KTH ROYAL INSTITUTE
OF TECHNOLOGY

Degree Project in the Field of Technology and the Main Field of Study Computer Science and Engineering

Second cycle, 30 credits

The Perception of Weight in Virtual Reality When Combining Velocity Limiting and Control/Display Ratio Modification

KÁRI STEINN AÐALSTEINSSON

Stockholm, Sweden, 2023

The Perception of Weight in Virtual Reality When Combining Velocity Limiting and Control/Display Ratio Modification

KÁRI STEINN AÐALSTEINSSON

Master's Programme, Interactive Media Technology, 120 credits
Date: June 12, 2023

Supervisors: Björn Thuresson, Adria Cruz
Examiner: Tino Weinkauf
School of Electrical Engineering and Computer Science
Host company: Ericsson
Swedish title: Uppfattningen om Vikt i Virtuell Verklighet vid Kombination av Hastighetsbegränsning och Ändring av Kontroll/Visningsförhållande

The Perception of Weight in Virtual Reality When Combining Velocity Limiting and Control/Display Ratio Modification

KÁRI STEINN AÐALSTEINSSON

This study explored the perception of weight in virtual reality using the combination of limiting the velocity with which an object can be lifted and control/display ratio modification and what affordances and challenges the combination might bring. Both interaction methods have thus far been explored independently, but not together. A virtual reality environment was designed that allowed each interaction method to be tested independently and together. A between-subjects user study was conducted where participants did multiple weight discrimination tasks along with a task involving absolute weight estimation. Results indicated that the combination of the two interaction methods increased the accuracy of weight discrimination considerably while not allowing for absolute weight estimation. It also increased the likelihood of noticing the interaction method with some participants that tested only one interaction method struggling to distinguish between the weights. Future work is required to reduce the frustrations that dropping objects due to the velocity limit caused among participants.

SAMMANFATTNING

Denna studie undersökte uppfattningen om vikt i virtuell verklighet genom att kombinera begränsning av lyftförmågan hos ett objekt med ändring av kontroll-/visningsförhållande och vilka möjligheter och utmaningar denna kombination kan medföra. Båda interaktionsmetoderna har hittills undersökts separat, men inte tillsammans. En virtuell verklighetsmiljö utformades som möjliggjorde testning av varje interaktionsmetod separat och tillsammans. En användarstudie med mellangruppsdesign genomfördes där deltagarna utförde flera uppgifter för viktdiskriminering samt en uppgift för att uppskatta absolut vikt. Resultaten indikerade att kombinationen av de två interaktionsmetoderna avsevärt ökade noggrannheten för viktdiskriminering samtidigt som den inte tillät uppskattning av absolut vikt. Det ökade också sannolikheten för att märka interaktionsmetoden hos vissa deltagare, där de som endast testade en interaktionsmetod hade svårt att skilja mellan vikterna. Framtida arbete krävs för att minska frustrationen då ett objekt tappas på grund av hastighetsbegränsningen vilket orsakades bland deltagarna.

CCS Concepts: • **Human-centered computing** → **Virtual reality**; **Empirical studies in interaction design**; **Interaction techniques**; **User studies**.

Keywords: Weight perception, virtual reality, velocity limiting, control/display ratio modification

Nyckelord: viktuppfattning, virtuell verklighet, hastighetsbegränsning, ändring av kontroll-/visningsförhållande

Author's address: Kári Steinn Aðalsteinsson, ksad@kth.se, KTH Royal Institute of Technology, School of Electrical Engineering and Computer Science, Stockholm, Sweden, SE 100 44.

Document date: June 12, 2023

Swedish title: Uppfattningen om Vikt i Virtuell Verklighet vid Kombination av Hastighetsbegränsning och Ändring av Kontroll/Visningsförhållande

1 INTRODUCTION

The perception of weight in virtual reality is a challenge due to the absence of force feedback. Virtual reality (VR) allows for immersive experiences as users can move through and manipulate a three-dimensional computer-generated virtual environment wearing a head-mounted display and holding a pair of handheld controllers [17]. Training is one of the sectors where VR has had the greatest impact due to several factors, including being able to simulate work situations using a one-to-one scale immersion and multi-sensory interactions, less risk for the learner as they can be taught how to handle dangerous situations, and savings on consumables and heavy equipment [2]. Many technical challenges prevent VR from fully immersing users in the experience [26]. These potential challenges include senses such as force feedback, smell, and taste [17]. The most prominent of these is the lack of kinesthetic (force feedback) cues. When lifting an object in the real world, the stretch and elongation receptors in the muscles and tendons convey to the brain how much load is being placed on the musculoskeletal system as well as the extension and rotation of the limb. However, when grasping a virtual object, the arm does not receive kinesthetic feedback. Thus, other means of conveying the perception of weight are necessary for VR, due to the lack of kinesthetic feedback [26]. The need for weight perception is especially important for the training field, where a lot of precision is necessary [15]. This includes feeling the weight of a tool in micro-surgery to be able to adjust the force applied to an area [17], or during assembly line simulations for the production of fragile objects [5, 17].

In recent years research has been done on how weight can be added to VR experiences. These include both haptic- and pseudo-haptic weight simulations. Haptic weight simulation is achieved using haptic interfaces such as handheld-, wearable, or ground-based interfaces [15]. The common thread for haptic interfaces is the requirement for additional hardware which is not cost-effective and often is specific to each application. Some examples of haptic interfaces are controllers that fan out to create air resistance [35], and a liquid-based haptic feedback device that simulates realistic object weights and inertia [7]. Pseudo-haptic weight simulation, however, employs sensory modalities such as visual feedback to generate haptic perception without the need for haptic devices. Various pseudo-haptic cues enable weight perception, these include control/display ratio, velocity, and passive force [15]. An exploratory study from 2018 by Emma Bäckström [4] explored how limiting the velocity with which objects could be picked up affected the perception of heaviness of objects in a virtual reality environment. The method showed some promise and people were able to order objects according to their weight. However, participants in her study did not always realize why they were dropping the objects, which caused some frustration among them. To help ease this frustration some additional feedback could help participants perceive why they are dropping the objects and assist in the perception of weight. This can be done by adding some visual cues when the user is lifting the object to better indicate the weight differences. One such cue could be control/display (C/D) ratio modification where the movement offset between the virtual hand and the actual hand can be adjusted to create an illusion of weight perception [15]. Hand movements when lifting heavy objects could thus be delayed relative to the real-world hand, while hand movements when lifting lighter objects would appear similar in both the real world and in the virtual environment. Studies have shown that people are able to distinguish differences in weight using C/D ratio modification [26] and thus a combination of C/D ratio modification and velocity limiting could help with the perception of weight differences and limit the frustrations of having only velocity limiting.

1.1 Problem statement

The ability to perceive weight differences is important for many applications such as training in VR [15]. The perception of weight in VR is a challenge in modern VR headsets such as the Meta Quest 2, Sony Playstation VR2, and the HTC Vive Pro 2. This is due to the controllers' limitations when lifting objects where all objects appear to have the same weight. Recent studies have explored whether haptic- or pseudo-haptic weight simulation can help address these challenges in weight perception. Modern VR headsets currently have no haptic devices to simulate weight and therefore it is of interest to explore ways to simulate weight perception without the need for additional hardware. This study will expand on studies that have explored how limiting the velocity with which an object can be picked up [4, 26] and how modifying the C/D ratio can affect the perception of weight of objects in virtual reality [1, 13, 26]. More specifically whether the combination of the two can add to the perception, or whether one of them might be redundant. The accuracy of each method will be explored through weight discrimination tasks to evaluate which method performs the best. Additionally, this study will explore the potential of using these methods for absolute weight perception of virtual objects. Absolute weight perception in VR using such methods has been previously explored [23, 26]. This study expands on that work by seeing if people choose an object with the same weight when presented with virtual objects that look the same but have different virtual weights. The object will be a 1-liter milk carton as most people have experience with holding a 1-liter carton and thus have an idea of how heavy it should be.

1.2 Research questions

The following research questions have been devised to address the problem:

RQ1: How is the perception of weight in virtual reality affected by limiting the velocity with which objects can be picked up and modifying the C/D ratio when lifting it as measured by the accuracy of weight discrimination?

RQ2: What are the affordances and challenges when adding both velocity limiting and C/D ratio modification compared to adding one of them?

1.3 Hypotheses

H1: The hypothesis for RQ1 is that adding both velocity limiting and C/D ratio modification will allow users to distinguish weight differences between lighter and heavier objects, and with better accuracy than with just one of the methods. This is due to the slower velocity required for lifting heavy objects and the fact that the C/D ratio modification should assist the perception of weight through slower virtual hand movements for heavier objects.

H2: The first hypothesis for RQ2 is that having only velocity limiting will cause similar frustrations as Emma Bäckström found in her exploratory study [4]. However, adding C/D ratio modification alongside velocity limiting will address the initial frustrations among users with only velocity limiting. This is because having the C/D ratio modification might assist users in initially identifying that objects have different weights.

H3: The second hypothesis for RQ2 is that it is expected that C/D ratio modification will give the best results with regard to time to task completion as the user will not drop the objects involuntarily due to the velocity limiting which might affect the time.

H4: The third hypothesis for RQ2 is that when presented with a task where people must identify an object according to its absolute weight, there will not be an agreement on which one to choose. This is due to the limitations that these weight perception methods have when it comes to absolute weight compared to relative weight.

1.4 Delimitations

This study will not explore the influence of size on the perception of weight. All of the virtual objects will be of the same size and will look exactly the same. The object can only be picked up using one hand to limit participants from comparing two objects directly side-by-side. The virtual environment will not feature any sound, tactile, or haptic feedback. The development will be limited to the Meta Quest 2 VR headset as it is a popular VR headset that has all the required features for this study while being portable and not requiring any external trackers.

2 RELATED WORK

This section will focus on aspects that affect the perception of weight in VR. The first of which is immersion and presence. Adding weight perception to VR could increase immersion and presence by adding one of the real senses into virtual environments. The next part focuses on weight perception in the real world and aspects of it that might be applied to VR. Following that weight perception in VR is explored and the different methods that have been used to simulate it using pseudo-haptics. Most of the emphasis will be on the two methods being explored in this study, velocity limiting and C/D ratio modification.

2.1 Immersion and presence

Immersion and presence are important terms for VR. In technical terms, immersion is capable of producing a sensation of presence [19]. Immersion describes the extent that a computer display is able to deliver an inclusive, extensive, surrounding, matching, and vivid illusion of reality to the senses of a human participant [28]. This means that immersion is related to the multi-modal nature of the perceptual senses and the interactive aspects of a VR experience [19]. Presence, however, has both a subjective and objective definition. In its essence, it is a description of a person's state with respect to an environment. Subjective presence is a person's evaluation of their sense of "being there" in a virtual environment. Objective presence relates to how a person behaves in a virtual environment and is thus rather an observable behavioral phenomenon [28]. The idea of presence is that a highly present participant should experience the virtual environment as more of an engaging reality than the surrounding physical world and that what they are seeing should rather be a visited place than just images being seen [28]. Presence is thus a core component of VR and has often been thought of as a sign of the "ecological validity" of VR devices [19].

Immersion can be an objective and quantifiable measure of what a system provides. Each of the dimensions of immersion; inclusive, extensive, surrounding, matching, and vivid, has multiple levels. For example, the "surrounding" part of immersion can refer to whether a user is looking at the virtual environment through a computer screen or a head-mounted display (HMD). Each of the levels refers to the extent that reality is excluded, how many senses are replaced, and how the user's movements are mapped in the virtual environment. Subjective presence can be correlated with higher levels of immersion [28].

Researchers in the field of VR have developed questionnaires to evaluate subjective presence. One commonly used is the *Igroup Presence Questionnaire* which is a scale for measuring the sense of presence experienced in a virtual environment [27]. This questionnaire will be used to compare the different conditions that will be explored in this study.

The perception of weight in VR could increase presence and immersion and VR as the perception of weight is currently missing from VR applications. Weight is something that is experienced in real life and bringing that to virtual environments adds to how real they feel.

2.2 Weight perception

Weight is a measure of the force that gravity exerts on an object. It is the product of an object's mass and the acceleration due to gravity. The perception of weight relates to the subjective weight of an object. The human ability to perceive and discriminate weights is complex and studies have revealed that multiple factors influence the weight perception process. The mechanisms supporting these perceptual theories, however, are not widely agreed upon [15].

The perceived weight of an object is affected by its properties. These include the object's size [6, 20, 29, 34], material [3], shape [10], and color or brightness [9, 31, 32]. For example, in the size-weight illusion, a smaller object is perceived to be lighter than a larger one, even if the mass of both objects is identical [6, 20]. To reduce the likelihood of these properties affecting the present study, all objects will be identical except for their mass.

Studies have suggested that one of the main cues used to judge weight is the speed with which we lift an object [20, 21]. A study by Davis and Roberts [8] also suggests that an object lifted with a greater velocity is perceived to be lighter.

2.3 Pseudo-haptic weight simulation in VR

The perception of weight in VR can be simulated through either haptic- or pseudo-haptic weight. Haptic weight simulation requires haptic devices and they are limited by the device's weight and wearing comfort [15]. However, pseudo-haptic weight simulation does not require any haptic stimulus to simulate haptic sensations. Instead, other sensory modalities, mainly visual stimuli, are used to simulate the perception of weight [15]. Pseudo-haptic weight simulation can be achieved using various methods. These include control/display ratio, velocity, passive force, and audio [15]. Research has mostly been conducted by exploring one of these methods, and therefore it could prove beneficial to explore a combination of them. Thus, this study focuses on the combination of C/D ratio and velocity.

2.3.1 Control/display ratio modification. Control/display ratio is the ratio between the speed of hand movement (control) and the speed of cursor movement (display) [16]. In VR the cursor is often a virtual representation of the handheld controllers that the user holds in each hand. Modifying the C/D ratio involves creating a discrepancy between the translation and rotation of the real hand of the user and the cursor. A C/D ratio of one means that the virtual movement matches the real movement. If the ratio is less than one then the virtual movement is decreased, while a ratio above one means that the virtual movement is amplified.

Various studies have explored how C/D ratio modification can affect the perception of weight in virtual environments. Some studies [11, 13, 26, 30] used visual feedback along with a physical haptic stimulus. Samad et al. [26] explored how manipulating the C/D ratio could change the perceived weight of virtual objects. They did so by conducting multiple weight discrimination tasks with two motion-tracked cubes, where they changed the C/D ratio of the cubes, and asked participants in their study to place the heavier of the two at a specified location. The C/D ratio in the study ranged from 0.7 to 1.3. Their findings demonstrated that C/D ratio modification induces a genuine perception of weight while preserving ownership over the virtual hand. Their findings suggest that C/D ratio modification can be introduced without affecting the sense of presence. Dominjon et al. [13] reported that amplified virtual displacement with respect to the users' displacement can make users feel that objects are lighter than their actual weight. In their study, the participants lifted a ball that had its weight simulated using a PHANToM haptic interface [18]. Participants moved the ball while watching a virtual representation of it on a computer screen. Their results showed that in some cases it was even possible to make a heavy object feel lighter than a light object just by amplifying the offset.

Unlike the studies mentioned above, the study by Rietzler et al. [23] solely utilized visual feedback without actual physical stimuli. They did so by visualizing the offset between the real hand position and the virtual hand position to suggest different weights. Their study showed that perceivable tracking offsets led to significantly higher levels of presence, immersion, and enjoyment compared to having no weight simulation. They also reported perceptional thresholds and offset boundaries as design guidelines. Participants in their study were willing to accept an offset of 42cm between the real and virtual hand as a weight metaphor while being able to detect differences of around 2.5cm. The authors suggest using offsets below 24cm for longer interactions, as all participants agreed that those values made for a good and comfortable weight metaphor.

2.3.2 Velocity. An object's mass is proportional to its velocity when subjected to constant force. Humans usually associate an item's traveling speed with its perceived weight, where an object that moves slowly is perceived to be heavier than an object that moves fast [15]. Using velocity for weight perception in VR has been explored both by altering the velocity with which objects move [12, 25] and by letting the velocity of users' hands influence the moving experience in some way [4, 14, 33]. The research of this study will focus on the second method, however, it is of interest to mention the first method due to its relation to C/D ratio modification.

The visually represented speed of an object in combination with other stimuli has been the subject of various studies [12, 25]. In those studies, the additional stimuli are provided by a wearable haptic device. A study by Rosa et al. [25] explored the influence of visual cues on passive tactile sensations in VR. Participants in their study were represented by an avatar with their arm lying on a table. The participants copied the pose of the avatar in real-life while wearing a vibro-tactile sleeve. The study involved a sphere falling down onto the arm of the avatar in the virtual environment and adding a tactile vibration to the arm when the sphere hit the arm. The size and velocity of the falling sphere changed while the vibration provided by the vibro-tactile sleeve stayed the same. When the velocity of the falling object increased the object was perceived to be heavier and the tactile vibration was perceived to be lighter, even though the intensity of the vibration stayed the same.

Kim et al. [12] explored how multi-sensory pseudo-haptic feedback affects simulated weight perception. They did so by combining C/D ratio manipulation and electrical muscle stimulation (EMS). Their C/D ratio modification was different from what was mentioned in section 2.3.1 as it involved adding a variable time delay to the virtual movement based on the real hand velocity. This meant that over time the virtual hand reached the actual position of the real hand which differs from the C/D ratio modification mentioned above where the offset stays throughout the lifting of an object. Therefore it applies more to the velocity category of pseudo-haptic weight perception, although it combines both aspects in some ways. The closer their C/D ratio is to zero, the more time the virtual hand takes to reach the real position. In two experiments they independently manipulated the C/D ratio with delay and EMS status. They specified the range that their C/D ratio could induce the perception of virtual weight and demonstrated that it could be more than twice widened with the addition of EMS on top of the visual offset.

Some studies have tried to induce a perception of weight based on movement speed by forcing users to move heavy objects slowly and less so with lighter objects [4, 14, 33]. This method focuses more on the real hand movement speed rather than the movement speed of the virtual object. Lee et al. [14] introduced a pseudo-haptic interface called "Force Arrow" where an arrow is presented to the user that indicates to them the current lifting force and the standard force so they can adjust their lifting force to the standard value. The arrow points upwards if the standard force is being applied, but points to the left if it is too slow, and points to the right if it is too fast. Their experiments involved lifting virtual objects that looked like dumbbells. Participants had to complete 10 trials where they had to lift three dumbbells of different weights. The results from their study showed that users could distinguish between the different weights of the virtual objects.

Weser et al. [33] explored transforming the virtual object weight into maximum object speed. This meant that light objects could be lifted quickly, whereas heavy objects could not. When the controller exceeded the given maximum speed the virtual object would begin to lag behind the controller and eventually detach from the controller so that the object dropped to the ground. This study is thus a combination of both limiting the velocity and introducing a movement delay. Participants in their study first got to try lifting reference objects of a given weight and were later tasked with lifting objects of different weights and sizes and guessing their weight. Participants were encouraged to alternate between lifting the reference objects and the object they were guessing the weight of to give as accurate of a guess as possible. The results showed that participant estimates of virtual object mass and the programmed object mass were strongly positively correlated. Thus demonstrating that users can learn to accurately identify the "weights" of virtual objects without any physical mass present.

The main inspiration for this study is the exploratory study by Emma Bäckström [4]. She implemented an interaction method similar to the one previously mentioned where a velocity limit was given to each virtual object based on its "weight". She conducted a concept study where participants were tasked with identifying the heavier and lighter of two identical-looking kettlebells and ordering six kettlebells according to their weight. The results from her study indicate that the interaction method has the potential to give users a notion regarding the heaviness of virtual objects, without affecting their presence. The interaction method seemed to give participants a tool for distinguishing between different weights. However, participants often became frustrated because they kept dropping the object.

3 METHODOLOGY

To explore how C/D ratio modification and limiting the velocity with which a virtual object can be lifted affect the perception of weight, a virtual reality environment was developed with both interaction methods implemented. A user study was conducted where each interaction method was tested along with the combination of the two. First, a pilot study was conducted to test whether the interaction methods were working as intended and if there were any problems with the study. Following that a between-subjects user study was conducted to explore whether the combination of both interaction methods was beneficial or whether one of them was redundant.

3.1 Participants

Participants were recruited from KTH, Ericsson, and through word of mouth. The only requirement was that participants should not have significant movement, vision, or hearing difficulties. There was no requirement for previous VR experience as the interaction was simple and the experience of participants should not have any influence on the performance.

Three participants were recruited for the pilot study, two males and one female. The procedure in the pilot study was the same as the main study and thus featured the same tasks and questions. Each participant in the pilot study tested one of the conditions of the study and the purpose was to discover any potential issues that might arise so that they could be fixed before the main user study. No data analysis was done on the data from the pilot study.

The main study featured 30 participants, 19 males, and 11 females. The participants were between 20-57 years old (mean = 27.3, sd = 7.9). The VR experience ranged from 5 participants having no experience to 4 participants using it multiple times per week. Four participants were left-handed and 26 were right-handed.

3.2 Study design

To ensure consistency throughout the study the same structure was used for each participant. Each test was conducted with a moderator and the participant. The participants were told that the test was for a degree project and that the test would take place in a virtual reality environment. The specifics of the condition they were testing were not explained, only that they were to order objects according to their weight. Prior to starting the experiment, the participants signed a consent form stating that they were voluntarily taking part in the study, that they could withdraw from the study at any time, and that the data collected during the study was anonymized and data collection was in accordance with GDPR. Following the signing of the consent form participants answered demographic questions and what experience they have had with VR previously.

Before entering the virtual reality environment the moderator explained the buttons on the controllers that the participants would be using. Three buttons were used during the study, two on the dominant hand controller and one on the non-dominant hand controller. The trigger button on the dominant hand controller was used in the beginning to select which hand was used during the trial. The grip button on that controller was then used to pick up the objects in the virtual environment. The primary button on the non-dominant hand controller was used to continue to the next trial.

When entering the virtual environment, the participants were given both hand controllers and started a tutorial to get acquainted with VR. The environment was simple and featured a white floor with a brown table directly in front of the participant. At the back of the table was a white sign with black text. This table and sign would be present in every part of the study, so the participants became familiar with the environment before beginning the study. The default light blue Unity skybox surrounded the environment. The tutorial did not feature any of the objects that they would go on to lift in the actual study. The participants saw both of their hands and in front of them was a white sign with black text explaining that they had to select their dominant hand. Upon doing so, the virtual representation of their non-dominant hand disappeared. This was done to limit the participant to lifting the objects with one hand instead of two so that they were not able to directly compare the weight of two objects at the same time. Upon confirming the selected hand the text on the sign changed to indicate that the height of the tables would be calibrated the next time they would press the continue button, so they should look straight forward the next time it was pressed. The next part of the tutorial featured a hammer lying on the table. The participant could practice picking up said hammer and this was done to familiarize the user with picking up an object so it would not influence the actual study if they did not know how to pick up an object. The hammer was the only asset used during the study that was found on the Unity Asset Store[1]. The environment with the hammer can be seen in Figure 1. The tutorial also familiarized the participants with the environment that the study took place in so that they would be less inclined to look around during the study. Following that they continued on to the study itself.

Figure 1: The tutorial environment featuring the hammer that participants could practice picking up.

Three conditions were tested during the user study, and they were: just limiting the velocity, just C/D ratio modification, and limiting the velocity with C/D ratio modification. The condition that each participant received was pseudo-randomized so that each condition featured an equal amount of times. All conditions came up for every three participants but in a randomized order. This was done to ensure that all conditions were featured an equal amount of times during

[1] https://assetstore.unity.com/packages/3d/props/industrial/workplace-tools-86242

the study in case not enough participants would have been recruited. When the first task started the virtual environment featured a sign with instructions similar to how it was in the tutorial. The participants were not instructed on how to interact with the objects in the study other than the tutorial in the beginning when they could practice picking up objects. The participants were not instructed to think-aloud during the study in order for them to be able to concentrate on the task at hand. While the participants finished the tasks the moderator wrote down observations about the strategy they were using during the study or anything of note that they said out loud.

After finishing all tasks the participants filled in a presence questionnaire that quantified the level of presence experienced by participants and is common practice during VR studies (more in section 3.5). A semi-structured interview was conducted following the presence questionnaire, where the focus was on whether participants experienced weight during the study and the strategy they used for the different tasks. Some of the questions asked were "Did you notice any differences when lifting the virtual objects?" and "Did it feel like lifting objects of different weights?". As a majority of the participants would be from Iceland, like the moderator, the questions for the semi-structured interview were translated into Icelandic for the Icelandic participants. This was done to ensure that participants could express their feelings adequately and would not be limited by their knowledge of English.

The user study was between-subjects and featured 30 participants where 10 participants tested each condition. The study was between-subjects in order to get the cleanest data by limiting the learning effect that might occur if participants were to test each condition and to stop the participants from comparing the different conditions. This allowed for comparisons that could shed light on whether having both C/D ratio modification and velocity limiting helped the perception of weight in VR, and whether one is more important than the other.

3.3 Tasks

Three variations of tasks were done during the study. They will be referred to as the six-weight, two-weight, and milk carton tasks. The first two involved ordering objects that looked like square kettlebells with a handle at the top according to their weight. As was mentioned in section 2.2, color is one of the properties that can affect the perceived weight of an object. Thus, all of the objects in the first two trials were grey in color to limit the possible effect that their color could have on the perception of weight. The purpose of the first two tasks was to find the accuracy differences between the conditions for RQ1 and to find the affordances and challenges that the combination brings for RQ2.

The first two tasks are subjective to each participant where they are simply comparing objects and identifying weight differences between them. A third task was thus developed to explore the potential of using these methods to identify absolute weights. This task was therefore developed more out of curiosity, and as an addition to the first two tasks, and is not the primary focus of this study, and its purpose was not to answer the research questions specifically. The third task involved the participant picking a milk carton, from six possible milk cartons, that best matched their expectation of what picking up a 1-liter milk carton should feel like. The purpose of the milk carton task was to attempt to answer RQ3. The milk carton was white with blue features and it had a gray handle on the side to indicate to the user how it should be picked up. The white and blue colors of the milk carton were chosen due to the fact that milk cartons often feature such colors. The color was not grey like in the first two tasks as the user was already told that it should be a 1-liter milk carton and thus they should go by their own experiences of lifting such objects. Therefore it is more important that it looks visually similar to a real milk carton than limiting the effect that the color could have on the perception of weight.

Each task featured a similar setup where there was a brown table with a lower and upper part. The objects that participants had to lift up were always on the lower part of the table and they had to be lifted up onto the higher part to finish the task. This is done to force participants to lift the weights some distance, as the study by Samad et al [26], mentioned in section 2.3.1, showed that if the lifting movement was too small participants did not notice the C/D ratio modification.

Figure 2: The six-weight environment featuring the task description and the 6 different objects on the lower platform which had to be ordered according to their weight on the upper dark grey platforms.

Thus, forcing participants to lift each object some distance could increase the likelihood of participants noticing the interaction method they were testing. Behind the table in every task was a white sign with black text featuring the task description. These task descriptions never mentioned the interaction method being tested, only what had to be done to finish the task.

All tasks featured objects where the mass of the Unity RigidBody of the virtual objects were: 1, 3, 5, 7, 9, and 11. The six-weight and milk carton tasks featured all of them simultaneously, while the two-weight task only featured a pair of them at a time. The weights were paired as follows: $1 \Leftrightarrow 11$, $3 \Leftrightarrow 9$, $5 \Leftrightarrow 7$. This notation will be used throughout the study where, for example, $1 \Leftrightarrow 11$ represents the pair with the 1 and 11 masses. These pairs and masses are the same as in Emma Bäckström's [4] study to have similar trigger values for the velocity limit as were used in her study.

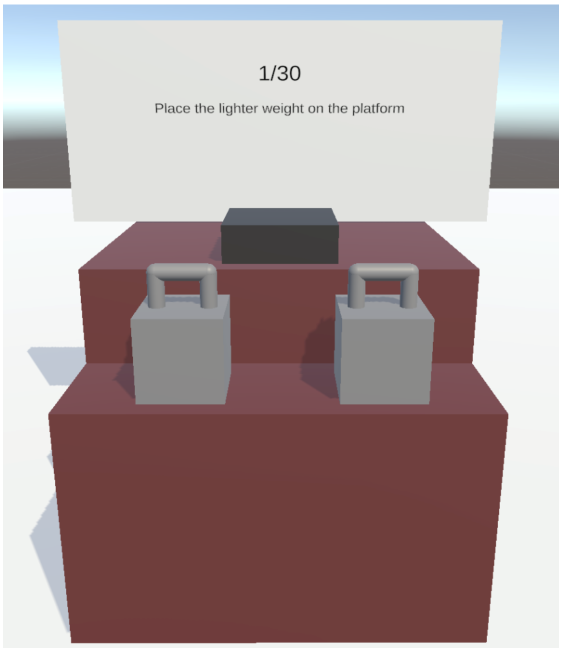

Figure 3: Two-weight environment with the description, counter, and 2 objects on the lower platform where the lighter one had to be placed on the dark grey platform.

3.3.1 Six-weight task. The six-weight task involved ordering six weights according to their weight from lightest (left) to heaviest (right). The weights had to be placed onto six dark grey platforms on the upper part of the table. The lightest weight was placed furthest to the left, and the heaviest furthest to the right. The virtual environment can be seen in Figure 2. The purpose of this task was to see how accurately participants could order the objects according to their weight relative to one another when presented with all of the possible weights used throughout the tasks.

3.3.2 Two-weight task. The two-weight task was similar to the six-weight task, except it only featured two objects and the task was to place the lighter object of the two on the upper dark grey platform. The two-alternative forced choice (2AFC) method was used for this task where the different pairs mentioned in section 3.3 featured 10 times each, for a total of 30 trials. This was done to get the probability of correct estimates by repeating each comparison multiple times. If a participant was not able to distinguish between the weights in one of the pairs, they would answer randomly and have a probability of 0.5. If they could reliably distinguish between them it would result in a higher number of correct estimates and a higher probability. A counter was placed on the task description sign to indicate to the participant that this task would feature 30 separate trials and in which trial they were at any given moment. The two-weight environment can be seen in Figure 3. This task aimed to explore how accurately participants could distinguish the lighter of two weights and to see whether the weight pair with the closest mass (5 ⇔ 7) would perform differently from the ones further apart (1 ⇔ 11, 3 ⇔ 9).

3.3.3 Milk carton task. The milk carton task featured six identical-looking milk cartons and the only difference was their weight. The task involved participants picking each milk carton up and placing the milk carton that best matched what they would expect a real 1-liter milk carton to feel like onto the dark grey platform on the higher table. The goal of this task was to see whether participants would choose the same milk carton as that would indicate that

Figure 4: The milk carton environment featuring the task description and the 6 different milk cartons on the lower platform where one had to be placed on the upper dark grey platform.

these interaction methods could possibly be used to represent the real weights of objects rather than just for comparison between two objects. The milk carton environment can be seen in Figure 4.

3.4 Implementation details

As mentioned in section 3.3, the only thing that differed between the objects was the mass of the Unity RigidBody (the masses were 1, 3, 5, 7, 9, and 11). Those values were then used to calculate the velocity limit and C/D ratio used for each object. Equation 1 was used to calculate the velocity limit for each mass.

$$velocityLimit = \frac{1}{log_{10}\left(mass + 1\right)} - 0.5 \tag{1}$$

Emma Bäckström's [4] implementation featured three different equations to calculate the velocity limit for each mass. equation 1 was created to give similar values to those used in Emma's study while giving a more consistent change between each mass. When an object was held and the velocity of the controller holding the object went above the threshold defined in equation 1 the object dropped from the hand of the user. This forced users to move slowly with the object if it was heavy as the user needed to maintain the velocity below the limit.

The value for the C/D ratio modification ranged from 0.7 to 1.0 where the lightest object with a mass of 1 was given the value of 1.0 so that the virtual movement of the hand would match the real movement exactly. For the heaviest object, the virtual object only moved 0.7 units compared with the real movements of the user's hand. This meant that the object lagged behind the real hand of the user and that the distance between the virtual and real hand increased over the distance that the hand moved. The C/D ratio only came into effect on the y-axis as that only affects the upwards and downward hand movements and if it was applied on all axes it would require the user to move the weight further when going from side to side. The formula used to calculate the C/D ratio for each mass can be seen in equation 2.

$$CDRatio = 0.7 + 0.3 \cdot \left(1 - \frac{mass - 1}{10}\right) \tag{2}$$

The C/R ratio went linearly down from the smallest mass to the biggest. It did not follow the same function as seen in equation 1 due to how small the differences would be for the heavier objects. To be able to distinguish between the weight of objects the linear function in equation 2 is better. The values ranged from 0.7 to 1.0 as Samad et al. [26] used values between 0.7 and 1.3. While their C/D ratio went above 1.0 it was decided not to do that in this study so that the lightest object's displacement would match the real movement exactly and not go beyond it. The LateUpdate function in Unity was used to add the C/D ratio displacement onto the position of the virtual hand. So in every frame when an object was grabbed equation 3 was used to calculate the new position of the virtual hand.

$$\overrightarrow{\text{virtualHandPosition}} = \begin{pmatrix} controllerPositionX - initialPositionX \\ (controllerPositionY - initialPositionY) \cdot CDRatio \\ controllerPositionZ - initialPositionZ \end{pmatrix} + \begin{pmatrix} initialPositionX \\ initialPositionY \\ initialPositionZ \end{pmatrix} \tag{3}$$

In equation 3 $\overrightarrow{\text{initialPosition}}$ is the initial position when the object was picked up, $\overrightarrow{\text{controllerPosition}}$ is the current position of the controller that holds the object, and $CDRatio$ is the C/D ratio from equation 2.

For each task, the mass of each object was randomly chosen from the available masses. No two objects got the same mass and in the two-weight task, a pair was randomly chosen and it was guaranteed that each pair appeared 10 times. In the two-weight task, it was random which mass, from the given pair, each object got assigned. Each time the user dropped an object on the floor, it returned to its original position. This was done to limit participants from having to pick the object up from the ground.

The height of the tables and objects was calibrated according to each participant's height during the tutorial at the beginning of the study. This was done in order to guarantee that all participants got the same experience when lifting the object. If the height had not been calibrated, taller people might have needed to bend down, or shorter people needed to reach up more. The y-position of the headset in the Unity scene was used as a reference point during the height calibration.

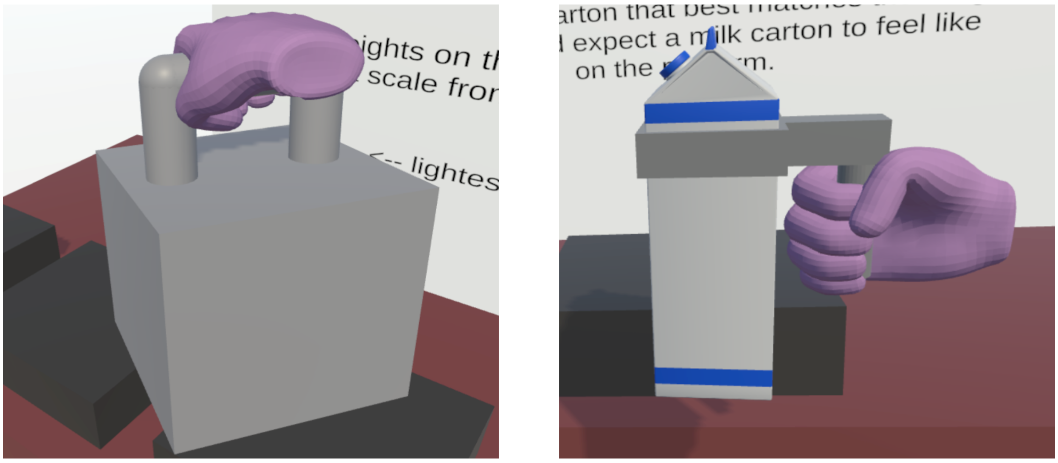

(a) The virtual hand model holding a grey weight. **(b) The virtual hand model holding a milk carton.**

Figure 5: The poses for the virtual hand model when holding an object.

A hand model was used to represent the user's hand. This hand model comes from the Virtual Reality Toolkit v4 package which was used for the virtual reality implementation. For further information on the software used see section 3.6. This hand was used to give a realistic representation of the user's hand and when the user pressed the button to grip an object the model closed its fingers to simulate a closed fist. Two specific poses were made to fit the hand model around the handles placed on top of the grey weights and on the side of the milk cartons. Having the handles and the hand model grip the handles correctly was done to have a somewhat intuitive affordance when grabbing the handle and to allow users to quickly recognize how they should grab the objects. The hand model can be seen holding the grey weight in Figure 5a and the milk carton in Figure 5b.

3.5 Data collection and analysis

Quantitative data was collected about how many times users dropped each object involuntarily after the velocity limit came into effect, the time to task completion from starting a given task until pressing the continue button on the controller, and which object was placed on which platform on the upper part of the table. This data was collected to analyze whether some weight differences were more distinguishable than others and to be able to compare the weight discrimination accuracy between conditions. For example, whether it was harder to distinguish the weight of lighter or heavier objects.

After the user study, the qualitative and quantitative data was analyzed deductively with the hypotheses in mind. Microsoft Excel was used during data analysis along with the Analysis ToolPak add-in to be able to perform statistical tests such as Analysis of Variance (ANOVA).

The Igroup Presence Questionnaire (IPQ) [2] was used to gather data related to the presence and immersion of each participant. The questions in the IPQ are sorted into 4 categories, they are spatial presence (SP), involvement (INV), experienced realism (REAL), and general presence (PRES). One-way ANOVA was used to see whether there were statistically significant differences between the conditions when comparing the answers in each category of the IPQ. Welch's t-test was also used to test the statistical significance between any two conditions as the ANOVA is used to test the significance between all conditions. This was done to see if any of the conditions presented better presence results.

To analyze the qualitative data the framework method was used. The framework method involves five steps, they are familiarisation, identifying the thematic framework, indexing, charting, and finally mapping and interpretation [22, 24]. During the familiarisation stage, the audio transcripts from the interviews were listened to and transcribed into text. When identifying the thematic framework interesting insights with relevance to the research questions were identified. Each insight was translated to English if the original was in Icelandic and paraphrased into shorter versions that summarised the raw data. Similar insights were indexed using color codes to get a grasp of how often each insight appeared in the data. Following that the data was charted where similar insights were put together under groups on one axis, and the participants on the other axis. Finally, during the mapping and interpretation stage, similar groups of insights were put together into themes and sub-themes.

3.6 Software and hardware

The virtual environment was developed using Unity[3] version 2021.3.18f1. The models for the weights, tables, and other scene objects were created using a combination of the default Unity assets and Blender[4]. The Virtual Reality Toolkit v4[5] version 1.3 was used for VR interaction. The Meta Quest 2 was used during the user study and what the user saw was streamed to a laptop of the moderator using the Quest Link application on the headset and screen recorded. The mobile phone of the moderator was placed in front of the participant and recorded the participant doing the study. The interview after the study was audio recorded using the mobile phone of the moderator.

4 RESULTS

The three conditions that were tested in this study, C/D ratio modification, velocity limiting, and the condition with both implemented will be compared in this section. First, the results from the logged data will be compared, and following that the insights from the interviews will be presented. To shorten the names of each condition, in the same order as previously mentioned, the following abbreviations will be used: CD, VL, and BOTH.

Every condition had 10 participants for a total of 30 participants. Four participants that tried the CD condition did not notice any differences when comparing the weights and thus were guessing every time. Similarly, two participants in the VL condition did not attribute the interaction method as weight so they were guessing as well. One participant in the VL condition thought the velocity limit meant that if they dropped it quickly it was light, and if they did not drop it then it was heavy. The participant explained that they thought so because if you are lifting and moving a light item such as a feather, then you have to be careful and move slowly. All of the participants that tried the BOTH condition noticed differences between lifting different objects and used that for weight discrimination.

[2] http://www.igroup.org/pq/ipq/index.php
[3] https://unity.com/
[4] https://www.blender.org/
[5] https://www.vrtk.io/

4.1 Logged data

4.1.1 Six-weight task. In the six-weight task, the BOTH condition had the best accuracy when it came to ordering the objects according to their weight. On average the BOTH condition had 3.7 correct guesses out of 6 while the VL condition had 2.2 and CD 2.5. The number of correct guesses does not tell the whole story as the average distance from the correct guess varied a lot between the conditions. The distance from the correct guess was calculated as the distance from the platform that the object was placed on to the platform it should have been placed on according to its weight. So if an object with mass three was placed on the platform for the object with mass seven, then the distance would be two. The average distance over all guesses per participant for the BOTH condition was 2.6 with a standard deviation of 2.2, while for VL the average distance was 7.2 with a standard deviation of 4.83, and CD had an average distance of 6 with a standard deviation of 3.46. This distribution can be seen in Figure 6 where the bubbles on the diagonal from the bottom left to the top right represent the correct guesses for each object. The diagonal for the BOTH condition are bigger and the furthest incorrect guess was only two platforms away in two cases. The distribution was much greater in the CD and VL conditions. For example, all of the weights were at least once placed on the platform for the heaviest object in the CD condition.

The BOTH and VL conditions had similar numbers when it came to picking up the objects and dropping them. In the BOTH condition, the objects were picked up 53.8 times on average and dropped 32 times on average. The objects were picked up 53 times on average in the VL condition and dropped 32.6 times on average. The objects in the CD condition were picked up 40.5 times on average, and never dropped as what quantified a drop was the velocity limit being triggered. An overview of how many times each object was picked up and dropped can be seen in Figures 13 and 14 in Appendix A.

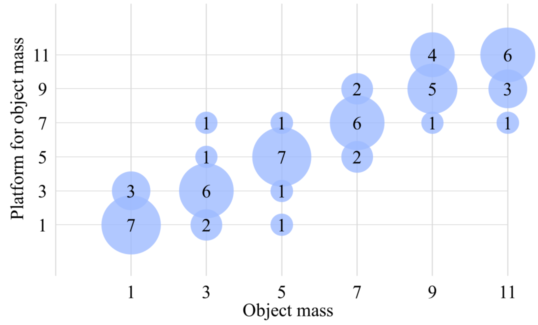

(a) Six-weight task object placement for the BOTH condition.

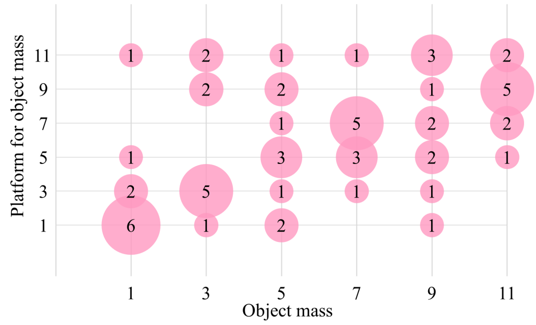

(b) Six-weight task object placement for the VL condition.

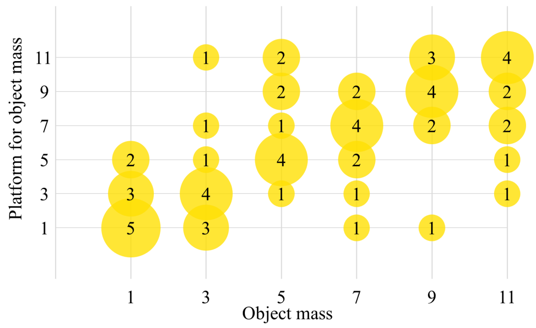

(c) Six-weight task object placement for the CD condition.

Figure 6: How many times a specific object (x-axis) was placed in a specific mass position (y-axis).

As can be seen in Figure 7a, the average time to task completion was similar between the three conditions. The BOTH condition had an average time of 277.77 seconds, VL had 310.33 seconds, and CD had 303.88 seconds. The BOTH condition had the smallest standard deviation of 65.10 seconds. VL had 97.76 seconds and CD 92.46 seconds.

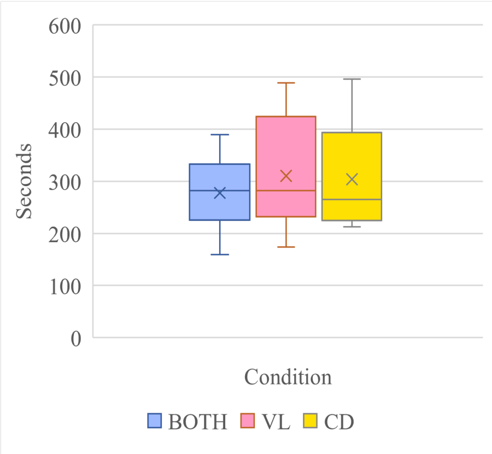

(a) Average time for the six-weight task.

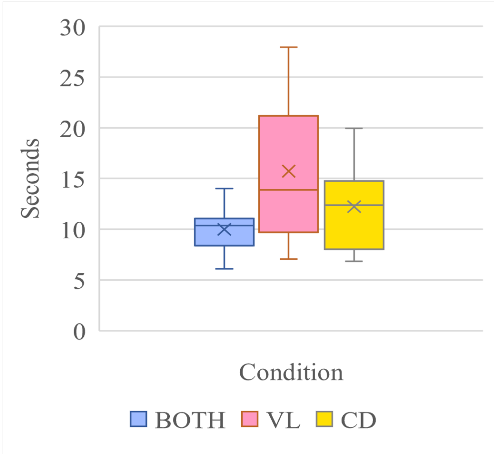

(b) Average time for the two-weight task.

Figure 7: Average time to finish the six-weight and two-weight tasks per condition.

4.1.2 Two-weight task. There were 300 trials for each condition and each participant conducted 30 comparisons. Five trials were accidentally skipped by participants by pressing the button to continue without placing an object on the platform, two of them were in the CD condition, and three in the VL condition. The results can be seen in Figure 8. The CD condition had 228 correct, 70 incorrect, and 2 skipped. The VL condition had 243 correct, 54 incorrect, and 3 skipped. The BOTH condition had 293 correct guesses and 7 incorrect ones.

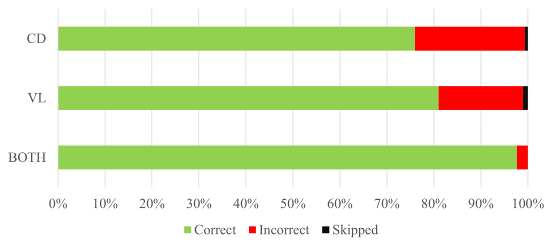

Figure 8: Percentage of correct, incorrect, and skipped trials in the two-weight task for each condition.

As there were 30 trials in total there was a possibility that a learning effect could occur later on in the trials. As can be seen in Figure 9 there was no difference in the BOTH condition, a marginal difference in the CD condition, but a large difference in the VL condition where only 70% of guesses were correct in the first 10 trials, but 91.8% were correct in the last 10. However, as the pair that a participant received each time was randomized, the difference could be caused by the participants receiving the easiest or hardest pair more often in the last 10. In the VL condition, the $5 \Leftrightarrow 7$ mass pair appeared 40 times out of 100 in the first 10 trials and participants had a 62.5% accuracy with that pair in the first 10 trials. In the middle 10 trials, they got it 29 times out of 99 (one skipped) with 79.3% accuracy, and 30 times out of 98 (two skipped) with 83.3% accuracy in the

Table 1: The average percentage of correct guesses for the first, middle, and last 10 guesses for each mass pair based on the condition.

	BOTH			**VL**			**CD**		
	$1 \Leftrightarrow 11$	$3 \Leftrightarrow 9$	$5 \Leftrightarrow 7$	$1 \Leftrightarrow 11$	$3 \Leftrightarrow 9$	$5 \Leftrightarrow 7$	$1 \Leftrightarrow 11$	$3 \Leftrightarrow 9$	$5 \Leftrightarrow 7$
First 10	96.8%	100%	97.2%	75.9%	74.2%	62.5%	77.8%	80%	67.6%
Middle 10	100%	100%	93.1%	85.4%	86.2%	79.3%	85.7%	66.7%	74.3%
Last 10	100%	97.2%	94.3%	96.7%	94.7%	83.3%	86.5%	82.4%	64.3%

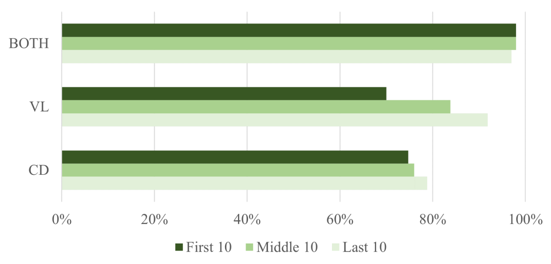

Figure 9: Average percentage of correct guesses during the first, middle, and last 10 trials in each condition of the two-weight task.

last 10 trials. Thus, the accuracy seemed to improve as time went on for the hardest mass pair, but it could be affected by the number of times that the pair appeared.

The three mass pairs ($1 \Leftrightarrow 11$, $3 \Leftrightarrow 9$, $5 \Leftrightarrow 7$) performed differently as can be seen in Figure 10. The percentages in Figure 10 only count the number of correct guesses out of the correct and incorrect guesses, thus the skipped trials are excluded. The $5 \Leftrightarrow 7$ mass pair performed the worst in all conditions, but the BOTH condition performed noticeably better than the other conditions. The VL and CD conditions were a lot closer in the number of correct guesses for all of the pairs compared to the BOTH condition, although the VL condition performed better than the CD condition. The VL condition performed 2.16% better than the CD condition in the $1 \Leftrightarrow 11$ mass pair, 8.9% better in the $3 \Leftrightarrow 9$ mass pair, and 4.73% better in the $5 \Leftrightarrow 7$ pair. The percentage of correct guesses stayed about the same for the $1 \Leftrightarrow 11$ and $3 \Leftrightarrow 9$ masses for the VL and BOTH conditions only going down for the toughest pair ($5 \Leftrightarrow 7$), although the dropoff was more extreme in the VL condition than the BOTH condition. However, the accuracy in the CD condition dropped down for each pair as they got closer together. The best performance was in the $1 \Leftrightarrow 11$ pair and the accuracy decreased for the $3 \Leftrightarrow 9$ pair, and even more in the $5 \Leftrightarrow 7$ pair. The average percentage of correct guesses for each pair during the first, middle, and last 10 pairs can be seen in Table 1.

The average times that each trial took can be seen in Figure 7b. As can be seen, the average (9.88 seconds) and standard deviation (2.04 seconds) were lower in the BOTH condition. The average for the VL condition was 15.72 seconds, and 12.22 seconds for the CD condition. The standard deviation for the VL condition was 6.81 seconds, and 3.87 seconds for the CD condition.

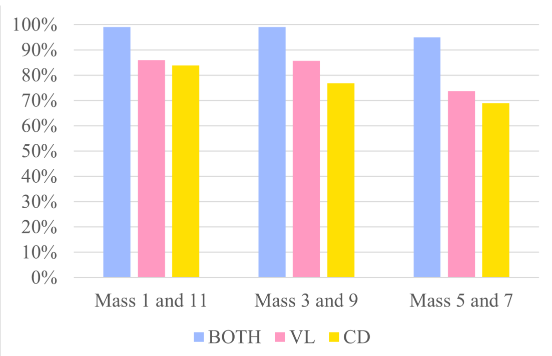

Figure 10: Percentage of correct guesses in the two-weight task for each mass pair in the different conditions.

4.1.3 Milk carton task. The number of times each milk carton was chosen can be seen in Figure 11. All of the participants in the BOTH condition picked one of the three lightest milk cartons, with six of them choosing the lightest. Eight participants in the VL condition chose one of the three lightest milk cartons, with 4 of them choosing the second lightest milk carton. The distribution varied the most in the CD condition, where every milk carton was chosen at least once. However, seven of the participants chose one of the three heaviest.

4.1.4 Presence questionnaire. The results from the IPQ can be seen in Figure 12. A one-way ANOVA was performed for the four categories in the Igroup Presence Questionnaire (IPQ). The categories were compared between the three conditions in this study. The categories were general presence (PRES), spatial presence (SP), involvement (INV), and experienced realism (REAL). The null hypothesis was that there is no difference between the means of the conditions (reject if p-value < 0.05, cannot be rejected if p-value ≥ 0.05), while the alternative hypothesis was that the means differ between conditions.

The results from the one-way ANOVA gave a p-value of 0.35 ($f(2,27) = 1.08$) for PRES, for SP it was 0.12 ($f(2,27) = 2.28$), for INV it was 0.28 ($f(2,27) = 1.32$), and for REAL it was 0.17 ($f(2,27) = 1.91$). Thus, the null hypothesis cannot be rejected and the results showed no differences between conditions.

Welch's t-test was also performed between each condition in all of the IPQ categories. The hypotheses were the same as for the ANOVA with the same p-value to reject the null hypothesis. Only one t-test returned statistical significance with a p-value of 0.02 and that was between the VL and CD conditions in the SP category ($t(18) = 2.55$). Thus the null hypothesis was rejected between those conditions for the SP category, while the null hypothesis could not be rejected in any other test.

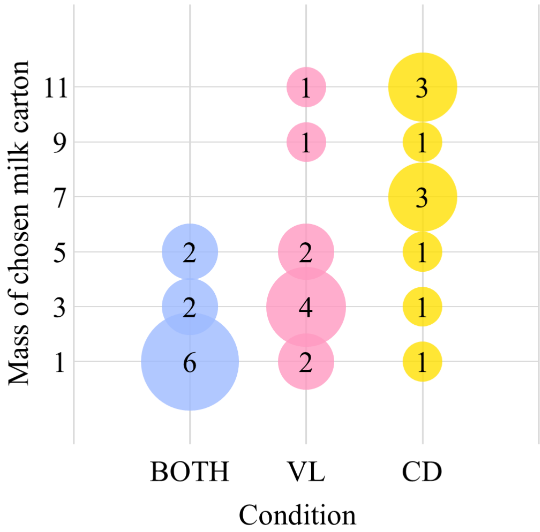

Figure 11: The number of times each milk carton was chosen based on their virtual mass in the milk carton task.

4.2 Interview results

Several themes and sub-themes were identified from the semi-structured interviews using the framework method. The main themes were related to the strategies used in the different tasks, with sub-themes related to the specific strategies, such as using the C/D ratio, using the velocity limit, and using no strategy at all. Other themes that were

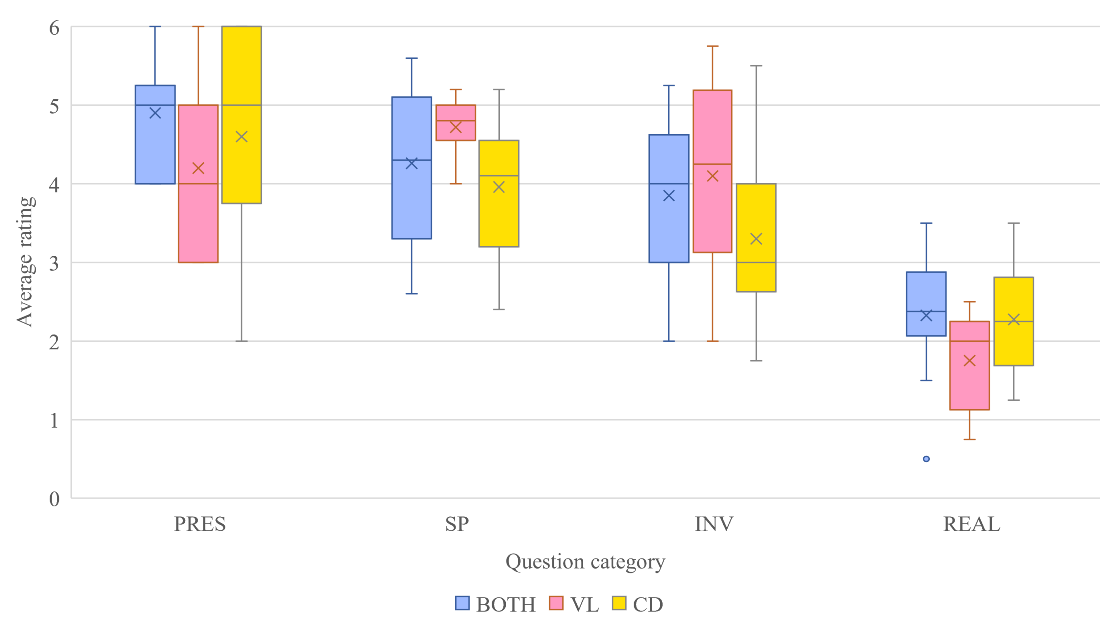

Figure 12: Boxplot showing the average rating within each category of the IPQ.

identified related to the immersion and realism of the interaction methods, such as the slow movement needed to move objects with the velocity limit adding the feeling of lifting objects of different weights. Another theme was frustration caused by dropping the objects when the velocity limit came into effect.

4.2.1 Weight discrimination strategy. All of the participants with the BOTH condition noticed the velocity limit and used that to distinguish the different weights. However, none of them specifically mentioned the C/D ratio modification and did not use that knowingly. The strategy in the BOTH condition for all participants in the six-weight task involved looking at when the objects fell from their hands, for example, participant 21 said:

> "I ordered it depending on how much I felt they [the objects] were falling out of my hands. That was the first thing I noticed so those that were stuck in my hand then I felt they were lighter, otherwise heavier." [translated]

Four participants in the BOTH condition felt like there was almost no difference between the closest weights in the two-weight task.

In the VL condition, eight participants noticed the velocity limit at some point, and seven of them used it in the same way that participants in the BOTH condition used it, that is that heavier objects dropped quickly, while you had to go faster to drop lighter objects. Five participants did not use the velocity limit in the six-weight task. Two of them used no strategy, two tried to distinguish between them using the size or shape of the object, and one participant was trying to see if they affected each other in some way. However, in the two-weight task, three of those participants realized that they could use the velocity limit for weight discrimination. The other two participants noticed that it was there, but thought it was just a glitch in the program and decided to use other methods for weight discrimination. One of those participants went from using the size and shape of the object in the six-weight task to trying to see differences by shaking the weights. The other participant continued to use what they used previously which was to see if they affected each other. One participant that had no strategy in the six-weight task but noticed the velocity limit in the two-weight task used it in the opposite way from what was intended. That participant thought that lighter objects dropped quickly and that they had to be more careful and stable when lifting them:

> "I felt like I needed to be more careful with the ones [objects] that were lighter. I felt like I was dropping them if they were light and I needed to be more stable when lifting them." [translated]

Six participants in the CD condition noticed the C/D ratio modification at some point and used it for weight discrimination. For example, participant 24 said the following when asked to describe the differences they saw when lifting an object:

> "It was the alignment of your real hand versus what it was visually represented. So if it was heavy it was lower than where I knew my hand was when lifting up an object."

Five participants did not notice the C/D ratio modification in the six-weight task. Four of them used no strategy during that task, and one of them tried to use differences when the object was falling down. The participant that was trying to use the differences when falling down realized that the offset was there after 7 trials in the two-weight task. The other four did not realize that the C/D ratio modification was there. One of them used no strategy throughout, while the other three went from using no strategy to using different methods. One of them used a combination of shaking the object, seeing if they affected each other and if there were differences when falling. Another participant also used a combination of shaking the objects and seeing if they affected each other. Finally, one participant tried throwing the objects to see if they flew differently in the air.

Every participant was asked if they could distinguish weight differences between all of the objects in the two-weight task. Four participants in the BOTH condition mentioned that there was a very small difference in some trials and that it was hard to distinguish between them. In the VL condition, two participants mentioned that there were some that were very hard, and three participants said that they could not distinguish between all of them. Five participants in the CD condition could not distinguish between all of the objects, and three of them even mentioned that all of the trials felt the same throughout.

4.2.2 Immersion and realism of the interaction methods. Six participants in the BOTH and five participants in the VL condition mentioned that being forced to move slowly with heavy objects required more effort and concentration which made them feel like they were lifting objects of different weights. Participant 29 from the BOTH condition said the following:

> "I had to put more effort in being able to successfully pick up a heavy object, and I feel like that would be the same case with picking up a real heavy weight."

Four participants in the BOTH condition and three in the VL condition felt that the velocity limit made the interaction feel more real and immersive. Participant 30 in the VL condition said: *"Yeah, I mean it felt quite real for me. So I think it really helped to get a feeling for the weight"*, and participant 29 in the BOTH condition said: *"...you get more actually involved you know, immersed in the world by having these things I guess"*. No participant in the CD condition specifically mentioned that the C/D ratio modification increased realism or immersion. However, four participants mentioned that the interaction method added to the experience and that they would prefer to have some weight representation. This was mentioned more frequently in the other conditions as seven participants in the BOTH condition and six in the VL condition felt that it added to the experience. Some participants mentioned the benefits that this might have in applications that require weight discrimination. For example, participant 17 in the CD condition said:

> "...if you give me like a really heavy dumbbell in VR, and also a bottle of water or something, and then the dumbbell acts like the heavier ones you just showed me, that just gives another perspective of how things are."

A common theme in the conditions with the velocity limiting was frustration caused by dropping the objects. Six participants in the BOTH condition and five participants in the VL condition mentioned frustration to some extent. Some of them mentioned that the frustration was most in the beginning when they did not understand the interaction method, but later on, the frustration went down as they realized that it was used for weight discrimination. For example, participant 11 said:

> "I was lightly frustrated. But then I was like okay maybe that is the weight [...]. Then I was less frustrated." [translated]

Some of the frustration stemmed from the fact that after participants had compared the different objects and had decided which object to put up on the platform, they still dropped it while moving it to the platform. Two participants in the VL condition even mentioned that they thought they were doing something wrong when the objects were dropping from their hands.

Three participants in the VL condition mentioned that they would not expect to drop a heavy object when lifting it and moving it around. Rather that they would not be able to lift it up to begin with or that they would grab the object harder if it was heavy. Another issue that one participant in the VL condition and two participants in the BOTH condition mentioned was that you had to drop the objects to realize how heavy they were. They felt that there was no indication of how heavy they were by just holding the object still or picking it up.

4.2.3 Absolute weight selection. Participants in every condition were asked about why they chose the specific milk carton they ended up putting up on the platform. Four participants in the BOTH and CD conditions and five participants in the VL condition mentioned that one of the milk cartons was just too light for a full 1-liter milk carton and that it felt more like an empty milk carton. However, three participants in the BOTH condition chose the lightest one as they felt that they would not drop a milk carton at any speed. Participant 22 in the BOTH condition said:

> "I felt very unrealistic to drop a milk carton no matter how quickly I would grab it somehow. I felt it was unrealistic that it would just fall from my hands somehow." [translated]

It was also the case that participants felt that some of the milk cartons were too heavy, although for different reasons. In the BOTH and VL conditions, participants immediately eliminated the heaviest ones because they dropped them way too quickly, which is not something that they expected when lifting a milk carton. However, in the CD condition participants were more focused on the feeling of lifting a milk carton, and the heaviest ones did not match that. For example, participant 13 in the CD condition said:

> "I completely put away like the ones that were super heavy. Because it just didn't match with like my real life experiences with a milk carton, like how hard it is to lift it."

4.2.4 Fatigue. Participants were asked at the end of the interview if they felt tired at any point during the study. No participant reported severe fatigue or tiredness, but there was some discomfort in the neck due to looking down because the VR headset was heavy (BOTH: 2, VL: 1). Participants also reported that they felt tired in their hand following the study due to lifting their hand a lot (BOTH 3, VL: 3, CD: 2.). There were also three participants (VL: 1, CD: 2) that mentioned that they got a bit mentally tired during the two-weight task when they realized that there were 30 trials in total.

5 DISCUSSION

In this section, the results will be discussed seeking to answer the research questions from section 1.2 and seeing if the hypotheses from section 1.3 can be confirmed or not. The research questions were:

RQ1: How is the perception of weight in virtual reality affected by limiting the velocity with which objects can be picked up and modifying the C/D ratio when lifting it as measured by the accuracy of weight discrimination?

RQ2: What are the affordances and challenges when adding both velocity limiting and C/D ratio modification compared to adding one of them?

And the hypotheses were:

H1: The hypothesis for RQ1 is that adding both velocity limiting and C/D ratio modification will allow users to distinguish weight differences between lighter and heavier objects, and with better accuracy than with just one of the methods. This is due to the slower velocity required for lifting heavy objects and the fact that the C/D ratio modification should assist the perception of weight through slower virtual hand movements for heavier objects.

H2: The first hypothesis for RQ2 is that having only velocity limiting will cause similar frustrations as Emma Bäckström found in her exploratory study [4]. However, adding C/D ratio modification alongside velocity limiting will address the initial frustrations among users with only velocity limiting. This is because having the C/D ratio modification might assist users in initially identifying that objects have different weights.

H3: The second hypothesis for RQ2 is that it is expected that C/D ratio modification will give the best results with regard to time to task completion as the user will not drop the objects involuntarily due to the velocity limiting which might affect the time.

H4: The third hypothesis for RQ2 is that when presented with a task where people must identify an object according to its absolute weight, there will not be an agreement on which one to choose. This is due to the limitations that these weight perception methods have when it comes to absolute weight compared to relative weight.

RQ1 will be answered in section 5.1, while RQ2 will be answered in all of the following sections as there are various affordances and challenges with the different interaction methods.

5.1 Weight discrimination accuracy

The results from the six-weight and two-weight tasks confirm that the BOTH condition had the best results when it came to the accuracy in weight placement. These results are the clearest when looking at the number of correct answers in the two-weight task where participants in the BOTH condition had only 7 incorrect answers out of 300 in total. There was the possibility that the BOTH condition achieved higher accuracy because of participants taking more time to finish each task and being more careful with their selection. However, the time to task completion says otherwise as the BOTH condition had the lowest average time to task completion and the lowest standard deviation in both the six-weight and two-weight tasks. So the combination of the two interaction methods seems to have enabled participants to perform the tasks with higher accuracy while achieving better times to task completion. This also goes against H3 as the CD condition did not have the lowest times to task completion. The fact that you dropped the objects in the other two conditions seems to not have influenced the time to task completion as originally thought. This is likely due to participants having difficulties discriminating the weight differences in the CD condition which led them to having to do more comparisons.

The VL condition performed better than the CD condition in the two-weight task by 5.31% but performed worse in the six-weight task as both the average number of correct guesses and the distance from the correct guesses were higher. This shows that the VL and CD conditions performed worse when it came to accuracy than the BOTH condition but between them, there was not a clear performance difference because they performed differently in the two tasks.

When looking at the results from the different mass pairs in the two-weight task the BOTH condition had 99% correct in both the $1 \Leftrightarrow 11$ pair and the $3 \Leftrightarrow 9$ pair, and 95% in the $5 \Leftrightarrow 7$ pair. However, the VL and CD conditions never got above 90% and decreased for each pair as their mass got closer together, even dropping below 80% for the masses that were closest together. Thus, the BOTH condition seems to have helped with weight discrimination for weights that were closer together.

Overall these results point towards the combination of the velocity limit and C/D ratio modification having a clear benefit when it comes to weight discrimination compared to having only one of them. This answers RQ1 and confirms H1. Along with that it partially answers RQ2 as it is a clear affordance of this combination that it increases the accuracy of weight discrimination.

5.2 Noticing the interaction methods

A challenge from the CD and VL conditions was that participants did not always notice and use the interaction methods for weight discrimination. Four participants that tried the CD condition did not notice the offset at all. Four participants out of 10 are a significant amount of people and this indicates that the C/D ratio modification might have been too subtle. Similarly, two participants did not attribute dropping the objects to the velocity limit or the weight of the objects, and one participant used it the opposite way from what was intended. Interestingly the two participants that did not use it said that they noticed it and thought it might be what was intended for the weight discrimination. However, they mentioned that they did not count on it for weight discrimination and chose to try to find other strategies to use. This shows that the velocity limit on its own might not be clear enough to be used for weight discrimination.

An affordance of the BOTH condition seems to have been that it reduced the likelihood of participants not noticing the interaction methods being tested. This is due to all participants using it for weight discrimination and noticing it immediately in the six-weight task. Interestingly no participant specifically mentioned the C/D ratio modification and every participant said they used the velocity limit for weight discrimination. Despite this, the C/D ratio modification seems to have played a part in the weight discrimination since the accuracy of correct guesses was much higher in the BOTH condition and every participant was able to distinguish between even the closest weights in the two-weight task. This might be due to some subconscious effect from the C/D ratio modification as the velocity limit was more noticeable since the objects dropped from your hand. This caused participants to focus on the velocity limit in order to distinguish between weights. In tougher situations, the C/D ratio modification might have been used subconsciously by participants to distinguish between the objects that dropped using similar velocities.

There only seemed to be a learning effect in the VL condition When looking at the accuracy between the first 10, middle 10, and last 10 trials in the two-weight task (see Figure 9). As mentioned in Section 4.1.2, the pair that was closest together appeared most often in the first 10 trials for the VL condition. That might have caused participants to struggle earlier on. However, participants did seem to improve later on as the accuracy for the toughest pair went from 62.5% in the first 10 trials to 83.3% accuracy in the last 10. It is interesting that this learning effect was not as prevalent in the CD condition as the accuracy only increased by 4.04% from the first 10 to the last 10. The performance in the BOTH condition was consistently high throughout and even decreased in the last 10 trials. This absence of the learning effect in the BOTH condition might be attributed to the fact that all participants noticed and used the interaction methods in the six-weight task prior to starting the two-weight task, while participants in the other two conditions might have learned to use the interaction method more later on in the trials.

5.3 Comparing immersion and presence

The non-significant results from the Igroup Presence Questionnaire point towards the combination not having a positive or negative effect on immersion and presence as there was no difference between the conditions in the ANOVA results. The only significant result came in the spatial presence (SP) category between the VL and CD conditions and even then the differences in the means were less than 1. If there had been a clear difference between the VL and CD conditions in all categories the combination might have had a more profound effect, but since this was not the case there was no obvious benefit when it came to immersion and presence from combining them. Since the BOTH condition did not negatively affect the results from the questionnaire, the combination can be incorporated without worrying about the effect it could have on the immersion and presence.

All of the conditions achieved high scores in the PRES, SP, and INV categories. The REAL category was considerably lower than the other three, which was to be expected as the virtual environment was simple and specifically designed that way to make participants focus on the task at hand rather than the environment.

5.4 Frustration when dropping objects

The results from Section 4.2.2 showed that dropping the objects based on their weight caused frustrations both in the BOTH and VL conditions. It was mainly due to the unexpected nature of it, as there was no indication that you were about to drop it prior to it happening. As six participants in the BOTH condition and five participants in the VL condition expressed this frustration it appears that the combination of velocity limiting and C/D ratio modification did not reduce the frustration that Emma Bäckström [4] found during her study. Thus, H2 is rejected as it was expected that having the C/D ratio modification would reduce the frustration. Of course, since participants did not specifically notice the C/D ratio modification they could not use that to recognize heavier objects, which might have helped indicate that they would drop the objects. Multiple participants did however mention that the frustration went down as time went on due to them learning what it was for. Thus, if the velocity limit was an established way to distinguish between weights it might reduce the frustration of using it.

5.5 Absolute weight

The results from the milk carton task (section 4.1.3) indicate that the current implementation of the interaction methods cannot be used for absolute weight estimation. The distribution of guesses as seen in Figure 11 shows that participants did not agree on what a 1-liter milk carton feels like. It is interesting however that the majority of participants in the BOTH and VL conditions chose a milk carton in the lighter half of the possible masses, while the majority in the CD condition chose objects on the heavier side. An explanation of this is that participants did not expect to drop a milk carton, no matter how fast they went. This is due to a 1-liter milk carton being quite light in real life. That was the common explanation for their choices, that they eliminated the heaviest ones due to dropping them too quickly. However, multiple participants said that one milk carton was too light and that they did not choose that one because it felt more like an empty milk carton. This means that participants were expecting the milk carton to have some weight and that the lightest one did not feel like it had any weight to it at all. That points towards these methods having the potential to express weight in some way, perhaps not for absolute weight estimation, but at least for expressing weights relative to one another. For example to indicate that one milk carton is empty while another is full.

Thus, absolute weight estimation remains a challenge in VR and these results confirm H4 as participants were not in agreement with which milk carton they chose.

5.6 Method criticism

The implementation of the C/D ratio modification used in this study could have been more exaggerated. Multiple participants did not notice the C/D ratio modification at all in the CD condition and no participant in the BOTH condition mentioned using it for weight discrimination. The values used in the study were motivated by a study done by Samad et al. [26] as the C/D ratio went down to 0.7 in their study. However, they also allowed the virtual hand to go faster than the real hand by having the value go above 1, al. the way up to 1.3. This meant that they had a wider range for testing and larger differences between the objects. Either the values in this study should have gone further down, perhaps to 0.4 or 0.5, or gone above 1 as well.

Another issue with the implementation is that to realize the weight of an object using the velocity limit you must drop it first. One participant in the VL condition and two in the BOTH condition mentioned this issue as there was no indication of weight when just holding the object still. Of course, participants in the BOTH condition could have used the C/D ratio modification to get an indication of the weight without dropping it, but as no participant noticed it that was not the case. Perhaps with exaggerated C/D ratio modification values as mentioned before this would not be an issue.

Some participants mentioned getting tired either in the hand they were using due to moving their hand so much or getting tired in their neck due to looking down. The tiredness in the hand could have been addressed by allowing participants to use both hands in some way. That would however increase the potential of getting different results from using their non-dominant hand and doing direct comparisons by holding two objects at the same time. Getting tired in the neck was due to the lower table being quite far down so participants had to either look down or bend down a bit to grab the objects. This could be addressed by moving the lower and higher parts of the table up so they would be at a height that would not require participants to look down as much.

The variations in strategies when doing the weight discrimination tasks could have influenced the results in some way. The only instructions that participants received were that they should either order the objects according to their weight or choose the lighter of the two, there was no mention of how the weight discrimination should be conducted. This was to limit the chance of knowing the interaction method beforehand. However, clearer instructions that participants should only pick up and put down the weights could have limited this need for different strategies.

5.7 Future work

The combination of velocity limiting and C/D ratio modification has yielded promising results, yet further additions could be made to make the perception of weight even more apparent. A possible addition could be to add some indication when you are about to drop the object because of the velocity limit. The object could for example start to shake when close to dropping it.

Another possible future work would be to require a specific velocity when beginning to lift the object so that the object and virtual hand would not move at all if you lifted it too fast to begin with. This could address the fact that participants could not identify the weight of an object prior to dropping it using the current implementation of velocity limiting.

An interesting future research opportunity that this study presents is to explore absolute weight estimation further. A study could be designed specifically about absolute weight estimation, using more objects and exploring what values of the C/D ratio modification along with the velocity limit work best to represent specific weights.

6 CONCLUSION

The questions this study aimed to answer were whether combining velocity limiting and C/D ratio modification improved the accuracy of weight estimation and what affordances and challenges it might bring. A study environment was designed where the interaction methods could be tested individually and combined. Participants were tasked with ordering and distinguishing between objects according to their weight, and to choose a milk carton based on what they thought a 1-liter milk carton would feel like.

Results indicate that the combination of the two interaction methods increased the accuracy of weight discrimination considerably. It also increased the likelihood of noticing the interaction method with some participants struggling to notice the interaction methods when testing a condition with only one of them present. The velocity limit caused frustration among participants as they did not expect to drop an object when lifting it. The interaction methods did not allow for absolute weight estimation as different milk cartons were chosen between participants. Finally, to improve the interaction methods the C/D ratio modification should be exaggerated and some indication that you are about to drop an object could be added.

REFERENCES

[1] Karl Andersson. 2017. *Manipulating Control-Display Ratios in Room-Scale Virtual Reality*. Student thesis. http://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-211675 Volume: Independent thesis Advanced level (degree of Master (Two Years)).

[2] Bruno Arnaldi, Stéphane Cotin, Nadine Couture, Jean-Louis Dautin, Valérie Gouranton, François Gruson, and Domitile Lourdeaux. 2018. New Applications. In *Virtual Reality and Augmented Reality*, Bruno Arnaldi, Pascal Guitton, and Guillaume Moreau (Eds.). John Wiley & Sons, Inc., Hoboken, NJ, USA, 1–71. 978-1-119-34103-1 978-1-78630-105-5 https://doi.org/10.1002/9781119341031.ch1

[3] Gavin Buckingham, Jonathan S Cant, and Melvyn A Goodale. 2009. Living in a material world: how visual cues to material properties affect the way that we lift objects and perceive their weight. *Journal of neurophysiology* 102, 6 (2009), 3111–3118. Publisher: American Physiological Society Bethesda, MD.

[4] Emma Bäckström. 2018. Do you even lift? An exploratory study of heaviness perception in virtual reality. (2018).

[5] Patrick Carlson, Judy M. Vance, and Meisha Berg. 2016. An evaluation of asymmetric interfaces for bimanual virtual assembly with haptics. *Virtual Reality* 20, 4 (Nov. 2016), 193–201. https://doi.org/10.1007/s10055-016-0290-z

[6] Augustin Charpentier. 1891. Analyse experimentale: De quelques elements de la sensation de poids. *Archives de Phisiologie Normale et Pathologique* 3 (1891), 122–135.

[7] Chih-Hao Cheng, Chia-Chi Chang, Ying-Hsuan Chen, Ying-Li Lin, Jing-Yuan Huang, Ping-Hsuan Han, Ju-Chun Ko, and Lai-Chung Lee. 2018. GravityCup: a liquid-based haptics for simulating dynamic weight in virtual reality. 1–2.

[8] Christopher M Davis and William Roberts. 1976. Lifting movements in the size-weight illusion. *Perception & Psychophysics* 20 (1976), 33–36. Publisher: Springer.

[9] J Edgar De Camp. 1917. The influence of color on apparent weight. A preliminary study. *Journal of experimental psychology* 2, 5 (1917), 347. Publisher: Psychological Review Company.

[10] Fletcher Bascom Dresslar. 1894. Studies in the psychology of touch. *The American Journal of Psychology* 6, 3 (1894), 313–368. Publisher: JSTOR.

[11] David Antonio Gomez Jauregui, Ferran Argelaguet, Anne-Helene Olivier, Maud Marchal, Franck Multon, and Anatole Lecuyer. 2014. Toward "Pseudo-Haptic Avatars": Modifying the Visual Animation of Self-Avatar Can Simulate the Perception of Weight Lifting. *IEEE Transactions on Visualization and Computer Graphics* 20, 4 (2014), 654–661. https://doi.org/10.1109/TVCG.2014.45

[12] Jinwook Kim, Seonghyeon Kim, and Jeongmi Lee. 2022. The Effect of Multisensory Pseudo-Haptic Feedback on Perception of Virtual Weight. *IEEE Access* 10 (2022), 5129–5140. https://doi.org/10.1109/ACCESS.2022.3140438

[13] L. Dominjon, A. Lecuyer, J. . -M. Burkhardt, P. Richard, and S. Richir. 2005. Influence of control/display ratio on the perception of mass of manipulated objects in virtual environments. In *IEEE Proceedings. VR 2005. Virtual Reality, 2005.* 19–25. 2375-5334 https://doi.org/10.1109/VR.2005.1492749 Journal Abbreviation: IEEE Proceedings. VR 2005. Virtual Reality, 2005..

[14] Jun Lee, Jee-In Kim, and HyungSeok Kim. 2019. Force arrow 2: A novel pseudo-haptic interface for weight perception in lifting virtual objects. IEEE, 1–8. 1-5386-7789-X

[15] Woan Ning Lim, Kian Meng Yap, Yunli Lee, Chyanna Wee, and Ching Chiuan Yen. 2021. A Systematic Review of Weight Perception in Virtual Reality: Techniques, Challenges, and Road Ahead. *IEEE Access* 9 (2021), 163253–163283. https://doi.org/10.1109/ACCESS.2021.3131525

[16] Anatole Lécuyer, Jean-Marie Burkhardt, and Laurent Etienne. 2004. Feeling Bumps and Holes without a Haptic Interface: The Perception of Pseudo-Haptic Textures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04)*. Association for Computing Machinery, New York, NY, USA, 239–246. 1-58113-702-8 https://doi.org/10.1145/985692.985723 event-place: Vienna, Austria.

[17] Sharmistha Mandal. 2013. Brief introduction of virtual reality & its challenges. *International Journal of Scientific & Engineering Research* 4, 4 (2013), 304–309. Publisher: Citeseer.

[18] TH Massie. 1994. The PHANTOM Haptic Interface: A Device for Probing Virtual Objects, ASME WAM, Vol. 55. 295–300.

[19] Daniel Mestre, Philippe Fuchs, A Berthoz, and JL Vercher. 2006. Immersion et présence. *Le traité de la réalité virtuelle. Paris: Ecole des Mines de Paris* (2006), 309–38. Publisher: Citeseer.

[20] David J. Murray, Robert R. Ellis, Christina A. Bandomir, and Helen E. Ross. 1999. Charpentier (1891) on the size—weight illusion. *Perception & Psychophysics* 61, 8 (Dec. 1999), 1681–1685. https://doi.org/10.3758/BF03213127

[21] Georg Elias Müller and Fr Schumann. 1889. Ueber die psychologischen Grundlagen der Vergleichung gehobener Gewichte. *Archiv für die gesamte Physiologie des Menschen und der Tiere* 45, 1 (1889), 37–112. Publisher: Springer-Verlag Berlin/Heidelberg.

[22] Catherine Pope, Sue Ziebland, and Nicholas Mays. 2000. Qualitative research in health care: Analysing qualitative data. *BMJ: British Medical Journal* 320, 7227 (2000), 114. Publisher: BMJ publishing group.

[23] Michael Rietzler, Florian Geiselhart, Jan Gugenheimer, and Enrico Rukzio. 2018. Breaking the Tracking: Enabling Weight Perception Using Perceivable Tracking Offsets. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–12. 978-1-4503-5620-6 https://doi.org/10.1145/3173574.3173702 event-place: Montreal QC, Canada.

[24] Jane Ritchie and Liz Spencer. 2002. Qualitative data analysis for applied policy research. In *Analyzing qualitative data*. Routledge, 187–208. 0-203-41308-3

[25] Nina Rosa, Wolfgang Hürst, Wouter Vos, and Peter Werkhoven. 2015. The influence of visual cues on passive tactile sensations in a multimodal immersive virtual environment. 327–334.

[26] Majed Samad, Elia Gatti, Anne Hermes, Hrvoje Benko, and Cesare Parise. 2019. Pseudo-Haptic Weight: Changing the Perceived Weight of Virtual Objects By Manipulating Control-Display Ratio. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*.

ACM, Glasgow Scotland Uk, 1–13. 978-1-4503-5970-2 https://doi.org/10.1145/3290605.3300550

[27] Thomas Schubert, Frank Friedmann, and Holger Regenbrecht. 2001. The experience of presence: Factor analytic insights. *Presence: Teleoperators & Virtual Environments* 10, 3 (2001), 266–281. Publisher: MIT Press One Rogers Street, Cambridge, MA 02142-1209, USA journals-info ....

[28] Mel Slater and Sylvia Wilbur. 1997. A framework for immersive virtual environments (FIVE): Speculations on the role of presence in virtual environments. *Presence: Teleoperators & Virtual Environments* 6, 6 (1997), 603–616. Publisher: MIT Press One Rogers Street, Cambridge, MA 02142-1209, USA journals-info ....

[29] Joseph C Stevens and Lee L Rubin. 1970. Psychophysical scales of apparent heaviness and the size-weight illusion. *Perception & Psychophysics* 8, 4 (1970), 225–230. Publisher: Springer.

[30] Yuki Taima, Yuki Ban, Takuji Narumi, Tomohiro Tanikawa, and Michitaka Hirose. 2014. Controlling fatigue while lifting objects using Pseudo-haptics in a mixed reality space. In *2014 IEEE Haptics Symposium (HAPTICS)*. 175–180. https://doi.org/10.1109/HAPTICS.2014.6775451

[31] Michele Vicovaro, Katia Ruta, and Giulio Vidotto. 2019. Influence of visually perceived shape and brightness on perceived size, expected weight, and perceived weight of 3D objects. *Plos one* 14, 8 (2019), e0220149. Publisher: Public Library of Science San Francisco, CA USA.

[32] Peter Walker, Brian J Francis, and Leanne Walker. 2010. The brightness-weight illusion. *Experimental psychology* (2010). Publisher: Hogrefe Publishing.

[33] Veronica Weser and Dennis R. Proffitt. 2018. Making the Visual Tangible: Substituting Lifting Speed Limits for Object Weight in VR. *Presence: Teleoperators and Virtual Environments* 27, 1 (Feb. 2018), 68–79. https://doi.org/10.1162/pres_a_00319

[34] Harry Kirke Wolfe. 1898. Some effects of size on judgments of weight. *Psychological Review* 5, 1 (1898), 25. Publisher: The Macmillan Company.

[35] André Zenner and Antonio Krüger. 2019. Drag:on: A Virtual Reality Controller Providing Haptic Feedback Based on Drag and Weight Shift. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, Glasgow Scotland Uk, 1–12. 978-1-4503-5970-2 https://doi.org/10.1145/3290605.3300441

A ADDITIONAL RESULTS FROM THE SIX-WEIGHT TASK

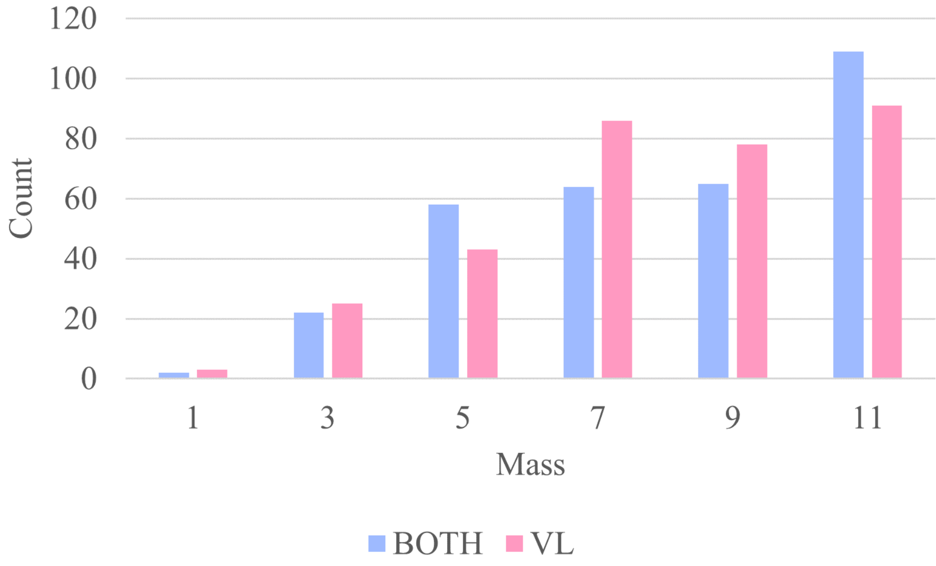

Figure 13: Number of times each object was dropped for the conditions with the velocity limit in the six-weight task.

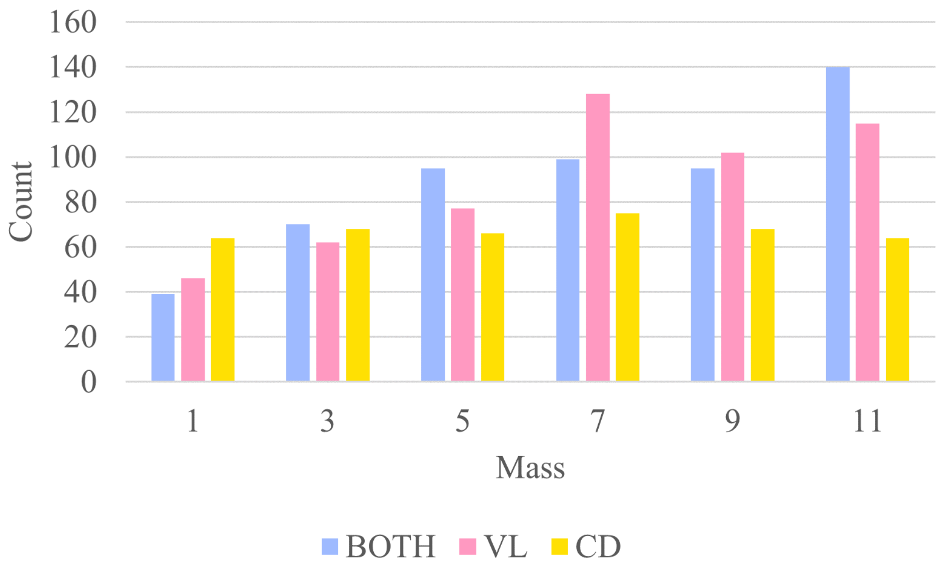

Figure 14: Number of times each object was picked up for each condition in the six-weight task.

TRITA-EECS-EX-2023:505