at a time. (c) User's perspective of dragon enemies and the mystical staff. (d) Swamp crawlers

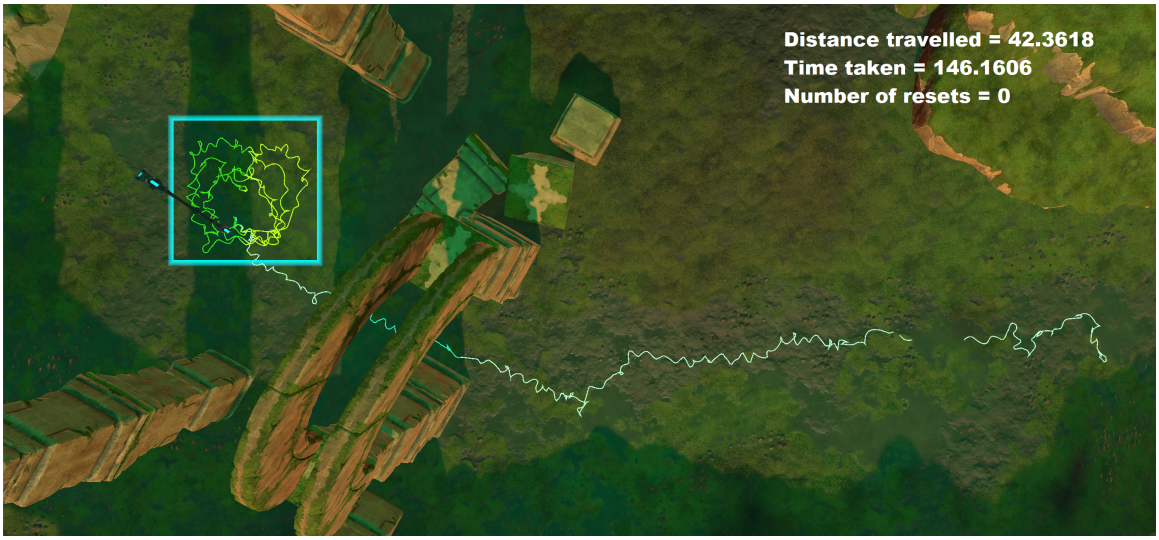

Figure D.6: Physical and virtual paths taken by a user during evaluation. The current statistics are shown on the top right. Cyan-coloured box indicates the 3.5 × 3.5m^2 PTS, while physical and virtual paths are marked with gray and yellow, respectively.

Figure D.7: Final User Study: Giant mother dragon before escaping the mysterious island

References

Akcayır, M., & Akcayır, G. (2017, February). Advantages and challenges associated with augmented reality for education: A systematic review of the literature. Educational Research Review, 20, 1-11.

al., A. C. (2006, June). Real-time markerless tracking for augmented reality: the virtual visual servoing framework. IEEE Transactions on Visualization and Computer Graphics, 12(4), 615-628.

Andujar, J. M., Mejías, A., & Márquez, M. A. (2010, October). Augmented reality for the improvement of remote laboratories: an augmented remote laboratory. IEEE transactions on education, 54(3), 492-500.

Armstrong, G., Hodgson, J., Manista, F., & Ramirez, M. (2012). The scarlet project: Augmented reality in special collections. Society of College, National and University Libraries focus, 54, 52-57.

Arnhem, J., & Spiller, J. (2014, May). Augmented reality for discovery and instruction. Journal of Web Librarianship, 8(2), 214-230.

Azmandian, M., Grechkin, T., Bolas, M., & Suma, E. (2016). The redirected walking toolkit: a unified development platform for exploring large virtual environments. In 2016 ieee 2nd workshop on everyday virtual reality (wevr) (p. 9-14). doi: 10.1109/WEVR.2016.7859537

Azmandian, M., Grechkin, T., & Rosenberg, E. S. (2017). An evaluationof strategies for two-user redirected walking in shared physical spaces. IEEE Virtual Reality, 91–98.

Azuma, R., Baillot, Y., Behringer, R., Feiner, S., Julier, S., & MacIntyre, B. (2001a). Recent advances in augmented reality. In Ieee computer graphics and applications (pp. 34–47).

IEEE.

Azuma, R., Baillot, Y., Behringer, R., Feiner, S., Julier, S., & MacIntyre, B. (2001b). Recent advances in augmented reality. IEEE Computer Graphics and Applications, 21(6), 34-47.

Azuma, R. T. (1997, August). A survey of augmented reality. PRESENCE: Teleoperators and Virtual Environments, 6(4), 355-385.

Bahill, T., Clark, M., & Stark, L. (1975). The main sequence, a tool for studying human eye movements. Mathematical biosciences, 24(3-4), 191–204.

Barandiaran, I., Paloc, C., & Gra na, M. (2009). Real-time optical markerless tracking for augmented reality applications. Journal of Real-Time Image Processing, 5(2), 129-138.

Bentivoglio, A. R., Bressman, S. B., Cassetta, E., Carretta, D., Tonali, P., & Albanese, A. (1997). Analysis of blink rate patterns in normal subjects. Movement Disorder, 12, 1028–1034.

Biocca, F. (1997, September). The cyborg's dilemma: Progressive embodiment in virtual environments. Computer-mediated communication, 3, 2.

Biocca, F., Harms, C., & Gregg, J. (2001, May). The networked minds measure of social presence: Pilot test of the factor structure and concurrent validity. In 4th annual international workshop on presence (p. 1-9).

Bolte, B., & Lappe, M. (2015). Subliminal reorientation and repositioning in immersive virtual environments using saccadic suppression. IEEE Transactions of Visualization and Computer Graphics, 21(4), 545–552.

Bradley, J., Henshaw, N., McVoy, L., French, A., Gilbertson, K., Becksford, L., & Givens, E. (2016, April). Creation of a library tour application for mobile equipment using ibeacon technology. Code4lib Journal, 32.

Burr, D. C., Morrone, M. C., & Ross, J. (1994). Selective suppression of the magnocellular visual pathway during saccadic eye movements. Nature, 371(6497), 511–513.

Chatzopoulos, D., Bermejo, C., Huang, Z., & Hui, P. (2017). Mobile augmented reality survey: From where we are ti where we go. IEEE Access, 5, 6917–6950.

Chen, C. M., & Tsai, Y. N. (2012, September). Interactive augmented reality system for enhancing library instruction in elementary schools. Computers and Education, 59(2), 638-652.

Cheng, L.-P., Roumen, T., Rantzsch, H., Kohler, S., Schmidt, P., Kovacs, R., ... Baudisch, P. (2015).

Turkdeck: Physical virtual reality based on people. In Proceedings of the 28th annual acm symposium on user interface software & technology (pp. 417–426).

Cherni, H., Métayer, N., & Souliman, N. (2020). Literature review of locomotion techniques in virtual reality. International Journal of Virtual Reality, 20, 1-20. doi: 10.20870/IJVR.2020 .20.1.3183

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

Christou, C., Tzanavari, A., Herakleous, K., & Poullis, C. (2016a). Navigation in virtual reality: Comparison of gaze-directed and pointing motion control. In 2016 18th mediterranean electrotechnical conference (melecon) (pp. 1–6).

Christou, C., Tzanavari, A., Herakleous, K., & Poullis, C. (2016b, April). Navigation in virtual reality: Comparison of gaze-directed and pointing motion control. In 18th mediterranean electrotechnical conference (melecon) (p. 1-6). IEEE.

Cieza, E., & Lujan, D. (2018). Educational mobile application of augmented reality based on markers to improve the learning of vowel usage and numbers for children of a kindergarten in trujillo. Procedia Computer Science, 130, 352-358.

Craig, A. B. (2013). Understanding augmented reality: Concepts and applications. Amsterdam. Retrieved from https://www.sciencedirect.com/book/9780240824086

Darken, R. P., Cockayen, W. R., & Carmein, D. (1986). A hand gesture interface device. In Acm sigchi bulletin (pp. 189–192).

Darken, R. P., Cockayen, W. R., & Carmein, D. (1997). The omni-directional treadmill: A locomotion device for virtual worlds. In Proceedings of the 10th annual acm symposium on user interface software and technology (pp. 213–221).

Dash, A., Behera, S., Dogra, D., & Roy, P. (2018). Designing of marker-based augmented reality learning environment for kids using convolutional neural network architecture. In Elsevier/displays (pp. 46–54).

Dash, A. K., Behera, S. K., Dogra, D. P., & Roy, P. P. (2018). Designing of marker-based augmented reality learning environment for kids using convolutional neural network architecture

(Vol. 55). pp. 46-54: Elsevier/Displays.

Dichgans, J., & Brandt, T. (1978). Visual vestibular interaction: Effects on self-motion perception and postural control. In Perception. Handbook of Sensory Physiology (pp. 755–804).

Dong, Z.-C., Fu, X.-M., Zhang, C., Wu, K., & Liu, L. (2017). Smooth assembled mappings for large-scale real walking. ACM Transactions on Graphics (TOG), 36(6), 211.

Eike, L., Paul, L., Gerd, B., & Frank, S. (2017). Application of redirected walking in room-scale vr. IEEE Virtual Reality.

Fang, Y., Nakashima, R., Matsumiya, K., Kuriki, I., & Shioiri, S. (2015). Eye-head coordination for visual cognitive processing. PLOS ONE, 10.

Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G* power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. Behavior research methods, 39(2), 175–191.

Fernandes, K. J., Raja, V., & Eyre, J. (2003). Cybersphere: The fully immersive spherical projection system. Communications of the ACM, 46(9), 141–146.

Fisher, S., Wenzel, E., Coler, C., & Mcgreevy, M. (1988, 02). Virtual interface environment workstations. Proceedings of the Human Factors Society Annual Meeting, 32. doi: 10.1177/154193128803200219

geun Kim, Y., & jung Kim, W. (2014). Implementation of augmented reality system for smartphone advertisements. Intl. Journal of Multimedia and Ubiquitous Engineering, 9(2), 358-392.

Glaholt, M. G. (2016). Field of view requirements for night vision devices (Tech. Rep. No. DRDC-RDDC-2016-R032). Defence Research and Development Canada. Retrieved from https://cradpdf.drdc-rddc.gc.ca/PDFS/unc243/p804476_A1b.pdf

Gärtner, G., & Schönherr, S. (1997). Smallest enclosing ellipses – fast and exact.

Hahn, J. (2012, September). Mobile augmented reality applications for library services. New Library World, 113(9/10), 429-438. Retrieved from https://doi.org/10.1108/03074801211273902

He, Y., Gu, Y., & Fatahalian, K. (2014, July). Extending the graphics pipeline with adaptive, multi-rate shading. ACM Trans. Graph., 33(4), 142:1–142:12. Retrieved from http://doi.acm.org/10.1145/2601097.2601105 doi: 10.1145/2601097.2601105

Hodgson, E., & Bachmann, E. (2013). Comparing four approaches to generalized redirected walking: Simulation and live user data. IEEE transactions on visualization and computer graphics, 19, 634–643.

Hodgson, E., Bachmann, E., & Thrash, T. (2014). Performance of redirected walking algorithms in a constrained virtual world. IEEE transactions on visualization and computer graphics, 20, 579–587.

Huang, J. Y. (2003). An omnidirectional stroll-based virtual reality interface and its application on overhead crane training. IEEE Transactions on Multimedia, 5(1), 39–51.

Ibbotson, M. R., & Cloherty, S. L. (2009). Visual perception: Saccadic omission- suppression or temporal masking? Current Biology, 19(12), R493–R496.

Infinadeck - an omni-directional treadmill. (n.d.). https://infinadeck.com/. (Accessed: 2020-10-01)

Iwata, H. (1999). Walking about virtual environments on an infinite floor. IEEE Virtual Reality, 286–293.

Iwata, H., & Fujii, T. (1996). Virtual perambulator: a novel interface device for locomotion in virtual environment. In Proceedings of the ieee 1996 virtual reality annual international symposium (pp. 60–65).

Joshi, Y., & Poullis, C. (2020a). Dynamic foveated rendering for redirected walking in virtual reality. In Acm siggraph 2020 posters. New York, NY, USA: Association for Computing Machinery. Retrieved from https://doi.org/10.1145/3388770.3407443 doi: 10.1145/3388770.3407443

Joshi, Y., & Poullis, C. (2020b). Inattentional blindness for redirected walking using dynamic foveated rendering. IEEE Access, 8, 39013-39024. doi: 10.1109/ACCESS.2020.2975032

Joshi, Y., & Poullis, C. (2020c). Portal to knowledge: a virtual library using marker-less augmented reality system for mobile devices. In Proceedings of optical architectures for displays and sensing in augmented, virtual, and mixed reality (ar, vr, mr).

Joshi, Y., & Poullis, C. (2023a, January 17). Managing real world and virtual motion. US Patent US11557105B2 [Online]. Retrieved from https://patents.google.com/patent/US11557105B2

Joshi, Y., & Poullis, C. (2023b). Saccadenet: Towards real-time saccade prediction for virtual reality infinite walking (invited). Journal of Computer Animation and Virtual Worlds.

Kao, T. W., & Shih, H. C. (2013). A study on the markerless augmented reality for picture books. In Proceedings of the ieee international symposium on consumer electronics (isce (p. 197-198).

Kennedy, R., Drexler, J., Compton, D., Stanney, K., Lanham, D., & Harm, D. (2003). Configural scoring of simulator sickness, cybersickness and space adaptation syndrome: Similarities and differences. Virtual and Adaptive Environments: Applications, implications and human performance issues, 247.

Kennedy, R. S., Lane, N. E., Berbaum, K. S., & Lilienthal, M. G. (1993). Simulator sickness questionnaire: An enhanced method for quantifying simulator sickness. The International Journal of Aviation Psychology, 3, 203–220.

Kothari, R., Yang, Z., Kanan, C., Bailey, R., Pelz, J., & Diaz, G. (2020). Gaze-in-wild: A dataset for studying eye and head coordination in everyday activities. Sci Rep, 10.

Langbehn, E., Bruder, G., & Steinicke, F. (2016). Subliminal re-orientation and re-positioning in virtual reality during eye blinks. In (pp. 213–213).

Langbehn, E., Lubos, P., & Steinicke, F. (2018). Evaluation of locomotion techniques for room-scale vr: Joystick, teleportation, and redirected walking. In Proceedings of the virtual reality international conference - laval virtual. Association for Computing Machinery. doi: 10.1145/3234253.3234291

Langbehn, E., & Steinicke, F. (2018a). In the blink of an eye - leveraging blink-induced suppression for imperceptible position and orientation redirection in virtual reality. ACM Transactions on Graphics, 37, 1–11.

Langbehn, E., & Steinicke, F. (2018b). Redirected walking in virtual reality. Encyclopedia of Computer Graphics and Games. doi: 10.1007/978-3-319-08234-9_253-1

Langlotz, T., Degendorfer, C., Mulloni, A., Schall, G., Reitmayr, G., & Schmalstieg, D. (2011, August). Robust detection and tracking of annotations for outdoor augmented reality browsing. Computers and graphics, 35(4), 831-840.

Lathrop, W., & Kaiser, M. (2002). Perceived orientation in physical and virtual environments: Changes in perceived orientation as a function of idiothetic information available. Presence:

Teleoperators and Virtual Environments, 11(1), 19-32.

LaViola Jr., J. J. (2000). A discussion of cybersickness in virtual environments. ACM SIGCHI, 32(1), 47–56.

Levene, H. (1960). Robust testes for equality of variances. Contributions to Probability and Statistics, 278–292.

Li, H., & Fan, L. (2020). Mapping various large virtual spaces to small real spaces: A novel redirected walking method for immersive vr navigation. IEEE Access, 8, 180210-180221. doi: 10.1109/ACCESS.2020.3027985

Li, J., Slembrouck, M., & Deboeverie, F. (2015, September). A hybrid pose tracking approach for handheld augmented reality. In Proceedings of the 9th international conference on distributed smart camera (p. 7-12).

Loomis, J. M., Blascovich, J. J., & Beall, A. C. (1999, December). Immersive virtual environment technology as a basic research tool in psychology. Behavior research methods, instruments, and computers, 31(4), 557-564.

Loomis, J. M., Golledge, R. G., & Klatzky, R. L. (1998, April). Navigation system for the blind: Auditory display modes and guidance. Presence: Virtual and Augmented Reality, 7(2), 193-203.

Massis, B. (2015, November). Using virtual and augmented reality in the library. New Library World, 116(11/12), 796-799. doi: http://dx.doi.org/10.1108/NLW-08-2015-0054

Maxwell, S. E., & Delaney, H. D. (2018). Designing experiments and analyzing data: A model comparison perspective. Routledge.

Medina, E., Fruland, R., & Weghorst, S. (2008). Virtusphere: Walking in a human size vr "hamster ball". In Proceedings of the human factors and ergonomics society annual meeting (Vol. 52, pp. 2102–2106).

Meredith, T. R. (2015, April). Using augmented reality tools to enhance children's library services. Technology, Knowledge and Learning, 20(1), 71-77.

Michael, S., Chin-Yih, U., Markus, Z., Max, M., Andreas, K., Christina, H., & Marino, M. (2019, March). Consumers' food selection behaviors in three-dimensional (3d) virtual reality. Food Research International, 117, 50–59.

Milgram, P., Takemura, H., Utsumi, A., & Kishino, F. (1995). Augmented reality: a class of displays on the reality-virtuality continuum. In Proceedings of spie, telemanipulator and telepresence technologies (Vol. 2351).

Moshell, M. (1999). Infinite virtual walk. In Personal communication.

M. Sagayam, K., Ho, C. C., Henesey, L., & Bestak, R. (2018, Aug). 3d scenery learning on solar system by using marker based augmented reality. In 4th international conference of the virtual and mented reality in education (p. 139-143).

Mulch, B. E. (2014, April). Library orientation transformation. Knowledge quest, 42(4), 50-53.

Nagamori, A., Wakabayashi, K., & Ito, M. (2005). The ball array treadmill: A locomotion interface for virtual worlds. IEEE Virtual Reality, 3–6.

Nilsson, N. C., Peck, T., Bruder, G., Hodgson, E., Serafin, S., Whitton, M., ... Rosenberg, E. S. (2018). 15 years of research on redirected walking in immersive virtual environments. IEEE Computer Graphics and Applications, 38(2), 44-56. doi: 10.1109/MCG.2018.111125628

NVIDIA. (2018). VRWorks - Variable Rate Shading. https://developer.nvidia.com/vrworks/graphics/variablerateshading (last accessed 4th Nov. 2019).

O'Connor, Chris. (2018). ZeroLight Improves Automotive Product Visualisation Quality and Performance with VRS. https://developer.nvidia.com/vrworks/graphics/variablerateshading (last accessed 4th Nov. 2019).

Onime, C., & Abiona, O. (2016, April). 3d mobile augmented reality interface for laboratory experiments. International Journal of Communications, Network and System Sciences, 9(4), 67.

O'Regan, J. K., Deubel, H., Clark, J. J., & Rensink, R. A. (2000). Picture changes during blinks: Looking without seeing and seeing without looking. Visual cognition, 7, 191–211.

Pang, Y., Yuan, M. L., Nee, A. Y. C., Ong, S. K., & Youcef-Toumi, K. (2006, January). A markerless registration method for augmented reality based on affine properties. In Proceedings of the 7th australasian user interface conference (p. 25-32). Vol. 50.

Patney, A., Kim, J., Salvi, M., Kaplanyan, A., Wyman, C., Benty, N., ... Luebke, D. (2016). Perceptually-based foveated virtual reality. In Acm siggraph 2016 emerging technologies

(pp. 17:1–17:2). New York, NY, USA: ACM. Retrieved from http://doi.acm.org/10.1145/2929464.2929472

Peck, T., Fuchs, H., & Whitton, M. (2011). An evaluation of navigational ability comparing redirected free exploration with distractors to walking-in-place and joystick locomotion interfaces. Proceedings of IEEE Virtual Reality, 56–62.

Peck, T., Whitton, M., & Fuchs, H. (2008). Evaluation of reorientation techniques for walking in large virtual environments. In 2008 ieee virtual reality conference (p. 121-127). doi: 10.1109/VR.2008.4480761

Poullis, C., & You, S. (2009). Automatic creation of massive virtual cities. In 2009 ieee virtual reality conference (pp. 199–202).

Powell, W., Powell, V., Brown, P., Cook, M., & Uddin, J. (2016, March). Getting around in google cardboard – exploring navigation preferences with low-cost mobile vr. In 2nd workshop on everyday virtual reality (wevr) (p. 5-8). Greenville, S. C.: IEEE.

Ramot, Daniel. (n.d.). Average duration of a single blink. https://bionumbers.hms.harvard.edu/bionumber.aspx?id=100706&ver=0 (last accessed 1st Nov. 2019).

Razzaque, S., Kohn, Z., & Whitton, M. C. (2001). Redirected walking. Proceedings of Eurographics, 9, 105–106.

Razzaque, S., Kohn, Z., & Whitton, M. C. (2005). Redirected walking. Citeseer.

Razzaque, S., Swapp, D., Slater, M., Whitton, M. C., & Steed, A. (2002). Redirected walking in place. Eurographics Symposium on Virtual Environments, 123–130.

Reder, S. M. (1973). On-line monitoring of eye-position signals in contingent and noncontingent paradigms. Behavior Research Methods & Instrumentation, 5, 218–228.

Rensink, R. A. (2002). Change detection. Annual review of psychology, 53, 245–277.

Rensink, R. A., O'Regan, J. K., & Clark, J. J. (1997). To see or not to see: The need for attention to perceive changes in scenes. Psychological science, 8, 368–373.

Rodríguez, R. A. M., & Rivero, M. O. M. (2016, February). Information skills training through mobile devices: Practical applications of qr codes in academic libraries. The Electronic Library, 34(1), 116-131.

Rothacher, Y., Nguyen, A., Lenggenhager, B., Kunz, A., & Brugger, P. (2018). Visual capture of gait during redirected walking. Scientific Reports, 8. Retrieved from http://doi.org/10.1038/s41598-018-36035-6 doi: 10.1038/s41598-018-36035-6

Ruddle, R., & Lessels, S. (2009). Walking interface to navigate virtual environments. ACM Trans. on Computer-Human Interaction, 16, 5:1–5:18.

Ruddle, R. A., Volkova, E. P., & Bulthoff, H. H. (2011). Walking improves your cognitive map in environments that are large-scale and large in extent. ACM Trans. on Computer-Human Interaction, 18, 10:1–10:22.

Sakono, H., Matsumoto, K., Narumi, T., & Kuzuoka, H. (2021). Redirected walking using continuous curvature manipulation. IEEE Transactions on Visualization and Computer Graphics, 27(11), 4278-4288. doi: 10.1109/TVCG.2021.3106501

Sansone, R. A., & Sansone, L. A. (2013, January). Cell phone: The psychological risks. Innovations in clinical neuroscience, 10(1), 33-37.

Shatte, A., Holdsworth, J., & Lee, I. (2014, April). Mobile augmented reality based context-aware library management system. Expert Systems with Applications, 41(5), 2174-2185.

Sidenmark, L., & Gellersen, H. (2019). Eye head and torso coordination during gaze shifts in virtual reality. ACM Transactions on Computer-Human Interaction (TOCHI), 20, 1–40.

Siegrist, M., Ung, C. Y., Zank, M., Marinello, M., Kunz, A., Hartmann, C., & Menozzi, M. (2019, March). Consumers' food selection behaviors in three-dimensional (3d) virtual reality. Food Research International, 117, 50-59.

Sikström, E., de Götzen, A., & Serafin, S. (2015). Wings and flying in immersive vr — controller type, sound effects and experienced ownership and agency. In 2015 ieee virtual reality (vr) (p. 281-282). doi: 10.1109/VR.2015.7223405

Simons, D., & Rensink, R. (2005). Change blindness: past, present, and future. Trends in cognitive science, 9(1), 16–20. doi: 10.1016/j.tics.2004.11.006

Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattentional blindness for dynamic events. perception, 28(9), 1059–1074.

Souman, J. L., Giordano, P. R., Frissen, I., Luca, A. D., & Ernst, M. O. (2010). Making virtual walking real: Perceptual evaluation of a new treadmill control algorithm. ACM Transactions

on Applied Perception (TAP), 7(2), 11.

Stahl, J. (1999). Amplitude of human head movements associated with horizontal saccades. Exp Brain Res., 126, 41–54.

Steed, A., & Julier, S. (2013, March). Design and implementation of an immersive virtual reality system based on a smartphone platform. In Ieee symposium on (p. 43-46). 3D User Interfaces (3DUI).

Steinicke, F., Bruder, G., Hinrichs, K., & Willemsen, P. (2011). Change blindness phenomena for virtual reality display systems. IEEE Transactions on Visualization and Computer Graphics (TVCG), 17, 1223–1233.

Steinicke, F., Bruder, G., Kohli, L., Jerald, J., & Hinrichs, K. (2008). Taxonomy and implementation of redirection techniques for ubiquitous passive haptic feedback. IEEE International Conference on Cyberworlds, 217–223.

Stengel, M., Grogorick, S., Eisemann, M., & Magnor, M. (2016). Adaptive Image-Space Sampling for Gaze-Contingent Real-time Rendering. Computer Graphics Forum. doi: 10.1111/cgf.12956

Suma, E., Bruder, G., Steinicke, F., Krum, D., & Bolas, M. (2012). A taxonomy for deploying redirection techniques in immersive virtual environments. In Proceedings of the 2012 ieee virtual reality (p. 43–46). USA: IEEE Computer Society. Retrieved from https://doi.org/10.1109/VR.2012.6180877 doi: 10.1109/VR.2012.6180877

Suma, E. A., Clark, S., Finkelstein, S. L., & Wartell, Z. (2010). Exploiting change blindness to expand walkable space in a virtual environment. Proceedings of IEEE Virtual Reality (VR), 305–306.

Suma, E. A., Clark, S., Krum, D., Finkelstein, S., Bolas, M., & Warte, Z. (2011). Leveraging change blindness for redirection in virtual environments. In 2011 ieee virtual reality conference (pp. 159–166).

Suma, E. A., Lipps, Z., Finkelstein, S., Krum, D. M., & Bolas, M. (2012). Impossible spaces: Maximizing natural walking in virtual environments with self-overlapping architecture. IEEE Transactions on Visualization and Computer Graphics, 18(4), 555–564.

Sun, Q., Patney, A., Wei, L.-Y., Shapira, O., Lu, J., Asente, P., ... Kaufman, A. (2018). Towards

virtual reality infinite walking: dynamic saccadic redirection. ACM Transactions on Graphics (TOG), 37(4), 67.

Sun, Q., Wei, L.-Y., & Kaufman, A. (2016). Mapping virtual and physical reality. ACM Transactions on Graphics (TOG), 35(4), 64.

Sutcliffe, A. G., Poullis, C., Gregoriades, A., Katsouri, I., Tzanavari, A., & Herakleous, K. (2019). Reflecting on the design process for virtual reality applications. International Journal of Human–Computer Interaction, 35(2), 168-179.

Sutherland, I. (1968). A head-mounted three dimensional display. In Proceedings of the afips fall joint computer conference, washington, d.c. (pp. 757–764). Thompson Books.

Sutherland, I. E. (1968). A head-mounted three dimensional display. Washington, D.C., Thompson Books, pp. 757-764: Proceedings of the AFIPS Fall Joint Computer Conference.

Usoh, M., Arthur, K., Whitton, M. C., Bastos, R., Steed, A., Slater, M., & Jr., F. P. B. (1999). Walking ¿ walking-in-place ¿ flying, in virtual environments. SIGGRAPH'99 Proceedings of the 26th annual conference on Computer graphics and interactive techniques, 359–364.

Volkmann, F. C. (1986). Human visual suppressions. Vision Research, 26, 1401–1416.

Volkmann, F. C., Riggs, L. A., & Moore, R. K. (1980). Eyeblinks and visual suppression. Science, 207(4433), 900–902.

Walther-Franks, B., Wenig, D., Smeddinck, J., & Malaka, R. (2013). Suspended walking: A physical locomotion interface for virtual reality. In International conference on entertainment computing (pp. 185–188).

Weber, N., Lennartz, R., Knitza, J., Bayat, S., Sadeghi, M., Ibrahim, A., ... Kleyer, A. (2022, 06). Ab1528-hpr full body haptic bodysuit - an instrument to measure the range and speed of motion in patients with axial spondyloarthritis (axspa) - preliminary results. Annals of the Rheumatic Diseases, 81, 1866.2-1867. doi: 10.1136/annrheumdis-2022-eular.3069

Whitchurch, M. (2012, October). A quick response: Qr code use at the harold b. lee library. The Reference Librarian, 53(4), 392-402.

Williams, B., Narasimham, G., Rump, B., McNamara, T., Carr, T., Rieser, J., & Bodenheimer, B. (2007). Exploring large virtual environments with an hmd when physical space is limited. In (pp. 41–48). ACM.

Williams, N. L., Bera, A., & Manocha, D. (2021a). Arc: Alignment-based redirection controller for redirected walking in complex environments. IEEE Transactions on Visualization and Computer Graphics, 27(5), 2535-2544. doi: 10.1109/TVCG.2021.3067781

Williams, N. L., Bera, A., & Manocha, D. (2021b). Redirected walking in static and dynamic scenes using visibility polygons. IEEE Transactions on Visualization and Computer Graphics, 27(11), 4267-4277. doi: 10.1109/TVCG.2021.3106432

Won, S. H. P., Melek, W. W., & Golnaraghi, F. (2009, September). A kalman/particle filter-based position and orientation estimation method using a position sensor/inertial measurement unit hybrid system. IEEE Transactions on Industrial Electronics, 57(5), 1787-1798.

Xiong, Q., Liang, X., Wei, D., Wang, H., Zhu, R., Wang, T., ... Wang, H. (2022). So-eaglove: Vr haptic glove rendering softness sensation with force-tunable electrostatic adhesive brakes. IEEE Transactions on Robotics, 1-13. doi: 10.1109/TRO.2022.3172498

Yang, X., & Cheng, K. T. T. (2012). Accelerating surf detector on mobile devices. In Proceedings of the acm international conference on multimedia (p. 569-578).

Zank, M., & Kun, A. (2015). Using locomotion models for estimating walking targets in immersive virtual environments. IEEE International Conference on Cyberworlds.

Zank, M., & Kun, A. (2016). Eye tracking for locomotion prediction in redirected walking. IEEE Symposium on 3D User Interfaces.