Institute for Visualization and Interactive Systems

University of Stuttgart
Universitätsstraße 38
D–70569 Stuttgart

Masterarbeit

LiftVR: VR-Based Motion Guidance System Teaching Back-Friendly Lifting

Andreas P. Farley

Course of Study: Softwaretechnik

Examiner: Prof. Dr. Michael Sedlmair

Supervisor: Xingyao Yu M. Sc,
Dr. André Tomalka ,
Prof. Dr. Tobias Siebert ,
Sebastian Rigling M. Sc,

Commenced: November 28, 2022

Completed: July 10, 2023

Abstract

A third of the European Union's workforce is frequently tasked with heavy lifts during their shifts, risking the development of musculoskeletal disorders (MSD). The prevalence of MSD resulting from improper lifting techniques highlights the need for effective training methods. Virtual reality offers promising opportunities to improve motion training through immersive experiences, however, no system is currently available for teaching lifting exercises while its users interact with real weight. Additionally, gaps in the domain of virtual full body motion guidance necessitate further exploration. The effects of different concurrent guidance visualizations on user performance are not known for full body systems. Furthermore, it has not been investigated how terminal feedback affects a students improvement when used contemporaneously with concurrent systems.

This thesis presents LiftVR, an accessible, single camera, full body motion guidance system which implements two concurrent visual guidance systems (state of the art and novel), and multiple terminal feedback systems. In the pursuit of improving visual guidance, a novel abstract guidance system named Zone has been implemented, contrasting the currently dominant approach of using humanoid entities for expert visualization. We made use of the symmetrical nature of lifting motions, which results in Zone having less visual clutter and fewer simultaneously active entities. Additionally, we provide the user with multiple terminal feedback systems. For immersion, LiftVR places students in a fully equipped gym environment, including a barbell visualization. Finally, students can chose between different practice modes, perspectives, and speeds.

To evaluate LiftVR we performed a study in which participants had to practice using each of LiftVR's practice modes and speeds. Overall, LiftVR had a significant positive effect on students lifting techniques, regardless of age, height or weight. Our analysis shows that Zone outperforms the state of the art system in posture driven practice modes. Overall Zone students achieved lower errors and greater improvement across all of our measured variables. User feedback also verified Zone's better scores in terms of understandability, mental effort, and helpfulness. Regarding terminal feedback, the study reveals no statistically significant effects on student performance when used in conjunction with concurrent guidance systems. However, the absence of prescriptive or corrective terminal feedback in LiftVR leaves room for future exploration of its potential impact. The findings of this thesis offer valuable insights for designing and developing concurrent guidance and feedback systems in the field of full body motion guidance.

Contents

List of Figures

List of Tables

List of Algorithms

Acronyms

1PP first-person perspective. 32

3D three-dimensional. 20

3PP third-person perspective. 32

AR augmented reality. 19

CAVE cave automatic virtual environment. 20

DK development kit. 22

EU European Union. 15

EWMA exponential weighted moving average. 42

FOV field of view. 33

FPS frames per second. 22

HMD head-mounted device. 16

HSM half-silvered mirror. 34

HUD Heads-Up Display. 40

INSPO Institute of Sport and Movement Science. 17

LSD least significant difference. 70

M mean. 72

MR mixed reality. 19

MSD musculoskeletal disorders. 15

SD standard deviation. 72

SDK software development kit. 22

UI user interface. 39

VR virtual reality. 15

1 Introduction

In this introductory chapter we first describe the motivation for this thesis followed by our research objectives. Subsequently we describe the working environment and finally give an outline of the thesis' structure.

1.1 Motivation

Lifting objects off the ground incorrectly is one of the physical factors which purportedly contributes to back pain [WSS23]. Simultaneously, a third of the current workforce in the European Union (EU) is tasked with lifting heavy loads frequently during their shifts [PBC+17]. In consequence, musculoskeletal disorders (MSD) form the most frequent health related problem in Europe [WSS23]. The rising prevalence of MSD resulting from improper lifting techniques has raised concerns about the need for effective training methods that promote back-friendly practices.

Concomitantly, virtual reality (VR) technology has emerged as a promising tool for enhancing motion training experiences by providing immersive virtual environments. However, no VR motion guidance system currently exists to teach students lifting exercises whilst they interact with real weight. Furthermore, the significant knee and hip flexion during the lowest portion of a lifting motion make students difficult to track. Expensive, multi-camera, tracking setups are not accessible to most and thus not applicable to solve the problem at hand. The challenge is to develop an accessible system which can qualify for wide-spread personal use whilst ensuring high-quality tracking and visual performance. As such a system does not currently exist, it is unknown which effects VR could have to improve lifting techniques and prevent MSD.

In addition to these challenges, gaps in the domain of motion guidance make the effects of certain design decisions uncertain. Effects of different perspectives and devices are well understood, however, the actual stimuli providing guidance (be it visual, auditory, or haptic) are rarely evaluated against each other. Whilst researchers confirm their own systems to provide effective guidance, they do not compare the performance of their visual guidance mechanism to other state of the art approaches. The currently dominant approach for visual guidance is to superimpose a visualization of the student onto a humanoid entity performing the motion. Examples of this include skeletons, outlines, videos, or full body models. Nonetheless, they all assume human form or characteristics. In the domain of full body guidance, it has yet to be investigated how different approaches to this would affect student performance.

The effect of terminal feedback implemented into motion guidance systems is another source of uncertainty. Feedback is assistance dependent on the contemporaneous or prior performance of a student, and terminal indicates it is presented post motion. We find two types of systems using terminal feedback, those who are entirely reliant upon it and those who present it after concurrent motion guidance ceases. For the latter, no evaluations of the systems with and without terminal feedback enabled exist. Thus, it is unknown how effective terminal feedback is when used in conjunction with a concurrent motion guidance system. Furthermore, we do not know how to integrate such a system in a way that makes it effective as we currently know nothing of its effect.

1.2 Research Objectives

The goal of this thesis is to implement a full body motion guidance system named *LiftVR*. It must teach a back-friendly lifting technique without the necessity for body markers, tracking suits or bands, multiple cameras, or base stations. Thereby, costs and hardware dependence are kept minimal. Additionally, the system must work regardless of user height, weight, or age. This would fulfill our requirements for an accessible system.

Furthermore, it must implement two concurrent visual guidance systems, one akin to state of the art approaches and one designed specifically for the motion at hand. Finally, terminal feedback systems, state of the art and novel, must be made available to students. As the requirements are steep and plentiful, implementing LiftVR is a considerable challenge. Nevertheless, its completion allows us to answer the following research questions:

RQ 1 - What effect does accessible full body motion guidance in virtual reality have on students lifting techniques?

Our first question addresses the lack of knowledge for the effectiveness of VR based motion training for lifting motions. Furthermore, no system currently exists in which the student interacts with weights whilst wearing a head-mounted device (HMD). It is unclear if current single-camera, no marker technology can be used effectively to implement a system to teach such a motion, nevertheless cause a positive effect.

RQ 2 - What are the effects of different concurrent guidance visualizations on student performance?

Humanoid entities are the predominant visual choice for state of the art full body guidance systems. However, it is unknown how these compare to other, perhaps more abstract visualizations, as these have never been implemented in full body systems. As such, we aim to challenge the existing state of affairs by investigating the effects of different visual systems on performance.

RQ 3 - How does the presence or absence of terminal feedback affect student performance?

Whilst we know that terminal feedback, when used standalone, helps students tremendously, its interactions with concurrent guidance systems are unknown. Thereby, we are uncertain of the differences between students using LiftVR with and without feedback, as both still have access to the other systems.

1.3 Working Environment

LiftVR is developed in close cooperation with the Institute of Sport and Movement Science (INSPO). Critically, they provide the lifting technique taught by the system. A data set of three lifts with varying weights was performed by an expert and recorded in INSPO's high quality motion capture laboratory. This allows us to forgo independent research, development, and capture of a proper lifting technique and focus on the implementation of LiftVR and the evaluation of its systems.

Whilst INSPO provides us the technique, a requirement of this thesis is to help them create new data sets of students performing lifts during various stages of training. As such, part of our study must be performed in the motion capture laboratory with students being recorded by a second, highly accurate system.

It was emphasized that LiftVR shall be an ongoing solution, and development is to be continued based on findings made by the thesis and general progress in the domain of full body motion guidance. Thereby, we must ensure understandability, maintainability, flexibility and scalability when engineering the software, to pass it to others in good faith. We will not evaluate our system based on these criteria, as the focus of this work is on motion guidance, however, during implementation we will take care to ensure these criteria are as much as possible.

1.4 Outline

We provide a brief outline of the thesis to come before going into any detail. Seven chapters form the entirety of our work, in the following order with content as described:

Chapter 2 (Foundations) gives important background information on technologies, concepts, and terminology required to understand the content of our thesis.

Chapter 3 (Related Work) grants insight into related full body motion guidance works in addition to challenges and gaps thereof. Works that contribute design principles are also included.

Chapter 4 (Implementation) explains design decisions and describes implementations of the various and systems which form LiftVR.

Chapter 5 (Evaluation) contains our research methodology and approach, the resulting study and subsequent result presentation and analysis.

Chapter 6 (Discussion) addresses our research objectives in relation to findings from our analysis, presents implications based upon these and finally discusses limitations to our work.

Chapter 7 (Conclusion) presents a summary of our key findings and provides an outlook on future development of LiftVR and general work.

2 Background

This chapter serves as a critical foundation and presents background information crucial to comprehend the systems, design choices, and terminology utilized in the thesis. We begin with an introduction to mixed reality (MR) and the technologies behind it. Importance is given to understanding the different tracking systems of VR devices.

Afterwards, we introduce the concept of motion capture and describe systems relevant to this thesis. The type of motion capture system utilized brings considerable design implications and restrictions, which is why we lay importance on knowing how they work.

Subsequently, we give a brief introduction to motion training in reality, its challenges, and potential solutions from MR. This is followed by an extensive dive into motion guidance terminology and definitions, which is crucial to differentiate and understand the classification schemes used in this thesis.

Finally, a short explanation of what entails a injury-preventive lifting technique is given.

2.1 Mixed Reality

Advances in technology make it increasingly possible to extend real environments by digital means. Digital enhancement or even complete replacement of reality is known as MR [MTUK95]. Increasing levels of virtual replacement can be placed into a continuum of reality-virtuality as seen in Figure 2.1.

Augmented reality (AR) is closest to reality and entails integrating virtual objects into an real environment.

Augmented virtuality steps closer to virtuality by providing a virtual environment and virtual objects, however, real objects are still integrated into the scene.

Virtual reality is the complete immersion of users into a computer-generated environment [SH16]. Immersion is determined by the quality of the visuals, interactions of the user with the virtual environment and system response to user action [SW97].

2.1.1 Devices

The devices used to deliver MR are plentiful and have changed vastly throughout time. For instance, one of the oldest VR systems had users sitting in a chair , staring into a large box watching moving pictures whilst being stimulated by vibrations, wind, and even smell [Hei62]. More modern

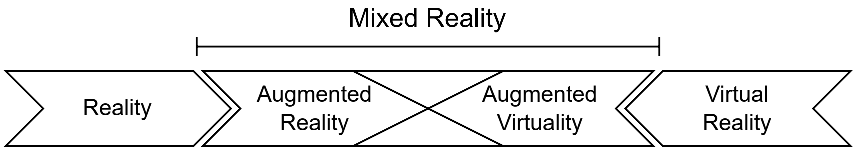

Figure 2.1: Visualization of the Reality-Virtuality Continuum based on Schmalstieg and Hollerer [SH16]

.

approaches include the cave automatic virtual environment (CAVE), a room which simulates a virtual environment through projection of high quality video [Cre03]. Users must wear shutter or active stereo glasses to receive a three-dimensional (3D) effect.

For LiftVR, we are interested in HMDs, which modifies images from near-eye displays through special lenses to display these in an immersive way [KBBD17]. These images could either be entirely computer generated or simply footage sourced from a cameras attached to the device. Other technologies, such as eye-tracking, latency, and input devices are crucial for the functionality of HMDs, However, we will only give further detail on the head-tracking systems, as only these have direct design implications on our implementation of LiftVR and must be understood.

2.1.2 Outside-In Tracking

Outside-in tracking is a technology utilized in VR and motion capture systems to accurately measure the user's movements and position within a real environment. In the case of VR, this movement is mapped into the virtual environment, providing the user with the immersive experience of physically navigating within the virtual environment. It involves external sensors or cameras placed in the physical environment, which track the movements of markers on the HMD or controllers, allowing the system to precisely calculate their position and orientation in real-time. This method offers several advantages, such as high tracking accuracy, reduced latency and tolerance to marker obfuscation when multiple cameras are deployed. Drawbacks include dependence on additional hardware and increased costs therefrom, as well less flexibility on when and where you can deploy the HMD.

2.1.3 Inside-Out Tracking

In contrast, inside-out tracking captures the user's movements and position without the need for external hardware. Instead, it relies on sensors and cameras integrated directly into the HMD which constantly monitor the surrounding environment. This data is then used to estimate the devices position and orientation in real-time. One requirement for this is the analysis of visual information for the system to precisely determine the user's movements. This can make inside-out systems more sensitive to the real-world environments they are employed in. Furthermore, if the controller positions are tracked outside-in by the inside-out HMD, the system has no information on their position when not held in the vicinity of the HMD's cameras. However, the independent tracking and reduced hardware makes these systems cheaper and faster to deploy.

Grasping the differences between these tracking technologies is crucial to understand our design decision to restrict LiftVR of using an inside-out HMD. These are cheaper and require less effort to set up, which increases their accessibility. As such, no systems that would require base stations are permitted. The only outside-in tracking system to be used in tandem with LiftVR is related to the motion capture technology, which we discuss next.

2.2 Motion Capture Systems

Motion capture is the process of recording an entities movement. The focus lies not on the visual appearance of the entity, but rather its motion. Reciting all capture technologies is vast and out of scope for this work, however, we will introduce those relevant for this work and highlight the differences between them.

2.2.1 Marker-based motion capture

As mentioned in Section 2.1.2, outside-in tracking of markers attached to objects in a physical environment is widely used by motion capture systems. Popular options for attaching markers to humans is through the use of suits adorned with them, or sticking them onto the targets using tape. One such motion capture system is Vicon's Nexus [Vic23]. Due to the multiple cameras and plentiful markers, Vicon offers precision and is thus popular for biomechanics and sports science in addition to virtual entertainment production. The expert lifting demonstrations provided by INSPO were recorded with such a system.

Nonetheless, this precision comes at a high cost in usability. Users must attach trackers to their body, as seen in Figure 2.2, and be calibrated to a multi-camera system. Whilst tracking suits can be made for an individual, there is no "one-fit-all" size, complicating the integration of new users. The process of taping the markers to participants can take several minutes and potentially require the participant to strip certain clothing items. High precision equipment makes the systems expensive. In conclusion, these form of motion capture systems are impractical for wide-spread personal use, as there were not intended for them. As such, we investigated an other, more accessible option for motion capture in LiftVR.

2.2.2 Machine learning based motion capture

As an alternative to marker-based motion capture, machine learning approaches are becoming increasingly popular. These use a model trained on vast sets of posture data to deduce joint positions from camera images. Furthermore, they are not reliant on multiple cameras and can be used with low cost hardware such as a webcam.

Figure 2.2: Figure showing a Vicon body model with required markers placed upon it. Taken from the Vicon Nexus documentation [Vic23]

Azure Kinect Microsoft provides a low cost solution for accurate motion capture with the Azure Kinect development kit (DK) [NHA+20]. It includes a device, the Azure Kinect, which is equipped with a depth, infrared, and color camera, while also containing orientation sensors and a spatial microphone array. Whilst it is perfectly capable on its own, multiple Azure Kinect's can be combined to increase tracking accuracy. The plethora of different settings, setups, and software development kits (SDKs) has made the Azure Kinect a popular option for motion guidance research [EHG+21]. For the remainder of this work, we will refer to this device as the "tracking camera".

The most important feature for LiftVR in the Azure Kinect DK is the Body Tracking SDK. It provides access to a machine learning model which, with the Azure Kinect's sensor data, recognizes human bodies before it and deduces their joint positions. With the SDK, we attach a "tracker" to the Azure Kinect, which provides us with joint positions at roughly 30 frames per second (FPS). Using the joint and bone hierarchy model as seen in Figure 2.3, developers can reconstruct the tracked users skeleton.

OpenPose The Azure Kinect does not represent the pinnacle of accessible low budget motion capture, as OpenPose [Oso18] can be run on modest hardware using any webcam as the capture technology. Through combination with PoseFix post-processing [MCL19], higher accuracy can be achieved than using an Azure Kinect, however, with an average of eight FPS, it does not fulfill the minimum requirement of 25 FPS to be categorized as a real-time capable system [IOA21].

When working with good lighting conditions, low-clutter environments and upper body motions, OpenPose can provide better accuracy than the Azure Kinect [IOA21]. Generally, choosing one over the other is a matter of the experiment setting and the motion. Crucially, capture accuracy has not been evaluated for lifting motions, in which major knee and pelvis occlusion could exhibit unique challenges, which may be handled with varying degrees of success by the systems.

In conclusion, we decided to use the Azure Kinect over OpenPose. This is due to uncertainties of OpenPose performance with lifting motions and our requirement to have the system work well regardless of the environment. As we will perform part of our study at the INSPO motion capture laboratory, the environment should have no impact on LiftVR's effectiveness. With our choice of capture technology explained, we move on to discuss motion guidance and its terminology used throughout this thesis.

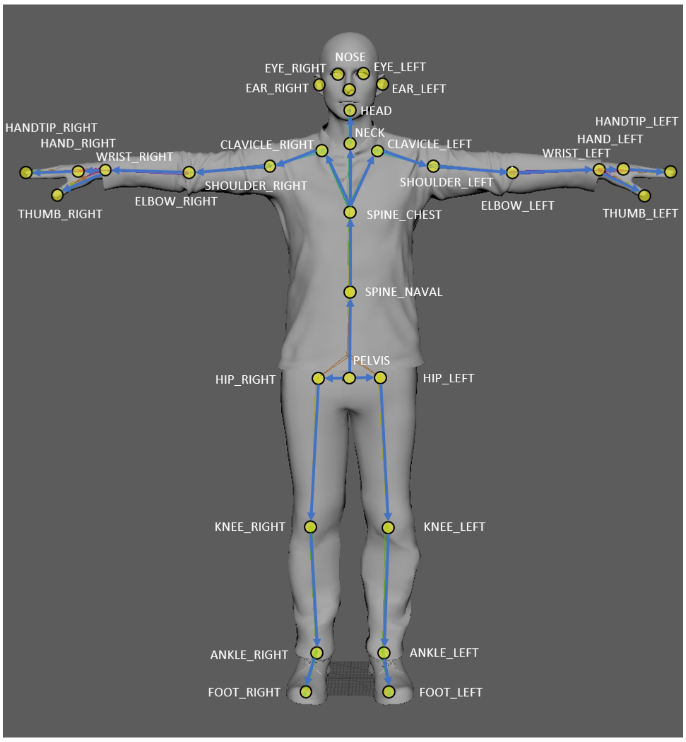

Figure 2.3: The skeleton provided by the Body Tracking SDK, with joint locations and hierarchies as seen in the figure, taken from the official Azure Kinect DK documentation. [Mic23]

2.3 Motion Guidance

The process of guiding a student to learn and possibly retain a posture or movement is regarded as motion guidance. The term guidance encompasses various stimuli provided to a student to facilitate this process [SFO+12].

Reality based examples of such stimuli would be a trainer physically positioning a students limbs to have them assume a posture or giving them active verbal feedback. As such, motion guidance is an integral part of teaching movements in sports, physiotherapy or general physical tasks. However, the question is how to provide motion guidance when these traditional stimuli are unavailable. For instance, a trainer may not be available, a gym may not be reasonably distanced or, for specialized activities, available in general. In these cases, students often turn to other forms of stimuli, such as watching videos or reading guides from home. These are less effective at teaching and do not immerse the student in an environment he expects to perform the motion. Furthermore, a video recording is unable to adjust to the anatomical circumstances of a student.

Increasing virtuality offers solutions to all of these problems, which has led to MR-based motion guidance to become increasingly popular [YAM+20]. It has been conclusively shown that students learning with MR systems improve more than with video tutorials. Virtual trainers are always available, students are immersed in their environments and systems can adjust to the individual circumstances of their users. These benefits have led to a plethora of different systems used for various recreational, educational, or rehabilitation purposes.

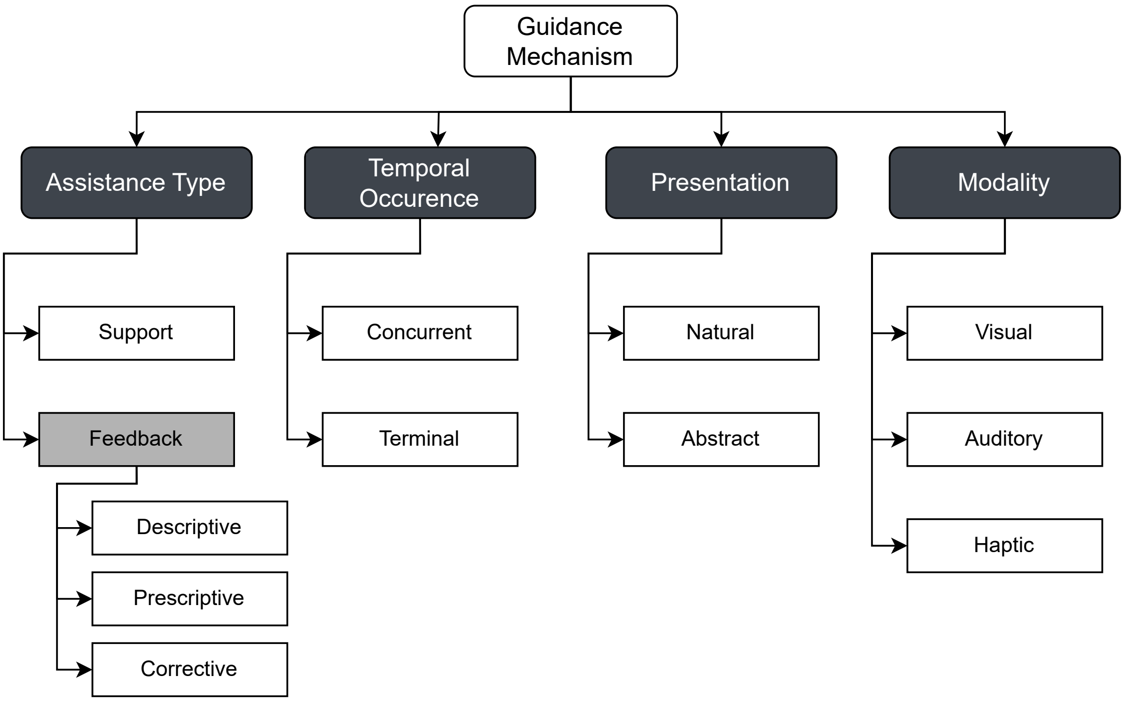

Figure 2.4: Classification scheme for guidance mechanisms.

As the comparison to video based training has been made extensively, we see no value in performing such an analysis for LiftVR. Instead, our focus lies on the stimuli used to facilitate guidance. Our goal is to evaluate different combinations of these to each other, thereby finding their effects on the learning process.

In order to compare and evaluate different guidance stimuli, it is essential to establish a systematic categorization or classification. However, the current state of research in this domain exhibits inconsistent and interchangeable terminology, making the task challenging. To address this issue, we extracted the most prevalent classifications to develop our own categorization scheme. The subsequent paragraphs provide an overview of this scheme, ensuring clarity and a common understanding of the terminology used throughout this work.

In our final subsection, we list the groups of motion guidance based on practice mode.

2.3.1 Stimuli, Mechanisms and Systems

We define the term “guidance mechanism” to describe the implementation of a single stimuli used to facilitate guidance. This substitution allows for a clearer and more explicit link between the abstract concept and its practical application. There are multiple factors which are used to classify mechanisms, a scheme of which is visualized in Figure 2.4. We distinguish them by **Assistance Type**, **Temporal Occurrence**, **Presentation** and **Modality**, all of which are explained in subsequent sections. To increase understandability and prevent excessive definitions, we intentionally excluded categories which do not occur in LiftVR from the scheme.

Continuing relevant terminology, multiple mechanisms can be used simultaneously or cooperatively, which we define in this work as a “guidance system”. Furthermore, we recognize that these systems themselves can be integrated into larger collections. Hence, we define collective systems to be devised of individual “guidance sub-systems”. Importantly, systems can be simultaneously comprised of sub-systems and individual mechanisms. This terminology allows one to distinguish and examine the interplay between different levels of guidance within a comprehensive framework.

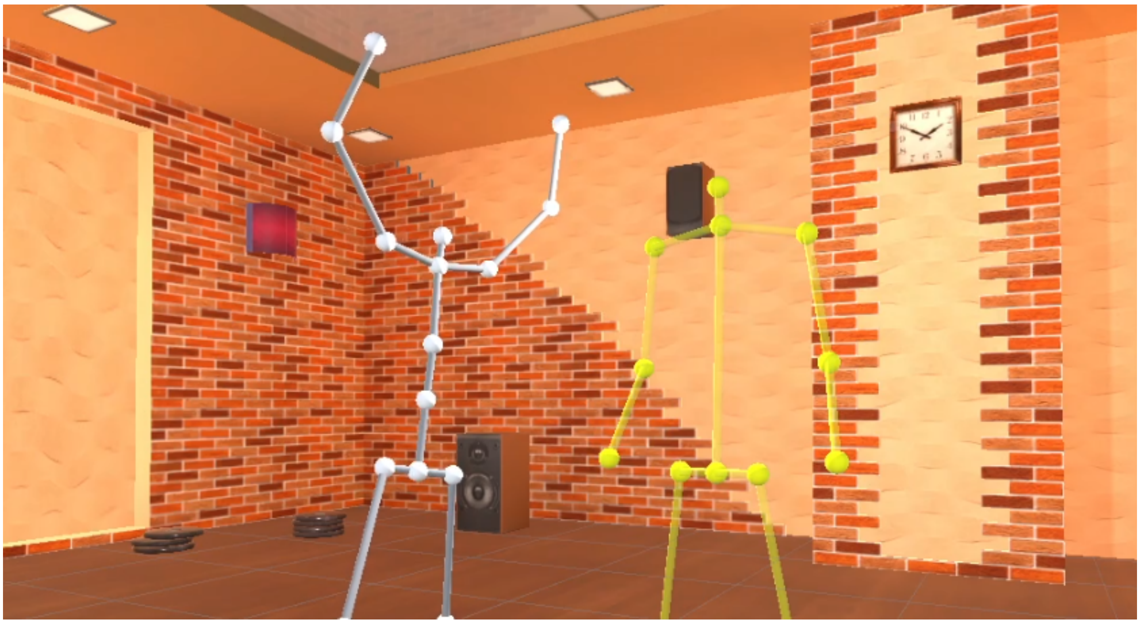

Figure 2.5: Example of feedback and support mechanisms side-by-side. This image shows two stick-and-balls skeletons from LiftVR. The white skeleton visualizes the student and provides feedback, the yellow skeleton visualizes the expert and provides support.

2.3.2 Assistance Types

Guidance mechanisms are classified based on their reliance or independence from the contemporaneous or prior performance of a student [MVKP10]. We name this category the **Assistance Type**. It consists of two different types, namely *Feedback* and *Support*.

Feedback Guidance mechanisms which fall in the category of feedback are distinguished through their dependence on the contemporaneous or prior performance of the student. They cannot assist the guidance process without it [DB10]. In consequence, feedback given concurrently during a performance is adjusted in response to the student’s actions, whilst supporting functions typically remain constant throughout the motion. Figure 2.5 shows a feedback guidance system in LiftVR. The white stick-and-ball skeleton visualizes the student’s position and built from several feedback mechanisms, which require active tracking information from the student to be displayed. If feedback merely describes or reacts to an error it is *descriptive*, if it suggests how to improve it is *prescriptive* and if it does both it is *corrective* [STK92].

Support These kinds of guidance mechanisms assist the student without any reliance on their contemporaneous or prior performance. Accordingly, a support guidance mechanism can be operational irrespective of the presence or absence of a student. Figure 2.5 shows how LiftVR implements a support guidance system in the form of a yellow stick-and-ball skeleton. This skeleton visualizes a posture of the expert, and exists regardless of the student's present state of being.

2.3.3 Temporal Occurrence

At the next level of categorization, we differentiate guidance mechanisms based on their temporal occurrence [SRRW13]. These can either be concurrent to the performance or appear when it terminates.

Concurrent Both feedback and support mechanisms can occur concurrently to the student's motion. Feedback is typically enabled or adjusted dynamically in reaction to a student's active performance, whilst support is enabled continuously [DB10]. The author has found that, while it is common that audio and haptic guidance mechanisms are dynamically enabled or disabled, this is not the case for visual mechanisms.

Terminal Terminal mechanisms require the performance of the student to cease, making the distinction only valuable for feedback mechanisms. As such, terminal feedback is mostly based on metrics calculated from the performance on a whole, whilst concurrent feedback reacts to errors as they happen. Examples for such terminal mechanisms include scores based on the student's performance [VBG13] or hints on how to improve their overall technique [NHA+20]. By definition, support mechanisms can always be displayed concurrently or terminal to the motion, making their categorization purely dependent on how the system designer restricts their use.

2.3.4 Presentation

The next categorization for mechanisms is based on how they are presented to the user. Determining precise classification of guidance mechanisms as either natural or abstract can pose challenges due to the absence of a distinct boundary between the two. Certain cases, however, offer clearer categorization. A terminal feedback mechanism that conveys a student's performance score is typically classified as abstract, while visually representing the student as a stick-and-balls skeleton is considered natural. Still, nuances emerge, as a stick-and-balls representation may still be deemed abstract when contrasted with a more detailed 3D model representation.

Abstract Abstract mechanisms are not restricted to the feedback category; they can also include support mechanisms. Yu et al. [YAM+20] provide different support mechanisms which display the expert's motion over time. One of these indicates the movement using merely a flat plane in 3D space, which we consider an abstract support mechanism. In his classification, Sigrist et al. [SRRW13] describe lines, curves, gauges, bars, or points as further examples of abstract mechanisms.

Figure 2.6: Example of abstract and natural feedback mechanisms taken from LiftVR. The white stick-and-balls skeleton is an instance of natural feedback, while the evaluation screen contains the abstract textual "score" and "worst performing joints" mechanisms.

Figure 2.6 shows the evaluation menu from LiftVR, which contains two abstract feedback mechanisms. Most pertinent are the scoring system and the accompanying textual description of joints producing low scores.

Natural Little information exists pertaining to the classification of guidance mechanisms as natural. Sigrist et al. [SRRW13] gives two examples for natural guidance systems, one which incorporates superposition of the expert and student, and one which shows them side-by-side. The authors could not find any definitions for natural guidance mechanisms covering all modalities.

To give clarity for visual guidance systems, we define that the term natural pertains to visualizations that bear resemblance to the real-world occurrences of living entities, given a moderate allowance of discrepancy. With this definition, visualizing the student as a stick-and-ball skeleton, body outline, or 3D body model of the actual student, would all be classified as natural. While the 3D body model is determined a more natural visualization than the stick-and-ball model, both are still natural feedback mechanisms.

Given this definition, both stick-and-balls skeletons shown in Figure 2.5, in addition to the skeleton shown in Figure 2.6 are classified as natural guidance mechanisms.

2.3.5 Modality

Humans perceive reality at any given time through a plethora of different senses. This intake of concurrent information at every moment makes existence a multimodal experience [SFO+12]. While the senses do not equally impact human perception [Hei92], their combination is important to fully perceive reality. As such, guidance mechanisms target different modalities and are classified based upon this.

Visual The foremost and unequivocally prevailing mechanism is one which provides visual guidance [KWM+10]. Visualizing the expert provides support to the student, while a visualization of himself offers feedback. No matter if abstract or natural, concurrent or terminal, visual guidance mechanisms provide definitive assistive power. The effect of visual guidance is, however, dependent on the medium on which the visuals are displayed upon [EHG+21; NTU+05].

Auditory With the second strongest influence on a humans perception [Hei92], auditory mechanisms hold similar assistive value. An example of an auditory feedback mechanisms is communicating joints in incorrect positions to the student during the performance [SFO+12].

Haptic While not as prevalent to the human experience as sound or visuals, haptics, which include the perception of muscles, joints, tendons, and the skin, improve motion guidance systems and are popular with students [Wen21]. As the focus of LiftVR is the comparison of visual guidance systems, the conscious decision was made not to include any haptic mechanisms. This is also to prevent the increased overhead resultant from combining multiple guidance modalities.

2.3.6 Practice Modes

In the domain of MR-based systems, motion guidance is classified into three categories based upon the type of practice [EKS+22]:

- **Posture Practice** requires students to assume positions independent of any continuous movements.
- **Path Practice** guides students along a predefined path by having them enter key postures without any time constraints.
- **Movement Practice** has students move alongside a continuously moving guidance system and the performance ceases when the guidance system reaches the final position in the motion. This restricts students time to assume positions, as the system will not wait for them.

To increase the variety of learning methods, a combination of the aforementioned practice modes can be used [AGMF13]. With this said, let us briefly investigate the type of motion LiftVR will teach.

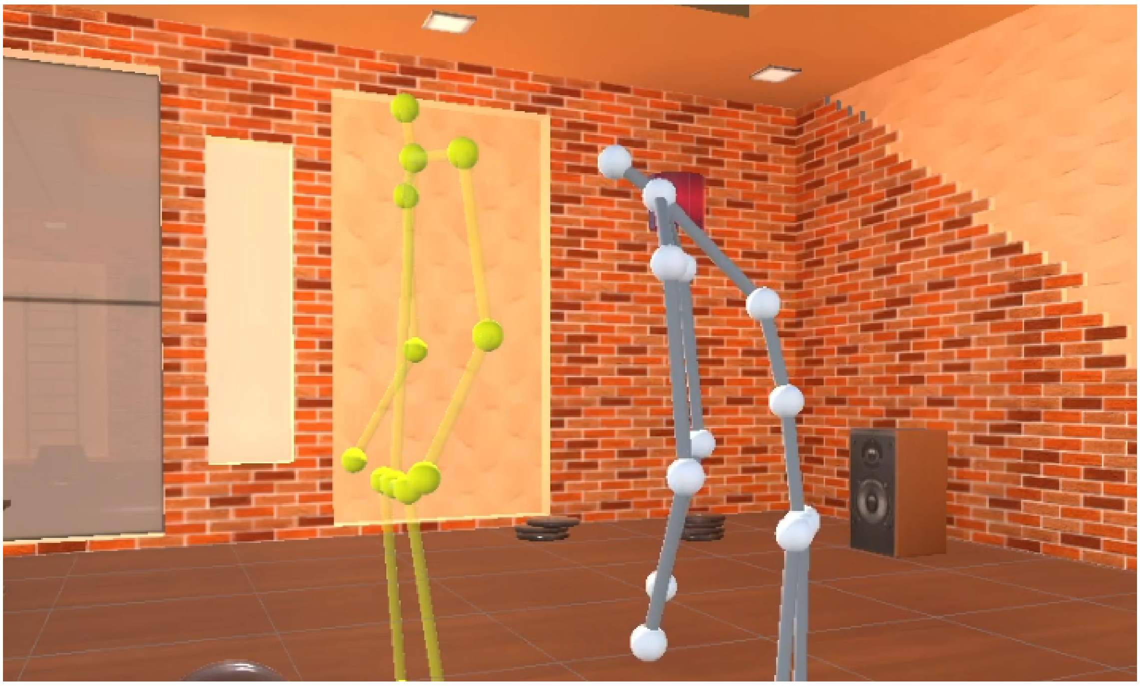

Figure 2.7: Example of a student (white stick-and-ball skeleton) assuming an incorrect posture.

2.4 Injury-Preventative lifting

Our goal is to teach a lifting motion which minimize injuries whilst maximizing muscle activation. Characteristics of an expert performing a lifting motion will include a wide stance, a natural feeling foot position, and uninhibited knees [CK07]. During the movement, he will gaze upward or forward to maintain the natural curvature of his spine, as doing so prevents spinal injuries [TMS+01]. Based on our expert recording from INSPO, further characteristics of a healthy lift are equal knee and hip flexion. A student must not bow forwards to retrieve an item without flexing his knees simultaneously. In contrast, he should not squat and lift without bowing forwards, as doing so compresses the spine. Whilst doing so, attention must be payed to maintain natural spine curvature at all times.

Figure 2.7 shows a student in LiftVR (white skeleton), next to the expert (yellow skeleton). The experts back remains perfectly straight at all times, promoting healthy lifting habits. LiftVR is capable of detecting excessive back curvature, as can be seen by the students assumes position.

Armed with all prerequisite knowledge and terminology to understand motion guidance systems, we next take a look into the current state of the art.

3 Related Work

Scientific research in the domain of VR assisted motion guidance is extensive and diverse. We lay special focus on full body motion guidance systems, as these are closely related to our own approach. However, other works we found which contribute design principles and guidelines are also included. Another limitation was to prioritize systems primarily built on visual guidance. This excludes a large portion of work and thus only a partial understanding into the overall domain of motion guidance is granted here.

In the subsequent sections we first present works deemed closely related based on our area of focus. For each work, we give a brief summary which includes a list of their guidance systems, highlight any contributions to our design and finally mention differences to our thesis. To ensure continuity, terminology used by related work has been adjusted to match our classification schemes and definitions given in Section 2.3. Subsequently we perform a similar analysis for work that is not considered closely related, but that all contributed to our design and implementation. Thereafter we present the main challenges and gaps currently identified in the domain.

Given the extensive nature of the domain, we discovered a substantial amount of literature beyond what was ultimately reviewed. In addition, some works with important findings were not found until after LiftVR's was implemented, due to them being released after development was complete. For other cases, we were unable to find them until after our literary review had ended due to inconsistent terminology in the domain. We explicitly mention works that fall into these categories in their respective paragraphs. Important findings are included as implications for future work on LiftVR in Section 7.2 on page 94.

3.1 Full Body Motion Guidance

Superimposed Motions Improve Student Performance

In their work "*Superimposed Skilled Performance in a Virtual Mirror Improves Motor Performance and Cognitive Representation of a Full Body Motor Action*" Hülsmann et al. [HFS+19] implement a full body motion guidance system using CAVE and a body tracking suit. They created two distinct concurrent subsystems: one that presents the student's virtually using an opaque avatar, and the other that visualizes the expert as a transparent avatar. With these they investigated the effects of superimposing a students performance onto the active playback of an experts recording.

Contribution It was found that superimposing the recording of an expert onto a student's virtual avatar improves their performance, regardless of the motion or perspective in play. Furthermore, a side view yields better learning retention than a front view. Although both perspectives demonstrate comparable learning rates, a side view minimizes errors related to the back-front axis, while a front view reduces errors associated with the center of mass. This work could not be found during literary review due its us of "motor action" instead of "motion" or "movement" and various other substitutions in terminology. Nonetheless, during LiftVR's design, we independently concluded superimposing motions is the best approach.

Differences The system requires a CAVE environment, we use VR provided by an HMD. Furthermore, a body tracking suit is required. In terms of guidance systems, besides the student visualization, no feedback is implemented. The effects of different approaches to the expert support system's design are not investigated.

Effects of Perspectives on Motion Guidance

Elsayed et al. [EKS+22] evaluate the effects of three different perspectives for motion guidance in their work"*Understanding Perspectives for Single- and Multi-Limb Movement Guidance in Virtual 3D Environments*". These being first-person perspective (1PP), over-the-shoulder third-person perspective (3PP) and seeing side views with over-the-shoulder simultaneously (multi-3PP). Using a Microsoft Kinect v2 (Azure Kinect's Predecessor) they implement a feedback system visualizing the student as a skeleton with cube joints and rectangular bones. The expert is visualized in similar fashion. During the performance, multiple postures of the expert skeleton are visible to guide the student along a movement path. For concurrent descriptive feedback, they use color indicators on the student skeleton and vibration bands, which would trigger when joints exceeded 15cm of their intended positions. In their evaluation, interaction effects between concurrent feedback, perspectives, and different motion parameters are investigated. They find that multi-3PP perform the best and concurrent feedback has no significant impact on performance.

Contribution 1PP is outperformed by 3PP, which in turn is outperformed by multi-3PP. This work was discovered after critical implementations of our system were already complete, which is why multi-3PP is not available in LiftVR. Nonetheless, as a result of these findings we removed our 1PP mode and removed perspectives from the independent variables in our evaluation.

Differences Whilst one concurrent feedback system (color indicators) is investigated for effects, no terminal feedback was implemented. Similarly, only one concurrent support system representing the expert was implemented, as the effect of different implementations was not the focus of their work.

Influences of Humanoid Visualizations on Motion Guidance

In their work "*CameraReady: Assessing the Influence of Display Types and Visualizations on Posture Guidance*", Elsayed et al. [EHG+21] investigated influences of different displays and humanoid visualizations on motion guidance. They implemented three humanoid visualization, namely stick-and-ball skeleton, silhouette, and a 3D virtual avatar. These visualizations were used both for the student feedback and expert support systems. Interaction effects between these visualizations and display types were evaluated, finding that small displays can be used effectively for posture motion guidance, though increasing size lowers error. Whilst students prefer less abstract visualizations such as silhouettes or the 3D virtual avatar, they cause an increase in task completion time for no other benefit.

Contribution For LiftVR we need to consider a humanoid visualization for the student feedback system and our expert support system intended to mimic the current state-of-the-art. The findings of this work led us to dismiss silhouettes or 3D virtual avatars as possible designs due to their negative effect on performance.

Differences The work implements no concurrent or terminal feedback systems. Whilst the effects of different expert support system visualizations were investigated, all of them were humanoid and not tailored to specific motions.

Virtual Reality Dance Training System

Chan et al. [CLTK11] developed a dance training system in their work "*A Virtual Reality Dance Training System Using Motion Capture Technology*". A concurrent feedback system visualizes students wearing a body tracking suit as a stick-man built from cylinders. The expert is visualized as a virtual avatar. Whilst the systems are not superimposed, the stick-man cylinders will change color if the body segment position is incorrect. We categorize this form of descriptive feedback as a color indicator. Two additional forms of terminal feedback are implemented, one which shows textual scores for each joint, the other a slow motion replay of the last performance with more distinct color indication. In their evaluation, they investigated the effects of different performance evaluation functions, the learning outcome of students and a comparison with video based self-learning.

Contribution The main contribution to the design of LiftVR comes from the investigation of performance evaluation functions. While joint positions, velocities, and angles can be used for measuring motion differences and the classification of similar and dissimilar motions, joint positions hold the most discriminatory power. As such, we opted to use joint positions for our own evaluation measurements.

Differences Guidance systems are visualized on a large screen, not a HMD, and participant performances are not superimposed on the expert recording. Furthermore, participants are required to wear body tracking suits. Different approaches to expert visualizations are not evaluated and the effect of terminal feedback is not investigated.

Learning Taichi Motions in VR

In their work "*Immersive and Collaborative Taichi Motion Learning in Various VR Environments*", He et al. [HCC+17] implement ImmerTai, a Taichi learning environment which utilizes VR and full body tracking. Using a Kinect camera they visualize 3D virtual avatars of the expert and student side by side. For their evaluation, the system is deployed in three different environments, with students either using a screen, a HMD, or a CAVE system. They found that whilst participants prefer HMD's, performance was best in the CAVE system. Students reported difficulties with the field of view (FOV) and HMD cable.

Contribution The findings of this work did not contribute to the design of LiftVR. Whilst choosing an HMD would have been preferable, due to hardware limitations, this was not feasible.

Differences The system implements no concurrent or terminal feedback besides the student's 3D virtual avatar. The expert is also visualized as such an avatar, with no evaluation of the effects of different implementations.

Motion Modeling, Performance, and Analysis System

"*MotionMA: Motion Modelling and Analysis by Demonstration*" is a system by Velloso et al. [VBG13] which allows users to create their own demonstrations, perform them, and receive feedback. Demonstration analysis detects repetitions and users can further modify these through a user interface. This includes a selection of the key joints monitored during any performances. For guidance systems, concurrent feedback visualizing the student is given in the form of a stick-man. If bones are positioned incorrectly, color indicator feedback is activated. Abstract concurrent feedback is given, as descriptive systems show the students current arm angles using a gauge. Two prescriptive systems are also implemented, one textual telling the user to slow down or speed up, the other using arrows to give corrections when arms are not in position. They analyze this motion modeling system and system usability in their evaluation. Therefrom they conclude the extracted motion model to be accurate and users satisfaction with MotionMA.

Contribution Motion extraction and modeling from demonstrations is a strong technology, however, it is also difficult to implement. Though LiftVR would certainly benefit from such an implementation, the conscious decision was made not to consider this in our design. Nonetheless we mention this technology in our implications for future features of LiftVR.

What did impact our design is the fact that whilst students enjoyed the abstract concurrent feedback systems, they found them to be visually overwhelming. As such, any abstract concurrent systems are to be designed not to overpower the users visual senses.

Differences The system does not superimpose the student and expert visualizations. Furthermore they are shown on a large screen, not an HMD. Whilst various additional abstract concurrent feedback systems were implemented, their effect on overall performance was not investigated. Finally, additional markers are required to be placed upon the participants body.

Motion Authoring and Training Guidance System

"*YouMove: Enhancing Movement Training with an Augmented Reality Mirror*" by Anderson et al. [AGMF13], also allows users to author motions, however, unlike the previous work, an editing interface requires authors to specify key frames and important parameters. Further differences include students not performing before a screen but rather a half-silvered mirror (HSM). During performance, the feedback system visualizing the student uses a stick-man representation. The support system showing the expert is a video recording.

Two practice modes are implemented, namely posture and movement practice. During posture practice, an additional concurrent feedback system displays color indicators on the student skeleton. In movement practice, ribbons are attached to joints give prescriptive feedback to cue upcoming movements. Finally, terminal feedback systems allow students to go through their performance frame-by-frame whilst a stick-man visualization of the expert is superimposed on their recording. To utilize gamification [HKS14], a score calculated from average joint euclidean distance is given. In their evaluation, they compare the performance of students with the system to others relying on videos only. Users first see a demonstration of the motion and must replicate it for the baseline measurement. Subsequently, posture training is carried out, followed by movement training. To test short-term retention, students must perform the motion one final time without any guidance systems.

Contribution Systems from this work, such as the superimposed replay terminal feedback and the scoring system for gamification, contributed to our design. Furthermore, the development of our scaling algorithm is based on YouMove's implementation. Most importantly, we base the training procedures in our study design upon the study carried out for YouMove.

Differences Does not use a HMD, but rather a HSM. The expert support system is merely a video recording. Effects of different concurrent visualizations are not investigated. Whilst terminal feedback systems are implemented in YouMove, their overall effect on student performance was not evaluated.

Two-Party VR Posture Motion Guidance System

In their work "*Onebody: Remote Posture Guidance System using First Person View in Virtual Environment*", Hoang et al. [HRVT16] implemented a system in which an expert and student are tracked simultaneously. Both are visualized using stick-and-ball skeletons and are joined at the hip during performance. Color indicator concurrent feedback is used to telegraph when student joints do not overlap the expert. The expert skeleton is scaled to match the student using normalized vectors. In their evaluation, they investigated the systems posture accuracy and user preference. To do this, they had subjects train with three different systems, finding that whilst users of Onebody produced lower errors than standalone VR motion guidance, task completion times increased significantly.

Contribution As cooperative VR was not an option for LiftVR, the findings of this paper did not contribute to our design.

Differences Implements no terminal feedback systems or evaluations of different approaches to expert and student visualizations.

Real-Time AR Rowing Coach

"*ARrow: A Real-Time AR Rowing Coach*" by Iannucci et al. [ICA+23] is an highly accessible system which uses a tablet to provide real-time feedback on rowing motions. As the main visual feedback system, students are implemented as stick-men. For support, experts are superimposed using a similar visualization. These are aided by multiple additional concurrent and terminal feedback systems. Text is used to describe mistakes during the performance. Current stroke rate and average drive time is calculated and displayed. A line is drawn to highlight wrist and elbow paths, gradually fading over time.

For their evaluations, they measured the accuracy of their feedback systems measurements and gathered qualitative user feedback. Their system achieved high accuracy and students were pleased with the feedback provided, though they wished for more metrics showing change over time.

Contribution As many of the feedback systems implemented in ARRow are specific to the motion, they do not affect LiftVR's design. An exception to this are the lines drawn to highlight wrist and elbow paths. Motivated by this, we developed our own terminal feedback joint path system.

Differences As no evaluation of the feedback systems took place, their effect on guidance is unknown. In general, systems were not evaluated against each other, giving no indication to their importance in the overall implementation. Furthermore, the system does not use VR.

Ballet Arm Posture Guidance System

As our final related work we present "*BalletVR: a Virtual Reality System for Ballet Arm Positions Training*" by Barioni et al. [BCAT19]. Student positional feedback is given through a 3D virtual avatar. The data is provided by a Kinect camera (not to be confused with the Azure Kinect). The expert support system is the same avatar, though slightly transparent and colorless. For ballet training, it is important for students to hold positions for several seconds, and most importantly remember the postures themselves. As such, the expert support system is only displayed in response to the student not assuming the posture for twenty seconds, ergo, when he has forgotten it. Once the posture is assumed, a concurrent abstract feedback system shows a visual countdown to indicate how long the posture must be maintained. The system was then evaluated based on qualitative feedback from students.

Contribution We derived no direct design implications from this work for LiftVR.

Differences No terminal feedback systems were implemented and no quantitative analysis of student performance was performed. Furthermore, concurrent feedback systems were limited to 3D virtual avatars.

3.2 Implications for Motion Guidance Design

Design Implications for Motion Guidance in Mixed Reality

In "*Perspective Matters: Design Implications for Motion Guidance in Mixed Reality*", Yu et al. [YAM+20] evaluate the effects of two different expert support visualizations and two perspectives on the performances of participants. A feedback system responsible for visualizing the students arm did so with a stick-and-ball skeleton. Two different support systems, tasked with displaying the recorded motion to the user, were implemented. The first implementation draws a plane through 3D space which corresponds to the arms movement in the recording. This plane is then attached to the students shoulder joint, superimposing the two. The second implementation showed multiple stick-and-ball arm postures with increasing transparency (called *GhostArm*), intended to guide the student along a movement path. Akin to the first system, these are also attached to the students shoulder. Finally, another concurrent feedback system draws lines between the student's stick-and-ball joints and their currently intended positions (called *Rubber Bands*). In their evaluation, the effects of these separate visualizations and two perspectives was evaluated. They found the less abstract *GhostArm* had a larger positive effect on student performance.

Contribution As the implementation of ghost limbs with rubber bands outperformed the other systems, the design of our state-of-the-art support guidance system uses these findings.

Differences As all the works in this category, tracking and motions were not full body. Additionally, no terminal feedback systems were implemented.

AR Table Tennis Training System

Nabil et al. [NHA+20] developed a motion guidance system for table tennis in their work "*Usability Study of a comprehensive table tennis AR-based training system with the focus on players' strokes*". The main focus was on a classification system to generate descriptive feedback. As such, no concurrent motion guidance systems are implemented, be it feedback or support. Students were trained exclusively with terminal feedback. After a performance, recordings are smoothed with a Kalman filter [Wel97]. In their evaluation, they investigated the accuracy of their classification system and the effects on student errors. Classification was deemed highly accurate and students improved significantly with terminal feedback (86% after three sessions).

Contribution Due to the impact of terminal feedback shown in the paper, we hypothesize that it has an high effect when employed interactively with concurrent guidance systems. Furthermore, when implementing smoothing in LiftVR, we used a Kalman filter based on this work utilizing it successfully. Finally, classification could help in the generation of corrective feedback, however, similar to authoring , the implementation of such a system was deemed out-of-scope for LiftVR. We do mention it as potential future improvement to LiftVR.

Differences Being in this category, no full body tracking or motion was used. Has no form of concurrent guidance systems. As such, the effects of their terminal guidance systems in combination with concurrent systems is not investigated.

Multimodal Feedback Motion Guidance System in VR

In "*Developing a Multimodal Feedback Motion Guidance System in VR for People with Motion Disabilities*", Wennrich [Wen21] interviewed physiotherapists and their patients to find requirements for VR motion guidance using multiple modalities. Based on their findings, they implemented visual and haptic concurrent guidance systems. In their evaluation, they presented the developed systems to users and asked them which they would perform. Participants prefer multiple guidance modalities, with haptics being especially popular. They wished the environment to look like a physiotherapists office or gym.

Contribution From this work, we derived the requirement of visualizing a gym environment to students to increase immersion.

Differences No full body guidance. Whilst visual systems in the form of stick-and-ball limbs were implemented, effects of different implementations were not investigated. No terminal feedback systems were developed.

Influences of Processing Modes on Azure Kinect Tracking

Whilst not specifically aimed towards the domain of motion guidance, Büker et al. [BQH+23] investigate the effects of processing modes on the Azure Kinect's tracking results. In their work "*How the Processing Mode Influences Azure Kinect Body Tracking Results*", they find that the *DirectML* processing mode should be used instead of *CUDA* or *TensorRT*. As this paper was released after our literary review, we discovered these findings only after our study was completed.

Contribution The *DirectML* processing mode is to be used for future use of LiftVR in combination with the Azure Kinect.

Effects of Feedback in Physical Education

The work by Silverman et al. [STK92] is targeted towards the domain of physical education, but bears important findings for the design of motion guidance systems, as these can be seen as “digital” teachers. “*Teacher Feedback and Achievement in Physical Education: Interaction with Student Practice*” firstly provides us with the categorization of feedback into *descriptive*, *prescriptive*, *corrective*, and many more. They find that providing one form of feedback is enough to have an effect on performance, however, the true power of feedback stems from combining its categories. As such, students receiving descriptive, prescriptive, and corrective feedback with an encouraging attitude perform significantly better then those with singular feedback.

Contribution This work bears implications on design decisions of motion guidance systems, such as not relying on descriptive feedback, as the combination with prescriptive feedback. We found this work after our study had completed in response to our findings. As such, design implications from this work are included as future improvements to LiftVR in Section 7.2 on page 94.

3.3 Challenges and Gaps

Based on our literature review, we identified key gaps in the current domain of full body motion guidance. The current dominant approach of full-body motion guidance is to superimpose a visualization of the student onto a humanoid entity. Examples of this include skeletons, outlines, videos, or full body models. Nonetheless, they all assume human form or characteristics. Only Yu et al. [YAM+20] provide a non-humanoid visualization of the expert support system. None of the works deemed closely related implemented anything other than humanoid entities. This leaves a gap in the domain of motion guidance, as the effects of different expert visualizations have not been broadly investigated, especially for full body systems.

Furthermore, we find no evaluations of the effects of terminal feedback systems utilized alongside concurrent guidance. Elsayed et al. [EKS+22] discovered no effects of concurrent feedback, leading us to believe that such effects will not exist for terminal feedback either. This conflicts the results of Nabil et al. [NHA+20], in which terminal feedback had a significant impact, though no concurrent guidance was present. Thereby, no clear investigation of these effects exists.

Finally, whilst we found motion authoring systems and trainers for ballet, Taichi, rowing, and various other sports, no system for lifting motions exist. Additionally, none of the aforementioned systems include interactions between the student and real weights. Closest to this is the work by Iannucci et al. [ICA+23], although the choice of AR instead of VR makes the rowing machine easy to integrate. Students do not wear HMD’s and thus have no difficulty interacting with the real world device. The implementation of such a system would thereby be the first of its kind.

To address these challenges and provide the first system for lifting motions with weight interactions, we developed LiftVR. Its implementation is subsequently described.

4 Design and Implementation

This chapter contains detailed information on the design and implementation of LiftVR. Following a brief overview of the system, data acquisition and processing techniques for the student and expert are explained. The calibration technique used by LiftVR is discussed next. Thereafter, the various guidance mechanisms, subsystems and systems which constitute LiftVR are presented. We commence with the concurrent systems followed by the terminal feedback implementations. Concluding, a brief explanation of the user interface (UI) and the software architecture is given.

4.1 Overview

LiftVR is a VR motion guidance system teaching students a lifting technique performed by an expert in an injury-preventative fashion. It accomplishes this by immersing students in a fitness studio environment and providing multiple options to individually tailor their learning experience. LiftVR was implemented in Unity 2021.3.22f1 using the standard render pipeline. For VR integration we chose SteamVR [Val23]. It is developed by Valve and available for free through the Unity asset store.

Environment The immersive environment, depicted in Figure 4.1, showcases the student, visualized as a white stick-and-ball skeletal structure, assuming a pose. We procured a fitness studio environment from the Unity asset store [IG-17], though all other assets were created exclusively by the author. The darkly hued floating planes form the UI, with the left plane serving as the main menu and the right plane functioning as the evaluation menu. Configuration of the various concurrent guidance systems, practice modes and options is enabled through the main menu. Access to the terminal systems is granted to the students through the evaluation menu.

Guidance Systems LiftVR employs 18 guidance mechanisms, some combined to form ten subsystems with a total of five main guidance systems. For each of the five main systems, we give a brief introduction:

- *Ghost-Skeleton* is the first of two concurrent guidance systems and our state of the art implementation. Students must follow an expert skeleton that has been adorned with spherical zones in which the student must position their joints.
- *Zone* is the second concurrent guidance system and our challenge to the currently dominant natural humanoid visualization. *Zone* omits any display of an expert skeleton and instead focuses solely on the presentation of capsule-shaped zones. These serve as abstract references for the placement of the student's corresponding joints. Furthermore, it utilizes the symmetrical nature of lifting motions to pair multiple joints into singular zones, reducing visual clutter.

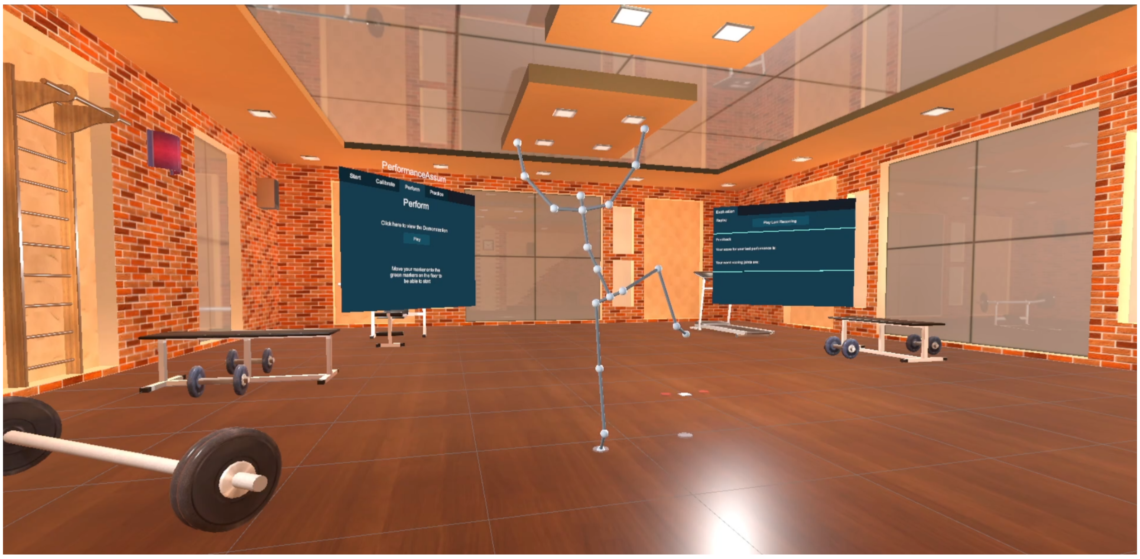

Figure 4.1: The immersive fitness studio environment of LiftVR. The two planes hovering in space are the main menu (left) and the evaluation menu (right). Centrally, the student can be seen visualized as a white stick-and-ball skeleton.

- *Foot Location Manager* helps students position their feet correctly before their performances. This is done to prevent novices unfamiliar to the motion from injuring themselves, as a wide, natural feeling stance is one the requirements for injury-preventative lifting.
- *Evaluation Manager* provides students with multiple descriptive terminal feedback systems, including a replay of their last performance, textual feedback on their score and the joints with poor performance, as well as the ability to inspect the paths traveled by individual joints. Together, these enable the student to identify his mistakes and correct them.
- *Heads-Up Display* serves as the primary means for providing students access to abstract concurrent feedback systems. Displaying this abstract feedback on menus during the performance could lead to frequent head rotation by the student, diverting their attention from the natural concurrent systems. To address this issue, we have attached the abstract feedback systems to a Heads-Up Display (HUD) that continuously remains within the user's FOV. The primary purpose of these systems is to ensure students have clear indication of when a performance starts and ends. This is to prepare them mentally, mitigating errors resultant from them being surprised by the concurrent guidance systems appearing. To facilitate this, a textual countdown is initiated, starting from three. During the motion, students observe a prominent red dot, signifying that the motion is active and being recorded. Additionally, a back angle feedback system continuously visualizes the current curvature of the back and alerts students when their back is not straight.

Detailed explanations and figures for the systems and their respective subsystems can be found in Section 4.5. Before this, we must address the various processing steps data requires to be usable by LiftVR's system, beginning with student data received from the tracker.

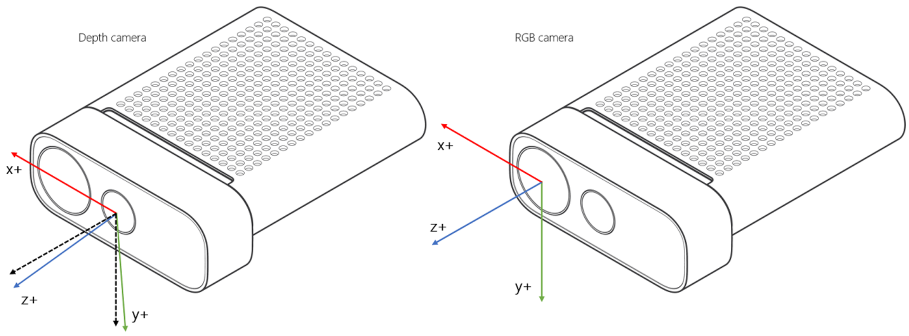

Figure 4.2: Depth and color camera coordinate systems of the Azure Kinect. Taken from the official Azure Kinect DK documentation [Mic23].

4.2 Student Data

To enable smooth and real-time visualization of the student, LiftVR goes through a multifaceted process to clean motion capture data. Three major challenges resulting from the Azure Kinect camera's hardware specifications are overcome.

Depth Camera Tilt Firstly, the depth camera of the Azure Kinect is tilted downwards by six degrees, as seen in Figure 4.2. Consequently, if the Azure Kinect is appropriately positioned in a horizontal manner, moving away from it will induce an upward shift of the tracker skeleton. Similarly, moving toward it will shift the skeleton into the ground. Furthermore, the skeleton will appear tilted forwards by six degrees.

We address this by placing the Azure Kinect on a level surface, which leads to an inclination of the depth camera of approximately eight degrees, mostly mitigating the aforementioned effects. This process is quick and does not necessitate the use of precise angular equipment. As a slight inclination still remains, one can not entirely prevent the student from rising and sinking into the floor, breaking immersion. This is troublesome, as LiftVR encourages the student to walk around the playable space rather than remain in one fixed position. As a countermeasure, LiftVR determines the foot nearest to the floor, calculates its distance to the ground and applies this offset to all joints. Regardless of the motion performed or the camera's distance, the student will always remain perfectly grounded. We name this system *FloorFix*. One drawback to this system is that jumping is not supported by LiftVR, as the student always remains on the ground. Nevertheless, as LiftVR does not teach jumping motions, this countermeasure can be employed without breaking immersion.

Figure 4.3 shows three depictions of the student skeleton with the user standing at various distances to the camera. In Figure 4.3a, the student is standing one meter away from an Azure Kinect placed perfectly horizontally at a height of precisely one meter. After increasing the distance to two meters, Figure 4.3b shows the skeleton lifted high above the ground, breaking immersion. Utilizing our aforementioned countermeasures, the skeleton, as seen in Figure 4.3c, appears firmly on the ground at the same two meter distance.

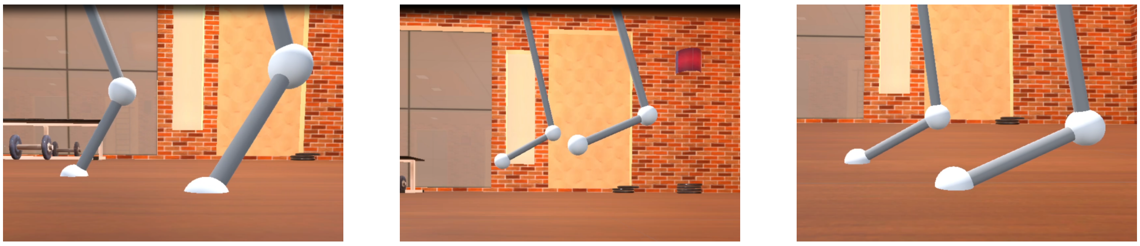

(a) One meter distance. **(b)** Two meter distance. **(c)** LiftVR FloorFix.

Figure 4.3: a) and b) demonstrate the Azure Kinect's upward shift whilst moving away from the camera. In c), the student has the same distance as in **b)** but remains on the grounds due to LiftVR's FloorFix.

LiftVR's countermeasures offer the additional advantage of facilitating swift height adjustments of the Azure Kinect without the need for re-calibration or resets. Ordinarily, lowering the Azure Kinect would cause the entire skeleton to shift downwards, as the tracker assumes a fixed camera height of one meter. However, LiftVR ensures that the skeleton remains anchored to the ground, regardless of changes in the Azure Kinect's height.

Interpolation and Smoothing Two further challenges imposed by the Azure Kinect necessitate intervention to enable immersive viewing of tracking data in VR. Firstly, the tracker operates by queuing captured sensor data for processing by the underlying machine learning model. Since depth information is crucial for generating results, the tracker's speed is constrained by the frame rate of the depth camera, which is limited to 30 FPS. Consequently, the update frequency for joint positions averages around 33 milliseconds. In comparison, LiftVR operates at 90 FPS, resulting in a frame render time of approximately 11 milliseconds. This means LiftVR updates significantly faster than data can be provided by the tracker. As a result, directly visualizing the tracking results leads to noticeable lag and stuttering.

Secondly, due to the reduced accuracy compared to marker motion capture systems, tracker data is noisy and sensitive to changes in the environment. LiftVR's intention is for students to view the skeleton as if it were their own body, such that they receive close up views of their joints. Directly visualizing the skeleton with tracker data causes immense stuttering, breaking immersion. Furthermore, partial occlusions when squatting can break joint confidence, which results in the skeleton twitching and rotating violently. This also happens when looking straight downward with an HMD or generally wearing baggy or unicolor clothing.

To mitigate these effects, LiftVR uses a mixture of smoothing, prediction, and interpolation to provide a smooth and immersive experience. The challenge hereby is to prevent the aforementioned abnormalities without introducing intolerable amounts of lag. After applying FloorFix, a smoothed position is retrieved from a Kalman filter, developed specifically for Unity by a third party [Fos23]. It works by generating predictions from prior information and outputting an estimate of the current state after receiving new measurements. The resulting estimate is then fed forward and smoothed again by an exponential weighted moving average (EWMA) algorithm [Hun86].

Whilst the Kalman filter and EWMA reduce noise at the cost of tolerable lag, the skeleton still stuttered due to the conflicting update rate of the tracker and LiftVR's FPS. To solve issue we added interpolation to LiftVR. In the context of frames, when three frames were rendered by the system, only one of them could contain new information from the tracker. To make use of these unused frames, we fill a frame buffer with three frames and use these positions to form a three-control-point Bézier curve for each joint [HN82]. We then calculate and store two positions on the curve, namely the quarter and half-way points. In the same frame each joint is then immediately moved to its Bézier quarter position. During the next two frames, LiftVR checks if tracker data is available and in its absence will move joints to their Bézier center positions. This improves the intermittent skeleton update time from 33 to 22 milliseconds. As a result, LiftVR effectively doubles the framerate from 30 FPS to 60 FPS, noticeably reducing skeleton stutter, whilst contributing only 90 milliseconds to general lag. Do to the nature of Bézier curves being smooth and continuous, this interpolation method provides smoothing simultaneously. Intentionally, no processing is done during one of the two unused frames. This is due to the tracker occasionally providing tracking data in short bursts of approximately 22 milliseconds. If this should occur, the third frame will immediately update joints using tracker information, rather than perform an interpolation step. Algorithm 4.1 shows the full student skeleton update algorithm, including FloorFix and all smoothing steps. Note that LiftVR will not call *UpdateStudentSkeleton* when the tracker is null and the interpolation step has already been performed.

Algorithm 4.1 LiftVR Smoothing and Interpolation

```
procedure UpdateStudentSkeleton(student, tracker)
    if tracker = null then
        for all joint ∈ student do
            SetJointToPosition(joint, curveCenterPosition)
        end for
    else if tracker ≠ null then
        heightOffset ← CalcFloorOffset(tracker)
        for all joint ∈ student do
            trackingJoint ← GetTrackingJointForJoint(tracker, joint)
            AddHeightOffset(trackingJoint, heightOffset)
            newPosition ← KalmanFilter(trackingJoint)
            newPosition ← ExponentialWeightedAverage(newPosition)
            if !IsBezierBufferFilledForJoint(joint) then
                InitializeBezierBuffer(newPosition)
            else if IsBezierBufferFilledForJoint(joint) then
                bezierPositionA ← CurrentPosition(joint)
                bezierPositionB ← bezierPositionC
                bezierPositionC ← newPosition
                curveQuarterPosition ← CalcBezierPoint(joint, 0.25)
                curveCenterPosition ← CalcBezierPoint(joint, 0.5)
                SetJointToPosition(joint, curveQuarterPosition)
            end if
        end for
    end if
end procedure
```

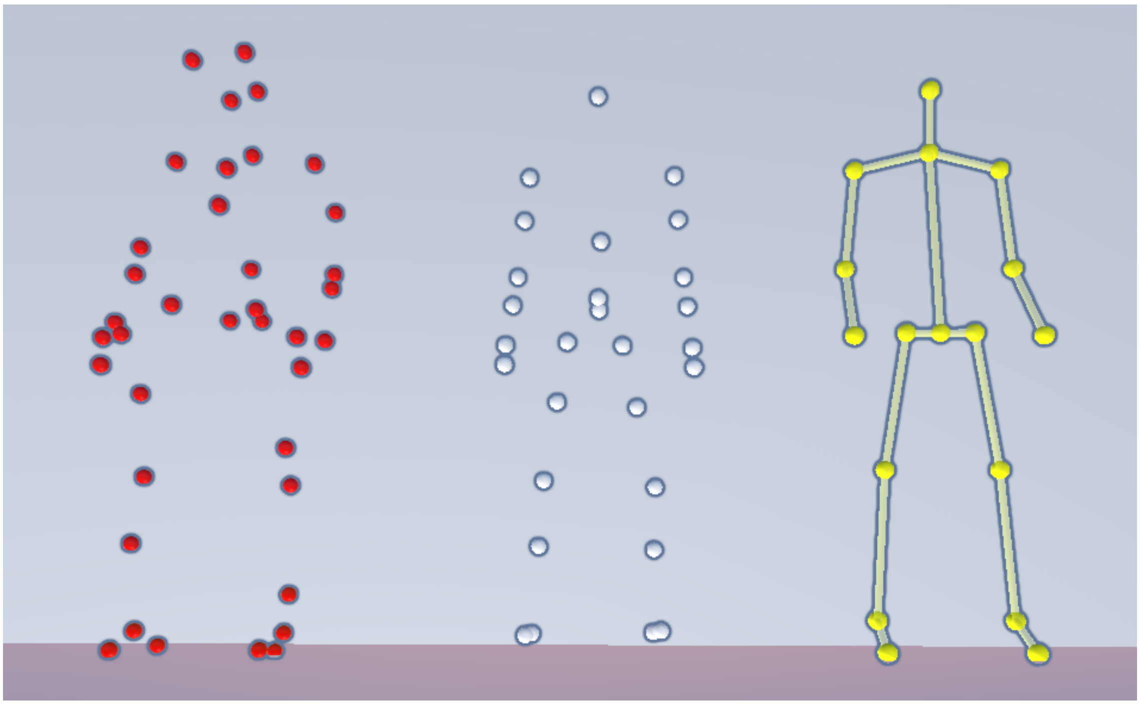

Figure 4.4: The red spheres visualize Vicon markers placed on the body as seen in Figure 2.2 on page 22. The white spheres are internal markers automatically calculated by Vicon. These two are processed to construct the expert's Azure Kinect skeleton (yellow).

4.3 Expert Data

Motion guidance can not exist without motion, and LiftVR obtains this motion from the recording of INSPO's expert performing injury-preventative lifts. The expert's motion was captured by a Vicon motion capture system, with the expert having attached trackers to his body as explained in Section 2.2.1 on page 21. However, LiftVR visualizes students using the Azure Kinect's body model as seen in Figure 2.3 on page 23. This poses a fundamental incompatibility. Consequently, transformations of the models to a unified system must be made. As the Azure Kinect's body is simpler, the Vicon recording is transformed into it.

Joint Mapping Most joint types can be mapped directly from Vicon's internal markers to the Azure Kinect's system. Table 4.1 shows how all but two joints are be mapped using a single correspondent internal Vicon marker. Azure Kinect joints not listed are either not used by the student visualization, the expert visualization, or both, and are therefore omitted. As seen in Figure 4.4, the white internal Vicon markers lack neck and hip information. Whilst the hip is easily reconstructed by taking the average positions of the left and right hip markers, the neck joint must be constructed using the external Vicon markers. Specifically, the average positions of the *C7* and *CLAV* markers are taken to approximate a neck joint. The expert intentionally does not have the Azure Kinect's *Spine Naval* and Spine Chest joints, as we found no reliable way to map these. However, since the expert should always have a perfectly straight back, we deemed these not necessary.

Vicon Markers	Azure Kinect Joint
RightFootCOM	FootRight
LeftFootCOM	FootLeft
RAJC	AnkleRight
LAJC	AnkleLeft
RKJC	KneeRight
LKJC	KneeLeft
RHJC	HipRight
LHJC	HipLeft
RHJC, LHJC	Pelvis
C7, CLAV	Neck
HeadCOM	Head
RSJC	ShoulderRight
LSJC	ShoulderLeft
REJC	ElbowRight
LEJC	ElbowLeft
RWJC	WristRight
LWJC	WristLeft

Table 4.1: Vicon to Azure Kinect Mapping

Motion Cleaning Certain imperfections and irregularities in INSPO's recording exist, as inherent limitations of human capabilities prevent achieving absolute perfection. The expert's left foot exhibits a slight backward offset, resulting in a left-veering effect on their body. Additionally, a flawless lifting motion would involve consistently keeping the back, neck, and pelvis on planes perpendicular to the bar. Finally, an abnormality of unknown cause has the expert's left knee bend backwards, causing it to be an unreachable position. As we import the experts recording into LiftVR, we edit the motion to fix the aforementioned issues and create a lift so flawless no human could have performed it. We do this by disposing of the expert's left body half and instead mirroring the right half, creating perfect symmetry in arm and leg movement. Furthermore, LiftVR locks the pelvis, back, neck, and head joints to a plane perfectly perpendicular to the bar. As a result, these joints only move up and down with no swaying movements. Finally, the feet and ankles are frozen to their positions in the first frame of the motion.

Scaling Humans are diverse in their physical appearances, varying in height, body-weight, and appendage lengths. To allow the expert to guide students regardless of their physical characteristics, LiftVR employs a scaling algorithm based on Anderson et al. [AGMF13]. During calibration, the bone lengths in the students body are estimated with tracking data. We then iterate through all recording frames and rebuild them based on these lengths. Rebuilding begins from the left foot and iterates upwards throughout the body. Consequently, the experts feet remain on the ground after scaling, however, there is a horizontal fixed-value shift. As the feet positions remain consistent throughout the motion, we can correct this shift by simply moving the entire expert in the opposite direction. This allows us to fix the shift without having to correct the positions of each individual joint.

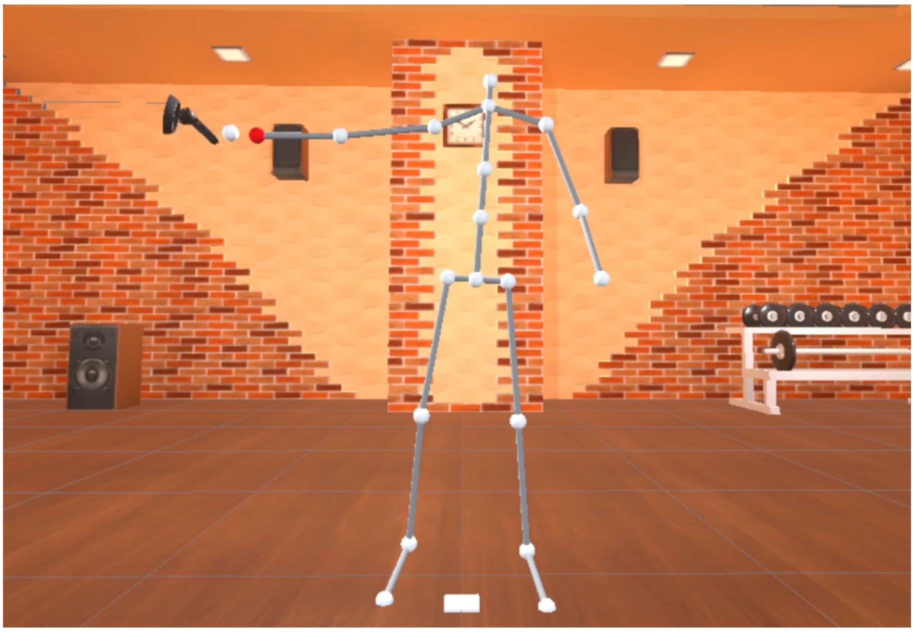

Figure 4.5: A student in the calibration process. His goal is to move the red sphere, his Azure Kinect hand, into the white sphere, his HMD hand, using buttons on the controller.

4.4 Calibration and Positioning

LiftVR is designed to be used with an Azure Kinect and an HMD without base stations in order to minimize cost, hardware requirements and setup time. As a result, LiftVR must deal with two separate tracking systems, namely the HMD position (inside-out-tracking) and tracker output (outside-in-tracking). The correct alignment of these two systems is crucial for the student to perceive the skeleton as their own body. LiftVR accomplishes this in a two step calibration process.

Algorithm 4.2 LiftVR Rough Alignment

procedure ROUGHALIGNMENT(*kinectBody*, *vrRig*, *vrCamera*)
  calibPoint ← GETCURRENTHEADPOSITION(*kinectBody*)
  vrRig.position ← VECTOR3(*calibPoint*.x, *calibPoint*.y, *calibPoint*.z - 0.15)
  cameraToVrRig ← *vrCamera*.position - *vrRig*.position
  cameraToVrRig.y ← 0
  vrCamera.position ← vrCamera.position - *cameraToVrRig*
  vrCamera.y ← −0.07
end procedure

Calibration For LiftVR's calibration phase we designed an alignment algorithm as seen in Algorithm 4.2. We first move the VR-Rig (theVR-Camera is contained therein) onto the head joint position reported by the tracker. Unfortunately, the camera itself cannot be moved directly. This means its container, the VR-Rig, must be repositioned. During this repositioning, the rig is shifted backwards by an additional 15 centimeters, which corresponds to the distance between the sensor in

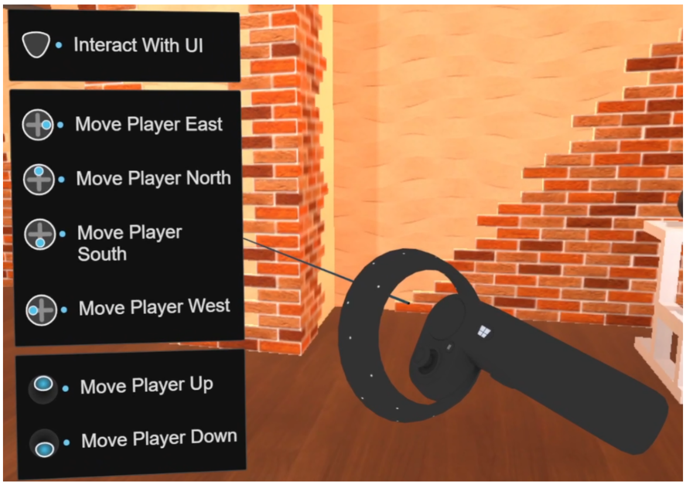

Figure 4.6: A student looking at the control scheme while within LiftVR

the HMD and the student's actual head. Next, the offset between the VR-Camera and the VR-Rig is calculated, and the VR-Rig is shifted to place the VR-Camera at its intended position. Finally, we shift the VR-camera downward by an additional seven centimeters. As a result, the two tracking systems almost entirely overlap. Figure 4.5 illustrates the achieved position (red sphere), which is not far from the intended position (white sphere nearest the controller).

Following this rough alignment, a student must manipulate a red sphere, representing his wrist. Their goal is to maneuver it towards a white sphere near the controller. They can do this using the buttons shown in Figure 4.6. While the student may perceive that pressing the *Move Player North* button moves the red sphere towards the goal, in reality, LiftVR shifts the VR-Rig backwards, effectively bringing the goal closer to the student. Though an external operator could perform this precise calibration faster than a student, we made the conscious decision not to design it that way, as it would make LiftVR unusable as a standalone product. Furthermore, interaction with the system increases student's immersion.

Positioning The depth mode of the Azure Kinect camera used for LiftVR has an operational range of 0.5 to 3.86 meters. As we are implementing a full body motion guidance system, the camera must be able to clearly see the students entire body. One challenge is that during the lifting motion, joint confidence drops sharply when students squat. Thereby, the camera must be positioned low and near to the student, which improves joint recognition during squats.

Through trial and error we determined the ideal camera height to be 0.5 meters for students smaller than 170cm and 0.75 meters for all those who are taller. Regardless of their height, we have students stand at at a *Performance Point*, which is 1.8 meters away from the camera. We found that this distance enables the tracker to maintain good confidence during all sections of the lift. In Figure 4.5, the *Performance Point* is visible as a white cube on the ground. This visual indicator lets players always know where approximately they are to stand.

Nr.	Mechanism	Assistance	Temporal	Presentation	Modality
[1]	Sphere Joints	Feedback/Support	All	Natural	Visual
[2]	Cylinder Bones	Feedback/Support	All	Natural	Visual
[3]	Capsule Zone	Support	Concurrent	Natural	Visual
[4]	Zone Color Indicator	Feedback	Concurrent	Natural	Visual
[5]	Sphere Zone	Support	Concurrent	Natural	Visual
[6]	Sphere Color Indicator	Feedback	Concurrent	Natural	Visual
[7]	Rubber Bands	Feedback	Concurrent	Natural	Visual
[8]	Band Color Indicator	Feedback	Concurrent	Natural	Visual
[9]	Foot Zone Cylinder	Feedback/Support	Concurrent	Natural	Visual
[10]	Foot Color Indicator	Feedback	Concurrent	Natural	Visual
[11]	Audio Countdown	Support	Preliminary	Natural	Audio
[12]	Text Countdown	Support	Preliminary	Abstract	Visual
[13]	Joint Path Line	Feedback	Terminal	Natural	Visual
[14]	Score Text	Feedback	Terminal	Abstract	Visual
[15]	Worst Joints Text	Feedback	Terminal	Abstract	Visual
[16]	Back Angle Visualizer	Feedback	Concurrent	Abstract	Visual
[17]	Angle Color Indicator	Feedback	Concurrent	Abstract	Visual
[18]	Recording Indicator	Support	Concurrent	Abstract	Visual
[18]	Digital Barbel	Support	Concurrent	Natural	Visual

Table 4.2: LiftVR Guidance Mechanisms.

4.5 Guidance Systems

LiftVR implements a total of 18 guidance mechanisms, all of which are listen in Table 4.2. All except one are visual, with most occurring as concurrent feedback. Note that the classification of mechanisms is partially dependent on their usage, which is why mechanisms such as *Sphere Joints* serve both as feedback and support. These have been combined to form the various guidance systems seen in Table 4.3. For example, the *Expert Skeleton* (Table 4.3 row [2]), is visualized with *Sphere Joints* and *Cylinder Bones* (Table 4.2 rows [1,2]). The Ghost Skeleton Guidance (Table 4.3 row [9]) is visualized through the subsystems *Expert Skeleton*, *Ghost Expert*, and *Sphere Indicators* (Table 4.3 rows [2,3,4]).

Noticeably, we classified *Audio Countdown* and *Text Countdown* as preliminary, since they occur before the motion and are thus neither concurrent nor terminal.

In the following subsections, LiftVR's most important systems (Table 4.3 rows [9,13]) are discussed in detail, beginning with concurrent guidance, followed by terminal guidance. The inclusion of dedicated sections for individual subsystems and mechanisms has been omitted due to their self-evident nature and their comprehensive explanations provided within the context of the main systems.

Nr.	System	Subsystems	Mechanisms	Figure
[1]	Student Skeleton	-	[1,2]	-
[2]	Expert Skeleton	-	[1,2]	-
[3]	Ghost Expert	-	[1,2]	-
[4]	Zone Indicator	-	[3,4]	-
[5]	Sphere Indicator	-	[5,6]	-
[6]	Rubber Band Indicator	-	[7,8]	-
[7]	Expert Foot Indicator	-	[9,10]	-
[8]	Student Foot Indicator	-	[9,10]	-
[9]	Back Angle Indicator	-	[16,17]	-
[10]	Joint Path Visualizer	-	[13]	-
[11]	Last Motion Replay	[1,2]	[1,2]	-
[9]	Ghost Skeleton Guidance	[2,3,5]	[1,2,5,6,7]	-
[10]	Zone Guidance	[4,5]	[3,4,7,8]	-
[11]	Foot Location Manager	[7,8]	[9,10]	-
[12]	Evaluation Manager	[10,11]	[1,2,13,14,15]	-
[13]	Heads-Up Display	[9]	[12,16,17,18]	-

Table 4.3: LiftVR Guidance Systems

4.5.1 Concurrent Guidance

In this subsection, the most important concurrent guidance systems, *Student Skeleton*, *Foot Location Manager*, *HUD*, *Ghost-Skeleton*, and *Zone*, *Digital Barbell*, are presented and explained.

Student Skeleton

Students need concurrent feedback on their position, which LiftVR implements with the *Student Skeleton* guidance system. Firstly, Student Skeleton visualizes users by rendering white spheres, with a diameter of five centimeters, at joint positions. These are then connected by white cylinders with two centimeter width. The stick-and-ball skeleton resulting from these two steps can be seen in Figure 4.7.

We made the conscious decision not to visualize all joints available in the Azure Kinect's body model (Figure 2.3 on page 23 as reference). The ears, eyes, nose, hands and hand-tips have been removed. As we attempted to keep visuals simple, facial information seemed purposeless. During development, we quickly noticed the hand and hand-tip tracking performance to be too inaccurate. Inaccuracy was also given for the feet an ankles when students entered a squatting posture. Students would slowly enter a "superman" pose as their legs drifted backwards (Figure 4.7a). In general, tracker lower body confidence would cause foot and ankle joints to violently stutter. To combat this, LiftVR freezes the feet and ankle joints when performances begin (Figure 4.7b). As students are instructed not to move their feet once the motion begins, most of them are unaware of this mechanic. Whilst knees can still shift somewhat unnaturally, violent stutters and "superman" poses are avoided, preventing users from losing immersion.

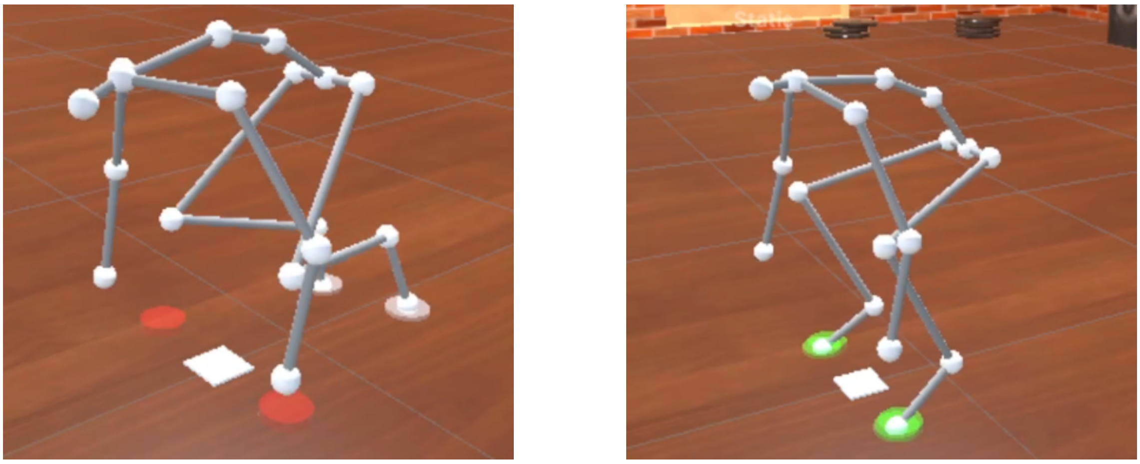

(a) Student without foot-lock enabled. **(b)** Student with foot-lock enabled.

Figure 4.7: Shows the student skeleton, the foot location indicators, and how the foot-lock mechanism prevents foot and ankle backwards shift.

Foot Location Manager

We designed a guidance system which ensure students correctly position their feet before every lift, guaranteeing a natural, wide stance. *Foot Location Manager* achieves this by combining the *Expert Foot Indicator* support system and *Student Foot Indicator* feedback system. *Expert Foot Indicator* places two red cylinders on the ground at the positions of the expert's feet joints, as seen in Figure 4.7a. As the expert has already been scaled to the student's proportions, this distance represents a perfect student stance. Through the *Student Foot Indicator*, two additional cylinders highlight the students current feet positions in white.

Upon placing a foot into the associated *Expert Foot Indicator*, the *Student Foot Indicator* turns green, signaling it is in position. In Figure 4.7a the student has brought both of their feet into the correct position. The indicators have a five centimeter tolerance, making it easy to get in position whilst preventing large discrepancies.

HUD

During the motion, a student's should be duly focused on the concurrent guidance systems. This discouraged us to attach abstract concurrent feedback to a fixed position, as this would require students to break visual focus to look at it. Instead, we designed a HUD, which contains all of LiftVR's abstract concurrent feedback, ergo the *Back Angle Indicator*, *Countdown Text*, and *Recording Indicator*. This is implemented via a canvas attached to the students VR-camera, causing the guidance systems to always remain in the students FOV. To improve readability we slightly angled systems towards the center, as seen in Figure 4.8b. Furthermore, systems are only displayed at the edge of the students FOV, as not to distract from the main concurrent guidance system.

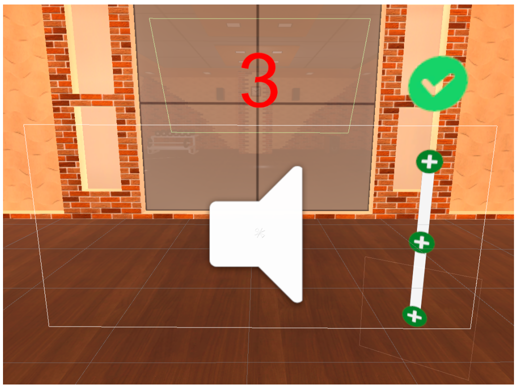

(a) Frontal view of the Heads-Up Display (HUD)

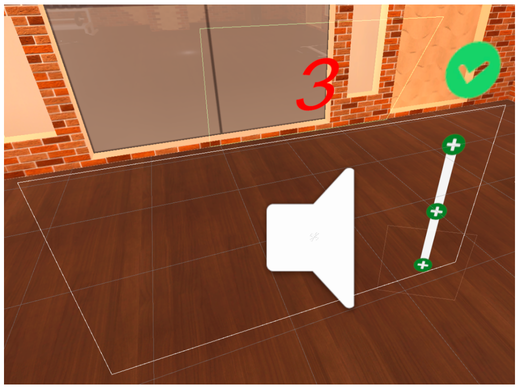

(b) Angled view of the Heads-Up Display (HUD)

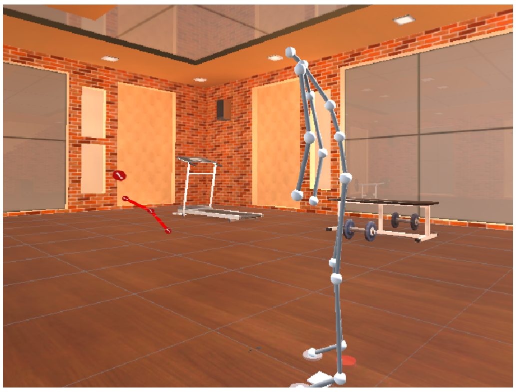

(c) Spectator's view of a back angle violation.

Figure 4.8: The HUD is attached to the main camera and displays its elements angled towards the player. It is also the source of the countdown audio.

Countdown Text and *Recording Indicator* both share the same text field on the northern quadrant of the students vision. As the name implies, the first system counts down from 3 and is replaced by a single red dot, the *Recording Indicator*, once the countdown has concluded. This dot remains until the performance is completed.

The *Back Angle Indicator* is LiftVR's most abstract and novel concurrent feedback feature. It is designed as a green bar with three joints and two bones connecting these (Figures 4.8a and 4.8b). A background manager named the *Back Angle Manager* continuously measures the students back curvature. It does this by retrieving the *Neck*, *Pelvis*, *SpineChest*, and *SpineNaval* joints. From these it calculates a new fictitious joint named *CenterBack* from the average positions of *SpineChest* and *SpineNaval*. Afterwards it determines two vectors going from *CenterBack* to *Neck* and *Pelvis*. The retrieved angle is then increased by a static value of five, due to the tracker reporting too low curvature. Finally, we rotate the upper joint of the *Back Angle Indicator* anti-clockwise by 180°, subtracted by the retrieved angle. If this value drops below the threshold of 170°, the joints, bones, and status symbol of the *Back Angle Indicator* turn red to indicate that the back must be straightened. Figure 4.8c shows a spectators view of a student committing such a back angle violation.

Ghost-Skeleton

The first of LiftVR's two concurrent guidance systems, *Ghost-Skeleton* (Figure 4.9) combines two support systems, one feedback system and a standalone mechanism. We discuss their designs and implementations one by one:

- **Expert Skeleton** *Ghost-Skeleton's* design is based on the most common approach currently found in full body motion guidance. As such, we implemented a humanoid, yellow stick-and-ball skeleton to represent the expert, aptly named *Expert Skeleton*. It is constructed from the same spherical joints and cylinder bones that are used for the *Student Skeleton*. We visualize it nearly fully opaque, as seen in Figure 4.9.
- **Ghost Expert** We designed a *Ghost Expert*, which is simply a highly transparent second visualization of the Expert Skeleton, which can clearly be seen in Figure 4.9b. This system was added due to the work by Yu et al. [YAM+20], who found that additional transparent limbs are the best visual indicators for upcoming postures. It makes apparent whether the next position will involve lowering oneself into a squat or rising therefrom.
- **Sphere Indicators** As the only concurrent feedback system, *Sphere Indicators* were designed to help students determine whether or not their joints are overlapping the expert joints. This was implemented through transparent spheres with a diameter of fifteen centimeters, which are placed on the neck, elbow, wrist, pelvis, and knee. When a student joint enters the correct sphere, it will turn from red to green, indicating the correct position is achieved. Figure 4.9c shows a student attempting to assume the *Expert Skeleton's* posture, however, only his knees and pelvis are in the correct positions.
- **Rubber Bands** Yu et al. [YAM+20] implemented a concurrent feedback system which connects the students joints to their intended positions with lines. For Ghost-Skeleton, we decided not to join the student to the expert, but rather the *Ghost Expert* to the *Expert Skeleton*. This adds additional support to indicate the next upcoming position. No *Rubber Bands* were added to the student in *Ghost-Skeleton* to reduce the amount of simultaneous visuals.

Ghost-Skeleton requires the student to overlap the aforementioned eight joints with a tolerance of fifteen centimeters. Initially, *Ghost-Skeleton* was designed with *Sphere Indicators* on all joints above the ankles with a tolerance of ten centimeters. This was quickly deemed impractically difficult and replaced with the current implementation. Individual shoulder and hip joints were removed from the feedback process, as their close proximity to each other would cause overlapping sphere and confusion. As rotating the hip or shoulder would misplace the arm or leg joints, this design has no negative effects on student posture.

This design is our attempt to implement a state of the art full body guidance system with a humanoid expert. The only novel design contribution from LiftVR are the *Sphere Indicators*, as we strictly rejected concurrent feedback systems which change the appearance of the student. Therefore, visual indicators of correct positioning were adorned on the expert, keeping the student consistent for both of the competing concurrent feedback systems. With these rules, we went on to design our own concurrent system, *Zone*.

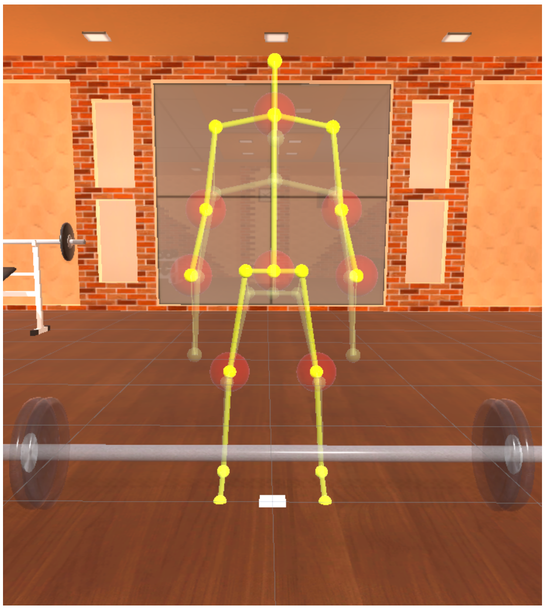

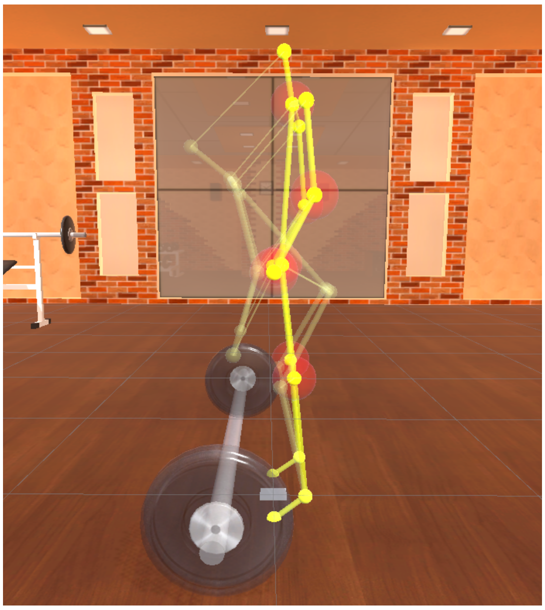

(a) *Ghost-Skeleton* front view **(b)** *Ghost-Skeleton* side view

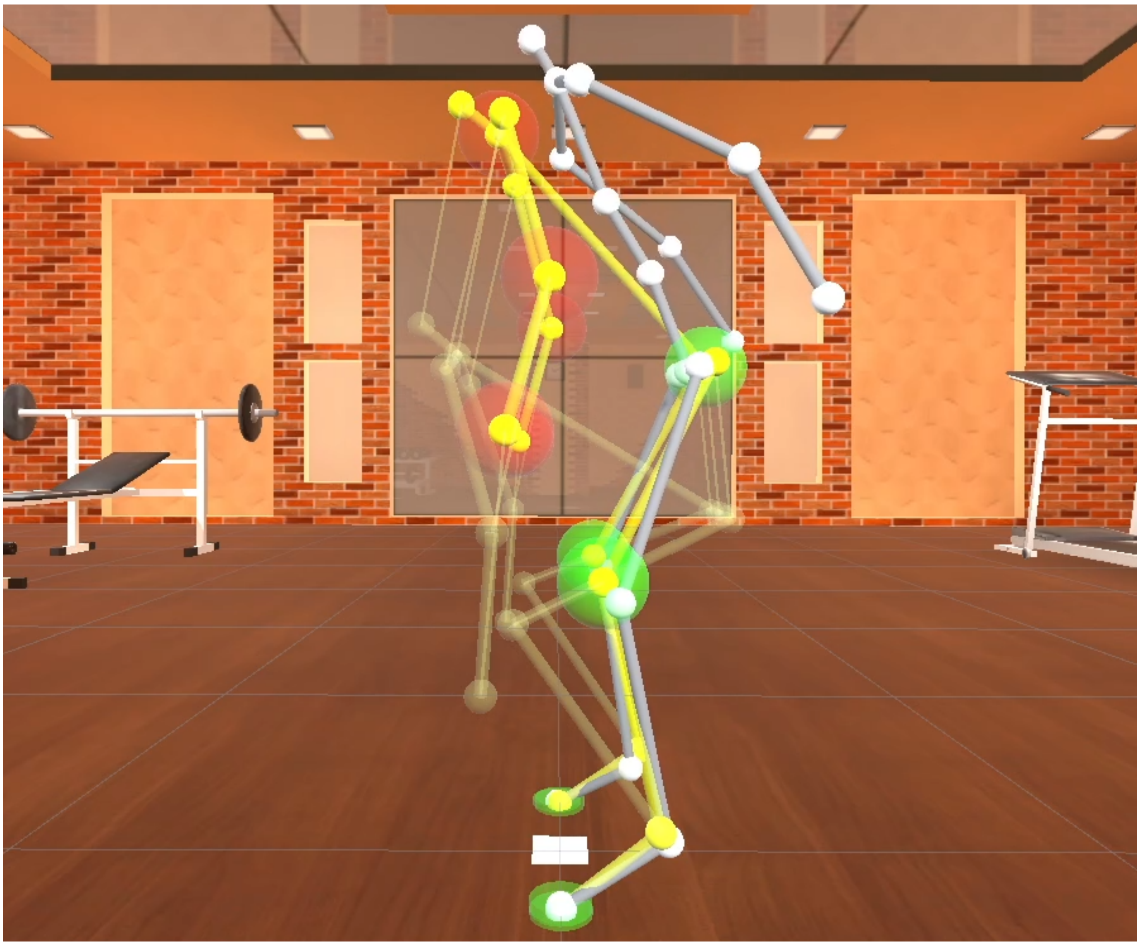

(c) Side view of a student trying to get into the next *Ghost-Skeleton* position.

Figure 4.9: Multiple views of the *Ghost-Skeleton* system.

Zone

Offering an alternative to the state of the art approach of full-body motion guidance by humanoid expert visualization, *Zone* is the second concurrent guidance system implemented in LiftVR. It displays only one abstract support system and two feedback systems to guide students through the motion (Figure 4.10). These are the following:

- **Zone Indicator** We designed capsules that float in space and are adorned with a special material to maximize visibility named *Capsule Zones*. As seen in Figure 4.10b, we implemented a shader that roughly outlines an object by increasing occlusion for surfaces aligned with the camera, whilst keeping surfaces perpendicular to it transparent. When external objects, such as the *Student Skeleton*, enter the zone, the occlusion effect is generated around their point of entry, making it clear when limbs enter the capsule.

  LiftVR utilizes the fact that for a lifting motion, a perfect demonstration mirrors the body on a vertical plane going through the spine. Consequently, joints which are mirrored can be combined to form one zone instead of two. We implemented this design for *Zone*, and Figure 4.10a shows how knees, wrists, and elbows are paired into singular capsules.

  Another design decision was that capsules will linearly interpolate between red and yellow when both limbs are outwith the area and from yellow to green when one joint is positioned correctly. Similarly to the *Sphere Indicators* from *Ghost Skeleton*, capsules have a width of fifteen centimeters, with the length being determined by the distance of the paired joints. In contrast to *Ghost-Skeleton*, *Zone* does not care where in the capsule the relevant student joints are, merely them being inside satisfies the system.

- **Rubber Band Indicator** As *Zone* does not display an superimposed expert skeleton and relies only on *Capsule Zones* as a support mechanism, additional feedback is required to help students understand the guidance system. For this, the *Rubber Bands* system was added which has lines extrude from the students skeleton to their intended positions within the *Capsule Zones* (Figure 4.10c). Our design addition is that these shrink in size the closer they get to the intended joint position, indicating the joints should be moved into the *Capsule Zones*. Furthermore, the bands interpolate their color from red to green based on the bands current length.

Compared to *Ghost Skeleton*, *Zone* is more abstract and was designed with an additional feedback systems to help students understand the objective. Nevertheless, the reduction in visual clutter and the pairing of mirrored joints enables *Zone* to display the motion more understandable and visually indicative. This was achieved by designing the system with the motion in mind and not relying on a humanoid visualization. As a result, *Zone* only has five points of interest, whilst *Ghost-Skeleton* has eight. With our concurrent guidance systems completed, we moved on to design the visualization of the bar students would lift.

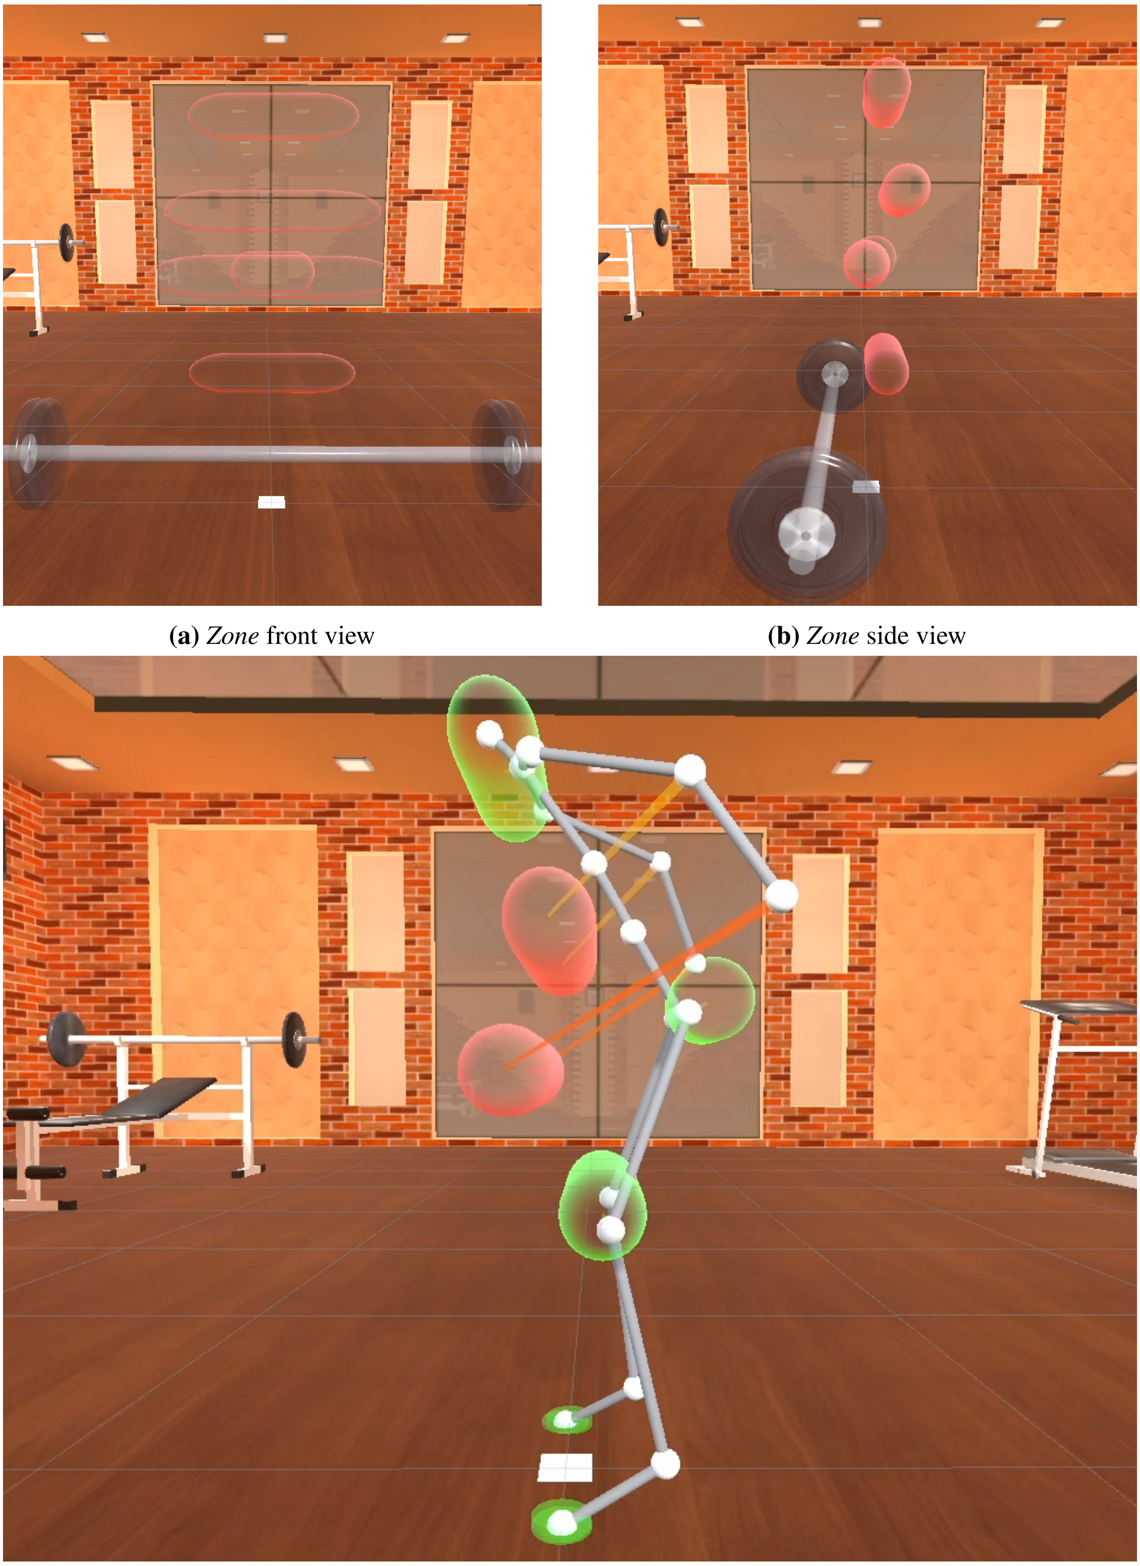

(a) *Zone* front view

(b) *Zone* side view

(c) Side view of a student trying to get into the next *Zone* position.

Figure 4.10: Multiple views of the *Zone* system.

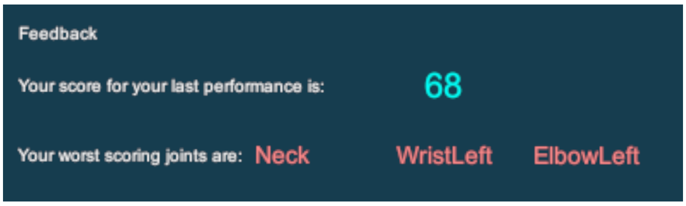

Figure 4.11: *Score Text* and *Worst Joints Text* mechanisms shown in the *Evaluation Manager*.

Digital Barbell

LiftVR was designed with the intention of students performing the lifting motion whilst interacting with a real barbell. However, due to our restraints on hardware and external trackers, no solution was found to actively track and visualize the weight. Instead, we added a barbell 3D model onto the performance point which sits at the exact height and distance as the real bar. It is intentionally highly transparent to clarify to the user that it is but a indication, not an active visualization. This *Digital Barbell* can be seen in Figures 4.9a, 4.9b, 4.10a and 4.10b. As explained in Section 2.4 on page 29, someone performing a lift should gaze upward or forward. Thus, students should be able to grab the bar without looking at it, which makes grabbing the bar in LiftVR not unnatural. Therefore, the *Digital Barbell* is designed to get students familiar to the system and help immersion by showing them where they can grab the bar.

4.5.2 Terminal Guidance

We designed and implemented an array of terminal guidance systems in order to investigate the effect of these in combination with concurrent motion guidance systems. All of the terminal systems are contained in the *Evaluation Manager*, which we now discuss.

Evaluation Manager

LiftVR's *Evaluation Manager* contains descriptive terminal feedback systems which enable students to review their performance, recognize mistakes and improve.

- **Last Motion Replay** The first terminal feedback system allows students to review their last performance. Specifically, the *Expert Skeleton* and *Student Skeleton* are overlapped and the last motion recording is played. To get a better overview, the student can walk around the active replay to see his performance from various angles.
- **Text Feedback Section** Scoring a students performance implements gamification into LiftVR. We designed an evaluation function which provides the user a score of his last attempt. The average distances of the knees, elbows, wrists, pelvis and neck to their intended position is multiplied by 100. Subsequently, the result is squared in order to increase the score rapidly when the error increases. Figure 4.12 shows a plot of the scoring function for distances between zero and 15 centimeters. Common scores of 36, 50, and 100 are highlighted with black dots. A student achieved a score of 68 in Figure 4.11, thus, his average joint distance was 0.08239 centimeters.

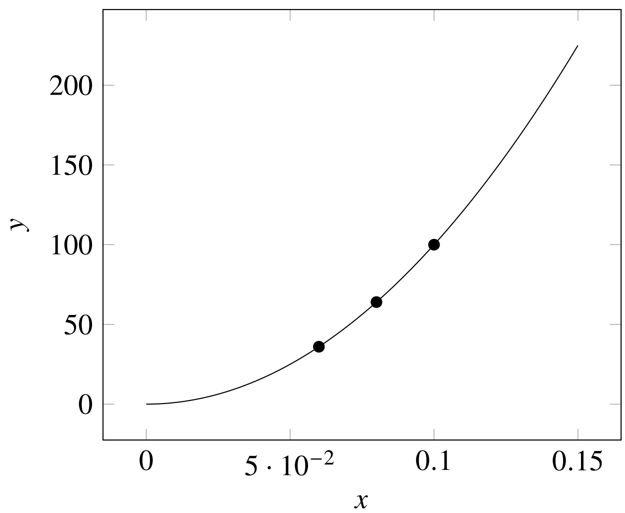

Figure 4.12: Plot of LiftVR's scoring function with highlights $f(0.06) = 36$, $f(0.08) = 64$, and $f(0.1) = 100$.

Average distances are then compared between joints, with the worst three performers being sent to the *Worst Joints Text*. Due to this feedback mechanism, students will know what to look for when viewing the replay. In Figure 4.11, the student's worst joints were his neck, left wrist and elbow . From the combination of wrist and elbow, they should be able to conclude that his left arm was misaligned. His neck having scored poorly indicates that they struggled to keep pace with the experts recording.

- **Joint Path Visualizer** When students receive feedback on their worst joints, they are able to watch their replay to deduce the problem. However, it may be difficult to judge if the error happened during a short time period or over a longer duration. To aid this process, we designed and implemented a novel terminal feedback system named *Joint Path Visualizer* (Figure 4.13). A line is drawn through all joint positions captured during the performance, visualizing the joint's traveled path for the entire motion. Simultaneously, an additional line is drawn for the expert's joint positions. The *Joint Path Visualizer* can be kept active whilst the *Last Motion Replay* is running. This enables student's to clearly recognize which portions of the path belong to which time in the motion and compare themselves to the expert.

Combined, the terminal feedback systems implemented in LiftVR offer powerful insight into the students last performance. A typical use case would be for a student to initially determine whether they improved by comparing their score to the previous performance using the *Score Text*. If their score diminished, they would then visualize the paths that their worst performing joints took *Joint Path Visualizer*. After viewing the paths from multiple angles, they then start the *Last Motion Replay* to determine during which period of the lift deviations occurred. Knowing now how, when, and why they performed poorly, they can start another performance and avoid their previous mistakes.

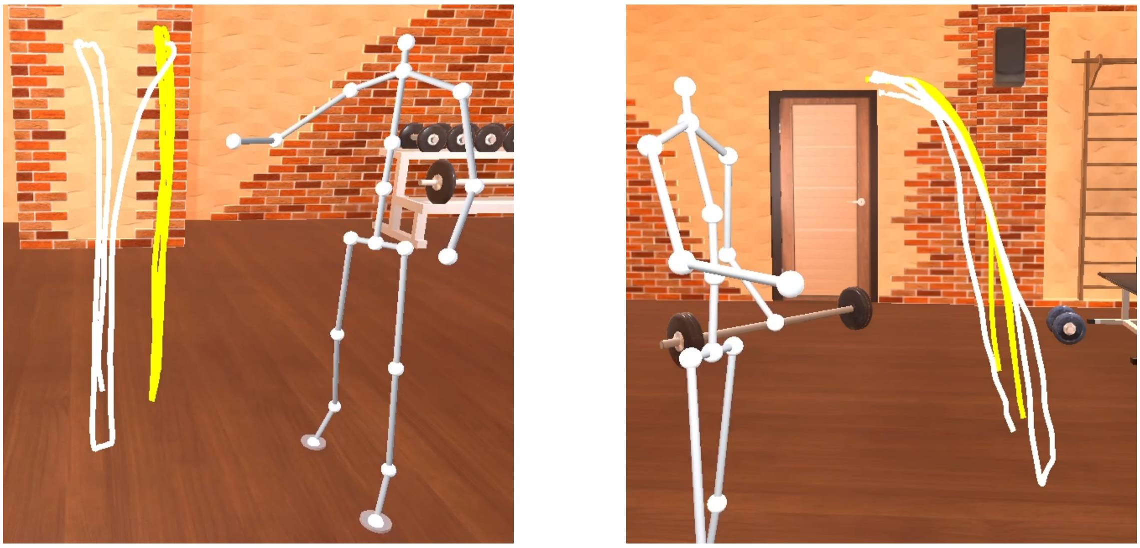

(a) Front view of student studying neck path

(b) Side view of student studying neck path

Figure 4.13: Student evaluating neck joint paths from his last motion. a) clearly shows how the students neck (white) drifted horizontally compared to the experts neck (yellow). However, b) indicates that the joints mostly overlapped vertically.

4.6 Guidance Settings

LiftVR implements multiple perspectives, guidance modes and speed settings to ensure each student can learn by their preference. All settings can be changed by the student with the UI.

4.6.1 Guidance Modes

For our study, we wished to include path and movement practice sections, akin to Anderson et al. [AGMF13]. Their distinction lies in the concurrent guidance system's dependence on the student achieving intermittent postures during the performance.

Path Mode In this posture driven mode, the active concurrent guidance system will only progress its motion when the student fulfills a set of requirements. LiftVR necessitates all positional indicators to contain their respective joints, ergo show green. Consequently, the student has assumed the posture demonstrated by the guidance system. Immediately, the guidance system will assume the next key posture in the motion. LiftVR has identified six of these which are separated through equal time steps. Figure 4.14 shows all six postures, beginning with the expert simply standing. Figures 4.14b to 4.14d describe the path downwards to the bar. Subsequently, the bar is lowered in Figures 4.14e and 4.14f. One full motion, therefore, entails lowering oneself to retrieve the bar, lifting it and returning to the starting position.

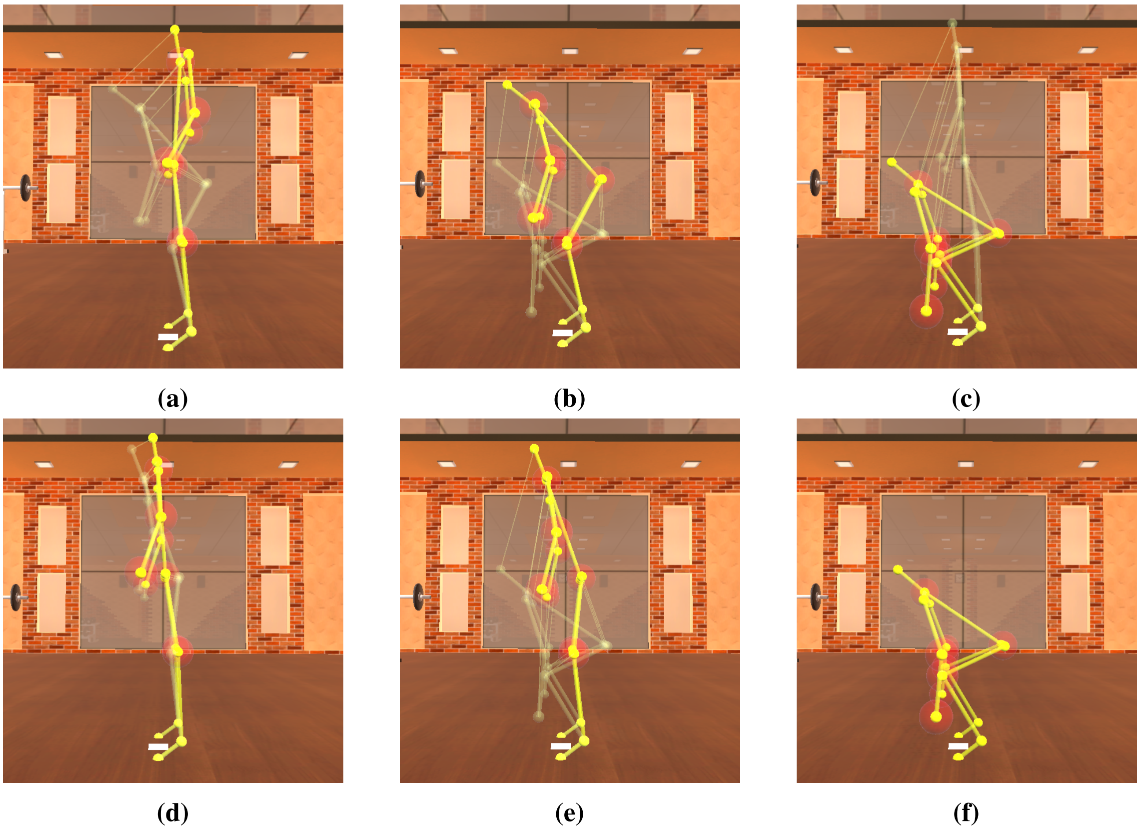

Figure 4.14: Six different postures used for path guidance mode, visualized with *Ghost Skeleton*

Movement Mode Unlike the aforementioned mode, here no requirements must be fulfilled for the active concurrent guidance systems to continue. They will, in fact, move continuously, regardless if the student follows it or not. One potential design was to combine this with path mode. If the student is not in position, it would wait, and then continue movement once the student has assumed the posture. Critically, it would not snap immediately to the next posture. However, as the effects of new practice modes are not part of our evaluation, we did not implement this into LiftVR.

Movement Speed To give students the option of self pacing, we implemented two speed settings into LiftVR. *Slow Mode* will have the motion playing at half speed, whilst *Regular Mode* has it at full speed. We intentionally designed the speeds not to be changed in a scalar fashion, as we do not wish speed to cause an effect in the evaluation of our system. Thus, the decision was made to implement two fixed speeds.

4.6.2 Perspectives

LiftVR was designed around three perspectives, *First Person*, *Side*, and *Mirror*. Of these, only *Side* and *Mirror* are available to the student, as 1PP has been shown to be inferior. 1PP remains hidden to the student and is used before and after guidance sessions for resetting.

In the implementation, the *Student Skeleton* and the selected guidance system are shifted to separate positions. The *Student Skeleton* is visualized from tracker data, which sees the Azure Kinect camera position as the center, with student joints being offset therefrom. For 1PP, the *Student Skeleton* is therefore located at $(0, 0, 0)$. In contrast, the *Expert Skeleton* is from Vicon data which assumes the tracked entity to be the center. Consequently, in 1PP, the *Expert Skeleton* is located at $(0, 0, 1.8m)$, directly on the Performance Point.

To shift and rotate the visualizations for the *Mirror* perspective, the expert can simply be placed at $(0, 0, 0)$ and rotated by 180 degrees. The *Student Skeleton*, however, must be first rotated by the same amount and then offset to $(0, 0, -1.8m)$ to counteract the rotation. A similar adjustment is made for *Mirror*. Whilst *Mirror* indicates a mirrored view, we merely rotate the *Student Skeleton* to enable viewing from the front. This creates a point reflection rather than a mirrored one. Whilst this design is inferior to a truly mirrored view, it was implemented due to time constraints.

4.7 User Interface

LiftVR was designed with the intention of being usable by a single person. As a consequence, all systems, settings, and steps must be available to the user through the UI. Two darkly hued planes are shown hovering in space and are intractable through a laser protruding from the controller. Buttons are highlighted when the laser enters them. We implemented clickable and toggle-able buttons.

4.7.1 Main Menu

The first plane holds the main menu, which contains the *Start*, *Calibrate*, *Perform*, and *Practice* screens (Figure 4.15). To swap between them, the topmost menu buttons are used. The active screens button assume a bright color, melding into the active sub-menu.

Start LiftVR automatically displays this screen upon startup (Figure 4.15a). Its only purpose is to indicate that the tracker has not begun providing information. Therefore, the camera may be disconnected or the student simply has not walked before it yet. Once tracking information is available, the LiftVR State-Machine discussed in the subsequent Section 4.8 will trigger a state update throughout the system. In consequence, the start screen will automatically proceed to the next step, calibration.

Calibrate This screen (Figure 4.15b) holds the calibration controls described in Section 4.4. By pressing the *Rough Alignment* button, the process described by Algorithm 4.2 takes place and the LiftVR State-Machine fires a state update. Consequently, the *Student Skeleton* becomes visible and precise calibration with the controller is enabled. Once the student has concluded the precise alignment process, he can press the *Proceed* button which yet again updates the systems state. Similarly to before, the calibrate screen will disappear and the performance screen becomes active. It is impossible for the student to become visible without completing the alignment procedure at least once.

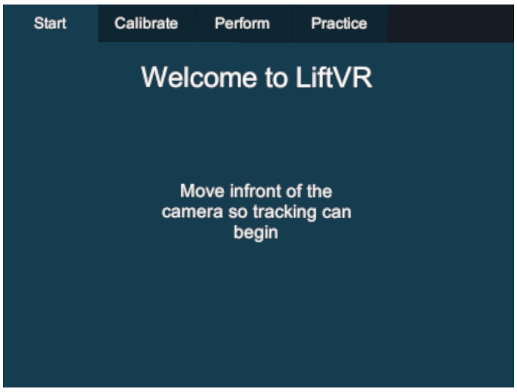

(a) Start Screen.

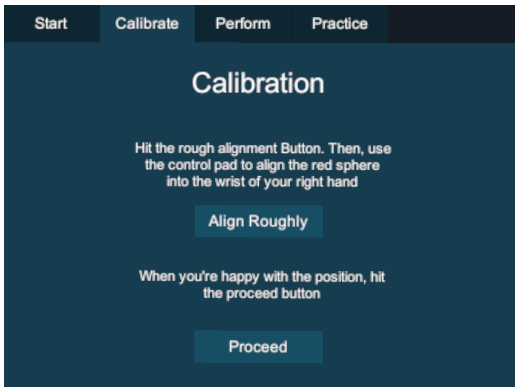

(b) Calibration Screen.

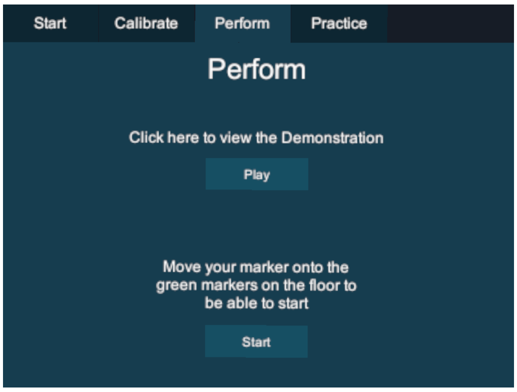

(c) Perform Screen.

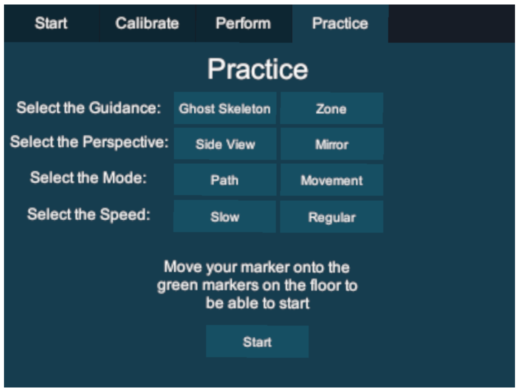

(d) Practice Screen.

Figure 4.15: LiftVR's main menu with its sub-screens.

Perform For our study, the student must perform the motion without concurrent visual guidance after calibration is completed. For this, the performance screen is designed. Before his first attempts, he can watch the *Expert Skeleton* perform the motion on its own by pressing the *Play* button (Figure 4.15c). To begin his performance, he can press the *Start* button, however, it only appears once both of his feet are in position. This logic is controlled by the *Foot Location Manager*. Once pressed, the HUD displays the Countdown Text and *Countdown Audio* engages.

Practice To practice the performance and modify the various settings implemented in LiftVR, we designed the *Practice* screen (Figure 4.15d). Though the figure shows all buttons as visible, when first visited, the screen only offers a choice between *Ghost Skeleton* and *Zone* guidance. Once the student has made his decision, the *Perspective* buttons become visible. Following a selection of these, the *Guidance Mode* must be chosen. If *Path* was chosen, the student can proceed to align his feet, otherwise the *Speed* must be decided upon. Finally, after all options are selected and the feet are in position, the student can begin practice with the *Start* button. This design choice was made to ensure no practice can be started without selecting all necessary options. It is important to mention that none of these visibility adjustments are reactions to buttons being pressed, but rather to changes of the LiftVR-State-Machine.

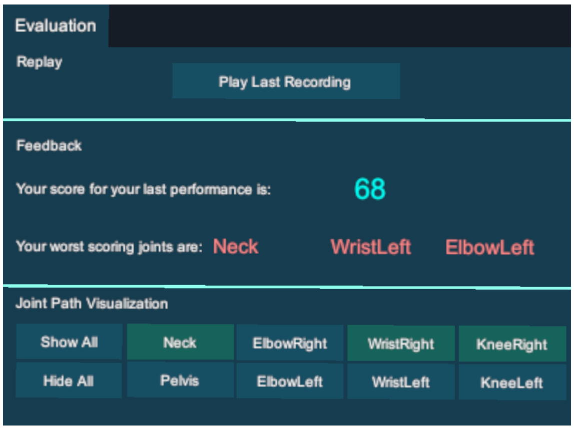

Figure 4.16: Evaluation Screen

4.7.2 Evaluation Menu

The second plane contains only one screen, the *Evaluation Menu* (Figure 4.16). Here, the systems offered by the *Evaluation Manger* (Section 4.5.2) are made interactable.

The upper section contains a button to start the *Last Motion Replay*. The center section visualizes the textual terminal feedback *Score Text* and *Worst Joints Text*. Finally, the lowest section contains buttons to toggle individual *Joint Path Visualizers*. These were designed as toggle-able buttons, which is why we implemented support for these in LiftVR. This makes it apparent which visualizers are currently on or off. In Figure 4.16, the student has currently selected to see the neck and right wrist/knee joint paths. This can be deduced from the respective buttons having a light-green coloration, indicating they are toggled. For ease of use, the student can optionally engage all visualizers at once or remove any active ones with the *Show All* and *Hide All* buttons.

4.8 Software Engineering

As a software engineering student, performance, understandability, and maintainability are of utmost importance to the author. Performance is crucial, as low FPS when working with VR reduces immersion and most critically leads to motion sickness. Therefore, LiftVR was designed to achieve a framerate which matches the refreshrate of the HMD at all times (90 FPS). This requirement was made non-negotiable. Any system which would jeopardize the performance would not be implemented.

As mentioned in Section 1.3 on page 17, LiftVR must be developed with the intent of it being an ongoing solution. Continued development is to be assumed, making understandability and maintainability crucial. Therefore, we designed LiftVR with coding patterns which were strictly followed during development. The most important of these are discussed here.

State	Description
NotTracking	Tracker not active or no data being provided.
UncalibratedTrackingLowConfidence	Student not calibrated, low tracking confidence
UncalibratedTrackingHighConfidence	Student not calibrated. high tracking confidence.
BeginCalibration	High tracking confidence, calibration process begins.
CalibrationRoughAlignment	Student has been roughly aligned.
TrackingAndCalibrated	Student calibrated, high tracking confidence.
NotTrackingButCalibrated	Student calibrated, low tracking confidence.
PracticeSelectGuidance	Guidance must be selected.
PracticeSelectPerspective	Perspective must be selected
PracticeSelectGuidanceMode	Guidance-Mode must be selected
PracticeSelectSpeed	Speed must be selected
PracticeAssumePosition	Selection process complete. Student not in position.
PracticeReadyToStart	Student is in position for practice.
PracticeStart	Practice session has started.
PracticeStop	Practice session has stopped.
PerformanceAssumePosition	Student must assume position for performance.
PerformanceReadyToStart	Student in position for performance.
PerformanceStart	Performance has started.
PerformanceStop	Performance has ended.

Table 4.4: State-Machine States

4.8.1 LiftVR State-Machine

At the core of the application lies the *State-Machine Manager*. LiftVR's system communication was designed to use the publish/subscribe pattern [EFGK03]. Instead of having systems continuously check whether or not they should be active, they subscribe to the *State-Machine Manager*. Whenever the state is changed, all subscribers are notified of the current and previous state. This allows most of LiftVR's systems to lie dormant and waiting, saving computing time. All states with brief descriptions are listed in Table 4.4. To limit the amount of subscribers, so-called *Manager* classes, which are global singletons, distribute actions to systems assigned to them. For example, the *UI-Manager* class will automatically swap from the *Calibrate* to the *Perform* screen when the state *TrackingAndCalibrated* is broadcast. The publish/subscribe pattern integrates perfectly into LiftVR's second core design pattern. *Call Down - Signal Up* does not hail from traditional software engineering, but rather from game development. Objects may call methods on their children, like the UI-Manager enabling and disabling menus, however, the children may only communicate upwards through signals. In LiftVR, these signals are state changes published to the *State-Machine Manager*. Consequently, when swapping from the *Calibrate* to the *Perform* menu, the *CalibrateMenu* class is prohibited from calling the corresponding swapping method on the UI-Manager. Instead it must call *SetState* on the *State-Machine Manager*, to which the UI-Manager is subscribed. The advantages of *Call Down, Signal Up* include less coupling of systems, simplicity, readability, and easier introduction of new systems.

4.8.2 Managers

Managers were designed in LiftVR to reduce the amount of listeners to the state-machine and provide easy global access to data. For the first case, they are responsible for invoking a number of subsystems in their area of responsibility. However, some are not subscribed to the state-machine and merely serve as a data platter. All of the, are implemented as singletons which allows quick, global access. The most important managers are:

BackAngleManager Is disabled until the systems achieves the state *TrackingAndCalibrated*. After that, it runs continuously, being one of the few classes that does so. To optimize for performance, its update loop is limited to 60 FPS. Since new tracking data enters at that speed, there is no point in the *BackAngleManager* running at 90 FPS. Is responsible for checking the students current back curvature and updating the *Back Angle Indicator*.

FeedbackManager Listens for the states *PerformanceStop* and *PracticeStop*. When receiving one of them, the last motion performed by the student is evaluated and the feedback is distributed to other systems.

FootLocationManager Does not subscribe to the main LiftVR state-machine, but runs its own internal state-machine. This design was done to allow it to be used the same way regardless whether the student is practicing or performing, reducing lines of code. The *Foot Location Indicators* subscribe to it and activate when receiving *NoFeetInPosition*. They continuously check and update the *FootLocationManager* until disabled due to *PracticeStart* or *PerformanceStart*.

JointPathVisualizationManager Listens for *PerformanceStop* and *PracticeStop*. Updates line renderers used to display the *Joint Path Visualizer* with positions from the students last motion.

PerformanceManager Is responsible for initiating and running a student performance. Upon receiving *PerformanceAssumePosition*, it activates itself and the *FootLocationMangager*. *PerformanceStart* triggers the *Countdown Text* and *Countdown Audio*. The update loop waits until the countdown has subsided, after which it starts the *RecordingManager*. Finally, it disables itself, the *FootLocationMangager* and the *RecordingManager* upon receiving *StopPerformance*.

PracticeManager Works similarly to the *PerformanceManager*, except that *PracticeStart* will start the concurrent guidance systems, receiving the current setting from the *GuidanceManager*. Once the guidance system has completed the motion, it sends *PracticeStop* which disables the *PracticeManager*.

RecordingManager Enabled or Disabled by the *PracticeManager* and *PerformanceManager*, this system captures the students and experts current position at 60 FPS. Once disabled, the captured information is stored as a completed motion on the disk. All of LiftVR's evaluations work by comparing the experts last recording to the students last recording. This causes the *RecordingManager* to re-record the expert for every performance. Typically we would have not designed it this way, however, we wanted to save this data in case we needed it during any point in our study.

4.8.3 Abstract Guidance System

As we had to develop two concurrent guidance systems in two practice modes, we had to design in such a way that we would complete the implementations in time. Therefore, limiting code duplication and maximizing shared logic was a core concept of the concurrent guidance systems design. For this, we used an abstract class to develop a *BaseGuidanceSystem*. This system contains all the functionality and control to run a concurrent guidance system in path or movement practice. in its update loop, it calls seven distinct abstract methods which handle certain steps in the concurrent update pipeline. With this base class developed, all we needed to do was have our individual systems implement the abstract methods and all the control and update logic would be working and consistent throughout systems. The seven methods are:

HandleVisualizationUpdateModeContinuous Invoked every update if the system is in movement mode. For Ghost-Skeleton the sphere indicators are colored, for Zone the capsules.

HandleVisualizationUpdateModeStepwise Invoked every update if the system is in path practice mode. The crucial difference to the aforementioned method is that here the active guidance system must check if the student fulfills its progression requirement. If so, posture progression is triggered in the next frame.

HandleVisualizationUpdateWaitingForCountdown This method is called exclusively whilst the student waits for the countdown to finish. Both implemented systems use this to have their concurrent feedback systems already running during the countdown.

HandleMotionFrameIncreased This method gets called when the system is in movement mode and we just progressed the expert to the next frame. For example, Ghost-Skeleton moves the *Rubber-Bands* to follow the experts current position.

HandleBeforeVisualizationToStepMovement Called before the expert's joints are moved into the next posture of the motion. Only used by Ghost-Skeleton, this method is used to determine whether or not to show the *Ghost-Expert* and *Rubber-Bands*.

HandleJointAfterBaseStepMovement Invoked for every single expert joint when it is moved into a new posture. Though not used by any of our currently implemented guidance systems, this hook is offered for future developers.

HandlePostVisualizationToStepMovement Called after all of the experts joints have been moved into a new posture.

To add new concurrent guidance system option to LiftVR, a developer simply needs to implement these methods for himself, saving time and reducing code. With our system design and implementation complete, we now move on to the evaluation.

5 Evaluation

This chapter describes the execution of our study, results obtained therefrom and subsequent analysis.

5.1 Methodology and Approach

5.1.1 Study Design

We used a between-subject multivariate design with the two independent variables *Guidance System* and *Feedback Mode*. Thereby, we had a 2×2 design with 4 distinct groups.

Our LiftVR study design is based on the study executed by Anderson et al. [AGMF13]. Firstly, students entered LiftVR's gym environment and learned the controls. Secondly, we explained how the UI works and asked students to get their feet into position. They were instructed to drop the controller once a performance begins, followed by a squat to grab the *Digital Barbell*. Retrieving the bar was practiced until participants were conformable grabbing it without looking downwards, trusting it was there. Only after these steps did the main phase of the study begin.

To receive a baseline measurement, the student watched the *Expert Skeleton* demonstrate the motion. They were permitted to do so a maximum of three times, after which they performed the motion themselves without aid. Three performances were captured to construct an average baseline measurement.

Subsequently, the student completed five repetitions of path practice using one of the two concurrent *Guidance System* implementations. If the student was in the group with terminal feedback enabled, they watched the replay of their last performance. Importantly, students were told to ignore score and joint path feedback options until movement practice mode begins. The *Evaluation Menu* was invisible to students not in the feedback group.

After completing path practice, the student performed five repetitions of slow movement practice. For this mode, students with access to terminal feedback were instructed to lower their score and were shown how to use the systems to do so. Importantly, whilst all students were made familiar with these systems, they were not explicitly instructed to use them after every attempt. The usage of the systems was up to the individual user. To complete movement practice, the student completed five additional repetitions in regular speed.

Finally, students performed three times without any of the concurrent or terminal guidance system enabled. This was done to investigate short-term retention and receive a comparison to the baseline measurement

Variable	Baseline	Path	Slow Movement	Regular Movement	Final
Error	X		X	X	X
Frame-Error	X	X			X
Best Error		X	X	X	
Improvement		X	X	X	
First Improvement		X			

Table 5.1: Shows which variables where calculated from the performances during the study.

5.1.2 Data Collection and Processing

After each performance, the motion captured by the *RecordingManager* was saved to disk. These files document the positions of every joint for each frame in a student's performance. Furthermore, metrics calculated from LiftVR's evaluative systems are also stored within them. These are the average joint distance, the worst performing joints and the sum of frames with back angle violations. From these measurements, we calculated the dependent variables used in our analysis.

Performance Variables We distinguished four different variable categories calculated from performance recordings:

- **Error** describes the average distance of the students joints to the experts joints during the performance and is measured in meters. As such, an error of $.17m$ expresses that joints averaged 17 centimeters of distance to the expert.
- **Frame-Error** describes the absolute difference in frames between the experts performance and the student. The expert recording has a duration of 384 frames, which when played back at 60 FPS equates to 6.4 seconds. For example, a frame-error of $600f$ indicates that the student continued his performance for 10 seconds after the expert was already done.
- **Best Error** is the lowest error or frame-error the student produced during the performance.
- **Improvement** is the difference, in percent, of the students best and worst performance.
- **First Improvement** is the difference, in percent, of the students first and second performance.

Not all variable categories were calculated and analyzed for each performance, as some variables hold no value for specific performances. For instance, in movement practice, the performance ends instantly when the expert completes his motion. Thereby, the frame-error is always zero. Similarly, the average error in path practice holds no value, as requiring extensive time to enter a posture lowers the average error. This is why students were told not to pay attention to their scores during path practice. Table 5.1 shows precisely which variables were calculated and analyzed from which performances. Noticeably, we were only interested in the first improvement during path mode, as this was users first introduction to the terminal feedback systems. Effects on their performances would thus be most apparent there.

Two additional dependent variables were used to compare the baseline performance with the final measurements. These are called **Overall Error Improvement** and **Overall Frame Improvement** and were formed by subtracting the average errors and frame-errors of the final performance from those of the baseline measurement.

User Demographics and Feedback Before the study commenced, users filled out a demographic questionnaire to use as independent variables in our analysis. This was to investigate interaction effects between different types of users and the system. We ensured each level of the resulting groups had at least three occurrences in our data:

- **Gender** - two levels (male, female).
- **AgeGroup** - three levels (age 18-24, 25-34, and 55-65).
- **LiftExperience** - two levels (experience, no experience).
- **ExerciseGroup** - four levels (< 1 per week, 1 per week, 2-3 per week, 4+ per week).
- **VRExperience** - two levels (experience, no experience).
- **HeightGroup** - four levels (155-165 cm, 166-175 cm, 176-185 cm, > 186 cm).
- **WeightGroup** - four levels (50-60 kg, 61-70 kg, 71-80 kg, 81-90 kg, > 91 kg).

After they completed training, we asked students to fill our a questionnaire to receive feedback on the systems. Our question categories are based on the study by Moesgen et al. [MSP+22]:

- **Understandability** How easy was it to understand how the guidance system works?
- **Mental Effort** During practice, how mentally demanding was it to follow the guidance system?
- **Helpfulness** When in an incorrect position, how helpful was the guidance system in highlighting the discrepancy and correcting the mistake?

Unlike the study our questions are based on, we chose a 1-7 Likert scale for scoring due to advantages over the 1-5 version [JKCP15]. From the individual categories, we calculated a **Total Score**.

We were also interested in qualitative results. Specifically, we wanted to know what users liked or disliked about LiftVR and gather suggestions for future systems. All users were asked to imagine additional concurrent guidance implementations. Those in the group with terminal feedback were asked which additional systems they would like. Others were asked to imagine feedback systems in general.

5.2 Analysis Techniques

We used one-way ANOVA to investigate the individual effects of independent variables. Two-way ANOVA was utilized to search for interaction effects. Every two-way ANOVA was accompanied by least significant difference (LSD) post hoc analysis. As *Guidance System* and *Feedback Mode* contain too few levels, they were combined to form a four-level group (*Guidance*×*Feedback*) for post hoc testing. During the investigation of the dependent variables, we checked Levene's homogeneity of variance to ensure equal variances across the samples. If the assumption of equal variances was violated, we used Welch's ANOVA for more robust results. Two-way ANOVA's were not performed on variables with unequal variances and we relied on post hoc analysis. To check for a significant difference between the users baseline and final performances, we used s one-sided paired-sample t test.

5.2.1 Reporting Terminology and Scheme

We differentiated between three levels of significance ($p \leq .05$ *, $p \leq .01$ **, $p \leq .001$ ***). The term "highly" significant was used merely as an indication that $p \leq .01$. If the p value did not show significance, we investigated the means and standard deviations for differences. Consistent differences between levels of independent variables on multiple dependent variables were seen as reference points to be conclusively confirmed or rejected from future higher powered studies.

For brevity, we did not explicitly state every non significant result. If results for a specific variable were not reported, they were not significant either individually or in interactions. Unless explicitly stated, variables fulfilled homogeneity assumptions and could be analyzed with regular ANOVA.

5.3 Study

Through our collaboration with INSPO, we were required to conduct a joint study at their motion capture laboratory. This was to aid in the construction of a lifting motion data set. As such, the study took place in two locations over a span of ten days. Overall, 24 participants (15 male, 9 female) used LiftVR. Of these, 9 were recruited at the university during the joint study, and 15 participated in the home study.

5.3.1 Joint Study

For the joint study, participants were required to strip their clothing and adorn markers for the Vicon system as seen in Figure 5.1. Tracker placement and Vicon system calibration would require around 30 minutes. Figure 5.1a shows a participant being calibrated to the Vicon system, whilst a ladder with LiftVR's Azure Kinect placed upon it can be seen in the distance. The participant in Figure 5.1b has grabbed the bar and is trying to assume a posture with it. For our study, we wanted to ensure the bars weight does not affect participants lifting abilities. Furthermore, loaded bars would have tired users significantly due to the high amount of practice repetitions. As such, we used a light, yet tangible curtain rod for users to grab instead. This had no impact on users immersion, as the *Digital Barbell* overlaps with the real-world entity on the performance point in LiftVR.

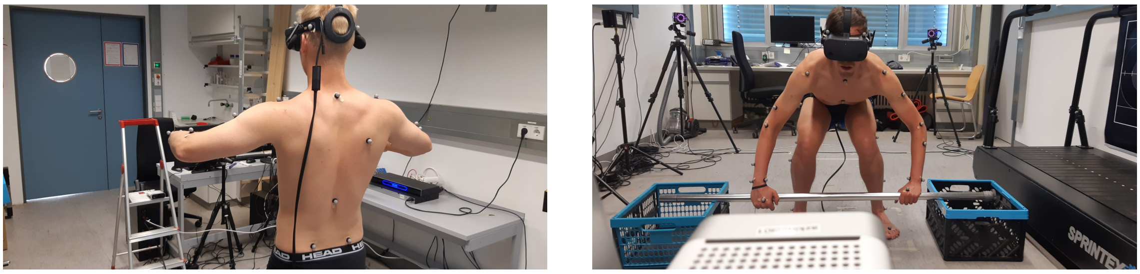

(a) Vicon Calibration. **(b)** Active Path Practice

Figure 5.1: Shows one participant being calibrated to the Vicon system and another currently attempting to assume a posture in path practice.

5.3.2 Home Study

For the home study, participants would visit a LiftVR setup in the authors apartment. This meant no body marker placement had to take place, shortening the overall procedure by 30 minutes. A participant of the LiftVR study would take between 30-60 minutes to complete it. This was dependent on their assigned group. Those without terminal feedback completed the study much faster, as they proceeded from one repetition to the next without delay. Some users would refuse to participate in the joint study due to finding the necessity to strip unpleasant. The widest variety, in terms of demographics, was achieved during the home study, as the joint study was mostly comprised of athletic individuals from INSPO.

5.4 Result Presentation and Analysis

In the following subsections, we analyze all data collected according to our methodology. To restate our intent, we are searching for individual and interaction effects of *Guidance System* and *Feedback Mode* during all training phases. Furthermore, our objective is to determine that LiftVR can effectively teach back-friendly lifting techniques irrespective of body weight, height, age, previous lifting or VR experience, and exercise levels. To do so, we perform individual and interaction analysis on our independent demographic variables. Following this, we present and analyze results from the questionnaires.

5.4.1 Omission of Back Error Measurements

During the study it became apparent that spine joint positions provided by the tracker were unreliable. Occasionally, participants standing perfectly still with straight backs would trigger back errors, whilst others with heavy curvature would not. As the camera is only provided a frontal view, the actual curve can not be observed. We suspect this induces a predisposition to incorrect joint positions based on a students anatomical build. For example, participants with broader shoulders, large chests, or “stocky” builds were observed to trigger back errors repeatedly when standing straight. In contrast, some thin or lean individuals would curve their backs by excessive amounts without being detected.

Practice Mode	Experience	Mean	Std. Deviation
Path	No Experience	975.63f	943.62f
Path	Experienced	401.19f	516.74
Slow Movement	No Experience	190.94f	172.74f
Slow Movement	Experienced	71.69f	66.04f
Regular Movement	No Experience	104.63f	85.00f
Regular Movement	Experienced	46.05f	33.38f

Table 5.2: Back Error Measurements.

The author noticed these inconsistencies quickly and attempted to adjust the system manually during the study, but to no avail. As such, the information retrieved is not deemed precise enough to evaluate guidance system or feedback mode effects. Instead, we only evaluate back errors in regards to demographic variables, to investigate what could be causing these inconsistencies. For this we use one-way MANOVA on back error measurements obtained during practice.

As our demographic inquisition was not intended for such detailed inspection of anatomy, we could reveal no anatomically based significant effects or general differences. The only non significant difference we could find, which is concordant with our own observations during the study, is that experienced lifters know to keep the back straight and thus have less back errors. During practice, students experienced with the motion would consistently score fewer back errors. The actual measurements to present the differences are in Table 5.2. Nonetheless, we find no statistical significance in a MANOVA over all back errors, prompting us not to investigate further.

5.4.2 Baseline Measurement

To form our baseline, we measured the variables *PerfA_Av_Err* and *PerfA_Av_FrameErr*. Participants have an average error of mean (M) = .163m, standard deviation (SD)= 0.037f and frame-error of (M = 70.58f, SD= 35.52f).

Demographic Interactions As no guidance system or feedback interactions are possible with the baseline measurement, we investigate demographic effects. The variables in question are *AgeGroup*, *Gender*, *LiftExperience*, *ExerciseGroup*, *VRExperience*, *HeightGroup* and *WeightGroup*. No effects on error are found for participant age, lifting experience, exercise frequency, VR experience, or height. Through one-way ANOVA we do find significant effects from weight ($F_{4,19} = 3.53, p = .026, \eta^2 = .426$) and gender ($F_{1,22} = 4.68, p = .042, \eta^2 = .175$) Figure 5.2 shows how females outperformed males and weight group 61-70kg beats all others.

We do not find effects on frame-error from age, gender, VR experience, height or weight. However, we do find significant effects from lift experience ($F_{1,22} = 7.19, p = .014, \eta^2 = .246$) and exercise group ($F_{3,20} = 4.16, p = .019, \eta^2 = .384$). With lift experience and increasing exercise frequency, participants frame-error decreases, as seen in Figure 5.2.

This confirms that the baseline measurements are mainly affected by participant experience and routine, not their anatomical features.

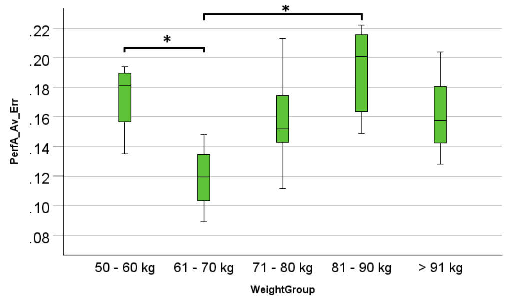

(a) Weight × Error.

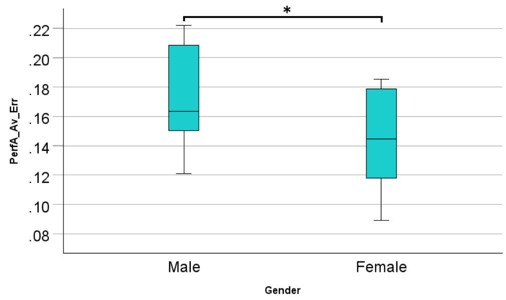

(b) Gender × Error.

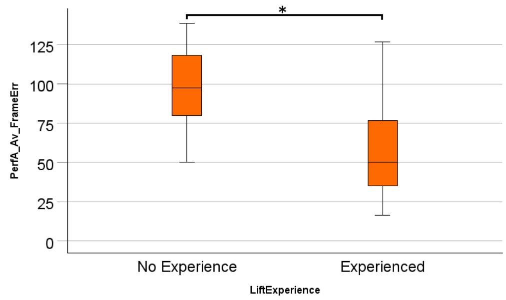

(c) Lift-Experience × Frame-Error.

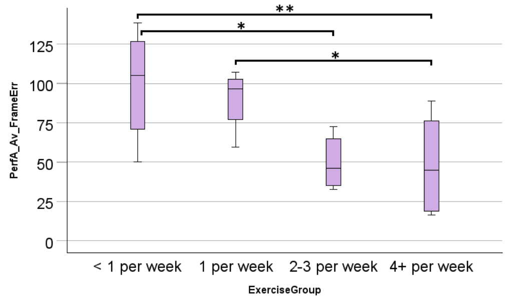

(d) Exercise × Frame-Error.

Figure 5.2: Demographic effects on baseline measurement.

5.4.3 Path Practice

In path practice we measured the variables *PracA_Av_FrameErr*, *PracA_Best*, *PracA_Total_Improve* and *PracA_First_Improve*. These were calculated from the individual frame-error measurements. Figure 5.3 shows the results for both guidance systems categorized by the feedback mode.

Guidance System *PracA_Av_FrameErr* failed the test for homogeneity of variances, which is why we perform Welch's ANOVA. We found a highly significant statistical effect ($F = 3.74$, $p = .005$, $\omega^2 = .069$). The Zone guidance system (M = 2013.97f, SD= 1004.83) allows students to complete path mode much quicker than Ghost-Skeleton (M = 5021.43f, SD=2974.99f).

This trend continues when investigating *PracA_Best*, which also reveals a highly significant effect ($F_{1,22} = 10.40$, $p = .004$, $\eta^2 = .321$). Zone (M = 592.08f, SD= 449.26f) students best attempts are superior to Ghost-Skeleton's (M = 1694.42f, SD= 1095.45f).

The effects on improvement are investigated next, and reveal no statistical significance. As we investigate the means for *PracA_Total_Improve* and *PracA_First_Improve*, we find that Zone outperforms Ghost-Skeleton in both. The respective total improvement is (M = 88.34%, SD= 6.51%) for Zone and (M = 78.72%, SD= 15.28%) for Ghost-Skeleton. Inspecting the first improvement reveals Zone (M = 53.77%, SD= 41.17%) and Ghost-Skeleton (M = 39.86%, SD= 49.43%). Consequently, we observe a consistent difference between Zone and Ghost-Skeleton in improvement.

In conclusion, Zone significantly outperforms Ghost-Skeleton in terms of frame-error. A consistent difference in improvement is present with Zone students improving more.

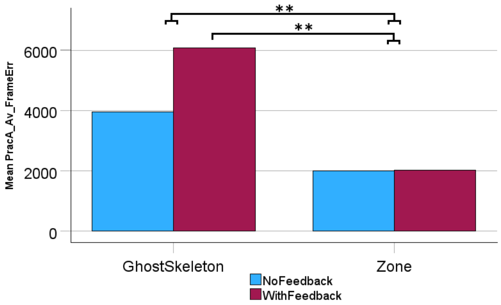

(a) Average Frame-Error.

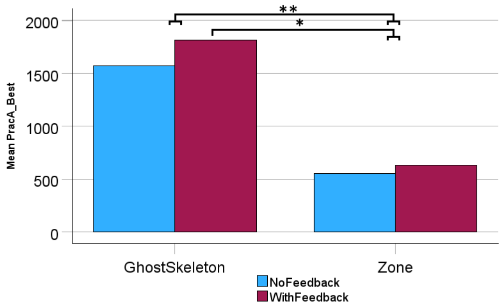

(b) Best Frame-Error.

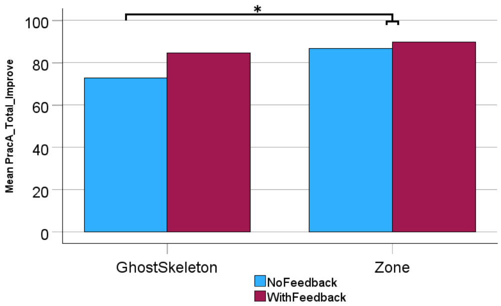

(c) Overall Improvement.

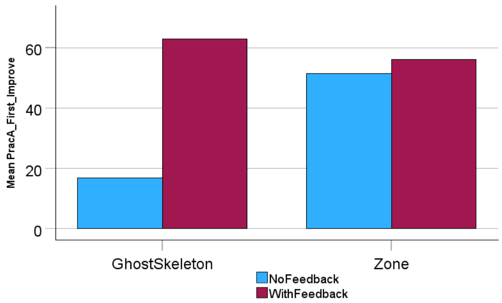

(d) First Improvement.

Figure 5.3: Path mode practice results.

Feedback Mode We find no significant effects from the feedback mode on path practice. Investigation of the means reveals that With-Feedback (M = 4052.62f, SD= 3345.68f) was beaten by No-Feedback (M = 2982.78f, SD= 1725.09f) on *PracA_Av_FrameErr*. This is consistent with the means of *PracA_Best*, as With-Feedback (M = 1221.67f, SD= 1292.77f) is outperformed by No-Feedback (M = 1064.83f, SD= 616.42f).

Improvement, however, indicates inverse differences. *PracA_Total_Improve* shows that With-Feedback (M = 87.24%, SD= 9.34%) beats No-Feedback (M = 79.75%, SD= 14.43%). Similarly the means for *PracA_First_Improve* are better for With-Feedback (M = 59.51%, SD= 46.90%) than No-Feedback (M = 34.14%, SD= 41.16%).

For our analysis we conclude that feedback creates no difference in overall or best performance, though mean differences indicate those with it are able to improve more and faster. However, the lack of statistical significance makes us not confident of this.

Guidance × Feedback The analysis of *PracA_Av_FrameErr* and *PracA_Best* fails the assumption of equality of error variances. As both show no significant interaction effects (($F_{1,20} = 1.42, p = .248, \eta^2 = 0.07$) and ($F_{1,20} = .054, p = .818, \eta^2 = .003$)), we move on to the improvement variables.

While *PracA_Total_Improve* and *PracA_First_Improve* do not fail error variance equality, both show no significant interaction effects (($F_{1,20} = .897, p = .355, \eta^2 = .043$) and ($F_{1,20} = 1.33, p = .263, \eta^2 = 0.062$)).

After we investigate the means we find that guidance systems do differ in total and first improvement values when performed with feedback. *PracA_Total_Improve* for Ghost-Skeleton has values No-Feedback (M = 72.74%, SD= 17.25%) and With-Feedback (M = 84.71%, SD= 11.44%). In comparison, Zone scores No-Feedback(M = 86.76%, SD= 6.51%) and With-Feedback (M = 89.84%, SD= 6.70%).

Ghost-Skeleton student's *PracA_First_Improve* has values No-Feedback (M = 16.79%, SD= 52.59%) and With-Feedback (M = 62.94%, SD= 36.52%). Zone has (M = 51.48%, SD= 15.49%) and With-Feedback (M = 56.51%, SD= 58.97%).

Though we have no findings we can claim confidence of, students using Ghost-Skeleton, the system we are confident is worse for performance, improved better with feedback than without.

Demographic Interactions The effects of all demographic variables on path practice and interactions with *Guidance System* and *Feedback Mode* are investigated next.

AgeGroup: Age has a significant effect on *PracA_Av_FrameErr* ($F_{2,21} = 4.20, p = .029, \eta^2 = .286$). Group *Age55-64* performs the worst (M = 6478.13f, SD= 4052.25f), group *Age25-34* the best (M = 2468.77f, SD= 1163.77f).

The interaction between *Guidance System* and age is not significant for *PracA_Av_FrameErr* ($F_{2,18} = 3.46, p = .053, \eta^2 = .278$). When inspecting the means for differences, we find all age groups perform better with Zone than Ghost-Skeleton. Most impressively, group *Age55-64* achieves M = 8789f with Ghost-Skeleton and M = 1846f using Zone.

No interaction effect is found with *Feedback Mode*, though mean differences reveal that group *Age55-64* performs worse with feedback than without. In general, all age groups deviate their frame-errors more with feedback than without.

Age affects *PracA_Best* significantly ($F_{2,21} = 5.64, p = .011, \eta^2 = .349$), with group *Age55-65* (M = 2635.33f, SD= 1913.95f) performing worse than *Age18-24* (M = 1099.29f, SD= 581.28f) and *Age25-34* (M = 845.50f, SD= 645.92f). Highly significant interaction is given both for *Guidance System* ($F_{2,18} = 6.77, p = .006, \eta^2 = .429$) and *Feedback Mode* ($F_{2,18} = 9.58, p = .001, \eta^2 = .515$). Groups *Age18-24* and *Age25-34* both achieve better bests with feedback, *Age55-64* does so without. In regards to the *Guidance System*, all age groups achieve better bests with Zone than Ghost-Skeleton.

Continuing our analysis, no individual or interaction effects are found for *PracA_Total_Improve* and *PracA_First_Improve*.

In conclusion, we reveal that group *Age55-64* accounts for our previous finding of *PracA_Best* being lower on students in the With-Feedback group. We can confidently state that for all other age groups, feedback clearly improves ones best performance. The same can not be said with confidence for *PracA_Av_FrameErr*, as *Age18-24* performed worse when in the With-Feedback group. When speaking about the *Guidance System*, we are confident that Zone is better for all ages.

Gender: No individual or interaction effects are found for gender, though mean differences reveal that females (M = 89.76%, SD= 8.03%) total improvement is higher than males (M = 79.80%, SD= 13.42%). This difference continues for the first improvement (females (M = 69.35%, SD= 27.05%), males (M = 33.31%, SD= 48.95%)).

Generally, females improvement is higher and more stable than males, though we lack the statistical significance to say so confidently.

LiftExperience: *PracA_Av_FrameErr* fails the homogeneity of variances, which leads us to perform Welch's ANOVA. Unlike the standard ANOVA ($F_{1,22} = 5.86, p = .024, \eta^2 = .210$), it shows no significance ($p = .10, \omega^2 = .111$).

We investigate differences in means to reveal those who are experienced with lifting motions (M = 2751.51f, SD= 1853.04f) are better than the inexperienced (M = 5378.66f, SD= 3501.62f). They also have better best performances (With-Experience (M = 859.24f, SD= 753.20f), No-Experience (M = 1745.57f, SD= 1294.15f)), though this finding is not statistically significant ($F_{1,22} = 4.12, p = .055, \eta^2 = .158$). There are no interaction effects in regards to performance and the improvement variables show no individual or interaction effects.

Though Welch's ANOVA was not significant, differences indicate that those experienced with similar motions perform better and deviate less.

ExcerciseGroup: Against our expectation those who exercise more do not produce better performances. No individual or interaction effects are found for all variables. Our post hoc tests reveal that the first improvement of group *4+PerWeek*(M = 77.99%, SD= 14.81%) is significantly higher than the group *<1PerWeek* (p = .028, M = 21.43%, SD= 43.35%).

When interpreting the differences in deviations and means, we can confidently claim that those who exercise more than four times a week improve vastly upon their second try, much more so than those who exercise less than once per week.

VRExperience: Having used VR before shows no individual or interaction effects. We conclude, with confidence, that it is entirely irrelevant to path practice performance in LiftVR.

HeightGroup: Participant height shows no individual or interaction effects, with exception of *Guidance System* and *PracA_Total_Improve* ($F_{2,17} = 3.73, p = .045, \eta^2 = .305$). Post hoc testing reveals no significant distinction between groups. Investigating mean differences shows that as height increases, frame-error decreases when using Ghost-Skeleton. Due to the lack of individual interactions and the aforementioned non distinct groups, we are confident that height is irrelevant to performance or improvement in path practice.

WeightGroup: Similarly to height, weight has no individual or interaction effects on performance or improvement. Differences in means and deviations show a trend that *PracA_Total_Improve* lowers and deviates higher when participants cross 80kg (*Under-80* (M = 87.02%, SD = 8.29%), *Over-80* (M = 76.69%, SD = 11.60%). As this is not statistically significant, we conclude that weight has no impact on path practice.

5.4.4 Slow Movement Practice

In slow movement practice we measured the variables *PracB_Av_Err*, *PracB_Err_Improve*, and *PracB_Best*. These are calculated from the individual error measurements. Figure 5.4 shows the results for guidance systems with and without feedback in addition to student averages as they improved attempt by attempt.

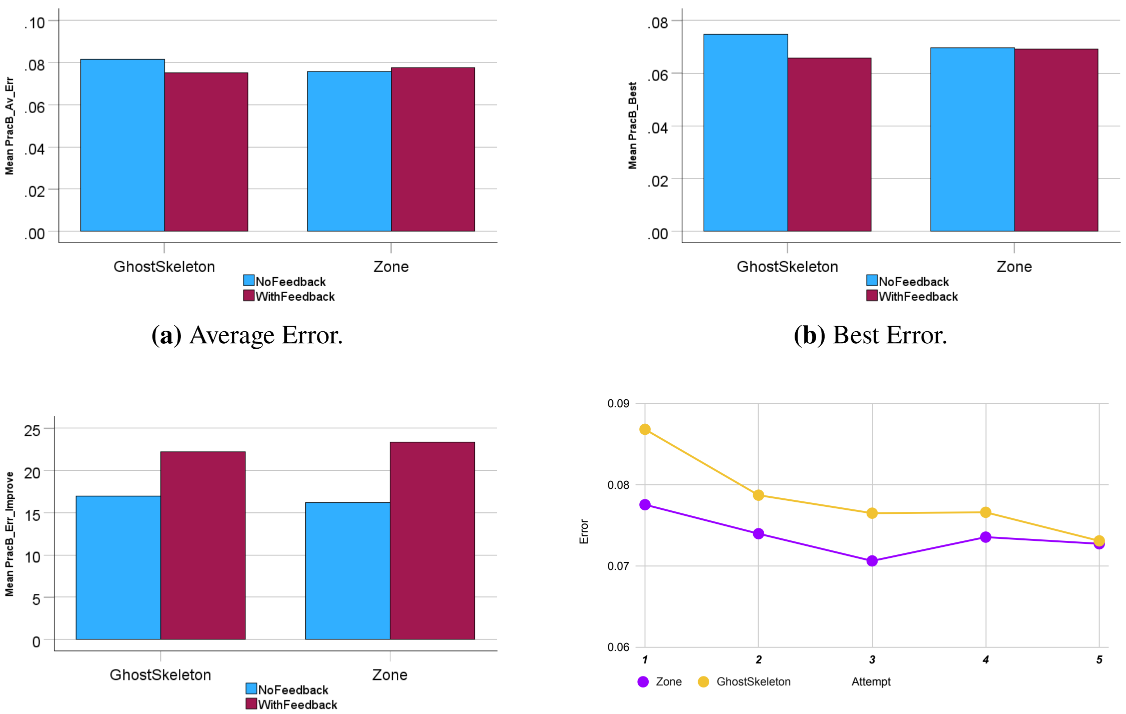

(a) Average Error.

(b) Best Error.

(c) Improvement.

(d) Error Progression.

Figure 5.4: Slow movement practice results.

Guidance System There are no statistically significant findings for all variables measured during slow movement practice in relation to *Guidance System*. When inspecting the differences of means, we discover that Zone students average lower errors for every attempt. This is clearly visible in Figure 5.4d. Nonetheless, we can not claim with confidence that the guidance system is affecting slow movement practice errors.

Feedback Mode Similarly to *Guidance System*, the *Feedback Mode* does not show significant effects in slow movement practice. Comparing improvement means, those with feedback (M = 22.76%, SD= 13.36%) beat those without (M = 16.61%, SD= 5.30%). *PracB_Av_Err* shows no differences, but *PracB_Best* does (With-Feedback (M = .067m, SD= .009m), No-Feedback (M = .072m, SD= .008m)).

In conclusion, the impact of feedback is not significant, yet creates differences in student improvement and best performances.

Guidance × Feedback Following the trend in slow movement practice, no interaction effects are found between *Guidance System* and *Feedback Mode*. We find differences when investigating the means of *PracB_Err_Improve*. Students of both guidance systems improve more with feedback (Ghost-Skeleton (M = 22.19%, SD= 16.10%), Zone (M = 23.34%, SD= 11.52%)) than without (Ghost-Skeleton (M = 17.00%, SD= 3.32%), Zone (M = 16.22%, SD= 7.13%)).

Though these differences are present, we must reject the notion that guidance systems and feedback interact in slow movement practice.

Demographic Interactions As with path mode practice, we check all demographic variables for individual or interaction effects on slow movement practice.

AgeGroup: Participant age has no significant individual or interaction effect on *PracB_Av_Err*. Post hoc analysis reveals that group *Age55-64* (M = .087m, SD= .008m) performs worse than group *Age25-34* (M = .074m, SD= .012m). The difference is not as severe compared to path practice. Further differences show that *PracB_Err_Improve* increases with age (*Age18-24* (M = 16.43%, SD= 2.84%), *Age25-34* (M = 20.28%, SD= 8.70%), *Age55-64* (M = 24.55%, SD= 25.63%)), but not significantly.

Looking at statistical significance, we can conclude age did not impact slow movement practice. Differences indicate that the older participants start out with higher errors, but are able to quickly close the gap regardless of feedback or guidance. The other age groups start closer to their bests and thus have less overall improvement.

Gender: Females (M = .064m, SD= .006m) achieved significantly better *PracB_Best* ($F_{1,22} = 6.04, p = .022, \eta^2 = .215$) than men (M = .072m, SD= .009m). *PracB_Err_Improve* and *PracB_Av_Err* show no significant differences. Nevertheless, our inspection of the means shows that females also outperform males in these categories. Interactions with guidance system or feedback are not found.

We conclude that gender has significant effects on slow movement practice, as we are confident females outperform men when comparing their best attempts. Differences in means indicate that this is also true for overall error and improvement, though not confidently so.

LiftExperience: No significant individual or interaction effects on slow movement practice from lifting experience are found.

When inspecting the differences of means, experienced lifters perform better. For *PracB_Av_Err* Experience (M = .076m, SD= .009m) beats No-Experience (M = .081m, SD= .017m). The same is true for *PracB_Best* (Experience (M = .068m, SD= .007m), No-Experience (M = .073m, SD= .013m)). Improvement shows no differences between the groups. Those with no experience differ depending on their guidance system, with Zone (M = 21.69%, SD= 13.24%) increasing improvement over Ghost-Skeleton (M = 15.41%, SD= 5.64.70%).

In concordance with our findings from path practice, differences indicate that those experienced with similar motions perform better, though not significantly so.

ExcerciseGroup: Post hoc tests reveal that, for *PracB_Av_Err*, those in group *2-3PerWeek* (M = .072m, SD= .007m) significantly outperform (p = .032) group *<1PerWeek* (M = .086m, SD= .015m). They are also beat by the group *4+PerWeek* (M = .073m, SD= .008m), though not significantly. This trend continues with the best attempts (*<1PerWeek* (M = .077m, SD= .011m) versus *2-3PerWeek* (p = .024, M = .066m, SD= .007m) and *4+PerWeek*(p = .056, M = .067m, SD= .008m)). Improvement has no statistically significant findings, though inspection of the means reveals that those who exercise less improve more, as others start closer to their best performances. No interaction effects with the guidance system or feedback are found.

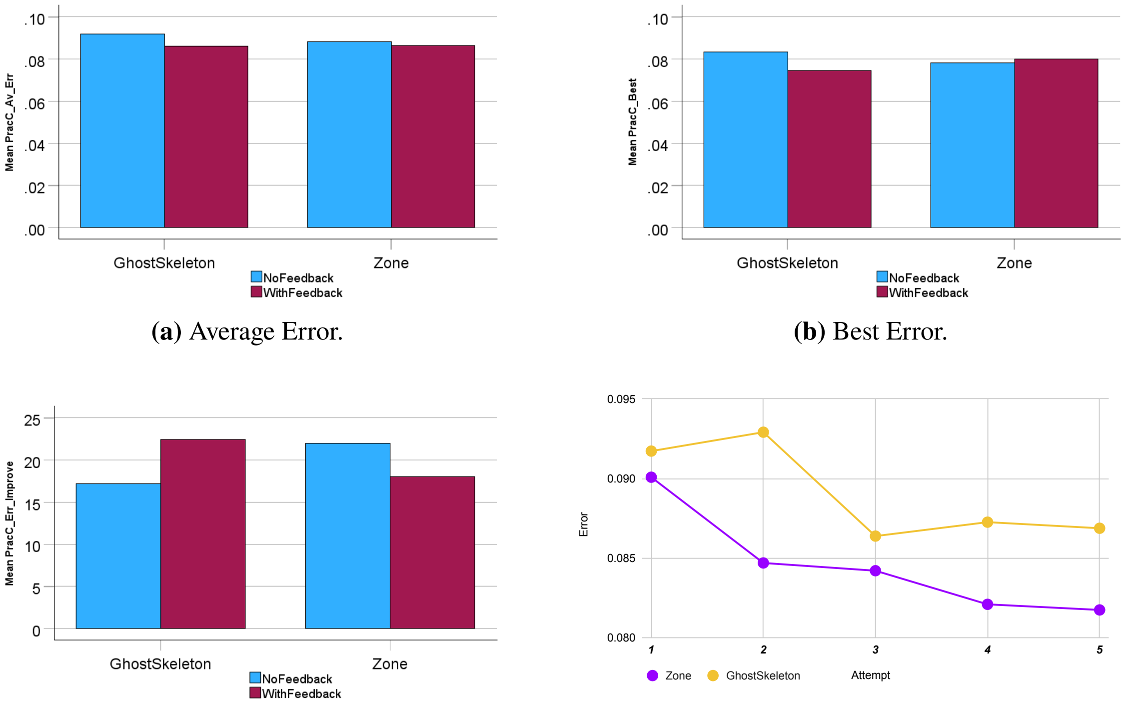

(a) Average Error.

(b) Best Error.

(c) Improvement.

(d) Error Progression.

Figure 5.5: Regular movement practice results.

We can confidently conclude that, in slow movement practice, those who exercise two to three times per week will have less error those who exercise less than once. Differences of means indicate that more exercise generally leads to lower errors and less improvement, as one starts closer to their best performance, though we cannot say so with confidence.

VRExperience, HeightGroup, WeightGroup: No individual or interaction effects based on any of the titled groups are found for slow practice mode. With the exception of gender and age, no anatomical variables have an impact on LiftVR's slow movement practice.

5.4.5 Regular Movement Practice

For regular path movement practice we measured the variables *PracC_Av_Err*, *PracC_Err_Improve*, and *PracC_Best*. As with slow movement practice, we calculate these from the individual error measurements. Figure 5.5 shows the results for guidance systems with and without feedback in addition to student averages of each individual attempt.

Guidance System Consistent to slow movement practice, regular movement practice yields no significant findings for Guidance System. Following the trend, differences in means show that Zone outperforms Ghost-Skeleton for each individual attempt in student progression (Figure 5.5d). The lack of statistical significance prevents us from claiming with confidence that the guidance system has an effect.

Feedback Mode Unlike slow movement practice, *Feedback Mode* does not show differences in means in regular movement practice. Our explanation of this observation is regular movement practice's increase in speed and last position in the study. Having mastered the motion in slow movement practice, improvement is then only achieved in regards to overcoming the initial shock from speed change. This explains the steep improvement between attempts one and two of Zone or two and three of Ghost-Skeleton in Figure 5.5d, which is not present in Figure 5.4d. Though this seems intuitive, the lack of significant results means we are not confident this is the full explanation.

Guidance × Feedback Significant interactions from *Guidance System* and *Feedback Mode* are not found. Post hoc testing and mean analysis yield no further results.

Demographic Interactions We continue our investigation of demographic interactions through individual and interaction analysis for regular movement practice.

AgeGroup: Whilst group *Age55-64* (M = .095m, SD= .008m) does perform worse than the younger groups (M = .088m, SD= .012m), the effect is not significant. This trend continues for *PracC_Best* with the same non significant result. Only for *PracC_Err_Improve* was no difference noticeable between groups. Interaction effects are none to be found.

In conclusion, we are confident that age has no significant effect on regular movement practice. Differences in means indicate older participants have slightly higher errors and worse bests.

Gender: In regards to the error, there are noticeable differences between the groups. Females (M = .082m, SD= .012m) still beat males average error (M = .092m, SD= .011m), however, not significantly so ($F_{1,22} = 3.95, p = .060, \eta^2 = .152$). In contrast, females best attempts (M = .072m, SD= .009m) significantly ($F_{1,22} = 5.18, p = .033, \eta^2 = .190$). outperform males (M = .083m, SD= .012m). Improvement shows no significant effects or indicative means. Furthermore, significant interactions with the *Guidance System* or *Feedback Mode* are not found. Analysis of the means did show that females with feedback have better means than those without, whilst men show no difference.

As we now have conclusive evidence of females producing better bests and overall errors in both slow and regular practice, we are confident that they are overall better at movement practice. Our investigation of the means indicates that this may stem from them utilizing feedback more (something we did observe during the study). However, we are not confident this is true, as we did not record how often participants used certain feedback features and thus cannot show it statistically.

LiftExperience: *PracC_Av_Err* is not significantly affected by lift experience ($F_{1,22} = 2.87, p = .104, \eta^2 = .115$). Investigation of the means does reveal that Experience (M = .086m, SD= .01m) leads to less error than No-Experience (M = .094m, SD= .015m)). To the contrary, *PracC_Best* did prove significantly affected ($F_{1,22} = 5.18, p = .033, \eta^2 = .191$) with Experience (M = .076m, SD= .009m) beating No-Experience (M = .087m, SD= .015m)).

The only difference measured in interactions was from *Guidance System*, though not significant ($F_{1,20} = 2.07, p = .166, \eta^2 = .094$). Those in No-Experience improve more with Zone (M = 22.77%, SD= 9.60%) than with Ghost-Skeleton (M = 13.97%, SD= 3.60%). These findings roughly coincide with those of slow movement practice and indicate that Zone helps inexperienced lifters improve more during movement practice, though not confidently so.

ExcerciseGroup: The effect of different amounts of exercise is not significant on *PracC_Av_Err* ($F_{3,20} = 2.13, p = .128, \eta^2 = .243$). Post hoc tests reveal that group *<1PerWek* (M = .098m, SD= .013m) has significantly higher errors than *2-3PerWeek* (p = 0.29, M = .084m, SD= .008m). Though not significant, *4+PerWeek* is not far off (M = .085m, SD= .019m).

The effect on *PracC_Best* is now significant ($F_{3,20} = 2.87, p = .104, \eta^2 = .115$). Post hoc tests reveal significant differences between *<1PerWeek* (M = .091m, SD= .012m), *1PerWeek* (p = .033, M = .076m, SD= .012m), *2-3PerWeek* (p = .008, M = .075m, SD= .008m), and *4+PerWeek* (p = .033, M = .075m, SD= .009m).

No significant interaction effects are found, except for *PracC_Err_Improve* with *Guidance System* ($F_{3,16} = 3.51, p = .040, \eta^2 = .397$). Careful analysis of this significance reveals that exercise groups "take turns" at which guidance system they improve more in. For example, those who exercise less than once per week excel with Zone over Ghost-Skeleton, however, the opposite is true for once per week. As this is not indicative of anything, we draw no conclusions from it.

We have found conclusive evidence that exercise groups differ selectively for error and best performances, which leads us to be confident of an overall effect on movement practice.

VRExperience, HeightGroup, WeightGroup: Coinciding with our findings from slow movement practice, no individual or interaction effects for these groups are found.

5.4.6 Final Measurement and Baseline Comparison

From the students final performances, we measured *PerfB_Av_Err* and *PerfB_Av_FrameErr*. Overall improvements are calculated as *Overall_Err_Improve* and *Overall_Frame_Improve*. Figures for these variables can be seen in Figure 5.6. Only overall improvement is considered for effect and interaction analysis. This is due to students naturally gravitating towards their original performances without active motion guidance. Consequently, it is only important how they compare to their baseline measurements.

We conduct a one-sided paired-sample t test to investigate if LiftVR practice had a significant effect on student performance. For this, we pair *PerfA_Av_Err* with *PerfB_Av_Err* and *PerfA_Av_FrameErr* to *PerfB_Av_FrameErr*. We achieve significance for both the error ($t(23) = 3.21, p = .002, \eta^2 = .115$) and the frame-error ($t(23) = 2.18, p = .020$). Therefore, we are confident that LiftVR ca teach students the motion effectively.

Guidance System Effects based on the guidance system are not significant. Mean inspection reveals differences for error and frame-error improvements. Zone beats Ghost-Skeleton for error improvement ((M = .032m, SD= .039m) versus (M = .025m, SD= .048m)) and frame-error

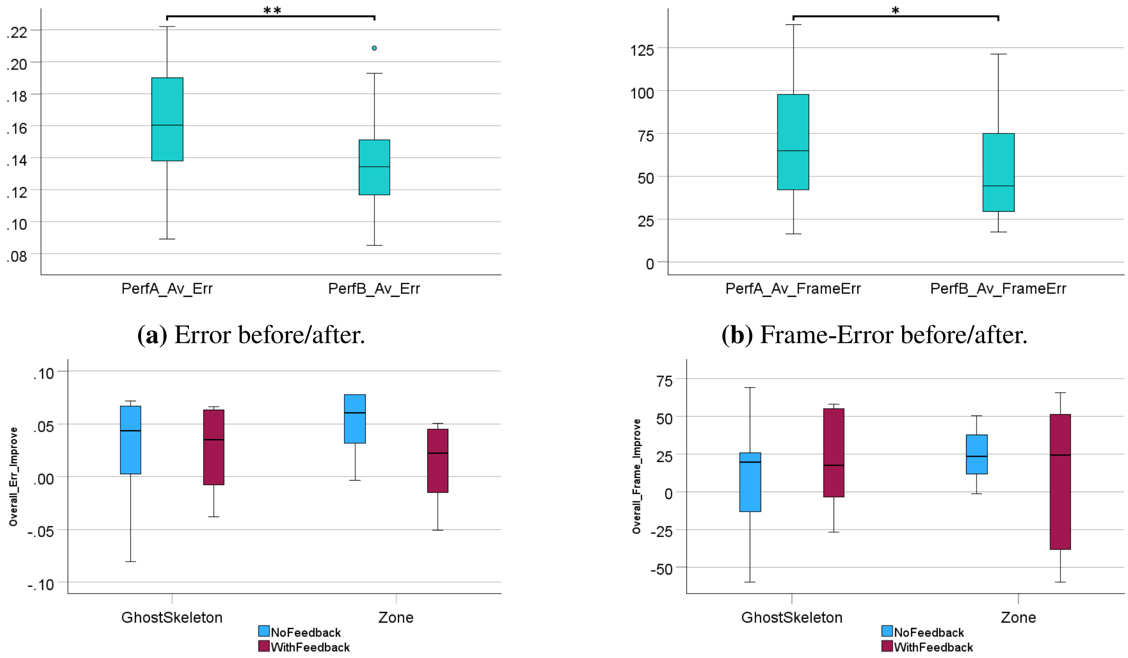

(a) Error before/after.

(b) Frame-Error before/after.

(c) Error improvement.

(d) Frame-Error improvement.

Figure 5.6: Path mode practice results.

improvement ((M = 17.75f, SD= 37.68f) versus (M = 14.86f, SD= 37.21f)). Though we lack statistical confidence, the differences indicate that students learn the motion more effectively with Zone than Ghost-Skeleton.

Feedback Mode Similarly to guidance systems, feedback effect on overall improvement is not statistically significant. Investigating the means reveals that those without feedback have higher error improvements((M = .038m, SD= .047m) versus (M = .019m, SD= .039m)). This continues for frame-error improvement, though less severe ((M = 17.22f, SD= 32.97f) versus (M = 15.39f, SD= 41.48)). This indicates feedback has a negative effect on overall improvement, though we cannot claim this with confidence.

Guidance x Feedback Whilst investigating interaction effects, no statistically significant findings are made. Investigation of the means, however, discovers that the unexpected discrepancy of students with feedback performing worse was exclusive for Zone. The frame-error deviation of Zone with feedback is tremendous (M = 11.17f, SD= 51.28f), though less severe for standard error. Nonetheless, this coincides with our general observation that feedback tends to increase deviations of most measurements.

Demographic Interactions In line with our previous analysis we investigate individual and interaction effects on overall improvement from demographic variables.

AgeGroup: Unlike in practice investigations, age plays no part on overall improvement and no indicative differences from means are found. Furthermore, the guidance system or feedback does not interact significantly with age. In conclusion, we can confidently claim that LiftVR has participants of all ages improve equally.

Gender: Contrary to practice mode results, females do not outperform males in overall improvement. Interaction effects are also not present. Consequently, females and males improved equally, regardless of the guidance system or feedback mode.

LiftExperience: We find a significant effect of lift experience on overall frame-error improvement ($F_{1,22} = 4.91, p = .037, \eta^2 = .182$). Those with experience barely improve their frame-error (M = 6.47f, SD= 36.93f) compared to those with no experience (M = 40.19f, SD= 23.94f). Overall error is not significant between experience groups ($F_{1,22} = 3.73, p = .066, \eta^2 = .145$). Though differences of means reveal experienced students improve less with high deviation (Experience (M = .018m, SD= .047m), No-Experience (M = .053m, SD= .016m)).

The guidance system or feedback mode has no interaction effect with lifting experience. This allows us to confidently conclude that those who have previously performed lifting motions barely improve their frame-error, regardless of the guidance system or feedback. Means indicate this is also the case for general error.

ExcerciseGroup: The exercise levels of students has no significant effect on overall frame-error or error improvement. Interaction effects are not found. However, post hoc analysis reveals interesting differences.

Looking at frame-error improvement, group *4+PerWeek* deteriorates (M = -2.55f, SD= 46.12f). and *2-3PerWeek* barely improves (M = 9.71f, SD= 36.76f). The means rise and deviation sinks for groups *1PerWeek*(M = 37.92f, SD= 24.90f) and *<1PerWeek* (M = 29.55f, SD= 26.59f).

The trend continues for error (*4+PerWeek* (M = .007m, SD= .071), *2-3PerWeek* (M = .023m, SD= .028m), *1PerWeek* (M = .038m, SD= .036), *<1PerWeek* (M = .050m, SD= .020m)).

The trends in differences between exercise groups continue from the baseline measurement and slow/regular practice. This indicates, without statistically given confidence, that those who exercise more frequently improve less overall. Deviation rises with the exercise frequency.

VRExperience, HeightGroup, WeightGroup None of these demographic variables interact individually or as interactions in regards to overall improvement. We can confidently conclude that previous VR experience or participant height and weight is irrelevant to improve with LiftVR.

5.4.7 Questionnaire

In this subsection we discuss results from the questionnaire. We asked students of LiftVR to rate their guidance system based on understandability, mental effortlessness, and helpfulness. These were combined to form a total score. Furthermore, we asked students with feedback how they liked the systems and those without how and where they would imagine feedback.

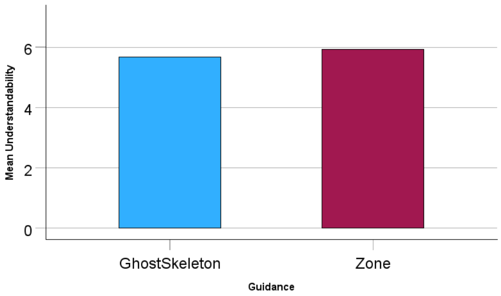

(a) Understandability.

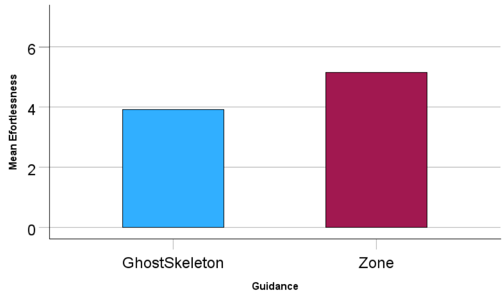

(b) Mental Effortlessness.

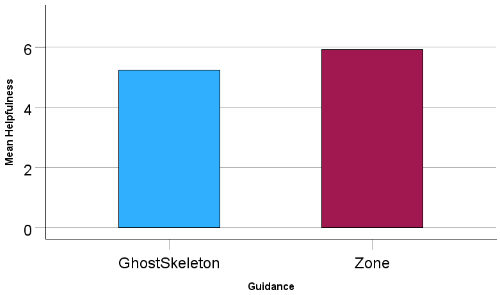

(c) Helpfulness.

(d) Overall Score.

Figure 5.7: Questionnaire results regarding the guidance systems.

Quantitative Results - Guidance System Students found both guidance systems easy to understand and helpful in correcting their mistakes. Ghost-Skeleton was deemed to be somewhat mentally demanding , whilst Zone was judged mostly effortless. The results from the guidance system questionnaire can be seen in Figure 5.7.

Beginning our analysis with understandability, we asked students how quickly they understood how their guidance system works. Ghost-Skeleton students rate their system with (M = 5.67, SD= 1.37), Zone students (M = 5.92, SD= 1.00). We find no significant effect, rejecting a confident stance that Zone is understood significantly better than Ghost-Skeleton.

Continuing with mental effortlessness, we do find a significant guidance system effect ($F_{1,22} = 5.08, p = .034, \eta^2 = .188$). Zone students rate their system with (M = 5.17, SD= 1.33) compared to Ghost-Skeleton's (M = 3.92, SD= 1.38). We can confidently say that Zone requires less mental effort to follow during performances.

Next we investigate how helpful the systems are in correcting errors and mistakes. Whilst there is a difference in means (Zone (M = 5.92, SD= .793), Ghost-Skeleton (M = 5.25, SD= 1.22)), we found no significant effect. Consequently, we can reject with confidence that Zone was more helpful in correcting students during the motion.

Finally, the overall score reveals no significant differences between Zone and Ghost-Skeleton ($F_{1,22} = 4.25, p = .051, \eta^2 = .162$). Nonetheless we can report a difference in means consistent with previous findings (Zone (M = 17.00, SD= 2.22), Ghost-Skeleton (M = 14.83, SD= 2.89)).

In conclusion, Zone was significantly better than Ghost-Skeleton in terms of mental effort and consistently outperformed the latter in all other categories.

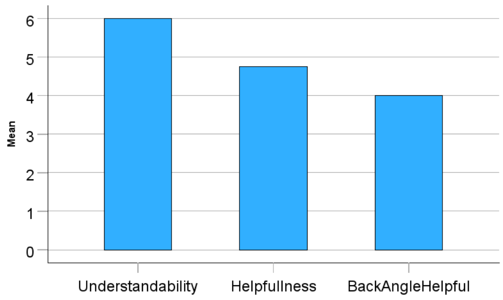

(a) Feedback-Students

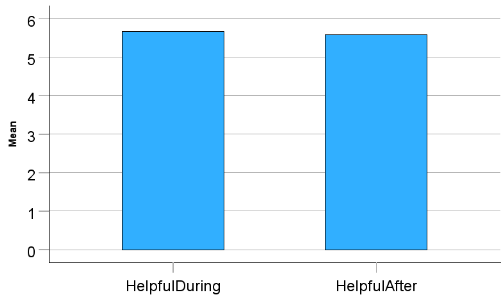

(b) No-Feedback-Students

Figure 5.8: Questionnaire results regarding the feedback modes.

Quantitative Results - Feedback Mode We asked students with terminal feedback to judge the systems for understandability and helpfulness. As the back-angle indicator was the most novel system, we asked separately if they found it helpful. The results reveal that participants found the feedback easy to understand and mostly helpful. We find that the back-angle indicator's helpfulness was perceived neutrally. The results can be seen in Figure 5.8a.

Participants without feedback were asked if they would prefer feedback during or after the motion. Figure 5.8b shows that they were indifferent to the idea.

Qualitative Results We asked all students to reveal what they liked or disliked about LiftVR and which additional guidance systems or mechanisms they could imagine. Furthermore, we asked terminal feedback students to propose additional feedback systems. Those in the No-Feedback group were asked to generally imagine feedback systems instead.

The most common positive feedback for LiftVR is the system's intuitive and easy to understand design. Participants like how joints and bones are visualized distinct from one another and find the feedback system good and precise. Superimposing the guidance system onto the *Student Skeleton* is received well. One praises LiftVR for the low latency of *Student Skeleton*. Three students say they like the Zone system, one calling it "cool". The color and foot location indicators are deemed helpful. One student stated they enjoyed the virtual gym environment.

Negative feedback was less plentiful, with most complaining that the linear interpolation of Zones color indicators is too weak. For example, Zone's capsules are slightly yellow when joints are nearly in position, which confuses participants into thinking the joints are already correctly positioned. One complained that it is hard to deduce corrections or current errors from Ghost-Skeleton. The foot indicators are deemed to be too sensitive. Another feedback was disappointment from the student skeleton calibration not ensuring 100% overlap accuracy. Finally, path practice was deemed difficult, especially to start out with, and the VR headset caused some students to sweat profusely.

When asked to imagine additional non-terminal guidance systems or improvements to the existing ones, most suggest implementing multi-3PP. As we already know this to be an improvement over LiftVR's current perspective options, this feedback was expected. Another expected feedback was the wish for a 3D body model, though only one participant expressed this desire. Two participants wish to view themselves from over-the-shoulder rather than frontally. Four of the participants suggest an additional guidance system to help them identify different time points in the motion.

One suggests a metronome that ticks continuously throughout the motion. Another would like a continuous sound, that changes pitch dynamically to indicate upward and downward motion. For example, the pitch lowers when squatting to grab the bar, and rises when the bar is moved upward. One wishes for a human voice to give them general instructions such as “now we go down to grab the bar” and “okay, now we rise again”. In general, additional auditory guidance was the most requested feature. For visual guidance, one suggestion is to only show the expert’s joint path lines. Another imagines arrows which point the joints towards the target positions. Only one participant suggested haptic feedback. They wish for an exoskeleton to force them into position.

When asking the With-Feedback group about ideas for additional systems, most expressed their wish to be notified when joints are “grossly” out of position. Two suggest arrows to pop up and point at the incorrect joints. Another two wish for the arrows to point from the triggering joint to its intended position. One wishes for brightly flashing colors on said joints. The second most expressed desire was the integration of auditory feedback into the systems. Two wish for a human voice to tell them which joints are currently wrong and how to correct them. For example the voice could say something in the lines of “Your right hand needs to go a bit lower and to the left”. One wishes for less precise feedback in the form of an error sound. A more novel system requested by one participant involves active encouragement. A human voice should routinely play lines such as “you can do it” and “you are doing great” during the performance. The wish for more extensive gamification is expressed by another, though they explain not how that would be achieved.

Finally we asked the No-Feedback group to image feedback systems that would help them. We expected them to name some of the systems already implemented by LiftVR, which did indeed happen. Three wish for a replay functionality, which is already present. However, two of them wish the ability to jump towards “key moments” in which they performed worst. Two want a list of their worst performing joints, a feature LiftVR does have. LiftVR lacks the wish of three participants, who request correction suggestions for their worst performing joints. For example, upon seeing their neck to be their worst joint, the system should say “you are leaning too far back during the lift, causing your neck to perform poorly”, if this was indeed the cause. One wishes for haptic feedback in the form of vibrating devices on joints which are currently not in position. Similarly to the With-Feedback group, two wish a human voice to tell them how to correct their incorrect joints during the motion.

Having completed our evaluation, we now discuss our findings in relation to the research objectives.

6 Discussion

In our penultimate chapter we discuss our findings from the results and analysis. With it, we aim to gain a deeper understanding of LiftVR's effects on lifting techniques, differences between its two concurrent guidance systems and how feedback affects student performance. Based on our findings, we suggest additional design considerations for future full body motion guidance systems and their respective studies.

6.1 RQ1: Effect of LiftVR on lifting techniques

Our first research question addresses the effect of accessible full body motion guidance in VR on students lifting techniques. To answer the overall question, we subdivided it into three separate items designed to address the systems validity and effect. Subsequently, we form our final conclusion based on the individual results.

Is the system accessible and provides full body motion guidance?

LiftVR requires no attached body markers, tracking suits, or base stations. The headset and camera configuration used for the study (HP Reverb + Azure Kinect DK) can, at the time of this writing, be replicated for a cost of roughly 600€ (buying used hardware). It is possible to cut the total cost even further by using a modern standalone headset such as the Oculus Go.

Due to LiftVR's corrective systems, precise camera placement is not required. We were able to get LiftVR running and functional in a new location in less than fifteen minutes, with most of the time being spent setting up the desktop. This enabled us to perform the study effortlessly from the home setup or at the INSPO motion capture laboratory.

LiftVR requires only the joint data provided by the Azure Kinect and does not rely on its many other features. This eases a potential transition to OpenPose, which works with any webcam, drastically reducing cost.

Even though LiftVR uses lower accuracy accessible hardware, students did not complain about a lack of tracking accuracy or visual quality. To the contrary, visualizations of the student, guidance systems, and environment were praised. A result of LiftVR's smoothing and interpolation algorithm, *Student Skeleton's* smooth low latency visualization was commended.

In conclusion, LiftVR is accessible in the field of full body motion guidance and does not impact user satisfaction doing so.

Do we find an effect on students lifting techniques?

We received motion capture data of an expert from INSPO performing a back-friendly lifting technique. Students were tasked to learn this motion using LiftVR and our analysis reveals that they improved significantly compared to their baseline performances. Consequently, LiftVR shapes their lifting technique to match the expert's.

We find the effect to be weak for students who exercise regularly or are familiar with healthy lifting techniques. Our hypothesis is that they are accustomed to their own lifting motion, making it difficult to adjust them towards a new one. In our short-term retention test, most reverted back to their previous lifting habits.

Nonetheless, overall improvement was statistically significant, which lets us claim that LiftVR has a positive effect on students lifting techniques.

Does the system provide its effects regardless of student age, height, or weight?

Merely providing a positive effect and being accessible is insufficient. LiftVR must provide its effects irrespective of the users age, height or weight. Our analysis reveals that the aforementioned demographic denominators are irrelevant to the positive learning effect. Therefore, we claim with confidence that LiftVR can be used regardless of user height, weight, or age.

Conclusion
In summary, LiftVR is an accessible and confirmed effective system for full-body motion guidance, providing users the opportunity to improve their lifting techniques. Its low cost, easy setup, and compatibility with different users makes it a promising tool for promoting back-friendly lifting to a wide-spread audience.

6.2 RQ2: Differences between concurrent guidance systems

We quantitatively evaluated two concurrent guidance systems to study the effects of different expert visualizations on student performance. Ghost-Skeleton has been built to implement the most dominant design principles and findings currently found in the domain of full-body motion guidance. Zone was designed specifically with a lifting motion in mind, gaining advantages therefrom. Through our analysis of students path practice, movement practice, and overall performances, we found multiple effects.

Practice Effects
The choice of implementation has an immense effect on path guidance performance. Zone students assume postures faster and produce better bests. This is due to Zone taking advantage of the symmetrical nature of the motion. Whilst Ghost-Skeleton students must mentally keep track of eight individual indicators or joints, Zone merely visualizes five. As a result, the system is more

tolerant of when to accept the posture of being assumed, naturally lowering the difficulty of Zone. Furthermore, visual clutter is greatly reduced. It is unsurprising that participants reported higher mental effort when using Ghost-Skeleton, as they must simply pay attention to more things.

The effect on movement guidance is not as conclusive. What we can claim with confidence is that Zone's more abstract visualizations do not lead to worse performance. Though it is more tolerant than Ghost-Skeleton in path guidance, no adverse effect is observed during movement practice. We hypothesize that finding significant differences between systems in movement guidance is less likely, as correct timing is the most crucial aspect of it. Improving mostly stems from being able to adjust to the speed, rather than incorrectly positioned joints. This is especially noticeable when students transition from slow to regular speed. Error sharply rises due to students being shocked by the unfamiliar speed, but is rapidly corrected in just one or two attempts. Nonetheless, Zone student performances produce lower means than Ghost-Skeleton in every attempt and speed. The difference, however, is not significant, which suggests a higher powered study is required to confirm or dismiss these indications.

In regards to overall improvement, our study confirms that a guidance system can forgo humanoid expert visualizations without causing a negative effect on learning. Zone's more abstract visualizations grant the same benefit to a student's lifting technique as Ghost-Skeleton. Similarly to our practice results, mean analysis indicates that Zone students most likely improve more, though this remains to be evaluated further.

Conclusion

Our work shows that the effectiveness of separate concurrent guidance systems can differ significantly based on the context they are used in. More specifically, the choice of implementation has a profound effect on posture focused practice. Differences are less pronounced on when performing movement guidance.

Based on our aforementioned findings, we are confident Zone is the better overall system. Zone achieved superior results in eleven of our twelve computed dependent variables. This consistent difference between the systems leads us to hypothesize that future higher powered studies will confirm these effects statistically. Furthermore, whilst students praised Zone in the qualitative evaluation, no such feedback was received for Ghost-Skeleton.

6.3 RQ3: Performance affection from terminal feedback systems

We expected terminal feedback to have a statistically significant effect on performance based on findings by Nabil et al. [NHA+20]. Whilst differences are present, we must dismiss that these are statistically significant individually or in interaction with the concurrent guidance systems. Our findings are more in line of Elsayed et al. [EKS+22], who found that concurrent feedback showed no effects on student performance. Unlike the dependent variable differences between Zone and Ghost-Skeleton, mean divergences of the two feedback groups are inconclusive. During practice modes, we did find that those with terminal feedback improved more, though not significantly so. A contradiction is found when inspecting the total improvement, which is worse for students using the feedback systems.

Our interpretation of these results is that terminal feedback loses its statistically significant effect when used contemporaneously to a concurrent guidance system. When both are present, the latter primarily affects the students performance, whilst terminal feedback has minimal to non existent effects. This would explain why Nabil et al. [NHA+20] did find significant effects, as their system is solely reliant on terminal feedback.

Another hypothesis for the lacking effect of terminal feedback is the category of terminal feedback employed. LiftVR's systems are all classified as descriptive feedback, thereby errors are identified and presented to the user. Prescriptive terminal feedback would, however, tell students how to correct these identified errors. Silverman et al. [STK92] discovered that a mixture of descriptive, prescriptive, and corrective feedback yields the best results. During the study, we observed that some students would diligently use the terminal feedback systems, whilst others did not. Creating a self-prescription based solely on a description requires effort, which some students appeared unwilling to invest. This coincides with our qualitative results of students wishing for the system to actively **tell them what to do**, especially with audio. As such, some students may need some form of authority giving them prescriptive feedback and encouragement to improve faster.

Based on our findings, we draw the conclusion that in systems combining concurrent guidance and terminal feedback, the concurrent guidance holds greater importance. This is due to terminal feedback showing no statistically significant effect on performance when used in combination with concurrent guidance. We hypothesize that descriptive feedback being effective is highly dependent on the user, as some do not bother to expend the effort to evaluate how they could improve themselves. For these students, corrective feedback in the form of an encouraging, yet authoritative tutor could show significant effects on their performances.

6.4 Implications for future design and studies

Based on our results, we formulate additional system and study design suggestions for future work in the domain of full body motion guidance.

Concurrent Guidance System Design We have shown that the current standard in the full-body motion guidance domain of visualizing experts in a humanoid form offers less performance compared to abstract visualizations tailored to specific motions. Especially for posture teaching systems, the choice of guidance system dramatically affects student performance. Being tolerant on joint positions and inducing less visual clutter has positive effects. If motions contain symmetry, pairing joints as Zone does should be considered.

Color indicators on out of position joints should use clear, yet distinguishable interpolation to be easily differentiable. Guidance systems which only appear when students are grossly off target should be implemented. These could be arrows highlighting the responsible joint or a general audio warning. If timing is a critical component of the motion, one should include audio or visual guidance mechanisms to indicate different key points in the movement.

Feedback System Design Descriptive terminal feedback systems show little to no effect when a well-designed concurrent guidance system is present. Furthermore, the effectiveness of descriptive feedback is dependent on the students personality. It has been shown that a combination of descriptive, prescriptive and corrective feedback is best.

Regarding prescriptive feedback, students wish for a “mentor” or “teacher” system. An example of this would be having a human voice give terminal or concurrent feedback. This system can also give encouragement or interact with gamification systems. Additionally, students express a desire to access a replay featuring annotated key frames, which highlight the weakest segments of their performance.

Study Design Based on the results of our demographic analysis, those of age 25 to 30 with no previous experience of the motion and an exercise frequency of once or less per week improve the most and show the lowest deviations. Users having previous experience with a motion has a significant effect on their interaction with the guidance system teaching it, which must be considered when choosing study participants.

6.5 Limitations and Challenges

Due to limits in scope and time, not all design principles found from related work could be implemented into LiftVR. Students could freely choose between two different 3PP’s, which may have had an impact on the studies results. For this we rely on the findings by Hülsmann et al. [HFS+19] that side and frontal views affect performances similarly. Motion classification and authoring systems which would have drastically improved LiftVR’s feature set and allow more precise feedback could not be included.

As only descriptive terminal feedback has been implemented, the effects of prescriptive or corrective feedback systems are still unknown. Based on the work by Silverman et al. [STK92], we assume these systems will increase the performance differences we found between students with and without feedback, but are uncertain if this will push the effect to statistical significance.

As we investigated a variety of different demographic groups, our sample sizes of the individual levels are small. In general, our study lacks high power due to limited participants, which may have led to type I errors. Furthermore, we discovered that those who are familiar with the motion or exercise frequently significantly impact the apparent effectiveness of a guidance system. As such, these groups may have caused individual or interaction effects to appear not significant, though they would be for the general population.

7 Conclusion and Outlook

In our final chapter we present a summary of our findings and the contributions they bring. To conclude the thesis, we give an outlook of potential improvements to LiftVR and future work in the full body guidance domain.

7.1 Summary

We developed LiftVR, an accessible full body motion guidance system to investigate the effects of VR based training on lifting techniques. Furthermore, we implemented two competing guidance systems to evaluate the effects of different designs on student performance. Finally, we created a multi-faceted feedback system including novel mechanisms to investigate how terminal feedback affects students.

Our evaluation of LiftVR shows that it has a statistically significant and positive effect on students lifting techniques. The system has been received well and achieved effective results regardless of user height, weight, or age.

We challenged the current dominant approach of full-body motion guidance to visualize experts as natural humanoid entities. This resulted in us designing an abstract guidance system named Zone. It utilizes the symmetric nature of lifting motions to pair joints into singular zones, leading to less visual clutter whilst making it easier for students to assume postures. To compare our design to the state of the art approach, we developed a second concurrent guidance system named Ghost-Skeleton. Our analysis revealed that in path practice, Zone students significantly outperform those using Ghost-Skeleton. Though it is more tolerant to student errors and conveys less simultaneous information than its counterpart, Zone students are able to match Ghost-Skeleton student performances. Analysis revealed that Zone students achieved lower errors and higher improvement in practically every measured variable, however, due to the small sample size in our study, we were unable to establish statistical significance.

We asked users to rate the systems based on understandability, mental effort, and helpfulness. Zone achieved better scores in every category, including statistical significance for mental effort. As such, the reduced visual clutter and fewer simultaneously visible entities made Zone easier to follow.

In conclusion, guidance system expert visualization has an profound effect for path or posture guidance systems, though less so for movement guidance.

We evaluated the effects of students with or without feedback to investigate whether terminal feedback impacts performance. For this, we implemented a score, replay and joint path visualization system. Our analysis revealed no statistically significant effects on student performance, revealing that when terminal guidance systems are used together with concurrent systems, only the latter affect the students performance.

LiftVR does not, however, incorporate prescriptive or corrective terminal feedback systems, which have been shown to affect student performance more than purely descriptive systems. We hypothesize implementing such terminal feedback in LiftVR could push the overall effect to significance.

Based on our findings and qualitative user feedback, we have derived design implications for future studies, concurrent guidance systems, and feedback systems in the realm of full body motion guidance.

7.2 Improvements to LiftVR

Based on our findings, the following paragraphs describe improvements which could be made to LiftVR.

OpenPose and PoseFix - Using OpenPose would bring LiftVR increased accessibility, higher FPS and similar performance to the Azure Kinect. Moreover, using PoseFix to process recorded motions enhances their quality, subsequently leading to improved feedback calculated from them.

Authoring and Classification - Authoring and editing tools could be implemented into the LiftVR ecosystem to allow additional motions to be added without requiring a Vicon system. Furthermore, repetition and motion classification would remove the necessity for timers and allow higher quality feedback.

Prescriptive and Corrective Feedback - Evaluation of our qualitative results revealed the necessity for prescriptive and corrective feedback systems, which LiftVR currently does not offer.

Critical Error Guidance - Feedback systems which respond only to significant errors could be included. For example, a large red arrow could appear when a singular joint is critically out of position or for a joint that failed to reach its intended position for several seconds.

Multi-3PP - LiftVR should remove the independent side and frontal perspectives for multi-3PP. However, students should still be able to choose between over-the-shoulder or frontal view.

Virtual Tutor - A virtual tutor could be implemented to give both encouragement and feedback. This provides LiftVR with humanity, something users wished for.

Audio Guidance - More inclusion of audio, such as a virtual tutor voice to deliver corrective feedback, or beeping noises when errors are being made.

Time Guidance - Audio or visual indication of different key time points in the lift.

Improved Replay - The replay feedback system could be improved to find key frames to highlight and how to correct the errors in them. In line with encouragement, it should also highlight positive sections of the performance.

Better Color Indicators - Whilst the color interpolation was generally received well, some complained the colors were not distinct enough.

Evaluation Function for Path Guidance - The current scoring function does not provide useful results during path guidance. As such, path guidance requires a new scoring function based on frame-error rather than error.

7.3 Future Work

Independent Object Tracking Currently, systems such as Azure Kinect DK and OpenPose provide excellent capabilities for full body tracking without markers. However, tools such as LiftVR would benefit from an independent object tracking system. For example, a developer would create a shape and feed its geometry to the tracker. The tracker would then report the locations of similar shapes in the scene. In LiftVR, this could be used to track barbells and weights without the need to attach trackers to them. The Azure Kinect's depth camera could prove critical for this functionality, giving it an advantage over OpenPose for LiftVR.

Improved Body Tracking Models During our study, we observed several limitations to Azure Kinect's body tracking accuracy. Baggy or unicolor clothing would cause noticeable issues. As such, students were instructed to wear skin-tight clothing articles of separate color. Furthermore, the HMD also caused the tracker to report inaccurate results. When looking downwards, the tracker would confuse the overall body rotation, causing the skeleton in LiftVR to spin violently. As a results, we instructed participants not to look straight down at the ground. Thereby, body tracking machine learning models could be improved to lessen these issues. We do not know if these issues also exist with OpenPose.

Bibliography

[AGMF13] F. Anderson, T. Grossman, J. Matejka, G. Fitzmaurice. "YouMove: enhancing movement training with an augmented reality mirror". In: *Proceedings of the 26th annual ACM symposium on User interface software and technology*. UIST '13. New York, NY, USA: Association for Computing Machinery, Oct. 2013, pp. 311–320. ISBN: 978-1-4503-2268-3. DOI: 10.1145/2501988.2502045. URL: https://doi.org/10.1145/2501988.2502045 (visited on 02/22/2023) (cit. on pp. 28, 34, 45, 58, 67).

[BCAT19] R. R. Barioni, W. Costa, A. Aleluia, V. Teichrieb. "BalletVR: a Virtual Reality System for Ballet Arm Positions Training". In: *2019 21st Symposium on Virtual and Augmented Reality (SVR)*. Oct. 2019, pp. 10–16. DOI: 10.1109/SVR.2019.00018 (cit. on p. 36).

[BQH+23] L. Büker, V. Quinten, M. Hackbarth, S. Hellmers, R. Diekmann, A. Hein. "How the Processing Mode Influences Azure Kinect Body Tracking Results". en. In: *Sensors* 23.2 (Jan. 2023). Number: 2 Publisher: Multidisciplinary Digital Publishing Institute, p. 878. ISSN: 1424-8220. DOI: 10.3390/s23020878. URL: https://www.mdpi.com/1424-8220/23/2/878 (visited on 06/16/2023) (cit. on p. 37).

[CK07] P. Comfort, P. Kasim. "Optimizing Squat Technique". en-US. In: *Strength & Conditioning Journal* 29.6 (Dec. 2007), p. 10. ISSN: 1524-1602. URL: https://journals.lww.com/nsca-scj/Abstract/2007/12000/Optimizing_Squat_Technique.1.aspx (visited on 06/12/2023) (cit. on p. 29).

[CLTK11] J. C. Chan, H. Leung, J. K. Tang, T. Komura. "A Virtual Reality Dance Training System Using Motion Capture Technology". In: *IEEE Transactions on Learning Technologies* 4.2 (Apr. 2011). Conference Name: IEEE Transactions on Learning Technologies, pp. 187–195. ISSN: 1939-1382. DOI: 10.1109/TLT.2010.27 (cit. on p. 33).

[Cre03] H. Creagh. "Cave Automatic Virtual Environment". In: *Proceedings: Electrical Insulation Conference and Electrical Manufacturing and Coil Winding Technology Conference (Cat. No.03CH37480)*. ISSN: 0362-2479. Sept. 2003, pp. 499–504. DOI: 10.1109/EICEMC.2003.1247937 (cit. on p. 20).

[DB10] D. Drobny, J. Borchers. "Learning basic dance choreographies with different augmented feedback modalities". In: *CHI '10 Extended Abstracts on Human Factors in Computing Systems*. CHI EA '10. New York, NY, USA: Association for Computing Machinery, Apr. 2010, pp. 3793–3798. ISBN: 978-1-60558-930-5. DOI: 10.1145/1753846.1754058. URL: https://dl.acm.org/doi/10.1145/1753846.1754058 (visited on 06/11/2023) (cit. on pp. 25, 26).

[EFGK03] P. T. Eugster, P. A. Felber, R. Guerraoui, A.-M. Kermarrec. "The many faces of publish/subscribe". In: *ACM Computing Surveys* 35.2 (June 2003), pp. 114–131. ISSN: 0360-0300. DOI: 10.1145/857076.857078. URL: https://dl.acm.org/doi/10.1145/857076.857078 (visited on 06/21/2023) (cit. on p. 63).

[EHG+21] H. Elsayed, P. Hoffmann, S. Günther, M. Schmitz, M. Weigel, M. Mühlhäuser, F. Müller. "CameraReady: Assessing the Influence of Display Types and Visualizations on Posture Guidance". In: *Designing Interactive Systems Conference 2021*. DIS '21. New York, NY, USA: Association for Computing Machinery, June 2021, pp. 1046–1055. ISBN: 978-1-4503-8476-6. DOI: 10.1145/3461778.3462026. URL: https://doi.org/10.1145/3461778.3462026 (visited on 12/08/2022) (cit. on pp. 22, 28, 32).

[EKS+22] H. Elsayed, K. Kartono, D. Schön, M. Schmitz, M. Mühlhäuser, M. Weigel. "Understanding Perspectives for Single- and Multi-Limb Movement Guidance in Virtual 3D Environments". In: *Proceedings of the 28th ACM Symposium on Virtual Reality Software and Technology*. VRST '22. New York, NY, USA: Association for Computing Machinery, Nov. 2022, pp. 1–10. ISBN: 978-1-4503-9889-3. DOI: 10.1145/3562939.3565635. URL: https://dl.acm.org/doi/10.1145/3562939.3565635 (visited on 05/29/2023) (cit. on pp. 28, 32, 38, 89).

[Fos23] D. Foster. *Simple Kalman filtering in Unity*. en. June 2023. URL: https://gist.github.com/davidfoster/48acce6c13e5f7f247dc5d5909dce349 (visited on 06/17/2023) (cit. on p. 42).

[HCC+17] T. He, X. Chen, Z. Chen, Y. Li, S. Liu, J. Hou, Y. He. "Immersive and collaborative Taichi motion learning in various VR environments". In: *2017 IEEE Virtual Reality (VR)*. ISSN: 2375-5334. Mar. 2017, pp. 307–308. DOI: 10.1109/VR.2017.7892299 (cit. on p. 33).

[Hei62] M. Heilig. *Sensorama Simulator. USA Patent 3050870*. 1962 (cit. on p. 19).

[Hei92] M. L. Heilig. "EL Cine del Futuro: The Cinema of the Future". en. In: *Presence: Teleoperators and Virtual Environments* 1.3 (Jan. 1992), pp. 279–294. ISSN: 1054-7460. DOI: 10.1162/pres.1992.1.3.279. URL: https://direct.mit.edu/pvar/article/1/3/279-294/58782 (visited on 06/12/2023) (cit. on p. 28).

[HFS+19] F. Hülsmann, C. Frank, I. Senna, M. O. Ernst, T. Schack, M. Botsch. "Superimposed Skilled Performance in a Virtual Mirror Improves Motor Performance and Cognitive Representation of a Full Body Motor Action". In: *Frontiers in Robotics and AI* 6 (2019). ISSN: 2296-9144. URL: https://www.frontiersin.org/articles/10.3389/frobt.2019.00043 (visited on 05/29/2023) (cit. on pp. 31, 91).

[HKS14] J. Hamari, J. Koivisto, H. Sarsa. "Does Gamification Work? – A Literature Review of Empirical Studies on Gamification". In: *2014 47th Hawaii International Conference on System Sciences*. ISSN: 1530-1605. Jan. 2014, pp. 3025–3034. DOI: 10.1109/HICSS.2014.377 (cit. on p. 34).

[HN82] K. Harada, E. Nakamae. "Application of the Bézier curve to data interpolation". en. In: *Computer-Aided Design* 14.1 (Jan. 1982), pp. 55–59. ISSN: 0010-4485. DOI: 10.1016/0010-4485(82)90013-6. URL: https://www.sciencedirect.com/science/article/pii/0010448582900136 (visited on 06/17/2023) (cit. on p. 43).

[HRVT16] T. N. Hoang, M. Reinoso, F. Vetere, E. Tanin. “Onebody: Remote Posture Guidance System using First Person View in Virtual Environment”. In: *Proceedings of the 9th Nordic Conference on Human-Computer Interaction*. NordiCHI ’16. New York, NY, USA: Association for Computing Machinery, Oct. 2016, pp. 1–10. ISBN: 978-1-4503-4763-1. DOI: 10.1145/2971485.2971521. URL: https://doi.org/10.1145/2971485.2971521 (visited on 02/22/2023) (cit. on p. 35).

[Hun86] J. S. Hunter. “The Exponentially Weighted Moving Average”. In: *Journal of Quality Technology* 18.4 (Oct. 1986). Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/00224065.1986.11979014, pp. 203–210. ISSN: 0022-4065. DOI: 10.1080/00224065.1986.11979014. URL: https://doi.org/10.1080/00224065.1986.11979014 (visited on 06/17/2023) (cit. on p. 42).

[ICA+23] E. Iannucci, Z. Chen, I. Armeni, M. Pollefeys, H. Pfister, J. Beyer. “ARrow: A Real-Time AR Rowing Coach”. en. In: (2023) (cit. on pp. 35, 38).

[IG-17] IG-Blocks. *GymInterior | 3D Interior | Unity Asset Store*. en. 2017. URL: https://assetstore.unity.com/packages/3d/props/interior/gyminterior-86812 (visited on 06/15/2023) (cit. on p. 39).

[IOA21] E. Ivorra, M. Ortega Pérez, M. L. Alcañiz Raya. “Azure Kinect body tracking under review for the specific case of upper limb exercises”. Inglés. In: *MM Science Journal (Online)*. Vol. 2021. Accepted: 2022-02-14T19:02:50Z ISSN: 1805-0476. MM Publishing, June 2021, pp. 4333–4341. DOI: 10.17973/MMSJ.2021_6_2021012. URL: https://riunet.upv.es/handle/10251/180803 (visited on 06/14/2023) (cit. on p. 22).

[JKCP15] A. Joshi, S. Kale, S. Chandel, D. Pal. “Likert Scale: Explored and Explained”. en. In: *British Journal of Applied Science & Technology* 7.4 (Jan. 2015), pp. 396–403. ISSN: 22310843. DOI: 10.9734/BJAST/2015/14975. URL: https://journalcjast.com/index.php/CJAST/article/view/381 (visited on 07/01/2023) (cit. on p. 69).

[KBBD17] G.-A. Koulieris, B. Bui, M. S. Banks, G. Drettakis. “Accommodation and comfort in head-mounted displays”. In: *ACM Transactions on Graphics* 36.4 (July 2017), 87:1–87:11. ISSN: 0730-0301. DOI: 10.1145/3072959.3073622. URL: https://dl.acm.org/doi/10.1145/3072959.3073622 (visited on 06/30/2023) (cit. on p. 20).

[KWM+10] M. Klapdohr, B. Wöldecke, D. Marinos, J. Herder, C. Geiger, W. Vonolfen. “Vibrotactile pitfalls: arm guidance for moderators in virtual TV studios”. In: *Proceedings of the 13th International Conference on Humans and Computers*. HC ’10. Fukushima-ken, JPN: University of Aizu Press, Dec. 2010, pp. 72–80. (Visited on 06/12/2023) (cit. on p. 28).

[MCL19] G. Moon, J. Y. Chang, K. M. Lee. “PoseFix: Model-Agnostic General Human Pose Refinement Network”. In: 2019, pp. 7773–7781. URL: https://openaccess.thecvf.com/content_CVPR_2019/html/Moon_PoseFix_Model-Agnostic_General_Human_Pose_Refinement_Network_CVPR_2019_paper.html (visited on 06/14/2023) (cit. on p. 22).

[Mic23] Microsoft. *Azure Kinect DK Documentation*. en-us. June 2023. URL: https://learn.microsoft.com/en-us/azure/kinect-dk/ (visited on 06/13/2023) (cit. on pp. 23, 41).

[MSP+22] T. Moesgen, A. Salovaara, E. Pouta, R. Pyykko, Y. Xiao. “Vibrotactile Motion Guidance for Stroke Rehabilitation: A Comparative Study”. In: *Proceedings of the 6th International Conference on Medical and Health Informatics*. ICMHI ’22. New York, NY, USA: Association for Computing Machinery, Oct. 2022, pp. 147–152. ISBN: 978-1-4503-9630-1. DOI: 10.1145/3545729.3545759. URL: https://doi.org/10.1145/3545729.3545759 (visited on 12/08/2022) (cit. on p. 69).

[MTUK95] P. Milgram, H. Takemura, A. Utsumi, F. Kishino. “Augmented reality: a class of displays on the reality-virtuality continuum”. In: *Telemanipulator and Telepresence Technologies*. Vol. 2351. SPIE, Dec. 1995, pp. 282–292. DOI: 10.1117/12.197321. URL: https://www.spiedigitallibrary.org/conference-proceedings-of-spie/2351/0000/Augmented-reality--a-class-of-displays-on-the-reality/10.1117/12.197321.full (visited on 06/30/2023) (cit. on p. 19).

[MVKP10] T. McDaniel, D. Villanueva, S. Krishna, S. Panchanathan. “MOVeMENT: A framework for systematically mapping vibrotactile stimulations to fundamental body movements”. In: *2010 IEEE International Symposium on Haptic Audio Visual Environments and Games*. Oct. 2010, pp. 1–6. DOI: 10.1109/HAVE.2010.5623965 (cit. on p. 25).

[NHA+20] A. Nabil, H. Hegazy, M. Abdelsalam, M. Hussien, S. Elmosalamy, Y. Hassan, A. Atia. “Usability Study of a comprehensive table tennis AR-based training system with the focus on players’ strokes”. In: *Journal of Ubiquitous Systems and Pervasive Networks* 13 (Oct. 2020), pp. 1–09. DOI: 10.5383/JUSPN.13.01.001 (cit. on pp. 22, 26, 37, 38, 89, 90).

[NTU+05] A. Nakamura, S. Tabata, T. Ueda, S. Kiyofuji, Y. Kuno. “Dance training system with active vibro-devices and a mobile image display”. In: *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*. ISSN: 2153-0866. Aug. 2005, pp. 3075–3080. DOI: 10.1109/IROS.2005.1545565 (cit. on p. 28).

[Oso18] D. Osokin. *Real-time 2D Multi-Person Pose Estimation on CPU: Lightweight OpenPose*. arXiv:1811.12004 [cs]. Nov. 2018. DOI: 10.48550/arXiv.1811.12004. URL: http://arxiv.org/abs/1811.12004 (visited on 06/14/2023) (cit. on p. 22).

[PBC+17] A. Parent-Thirion, I. Biletta, J. Cabrita, O. Llave Vargas, G. Vermeylen, A. Wilczynska, M. Wilkens. *6th European Working Conditions Survey: overview report*. en. 2017 update. EF 16/34. Luxembourg: Publications Office of the Euroepan Union, 2017. ISBN: 978-92-897-1596-6. DOI: 10.2806/784968 (cit. on p. 15).

[SFO+12] C. Schönauer, K. Fukushi, A. Olwal, H. Kaufmann, R. Raskar. “Multimodal motion guidance: techniques for adaptive and dynamic feedback”. In: *Proceedings of the 14th ACM international conference on Multimodal interaction*. ICMI ’12. New York, NY, USA: Association for Computing Machinery, Oct. 2012, pp. 133–140. ISBN: 978-1-4503-1467-1. DOI: 10.1145/2388676.2388706. URL: https://doi.org/10.1145/2388676.2388706 (visited on 12/08/2022) (cit. on pp. 23, 28).

[SH16] D. Schmalstieg, T. Hollerer. *Augmented Reality: Principles and Practice*. en. Google-Books-ID: qPU2DAAAQBAJ. Addison-Wesley Professional, June 2016. ISBN: 978-0-13-315320-0 (cit. on pp. 19, 20).

[SRRW13] R. Sigrist, G. Rauter, R. Riener, P. Wolf. "Augmented visual, auditory, haptic, and multimodal feedback in motor learning: A review". en. In: *Psychonomic Bulletin & Review* 20.1 (Feb. 2013), pp. 21–53. ISSN: 1531-5320. DOI: 10.3758/s13423-012-0333-8. URL: https://doi.org/10.3758/s13423-012-0333-8 (visited on 06/11/2023) (cit. on pp. 26, 27).

[STK92] S. Silverman, L. Tyson, J. Krampitz. "Teacher feedback and achievement in physical education: Interaction with student practice". en. In: *Teaching and Teacher Education* 8.4 (Aug. 1992), pp. 333–344. ISSN: 0742-051X. DOI: 10.1016/0742-051X(92)90060-G. URL: https://www.sciencedirect.com/science/article/pii/0742051X9290060G (visited on 06/30/2023) (cit. on pp. 25, 38, 90, 91).

[SW97] M. Slater, S. Wilbur. "A Framework for Immersive Virtual Environments (FIVE): Speculations on the Role of Presence in Virtual Environments". In: *Presence: Teleoperators and Virtual Environments* 6.6 (Dec. 1997), pp. 603–616. DOI: 10.1162/pres.1997.6.6.603. URL: https://doi.org/10.1162/pres.1997.6.6.603 (visited on 06/30/2023) (cit. on p. 19).

[TMS+01] T. Tsuji, Y. Matsuyama, K. Sato, Y. Hasegawa, Y. Yimin, H. Iwata. "Epidemiology of low back pain in the elderly: correlation withlumbar lordosis". en. In: *Journal of Orthopaedic Science* 6.4 (July 2001), pp. 307–311. ISSN: 0949-2658. DOI: 10.1007/s007760100023. URL: https://www.sciencedirect.com/science/article/pii/S0949265815333315 (visited on 06/12/2023) (cit. on p. 29).

[Val23] Valve. *SteamVR Plugin | Integration | Unity Asset Store*. en. 2023. URL: https://assetstore.unity.com/packages/tools/integration/steamvr-plugin-32647 (visited on 06/16/2023) (cit. on p. 39).

[VBG13] E. Velloso, A. Bulling, H. Gellersen. "MotionMA: motion modelling and analysis by demonstration". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '13. New York, NY, USA: Association for Computing Machinery, Apr. 2013, pp. 1309–1318. ISBN: 978-1-4503-1899-0. DOI: 10.1145/2470654.2466171. URL: https://doi.org/10.1145/2470654.2466171 (visited on 02/22/2023) (cit. on pp. 26, 34).

[Vic23] Vicon. *Full body modeling with Plug-in Gait - Nexus 2.12 Documentation - Vicon Documentation*. June 2023. URL: https://docs.vicon.com/display/Nexus212/Full+body+modeling+with+Plug-in+Gait (visited on 06/14/2023) (cit. on pp. 21, 22).

[Wel97] G. Welch. "An Introduction to the Kalman Filter". en. In: (1997) (cit. on p. 37).

[Wen21] K. Wennrich. "Developing a multimodal feedback motion guidance system in VR for people with motion disabilities". en. Accepted: 2021-11-15T08:32:36Z ISBN: 9781777925994. MA thesis. 2021. DOI: 10.18419/opus-11776. URL: http://elib.uni-stuttgart.de/handle/11682/11793 (visited on 12/08/2022) (cit. on pp. 28, 37).

[WSS23] T. Walter, N. Stutzig, T. Siebert. "Active exoskeleton reduces erector spinae muscle activity during lifting". In: *Frontiers in Bioengineering and Biotechnology* 11 (Apr. 2023), p. 1143926. ISSN: 2296-4185. DOI: 10.3389/fbioe.2023.1143926. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10168292/ (visited on 06/12/2023) (cit. on p. 15).

[YAM+20] X. Yu, K. Angerbauer, P. Mohr, D. Kalkofen, M. Sedlmair. “Perspective Matters: Design Implications for Motion Guidance in Mixed Reality”. In: *2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. ISSN: 1554-7868. Nov. 2020, pp. 577–587. DOI: 10.1109/ISMAR50242.2020.00085 (cit. on pp. 23, 26, 36, 38, 52).

All links were last followed on July 10, 2023.

Declaration

I hereby declare that the work presented in this thesis is entirely my own and that I did not use any other sources and references than the listed ones. I have marked all direct or indirect statements from other sources contained therein as quotations. Neither this work nor significant parts of it were part of another examination procedure. I have not published this work in whole or in part before. The electronic copy is consistent with all submitted copies.

place, date, signature