Proprioceptive Puppeteering: Enhancing Lower-Body Control in Virtual Avatars

by

Ethan C. Bowmar

(Under the Direction of Kyle J. Johnsen)

Abstract

Virtual Reality (VR) and Mixed Reality (MR) present promising new solutions for pain management, particularly for phantom limb pain (PLP) in amputee patients. These technologies enable immersive experiences that promote the embodiment of missing limbs, offering novel therapeutic possibilities. However, fully embodying one's lower limbs with a virtual avatar can be difficult due to the lack of robust lower-body alignment techniques. To address this, two lower-body control interfaces were designed for use within both VR and MR environments. A user study with 40 healthy participants assessed how the immersive environment (VR vs. MR) influences task performance, sense of embodiment, and system usability during a body alignment task.

Index words: Software Development, Unity3D, Virtual Reality, Mixed Reality, Augmented Reality, Extended Reality, Embodiment, Thesis, Proprioception, Control Interfaces, Phantom Limb Treatment, Virtual Therapies

Proprioceptive Puppeteering:
Enhancing Lower-Body Control in Virtual Avatars

by

Ethan C. Bowmar

B.S. of Science, The University of Georgia, 2023

A Thesis Submitted to the Graduate Faculty
of The University of Georgia in Partial Fulfillment
of the
Requirements for the Degree

Master of Science

Athens, Georgia

2025

Proprioceptive Puppeteering:
Enhancing Lower-Body Control in Virtual Avatars

by

Ethan C. Bowmar

Approved:

Major Professor: Kyle J. Johnsen

Committee: Fred R Beyette Jr.
Mark Haidekker

Electronic Version Approved:

Ron Walcott
Dean of the Graduate School
The University of Georgia
August 2025

Acknowledgments

I am very grateful for the support provided by the Virtual Environments Research Lab (VEL), as well as the direct guidance and mentorship of my advisor, Dr. Kyle Johnsen. I would also like to thank my committee members for their valuable input and encouragement throughout this process.

Additionally, I extend my thanks to Dr. Sun Joo (Grace) Ahn, Dr. Yuqian Ni, and Dr. Joshua Baldwin from the Center for Advanced Computer-Human Ecosystems (CACHE) in the Grady College of Journalism and Mass Communication at the University of Georgia, whose assistance with the user study greatly enriched this work.

Table of Contents

Page

Acknowledgments iv

Chapter

1 Introduction 1

1.1 Background and Motivation 1

1.2 The Research Problem 2

1.3 Research Objectives and Approach 2

1.4 Outline of the Thesis 3

2 Literature Review 5

2.1 Introduction to the Literature Review 5

2.2 Foundations of Immersive Technologies 5

2.3 Embodiment and Presence in Extended Reality (XR) . . . 8

2.4 The Challenge of Lower Body Embodiment and Control in XR 12

2.5 The Impact of Immersive Environments on Performance and Embodiment 15

2.6 Proprioception and Alignment Accuracy 18

2.7 Summary of Key Findings and Identified Gaps 18

3 Methodology 21

3.1 Research Question 21

3.2 Experimental Design 23

3.3 Participants 25

3.4 Apparatus and Materials 26

3.5 Procedure 27

3.6 Measures 32

3.7 Data Pre-processing and Screening 40

3.8 Statistical Analysis Plan 43

4 System 50

4.1 High-Level System Overview 50

4.2 Development Environment and Core Technologies 50

4.3 Avatar and Interface Design 51

4.4 Implementation Journey: From IK to Physics 54

4.5 Final System Components 56

4.6 Data and Session Management 60

4.7 System Conclusion 61

5 Results 62

5.1 Descriptive Statistics 62

5.2 Missing Data and Outliers 63

5.3 Hypothesis 1: Effect of Immersion on Alignment Performance 70

5.4 Hypothesis 2: Effect of Immersion on Sense of Embodiment 75

5.5 Exploratory Findings 76

5.6 Results Summary 84

6 Discussion 85

6.1 Summary of Key Findings 85

6.2 Interpretation of Hypothesis 1 (Alignment Performance) 86

6.3 H2: The Trade-off Between Transformation and Disembodiment 89

6.4 Interpretation of Exploratory Findings 90
6.5 Broader Implications 94
6.6 Limitations and Future Work 96
6.7 Discussion Conclusion 99

7 Conclusion 101
7.1 Summary of the Thesis 101
7.2 Primary Contributions 103
7.3 Final Remarks and Future Outlook 103

Chapter 1

Introduction

1.1 Background and Motivation

1.1.1 Extended Reality and Phantom Limb Pain

Phantom Limb Pain (PLP) is a painful sensation that affects upwards of 85% of amputees[14]. The sensation is described as pain or discomfort in a limb that is no longer there, ranging from mild discomfort to a stinging, burning sensation [23]. PLP events can last from minutes to hours, and for many individuals, the pain is ongoing and periodic—a condition classified as *chronic PLP* [23]. The gold standard for managing this pain has been a mixture of the prescription of traditional opioids and therapies. As virtual reality (VR) devices become more accessible, there is growing interest in utilizing them for these therapies. Applications span from fully immersive VR to mixed reality (MR), which blends virtual elements with the user's physical environment. These technologies, collectively known as Extended Reality (XR), often achieve promising results by positioning a virtual limb where the user perceives their phantom limb to be. Doing so creates a sense of embodiment (SoE), particularly in the form of body ownership—an essential component in PLP therapies shown to foster healing and trick the brain into feeling or "sensing" virtual sensations [17, 24].

Traditionally, this same effect was achieved through methods such as mirror therapy (MT), where a physical mirror is used to create the illusion of having a once missing limb by mirroring the in-tact [13]. This therapeutic effect is achieved by effectively closing the loop between perception and intention by enabling a patient to experience control and body ownership over a *virtual phantom limb* [3]. Virtual mirror therapies replace the mirror and use virtual replicas of the phantom limb, rendered by the headset, to be embodied in the patients avatar [18]. These therapies are controlled using traditional XR interfaces such as controllers, motion capture cameras, or built-in hand tracking technology [7, 20]. This new approach to managing PLP has been studied and compared to traditional MT on several accounts and found to be just as effective [43], while also offering the potential to increase flexibility and motivation [43, 36].

1.2 The Research Problem

While this intervention has been shown to work, a distinctive challenge remains: creating a robust, accessible, *lower-body* control interface. The interface needs to be intuitive and easy to use, and it must also foster a strong SoE in either VR or MR environments.

The development of an effective lower body interface for XR environments poses an interesting design challenge. Existing literature shows a key tradeoff between VR and MR experiences, notably between embodiment and task performance. It is theorized that users in MR, who see both the real and virtual worlds, perform better in precision-based tasks but may experience a weaker SoE compared to users in fully immersive VR. However, this trade-off is not well understood when it comes to lower-body alignment tasks. To address this gap, we designed two novel lower-body interfaces and conducted a study to compare the impact of MR and VR environments directly.

1.3 Research Objectives and Approach

The primary objective of this research is to investigate how novel lower-body control interfaces and different immersive environments (MR vs. VR) impact user performance on a body alignment task and their resulting SoE. This objective is addressed through the following research question:

How do body alignment method and immersive environment (MR vs. VR) impact user preference and performance in virtual avatar control?

From this question, the following primary hypotheses were formulated:

H1: Users in the MR environment will demonstrate significantly greater accuracy during the avatar body alignment task compared to users in the VR environment.

H2: There will be a significant difference in the subjective SoE experienced by users in the VR environment compared to those in the MR environment.

In addition to these primary hypotheses, we also investigated several exploratory research questions concerning the role of proprioception, usability, and potential learning effects related to task completion time (RQ1-3).

To test these hypotheses, this thesis details the design and implementation of two novel lower-body control interfaces. A user study was then conducted to evaluate these interfaces and compare the effects of MR and VR environments on performance, embodiment, and usability.

1.4 Outline of the Thesis

The remainder of this thesis is structured as follows. Chapter 2 provides a comprehensive review of the literature on XR therapies, embodiment, and lower-body control. Chapter 3 details the methodology of the user study, including the experimental design, participants, and measures. Chapter 4 presents the results of the statistical analyses. Chapter 5 discusses the interpretation and implications of these findings. Finally, Chapter 6 concludes the thesis with a summary of the primary contributions and directions for future work.

Through these chapters, this thesis makes the following primary contributions to the field:

1. The design and implementation of two novel interfaces for lower-body avatar control and alignment.

2. Empirical evidence demonstrating the trade-off between performance and embodiment in MR versus VR for an alignment task.

3. An analysis exploring the role of proprioception and the nature of the user experience in MR environments.

Chapter 2

Literature Review

2.1 Introduction to the Literature Review

This chapter introduces virtual and mixed reality technologies through the lens of virtual therapies, modern applications and acceptance. The Chapter is divided into sections based on cohesive findings related to lower-body control, the importance of embodiment and alignment methods.

2.2 Foundations of Immersive Technologies

VR represents an immersive environment designed to replace the user's real-world surroundings with a fully digital one. The primary objective of these virtual spaces is to create the effect of an interactive, convincing, three-dimensional space not unlike the real world. When users use these devices and enter virtual spaces, they carry with them expectations of reality - and when executed, can convey the illusion of truly being 'inside' another world [11]. This *illusion of participation*, instead of simply being an observer, is what further emphasizes the power of the cognitive effect produced by the technology [32].

A defining characteristic of VR is the near complete occlusion of the real world. This is achieved through near-eye lenses inside the head-mounted display (HMD) in an attempt to fully control the user's vision. The system makes use of a variety of techniques to capture the user's intentions while inside the virtual environment. Inertial Measurement Units (IMUs) inside the controllers and headset detect motor movements and an array of cameras inside and outside of the HMD detect motion, capture the real world around the player, and track the controllers. All of this happens in real time in order to fully convince the user of their position in the virtual world. If any one aspect of this balanced system is out of sync, immersion can be disrupted and the experience could fall flat.

2.2.1 Mixed Reality and the Extended Reality Continuum

Extended Reality (XR), is an umbrella term used to describe the family of technologies that alter one's perception of reality by digitally adding or removing elements from the view of the real world. In short, these technologies blur the line between the real and the virtual world [48]. While exact definitions vary, XR is generally accepted to include three defining technologies: Virtual Reality (VR), Mixed Reality (MR), and Augmented Reality (AR). VR, MR, and AR each alter the perception of the real world through a wide variety of techniques. Consequently, each hold their own strengths and weaknesses that make them applicable in different scenarios.

AR is a technology that superimposes virtual objects over the real world [48]. Holograms and projectors are classic examples of AR in action. MR, on the other hand, varies slightly from AR in that it *allows not only the superposition of digital elements into the real world, but also their interaction* [48]. The Microsoft HoloLens is one of the most widely studied examples of a complete mixed reality system [34]. It functions similarly to a pair of glasses equipped with transparent lenses and embedded waveguides, allowing users to see digital holograms projected onto their real-world environment. The device enables digital interaction through built-in sensors, eye tracking, spatial mapping, and hand-tracking capabilities, enabling real-time interaction with virtual content while maintaining full awareness of the physical surroundings [34].

The Meta Quest 3 is a staple in the field of XR technologies. Not only does the HMD offer quality VR experiences, but it also offers a high-quality and affordable virtual passthrough mode as well. This feature enables users to see a real world projection inside the headset via front-facing colored cameras. This design also enables digital altering of the image captured by the cameras in real time. By combining this feature with depth sensing technology, virtual elements can be positioned in front of and behind the perceived real world objects, creating a truly immersive and unique experience [33]. Because of these features, the Quest 3 is an excellent choice for exploring virtual and real-world interactions and their effect on performance and user engagement.

2.2.2 Applications of XR in Medicine

Beyond entertainment, XR technologies are rapidly gaining traction as powerful tools in many fields, particularly within healthcare education and therapy [53]. These immersive systems offer novel solutions for rehabilitation, pain management, psychological treatments, and surgical training [37], leveraging their ability to create controlled and engaging environments [2]. For example, XR has been used in exposure therapies for treating PTSD as well as distraction-based pain management therapies. A powerful area of interest for applying these systems also in stroke and amputee therapies where patients can undergo virtual physical therapy or have virtual limbs superimposed over missing limbs.

Following a central nervous system injury such as a stroke or a peripheral injury such as a limb amputation, the brain undergoes significant change. This takes the form of a reorganization of areas that handle somatosensory and motor functionality [3]. Areas of the brain that once managed temperature and pain regulation for a now missing or paralyzed limb, are now no longer receiving the expected input signals. Due to the area of the brain now lacking any activity, surrounding areas slowly begin to invade the now dormant territory in a process called maladaptive plasticity [3]. The sensory motor mismatch theory aims to further expand upon this idea claiming that this lack of sensory feedback results in an 'error signal' in the form of chronic phantom limb pain (PLP) in amputees and motor dysfunction in stroke patients [3].

XR devices have been viewed as powerful instruments in addressing the sensorimotor mismatch. These HMDs allow for powerful and controllable environments and provide synchronous visual and auditory signals in situations where they would otherwise not be present. By providing these signals, it effectively *closes the loop* between intention and perception while promoting healing [3].

Mirror Therapy (MT), introduced in the early 1990s, serves as a precursor to these new techniques involving XR devices. MT uses a mirror positioned to reflect an intact limb, creating the illusion for a patient that they are seeing their missing limb and its corresponding movements [13]. While MT is widely accepted and used to treat PLP, it has its limitations. The use of a physical mirror limits mobility and any shifts in perspective or posture can break the illusion. Moreover, this approach limits the users movements to be exclusively symmetrical [18].

In contrast, a virtual limb created through extended reality (XR) mirror therapy overcomes many of these limitations. Depending on the system's implementation, the virtual limb can be controlled independently of the intact limb, offering greater flexibility. XR technology holds meaningful potential in the therapeutic space. It enables more immersive and interactive rehabilitation experiences while helping to resolve the conflict between missing visual feedback and movement intention [25].

2.3 Embodiment and Presence in Extended Reality (XR)

A common area of interest and experience pertaining to XR technologies and studies is Embodiment. Through the lens of XR, embodiment can be defined as the way individuals perceive their body in virtual space. In applications that induce greater levels of embodiment, interactions between the virtual and the physical world are seamless and carried out as though the perceived virtual body is one's own.

Research has examined the effects of embodiment in XR applications and have broken the concept down into three major parts:

- Sense of Body Ownership: The feeling that an external object or body part is an integral part of one's own biological self [47, 16].
- Sense of Agency: The feeling of being the author or initiator of and observed action [47, 16].
- Sense of Self-Location: The feeling of being physically located within the boundaries of a body representation [47].

Many studies have explored the relationship between these components and revealed a complex relationship between the three. For example, shifts in perspective—such as a first-person perspective (FPP) versus a third-person perspective (TPP)—affect body location and ownership. The effect is also present in shifts between FPP and HMD (VR) [47].

2.3.1 Body Representation and Proprioception

The core of the feeling of embodiment lies with the human brain and its ability for complex, parallel sensory processing [31]. This is known as multisensory integration and is fundamental to achieving a feeling of ownership over one's body and maintain spatial awareness [31] [12]. An early contribution, also referenced across many studies pertaining to multi sensory integration, is the Rubber Hand Illusion(RHI) [9]. This experiment involved placing a rubber hand in view of the participant while their real hand was hidden from sight. When both the rubber hand and the real hand were stroked simultaneously, participants reported feeling the touch on the rubber hand, demonstrating that visual and tactile stimuli can override proprioceptive signals and induce a sense of ownership over a non-biological limb. This experiment is telling in that in order to have a higher level of perceived body ownership, it is likely fundamental to have a synchronized inputs (sight, touch, sound, etc.). Another study explored the effects of the rubber hand illusion while inside an AR environment [45]. Instead of a physical "rubber" hand, an animation of a digital hand being stroked with a brush was rendered side by side a participants real hand. Reported body ownership resulted in showing signs of similar strength compared to the original study but still lesser than the physical rubber hand [45].

When visual, auditory, and tactile senses are effectively coordinated, studies have shown that this multisensory integration can help stroke patients rebuild damaged or lost sensory pathways [26]. The implications of this, as it pertains to XR, highlights some key variables to be aware of during these therapies in not just the environment but also the quality of the experience. Input latency, lag, and poor synchronization between a virtual avatar's movements and the user's can have the potential to distract the user and, consequently, significantly reduce the effectiveness of XR therapies [50]. Moreover, it has also been found that poor fidelity can negatively affect embodiment, further increasing the complexity of this problem [8]. Maintaining near-perfect synchronization between sensory inputs and outputs are crucial for achieving high levels of embodiment and introduce more complexities with virtual spaces where virtual and real elements merge [16].

Beyond therapeutic applications, several studies have investigated the extent that one's sense of embodiment (SoE) can be influenced across different immersive modalities. During one such study, virtual hands were superimposed on top of a participants real hands during a visuo-motor task in MR [20], as shown in figure 2.1. The task introduced both real and virtual interactions in order to invoke a SoE over the avatar. A widely used embodiment questionnaire [40] was utilized following a session and found mixed results regarding the real and virtual interactions with the avatar's hands. The results suggest that users experienced a stronger sense of ownership over virtual hands when both virtual and real objects were present, compared to when only the virtual hands were visible. Additionally, potential correlations between the avatar's visual content (i.e. hand model and texture) and immersion were identified and it was suggested that it be taken into account during embodiment experiments [20].

The ability of XR to induce an SoE has been identified as a fundamental reason for its efficacy in many therapeutic applications, particularly for conditions like post-stroke

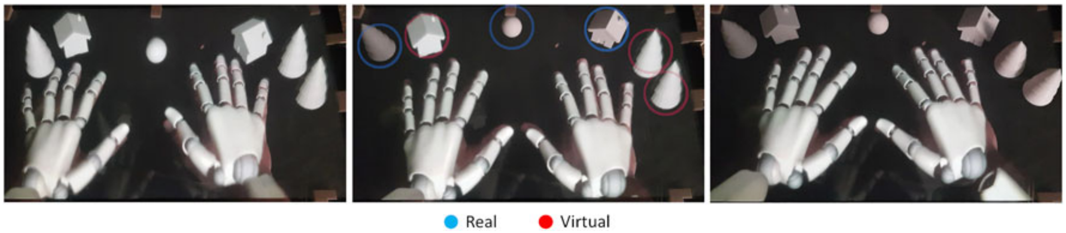

Figure 2.1: Genay, Lécuyer experimental design using superimposed virtual avatars during visuo-motor task [20]

motor impairment and PLP [1]. Studies and trials centered around healing these conditions have been shown to employ virtual bodies for patients to control. This directly addresses the aforementioned sensorimotor mismatch hypothesis. By enabling a patient to experience control over a virtual limb, XR therapies can further alter body representation by closing the loop between perception and intention [3].

Embodiment Across XR Modalities and Lower Body Focus

Studies like those by Skola et al. [45] (AR RHI study) and Genay et al. [20] (MR overlaid hand avatar) showcase the ability of XR systems (AR and MR, respectively) to induce feelings of embodiment. A comprehensive meta analysis of studies published between 2010 and 2021 by Mottelson et al., revealed that, in VR, the embodiment of virtual avatars generally leads to a modest sense of body ownership -a critical aspect of embodiment and XR [35].

A review of several studies related to lower-body avatar embodiment revealed that the presence of a virtual avatar representation can improve task performance in VR and assist with motor skill development [52]. A lower-body-focused VR study that measured foot accuracy placement showed that foot placement accuracy improved when using some form of virtual representation when compared to none at all [52]. This finding is intuitive, as a core feature of VR is the occlusion of the real world and any amount of occlusion would likely decrease accuracy relative to real world objects.

While VR excels in controlled, fully immersive experiences for targeted motor learning, MR offers distinct advantages for tasks requiring integration with the physical environment, particularly due to its ability to preserve real-world sensory feedback [7]. One example of this comes from the "self-avatar-follower effect" where the virtual avatar can influence the actions of the user and move beyond a simple unidirectional control model between the real and virtual [7]. This effect is similarly seen in visually guided motor behavior which is the driving motivation behind in some therapeutic procedures. The behavior of the user's body being drawn or influenced by the virtual was claimed to be the result of a *visuo-proprioceptive* coupling effect established between the physical and the virtual body during embodiment [21]. This bidirectional influence provides strong and unique opportunities for how MR/VR can be used for corrective motor behavior tasks or virtual therapies.

However, direct comparative studies, especially for specific tasks, has varying results pertaining to factors such as the degree of real-world occlusion as shown in the foot placement task [52]. A study comparing AR and VR for body weight perception noted that *"motor control effects were weaker than in similar VR studies"* in AR, *"likely due to the augmented reality setting, where the user's real body remains visible alongside the virtual one"* [27]. This suggests that the constant visibility of the real body in AR may weaken embodiment-related motor control effects, as suggested previously, by introducing visual mismatches.

2.4 The Challenge of Lower Body Embodiment and Control in XR

Phantom limb pain (PLP) can affect any amputated limb, but the residual limb's location and length often introduce unique challenges in pain management. Lower-limb amputees (LLAs), in particular, face additional difficulties related to gait (walking patterns), balance, stability, and fear of falling [13]. The most important of these functional movements is often considered to be walking, an inherently asymmetrical series of movements [18]. As mentioned previously, the ability to use asymmetrical movements during treatment is an area that XR therapies have potential to excel at when compared to traditional therapeutic counterparts such as MT.

Further analysis of the existing literature reveals disparity regarding XR and amputee therapies: Hali et al. explored a wide range of articles and found that 79 of the 109 patients included in examined studies targeted upper-body amputations and PLP [22]. Hali et al. claim this significant lack of representation *restricts the clinical utility of VR therapy* [22]. This under-representation is particularly interesting given that lower-body amputations are vastly more common than upper body amputations, making up nearly 85% of all amputations [39]. The distinct under-representation of lower-body amputees studied in VR therapies is likely a direct result of limited accessibility to effective lower-body XR therapies, ultimately stemming from the lack of robust and accessible lower-body control solutions. By providing an effective means of manipulating and aligning the lower-body of a user's avatar while using these XR systems, our research could drastically increase the presence and effectiveness of lower-body PLP studies and their broader clinical utility.

2.4.1 Approaches to Lower Body Tracking and Virtual Limb Control

Researchers have developed various techniques for achieving robust lower-body alignment and control in XR systems. These approaches often involve trade-offs in terms of precision, hardware complexity, cost, and their impact on user agency and overall usability.

Virtual Mirror Therapy: Mirrored Control from the intact limb enables an advanced, more flexible experience compared to traditional MT. A prominent example of this application is the Mr. MAPP (Mixed reality system for Managing Phantom Pain) system, which uses a Microsoft Kinect sensor to perform full-body tracking without requiring any sensors to be worn by the patient [4]. Users, while wearing an HMD, can see their full body avatar and engage in interactive activities. A pilot, at-home clinical trial of this system found it feasible and effective for reducing PLP[4]. This setup enabled control over a full body avatar, targeting sense of body ownership, and superimposed the avatar over the user's real body (such as a missing leg or arm), targeting sense of self location. However, similar to MT and other XR MT applications, the sense of agency was somewhat lacking, seeing as input relied on the opposite intact limb.

Independent Motion Tracking: This method provides more direct input scheme, by placing sensors, generally Inertial Measurement Units (IMUs), directly on the residual limb. For lower-body amputees, inputs are collected from movements about their stump (example: moving the thigh or swaying the hip) [3]. A recent clinical trial demonstrated that immersive VR interventions can significantly reduce phantom limb pain (PLP) in individuals with transtibial (below the knee) lower-limb amputation. Participants engaged in both passive and active VR treatments using the HTC Vive, with the active, movement-based sessions yielding a greater average pain reduction (39.6%) compared to passive ones (28%) [3]. These findings support the potential of targeted, game-like VR experiences as effective therapeutic tools for managing PLP.

Unlike the Mr. MAPP system, where users wore no additional sensors and relied on mirrored movements of the intact limb [4], this approach takes the opposite route—using direct input from the residual limb via wearable sensors. While Mr. MAPP prioritizes ease of use and accessibility by minimizing hardware, sensor-based methods are able to offer more agency by enabling the aforementioned synchronized movements of both virtual limbs. However, this approach, while more robust in terms of control, tends to compromise on usability and may face barriers to widespread, at-home adoption. Similarly, myoelectric control—using electromyography (EMG) signals from residual limb muscles—has also shown promise for enabling phantom motor execution and real-time limb control, offering another direct-input method with therapeutic potential for PLP [29].

2.5 The Impact of Immersive Environments on Performance and Embodiment

2.5.1 Performance in Avatar Alignment

When using a fully immersive virtual reality headset, such as the HTC Vive or Meta's Quest line, the HMD nearly fully occludes any visual information from the real world [30]. When a task demands precise body positioning or alignment, users must rely primarily on two sources of information: visual feedback from their virtual avatar and their proprioceptive sense [30]. For such tasks, VR creates a reliance on proprioceptive feedback. Much of the literature, however, suggests otherwise with VR being the driving factor that can degrade or disrupt this sense during tasks [49]. Immersive VR impairs the proprioceptive sense by creating sensory conflicts—especially when perceiving motion, leading to a disconnect between visual and proprioceptive inputs [49]. This is well documented and it is understood that the mismatch forces the brain to create predictions and estimate the orientation of the body and its surroundings [6, 30].

While evidence for inducing a SoE has been detected through studies comparing VR to traditional approaches, there are mixed and generally subtle findings in direct comparisons of embodiment between VR and MR/AR. This is likely explained by the general nature of MR. The constant visibility of the body, as claimed by Genay et al., is likely hindering the strength of the embodiment (SoE) illusion of superimposed avatars [20]. This conflict of discerning between the proprioceptive feeling of a body part and visual input of seeing the user's real body, forces the brain to recognize mismatch, likely creating an uncanny valley effect. For the amputee population using virtual therapies, this effect is likely less serious due to the absence of overlap. Instead, the lost sense of body ownership could be rediscovered in the virtual 'phantom limb'

Mixed Reality systems, particularly through modern HMDs equipped with video pass-through technology, such as the Meta Quest 3, offer a fundamentally different sensory experience by allowing virtual and visual data to work in tandem [41]. In the context of body alignment tasks, unlike VR, users have the capacity to see and work around their physical surroundings, which theoretically should increase spatial accuracy [49]. This is achieved through the availability of visual feedback which should improve the accuracy of precise movements [49] through a process called sensorimotor integration [44].

Despite the strong theoretical argument for MR's superiority in body precision related tasks, direct performance comparisons between VR and MR yield mixed results. Instead, when it comes to precision-based tasks, findings appear to be dependent on the specific task being performed. These results were found while assessing short-term spatial memory between VR, AR, and MR environments. The results yielded no significant difference in task accuracy between the different immersion modes, but the study did find that the MR group required significantly more time to complete the tasks [41]. On the other hand, a separate study focused around a pick-an-place task found the opposite result where the MR group performed much faster compared to VR [28]. These differences likely arise from the task itself instead of the modality. The authors understood this as well, concluding that the key to higher performance in the pick and place task was understanding the physical space surrounding the task [28].

These mixed findings highlight a gap in the literature: There are few, if any, studies that focus on a precise body alignment task and how accuracy is compared between VR and MR environments. By directly measuring perceived virtual and physical alignment between these two immersive states, our study should enable a more in-depth understanding of the impact of immersion mode on alignment accuracy. This leads to the formulation of the primary performance hypothesis: **H1: It is hypothesized that users in the Mixed Reality environment will demonstrate significantly greater accuracy during the avatar body alignment task compared to users in the Virtual Reality environment.**

2.5.2 Sense of Embodiment

SoE has been shown to be greatly influenced by XR environments [27, 52]. While the effect is known, inconsistencies are still prevalent between studies. As previously discussed, while incorporating both virtual and physical objects in an MR task (e.g., as in the visuomotor task by Genay et al. [20]) can increase a greater sense of body ownership, the additional real world elements can introduce complexities. The constant view of ones real body brings about visual mismatch between what is felt and what is seen [20].

On the other hand, the complete visual control and occlusion, offered by VR, might present a more coherent and less conflicting visual result, leading to a greater SoE. The current study is designed to investigate these differences by focusing solely on the alignment of the user in the real and the virtual space. This concentration on the overlap should offer a more targeted assessment of the complex relationships between the immersive environment and the influence it has on embodiment. This leads to the formation of hypothesis 2: **There will be a significant difference in the subjective SoE experienced by users in the VR environment compared to users in the MR environment during the avatar control task.**

2.5.3 Visual versus Internal Feedback

The influence of proprioceptive senses has been linked to biases in visual feedback, particularly among individuals with Body Image Disturbance (BID) disorders (e.g., body dysphoria, anorexia nervosa) [15, 19]. For these individuals, strong internal predictions about one's own body can sometimes override conflicting visual and tactile cues, similar to those who had little to no reactions from the rubber hand experiments[42, 45, 9]. Interestingly, research has shown that individuals with decreased interoceptive processing—the ability to sense internal bodily signals such as heartbeat, hunger, or breath— abilities, are often more susceptible to embodiment illusions [38]. This may be because weaker internal signals shift the user's reliance toward external cues like visual information, potentially making them more responsive to virtual body ownership effects.

However, not all research supports a purely vision-dominant account. Barca and Pezzulo [5], for instance, propose an active inference perspective in which individuals with BID may instead rely more heavily on top-down expectations than on sensory inputs, *regardless of modality*. In such cases, modifying visual feedback may not fully resolve the internal body conflict. Immersive systems that hide areas of discomfort or modify the expected feedback loop can temporarily suppress discomfort [42], offering a therapeutic opportunity for virtual embodiment in clinical populations. These mixed findings suggest that task performance or obtaining a strong SoE, might be dependent the innate processes that dictate how external inputs are perceived.

2.6 Proprioception and Alignment Accuracy

Proprioception is critical for motor control, and can be defined as the innate awareness of the spatial and mechanical status of one's body [49]. Proprioceptive abilities have also been shown to vary significantly between individuals [49]. As established previously, XR devices have demonstrated the potential to disrupt this sense and influence the actions of the user [30]. The common theme has been a measurable decrease in the accuracy of *proprioceptively-guided* movements particularly within VR environments [6]. This highlights several key considerations in this area: First, individuals have varying baseline proprioceptive abilities. Second, while VR tends to disrupt proprioception, MR appears to rely more heavily on visual feedback, thereby diminishing the reliance on proprioceptive abilities for spatial tasks. The question remains: is there a relationship between proprioceptive abilities and alignment-based task performance between immersive environments (RQ1)?

2.7 Summary of Key Findings and Identified Gaps

1. **Embodiment & XR Foundations**

- **Foundations of XR & Medical Applications:** XR technologies (VR, MR) provide immersive, customizable environments for rehabilitation and therapeutic interventions, particularly for sensorimotor mismatch conditions.
- **Embodiment Theory & Role in Therapy:** Core embodiment components (i.e. ownership, agency, self-location) are essential to XR's therapeutic efficacy and influence neuroplasticity and mismatch resolution.

2. **Lower Body Embodiment & Control Interfaces**

- **Lower Body Embodiment & Control Challenges:** Most XR research focuses on upper-limb embodiment; the lower-body poses unique challenges (biomechanical, perceptual), particularly for underrepresented groups like lower-limb amputees.
- **Existing Control Approaches:** Current lower-body control methods (e.g., motion capture, IMUs, mirrors) are often complex, expensive, or reduce user agency, especially in home or clinical use.

3. **Comparative Effects of VR vs. MR**

- **Embodiment Comparison (H2):** Mixed results in literature regarding whether MR or VR leads to stronger embodiment; real-world occlusion and task specifics influence outcomes.
- **Performance Comparison (H1):** MR may provide visual anchoring advantages for precise tasks like alignment, but performance benefits remain unclear across studies.
- **System Usability (RQ2):** Usability outcomes vary between VR and MR depending on task complexity, cognitive load, and interaction methods.

4. **Individual Differences**

- **Proprioception & Alignment Accuracy (RQ1):** Individual proprioceptive skill impacts alignment accuracy in XR; the relationship may differ between VR and MR, warranting targeted investigation.

2.7.1 Contribution of the Current Study

This study contributes to the field by investigating task alignment performance, sense of embodiment, the effect of proprioception, and the usability of two novel lower-body control interfaces within both virtual and mixed reality environments.

Chapter 3

Methodology

3.1 Research Question

The study was focused on the usability of two lower-body control interfaces, the performance of users in aligning a virtual body with their real body in Mixed Reality (MR) environments, and the resulting sense of embodiment throughout the session. This leads to the primary research question:

How do body alignment method and immersive environment (MR vs. VR) impact user preference and performance in virtual avatar control?

In order to address this question, the following outlined experiment was conducted:

3.1.1 Hypotheses

Based on the research question, "How do body alignment method and immersive environment (MR vs. VR) impact user preference and performance in virtual avatar control?", the following hypotheses will be investigated:

H1: **Performance in Avatar Alignment:** Users in the Mixed Reality (MR) environment will demonstrate significantly greater accuracy during the avatar body alignment task compared to users in the Virtual Reality (VR) environment.

- **Rationale:** In the avatar body alignment task, users are required to superimpose a virtual avatar over their real body. While VR users must rely primarily on proprioceptive cues, due to the lack of direct visual feedback of their own body, MR users can leverage direct visual feedback of their real body alongside the virtual avatar. This visual information in MR is expected to significantly enhance spatial alignment capabilities and reduce reliance on potentially disrupted proprioception within an immersive headset.

H2: **Sense of Embodiment:** There will be a significant difference in the subjective sense of embodiment (SoE) experienced by users in the Virtual Reality (VR) environment compared to users in the Mixed Reality (MR) environment during the avatar control task.

- **Rationale:** While immersive technologies generally foster greater embodiment than their non-immersive counterparts, direct comparisons of SoE between different immersive conditions (VR vs. MR) have yielded mixed and often subtle results in the existing literature. This hypothesis seeks to further investigate whether the distinct characteristics of full immersive VR (e.g., complete visual occlusion of the real body) compared to MR (e.g., blending real and virtual elements, and the persistent visibility of the real body) lead to a significant difference in the subjective feeling of embodiment.

3.1.2 Exploratory Research Questions

Beyond the primary hypotheses, this study also aims to explore additional relationships in terms of lower-body avatar control, usability, and the influence of baseline proprioception. To that end, and to increase understanding of the field as a whole, the following exploratory research questions will be investigated:

RQ1: **Proprioception and Alignment Accuracy:** Is there a relationship between a participant's baseline proprioceptive accuracy and their performance (alignment accuracy) during avatar control tasks in immersive environments? Does this relationship differ between the Virtual Reality and Mixed Reality environments?

- **Rationale:** Prior research shows that proprioceptive accuracy can be disrupted in immersive virtual reality. This question seeks to understand if any relationship between proprioceptive abilities and alignment performance are influenced by the distinct characteristics of the immersive environment (VR or MR).

RQ2: **User Preference (Overall System Usability):** Are there differences in perceived overall system usability (measured via System Usability Scale scores) between the VR environment and the MR environment during an avatar alignment task?

- **Rationale:** Both lower-body control interfaces were available and demonstrated to all participants within their assigned immersion environment. This question aims to explore the perceived usability of interfaces and how they compare between immersion modes (VR vs. MR) during the body alignment tasks. Given the fundamental differences in visual feedback (full occlusion in VR vs. passthrough in MR) and cognitive load, between the environments, an understanding of perceived overall system usability within these distinct contexts is valuable, given the mixed comparative findings of previous investigations.

In order to evaluate the claims listed above, the following experiment was conducted:

3.2 Experimental Design

The experiment was a mixed-factorial design incorporating both between-subjects and within-subjects factors. The between-subject factor is the immersion mode: Mixed Reality or Virtual Reality. This experimental design will be used to evaluate the impact of the different lower-body control methods and immersive environments on virtual avatar alignment and user experience.

The **between-subjects factor** was the **Immersion Environment**, with participants randomly assigned to either a Mixed Reality (MR) or a Virtual Reality (VR) condition. This factor was critical for investigating the role of visual feedback in lower-body avatar alignment. In a virtual environment, vision of the user's body is completely occluded. This allows for an exploration of how users rely on proprioception - similar to how amputees experience their phantom limb. On the other hand, the MR environment provides a digital overlay of virtual objects on top of the real world. This is referred to as a **passthrough** view and enables users to rely on direct visual cues from their surroundings.

The **within-subjects factor** was **Pose Type**, with participants required to achieve three distinct poses: sitting, standing, and laying. These specific poses were chosen because they represent common postures adopted by users in various VR applications and are particularly relevant for virtual rehabilitation therapies, especially for bedridden patients. By providing multiple common poses during the experiment, this will allow a rigorous test for the lower-body control interfaces' efficacy across a variety of lower-body configurations. To mitigate the order of the effects from each pose, the sequence was assigned to each participant using a Latin Square design.

Table 3.1: Independent Variables

Variable Name	**Possible Values**
Immersive Environment	Mixed Reality (MR), Virtual Reality (VR)
Pose Order	Seated, Standing, Laying - Counterbalanced using a Latin Square design

3.3 Participants

The study utilized the University of Georgia's student research participation system, SONA, recruiting from the Department of Communication Studies and the College of Journalism. This portal briefs students with a short description of various studies and the expected time commitment for participating. Students within the department sign up for time slots to participate in studies and, in return, fulfill course requirements or receive extra course credit, per instructor discretion.

Forty student volunteers from the UGA College of Journalism participated in the study. Eligibility criteria for participants include:

- No physical disabilities (e.g., amputations).
- No impairments to upper or lower-body motor function.
- Corrected-to-normal or 20/20 vision.
- Fluent in Written and Spoken English.

Exclusion Criteria: Participants prone to motion sickness, epileptic seizures, or vertigo, or those unable to perform the physical motions required to use the technologies, were excluded from the study. Study participants should be in good general health and have 20/20 vision or wear corrective lenses that do not impair the use of a virtual reality headset.

3.3.1 Ethical Considerations

This study was approved by our institution's Institutional Review Board (IRB), and all participants provided informed consent before the study began. Participants were informed that, if at any moment during the study, they felt nauseous or wished to stop, they were allowed to do so at no penalty or loss of incentive credit. To ensure confidentiality, all collected data were anonymized using participant IDs and stored on a secure, encrypted database.

3.4 Apparatus and Materials

3.4.1 Hardware

The experiment used the Meta Quest 3 and the Meta Quest 3 Touch controllers. The Quest 3 was chosen for this study due to its high-fidelity, color passthrough support for MR environments. The Quest 3 is also widely available and is a likely candidate for use in potential virtual therapies. Moreover, the Quest 3 offers high quality hand tracking which was used to measure the baseline proprioceptive ability of participants. In addition to the headset, a laptop was used to interact with the web dashboard that drove the program on headset and also recorded participant registration and data logging. Additionally, a tablet was used for the post-session survey.

For the experiment, a standard office chair was used for the seated and the laying task. The laying task also made use of an ottoman that participants reclined their feet onto to achieve the laying down pose. For the proprioception measurements, a white table was used for participants to lay their hands on during the procedure.

Lower Body Control Interfaces

The experiment featured two new lower-body control interfaces used to align the virtual avatar with the user's real body. Both interfaces were available for use simultaneously during each participant's trial. To clarify, the interfaces themselves are *not* an independent variable for the study and participants could freely use either interface to achieve the alignment task.

The first of the two is the *Direct Interface* which enabled users to directly interact with their virtual legs. This features additional spheres superimposed over the avatars hips, knees, and ankles. These joints are interactable and drive the virtual avatar's position and orientation. Users can interact with these by holding the primary trigger of the controller and dragging them through the scene.

The second lower-body control interface is the *Puppet Interface*. A smaller set of legs floated near the player and functioned as a virtual digital twin of the players virtual legs. Any changes in orientation and position of the puppet were also reflected onto the player's avatar.

Virtual Environment

At the start of the session, participants that receive the VR immersion condition were placed in a plane, empty VR world with a clear blue sky. They stand on a small brick platform and are presented with a large floating menu showing instructions for the current assignment. This includes instructions for the proprioception and the pose tasks. The menu also displays important information pertaining to using the lower-body control interfaces and how to use them. After a pose is completed, the menu will also display the in-headset embodiment survey question which users can answer using the Meta Touch controllers and a laser.

3.5 Procedure

3.5.1 Arrival and Instructions

This study investigated the effects of two lower-body control interfaces on task performance and the participants' sense of embodiment during a series of simple body posing task. After arriving at their registered time slot and filling out a consent form, participants were randomly assigned to either the MR or VR immersion condition (the between-subjects factor). Participants in the MR environment could see virtual objects, such as the avatar, overlaid over the real world using the Meta Quest 3 colored passthrough feature. Participants given the VR environment experienced full immersion.

Next, participants were informed of a quick background of the study and an overview of the tasks they were to complete. They were also informed on how to use the headset and the controllers during critical sections of the study.

3.5.2 Baseline Proprioception Task Procedure

Participants were then asked to complete a simple limb position matching task in order to calibrate their baseline proprioceptive ability (covariate). While wearing the VR headset, and not holding controllers, users were instructed to place their left hand palm down on the center of a table. With their eyes closed inside the headset, they would then slide their right hand forward across the table until they felt that both of their hands matched in distance from their body (z-axis). At this point, users would let the proctor know and the positions of both hands would be automatically recorded by the headset's hand tracking abilities. This was repeated two more times for a total of three measurements per participant.

3.5.3 Posing Task Procedure

After the baseline proprioception test,participants were handed the controllers and began the main experimental trials, starting with their first assigned pose (e.g., seated, standing, or laying down). This required them to position their virtual avatar to match the perceived position and orientation of their physical body. For example, if their first assigned pose was *seated*, then they would need to then adjust the hips, knees, and ankles of their virtual avatar to mimic a seated pose. Because the study focuses on comparing the posture of the virtual avatar with the user's real body, no strict pose guidelines were enforced. Instead, participants were given general guidelines to follow when achieving the poses and brought their virtual avatars to meet their pose. A simplified visualization of this is shown in figure 3.1.

After the participants felt that they were aligned to the best of their abilities, they would inform the study proctor. At which point, the proctor would then trigger the system to capture the positions and rotations of the virtual avatar through use of the web dashboard. Next, participants were asked to sequentially align the Meta Quest 3 controllers with anatomical landmarks on their person. The sequence of these landmarks is as follows:

- **Lateral Hip (Greater Trochanter):** Participants placed the controller on the prominent bony point on the outermost side of their hip.

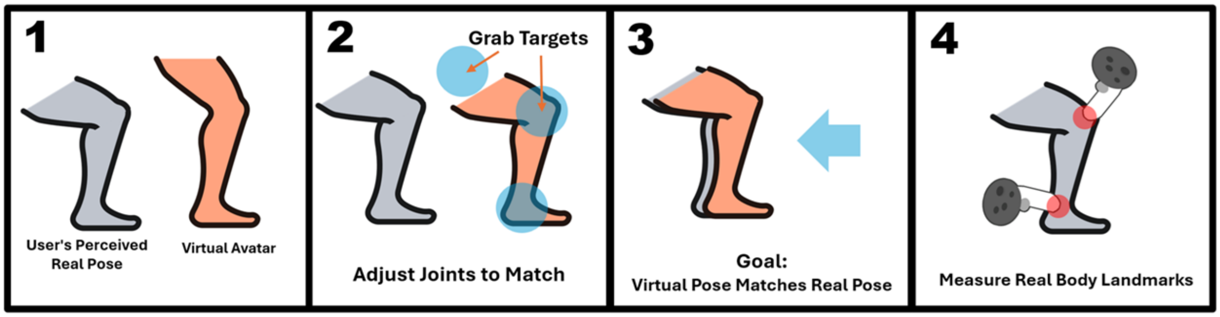

Figure 3.1: Simplified Pose Alignment Procedure

- **Anterior Hip (Anterior Superior Iliac Spine - ASIS):** Participants located the front-most point of their pelvic bone and placed the controller there.
- **Lateral Knee (Lateral Femoral Epicondyle):** Participants placed the controller on the bony prominence on the outer side of the knee joint.
- **Anterior Knee (Center of Patella):** Participants placed the controller on the center of their kneecap.
- **Medial Knee (Medial Femoral Epicondyle):** Participants placed the controller on the bony prominence on the inner side of the knee joint.
- **Lateral Ankle (Lateral Malleolus):** Participants placed the controller on the bony bump on the outer side of their ankle.
- **Anterior Ankle (Midpoint):** Participants placed the controller on the front of their ankle, centered between the inner and outer ankle bones (malleoli), typically over the underlying joint line. **Medial Ankle (Medial Malleolus):** Participants placed the controller on the bony bump on the inner side of their ankle.

- **End of Shoe:** Participants placed the controller against the very tip of their shoe (near the toes).

Upon completing the sequence, the proctor would then use the web dashboard to proceed to the next pose task.

Data Acquisition

To recall, each participant was tasked with using lower-body control interfaces to reposition their virtual avatar so that it aligned with how their actual bodies were positioned in the real world. In Unity, several game objects make up the player's avatar rig, each with their own transforms. These transforms contain 3D vectors representing their local position, local rotation, and local scale. In order to determine the avatar's position in space, all bone positions that make up the rig were captured at the completion of each pose.

Each participant was equipped with Meta Quest 3 Touch controllers and instructed to position the tips of the controllers according to the positions outlined in Section 3.5.3. The two hip measurements (per side) had their positions averaged, resulting in a generic hip position for both the left and right sides. This was also the case for the three points gathered for both the knees and ankles. Altogether, there were a total of eight final joint transforms that were extrapolated for use in the data analysis. These points represented the hip, knee, ankle, and toe for both legs 3.2.

These eight joint positions—`hip_L`, `hip_R`, `knee_L`, `knee_R`, `ankle_L`, `ankle_R`, `toe_L`, and `toe_R`—were selected to represent the critical points of lower-body posture. Each label corresponds to a left or right-side anatomical reference and was averaged or captured as described earlier.

3.5.4 Post-Task Questionnaires and Debriefing

After the last pose was completed, the session concluded and the participant removed the headset. A post-session survey was then administered to each participant. This survey was

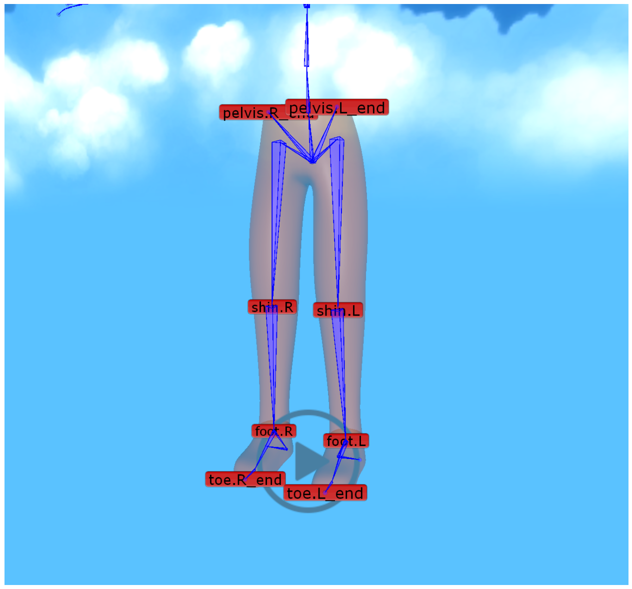

Figure 3.2: This shows an in-unity view of the player's virtual avatar along with the avatar bones inside its rig. Additionally, the eight joint transforms are labeled. Note: In the Unity rig, the bone transforms are named 'shin' and 'foot', but they correspond to the position and behavior of the knee and ankle joints, respectively.

hosted on Qualtrics, an online survey and data management system, and included several sections to collect subjective data:

- **System Usability Scale (SUS):** This questionnaire [10] assesses the perceived usability of the overall system.

- **Embodiment Scale:** The post-session Embodiment Scale [40] was used to capture participants' overall sense of embodiment with the virtual avatar. Disembodiment, Transformation, and Body Ownership subscale calculations can be derived from this survey to capture a more in-depth picture of SoE.

Participants were then debriefed about the study's objectives and thanked for their participation.

Table 3.2: SUS Questions and Their Descriptions

Q#	**SUS Question Summary**
Q1	I found the system unnecessarily complex.
Q2	I thought the system was easy to use.
Q3	I would need support of a technical person.
Q4	The various functions were well integrated.
Q5	There was too much inconsistency in the system.
Q6	Most people would learn to use this system very quickly.
Q7	I found the system very cumbersome to use.
Q8	I felt very confident using the system.
Q9	I needed to learn a lot before getting going with this system.
Q10	I think that I would like to use this system frequently.

3.6 Measures

This section clearly lists the dependent variables measured throughout this study. It is divided into subsections pertaining to the quantifying alignment, proprioception, and subjective measures such as sense of embodiment and system usability.

Below is a high level overview of the dependent variables 3.3

3.6.1 Alignment Accuracy

Alignment accuracy was measured through **positional error** and **rotational error**. These were both calculated separately based off the point cloud data collected from the **real pose data** and the **avatar pose data**. Additionally, it was decided that all joint transforms except the toes would be used during the data analysis. This was decided prior to analysis due to the inability for users to adjust toe positions during the alignment task. Moreover, these landmarks were also subject to confounding errors from variations in foot and shoe size

Table 3.3: Dependent Variables

Variable Name	Description
Task Performance (Accuracy)	Alignment accuracy: positional and rotational error between virtual avatar and user body
Task Performance (Time)	Task completion time for body alignment trials
Sense of Embodiment	Post-session Embodiment Scale [40]
System Usability	System Usability Scale (SUS) questionnaire [10]
Baseline Proprioceptive Ability	Mean offset (error) between the positions of the left and right hands during proprioception test

not accounted for by the avatar model. Removing them provides a more accurate measure of the user's alignment performance based on the joints they could actively manipulate.

Justification for Procrustes-Based Positional Error Analysis

To accurately quantify the positional alignment error between the avatar and the participant's real body, several considerations needed to be made and accounted for before. While a metric such as Mean Per Joint Position Error (MPJPE) is a reliable and widely used method of measuring the absolute positional error between the two point clouds [46], it is sensitive to *global translational and rotational differences* [46]. Tracking inaccuracies or system latency could introduce large positional offsets, leading to an inaccurate representation of the pose's internal geometry which would have large effects on the resulting error calculation. This would result in an inaccurate representation of the internal relationships between points - which are essential in measuring pose [46]. This calls for a method that can isolate the internal geometric relationship of the pose itself.

To address this, Procrustes Analysis (PA) was selected as the foundational method for calculating the positional error of each participant. PA is a statistical shape analysis technique designed to *optimally superimpose two sets of points by systematically removing confounding components of translation, rotation, and uniform scaling* [51]. The resulting data analysis pipeline, using PA, results in the **Procrustes-Aligned Mean Per Joint Position Error** (PA-MPJPE), which is calculated after the avatar's joint configuration has been optimally aligned to the *real body's* configuration [46].

This approach yields an advantage in that it works around any global positioning, rotational, or scale-related error [46, 51]. This is crucial for the accomplishing the objective of this study in that it isolates the virtual avatar's intrinsic shape (or pose) and returns the local alignment error [46]. To clarify, this process does not change internal scaling between measured joint positions the relative positioning and rotation is maintained. By not using this approach to error analysis, the potential to encounter global offsets from tracking or performance issues increase and results in an inaccurate representation of the users pose-matching performance.

Moreover, PA addresses a critical flaw in the design of the study: avatar scaling between participants. A limitation of the avatar design was the absence of dynamic scaling to match individual participant anthropometry. The current experiment does *not* provide functionality to scale the avatar or parts of the avatar. PA compensates for this to some degree by *uniformly* scaling the entire joint point cloud (for a single pose) to best match the avatar's overall size. Given that people have varying lengths of legs, shins, etc. and that they are not uniform, the assessment of positional error effectively asks: ”If we could make the avatar the same general size as the participant, how well does its shape align?” This allows a more fair comparison of the positional alignment accuracy between a tall participant and a short participant, removing the error component that is just due to a simple, global size differences.

Positional Alignment Measurement Pipeline

In this study, positional error is defined as the Procrustes Aligned Mean Per Joint Position Error (PA-MPJPE) of the difference between the virtual avatar and the user's real body's position in space. For all complete sets of avatar and real positional data, the aforementioned Procrustes alignment was applied to account for global systematic errors. Afterwards, the MPJPE was calculated from the resulting altered positions of each joint set. This results in a singular error associated for an individuals attempt at a pose task.

Rotational Alignment Measurement Pipeline

The calculation for rotational alignment error is more involved. In this study, rotational error is defined as the Root Mean Squared Error (RMSE) of the angular difference (in degrees) between the direction vectors of corresponding bones between the real-world skeleton and the virtual avatar. Each of the joints defined and measured in this study (i.e. hip, knee, and ankle) were paired up with their anatomically correct neighboring joint to create a *joint pair* (ex: left knee - left ankle). The full list of pairs are listed in table 3.4.

Table 3.4: Anatomical joint pairs used for rotational alignment analysis. The resulting shape essentially corresponds to a bone or body segment.

Joint Pair	**Description**
hip_L – knee_L	Left thigh
knee_L – ankle_L	Left shin
hip_R – knee_R	Right thigh
knee_R – ankle_R	Right shin

Each joint in the pair is represented by a 3D coordinate. For the virtual avatar, this was captured directly from the Vector 3 transform position in the Unity bone hierarchy. These are located in the exact center of the virtual avatar's joints. For the real avatar, as mentioned previously, these are derived from the average of three measurements taken around the matching joint on the participants real body. The joint pair is the resulting vector between the two positions, pointing away from the torso.

After the joint pair vectors were determined, the corresponding avatar and real vectors were first normalized to unit vectors and then compared using the angle between them. This angular difference, measured in degrees, quantifies the rotational alignment error between the participant's actual limb orientation and the virtual avatar's bone orientation. For each joint pair and *pose*, the angle between the real and avatar vectors was calculated using the **arccosine** of the dot product of their normalized vectors. Finally, the root mean square error (RMSE) of these angle differences was computed for each participant, pose, and immersion condition. This RMSE serves as a single metric representing the average magnitude of rotational misalignment across all joint pairs within a pose for a single user.

It should be noted that this method of angle calculated (using a direct angular difference between the two vectors) was chosen over a more complex 3D orientation analysis (such as using quaternions). This was chosen because the novel lower-body control interfaces used in this study prevented the user from twisting the virtual avatar's joints. This did not, however, interfere with the assigned pose alignment tasks. Future work into this area could potentially investigate aligning the twist of virtual limbs and real limbs.

$$\theta_i = \cos^{-1}\left(\mathbf{u}_{\text{real},i} \cdot \mathbf{u}_{\text{avatar},i}\right) \cdot \frac{180}{\pi} \tag{3.1}$$

$$\text{RMSE}_\theta = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\theta_i^2} \tag{3.2}$$

Table 3.5: Description of variables used in rotational alignment error calculation.

Symbol	Description
$\mathbf{u}_{\text{real},i}$	Unit vector of joint pair i from the real-world skeleton
$\mathbf{u}_{\text{avatar},i}$	Unit vector of joint pair i from the avatar skeleton
θ_i	Angular difference (in degrees) between real and avatar joint pair i
N	Total number of joint pairs for a given user and pose
RMSE_θ	Root mean square error of angular differences across all joint pairs
$\cdot$	Dot product between two unit vectors

3.6.2 Proprioception Baseline Score Calculations

To establish a baseline measure of proprioceptive ability, a hand-position matching task was administered at the start of each session. Participants placed their left hand on a table and, with their eyes closed, attempted to align their right hand with it. While this task was performed three times, only the data from the first trial was used for analysis to ensure the measurement reflected the participant's baseline ability without practice effects.

For this trial, the 3D positions of all ten fingertips were captured via the headset's hand tracking. The proprioceptive error was calculated as the distance between each corresponding finger pair (e.g., left index to right index) projected onto the user's forward-facing axis. This was achieved by calculating the dot product of the vector between the two fingertips and the camera's normalized forward vector, with the result converted to centimeters. Mathematically, the projected distance, d_i, for a single finger pair i is given by Equation 3.3:

$$d_i = (\mathbf{p}_{\text{L},i} - \mathbf{p}_{\text{R},i}) \cdot \mathbf{u}_{\text{forward}} \tag{3.3}$$

where:

$\mathbf{p}_{\text{L},i}$ is the 3D position vector of the left fingertip for finger pair i.

$\mathbf{p}_{\mathrm{R},i}$ is the 3D position vector of the right fingertip for finger pair i.

$\mathbf{u}_{\mathrm{forward}}$ is the user's normalized forward-facing unit vector.

To create a stable final metric, the error score was calculated by averaging the distances from the index, middle, and ring finger pairs; the thumb and pinky were excluded as their positions on the edge of the hand resulted in higher measurement variance. This final proprioceptive error score, E_{proprio}, is given by Equation 3.4:

$$E_{\mathrm{proprio}} = \frac{|d_{\mathrm{index},t=1}| + |d_{\mathrm{middle},t=1}| + |d_{\mathrm{ring},t=1}|}{3} \tag{3.4}$$

A visualization of the x-z projection is shown in figure 3.3

3.6.3 Subjective Measures

The exit survey includes three short questionnaires for the purpose of supporting H2's claims regarding embodiment as well as RQ2's investigation on system usability and user preference. SoE will be measured through Peck et al.'s proposed embodiment questionnaire [40]. This questionnaire is composed of ten questions relating to SoE and were designed to explore the components that influence embodiment. The overall system usability was measured using the System Usability Scale (SUS) curated by Brooke et al [10]. To create a single usability metric for each participant, the scores from the two separate SUS surveys—one for the Direct Interface and one for the Puppet Interface—were averaged together. All responses were collected using a seven-point Likert scale with the following options: *Strongly Disagree*, *Disagree*, *Somewhat Disagree*, *Somewhat Agree*, *Agree*, and *Strongly Agree*. For the SUS survey, participant responses on the seven-point Likert scale were mapped to a five-point scale, and SUS scores were computed and normalized to a range from 0 to 100.

3.6.4 Session Metrics

In addition to the primary measures, the time required to complete each pose alignment task was recorded from within the headset and controlled via the web dashboard. For each

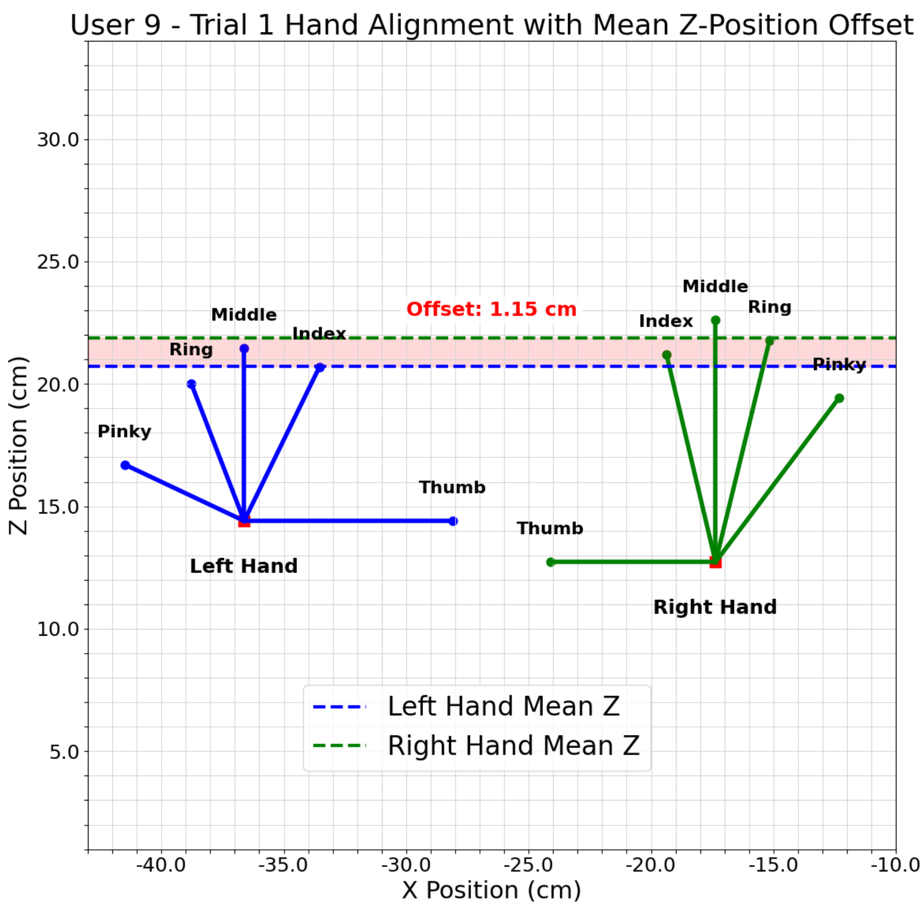

Figure 3.3: User 9's fingertip positions and thumb bases in the XZ plane for Trial 1. Colored lines connect fingertips to thumb base estimates (red squares). Dashed lines show mean Z-positions of index, middle, and ring fingertips per hand. The red shaded area between these lines highlights the mean Z-offset representing the proprioception error score for the respective trial. Note: lines connecting fingertips to the thumb base are estimates and used purely for visualization purposes.

pose, timing began when the proctor instructed the participant to begin and pressed the corresponding button on the dashboard. The timer stopped when the participant indicated they felt aligned and ready to proceed, at which point the proctor ended the timer and initiated the measurement process. This procedure was designed to exclude any setup time or transitions between poses, ensuring that only the active alignment phase was captured.

3.7 Data Pre-processing and Screening

Prior to the main statistical analysis, the collected data underwent a rigorous screening and pre-processing procedure to ensure validity and reliability. This involved the following steps:

3.7.1 Data Exclusion and Outlier Handling

The following criteria were established for data exclusion prior to the main statistical analysis:

- **Incomplete Task Data:** Any participant's data was excluded from performance-related analyses if they were unable to complete all experimental tasks due to technical failures, system errors, or withdrawal from the study.

- **Failure to Meet Screening Criteria:** Data was excluded if a participant, post-screening, was found to have a condition that violated the initial inclusion criteria (e.g., a motor impairment or a high susceptibility to motion sickness that affected their ability to perform the tasks as intended).

- **Procedural Non-Compliance:** Individual trials were excluded from the analysis if the proctor noted that the participant did not correctly follow the instructions for that specific trial.

- **Physiologically Implausible Poses:** Resulting pose data was visually inspected for gross anatomical errors or physiologically implausible configurations (e.g., joints bending beyond their natural range of motion). Such trials, reflecting either a significant tracking error or a failure to follow instructions, were discarded.

- **Statistical Outliers:** Outliers in the recorded performance metrics (e.g., alignment accuracy, task completion time) were identified. An outlier was defined as any data point that fell more than three standard deviations ($\pm 3SD$) from the mean of its respective condition.

Proprioception Exam Data Screening

While three trials of the proprioceptive hand-matching task were administered, only the data from the first trial was used to calculate each participant's baseline score. An initial inspection of the data revealed a pattern indicative of a learning effect. The first trial yielded errors with a relatively normal distribution, whereas the second trial showed the highest variance. Accuracy then improved in the third trial, suggesting participants were becoming accustomed to the task. To ensure the score reflected a true baseline of proprioceptive ability, free from the influence of practice or learning effects, the second and third trials were excluded from the final analysis.

3.7.2 Joint Position Pre-processing

The set of anatomical landmark positions (joint positions) was collected from each user during each pose, resulting in six sets of joint transforms per participant. Before the datasets could be used for the Generalized Procrustes Analysis (GPA), data cleanup was performed for each individual set of points. This process determined whether or not measured joint positions were feasible. Due to inconsistencies with the controller tracking, some of the measured real points were extremely inaccurate and considered to be a bad data point. This was made obvious by measured points that were located at the origin (0,0,0) and were unrealistic due to anatomical limitations (i.e. ankles being above knees, recorded knees bent at impossible angles, etc.).

When measuring the participants real body joint positions, the procedure called for multiple measurements for the hips, knees, and ankles. This resulted in two measurements (per side) for each hip and three for both the knees and ankles. To calculate the final positions of the left and right hips, a simple average was taken between the two hip measurements. If one of the two hip measurements was missing or deemed invalid, then the remaining valid point was used to represent the approximate joint position. If both points were missing, then the joint position received a **null** value.

For the knees and ankles, a circle was fit using the three points measured. The center of this derived circle then represented the center of the knee/ankle to be used in the analysis. This method was chosen to reflect the circle-like shape of the users' limbs. If fewer than three points were present, then a simple average was calculated. Similar to the hips, if no points were present for that joint, it would receive a **null** value.

In order to separate out the invalid points, a program was created that allowed points to be labeled as **invalid** (bad points). This process involved manually inspecting 186 datasets (31 participants $\times$ 3 poses $\times$ 2 point clouds [real and avatar] and marking each by hand. Of the 1488 points throughout all 186 datasets, 89 points were considered invalid due to impossible anatomical positioning decided upon by the user or points centered at the origin (result from controller tracking loss).

3.7.3 Questionnaire Data Quality Control

To ensure the integrity of self-report responses on the Sense of Embodiment (SoE) questionnaire, a screening procedure was used to detect inattentive or disengaged participants. Specifically, the distribution of responses across the nine SoE items for each participant was examined. Responses were flagged as suspicious if (1) all responses were identical (e.g., all 1s), or (2) all but one response were the same—a pattern indicative of response spamming or inattentiveness.

One participant exhibited the second pattern, selecting the lowest possible score (1.0) for eight of the nine items, with a single divergent score (5.0). This response pattern suggests a lack of thoughtful engagement with the questionnaire, and the data from this participant was therefore excluded from further quantitative analyses involving SoE scores.

To further justify the decision of removing their responses from the analysis, the SUS and embodiment sections made use of alternating positive and negative questions, specifically for the purpose of maintaining engagement and identifying invalid responses.This quality control step helps maintain the reliability of the questionnaire-based measures and ensures that only valid, interpretable data contribute to the statistical results.

3.8 Statistical Analysis Plan

To address the hypotheses and research questions, the following statistical analyses were planned and conducted.

3.8.1 Analysis of H1 (Performance in Avatar Alignment)

To test Hypothesis 1, two separate 2x3 mixed-design ANOVAs were conducted—one for each of the primary performance metrics. The significance level for all tests were set at $\alpha = .05$.

For each ANOVA, the variables were defined as follows:

- The **independent variables** were:
  - A between-subjects factor of **Immersion** with two levels (Mixed Reality vs. Virtual Reality).
  - A within-subjects factor of **Pose** with three levels (Laying, Sitting, Standing).
- The **dependent variables**, each measured in a separate test, were:
  - Positional Accuracy, measured as Procrustes-Aligned Mean Per Joint Position Error (PA-MPJPE).

- Rotational Accuracy, measured as the Root Mean Square Error (RMSE) of segment angle differences.

The analysis of each ANOVA focused on three key effects:

1. The **main effect of Immersion** to determine if there is an overall significant difference in accuracy between the MR and VR groups.
2. The **main effect of Pose** to determine if alignment accuracy differs significantly across the three poses.
3. The **Immersion × Pose interaction effect** to determine if the effect of the immersive environment on accuracy depends on the specific pose being performed.

3.8.2 ANOVA Assumptions

Prior to analysis, the alignment data were assessed for the assumptions of a mixed-design ANOVA. To correct for skewness, a log-transformation was applied to both the positional (PA-MPJPE) and rotational (RMSE) data. The positional data were also winsorized to handle extreme outliers. The assumptions of normality and homogeneity of variances were met for all groups after preprocessing. The assumption of sphericity was met for both analyses (Greenhouse-Geisser $\epsilon > 0.75$), so no correction was necessary.

H1 ANOVA Power Analysis

A power analysis was conducted for a mixed ANOVA using G*Power software to evaluate the effect of immersion mode on alignment accuracy and the sense of embodiment. Existing literature reports mixed results regarding the estimated effect size of immersion mode on task accuracy, particularly in the context of proprioceptive accuracy during spatial alignment tasks.

In the proposed experiment, users are tasked with aligning a virtual avatar to where they perceive their real body to be. VR users are expected to be at a disadvantage due to the lack of direct visual feedback, resulting in greater reliance on proprioception. Therefore, the anticipated effect size is a medium effect, specifically Cohen's $f = 0.20$, which corresponds to a partial eta-squared (η_p^2) of approximately 0.0385.

The power analysis was conducted with the following parameters: $\alpha = 0.05$, power = 0.80, number of groups = 2 (VR and MR), and number of measurements = 3 (three poses per participant). Given the expectation that individuals who perform well in one pose will likely perform similarly in the others, the correlation among repeated measures was set to 0.7.

The analysis indicated that a minimum total sample size of 26 was required to achieve the desired power, corresponding to a critical F-value of 3.19 with 2 numerator and 48 denominator degrees of freedom ($F(2, 48) = 3.19$).

3.8.3 Analysis of H2 (Sense of Embodiment)

To test Hypothesis 2, which predicted a significant difference in the subjective sense of embodiment between the immersive environments, an **independent samples t-test** was conducted.

The test was designed as follows:

- The **independent variable** (or grouping variable) was **Immersion**, which had two levels: Mixed Reality and Virtual Reality.
- The **dependent variable** was the overall **Sense of Embodiment (SoE) score**, calculated from the standardized Embodiment Questionnaire [40].

The purpose of this test was to compare the mean SoE score of the MR group against the mean SoE score of the VR group. The significance level for the test was set at $\alpha = .05$.

H2 Power Analysis

An a priori power analysis was conducted using G*Power for an independent samples t-test to determine the required sample size to test Hypothesis 2. The analysis was based on detecting a conventional medium effect size (Cohen's d = 0.6), with an alpha of .05, and a target power of 0.80 for a two-tailed test. This analysis indicated that a total sample size of 90 participants (45 per immersion group) would be required to reliably detect a medium-sized effect.

The final analysis for this hypothesis included 31 participants. A post hoc sensitivity analysis was therefore conducted to determine the minimum effect size the study was adequately powered to detect with this sample size. With an actual sample of N = 31, an alpha of .05, and power of .80, the analysis revealed that the study was only sensitive enough to reliably detect a large effect size of d = 1.04 or greater.

This indicates that while the study was sufficiently powered to find a very large effect for H2, it was underpowered to detect a more subtle small or medium-sized effect. The results pertaining to H2 should therefore be interpreted with this limitation in mind.

3.8.4 Analysis of Exploratory Research Questions

The following outlines the analysis plan of RQ1 and RQ2. These were intended as exploratory and did not receive a power analysis.

RQ1: Analysis of Proprioception

A linear mixed-effects modeling approach was used to examine the relationship between participants' baseline proprioceptive accuracy and their performance in avatar alignment tasks, accounting for repeated measures by including participant ID as a random effect. Separate models were constructed for positional and rotational alignment error, each incorporating an interaction between proprioception and immersion group (MR vs. VR). Diagnostic checks confirmed that the assumptions of linearity, homoscedasticity, and normality of residuals and random effects were satisfied, supporting the validity of the models.

RQ2: Analysis of System Usability

To address RQ2, an independent samples t-test was conducted to compare the mean System Usability Scale (SUS) scores between the MR and VR groups. For this analysis, the independent variable was Immersion (Mixed Reality vs. Virtual Reality), and the dependent variable was the overall SUS score. The significance level for the test was set at $\alpha = .05$.

RQ3: Analysis of Task Completion Time

To explore potential learning effects and task completion times between MR and VR, task completion time was also analyzed. The time taken to complete the first, second, and third pose for each participant was compared between the MR and VR groups using independent samples t-tests. To account for unreasonable times due to system error or experimenter error, a common filtering practice removed outliers beyond the 1.5 interquartile range (IQR) above the third quartile and below the first quartile. Specifically, session times that fell outside the range defined by

$$Q_1 - 1.5 \times \text{IQR} \quad \text{and} \quad Q_3 + 1.5 \times \text{IQR}$$

were excluded from further analysis. This method effectively identifies and removes extreme values likely caused by technical issues or procedural interruptions, ensuring a more accurate representation of typical task completion times.

3.8.5 Interface Performance and Usage

Additionally, the number of interactions and the time spent interacting with both the Direct and Puppet interfaces will be captured. This will provide additional information pertaining to user preference and the overall usability of the interfaces that SUS surveys cannot capture. This data will be collected on a purely exploratory basis and will track the contents of table 3.6. Due to an implementation oversight, the time spent interacting with individual joints will be measured by the number of frames spent interacting with the respective joint instead of a value of time.

Table 3.6: Summary of Session Time Fields

Field	**Type**	**Description**
sessionTime	Float	Total duration of the session in seconds.
pose1Time	Float	Time spent in pose 1 in seconds.
pose2Time	Float	Time spent in pose 2 in seconds.
pose3Time	Float	Time spent in pose 3 in seconds.
puppetLeftFoot	Integer	Frames spent interacting with the puppet's left foot.
puppetLeftKnee	Integer	Frames spent interacting with the puppet's left knee.
puppetRightFoot	Integer	Frames spent interacting with the puppet's right foot.
puppetRightKnee	Integer	Frames spent interacting with the puppet's right knee.
puppetWaist	Integer	Frames spent interacting with the puppet's waist.
playerLeftFoot	Integer	Frames spent interacting with the player's left foot.
playerLeftKnee	Integer	Frames spent interacting with the player's left knee.
playerRightFoot	Integer	Frames spent interacting with the player's right foot.
playerRightKnee	Integer	Frames spent interacting with the player's right knee.
playerWaist	Integer	Frames spent interacting with the player's waist.

3.8.6 Programs and Libraries Used

The experimental application was developed in the **Unity** game engine using **C#**. All real-time measurements and initial data processing were handled within the application. Vector calculations utilized the built-in methods of Unity's `Vector3` struct, part of the core scripting API. During the study, data were logged locally to CSV files, with periodic backups stored in a **SQLite** database to ensure redundancy and maintain data integrity.

For the power analysis, **G*Power** software was used. All subsequent data cleaning, statistical analysis, and visualization were conducted in **Python**, utilizing the libraries detailed in Table 3.7.

Table 3.7: Python Libraries Used for Data Processing and Analysis

Library	**Purpose**
Pandas	Data manipulation and analysis via DataFrames.
NumPy	Fundamental numerical operations and array handling.
Matplotlib & Seaborn	Creation of data visualizations and plots.
SciPy & statsmodels	Performing statistical tests (ANOVAs, t-tests, correlations).

Chapter 4

System

4.1 High-Level System Overview

The following chapter highlights critical design decisions made through the development process of the program as well as some challenges faced along the way.

This experiment was designed for the Meta Quest 3 using Unity Game Engine. The system features a full VR environment and has mixed reality passthrough support as well. The system is set to require both Hand Tracking and Controller support. This was chosen in order to: 1. measure proprioception alignment accuracy via hand tracking and 2. provide a reliable, easy to use, input method for participants via controllers.

4.2 Development Environment and Core Technologies

The Unity version 2022.3.17f1 was selected for this system due to its compatibility with another critical package, `VelUtils`. The package `VelUtils` is an open source library that provides solutions for developing virtual applications within unity and was developed by the University of Georgia Virtual Environments Lab. The most notable library used was the `World Mouse` library. `World Mouse` enables quick and easy interactions with virtual canvases present in the unity scene. This enables controllers the ability to interact with UI elements such as buttons, scroll bars, etc. right out the gate.

The Meta XR Core SDK version 72.0.0 was used as well. This package provides various prefabs and easy-to-use features for VR projects. The player `gameobject` makes use of the `OVRCameraRig` prefab provided by this SDK. The rig comes stock with the `OVRCameraRig`, `OVRManager`, and a Unity `InputModule`. The `OVRManager` is the main interface for the Meta Quest environment and manages critical features and build settings required for applications to be run on Meta headsets. Using the `OVRManager`, the project was configured to support the Meta Quest Pro, Quest 3, and Quest 3S for their support of colored passthrough. This project uses a tracking origin type set to the `Floor Level`. Floor level tracking enables the use of the real-world calibrations from the floor and also allows for simpler calculations when calibrating the users position.

Additionally, the system is set to require both Hand Tracking and Controller support. This is due to the proprioception task requiring hand tracking for the position measurements. To that end, and to ensure optimal performance and accuracy, the hand tracking frequency was also set to **MAX** and passthrough support is set to **Required**.

Universal Render Pipeline (URP) was chosen to be the render pipeline for this project due to to its additional camera and material settings. The ability to add additional cameras to the render stack is very useful for showing/hiding specific layers present in the scene. This did however lead to some difficulty integrating with libraries that did not have URP-compatible shaders. OOne such compatibility issue arose during the integration of World Mouse and its laser feature. The library was written to incorporate a hard-coded, Built-In render pipeline shader (Unity's default render pipeline). This issue was overcome by adjusting the `VelUtils` package to allow for custom materials to be used for the laser.

4.3 Avatar and Interface Design

Two novel lower-body interfaces were developed to allow users to align a virtual avatar with their own body by grabbing and moving the avatar's joints with the controllers. The design ensures only anatomically correct poses can be achieved and that the two interfaces work seamlessly together.

The design makes use of two distinct control mechanisms:

- *Direct Interface:* This interface overlays **interactive, physics-based joints** directly onto the hips, knees, and ankles of the virtual avatar, which can be manipulated using controller-based input. Joints have realistic constraints including mass and use a realistic force-to-target drive system taking advantage of Unity Engine's Articulation Bodies feature set. Each leg is able to be manipulated individually enabling non-symmetrical movements and poses.

- *Puppet Interface:* Complementing the *Direct Interface*, the *Puppet Interface* provides a miniature, proxy version of the player's avatar, visible within the immersive environment. This miniature avatar features the same realistic, physics-based joint manipulations as the *Direct Interface*. Crucially, changes in the position and orientation of the puppet's legs are reflected in real-time on the main avatar, and any direct manipulations of the main avatar's legs are simultaneously visible on the puppet. This bidirectional synchronization establishes a dynamic, bounded digital twin, offering users an alternative, and potentially more convenient, method for indirect manipulation of their virtual lower-body.

4.3.1 Proposed Approach to Lower Body Control Interfaces

The lower-body control interfaces presented in this study provide a low-profile and intuitive way to manipulate the lower-body of a virtual avatar. The setup uses the Meta Quest 3 headset and the two touch controllers. Outside of the headset, no other sensors or wearable devices are needed to use the interface. Additionally, the Quest 3 enables colored passthrough support, enabling virtual and mixed reality environments.

These interfaces are designed to encourage a stronger sense of embodiment by prioritizing both **multisensory integration and synchronization** and a robust **sense of agency** within a consumer-accessible setup. The use of a realistic physics engine for the avatar's legs is critical in this regard for achieving this multisensory effect. By enabling the virtual avatar to mimic the complex physical constraints of real legs, the controller aims to provide interactions that are closely aligned with the expectations that actual legs carry. This aligns with foundational principles of embodiment, where multisensory input (seeing the avatar move realistically in response to intent) is crucial for a strong sense of body ownership and self-location [16, 1]. Furthermore, by striving for natural and predictable responses, the interfaces enhance the user's sense of agency, as the perceived sensory consequences of their actions closely match their motor commands, effectively "closing the loop" between intention and perception [3].

4.3.2 Player Avatar

The player's avatar is a full body custom humanoid avatar. The avatar was designed using Blender, an open source 3D modeling software. In order to create a quality and modular in-house avatar, two human perspective reference images were used. During development, it was very important that the mesh be divided into several separate sections before being ported over to unity. This meant that the head, legs, torso, hands, feet, and arms all needed to be separate. This approach makes it trivial to enable or disable the respective components

After the model was completed, Blender's built-in humanoid skeleton was added and adjusted to fit the needs of the avatar. A full-body skeleton rig complete with unique bone IDs enables the potential for animating and implementing full body Inverse Kinematics.

Next, the completed player avatar was exported as an FBX (Filmbox) file—a widely used 3D model format—and imported into Unity, where the avatar was properly initialized. For the purposes of this project, the avatar needed to be able to move with the headset and respond to changes in orientation. This also meant that the arms and head needed to follow the controllers/hands and head in natural, predictable ways. To that end, an inverse kinematic component was added to support full-body, VR input.

4.3.3 Lower Body Control Interfaces

The study was centered around investigating the effect of two lower body control interfaces for the purpose of aligning the virtual avatar with the users real body. The central idea was to give players the ability to grab the joints of their lower body and move them with the controllers. The resulting design, enabling **direct** manipulation of the player's avatar, was dubbed the *Direct Interface*. This idea was taken a step further with the addition of a mini avatar **puppet**. This puppet controller was a scaled down version of the player's avatar and was dubbed the *Puppet Interface*. The legs of the *Puppet Interface* were bound to the larger avatar which resulted in both pairs of legs copying each others movements. This meant that poses forced on one, would be copied by the other resulting in a bidirectional relationship.

4.4 Implementation Journey: From IK to Physics

To achieve this intended behavior, three things needed to be implemented:

- **Grabbing System:** Users need the ability to start and end a grab. When grabbing objects, the objects must follow the direction of the player's controller.

- **Rigid Body Articulation:** In order to achieve the most immersive and responsive system, joints and limbs must act predictably, and behave using Unity's rigid body physics.

- **Avatar-Puppet Copying:** Both the Player Avatar and the Puppet should copy each others movements in real-time. The system must allow for real-time feedback when using either controller.

4.4.1 Initial Approach: A Purely Kinematic System

The system uses Inverse Kinematics (IK), the process of calculating joint angles to reach a target, to drive the positioning of the avatars. Initially, we planned to use a popular asset, Root Motion's Final IK, to handle the full-body avatar. This solution worked well for upper-body tracking but presented immediate challenges for our lower-body-focused task. The asset's procedural walking animations had to be disabled, and initial attempts to drive the legs using a custom animation blend tree proved overly complex and could not run concurrently with the main IK system.

4.4.2 Challenges with the IK Approach

Further testing revealed critical flaws in using a purely kinematic approach for the manipulation task. When a user-grabbed target was moved beyond the IK solver's joint constraints, the avatar's leg would attempt to follow and then immediately snap back to a valid position. This resulted in jittery, uncanny motion, and the manipulation felt instantaneous and lacked a sense of weight, making fine-tuned adjustments difficult. It became clear that a purely kinematic solution could not provide the natural, stable interaction required for the alignment task.

4.4.3 Final Solution: Physics-Driven Articulation Bodies

To address these challenges, the Final IK component was replaced with Unity's built-in Articulation Bodies. Articulation Bodies simulate joint hierarchies with physical properties, allowing for continuous, constraint-driven motion that feels more natural and weighted. This physics-based approach solved the jittering and snapping issues, as forces are applied gradually to move a joint towards its target. If a target is moved outside the joint's physical limits, the system responds by applying an opposing force—much like a real joint resisting over-extension—rather than breaking the simulation. This design decision was the key to achieving a stable and intuitive user experience.

4.5 Final System Components

4.5.1 Grabbing System

To handle interactions, a simple grabbing system was implemented using a C# interface `IGrabbable`. Objects implementing this interface respond to `OnGrabStart` and `OnGrabEnd` events. To ensure objects do not snap to the controller's center upon being grabbed, the system calculates and maintains a positional offset for the duration of the grab. This system provides the core mechanism for all user manipulation within the study.

4.5.2 Articulation Body Configuration

The avatar's lower-body hierarchy is driven by Articulation Bodies, with the transform object representing the *spine* as the root. Both thighs use a **Spherical** joint type to allow for realistic hip motion, while the shins use a **Revolute** joint type to correctly simulate a hinge-like knee joint. The specific angular limits for each joint, detailed in Table 4.1, were chosen based on standard anthropometric data and user comfort testing during development. Critically, all joints operate using `DriveMode.Force`, with custom stiffness and damping values. This ensures that all manipulation feels organic and physically grounded, introducing a natural sense of resistance and weight.

Table 4.1: Articulation types and constraints for each bone.

Bone	Articulation Type	Constraint (X, Y, Z)	Notes
Spine	None	None	Parent object
Thigh.L	Spherical	X: [-90, 90], Z: [-45, 45]	Left thigh
Thigh.R	Spherical	X: [-90, 90], Z: [-45, 45]	Right thigh
Shin.L	Revolute	X: [-10, 135]	Left shin (knee)
Shin.R	Revolute	X: [-10, 135]	Right shin (knee)

4.5.3 Interface Implementation

The Grabbing System and Articulation Body configuration come together to form the two user-facing interfaces. Large, interactive spheres, called *Grab Targets*, are placed on the joints of both the player avatar and the puppet. These targets are the objects that implement the `IGrabbable` interface, providing a clear visual affordance for interaction. When a user grabs and moves a target, its new position is fed to the corresponding Articulation Body drive, which then applies the necessary forces to move the avatar's limbs.

The Direct Interface

The *Direct Interface* is composed of the *Grab Targets* placed directly on the main player avatar's legs, as shown in Figure 4.1. This allows for immediate, one-to-one manipulation of the avatar. The articulation bodies used in the implementation create a realistic, weighted feeling while moving the joints. Similar to IK designs, the avatar's leg will follow the users gripping hand's motion and adjust according to the *Grab Target's* position and the anatomical limitations constrained to it by the system. The result of this design creates a sensation akin to using one's arms to move their legs.

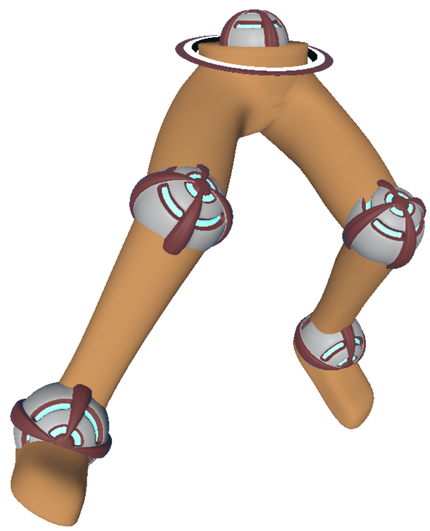

Figure 4.1: Lower body *Direct Interface* representing the player's virtual avatar. Participants were tasked with aligning the legs of this avatar with their real legs.

The Puppet Interface

The *Puppet Interface* is a miniature, proxy version of the avatar that can be freely moved around by the user (Figure 4.2). It features the same *Grab Targets* and interaction schema as the *Direct Interface*, allowing it to be manipulated to match realistic poses with smooth, weighted inputs.

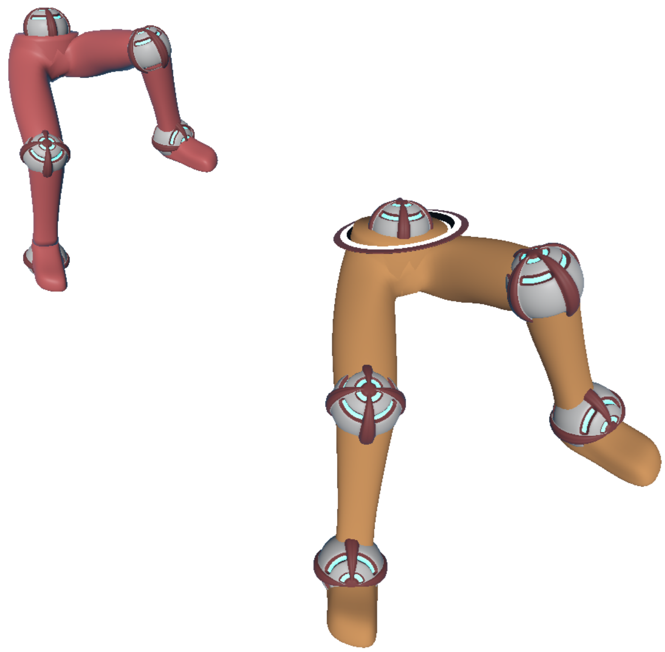

Figure 4.2: *Puppet Interface*(left) next to the *Direct Interface*(right). Note: due to the bidirectional relationship of the two, they will always share the same pose.

4.5.4 User Interaction and Workflow

Both the *Direct Interface* and the *Puppet Interface* are coupled and were made to mimic each others motions in real time. This design creates a unique opportunity for applications requiring a wide range of lower-body alignment poses, such as XR amputee therapies, that might be restricted by posture limitations. For bedridden patients either laying down or at an incline, the level of control offered by these interfaces offers a convenient solution for alignment.

First, the user can utilize the *Puppet Interface* to perform large, coarse adjustments needed to move the virtual avatar's limbs into a more reachable position. This method avoids wide, sweeping physical motions that would otherwise require the user to adjust their posture from a reclined or lying position. Afterwards, the user can fine-tune the alignment of their virtual avatar over their real body using the *Direct Interface*. Any further adjustments, or changes in posture can be done on the fly, or automatically through the use of preset poses by the application.

4.6 Data and Session Management

The system includes a custom web dashboard used by the proctor to register participants, assign them to an experimental condition (VR or MR), and trigger data capture events. At the end of each alignment task, the final state of the avatar's joint transforms, along with the ground-truth landmark data measured by the controllers, were serialized into JSON files and saved for offline analysis. Figure 4.3 shows an example of the three poses used in this study, extracted directly from the JSON packets.

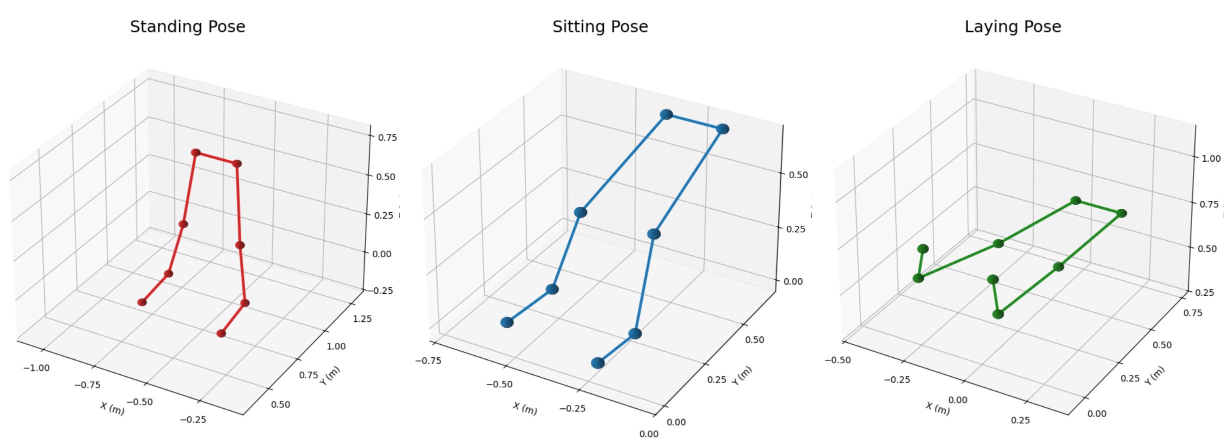

Figure 4.3: 3D representations of the pose point cloud data captured throughout this user study.

4.7 System Conclusion

In summary, the final design utilized a physics-based approach with Unity's Articulation Bodies to create two novel, robust user interfaces for lower-body avatar control. This approach successfully addressed the limitations of a purely inverse kinematics (IK) solution, resulting in a more realistic, flexible, and intuitive control scheme. The *Direct Interface* enables precise, fine-tuned manipulation of the avatar's legs, allowing users to match target poses with a natural arm-like interaction. The *Puppet Interface* facilitates broader adjustments and offers an efficient way to reposition the lower body. Changes made in one interface are reflected in the other, creating a seamless and adaptive workflow regardless of the user's physical pose.

Chapter 5

Results

5.1 Descriptive Statistics

This chapter presents an overview of the data collected in the study. A summary of the descriptive statistics for the primary dependent variables is presented first in Table 5.1, followed by more detailed breakdowns of the performance metrics.

Table 5.1: Descriptive Statistics for Key Dependent Variables by Immersion Group

Measure	Group	n	Mean	Std. Deviation
Positional Accuracy	MR	16	7.88 cm	2.80 cm
(PA-MPJPE, cm)	VR	15	7.03 cm	2.45 cm
Rotational Accuracy	MR	16	13.51 deg	5.24 deg
(RMSE, degrees)	VR	15	14.45 deg	7.96 deg
Sense of Embodiment	MR	19	3.82	1.04
(Overall Score / 7)	VR	18	4.56	0.82
System Usability	MR	19	61.40	16.56
(SUS Score / 100)	VR	18	62.78	10.97
Proprioceptive Error	NA	29	1.39 cm	1.01 cm
Task Completion Time	MR	19	5.53 min	1.34 min
(minutes)	VR	18	6.64 min	2.54 min

Note: n values reflect final participant counts after data cleaning for respective category. Accuracy measures are from valid pose datasets (H1); Embodiment and Usability are from valid survey responses (H2, RQ2); Proprioceptive Error values are from trial #1 valid results (RQ1); Task Time is based on recorded trials per group (RQ3).

5.2 Missing Data and Outliers

5.2.1 Initial Data Screening Results

After all participants had completed the sessions, it was found that 8 were missing session data, which includes the proprioception and alignment task, but did have valid survey responses. Additionally, one participant was found to be missing the survey data but did have session data. After further investigation into this, it was found that the individual without survey data also had an incomplete session dataset. These 9 were excluded from from the final analysis. This is outlined in table 5.2. This missing data was due to an oversight in the data collection procedure which resulted as well as sever technical difficulties, preventing one users session and survey data from being recorded properly.

Table 5.2: Post-study Data Overview

Category	**Entries**	**Description**
Total Participants	40	All entries
Survey and Session	31	Full, verified session data
Survey Only	8	Survey data only, no verified session data
Session Only	1	Session data only, no verified survey data

5.2.2 H1 Final Counts

After all alignment positional data was unwrapped and separated into both real and avatar joint positions, several inaccuracies and missing positions were noted. All avatar joint positions were present however, several real joint positions were missing.

To assist in the screening of the anatomical landmark data, a custom data visualization and outlier tagging tool was developed, as shown in Figure 5.1. This tool presented the joint data for each participant and pose in both a table and an interactive 3D scatter plot. This allowed for the visual inspection of each pose. Using the tool's interface, points deemed invalid could be flagged for exclusion from the analysis. Invalid points included those that were deemed anatomically impossible and tracking error where one of the joint point clusters was located far away from the others. Interestingly, during most tracking errors, joints would have their y values set to 0. In the figure 5.1, joint index 38 is shown to have this behavior. Consequently, it was marked and highlighted in red.

After all poses for all participants were examined, the list of outliers was exported and the raw dataset was filtered using the joint's index. Through this manual inspection process, 89 out of 1488 measured points were identified as invalid and excluded.

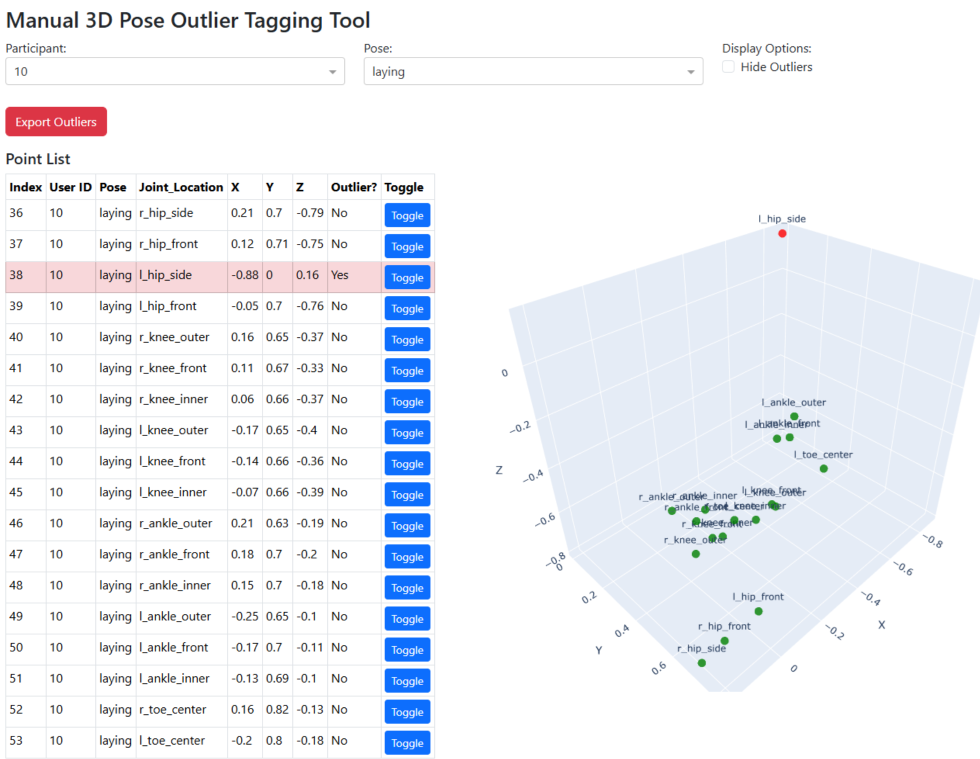

Figure 5.1: The custom-built manual outlier tagging tool used to screen the 3D pose data. The interface displays joint coordinates in a table alongside an interactive 3D visualization, allowing for the inspection and flagging of invalid data points. Ex: index 34 (in red), belonging to the left hip side measurement, is unusable and is flagged as an outlier.

Table 5.3 summarizes the data quality and filtering process for body alignment accuracy across different poses and immersion environments. This process began with an initial dataset of 93 group-condition pose tasks (31 participants × 3 poses). Within this dataset, any poses with an incomplete set—defined as missing one or more of the six lower-body joints, or having anatomically impossible joint positions, including those with extreme tracking failures—were identified as invalid and excluded. This ensured that only fully defined and plausible point configurations entered the Generalized Procrustes Analysis (GPA). Each pose was expected to contain data for the following joints: `ankle_L`, `ankle_R`, `hip_L`, `hip_R`, `knee_L`, `knee_R`.

Table 5.3: Final Valid Positional Data by Pose and Immersion

Pose	**Immersion**	**Valid**	**Excluded**	**Total**
Laying	MR	15	1	16
	VR	15	0	15
Sitting	MR	15	1	16
	VR	13	2	15
Standing	MR	12	4	16
	VR	8	7	15
Grand Totals		**78**	**15**	**93**

*Note: The study design originally called for 93 group-condition pose tasks (31 participants × 3 poses). The **Excluded** column accounts for all instances of invalid data, encompassing both initial extreme system errors (e.g., complete tracking failure) and subsequent checks for anatomically impossible joint positions. The **Valid** column represents the data points that proceeded to Generalized Procrustes Analysis. The **Total** column sums valid and excluded sets for each sub-condition, resulting in a grand total of 93 recorded and accounted-for tasks.

Figure 5.2 visualizes one participant's joint transforms before and after GPA alignment for the sitting pose. Real joint centers are shown in blue, avatar joints in orange, with a dashed red line indicating the error offset. The participant is facing the positive y-axis in Unity world space.

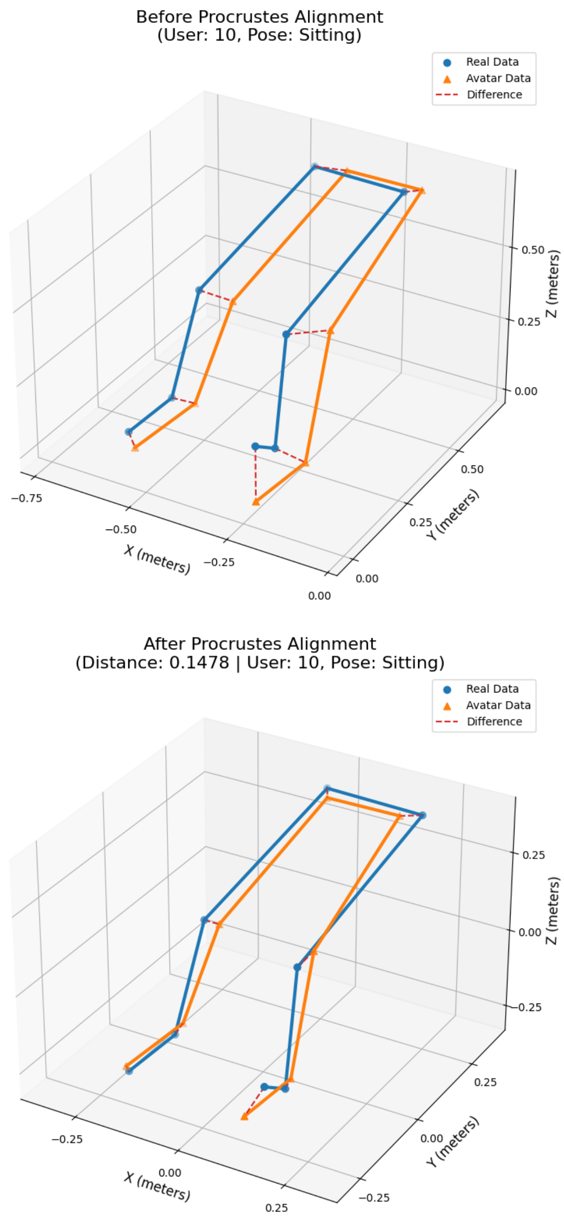

Figure 5.2: 3D visualization of joint transform data from a sitting pose task for one participant. Blue represents the participant's real body joint averages, while orange represents the virtual avatar's. The first panel shows the initial configuration, and the second panel shows the configuration after Generalized Procrustes Analysis (GPA) alignment, with dashed red lines indicating the error offset. All positions are in Unity world space (meters).

5.2.3 Survey Final Counts

The analysis for H2 relied on valid survey data from each participant. As mentioned previously, 1 user was identified to have missing survey data. This was a data organizational failure to keep track of each participants ID throughout the entire study.

In addition to having completed the survey, each participant's responses were screened. To ensure the integrity of self-report responses, a screening procedure was used to detect inattentive or disengaged participants. Survey responses for both the system usability and the sense of embodiment questionnaires used a positive/negative question pattern. Responses were flagged as suspicious if all responses or all but one response were the same—a pattern indicative of response spamming or inattentiveness. Those flagged as such were examined individually for signs of disengagement and were ultimately decided on whether or not to include in the final analysis.

The screening procedure found two users suspected of this violation. After examining all of their responses, it was decided to exclude their survey data entirely from all survey-related analysis - this includes the sense of embodiment score as well as the system usability. 37 responses were used in the final analysis, as shown in table 5.4.

Table 5.4: Participant Flow for Survey-Based Analyses

Category	**Count**
Total Participants Recruited	40
Missing Survey Data	1
Excluded for Inattentive Responses	2
Total Included in Final Analysis	**37**
Breakdown by Immersion Group:	
Mixed Reality (MR)	19
Virtual Reality (VR)	18

5.2.4 Proprioception Screening

Similar to the joint positional screening processed used for H1, RQ1 underwent a process to remove any outliers caused by tracking issues. Recall, there were 8 participants missing session data (including alignment task and proprioception task). There was also one participant who was found to have incomplete session data and was missing the survey entirely. This left 31 valid sets of proprioception data to be screened before analysis.

Because in-headset hand tracking was used to measure the location of the fingers, the likeliness of anatomical issues was low. A quick visualization of the data confirmed this. There were, however, two users found with severe tracking issues. These two datasets were unable to be salvaged and subsequently removed from the final analysis. The final proprioception analysis used 29 valid datasets.

5.2.5 Session Times Screening

Recorded session times were assessed to remove outliers due to issues during execution of the procedure as well as unreasonable times due to system errors. During the procedure, one participant was noted to have a headset issue which required the experimental proctor's assistance to fix, adding an inaccurate representation of the task completion time. Data was also analyzed for any times deemed 'unreasonable' such as being incredibly long or short with no explanation other than system error.

After screening the initial $31 - 1$ (headset issue) $= 30$ participants with fully valid session data, 2 more were removed using the IQR filtering method. Thus, task completion times from 28 valid participants were used in the final analysis.

5.2.6 Valid User Summary

Table 5.5 summarizes the final number of valid participants or data points used in the study for each primary analysis, reflecting the specific data cleaning and exclusion criteria applied at each stage.

Table 5.5: Summary of Final Data Counts Used for Each Analysis

Analysis	**Analysis Method**	**n**
Alignment Performance (H1)	Mixed-Design ANOVA	29
Sense of Embodiment (H2)	Independent Samples t-test	37
Proprioception (RQ1)	Linear Mixed-Effects Model	29 (72)
System Usability (RQ2)	Independent Samples t-test	37
Task Completion Time (RQ3)	Independent Samples t-test	28

Note: The final n for each analysis differs due to specific exclusion criteria, such as anatomical implausibility in pose data, inattentive survey responses, hand-tracking failures, and removal of outlier session times. RQ1 has 72 observations used during the analysis of the Linear Mixed-Effects Model.

5.3 Hypothesis 1: Effect of Immersion on Alignment Performance

To test the effect of the immersion environment on alignment performance (H1), we analyzed rotational and positional accuracy using mixed-design ANOVAs, with follow-up t-tests for each pose.

5.3.1 Rotational Accuracy

The analysis for rotational accuracy (RMSE), detailed in Table 5.6, revealed no significant main effect of Immersion, $F(1, 17) = 0.07, p = .80$, no significant main effect of Pose, $F(2, 34) = 0.62, p = .54$, and no significant Immersion × Pose interaction, $F(2, 34) = 0.28, p = .75$. Follow-up t-tests comparing MR and VR within each pose also showed no significant differences. These results are visualized in Figure 5.3.

Table 5.6: Mean and Standard Deviation of RMSE Angle Error (in Degrees) by Pose and Immersion. Note: n represents the number of valid pose datasets.

Immersion	**n**	**Pose**	**Mean (deg)**	**SD (deg)**
MR	15	Laying	13.09	5.54
MR	15	Sitting	14.55	5.88
MR	12	Standing	12.73	4.09
VR	15	Laying	13.82	7.31
VR	13	Sitting	15.54	10.29
VR	8	Standing	13.85	4.99
MR Overall	42		13.51	5.24
VR Overall	36		14.45	7.96

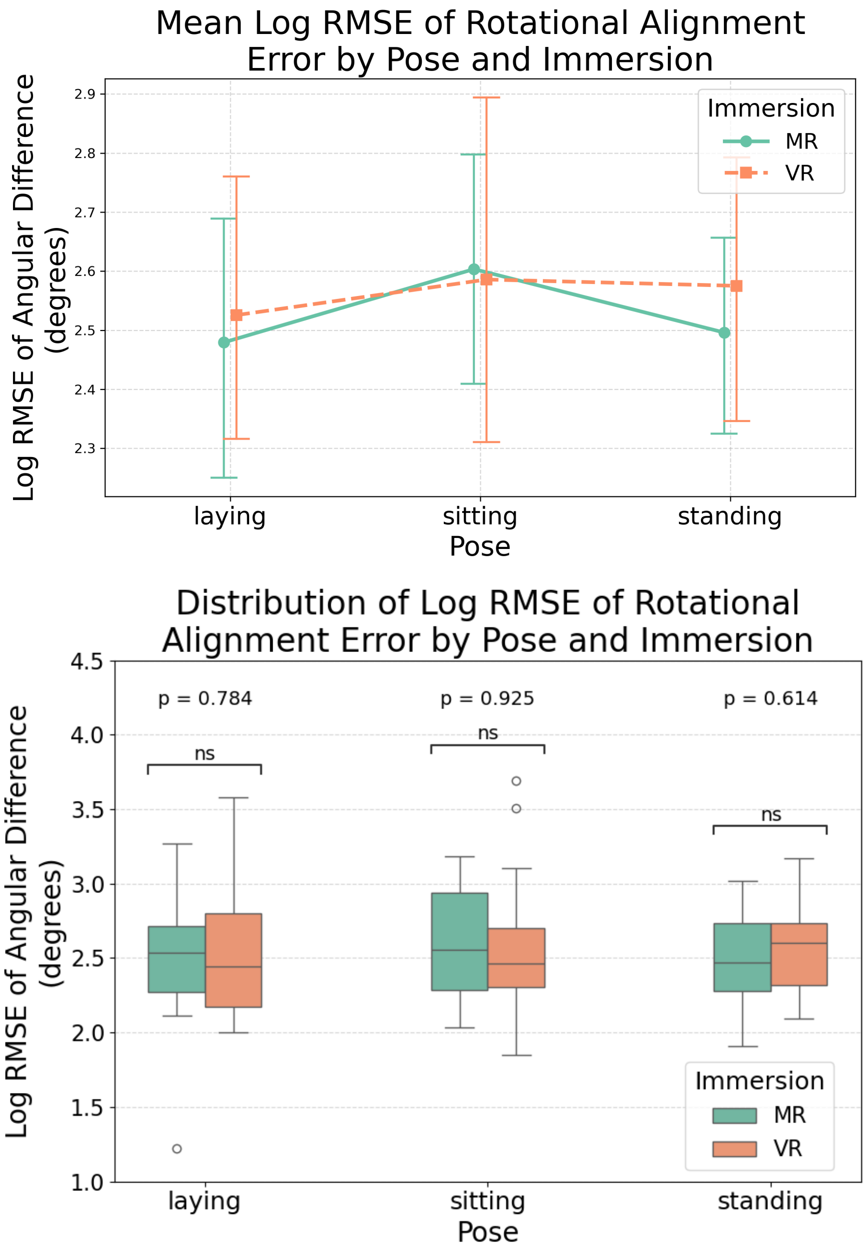

Figure 5.3: Rotational alignment error results. (Top) An interaction plot showing the mean log-transformed RMSE across poses and immersion conditions. (Bottom) Box plots comparing the rotational error between MR and VR for each individual pose. No significant differences were found.

5.3.2 Positional Accuracy

The ANOVA for positional accuracy (PA-MPJPE), detailed in Table 5.7, revealed no significant main effect of Immersion, $F(1,17) = 0.41, p = .529$, no significant main effect of Pose, $F(2,34) = 0.96, p = .393$, and no significant Immersion × Pose interaction, $F(2,34) = 1.92, p = .162$. While follow-up t-tests did not reveal any statistically significant differences between immersion conditions for individual poses, a marginal trend toward significance was observed for the **laying** pose ($p = .064$), as visualized in Figure 5.4.

Table 5.7: Mean and Standard Deviation of Positional Accuracy (PA-MPJPE) in cm. Note: Inferential tests were conducted on log-transformed, winsorized data for statistical validity.

Immersion	**n**	**Pose**	**Mean (cm)**	**SD (cm)**
MR	15	Laying	8.64	3.27
MR	15	Sitting	7.70	2.74
MR	12	Standing	7.16	2.17
VR	15	Laying	6.58	2.20
VR	13	Sitting	8.07	2.91
VR	8	Standing	6.17	1.59
MR Overall	42		7.88	2.80
VR Overall	36		7.03	2.45

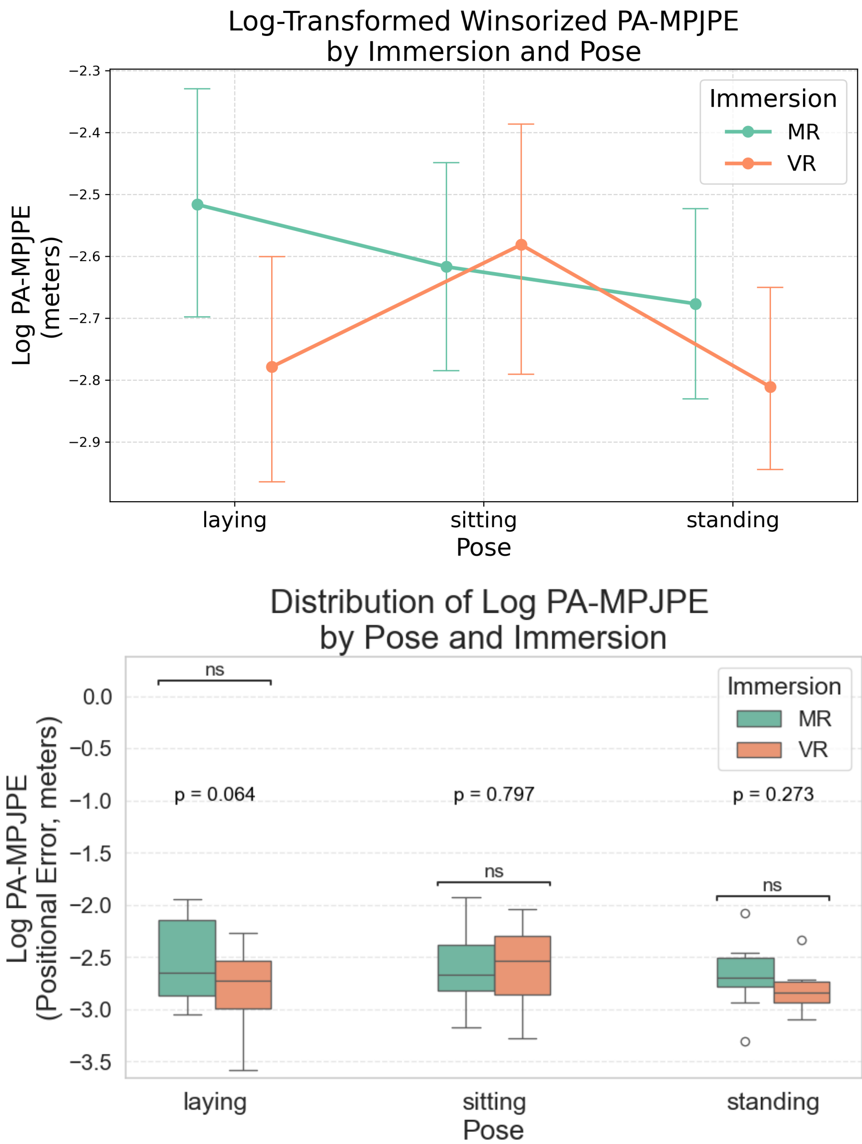

Figure 5.4: Positional alignment error results. (Top) An interaction plot showing the mean log-transformed, winsorized PA-MPJPE across conditions. (Bottom) Box plots comparing positional error between MR and VR for each individual pose. A marginally significant difference ($p = .064$) was found for the laying pose.

5.3.3 Summary for H1

Taken together, the analyses of both rotational and positional accuracy do not support H1. The data indicate that participants in the MR environment did not perform the alignment task with significantly greater accuracy than those in the VR environment.

5.4 Hypothesis 2: Effect of Immersion on Sense of Embodiment

To test the effect of immersion on the subjective sense of embodiment (H2), an independent samples t-test was conducted on the overall scores from the Embodiment Questionnaire. After removing data from two participants based on pre-defined criteria for inattentive responses, the analysis was completed on the remaining 37 participants.

The analysis revealed a statistically significant difference in the overall SoE score between the two groups, $t(34.0) = 2.42, p = .021$. In support of H2, participants in the VR condition ($M = 4.56, SD = 0.82$) reported a significantly higher sense of embodiment than those in the MR condition ($M = 3.82, SD = 1.04$).

Additionally, further analysis of the questionnaire's subscales revealed an interesting find. The higher overall SoE score in VR appears to be driven by a significantly stronger feeling of **transformation**, with the VR group scoring higher than the MR group ($t(34.7) = 2.50, p = .017$). Conversely, participants in the VR group also reported significantly higher feelings of **disembodiment** (a lower score on the reverse-scored subscale) compared to the MR group ($t(32.5) = -2.18, p = .037$). No significant difference was found for the body ownership subscale ($p = .131$). These results are visualized in Figure 5.5.

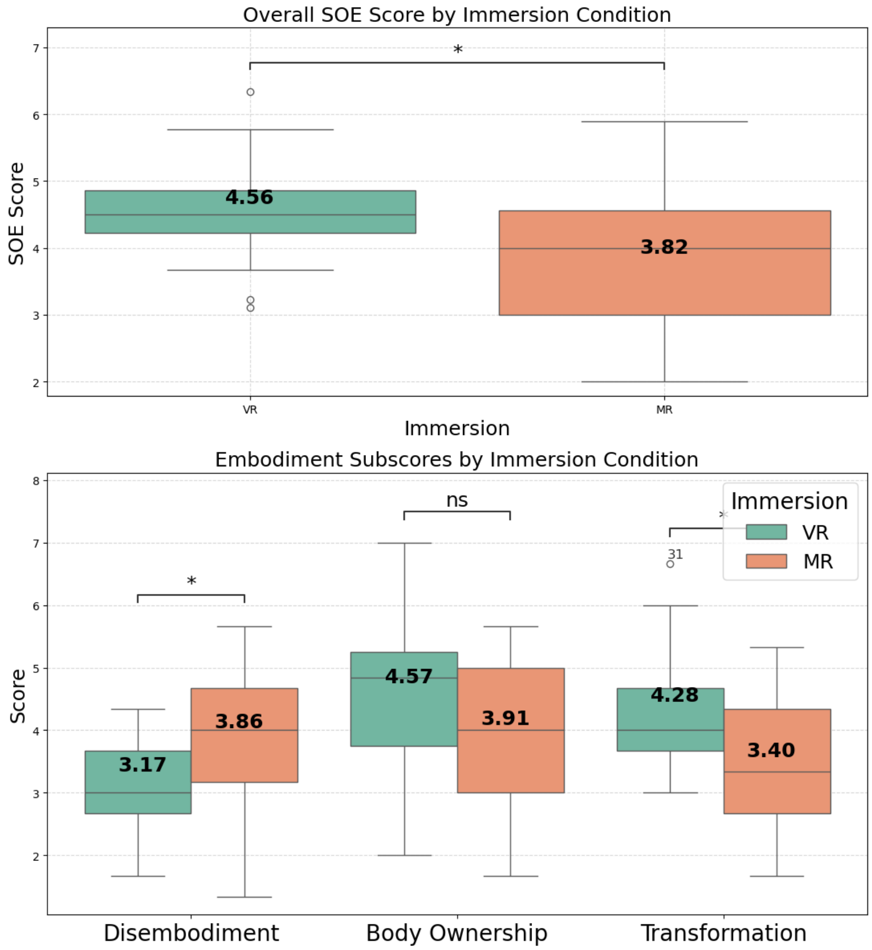

Figure 5.5: Sense of Embodiment (SoE) scores for the VR and MR groups. The significant difference in overall SoE (left) appears to be driven primarily by the transformation subscale (right). An asterisk (*) indicates p < .05.

5.5 Exploratory Findings

5.5.1 RQ1: Proprioception and Alignment Performance

To explore the relationship between baseline proprioceptive ability and alignment task performance (RQ1), we analyzed the data using linear mixed-effects models.

The model for positional accuracy (PA-MPJPE) revealed a significant interaction between proprioception and immersion condition ($\beta = -0.19$, $SE = 0.09$, $p = 0.047$). A follow-up analysis of this interaction showed that for participants in the MR group, there was a significant positive relationship between proprioceptive error and positional alignment error ($\beta = 0.12$, $SE = 0.06$, $p = 0.039$), indicating that poorer proprioception was associated with poorer performance. This relationship was not present for participants in the VR group.

The model for rotational accuracy (RMSE) showed no significant main effects or interaction effects involving proprioception.

These results suggest that baseline proprioceptive ability is a significant predictor of positional alignment performance, but only within a Mixed Reality environment where users must integrate their proprioceptive sense with external visual cues from the real world.

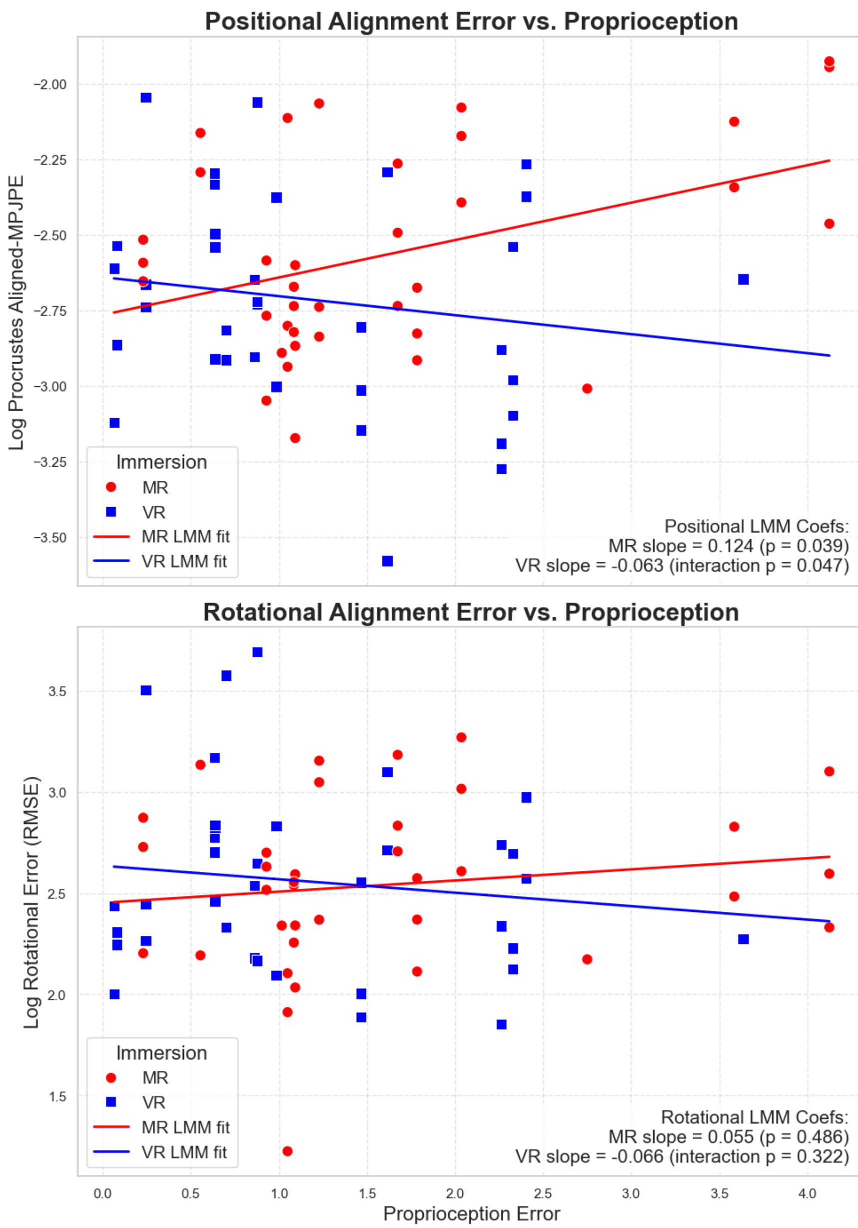

Figure 5.6: Positional and Rotational Alignment Error vs Proprioception Error. Lower values indicate greater accuracy (lower error) for both proprioception and alignment.

5.5.2 RQ2: System Usability

To explore differences in perceived usability (RQ2), an independent samples t-test on the average System Usability Scale (SUS) scores was conducted.

The results indicated that there was no significant difference in the average usability scores between the two groups ($t(29) = 0.299$, $p = .767$). However, it was observed that the variance in scores was larger in the MR group ($SD = 16.56$) compared to the VR group ($SD = 10.97$). While this difference in variance was not statistically significant ($F(18, 17) = 2.28$, $p = .096$), the trend suggests that the user experience in MR may have been more varied across participants. (see Figure 5.7).

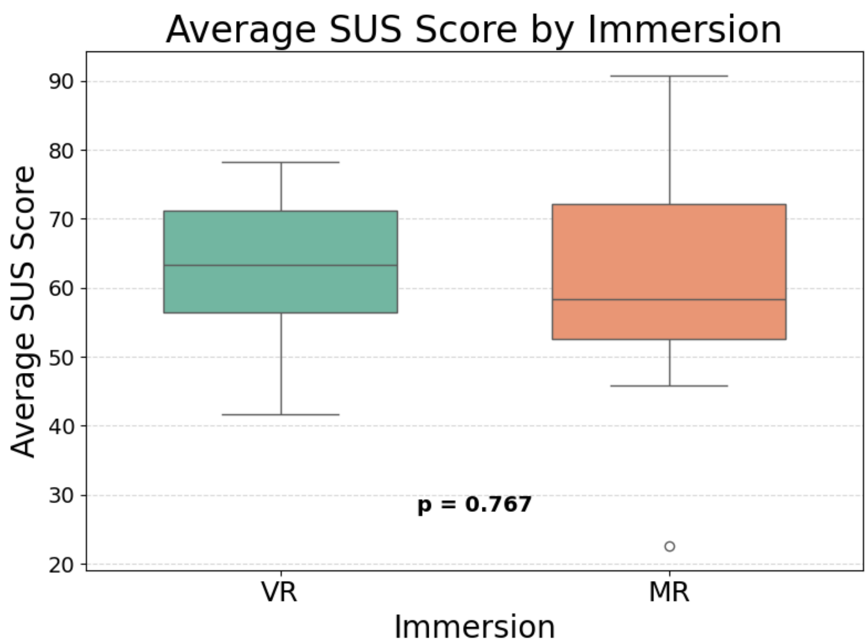

Figure 5.7: Average SUS scores by immersion condition. While mean scores were not significantly different, note the greater variability in the MR group.

5.5.3 RQ3: Task Completion Time and Learning Effects

An analysis of task completion time revealed a significant learning effect dependent on the immersion condition. A series of independent samples t-tests compared the time taken for the first, second, and third poses completed by each participant. Table 5.8 provides a summary of the descriptive statistics.

As shown in Figure 5.8, there was a statistically significant difference for the first pose. Participants in the VR group ($M = 2.80$ min, $SD = 1.23$ min) took significantly longer to complete their initial task than participants in the MR group (M = 1.93 min, SD = 1.07 min), $t(21.52) = -2.94$, p = .008. However, no significant differences in completion time were found for the second pose (p = .591) or the third pose (p = .305), suggesting that this initial difference disappeared after the first task.

Table 5.8: Descriptive Statistics for Pose Times by Immersion Group (in minutes)

Measure	Group	Mean	Std. Dev.	Min	Max
Pose Durations					
Pose 1 (First Pose)	MR	1.93	1.07	0.54	4.22
	VR	2.80	1.23	0.42	4.87
Pose 2 (Second Pose)	MR	1.73	0.39	1.14	2.66
	VR	1.95	1.13	0.79	4.24
Pose 3 (Third Pose)	MR	1.87	0.98	0.88	5.16
	VR	1.88	0.76	0.81	3.56
Per-Pose Type Durations					
Standing Time	MR	1.76	0.54	0.70	2.66
	VR	1.78	0.70	0.81	2.97
Sitting Time	MR	2.03	1.26	0.56	5.16
	VR	2.36	1.46	0.42	4.87
Laying Time	MR	1.74	0.59	0.54	2.95
	VR	2.50	1.01	1.16	4.70
Total Time					
Overall Pose Task Times	MR	5.53	1.34	3.45	8.46
	VR	6.64	2.54	2.58	11.30

Note: Pose durations are reported in minutes. "Pose 1/2/3" reflect the time spent per trial in that order. Standing, Sitting, and Laying reflect durations associated with each posture, and Total Pose Time sums the three poses for each participant.

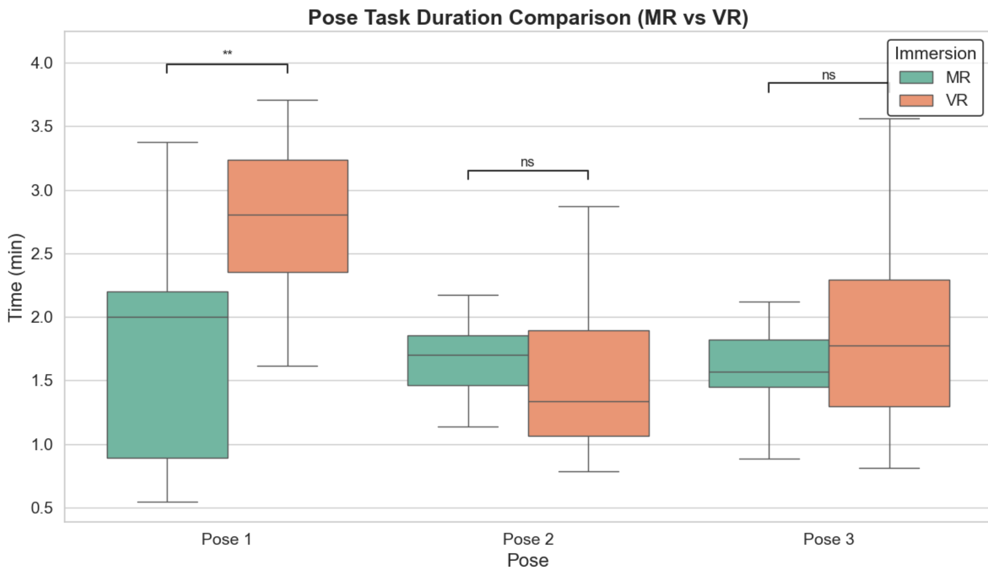

Figure 5.8: Comparison of task completion time (in minutes) for the first, second, and third poses completed by participants in each immersion group. A significant difference was found only for the first pose. (** indicates $p < .01$)

Furthermore, when comparing the total task time summed across all three poses, there was no statistically significant difference between the MR group (M = 198.22s, SD = 27.64s) and the VR group (M = 230.16s, SD = 102.08s), $t(13.63) = -1.09$, $p = .294$.

5.5.4 System Usage and Interaction

Table 5.9 highlights the average number of frames spent interacting with the joints of both the Direct and the Puppet interfaces. Overall, the Direct interface had far greater interactions on average compared to the Puppet interface. When excluding the waist from the totals, the differences is still apparent (Puppet average = 1909.72, Direct average = 3034.71). This was also seen when comparing the overall interaction time between immersion groups as seen in figure 5.10.

Table 5.9: Average Interaction Time (in frames) per Joint for Direct and Puppet Interfaces

Joint	**Puppet (frames)**	**Direct (frames)**
Left Foot	474.14	414.14
Left Knee	469.43	1169.71
Right Foot	508.61	509.07
Right Knee	457.54	941.79
Waist	503.36	2673.61
Total	**2413.08**	**5708.32**

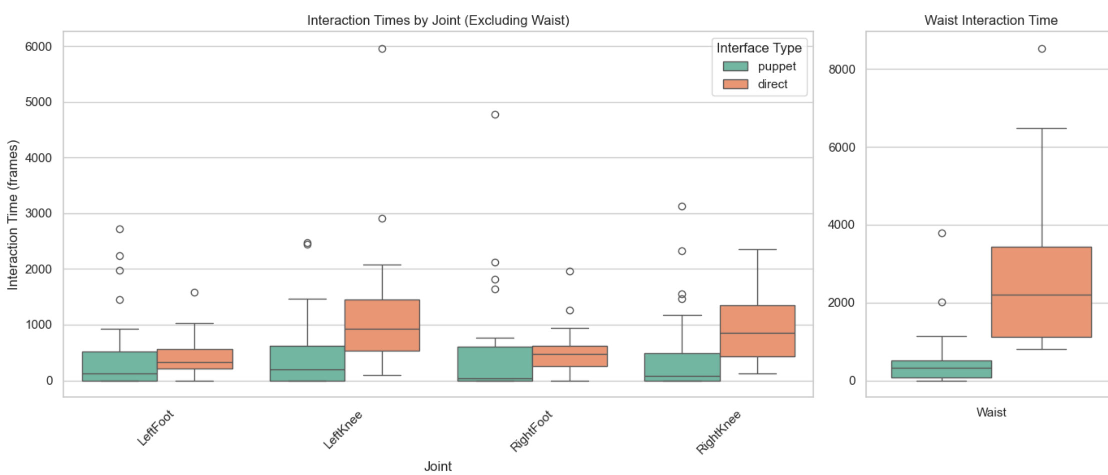

Figure 5.9: Breakdown of joint interaction times between Puppet and Direct interfaces.

Note: Waist was separated for clarity.

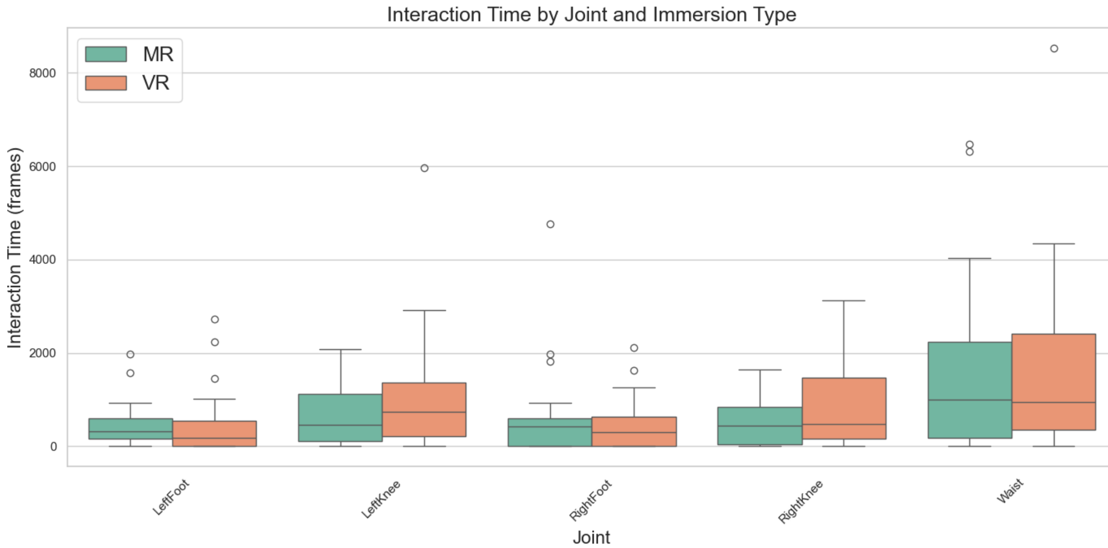

Figure 5.10: MR and VR joint interaction times between total Puppet and Direct interface interactions.

5.6 Results Summary

In summary, the statistical analyses did not support H1, revealing no significant differences in either rotational or positional alignment accuracy based on immersion mode. H2 was supported, with participants in the VR condition reporting a significantly higher sense of embodiment. The exploratory analyses revealed several key findings. For RQ1, a significant interaction effect showed that the relationship between proprioceptive error and positional alignment error was present only for the MR group. For RQ2, no significant difference in system usability was found between immersion groups, though the MR group showed substantially greater variance in user ratings. Finally, the analysis for RQ3 revealed a significant learning effect: participants in the VR group took considerably longer to complete their initial pose task compared to those in the MR group, a difference that disappeared in subsequent tasks.

Chapter 6

Discussion

6.1 Summary of Key Findings

The primary objective of this study was to explore how the body alignment method and immersive environment (MR vs. VR) impact user preference and performance in virtual avatar control. After analysis, it was found that the immersive environment had no effect on measured performance in body alignment tasks (H1) but did have a significant effect on the sense of embodiment (H2), with the VR group reporting higher scores. Exploratory analyses for RQ1 indicated a significant positive correlation between proprioceptive error and positional alignment error, but only for the MR group. For RQ2, no significant difference in system usability was found between the two immersion groups, though there was a non-significant trend toward higher variance in the MR group's usability ratings. Finally, the analysis for RQ3 revealed a significant learning effect: participants in the VR group took considerably longer to complete their initial pose task compared to those in the MR group, a difference that disappeared in subsequent tasks. This is summarized in Table 6.1.

Table 6.1: Summary of Key Findings Across Variables

Factor	**Outcome**	**Finding**
Immersion (MR vs. VR)	Body Alignment Accuracy (H1)	No significant effect
Immersion (MR vs. VR)	Sense of Embodiment (H2)	Significant effect (VR > MR)
Baseline Proprioception	Alignment Accuracy in VR (RQ1)	No significant correlation
Baseline Proprioception positive	Alignment Accuracy in MR (RQ1)	Significant correlation
System Usability	Across MR and VR (RQ2)	No significant effect on mean scores
Task Completion Time	By Immersion & Pose Order (RQ3)	Significant effect on first pose only (VR > MR)

6.2 Interpretation of Hypothesis 1 (Alignment Performance)

6.2.1 H1: The Surprising Parity in Alignment Performance

H1 claimed that users within the MR environment would have greater accuracy during the alignment task than users in the VR environment. To measure this, the rotational and positional errors were isolated and compared between each task in VR and MR.

6.2.2 Individual Pose Performance

Of the three poses, the standing task had the most unusable data. This was observed during data cleanup as a large number of joint coordinates were flagged as invalid for either being too far away from the rest or centered at the origin. This might be explained in two ways: (1) the controllers lost tracking as users bent down to record their anatomical landmarks, and (2) the downward head tilt may have disrupted headset tracking due to reduced visibility of environmental features needed for positional reference.

The alignment results for the standing pose showed the least amount of variance of the three poses in both VR and MR with the VR group specifically having the lowest *positional* error of all sub groups in all poses. The *rotational* error shows no distinctions and is very comparable to the other pose error between MR and VR. This is likely the result of the default state of the avatar during the test. At the start of the application, the virtual avatar is in an A-pose with its upper half hidden. Its legs are straight down and nearly replicate the standing pose. Between poses, the avatar does not reset back to this pose, requiring users to manipulate the legs from where they left off from the last pose.

For users starting in the first task, however, the avatar was nearly aligned and required very little adjustments aside from aligning the waist to the user's torso. This explains the extremely small variance for the positional error as many users likely felt no need to adjust the avatar at all. Unlike the rotational analysis, Procrustes alignment was used to determine positional error and account for systematic error and the *overall* size differences between participants. Differences in heights, however, still influenced rotational results and likely led to some rotation errors that were either unavoidable (due to scaling) or the lack of *need* users felt when aligning with a standing pose for task 1.

The sitting pose yielded the highest rotational errors with moderately high positional errors. This was likely because it was the most geometrically complex pose to align with. Users in both environments needed to align their virtual ankles with their real ones while preserving the correct knee bend which likely led to many users struggling or giving up early. This complex relationship between hip to knee and knee to ankle angle angular relationship needed careful attention from users. This most likely led to the much higher mean rotational error compared to the other poses in both VR and MR.

The laying pose had the most interesting turnout. Despite the lack of significant difference between VR and MR and accuracy, positional accuracy seemed to vary the most between the two. In all other poses, performance between the two conditions was remarkably similar. This was contrary to the initial hypothesis, which predicted a clear advantage for MR due to the availability of direct visual feedback. A potential explanation for this surprising turnout of VR performance lies in a ”grounding” hypothesis.

The Grounding Hypothesis

In MR environments, users are able to see not only virtual objects but also their surroundings. The benefit of this may be tied with the feeling of being physically *grounded* in their actual surroundings. For the sitting and standing poses, where participants' feet were on the floor, the passthrough view provided a stable frame of reference, aligning visual and proprioceptive cues for the user during the alignment task.

In the laying pose, participants were asked to prop their feet up on a stool and lean back in the chair. Doing so resulted in this stable anchor being removed. It could be speculated that seeing the real-world environment while feeling physically detached from it may have created a sensory conflict or increased the cognitive load which, in turn resulted in decreased positional accuracy than in VR ($p=0.064$).

For users experiencing the VR environment, by replacing the real world entirely, the system conveyed a more consistent sensory experience, through the control of sensory input (graphical consistency and simplicity of events). With no real-world visual cues to conflict with, participants may have had a lower cognitive load, allowing them to focus purely on the alignment task. This could explain why performance in VR remained so consistent across all poses, regardless of whether the user was ”grounded” or not. This suggests that for body alignment tasks, the consistency of the virtual environment may be as important as the realism of the visual feedback. These findings also support the mixed nature of MR and VR in tasks involving accuracy, consistent with findings from [41, 28].

6.2.3 H1 Concluding Remarks

When it comes to the performance of users between MR and VR, there was no statistical difference detected. After considering performance on a per-pose level, there was some nuanced behavior noted, especially when considering the unique experiences offered by the immersive environment and their cognitive demand. The results of this hypothesis hold true to the existing literature regarding the mixed findings when comparing user performance in MR and VR [7, 28].

6.3 H2: The Trade-off Between Transformation and Disembodiment

H2 claimed that users within the virtual environment would have a greater sense of embodiment (SoE) than users within the MR environment. The results from an independent samples t-test support this claim, revealing a significantly higher overall SoE score for the VR group (t(34.0) = 2.42,p = .021). When broken down, a look at the questionnaire's subscales suggests a nuanced picture. The overall difference appears to be driven by the *transformation* subscore, which was significantly higher in VR (t(34.7) = 2.50,p = .017). This finding aligns with the theory that the constant visibility of one's own body in MR can create a sensory conflict that hinders the feeling of fully ”becoming” the avatar [20].

Moreover, the analysis revealed that participants in the VR group also reported a significantly higher degree of *disembodiment*—a feeling of being disconnected from their physical body—than the MR group, $t(32.5) = -2.18$, $p = .037$. This suggests that the total sensory occlusion of VR creates a more potent experience overall: it more effectively transforms the user into the avatar, but at the cost of increasing the sensation of being disembodied or detached from one's own physical self.

The potency of this immersive experience in VR was particularly evident in an anecdotal observation from one participant (ID 29). While manipulating their avatar's leg, this participant was observed to move their own corresponding real leg and remarked, ”I felt like I was moving with the avatar.” This is a powerful real-world example of the ”self-avatar follower effect” discussed in the literature, where a strong visuo-proprioceptive coupling can cause a virtual avatar's movements to influence the user's own body [7, 21].

Notably, this specific participant had the second-highest baseline proprioceptive accuracy score recorded in the study, indicating a strong innate sense of their body's position. Furthermore, their questionnaire responses aligned perfectly with the VR group's overall trend: they reported a high degree of disembodiment (a score of 2.33) alongside a strong overall sense of embodiment (4.22). This case is particularly illustrative, as it suggests that the transformative power of VR is potent enough to induce feelings of disembodiment—and even influence motor actions—in individuals who are, by objective measures, highly attuned to their own physical body.

6.3.1 Implications for Different XR Applications

This phenomenon, where a virtual limb can trigger reactions in the physical body, has profound implications for XR therapies, particularly for phantom limb pain. It directly supports the goal of addressing the sensorimotor mismatch; by creating a believable connection to a virtual "phantom" limb, these systems can help re-engage dormant neural pathways [3, 1].

Given this potential, and the overall findings for H2, our results suggest that designers of PLP therapies should prioritize the use of VR. The goal should be to leverage VR's unique ability to foster a high degree of both **transformation** and **disembodiment**. By immersing patients in an experience where they can disembody from their physical reality and fully "become" a virtual avatar with intact limbs, designers can create the potent conditions necessary for this therapeutic bidirectional coupling to occur.

6.4 Interpretation of Exploratory Findings

6.4.1 RQ1: The Role of Proprioception in Mixed Reality

Our first exploratory research question (RQ1) centered on the role of proprioception in different immersive environments and how it might correlate with alignment task accuracy. To account for the repeated-measures design, we analyzed this using linear mixed-effects models.

The model for positional accuracy revealed a significant interaction between proprioception and the immersion environment ($\beta = -0.19$, $p = .047$). A deeper look at this interaction showed that for participants in the MR group, there was a significant positive relationship between their baseline proprioceptive error and their positional alignment error ($\beta = 0.12$, $p = .039$). In contrast, no such relationship was found for the VR group. The model for rotational accuracy found no significant effects related to proprioception.

This finding suggests that in MR environments—where users must reconcile what they see with what they feel—those with a better innate sense of body position perform the alignment task more accurately. This supports claims that understanding one's relationship to the physical space is key for performance in MR [28]. It is plausible that for users with poorer proprioceptive ability, the cognitive load of processing both real-world visual cues and proprioceptive signals becomes overwhelming, leading to decreased performance. In the fully occluded VR environment, this specific conflict does not exist, which may explain why baseline proprioceptive ability was not a significant factor for performance.

6.4.2 RQ2: System Usability in MR

The ability to recognize the location of oneself in space attributes some degree of cognitive load. Additionally, the tools used in that environment can vastly influence this as well. RQ2 investigated the overall system usability within the MR and VR environments. An independent t-test was used to assess if the immersive environment had an effect on the usability of the lower-body control interfaces. The results indicated that there was no significant difference detected ($t(29) = 0.299$, $p = 0.7669$; $VR : 62.78$, $MR : 61.40$). However, it was observed that the variance in the MR group was larger (SD = 16.56) than in the VR group (SD = 10.97). While this difference was not statistically significant (F(18, 17) = 2.28, p = .096), this trend may suggest a more inconsistent or polarizing user experience in MR. When the tests were run again on the separated interfaces, there was still no identifiable significance.

6.4.3 RQ3: Task Completion Time and the Initial Learning Curve

The analysis of task completion time revealed a notable learning effect dependent on the immersive environment. While no significant difference was found in the total time spent across all three poses, a targeted analysis of pose order showed that participants in the VR group took significantly longer to complete their first assigned pose compared to those in the MR group. This initial time disparity disappeared in the second and third poses, suggesting a steeper initial learning or adaptation curve in the fully immersive VR environment.

A likely explanation for this is the difference in subjective experience between the two modalities, particularly concerning time perception and immersion. It can be speculated that the world-occluding nature of VR led to a greater sense of ”time loss” or dissociation from the real-world experimental setting during the initial task. This interpretation is strongly supported by the findings for H2, where these same VR participants reported a significantly higher sense of embodiment and transformation. Once they completed the first task, they may have become more aware of the experimental process, leading them to normalize their completion speed in subsequent poses. Conversely, the MR participants, who remained visually ”grounded” in the real world, likely maintained a more consistent awareness of the setting, resulting in a more stable task duration from the outset.

This finding—that the effect of immersion on task time is nuanced—aligns with the mixed results in the literature. For instance, while one study found MR to be faster in a pick-and-place task [28], another found it to be slower in a spatial memory task [41]. Our results contribute to this by suggesting that for novel body-alignment tasks, VR may impose a higher initial time cost due to its powerful immersive properties.

It is also worth noting that when analyzing the data by pose **type** instead of pose **order**, the results were more complex, with the laying pose in particular showing a significant time difference between conditions. However, the most robust finding is the learning effect tied to the first exposure to the task, as this effect was the primary driver of the variance in completion times.

6.4.4 System Interaction Metrics

The average time spent interacting with either of the two interfaces proved to be similar between immersion modes. When comparing the two interfaces directly, the Direct interface was interacted with far more than the Puppet interface. This is likely due to the nature of the task itself: Participants were asked to align their virtual avatar (which was controlled via the direct interface) with their real body. In order to fully understand the interaction differences between both interfaces in VR and MR environments, a more robust data collection scheme and experimental procedure must be implemented that can more accurately capture the behavior of the interfaces without the bias of the task.

In addition to the system recorded metrics, a few observations were made regarding the individual use cases of the Direct and Puppet interfaces. Some users neglected to use the Puppet interface during the trial and only interacted with it during the initial instructions/demonstrations. Of the seven users noted to not be using it, five of them were in MR. These users were observed to focus primarily on manipulating the direct interface (their player avatar). This is likely due to the amount of distracting information seen in MR environments coupled with the task itself. Because their focus was drawn to to the direct interface, it would have been easy to neglect the Puppet interface all while trying to parse real and virtual information during the completion of the alignment task. Users in VR, however, were only able to see a predetermined set of virtual objects. The full immersion of the VR environment controlled the focus of the VR users to be more aware of what virtual tools were at their disposal throughout the task.

6.4.5 Synthesizing Proprioception, Embodiment, and Performance

A particularly intriguing narrative emerges when synthesizing the findings for proprioception (RQ1) and embodiment (H2). The results present a fascinating paradox between the MR and VR conditions. In the MR condition, the outcome was straightforward: our model found a significant positive relationship where participants with higher baseline proprioceptive accuracy performed the alignment task better. This suggests they were better able to integrate real-world sensory information with the virtual task.

In contrast, the VR condition revealed a more complex relationship. Our model confirmed that the link between proprioception and accuracy seen in MR was absent in VR, a difference confirmed by a **significant interaction effect** ($p = 0.047$). The trend for the VR group was a non-significant negative relationship, which becomes compelling when viewed alongside the qualitative experience. To recall, It suggests a speculative but powerful hypothesis: individuals with a more refined proprioceptive sense may not be worse at the task in VR, but are instead more susceptible to the immersive and transformative power of the virtual environment. Their brains, being highly attuned to bodily signals, might more readily "give in" and accept the proprioceptive requirements of the virtual world, leading to a higher degree of embodiment and disembodiment from their physical selves.

This has profound implications for PLP therapy. As noted, many XR therapeutic goals are not necessarily to achieve perfect objective alignment, but rather to foster the subjective "feeling" of connection to the virtual limb. Our findings raise the possibility that VR-based PLP therapy could be most effective for patients who have retained a strong proprioceptive map of their phantom limb, as they may have a greater capacity to embody the virtual replacement. This, in turn, opens several clear avenues for future research: Is PLP treatment in VR more effective for users with greater pre-existing proprioceptive abilities? Is their ability to embody a virtual limb greater because of this? And why might this effect appear to contradict the intuitive expectation that strong proprioception would lead to better physical-virtual alignment?

6.5 Broader Implications

Results from this study leave several implications and points of considerations for lower-body design and PLP-targeted XR therapies.

Implications for HCI Research

For tasks requiring precision in the real world, such as the body alignment task outlined in this study, it shouldn't be automatically assumed that MR will offer greater precision by itself. Just as shown from the literature and the results from this study, the interface design and the task itself are critical factors.

6.5.1 Implications for Therapeutic XR Design

For XR therapies targeting suggestive feelings, such as SoE or body ownership for PLP, our results suggest designers should use VR environments, even if it doesn't offer greater accuracy in real-world precision tasks. Our **grounded** hypothesis suggests that for therapies involving prosthetic training, where real-world precision and embodiment intertwine, the user's connection to the physical environment is paramount. Designers of such MR systems must carefully consider the cognitive load imposed on the user, especially in relation to their baseline proprioceptive abilities.

That said, for applications in VR, the user's baseline proprioception could be even more critical for effect therapies. Designers for these therapies could possibly consider two avenues for PLP treatment: For amputee patients with lower baseline proprioceptive abilities, treatment within MR environments could be *more effective* than those in VR. These individuals, as shown in the MR trend from RQ1, might rely more visual input from their surroundings **and** their real body for performance-based tasks. This however, needs more exploration and a more in-depth analysis on the relationship of proprioceptive abilities and having PLP symptoms in the first place. Likewise, for amputee patients with greater proprioceptive abilities, a strong, cohesive VR environment might be preferred for inducing a strong sense of embodiment and body ownership in the avatar.

6.6 Limitations and Future Work

This study, while reinforcing existing ideas and revealing new theories, carried several limitations that inhibited a fully comprehensive picture. The discoveries made in this research further highlight the complex nuances between VR and MR that need further analysis. The following sections outline these limitations and propose potential avenues for future research.

6.6.1 Limitations of the Current Study

Ground-Truth Measurement

The method for measuring ground-truth body pose relied on participants placing physical controllers on anatomical landmarks. This approach is prone to inconsistencies in placement, user error, and potential tracking failures, particularly when direct line-of-sight is lost. This introduces a source of measurement noise that could have reduced the statistical power to detect more subtle differences in alignment accuracy.

Avatar Personalization

The player avatar used in this study lacked the ability to scale dynamically between users. This meant that users with longer limb segments or who had overall height differences were not accounted for. While the effects of this discrepancy were accounted for in the positional error analysis through the use of Procrustes alignment, the effects of seeing a slightly *off* virtual avatar superimposed on ones body could potentially negatively influence feelings of embodiment in both MR and VR [20]. The avatar used in this study was also lacking in detail and not made to resemble a particular sex. Furthermore, the avatar used in this user study had its upper half hidden. Only the waist and legs were seen by the user to concentrate on the lower body alignment. Future adaptations or research should investigate whether these findings hold true beyond the lower body to verify consistency.

System Data Logging

Due to an implementation oversight, interaction counters were incremented in Unity's Update loop rather than `FixedUpdate`, and the actual runtime frame rate was not recorded. As a result, the frames spent interacting were used. This measurement is still reliable, but may have downsides in terms of unexpected lag or latency issues that could result in frame drop between participants. Because of this, the data was not used in any formal statistical analysis. The task assigned to the users was biased towards the Direct interface. To truly compare the Puppet and the Direct, a new experiment would need to be designed and its effects examined in both VR and MR enthronements.

Interpretation of Proprioception Correlation

The relationship between baseline proprioception and alignment accuracy needs more clarity. While a positive correlation was detected for positional accuracy in MR, this finding should be interpreted with caution, as it may not represent a robust trend without further investigation in a more controlled context. Moreover, a larger sample size and perhaps a 'simpler' alignment interface could provide a clearer picture on this trend. Additionally, the potential relationship between proprioception and the subjective experience of embodiment warrants a more dedicated investigation.

Additional Measurements and Future Work

This thesis focused on the primary hypotheses related to performance, embodiment, and usability. It should be noted that additional subjective data were collected during the study, including the Spatial Presence Experience Scale (SPES), the Inclusion of Other in the Self (IOS) Scale, and a single-item survey to assess embodiment between each alignment task. A full analysis of these complementary measures was beyond the scope of the current research but presents an opportunity for future work to explore the relationships between spatial presence, self-other overlap, and the core findings presented here. The decision to exclude the inter-session survey from the primary analysis was made to focus on the results from the more comprehensive and widely validated post-session Embodiment Scale [40].

6.6.2 Directions for Future Research

Improving Measurement Fidelity

A primary limitation of this study was the method for measuring ground-truth body pose. While using the built-in tracking of Quest 3 controllers placed on anatomical landmarks was practical, this approach is susceptible to tracking noise, drift, and occlusion, which limits its precision. To enhance measurement fidelity, future work should integrate external tracking technologies. For example, depth sensors like the Microsoft Kinect can estimate body pose without markers, offering a balance of accuracy and convenience, as has been demonstrated in therapeutic systems [4].

Alternatively, other technologies offer different advantages. Optical motion capture systems provide highly precise marker-based tracking, while limb-attached IMUs offer continuous orientation data. However, each of these technologies presents its own trade-offs in terms of user comfort, cost, and environmental constraints, making the choice of tracking system dependent on the specific goals and context of the future study.

Investigating Avatar Personalization

While prior work has investigated avatar fidelity, a crucial question remains regarding its role in therapeutic contexts, especially concerning the "uncanny valley" of embodiment. This is particularly relevant given our finding that VR produced both high transformation and high disembodiment. A future study could therefore investigate the trade-offs between different levels of avatar personalization for PLP therapy. For instance, would a hyper-realistic, custom-scaled avatar that perfectly resembles the patient enhance the sense of transformation, or would imperfections (or unintended uncanniness) create a worse experience that increases disembodiment? Additional explorations could investigate the usefulness of such an avatar in M environments and the effect it might have on cognitive load given that a more "realistic" avatar is expected to be more capable of blending in to the surroundings.

Proprioception and the Grounding Theory

The role of proprioception and embodiment could be paramount to the success of virtual therapies as it ties directly into inducing the sense of body ownership required to manage effects of PLP. Future investigations need to be made exploring the extent that proprioceptive abilities and cognitive load influence task performance in MR.

The grounding theory, found from the results of H1, theorizes that physical contact with the floor during MR experiences could potentially lead to increased cognitive loads for its users over those in VR leading to decreased performance. This could be due to the mismatch between the real world and the virtual world as well as the effect of proprioception in processing all of the information. A future study could examine exactly what limitations not being physically grounded might have between difference XR environments.

Proprioception and Embodiment

Furthermore, while the current study did not find a significant correlation between a participant's baseline proprioceptive ability and their subjective embodiment scores, exploring this relationship more deeply presents an excellent avenue for future research. A study with a larger sample would have greater statistical power to detect more subtle relationships. It would be particularly valuable to investigate if proprioceptive abilities can predict a user's capacity for embodiment, and how they relate to interoceptive feelings, which could have important implications for selecting candidates for certain types of XR therapy or training.

6.7 Discussion Conclusion

The user study presented in this thesis revealed an interesting trade-off between MR and VR environments for lower-body avatar alignment. Contrary to our hypothesis, we found no significant performance difference between the conditions; the expected real world accuracy benefit of MR's visual feedback did not manifest. However, VR demonstrated a significantly stronger sense of embodiment. Our exploratory findings suggest the experience in MR is more complex, with performance being correlated to a user's baseline proprioception. Taken together, these results suggest that for therapeutic goals centered on subjective experience, designers should favor the more consistent and highly embodying environment of VR.

Chapter 7

Conclusion

This chapter provides an overview on the motivation, design, and execution of the user study presented in this thesis as well as lists the key finds and critical takeaways from this work.

7.1 Summary of the Thesis

Phantom limb pain (PLP) is a painful sensation that affects over 85% of amputees, and can manifest from a mild discomfort to a painful, burning sensation in the position of the missing limb [23, 14]. Traditionally, PLP has been managed through a combination of opioid prescriptions and physical therapies. While these have been effective at mitigating the debilitating effects of PLP, persistent use of opioids and accessibility to effective, personalized therapies present clear disadvantages.

Recently, the rise of accessible virtual reality (VR) devices has sparked interest in using Extended Reality (XR) therapies as alternatives. Many of these approaches reduce PLP by placing a virtual limb where the phantom is perceived, promoting a sense of embodiment—especially body ownership—that can trick the brain into feeling virtual sensations [17, 24]. Classic mirror therapy (MT) achieves a similar effect using reflected images of the intact limb [13], but XR versions replace the mirror with a real-time rendered virtual limb. These therapies typically rely on motion controllers, optical tracking, or hand tracking to drive avatar movement [7, 20], and studies show they are as effective as traditional MT [43], with added benefits like improved engagement and flexible, gamified experiences [36].

However, while these methods are effective for upper-limb amputees, a distinct challenge remains for the majority of the amputee population: the lack of robust, accessible, and intuitive control interfaces for the lower-body [22, 39]. Developing such an interface requires careful consideration of the immersive environment. While existing literature points to a key trade-off between the high embodiment of Virtual Reality (VR) and the potential performance benefits of Mixed Reality (MR), this trade-off is not well understood in the specific context of lower-body alignment tasks. This gap in knowledge is the primary problem this research aimed to address.

To investigate this problem, two novel lower-body control interfaces were designed and implemented. A user study was then conducted with 40 (healthy, non-amputee) participants using a 2 (Immersion: MR vs. VR) x 3 (Pose: Laying, Sitting, Standing) mixed-factorial design. The study evaluated body alignment performance by comparing the positional and rotational accuracy of the virtual avatar to the user's real body, along with assessing the subjective sense of embodiment, system usability, and task completion time.

This research revealed an interesting trade-off between MR and VR environments for lower-body avatar alignment. Contrary to our H1, there was no significant performance difference between the conditions. The real world accuracy benefit of MR visual feedback did not lead to the expected higher alignment accuracy we theorized. On the other hand, VR *did* demonstrate a significantly stronger sense of embodiment, confirming H2. Our exploratory findings suggest the experience in MR is more complex when compared to VR, with performance possibly being correlated to the intrinsic traits and innate abilities of user. Taken together, these results suggest that, for therapeutic technologies centered around subjective experience, designers should favor the more consistent and higher SoE inducing environment of VR- the primary component behind PLP therapies.

The results of the experiment outlined in this thesis revealed interesting trade-offs between VR and MR environments. MR users, contrary to the hypothesis, did not outperform VR users in the body alignment task. Similar to the mixed results uncovered in the literature, these results suggest that while MR does grant users visual cues from the real world, it does not correlate to greater accuracy for every task. The second claim tested in this user study found that participants in VR did, in fact, experience a greater reported level of embodiment than those in MR. This came at the cost of a longer task completion time for their first pose. This finding supports claims that VR may foster a stronger sense of embodiment. However, our exploratory results suggest the MR experience is more complex: performance appears linked to a user's baseline proprioception, and usability ratings show a non-significant trend toward greater variance. Taken together, these results suggest that for therapeutic goals centered on subjective experience, designers should favor the more consistent and highly embodying environment of VR.

7.2 Primary Contributions

This thesis presents the following contributions to the field:

1. The design and implementation of two novel interfaces for lower-body avatar control and alignment.

2. Empirical evidence demonstrating the trade-off between performance and embodiment in MR versus VR for an alignment task.

3. n analysis exploring the role of proprioception and the nature of the user experience in MR environments.

7.3 Final Remarks and Future Outlook

This thesis detailed the design and implementation of two novel lower-body control interfaces, which were evaluated in a user study comparing the effects of MR and VR environments. The results revealed a critical trade-off: contrary to our initial claims, MR did *not* produce greater body-alignment accuracy than VR. However, participants in VR experienced a *significantly higher* subjective sense of embodiment (SoE), despite taking much longer to complete the initial task. Both interfaces received similarly moderate system usability scores, though responses from MR participants showed a non-significant trend toward greater variability, possibly suggesting a more inconsistent user experience. This variability, combined with the lack of the performance advantage MR was expected to have over VR, may be explained by individual differences. Our exploratory findings suggest that alignment accuracy in MR environments is influenced by users' innate proprioceptive abilities—linking performance to one's sense of body awareness. While further investigations are still needed to identify cognitive load as an additional factor, the discoveries made through this user study could be critical in the treatment and management of both upper and lower-body PLP.

Bibliography

[1] Nicole E Acerra, Tina Souvlis, and G Lorimer Moseley. “Stroke, complex regional pain syndrome and phantom limb pain: can commonalities direct future management?” In: *Journal of Rehabilitation Medicine* 39.2 (2007), pp. 109–114. DOI: 10.2340/16501977-0027.

[2] Shah Mahsoom Ali et al. “Medical Education, Training and Treatment Using XR in Healthcare”. In: *Proceedings of the 2023 25th International Conference on Advanced Communication Technology (ICACT)*. 2023, pp. 388–393. DOI: 10.23919/ICACT56868.2023.10079321.

[3] Elisabetta Ambron et al. “Virtual Reality Treatment Displaying the Missing Leg Improves Phantom Limb Pain: A Small Clinical Trial”. In: *Neurorehabilitation and Neural Repair* 35.12 (2021). Epub 2021 Oct 27, pp. 1100–1111. DOI: 10.1177/15459683211054164.

[4] D. Annapureddy et al. “A novel mixed reality system to manage phantom pain in-home: results of a pilot clinical trial”. In: *Frontiers in Pain Research* 4 (2023), p. 1183954. DOI: 10.3389/fpain.2023.1183954. URL: https://doi.org/10.3389/fpain.2023.1183954.

[5] Laura Barca and Giovanni Pezzulo. “Keep your interoceptive streams under control: An active inference perspective on anorexia nervosa”. In: *Cognitive, Affective, & Behavioral Neuroscience* 20 (2020), pp. 427–440. DOI: 10.3758/s13415-020-00777-6.

[6] Rena Bayramova et al. "The Role of Vision and Proprioception in Self-Motion Encoding: An Immersive Virtual Reality Study". In: *Attention, Perception, & Psychophysics* 83.7 (2021), pp. 2865–2878. DOI: 10.3758/s13414-021-02344-8. URL: https://doi.org/10.3758/s13414-021-02344-8.

[7] Loën Boban et al. "Unintentional synchronization with self-avatar for upper- and lower-body movements". In: *Frontiers in Virtual Reality* Volume 4 - 2023 (2023). ISSN: 2673-4192. DOI: 10.3389/frvir.2023.1073549. URL: https://www.frontiersin.org/journals/virtual-reality/articles/10.3389/frvir.2023.1073549.

[8] Niklas Bockelmann, Roland Fischer, and Gabriel Zachmann. "Embodiment in Virtual Environments - Analyzing the Effects of Latency and Avatar Representation". In: *2024 International Conference on Cyberworlds (CW)*. Los Alamitos, CA, USA: IEEE Computer Society, 2024, pp. 17–24. DOI: 10.1109/CW64301.2024.00012. URL: https://doi.ieeecomputersociety.org/10.1109/CW64301.2024.00012.

[9] Matthew Botvinick and Jonathan Cohen. "Rubber Hands 'Feel' Touch That Eyes See". In: *Nature* 391 (1998), p. 756. DOI: 10.1038/35784. URL: https://doi.org/10.1038/35784.

[10] John Brooke. "SUS: A "quick and dirty" usability scale". In: *Usability Evaluation in Industry*. Ed. by Patrick W. Jordan et al. London: Taylor & Francis, 1996, pp. 189–194.

[11] Steve Bryson. *Virtual Reality: A Definition History - A Personal Essay*. 2013.

[12] Jennifer L. Campos et al. "Virtual Hand Illusion in Younger and Older Adults". In: *J. Rehabil. Assist. Technol. Eng.* 8 (2021), p. 20556683211059389. DOI: 10.1177/20556683211059389.

[13] J.C.-W. Cheung et al. "X-reality for phantom limb management for amputees: a systematic review and meta-analysis". In: *Engineering Regeneration* 4 (2023), pp. 134–151.

[14] Krista L. Collins et al. "A Survey of Frozen Phantom Limb Experiences: Are Experiences Compatible With Current Theories". In: *Frontiers in Neurology* 9 (2018), p. 599. DOI: 10.3389/fneur.2018.00599. URL: https://www.frontiersin.org/articles/10.3389/fneur.2018.00599.

[15] Elif Eshkevari et al. "Increased plasticity of the bodily self in eating disorders". In: *Psychological Medicine* 42.4 (2012), pp. 819–828. DOI: 10.1017/S0033291711002091.

[16] James Coleman Eubanks et al. "A Preliminary Embodiment Short Questionnaire". In: *Frontiers in Virtual Reality* Volume 2 - 2021 (2021). ISSN: 2673-4192. DOI: 10.3389/frvir.2021.647896. URL: https://www.frontiersin.org/journals/virtual-reality/articles/10.3389/frvir.2021.647896.

[17] Charles Faure, Annabelle Limballe, and Hugo A. Kerhervé. "Fooling the Brain, Fooling the Pain: The Role of Mirror Therapy and Modern Uses in Virtual Reality". In: *Frontiers for Young Minds* 7 (2019), p. 91. DOI: 10.3389/frym.2019.00091. URL: https://kids.frontiersin.org/articles/10.3389/frym.2019.00091.

[18] Ranya El-Gabalawy et al. "Treating Lower Phantom Limb Pain in the Postoperative Acute Care Setting Using Virtual Reality: Protocol for a 4-Phase Development and Feasibility Trial". In: *JMIR Research Protocols* 14 (2025), e68008. DOI: 10.2196/68008. URL: https://www.researchprotocols.org/2025/1/e68008.

[19] Santino Gaudio, Samantha J. Brooks, and Giuseppe Riva. "Neurofunctional correlates of body image disturbance in anorexia nervosa and bulimia nervosa: A review of literature". In: *Journal of Psychosomatic Research* 77.3 (2014), pp. 169–175. DOI: 10.1016/j.jpsychores.2014.05.003.

[20] Adélaïde Genay, Anatole Lécuyer, and Martin Hachet. "Virtual, Real or Mixed: How Surrounding Objects Influence the Sense of Embodiment in Optical See-Through Experiences?" In: *Frontiers in Virtual Reality* Volume 2 - 2021 (2021). ISSN: 2673-4192. DOI:

10.3389/frvir.2021.679902. URL: https://www.frontiersin.org/journals/virtual-reality/articles/10.3389/frvir.2021.679902.

[21] Mar Gonzalez-Franco et al. "The Self-Avatar Follower Effect in Virtual Reality". In: *IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE. 2020. DOI: 10.1109/VR46266.2020.00034.

[22] K. Hali et al. "Use of virtual reality for the management of phantom limb pain: a systematic review". In: *Disability and Rehabilitation* 46.4 (2023), pp. 629–636. DOI: 10.1080/09638288.2023.2172222.

[23] Andrea A. Hanyu-Deutmeyer, Marco Cascella, and Marc A. Varacallo. "Phantom Limb Pain". In: *StatPearls*. Updated 2023 Aug 4. Available from: https://www.ncbi.nlm.nih.gov/books/NBK448188/. Treasure Island (FL): StatPearls Publishing, 2025.

[24] Daniel S. Harvie et al. "Using visuo-kinetic virtual reality to induce illusory spinal movement: the MoOVi Illusion". In: *PeerJ* 5 (2017), e3023. DOI: 10.7717/peerj.3023. URL: https://doi.org/10.7717/peerj.3023.

[25] Bartal Henriksen et al. "Comparison of Movements in a Virtual Reality Mirror Box Therapy for Treatment of Lower Limb Phantom Pain". In: *Proceedings of the 13th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2018) – Volume 1: GRAPP*. Ed. by Dominique Bechmann, Ana Paula Claudio, and Jose Braz. SciTePress, 2018, pp. 167–174. DOI: 10.5220/0006537801670174.

[26] Chaltu D. Kenea et al. "Immersive Virtual Reality in Stroke Rehabilitation: A Systematic Review and Meta-Analysis of Its Efficacy in Upper Limb Recovery". In: *Journal of Clinical Medicine* 14.6 (2025), p. 1783. DOI: 10.3390/jcm14061783.

[27] DoHyung Kim, Halim Yeo, and Kyoungju Park. "Effects of an Avatar Control on VR Embodiment". In: *Bioengineering* 12.1 (2025). ISSN: 2306-5354. URL: https://www.mdpi.com/2306-5354/12/1/32.

[28] Advay Kumar et al. *Mixed Reality Outperforms Virtual Reality for Remote Error Resolution in Pick-and-Place Tasks*. Feb. 2025. DOI: 10.48550/arXiv.2502.06141.

[29] Eva Lendaro et al. "Real-time Classification of Non-Weight Bearing Lower-Limb Movements Using EMG to Facilitate Phantom Motor Execution: Engineering and Case Study Application on Phantom Limb Pain". In: *Frontiers in Neurology* Volume 8 - 2017 (2017). ISSN: 1664-2295. DOI: 10.3389/fneur.2017.00470. URL: https://www.frontiersin.org/journals/neurology/articles/10.3389/fneur.2017.00470.

[30] Zhipeng Li et al. "Modeling the Impact of Visual Stimuli on Redirection Noticeability with Gaze Behavior in Virtual Reality". In: *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '25)*. CHI '25. Yokohama, Japan: Association for Computing Machinery, 2025. DOI: 10.1145/3706598.3713392. URL: https://doi.org/10.1145/3706598.3713392.

[31] C. Lopez, P. Halje, and O. Blanke. "Body ownership and embodiment: Vestibular and multisensory mechanisms". In: *Neurophysiologie Clinique/Clinical Neurophysiology* 38.3 (2008), pp. 149–161. ISSN: 0987-7053. DOI: https://doi.org/10.1016/j.neucli.2007.12.006. URL: https://www.sciencedirect.com/science/article/pii/S098770530800018X.

[32] Tomasz Mazuryk and Michael Gervautz. *Virtual Reality — History, Applications, Technology and Future*. Tech. rep. Institute of Computer Graphics, Vienna University of Technology, 1996. URL: https://www.cg.tuwien.ac.at/research/publications/1996/mazuryk-1996-VRH/mazuryk-1996-VRH-paper.pdf (visited on 06/15/2025).

[33] Meta Platforms, Inc. *Meta Quest 3*. Accessed: 2025-06-15. Meta. 2023. URL: https://www.meta.com/quest/quest-3/.

[34] Microsoft. *HoloLens 2 Hardware*. Microsoft Learn. 2025. URL: https://learn.microsoft.com/en-us/hololens/hololens2-hardware (visited on 06/15/2025).

[35] Aske Mottelson et al. "A Systematic Review and Meta-analysis of the Effectiveness of Body Ownership Illusions in Virtual Reality". In: *ACM Transactions on Computer-Human Interaction* 30 (2023). DOI: 10.1145/3590767.

[36] Ronni Nielsen et al. "Comparison of Body Positions in Virtual Reality Mirror Box Therapy for Treatment of Phantom Limb Pain in Lower Limb Amputees". In: *Proceedings of the Virtual Reality International Conference - Laval Virtual 2017*. VRIC '17. Laval, France: Association for Computing Machinery, 2017. ISBN: 9781450348584. DOI: 10.1145/3110292.3110307. URL: https://doi.org/10.1145/3110292.3110307.

[37] S. Oyama et al. "Insights and trends review: Use of extended reality (xR) in hand surgery". In: *Journal of Hand Surgery (European Volume)* 50.6 (2025), pp. 762–770. DOI: 10.1177/17531934241313208.

[38] HyungGoo Park and Olaf Blanke. "Interoception and the body: How signal accuracy and awareness shape body representations". In: *Cortex* 124 (2020), pp. 138–149. DOI: 10.1016/j.cortex.2019.10.021.

[39] Javad Parvizi and Gregory K. Kim. "Chapter 9 - Amputation of the Lower Limb". In: *High Yield Orthopaedics*. Ed. by Javad Parvizi and Gregory K. Kim. Philadelphia: W.B. Saunders, 2010, pp. 17–19. ISBN: 978-1-4160-0236-9. DOI: https://doi.org/10.1016/B978-1-4160-0236-9.00018-3. URL: https://www.sciencedirect.com/science/article/pii/B9781416002369000183.

[40] T. C. Peck and M. Gonzalez-Franco. "Avatar embodiment. A standardized questionnaire". In: *Frontiers in Virtual Reality* 1 (2021), p. 575943. DOI: 10.3389/frvir.2020.575943. URL: https://doi.org/10.3389/frvir.2020.575943.

[41] David Ponce et al. "Extended Reality to Assess Short-Term Spatial Memory—A Comparative Study of Mixed Reality, Augmented Reality, and Virtual Reality". In: *Sensors* 24 (Dec. 2024), p. 7938. DOI: 10.3390/s24247938.

[42] Charlie Portingale, Ioannis Kontaris, and Janet H. Bultitude. "Phantom limb discomfort reduced using virtual reality with body visualisation in people with limb loss". In: *Journal of NeuroEngineering and Rehabilitation* 20.1 (2023), pp. 1–14. DOI: 10.1186/s12984-023-01166-y.

[43] Christopher Rajendram et al. "Efficacy of mirror therapy and virtual reality therapy in alleviating phantom limb pain: a meta-analysis and systematic review". In: *BMJ Military Health* 168 (Jan. 2022), e002018. DOI: 10.1136/bmjmilitary-2021-002018.

[44] Rachael L. Shafer et al. "Visual Feedback During Motor Performance Is Associated with Increased Complexity and Adaptability of Motor and Neural Output". In: *Behavioural Brain Research* 376 (2019), p. 112214. DOI: 10.1016/j.bbr.2019.112214. URL: https://doi.org/10.1016/j.bbr.2019.112214.

[45] František Škola and Fotis Liarokapis. "Augmenting the Rubber Hand Illusion". In: *The Visual Computer* 32 (2016), pp. 761–770. DOI: 10.1007/s00371-016-1246-8.

[46] Chandrasekhar B. Sudhir. *Human Pose Estimation 101.* Accessed: 2025-06-18. 2021. URL: https://github.com/cbsudux/Human-Pose-Estimation-101.

[47] Alexander Toet et al. "Toward Enhanced Teleoperation Through Embodiment". In: *Frontiers in Robotics and AI* 7 (2020), p. 14. DOI: 10.3389/frobt.2020.00014.

[48] L. Tremosa. *Beyond AR vs. VR: What is the Difference between AR vs. MR vs. VR vs. XR?* Interaction Design Foundation (IxDF). 2025. URL: https://www.interaction-design.org/literature/article/beyond-ar-vs-vr-what-is-the-difference-between-ar-vs-mr-vs-vr-vs-xr (visited on 06/15/2025).

[49] Irene Valori et al. "Proprioceptive Accuracy in Immersive Virtual Reality: A Developmental Perspective". In: *PLOS ONE* 15.1 (2020), e0222253. DOI: 10.1371/journal.pone.0222253. URL: https://doi.org/10.1371/journal.pone.0222253.

[50] Sam Van Damme et al. “Impact of Latency on QoE, Performance, and Collaboration in Interactive Multi-User Virtual Reality”. In: *Applied Sciences* 14.6 (2024). ISSN: 2076-3417. DOI: 10.3390/app14062290. URL: https://www.mdpi.com/2076-3417/14/6/2290.

[51] Chang Wang and Sridhar Mahadevan. “Manifold Alignment using Procrustes Analysis”. In: *University of Massachusetts Amherst - ScholarWorks* (2008). Accessed: 2025-06-18. URL: https://scholarworks.umass.edu/entities/publication/5768f2cd-e235-4e0a-b844-2ba8b9882b77.

[52] Linfeng Wu and Karen Chen. “Embodiment of Virtual Body and Extremities With Movement Control in Reaching Tasks Using Virtual Reality”. In: *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 68 (2024). DOI: 10.1177/10711813241273495.

[53] J. Zhang, V. Lu, and V. Khanduja. “The Impact of Extended Reality on Surgery: A Scoping Review”. In: *International Orthopaedics (SICOT)* 47 (2023), pp. 611–621. DOI: 10.1007/s00264-022-05663-z. URL: https://doi.org/10.1007/s00264-022-05663-z.