AALBORG UNIVERSITET

AN EXPLORATION IN USING PASSIVE HAPTICS AND PERCEPTUAL ILLUSIONS FOR REALISTIC VIRTUAL TOOL SIMULATION

Oana Andreea Dogaru
Patrick Lindholt Strandholt

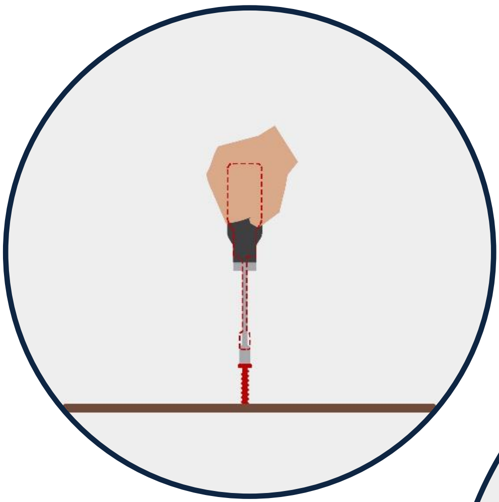

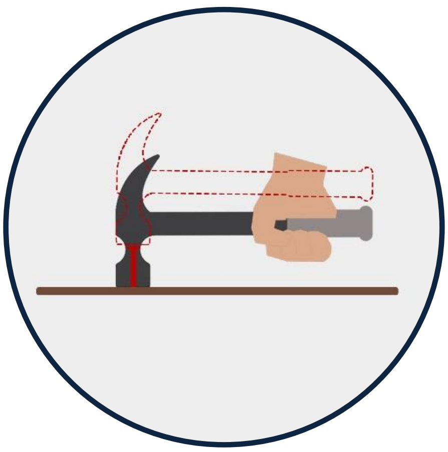

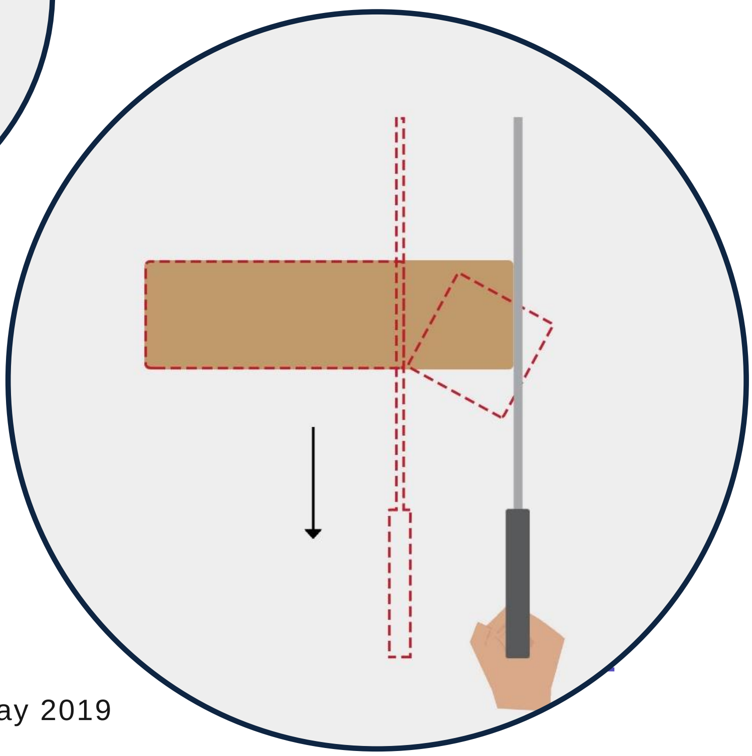

28 May 2019

AALBORG UNIVERSITY COPENHAGEN

Semester:

10th

Title:

An exploration in using passive haptics and perceptual illusions for realistic virtual tool simulation

Aalborg University Copenhagen

Frederikskaj 12,

DK-2450 Copenhagen SV

Semester Coordinator: Stefania Serafin

Secretary: Lisbeth Nykjær

Project Period:
1 February 2019 – 28 May 2019

Semester Theme:
Master Thesis

Supervisor(s):
Rolf Nordahl
Niels Christian Nilsson

Project group no.:
N/A

Members:
Oana Andreea Dogaru
Patrick Lindholt Strandholt

Abstract:

This thesis explore the use of passive haptics and perceptual illusions, and the degree to which they affect realism in a virtual tool simulation. Passive haptics and techniques for perceptual illusions existing research was analyzed and it has shown that both of these aspects can be used to enhance virtual experiences and interactions. Three tools were implemented in three conditions: controller, haptic-air, and haptic-surface. A within-group study was done to find out the extent to which the level of passive haptics and perceptual illusions implemented would affect the participants perceived realism. The results have shown that controller was considered the most unrealistic, pointing that passive haptics through props increased the perceived realism. The most realistic considered was the haptic surface as it provided support for the interaction.

Copies:1
Pages: 86

Contents

1 Introduction 3

1.1 Motivation 4

2 ANALYSIS 6

2.1 Challenges of Realistic Tool Interactions 6

2.2 Human Perception in a Virtual Environment 7

2.3 Criterion of Similarity 8

2.4 Criterion of Co-Location 13

2.5 Criterion of Plausible Interaction 16

2.6 Realism in VR 18

2.7 Design Requirements 21
2.7.1 Final Problem Statement 23

3 IMPLEMENTATION DESIGN 24

3.1 Chosen Tools 24

3.2 Building the Tools 26
3.2.1 Hammer 27
3.2.2 Screwdriver 29
3.2.3 Saw 30

3.3 Environment 33

3.4 Performance Measurement 34

4 METHODS 36

4.1 Experiment Objective 36

4.2 Experimental Design and Setup 36
4.2.1 Sample Management 36
4.2.2 Procedure 37
4.2.3 Setup 38

4.3 Measurements 39

4.3.1 Performance Data 39
4.3.2 Question Sets 40

4.4 Participants 41

5 RESULTS 43

5.1 Quantitative Results 43
5.1.1 Hammer 43
5.1.2 Screwdriver 52
5.1.3 Saw 61
5.1.4 Overall results 70

5.2 Qualitative Results 71
5.2.1 Hammer 71
5.2.2 Screwdriver 71
5.2.3 Saw 72

6 DISCUSSION 73

6.1 Perceived Realism 73

6.2 Passive Haptics 76

6.3 System Issues 77

6.4 Test Validity and Bias 79

6.5 Future Work 80

7 CONCLUSION 81

8 BIBLIOGRAPHY 82

9 ACKNOWLEDGEMENT 85

1 INTRODUCTION

Virtual environments (VE) are on the road to reach even higher fidelity, and virtual reality (VR) hardware such as the HTC Vive and Oculus Rift are still being developed, with new systems coming out this year[1,2], speaking to their longevity. In the last decade, VEs were researched for their possibility of being used as realistic simulations for task performance and learning (Leder, Horlitz, Puschmann, Wittstock, & Schütz, 2019). It has been shown that VR can be used to simulate a range of specific work tasks and job situations, something that can benefit adolescent students. Fewer school students graduating from 9th or 10th grade are choosing manual labour educations (e.g., carpentry) as their next step of education, both compared to those who choose high schools and the educations themselves going back to 2000 (Undervisningsministeriet, 2018). Missing information about the educations stands as one of the reasons for the dwindling choice (Undervisningsministeriet, 2018). This issue inspires us to look into low-cost, low-complexity approaches in assisting the virtual experience of using tools. Therefore a more exploratory technical path is taken in this thesis that does not seek to solve this broader problem. As a result, we ask the question if it is possible to use the current consumer VR hardware, combined with passive haptics in such a way that it provides accurate and realistic simulation in handling of different tools, and whether this performs better than only using controllers? This could pave the way for future full simulations that prove to not need any additional sensors or hardware than what is provided in the off-the-shelf box.

Passive haptics in virtual environments provide several solutions in recent research, that could be applied to create a credible illusion of tool handling. For example, some studies give insight into the usage of props (Simeone, Velloso, & Gellerson, 2015), and provide some information about which properties of an object are more important for a realistic handling of the virtual counterpart. While weight, for example, might have a low impact in a user's ability to suspend disbelief, shape might be of higher importance. However, depending on the number of props being used, it might not always be feasible to have several differently shaped objects for each virtual representation. Taking advantage of the human's visual sense dominance over the other senses, haptic shape illusions (Fujinawa et al., 2017; Kohli, 2013) can be used by having a single prop potentially work for several virtual representations. This approach solves some of the limitations of having to use a multitude of props, as shown by Kohli's (2013) three methods of using remapping and retargeting in order to redirect the user and create the illusion of interaction with several objects, when in reality it is a single one. He suggested rotations of virtual world or translations of the virtual object to align appropriately with their counterparts, using minimal angles that are unnoticeable while the user is moving in the VE. Additionally, world warping is also used in order to create discrepancies between the virtual-real object pairs, to create

[1] Get Ready for Rift S, https://www.oculus.com/blog/get-ready-for-rift-s/
[2] Introducing the VIVE Focus Plus for Premium Standalone VR Experiences, https://blog.vive.com/us/2019/02/21/introducing-vive-focus-plus-premium-standalone-vr-experiences/

mappings capable of redirecting the user's virtual hand without them noticing any differences. The basis of these techniques proved to be useful, as more recent research shows successful applications of redirection and remapping, where users are capable of actions such as stacking cubes or aligning them in rows in the virtual world while handling a single real cube prop (Azmandian, Hancock, Benko, Ofek, & Wilson, 2016).

In the current research project of this thesis, we will create a realistic handling of certain tools, by modifying and combining existing passive haptics techniques together with perceptual illusions. More specifically, we aim to compare interaction of tool handling between standard hardware and the real tool, with implementation of appropriate haptic feedback to simulate the tool's real usage, in terms of realism and performance of the interaction. The intent is to see to what degree realism can be kept by using a low number of passive haptics, while also preserving a realistic representation of the interaction from a manual labor job task (e.g., using a hammer, wrench, screwdriver etc.). We think that this study can help guide future decisions in regards to designing realistic tool simulations.

1.1 MOTIVATION

From 2000 to 2018 a clear tendency is shown in the decreasing number of school students choosing an education in manual labor, in one of the 102 options, when they graduate 9th or 10 grade (Undervisningsministeriet, 2018). The percentage has through the years fallen from 30% of students to 19% and the Danish Government blames lack of insight into the educations amongst the causes for this. This insight is best acquired through trainee programs where the student visits the education and tries out the labor. However, it is not feasible to try that many educations before making a choice. This is a clear inspiration for the work of this thesis. One could imagine a realistic virtual environment (VE), where students could be given the opportunity to try all the labor jobs in virtual reality (VR), before making a decision on where to study. By giving users the ability to try the jobs within a VE, they can quickly gain insight into what each job position is about, how it is to work there, and the type of task they perform. VR is a very valid choice for this as it has been shown useful in teaching people in different subjects (Leder et al., 2019; Chittaro & Buttussi, 2015) even compared to other audiovisual media. Which might be due to its ability to emit a high sense of presence (Slater, 2009) - the illusion of plausible events leading to a more realistic response and experience. Inspired by the research in passive haptics for VR, we want to explore tool simulations, augmented with passive haptics and perceptual illusions, for the purpose of creating a realistic depiction of a tool interaction. The research in passive haptics has also shown that adding several of these haptics - or essentially filling the whole VE with them - increases presence and arguably realism as well. However, including many props and systems in a simulation does decrease the utility and increases the cost of the creation, adding layers of complexity. Therefore it is our wish to keep the amount of passive haptics as low as possible, but still have the interaction feel real. Therefore, it is valid to consider using the current consumer VR hardware, combined with passive haptics in such a way that it provides accurate and realistic simulation in handling of different tools. However, switching the controller with passive haptics is one step in this simulation. The interaction with the tools also needs to be part of the experience. As such some perceptual illusions will have to be created for this to work. This will be the main focus of the thesis. From this an Initial Problem Statement (IPS) can be formed as such:

> *"Can a virtual reality tool simulation augmented with passive haptics and perceptual illusions provide a realistic experience of tool interaction?"*

2 ANALYSIS

The analysis will seek to specify the IPS and takes it into consideration for all of the subjects throughout. Points of focus being virtual reality, passive haptics, and perceptual illusions. Virtual reality and subsequently virtual environments can be used as a tool for play and enjoyment as evident by the many commercial games, but they also provide a platform for teaching people in different subjects in regards to learning outcome (Leder et al., 2019) and knowledge retention (Chittaro & Buttussi, 2015).

Passive haptics have - maybe not surprisingly - shown to have positive effects on the experience of the user. The feeling of presence has shown to be higher in VEs when passive haptics are in use (Barfield, 1995; Witmer & Singer, 1998). Further clarification of the sense of presence by Slater (2009), named “place illusion”, defined as a component of realistic user response in VE, together with “plausibility illusion” has also shown to contribute to a positive and more realistic behaviour in a virtual situation. Passive haptics provide feedback through physical properties such as shape, texture or size, and their relevance is to tie a physical object to a virtual one in the VE. One such study using passive haptics to elicit a fear response when walking up to a ledge (Meehan, Insko, Whitton, & Brooks, 2002). For simple selection tasks, haptic feedback increases presence and task performance (Viciana-Abad, Lecuona, & Poyade, 2010). For using tangible devices between multiple users to manipulate 3D objects compared to other non-haptic techniques, the haptic devices had a higher degree of realism and presence (Aguerreche, Duval, & Lécuyer, 2010). Adding passive haptics to a VE gives users the expectation that other objects in the virtual environment are real (Hoffman, 1998). Haptic feedback in general has seen similar effects. Their usage has seen increased task performance (perceived and measured), and virtual and social presence (Sallnäs, Rassmus-Gröhn, & Sjöström, 2000). Participants training with passive haptics in navigating a maze in a VE, had faster completion time and less collisions, but reported presence did not show significant difference in this study (Insko, Meehan, Whitton, & Brooks, 2001).
There is no doubt that passive haptics has an effect on the VR experience. What this means in regards to a realistic experience of tool interaction will be explored further in the following subchapters.

2.1 CHALLENGES OF REALISTIC TOOL INTERACTIONS

Creating a realistic experience of tool interaction requires certain factors to be taken into account. When considering the baseline of using a standard off-the-shelf controller as a tool for performing a certain task, it is clear that a major factor would be the interaction itself. It is unlikely for an interaction to feel realistic when you see the virtual controller go through surfaces and present no resistance. For a realistic interaction the user would expect to pick up the tool and be able to use it as intended for the desired task, and the tool would display the expected behaviour. Before this, the first step can be to map a 3D model of the tool in the VE. However, there will be clear discrepancies in terms of shape, size, and weight between the controller and the displayed tool, leading to a second factor - the problem of providing proper haptic feedback. Depending on the tool being depicted, the controller's physical properties can be either far from or close to it. In this case, a physical prop - or proxy - would be used instead, in order to provide the passive haptic feedback. The proxy would then match the tool being depicted as closely as possible. However, depending on the number of props used in a scene and how they are mapped, this system can get extensive quickly and could also lead to a mismatch of locations in relation to the proxy-virtual pair placements with each change of the virtual world. This leads to yet another factor of the interaction challenge, that of positioning and alignment.

To summarize, the presented challenges arising from implementing realistic tool interaction can be separated and described as criteria, as similarly done by Lohse et. al (2019). The authors present two criterions that need to be met in order to successfully implement physical props in VR:

(1) *The criterion of similarity*
    In terms of physical features such as shape, weight and size, the physical props used have to be sufficiently alike the virtual objects they serve as
(2) *The criterion of co-location*
    All interactable virtual objects in the scene should be represented by co-located props

These criterions do not mention the challenge of the interaction as described earlier, therefore a third criterion will be added to the list as follows:

(3) *The criterion of plausible interaction*
    The virtual object should interact with the world in a credible manner and behave as their real world counterpart

These criteria will be further analyzed and explored in the following subchapters, where possible and existing solutions will be presented and discussed.

2.2 HUMAN PERCEPTION IN A VIRTUAL ENVIRONMENT

Creating a realistic depiction of an interaction in a virtual environment implies that the user would be responding to it as they would respond to a real life situation of the same interaction. In VR these reactions are prompted by illusions, which have been studied since the 1980s. These are *perceptual illusions*, which are a cognitive phenomenon tied to how humans perceive the world around them. Compared to optical illusions, that are purely an optical phenomenon, virtual reality can provoke a variety of perceptual illusions: e.g., embodiment illusions - users feel their body has been replaced by an avatar; plausibility illusion - feeling that events unfolding in VR are real; or place illusion - feeling they have been taken to a new location (Gonzalez-Franco & Lanier, 2017). Today's virtual reality systems have capabilities of recording wide and accurate sets of data about the user's position, movement, and pose. Additionally they are also able to display and provide varied responses through the HMDs and controllers, making all these illusions possible. Gonzalez-Franco and Lanier (2017) present a compilation of studies in relation to how these VR illusions are perceived by the human mind and how VR systems are able to provoke realistic responses and behaviours. They mention how it has been clearly shown that a virtual world can elicit real responses, because of place and plausibility illusions, together with embodiment illusions if an avatar of the self is present. The parts of human cognition responsible for the extent to which these virtual worlds are perceived as believable and real, are found within three types of processes: bottom-up multisensory processing, sensorimotor self-awareness frameworks, and top-down prediction manipulations. The first process is responsible for picking up available information from all the senses as it comes in, and aggregating it. As the authors note, when multiple senses provide consistent information, it is more prone for the brain to believe it as true. Contrarily, when conflicting information is provided from different senses, some of the information will be chosen as false. In terms of implementation of a virtual reality illusion, one could take advantage of the visual-dominance human cognition is prone to. Additionally, the sense of touch provided by haptic feedback would further enhance the believability of the illusion. Studies and experiments pertaining to this will be discussed in the following subchapters in more detail.

Gonzalez-Franco and Lanier (2017) further explain how the sensorimotor self-awareness frameworks aid to reinforce the illusions and produce an expected reaction from the user. When moving in a virtual world voluntarily, the brain makes predictions of the following states. When the information provided through the senses matches the prediction, the illusion is strongly perceived. Therefore, the strength of the illusion can be possibly shown to increase in relation to the user's expectations of the tool interaction. More clearly, as later mentioned in subchapter *2.5 Criterion of Plausible Interaction*, the tool behaving in a predicted manner when the user is handling it, would be able to solidify the illusion of it being used.

2.3 CRITERION OF SIMILARITY

Passive haptics have in recent times seen a lot of use in VR research, where the focus is on enhancing the haptic experience. In a study on Substitutional Reality (see Figure 1), Simeone, Velloso, and Gellersen (2015) take a look at how great of a mismatch can be had between physical objects and virtual objects before it no longer feels right. They present a layered model of modification. These levels present how closely the physical proxy is related to its virtual counterpart:

- Replica: a true copy
- Aesthetic: differing appearance
- Addition/subtraction: elements have no counterpart, e.g. differing size or parts missing
- Function: mismatch in affordance
- Category: appearance has no connection and affordance is altered

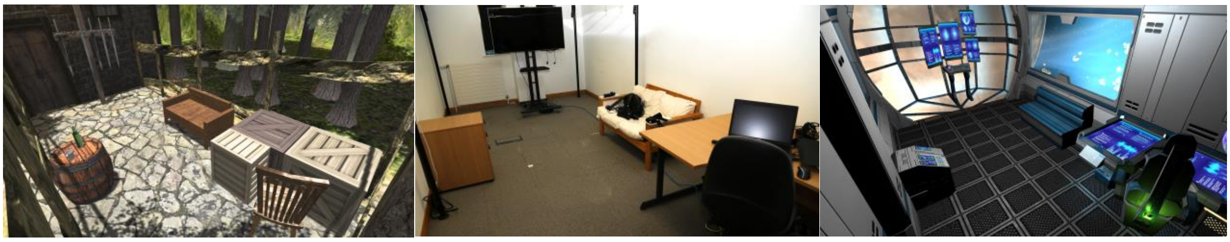

Figure 1. The two substitutional VEs based on the layout of a real room (middle); a medieval courtyard (left); and the bridge of a spaceship (right) (Simeone et al., 2015, p. 3307).

In a user study (n=20) they created several virtual objects in relation to a single physical object (prop), a coffee mug, based on these levels. Participants would interact with each virtual object and rate how similar the virtual object felt to the real one in terms of physical properties such as size, shape, etc. They also asked them to rate the ease of handling the object and how likely they were to believe they were manipulating that object (i.e., suspension of disbelief). They found that some physical properties have more leeway than other. Weight had less significant difference between the objects as compared to shape. Even though they physical felt different, the suspension of disbelief was not significantly altered. In a second user study (n=20), participants were tasked with hitting moving targets with a virtual lightsaber; a physical lightsaber replica acting as the baseline, compared to an umbrella and a torch. The virtual object would stay the same and the prop was changed. Between each trail they would be asked to rate engagement, preference, and exertion. The torch found significant positive results compared to the other props. The author's reasoning is that the lower weight may be the reason for this, making it easier to manipulate, based on comments from participants. In some cases the right proxy is not necessarily a replica - it comes down to expectations of the interactable object. Lightsabers are a work of fiction, we do not know how it would feel to actually wield one. In similar fashion, a convincing illusion of using a longsword would include its weight and it takes time to train to use it and be able to wield its weight, but playing the role of a swordsman in VR, the expectation might be that you know how to use it immediately, making people prefer something lighter.

Rather than using everyday props as proxies, haptic shape illusion can be used to computationally design hand-held VR controllers (Fujinawa et al., 2017). Humans can perceive spatial properties of an object by holding it and thus the researchers believe that a contradiction between the haptic and visual shape perception causes lower immersion and inappropriate object handling in VR. They create a haptic shape perception model based on the assumptions: (1) the perceived shape of a wielded object is represented by its length and width and (2) there exists a mapping between the perceived shapes and specific mass properties (see Figure 2). With this they are able to build haptic objects based on mass properties, such that it is perceived as the desired virtual objects. For example, rather than creating a prop of a sword 1:1, a smaller prop can be created to illicit the same perceived object. A user study was conducted (n=5) using five controllers designed with this model, each tied to a perceived virtual object. The perceived shapes were in turn superimposed on all controllers and participants were told to hold each controller that was laid in front of them, in any order. They were asked to choose the best match. Participants in general chose the corresponding controllers more than others. However, the test used controllers and perceived shapes for rectangular objects, and while they claim the model can be used to design more complex controllers (e.g., a sword, tennis racket, or guitar) it cannot be said how well these would be perceived. This is also something that only seems to become relevant when trying to create proxies for larger objects and depending on the interaction the object is used for, one also has to consider whether or not it can be used (e.g., simplifying a hammer might not make sense).

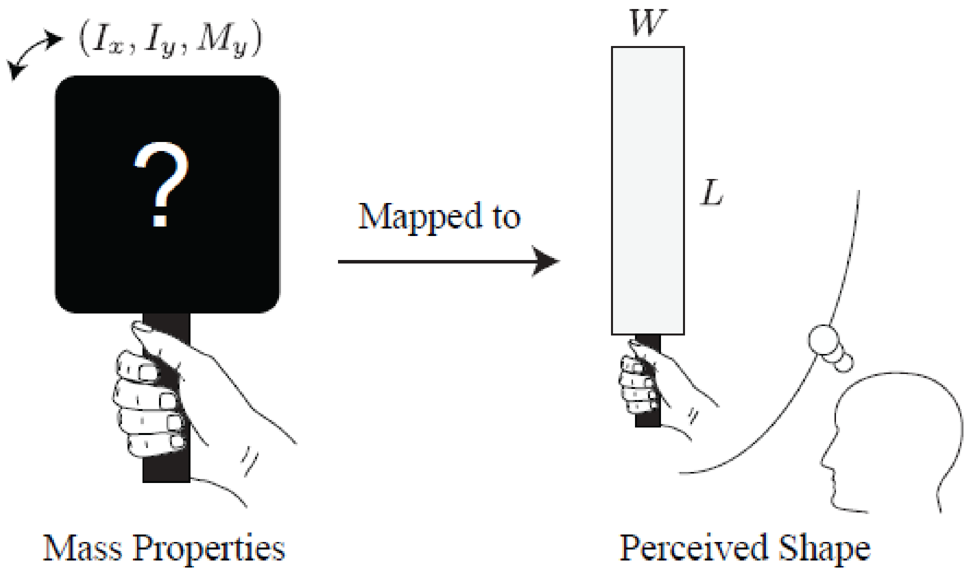

Figure 2. Illustration of the shape perception function that maps the mass properties (three parameters: Ix, Iy, My) to the perceived shape (parametrized perceived length L and perceived width W) (Fujinawa et. al., 2017, p. 31).

Another haptic property illusion - based on weight - has shown that it is possible to manipulate the perception of a mass of an object by controlling the control/display (C/D) ratio (Dominjon, Lecuyer, Burkhardt, Richard, & Richir, 2005). A study has been conducted (n=10), where participants were asked to weigh two virtual balls and say which was heavier. They were presented with a virtual ball on a display screen and a ball prop being dropped down from a mechanism, as seen in Figure 3. The comparison ball was systematically made heavier with four different C/D ratios - meaning that the virtual motions being presented were faster than the real motions made by the participants. The results have shown that the participants' perception of the object's mass was heavily influenced, although not with the same intensity for each individual.

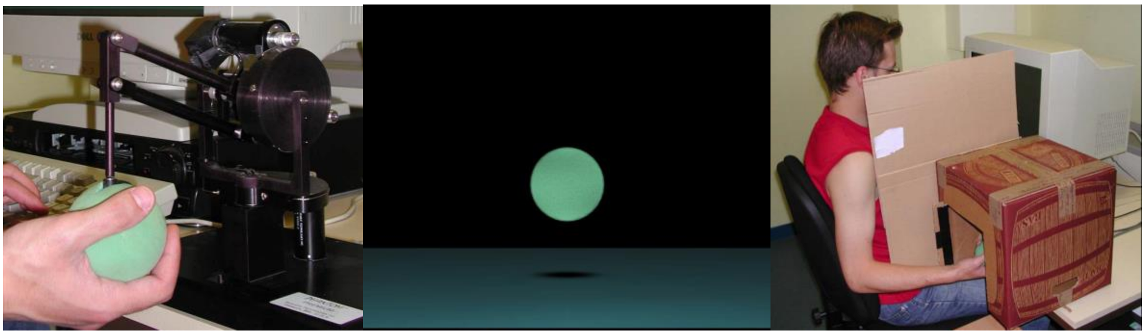

Figure 3. The grasping device setup (left); Screenshot of the visual display (middle); Experiment setup (right) (Dominjon et al., 2005, p. 318).

Another example of a single passive haptic prop being tied to several virtual objects, comes from McClelland, Teather, and Girouard (2017) with their creation HaptoBend. Here a prop was created that can switch between 2D plane-like shapes and 3D multi-surface ones; a rectangle with three hinges, as observed in Figure 4. For a user study (n=20) they created three 2D and 3D virtual objects respectively, and instructed participants to fold the prop into the virtual object they saw. They would handle the “newly” created prop and rate it in terms of goodness - how well the chosen shape would allow control of the object; and ease - ease of creating the chosen shape with the prop. Eight original shapes were created and a frequency of use for each compared to the virtual objects was noted. No shape was convincingly chosen over another, which is backed up by their calculated agreement score. The torch object received a score of 0.489, meaning there was not a high agreement on the prefered shape. Comparing agreement on the choice of 2D vs. 3D shapes the lowest score was 0.605 and highest 0.900 given to the torch. A significant amount agreed on using 3D shapes to manipulate the torch. There is a case for this prop to be able to take the form of many virtual objects, but it does seem that the ambiguity of the prop makes it fair worse than the more specialized props of for example Fujinawa et al. (2017). Unless the adaptability of the proxy is required, choosing a simple prop might be prefered.

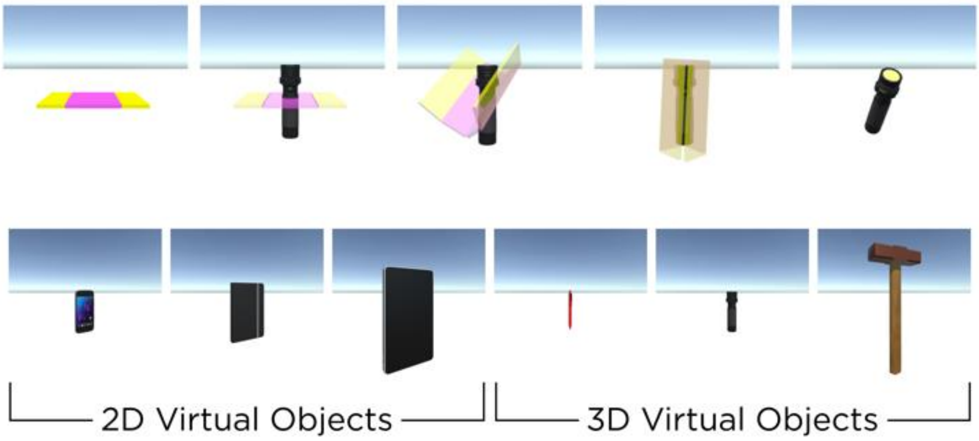

Figure 4. HaptoBend mapping process (top); Virtual objects used in the shape evocation (bottom) (McClelland et al., 2017, p. 85).

The haptic hand method (Kohli, 2013) presented a new approach of using a physical prop. It proposed using the nondominant hand as a haptic feedback interface. This hand would be represented as a panel of larger size than the palm, with different buttons and sliders in virtual reality, that the user has to interact with (see Figure 5). In a short usability study that Kohli (2013) performed, the participants interacted with their non-dominant hand, by having to press buttons and move sliders on the virtual panel. While this would not prove useful in a manual labor scenario where both hands might be actively needed, this experiment has proven that it is possible for users to interact with a virtual object that is shaped differently than its real counterpart. The idea behind this method could be applied when using a simple tool handle prop for various virtual tools of slightly different shapes and sizes.

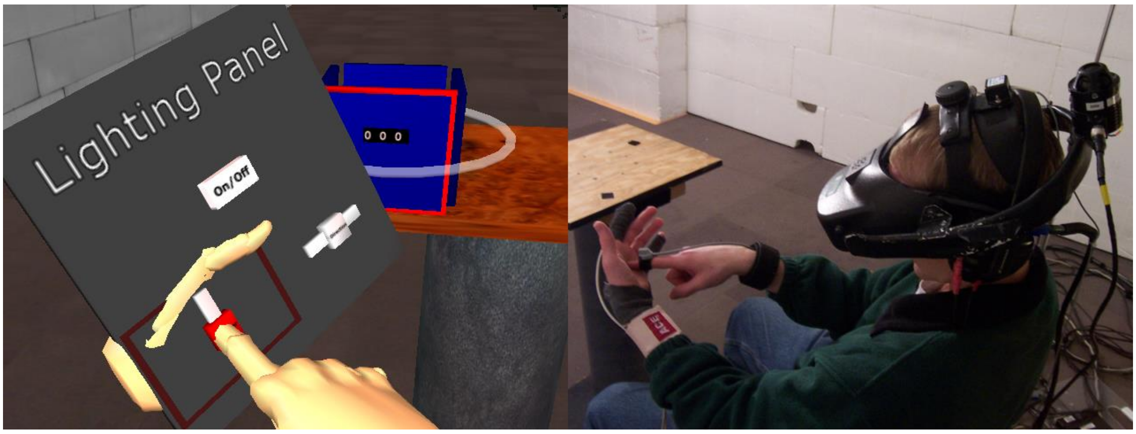

Figure 5. Representation of the virtual panel (left); Participant interacting with his non-dominant hand as a haptic surface (right) (Kohli, 2013, p. 5).

The combination of different redirection techniques has proven useful in providing credible haptic feedback and perceptual illusions in the virtual space. As a study by Matsumoto et al. (2017) has shown, users can even change the shape of an existing virtual object, through combined usage of rotational manipulation and body warping. In their experiment, they were able to provide a realistic illusion that the participants could transform the shape of a table through touching it's edges while walking around the corners. They shaped the virtual table into a triangle and a pentagon, while the real prop they walked around was squared. When creating tool interaction in a VE it is important to consider the environment as part of that interaction. This study show that it is possible to change the expectations of the interaction.

Zenner and Krüger (2017) created a proxy - called Shifty - that can dynamically change its weight distribution upon runtime, leading to an adaptive haptic feedback. The proxy provides a hybrid of passive and active haptic feedback, called *dynamic passive haptic feedback.* Shifty is a rod-shaped device that can change its physical properties - mainly weight distribution - making it able to represent various virtual objects. A sketch of its components can be seen in Figure 6 (left). A two phase experiment was performed by the authors on this device, which aimed to see if the realism and fun would be increased by the adaptive nature of their proxy in two conditions, one with shifting weight and the other with stationary weight at the grip section. Although the authors also evaluated the performance of Shifty. Participants had to extend and retract a virtual telescope step by step as phase one, and change the thickness as phase two, as seen Figure 6 (right).The results have shown that realism was enhanced with a significant difference for the case of the shifting weight, with participants highly preferring the dynamic feedback provided instead of the stationary weight scenario. While the dynamic nature of the haptic might not be relevant for the envisioned tool interaction of this thesis, this paper still show that appropriate weight for the proxy is important for realism (i.e., the one that did not shift weight faired worse), at least in regards to dynamic changes.

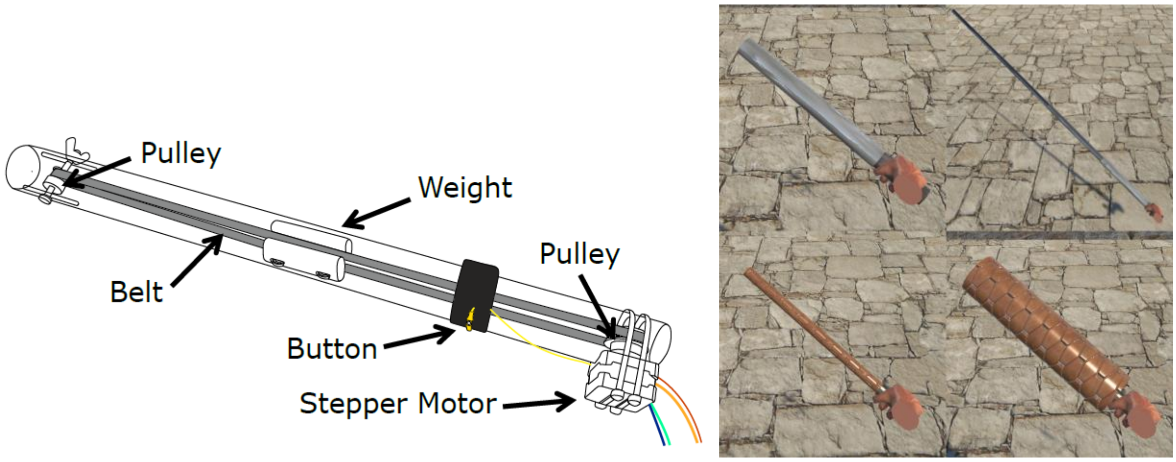

Figure 6. Sketch of Shifty's main components (left); Virtual objects of differing length and thickness wielded in the experiment (right) (Zenner and Krüger, 2017, pp. 1314, 1317).

To summarize, it is clear that it is entirely possible to induce the perception of similarity through proxies that don't exactly match 1:1 the virtual objects they serve as. Instead proxies that are considered just similar enough in shape and size can be used, while in terms of weight the opinions are varied, depending on the specific situation and the task being performed. In the case of this thesis, a focus on using consumer hardware along with minimal props for a scenario of tool interaction can point to the usage of the real tools as props being sufficient. While a custom proxy can be constructed for specific scenarios of tool interaction, able to dynamically change size and weight, it might not be necessary, since it would also add onto the complexity of the system needed for the simulation.

2.4 CRITERION OF CO-LOCATION

In this subchapter, techniques and methods will be explored, in relation to the placement and re-positioning of real-virtual objects pairs.
For the most realistic representation of passive haptics, one would need each interactable virtual object to have a physical prop counterpart. As mentioned earlier in subchapter *2.3 Criterion of Similarity*, this method has obvious limitations in terms of building a system that can have a complex environment with multiple objects.

The research done in this field of haptics contains proposals and systems that combine passive haptics with redirection, resulting in similarly successful methods, through using the human perception to create credible enough illusions of the positioning and movements of virtual objects. A set of three redirection techniques was implemented and tested in studies done by Kohli (2013). He proposed *moving the virtual world* to align the virtual objects to the real one (redirected passive haptics), *moving a virtual object* to align with the real one (the haptic hand, discussed earlier in *2.3 Criterion of Similarity*), and *mapping the real hand motion* to different virtual hand motions (redirected touching, will be discussed in *2.5 Criterion of Plausible Interaction*). The redirected passive haptics implementation by Kohli (2013) has the technique of redirected walking by Razzaque, Kohn, and Whitton (2005) as basis, by using rotation for aligning the desired virtual objects with the real counterpart. The virtual world would rotate at a different rate than the head mount display rotation, by undetectable small angles, in order to redirect the user from one virtual object to the next, when in reality they were being redirected to the same real object. The steps of this process are shown in Figure 7. A test of this method was performed in which the participants were redirected around a small room to touch five different virtual pedestals that were mapped to a single pedestal in the center of the physical room. The results have demonstrated that it is possible to successfully map a single real object to a number of different virtual ones, by using virtual world rotation.

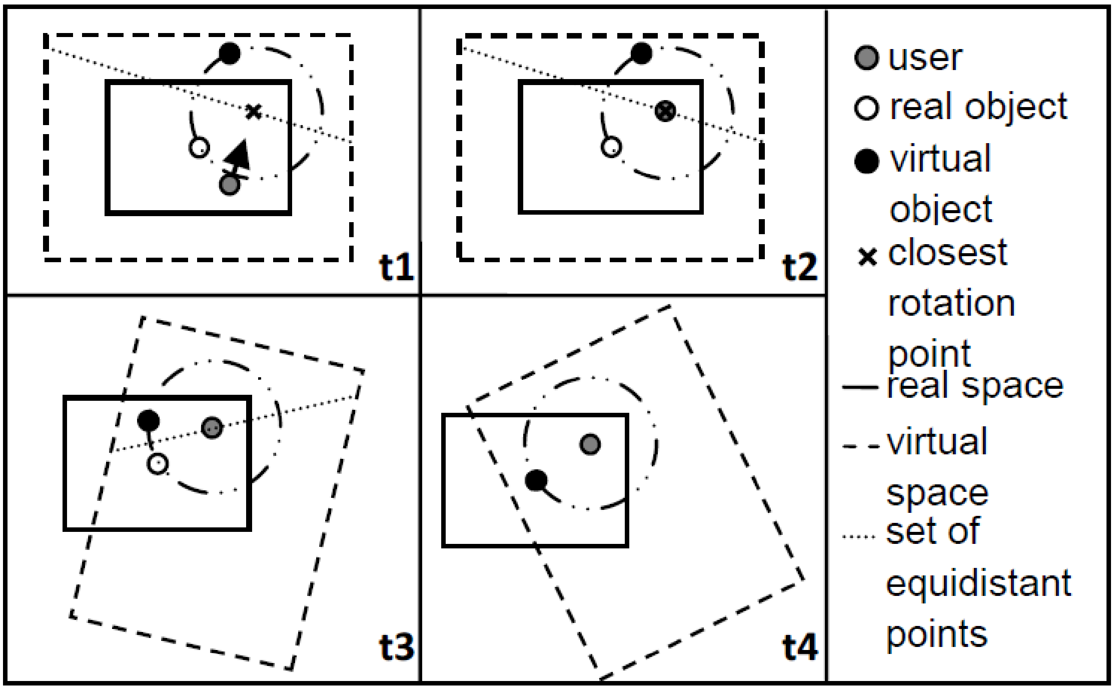

Figure 7. Detailed depiction of each step of aligning the virtual objects with the real objects in the case of redirected passive haptics (Kohli, 2013, p. 14).

Azmandian et al. (2016) developed a proposed solution to the aforementioned limitations of a haptic system - called haptic retargeting. It is based on taking advantage of human perception and the ability to alter it in such a way that unnoticeable changes are made while the user is performing movements. Certain translations of either the virtual body, the environment, or the manipulated object can be made, in order to redirect the user’s touch and sense of placement. The authors have conducted a series of experiments in order to test their approaches in regards to world and virtual self manipulation (see Figure 8). Their setup included a single real life tracked prop - a cube - being mapped to several virtual cubes in the scene. The tests required participants to perform a series of small visual-motor tasks, such as arranging several cubes in a ring, or building a small tower by placing a number of cubes on top of eachother, by using the single cube prop. The results have shown that a hybrid approach worked best: an equal combination of both world and body translation by certain degrees that were unnoticeable by the participants, essentially creating a credible enough illusion of the number of objects being handled and their position. This paper further points towards the concept of haptic retargeting as a useful technique to implement for a case of handling different tools that have similar grips or handles.

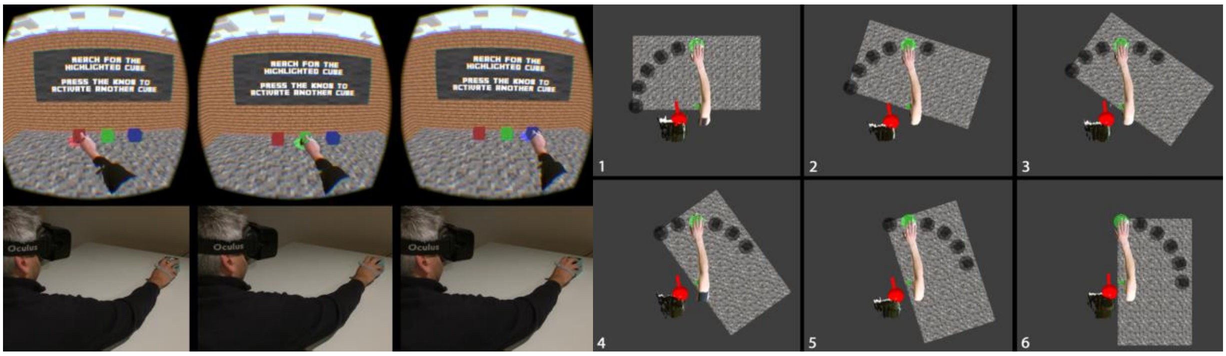

Figure 8. Depiction of Body Warping technique being used to direct the user to three different virtual cubes while using a single cube prop (left); World Warping method showing the mapping of a single cube prop to several virtual cubes arranged in an arc (right) (Azmandian et al., 2016, pp. 1972-1973).

A more recent study done by Han, Suhail, and Ragan (2018) has also proposed *redirected reach*, in order to solve the problem of the real prop not being in the correct position for interaction upon reach. If using one prop for several virtual objects, and if the user is being redirected or realigned from one to the other, it is possible that the real object will be out of position compared to how the user reaches for it. Their implementation applies offsets to the virtual hand’s position, essentially translating it according to how the real hand reaches for the real object. In their test, the authors defined an area for interaction, so that participants would only see the virtual hand when they were in the interaction zone. Upon reaching into this zone, the hand would appear with the translational offsets already applied, to minimize them noticing the difference. A single tool handle for many virtual tools could be used with such a technique, but even if each tool have their own proxy this study shows how much users rely on vision.

To sum up, the research shows that in the case of an interaction between the user’s hand and an object, several techniques can be successfully used to induce the illusion of placement and movement happening at different locations, when in reality they might happen in the same spot. When considering the usage of a tool, in many cases it would have to interact with various objects such as nails, bolts, screws etc. However, having proxies for all of these would increase the complexity of the system needed for such a simulation, therefore implementing software only solutions like the *haptic retargeting* or *redirected passive haptics* methods could allow for a single proxy to serve as an interaction surface, while preserving the perception for the user that interaction with multiple objects is taking place.

2.5 CRITERION OF PLAUSIBLE INTERACTION

A plausible interaction is produced when the the tool being handled behaves in an expected manner, and is able to actually interact with the virtual environment in a believable way. As mentioned in subchapter *2.3 Criterion of Similarity*, props are widely used to provide haptic feedback, both serving as the object being handled or surfaces to interact with. In the case of the current project, the interaction would refer to the tool itself moving and operating as expected, coupled with the objects it is meant to interact with, which may or may not be represented by props as well, but rather represented by perceptual illusions.

Of course, in an interaction between a tool and a surface, one can simulate forces such as friction and resistance to a certain extent. However, for the closest representation of the physics of an interaction, an extensive and complex setup can be used, such as motorized rigs or specially designed custom controllers - like the aforementioned *Shifty* (Zenner & Krüger, 2017) *or HaptoBend* (McClelland et al., 2017). For the usage of a simple consumer controller or just a standard real object prop, the necessary forces coming into play in a tool interaction would be achieved through visual manipulations. The study done by Rietzler, Geiselhart, Frommel, and Rukzio (2018) has shown an implementation of pseudo-haptics, as the authors called it, a software only solution able to communicate changes in resistance and weight. Their approach introduces an offset between the real user’s hand and the virtual representation, on impact with a virtual surface. If the surface is movable, the offset amount depends on the strength of resistance being portrayed. A comparison study was done, between this approach and the method of just letting the virtual hand visually clip through the objects (see Figure 9). As expected, the results have shown that participants showed a higher score of immersion and enjoyment in the pseudo-haptics situation, as well as an increased feeling of touch and resistance.

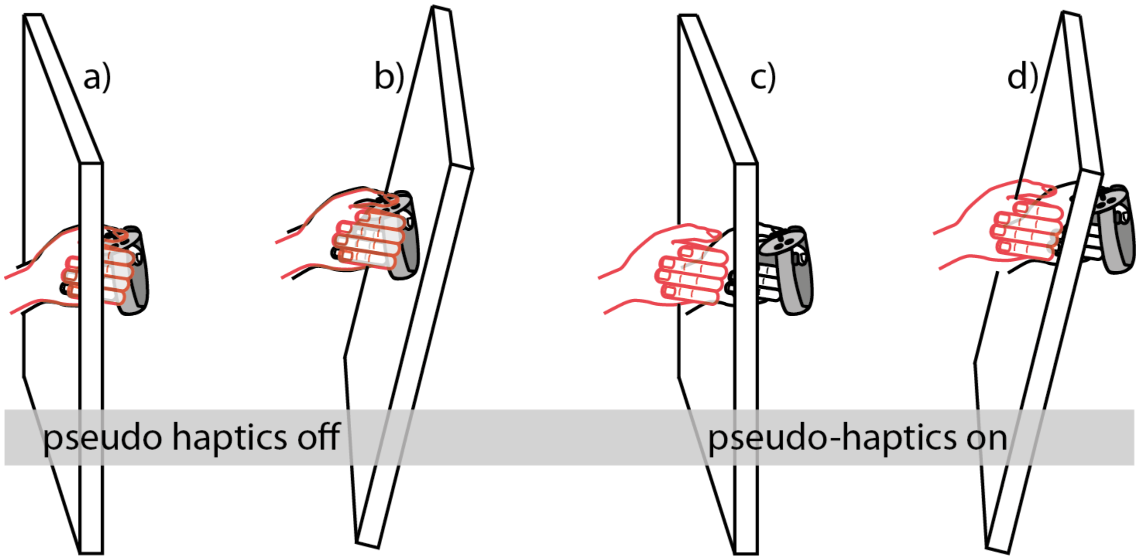

Figure 9. The visual effect of pushing a movable object with no pseudo-haptics (a and b) and with pseudo-haptics (c and d) (Rietzler et al., 2018, p. 463).

A credible illusion of interaction could be produced through using Kohli's (2013) *redirected touching* technique. Taking advantage of the visual sense dominance over the other senses this method allows for differences between virtual and real objects to go unnoticed. This is done through introducing discrepancies between the real hand and the virtual hand positions, such that both hands reach collision with the respective objects at the same time, regardless of the interactable surface's disparity in shape. To illustrate, the author presents a regular easel with a straight surface and a concave easel as the props. Using this technique, when the user would reach his hand to touch the second easel's surface, the virtual hand would be redirected in such a way that the extra distance to reach the surface given by the concave shape would go unnoticed by the user. Figure 10 illustrates this concept.

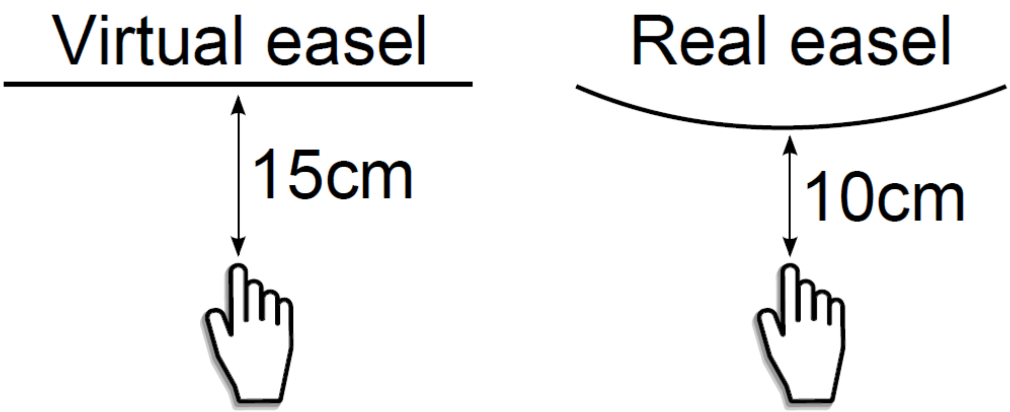

Figure 10. Flat (left) and curved (right) easels seen from above. Real hand will move 10 cm to touch the real easel while the virtual hand will move 15 cm to touch the virtual easel (Kohli, 2013, p. 36)

The implementation is based on warping the virtual world, essentially mapping the distorted virtual space to the real one in a way that introduces gradual unnoticeable real hand position differences to account for the object differences. The study of this method focused on the angular discrepancy of the real-virtual object pair, and how much it affected task performance. Participants had to perform a fast aiming task on a virtual board, placed at various angles compared to the real board (see Figure 11). The measured data from the results shows that seen and felt angular differences of up to 18° provided acceptable task performance when compared to the one-on-one (no discrepancy) condition. It was additionally shown that the threshold would lie at 18°, while discrepancies under 12° went virtually undetected, with the limit being 24° where people would mostly notice. Further studies done by the author in terms of adaptability to this disparity, have shown that training in those conditions was more ineffective, as participants in the discrepant virtual condition performed worse. However, adaptability was apparent after some time, from the participants error rate and throughput, showing that they have adapted to the discrepancy and have later performed no worse than the other virtual 1:1 condition.

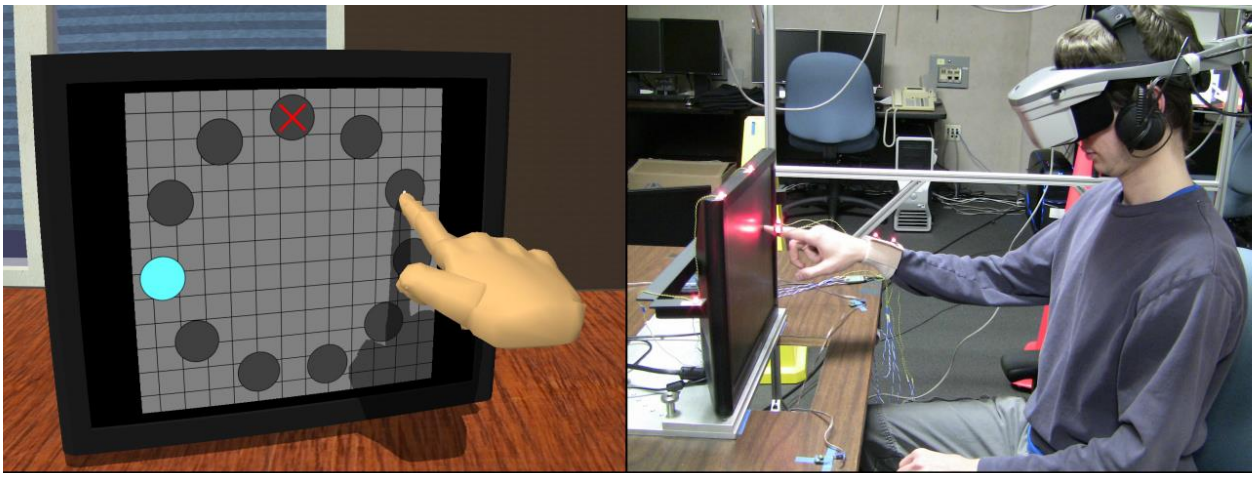

Figure 11. Virtual monitor oriented at 18° (left), while user touches a monitor angled at 0° (Kohli, 2013, p. 79).

Overall, in the specific case of a tool interaction, it is possible to implement these techniques in order to provide credible expectation of a tool usage. Using slightly discrepant virtual representations, some forces can be simulated successfully to certain degrees, offering a satisfactory depiction of how the tool is used in the real world.

2.6 REALISM IN VR

When speaking about a virtual experience that is supposed to provide realism, it is important to note what exactly from that experience has the highest impact on realism, and what subtracts from it. Especially when dealing with a scenario that is showcasing interaction of real life tasks, realism is an important aspect that needs to be investigated. In the first part of this analysis, it was shown how passive haptics affects things such as presence, task performance, and realism. A realistic response in virtual reality is closely tied to the sense of presence, as defined by Slater (2009). Presence is explained as a construct of two components named *"place illusion"* (PI) and *"plausibility illusion"* (Psi). While PI is explained as simply the user's sense of being in a real place; Psi refers to the illusion that the situation being represented is actually occurring. According to Slater (2009), these together lead to a realistic behavioural response to the event taking place in VR. In his research, it is also noted that these illusions depend on the virtual reality system and its affordances. PI is dependent on the level of how much of the sensorimotor contingencies are permitted in VR system, meaning how much of our implicit full body perception can be used and is allowed by the system. Equally, Psi is given by the degree to which the scenarios produced by the system are credible in comparison with the expectations, and the level of response of the VE to the user's presence. It seems that even though there is an awareness that what is being seen and experienced is not real, realistic response can be given when encountering known and expected behaviour in a virtual environment. Such a behaviour was shown by Meehan et al. (2002), where some participants would not walk over the edge or even near it.

Another aspect of realism is given by the graphics (i.e., visual realism). Slater, Steed, and Chrysanthou (2002) mention that the virtual representations of real life objects should have the same geometric properties - shape, size, and ratio - as their real life counterparts, as to be perceived as accurate representations of them. Illumination is another factor being discussed, but more in the context of prototyping and architectural design. Furthermore, Kohli (2013) has noted in his studies that some results of his tests has shown that having shadows for the virtual hands as opposed to having none increased the believability of the object being interacted with - especially when the participant was supposed to be in close contact with the object's surface. A study regarding visual realism conducted by Hvass et. al (2017) has shown that participants can experience a higher sense of presence (PI) along with stronger behavioural response in the form of fear. Their experiment (n=50) has compared two groups, being put in a virtual apartment made to induce fright, one group in a condition of low geometric realism and the other in high geometric realism. Self-reported and physiological measures have shown statistically significant higher sensation of presence and fear response.

Furthermore, examples of haptic realism are also present in studies, such as the study done by Hoffman (1998). He points out that adding physical properties to virtual objects might increase realism of a virtual scene. In his experiment (n=19), two groups were compared and given a virtual plate to interact with. One of the groups were given a real prop while handling the virtual plate, while the other group did not. The results have shown that the group that could feel and touch the real prop while seeing its' counterpart in the virtual world, were able to more accurately predict the physical properties of other objects around the virtual world they were placed in. The author has based this on the human brain's ability to unify disparities coming from several modalities, along with the virtual sense dominance - this allowed the real plates properties to be applied to the virtual object. Overall this experience enhanced the realism of the virtual reality scene, showing that in the case of handling a virtual tool, its prop or custom controller counterpart should provide physical properties as close as possible to the real object, as presented in the *2.3 Criterion of Similarity*.

However, there might be a thing such as too much realism, expressed as a factor of fidelity (Nilsson, Nordahl, & Serafin, 2017). The authors bring forth an interesting discussion about how a high enough fidelity can actually decrease realism. Fidelity, in this case would be comprised of three factors, as defined by McMahan (2011):

- Display fidelity: objective measurement of the displays producing real world sensory stimuli
- Interaction fidelity: user's reaction and interaction with the VE is akin to his real life response
- Simulation fidelity: relating to objective measurement of the system reproducing real world physics and properties.

If their hypothesis - that too much fidelity can actually decrease realism - holds true, it might be beneficial to look at the components of fidelity and adjust them according to the response of the user, because the realism of the experience is ultimately based on the degree to which they find the VE mistakable for real life.

When judging the level of realism of an interaction with a tool in a virtual environment it is relevant to look at interaction fidelity, as defined by McMahan (2011): *"the objective degree of exactness with which real-world interactions can be reproduced".* McMahan's (2011) framework breaks down interaction fidelity into three components, each defined by individual factors, which can be observed in Table 1.

Biomechanical symmetry	Control symmetry	System appropriateness
• Kinematic symmetry • Kinetic symmetry • Anthropometric symmetry	• Dimensional symmetry • Transfer function symmetry • Termination symmetry	• Input accuracy • Input precision • Latency • Form factor

Table 1. Components of interaction fidelity, per McMahan's framework.

The factors in the biomechanical symmetry have to do with the user's movements when they perform a task in virtual reality and the degree to which they match the movements performed in the same task in the real world. Control symmetry is in regards to the amount of control offered by the interaction, again matching the real life counterpart. Lastly system appropriateness deals

with the suitability of the system to implement a type of interaction based on the four factors in the third column (see Table 1).
When translating these definitions into the specific case of a tool interaction in a consumer VR system, the controllers offered by popular consumer hardware allow for a relatively high degree of interaction fidelity to be kept as long as the tools they represent have handles of similar shape and size.

2.7 DESIGN REQUIREMENTS

To summarize, this thesis is exploring if a VR tool simulation using passive haptics can provide a realistic experience in regards to tool interaction. In this subchapter the design requirements for the implementation are formed. Through the analysis we have found that in order for the realism to be high, certain criterions must be fulfilled. First up is the criterion of similarity. The virtual tool must be similar to the proxy. The implementation will use real tools mapped to a virtual replica. Next is the criterion of co-location. The system should avoid being too complex in regards to the number of proxies used. Therefore it should use techniques such as *haptic retargeting* or *redirected touching* to substitute the interactable objects. Last is the criterion of plausible interaction, which makes sure a credible interaction of the used tool is provided, with an appropriate representation of how it works in the real world, meaning the virtual objects should behave as their real world counterpart. As shown in the previous research, many studies uses some sort of surface for the interaction, which will be the case in this implementation as well. Perceptual illusions will be used to simulate the ongoing interaction.
As mentioned, the most realistic representation of passive haptics would need each interactable virtual object to have a physical prop counterpart. To see if this is the case in the implementation, a comparison will be made between a virtual tool mapped to a real tool and a virtual tool mapped to a controller. This implementation will be compared against the status quo of consumer VR. That is, a virtual tool mapped to a controller with no surface, as the controller was never intended for impact with other objects. To see how much adding a real tool alone affects realism, a virtual tool will also be mapped to a real tool with no surface. This leaves us with three conditions: *controller, haptic-air, haptic-surface*. The surface condition will therefore consist of two proxies. All conditions will have the same virtual tool mapped to them.

Based on the research, certain aspects of the implementation design can be outlined. First up is the tools used in the simulation. The choice of tool for the simulation should be chosen from a list of criteria. These criteria stem from the analysis as well as informal discussion between us. They are chosen to streamline the implementation of the thesis and ensure that different aspects are tested. To simplify the style of the VE and make a more coherent experience between the usage of the tools, the tools should be chosen from a single job (e.g., plumber, carpenter). The choice criterias are as follows:

(1) Dimensions. This is meant in both the size, but also form of the tool. It allows for a fairer comparison between the controller and proxies. The size means bigger tools, like some power tools (e.g., a chainsaw), are excluded. The form of controllers shipped with consumer hardware VR systems such as the HTC Vive or Oculus Rift are formed with handles for the user to hold and grab on to. It is important to note that the controller also only requires a single hand to operate. Tools that do not conform to this are excluded as well (e.g., a ruler, wood planer).

(2) Complexity. Some tools are made for complex interactions. Gardening shovels requires a garden with dirt to be dug, plunger requires a toilet with water, etc. If the assets needed to make up the perceptual illusion are too complex to implement or require too many passive haptics, they will be excluded. On the other end of the spectrum, if the interaction of the tool is too simple, it will be excluded as well. This includes such tools where little action is applied to them for the interaction to happen (e.g., tape measure, soldering iron).

(3) Type. Given that this thesis will explore different types of tools in the simulation it is important that these differ in the type of interaction that they are used for. This allows us to test different aspect of perceptual illusions. A way to differentiate could be through the motion made in the interaction (i.e., if the tool is used in motion or is stationary). Another way is through the timing (i.e., if the interaction is constant or brief). These are not binary requirements, but exist on a spectrum and must be far enough from each other to be chosen.

Realism is the focus of the interaction, but it should also be considered when creating the VE itself. The environment should mimic the location it is trying to portray, while being grounded in reality. This pertains to the assets used in the scene that are non-interactable and the assets used for purely decorative reasons. The interactable replicated tools should be realistic in appearance - avoiding stylistic shading and low poly assets where possible.
Additionally, part of a realistic interaction with a tool are the sounds produced upon contact and motion between objects. Real auditory feedback will inevitably be produced in the contact between the surface and tool proxies. However, in order to make the experience and conditions equal in all regards, the virtual auditory feedback will be included for all interactions.

To summarize the requirements:

- The implementation should respect the *criterions of similarity, co-location, and plausible interaction*
- The virtual assets must be non-stylistic replicas of the proxies
- Perceptual illusions should be used to simulate the haptic feedback of the tool interaction, when it's not given by the proxies themselves
- The interaction should match the real life counterpart, following the components of interaction fidelity:
    - User's physical movements in the virtual interaction should correspond with the real life interaction of the same task
    - Virtual interaction should allow the same amount of control for a task as when it is performed in the real world

    - The system used for implementing the interactions should be suitable in regards to precision, latency and input device
- Chosen tools must follow the choice criterias outlined in this section
- The VE should reflect the real world
- All conditions, for each tool respectively, should use the same appropriate audio feedback

It is important that the design of the implementation takes these requirements into consideration.

2.7.1 FINAL PROBLEM STATEMENT

The prior analysis and design requirements can be condensed into a Final Problem Statement:

> *"When comparing the three conditions of controller, haptic-air, and haptic-surface in a virtual reality tool simulation, to what degree can passive haptics and* perceptual *illusions enhance the interaction in terms of realism?"*

3 IMPLEMENTATION DESIGN

As outlined in the *2.7 Design Requirements* subchapter, there are a number of things to consider when creating the implementation. The implementation covers the creation of a simulation for three tools and the three conditions of all of them: controller, haptic-air, and haptic-surface. Though the implementation will mostly focus on the haptic-surface, as the others could be considered simplified versions of this. The design of the implementation was done through exploration while creating the implementation. Therefore there is little definite design to begin with. That is not to say that some things cannot be decided from the beginning. First a specific manual labor job must be chosen. Internal discussion and consideration of the tools, found carpentry to be the most sensible. The field has a great number of tools, most of which people are likely to have some level of familiarity with, perhaps more so than other fields. The next subchapter will discuss the chosen tool.

The implementation will be built using Unity[3], for the HTC VIVE Pro[4] VR system. SteamVR[5] is used to bridge this connection. Unity was chosen due to familiarity with the software and the VIVE due to their separate mountable trackers that can be attached onto the tools. The VR system is an important consideration in regards to the controller size.

3.1 CHOSEN TOOLS

For this implementation it was chosen to implement several different tools to test out different ways of utilising perceptual illusions, rather than focus on a single tool with different versions of such illusions. Having several tools would allow for a more varied approach of implementing the perceptual illusions, because more forces and types of interactions can be simulated in this case. This fit the explorative nature of the implementation. For the scope of this thesis, three tools were chosen.

There are several carpentry tools to look at, but holding them against the choice criterias stated in the design requirements leave only a few. Tools that are a permanent fixture (e.g., a table saw) are excluded from the beginning. Table 2 shows the consideration of several tools against the criterias.

---

[3] Unity, https://unity.com/
[4] HTC VIVE Pro, https://www.vive.com/eu/product/vive-pro/
[5] SteamVR, https://steamcommunity.com/steamvr

Tool	Dimensions	Complexity	Type
Hammer	Handle fits the controller Can be operated in on hand	Interaction requires a nail or another object to hit.	Used in motion Impact.
Tape measure	Size could fit controller, but changes in dimensions through use Operation sometimes requires two hands	Can be used for any purpose	Motion/steady Constant
Utility knife	Size could fit controller One-handed	Requires something to cut through/in	In motion Constant
Chisel	Handle fits controller Two-handed	Requires wood to chisel through Requires something to hit the chisel with	Steady Impact
Level	Doesn’t fit controller One-handed	Requires a surface Simple	Steady Constant
Screwdriver	Handle fits controller One- or two-handed	Requires a screw	Steady Constant
Block plane	Handle fits controller Two-handed	Requires a wooden surface	In motion Constant
Clamp	Hardly fits controller Two-handed to secure	Requires something to clamp	Steady Constant
Handsaw	Handle fits controller One-handed	Requires something to cut through	In motion Constant
Circular saw	Doesn’t fit controller Two-handed	Requires something to cut through	In motion Constant
Power drill	Handle fits controller One-handed	Requires something to drill into	Steady Constant
Electric screwdriver	Handle fits controller One-handed	Requires a screw	Steady Constant
Caulking gun	Handle fits controller Two-handed	Can anywhere	In motion Constant

Table 2. Considered tools against discussed criteria.

From the table, three tools have been chosen: *hammer, screwdriver, and hand saw (saw)*. All fit the dimension criteria, the complexity leaves room for perceptual illusions, and their interaction types are all different. Both the power drill and electric screwdriver could have substituted the screwdriver, but due to the manual nature of both hammer and saw, the screwdriver seemed more fitting. The screwdriver is also often used with two hands, which would be an interesting difference from the hammer and saw.

3.2 BUILDING THE TOOLS

This section will cover the design and implementation of the tools, mostly for the VE in regards to the haptic-surface condition.

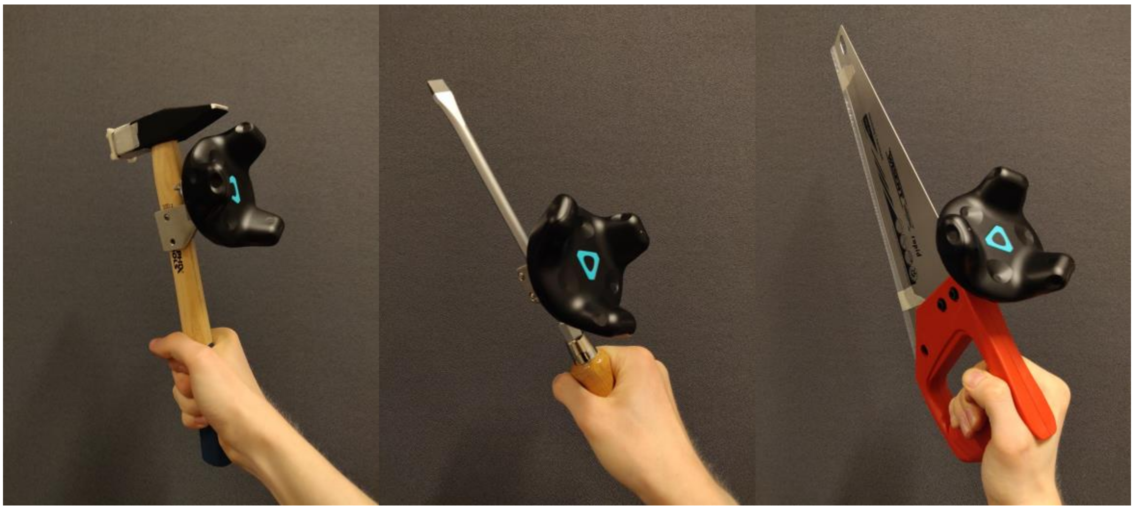

Figure 12. Trackers mounted on the tools: hammer (left), screwdriver (middle), saw (right). .

Figure 13. Assets of the virtual tools representations: hammer (left), screwdriver (middle), saw (right).

In regards to the proxies, all tools were bought at a local hardware store. Since the VIVE Trackers needed to be mounted on, some considerations had to be taken on the size and shape of the tools being used as proxies. While the hammer and saw are generally large enough tools, the screwdriver version chosen for this had to be longer, to have enough space for the mount, for the users to not touch it during handling. In Figure 12, the position of the trackers can be seen. For the hammer, this means that the upper part cannot be held (left figure), but weight and balance is hardly affected. As stated, the screwdriver is large enough such that the tracker could be affixed to the barrel rather than the handle (middle figure). This meant that the screwdriver can be held without issues, but it increases the possibility of throwing off the balance in some cases. Saw has the least issues, as it is placed on top, a part neither held nor used during operation of the tool (right figure). The added weight is insignificant. A small amount of foam was taped to the hammer head to allow for a softer impact due to the loud sound created during interaction. The safety guard was kept on the saw so that the teeth wouldn’t dig into the prop serving as the wood board.

All proxies have a virtual replica (see Figure 13). As decided through the design requirements, non-stylistic assets were chosen. The virtual tools had the same scale in all conditions, based on the scale of the proxies. In the haptic conditions, the virtual tools’ dimensions were made to fit the proxy as close as possible. For the controller conditions, the virtual tools positions were made to fit the handle - such that the hand placement would fit.

The next parts will each dive deeper into the implementation and design of the interaction for each tool.

3.2.1 HAMMER

Figure 14. In-motion depiction of hammer being used. Real world (left), virtual world view through HMD (right).

The hammer interaction will have participants use the hammer to hit nails. For the surface condition, the proxy for the hammer and the table will be real, but the nails will only have virtual representations, in order to apply the perceptual illusions of them being interacted with. As discussed in subchapter *2.3 Criterion of Similarity*, the most realistic setup would have each virtual nail represented by a real one, but the focus of the project is using the least number of props possible, and having the interaction enhanced through perceptual illusions.

In a real world interaction, the hammer would hit the head of the nail several times, each time lower than the last, until the top of the nail is at the same height as the surface. A decision was made that the simulated interaction should perform similarly for an effective realistic feeling to be transmitted. Some of the techniques described in *2.5 Criterion of Plausible Interaction* can be used to simulate the interaction. Inspiration can be taken from redirected touching (Kohli's, 2013) and haptic retargeting (Azmandian et al., 2016). The desired outcome for the interaction is that hitting the head of the nail should feel real. However, in this condition's case, the only thing to impact with is the surface. An offset of the virtual hammer is needed for when it reaches the nail, such that the hammer proxy makes impact with the surface proxy at the same time as the virtual surface. On every subsequent hit, this offsets' distance should be smaller. The offset is thus related to the height between the surface proxy and the virtual nail head. The implementation of this displacement is described in the next part.

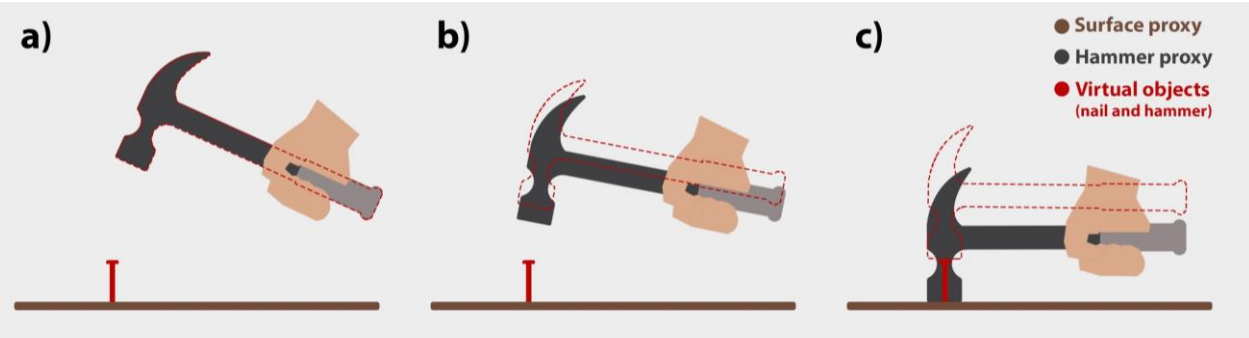

Figure 15. Depiction of hammer displacement: a) approaching the surface, virtual and proxy in the same position; b) within the interaction zone, virtual hammer displaced slightly by the offset; c) proxy touching the surface while virtual hammer touches the virtual nail, full amount of offset added.

Displacement

Creating the algorithm was an iterative process of trying out solutions that would reliably offset the virtual tool the desired amount. The displacement algorithm can be seen in CodeSnippet 1. When the script first runs, it gets a list of all interactable objects and finds the nearest one if there are any. As long as it is not currently interacting, it will look for the nearest object. If the physical tool is within a spherical sector - the interaction zone defined by 45° and 30 cm distance - the displacement will begin. The height of the object, defined as the distance between the physical surface and the virtual surface, is used to set the offset of the displacement. A vector is found moving from the where the surface intersects with the virtual nail, going through the tools' interaction point, out to the sphere's surface. The tools' interaction point is set in the middle of the hammer head. The displaced tool exists on this line, its current amount of offset decided by the height and distance from surface and nail intersection. The offset right at entering, will therefore be 0. Halfway it would be half the offset and when the hammer rests on the intersection point, the offset will be fully added. As only the virtual tool is visible, users will rely on this to hit the nail, feeling a surface when they impact with the virtual nail. When the tool gets close enough to the interactable object - 10 cm - the displacement shifts from the sphere to a direct upwards vector, as the angled displacement gave issues when close to the intersection point. The virtual tool is displaced through a Lerp function to give control over how smoothly is should move.

```
SurfaceToTracker = transform.position - RealSurface.position;

CalculateHeight();
GetAngle();

Vector3 EnterAreaVector = SurfaceToTracker.normalized * AreaEnterAdjustment;
Vector3 EnterAreaPosition = RealSurface.position + EnterAreaVector;
Vector3 EnterToTracker = SurfaceToTracker - EnterAreaVector;

...

float AreaEnterWeight = 1 - (Height / AreaEnterAdjustment);
Vector3 displacedPosition = EnterAreaPosition + (EnterToTracker * AreaEnterWeight);
VirtualTool.transform.position = Vector3.Lerp(
    VirtualTool.transform.position,
    displacedPosition,
    Time.deltaTime * DisplacementSpeed);
```

CodeSnippet 1. Displacement algorithm.

On impact, the nail will move down in five increment steps, based on a curve akin to $y = x^{1/4}$, meaning it will move a lot at the first steps, but less at the last. This is to give some illusion of resistance as the nail digs deeper. On impact, the sound of a hammer impact is played at random from a list, to give more variation on every hit. On non-surface conditions, this will function the same, but the displacement will be off.

3.2.2 SCREWDRIVER

Figure 16. In-motion depiction of screwdriver being used. Real world (left), virtual world view through HMD (right).

The screwdriver interaction will have participants using the tool to tighten screws. For the surface condition, the screws will be the simulated object. Participants will place the screwdriver on top of the screws and rotate to tighten. The same displacement as the hammer will be used, to have the screwdriver hit the physical table proxy at the same time as the virtual tip of the screwdriver hits the the top of the virtual screws. As the screwdriver is rotated clockwise to tighten, the screws will go down and the displacement will follow.

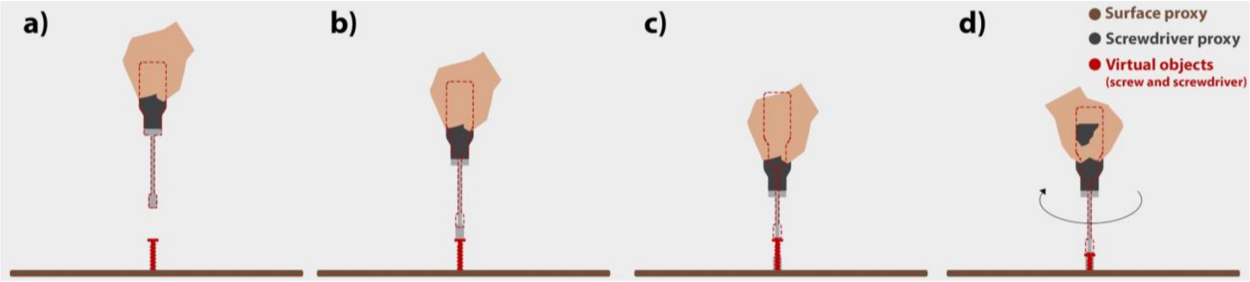

Figure 17. Depiction of screwdriver displacement: a) approaching the surface, virtual and proxy in the same position; b) within the interaction zone, virtual screwdriver displaced slightly by the offset; c) proxy touching the surface while virtual screwdriver touches the virtual screw, full amount of offset; d) proxy being rotated against the surface, virtual screwdriver follows, offset lowering equal to screw height.

The displacement of the screwdriver runs on the same script as the hammer. The tool interaction point is set at the tip instead. The feel is different as the screwdriver requires a more careful aim and the reliance on the virtual tool is higher than the hammer.
The screw is created to follow the rotation of the screwdriver. At every 45 degrees turn of the screwdriver the screw will move downwards. The distance from starting point to finish is given over 20 increments. The amount it moves down uses the same $y = x^{¼}$ curve that the hammer does, meaning the screws' first increments moves it a lot, but changes to less and less over time. Once more, this was done in order to simulate the added friction of the screw as it digs into the wood. In the real world the screw would always move the same amount on an equal turn due to the threads. Instead the screwdriver was envisioned to follow such a curve, requiring more and more real turns to turn the virtual tool, but issues with rotations in the scene made this too difficult. Similar to the hammer, on non-surface conditions the displacement will be off. A sound sample is added to play on a loop when the screwdriver is colliding with the screw and rotating.

3.2.3 SAW

Figure 18. In-motion depiction of saw being used. Real world (left), virtual world view through HMD (right).

For the saw interaction, participants will cut through a wooden board. The surface condition differs from the two prior interactions, and thus uses a different surface proxy to simulate the interaction - the wooden board prop. This is due to the nature of the simulation for this particular tool to be regarded as close to real life as possible. When sawing through a board, the resistance of the wood is felt until the last moment, where the board is cut all the way through and the cut off piece drops down. The resistance immediately disappears at this moment. Because we do want to limit the usage of proxies, the top edge of the table could have been used for this situation, but was later found as unsatisfactory for this simulation, because it would have the participants lean across the entire width of the table in an unnatural and possibly uncomfortable position. It was decided that the risk of this affecting the realism of the interaction was too great, thus the wooden board prop was introduced for this condition. Since there will be no actual cutting taking place on the prop, the most challenging part of the sawing interaction would be for the end of the cutting point, where the resistance disappears. With an incremental offset, we can push the participants hand over the edge the moment the saw goes through the board. That way the drop is felt without anything really dropping.

Unlike the hammer and screwdriver where the intention is to make the impact seem to appear at a different place, the impact of the saw is constant. The end of the interaction stops the contact, and it is where the feeling of dropping and cutting through is happening.

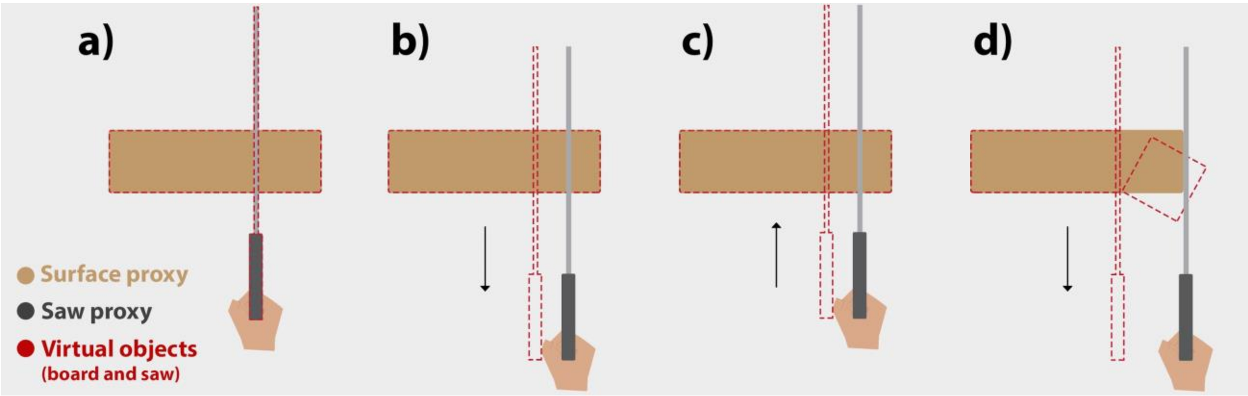

Figure 19. Depiction of saw displacement: a) start of interaction on the marked line, both the proxy and virtual saw in the same place; b) dragging the saw towards one self begins the displacement, user adjust to the marked line; c) sawing with saw proxy, displacement increases as user approaches edge of surface proxy; d) saw proxy reaches edge of surface proxy, virtual saw appearing to have cut through the board, virtual piece drops to the floor.

Displacement

Rather than moving the saw down through the board when sawing and trying to time it with the moment the hand goes over the edge of the board, the cut distance is reliant on how far the hand has moved. The CodeSnippet 2 below shows how the virtual tool's position is set based on a calibrated position from the tracker and an impact vector created from a set of variables. The impact vector is what pushes the virtual tool to redirect the participant's hand.

```
DesiredPosition = Tracker.position + CalibratedPosition;
Vector3 impactVector = new Vector3(xMod, yMod, zMod);
transform.position = DesiredPosition + impactVector;
```

CodeSnippet 2. Redirection algorithm.

For the interaction itself, a collision area defines the cutting space and when the participant begins sawing, a peak detection algorithm created to see when the user is sawing, will move the virtual tool by a given amount to the left on every upstroke - see that the variable xMod is being changed. This direction is hardcoded because the interaction area is always the same. The peak detection works by getting the current height - the saws' position on the y-axis - and comparing it against the last height. If the current height is larger, the saw has moved up since the last position. When sawing in real life, the strength should be put in the upstroke, so this direction was chosen. If the saw is in the collider for cutting the board, the offset can be changed. It ends by setting the last height equal to the current height. The peak detection runs again after 50 ms. This means the algorithm would check for a peak every 50 ms. Other values for displacement were adjusted to fit this, which is 0.75 for the DisplacementModifier. A value of 1 on the axis equates to 1 meter. So here the xMod is set to a base distance of 1 millimeter, adjusted by the modifier to 75 mm.

```
void PeakDetection()
{
    currentHeight = transform.position.y;
    if (currentHeight > lastHeight) //See if we are moving up
    {
        if (SawingAllowed)
        {
            xMod = xMod + 0.001f * DisplacementModifier;
            DropAllowed = true;
            ...
        }
    }
    ...
    lastHeight = currentHeight; //Save the height as what it was for next run
    Invoke("PeakDetection", PeakDetectionInterval); //Run it again,
}
```

CodeSnippet 3. Peak detection. Note that "…" is used where code is left out.

Next, the distance between the virtual and physical tool is calculated and the zMod variable is changed according to this, with an adjustment variable SawSpeed that is set to 1.4. When the distance between real and virtual tool is 50 mm, the virtual tool will have moved 70 mm along the z-axis. Next, if a drop is allowed - as long as the virtual tool is still in the collision area - and the physical saw moves over the edge of the board, the end piece of the wooden board will have gravity enabled and drop to the ground.

```
float boardDistance = StartPoint.position.x - EndPoint.position.x;
float trackerDist = transform.position.x - Tracker.position.x;
boardDistance = Mathf.Abs(boardDistance); //We don't want negative values
trackerDist = Mathf.Abs(trackerDist); //We don't want negative values

if (SawingAllowed && Tracker.position.x < StartPoint.position.x)
{
    zMod = trackerDist * SawSpeed;

    ...
}

if (DropAllowed && SawFixer.position.x < EndPoint.position.x)
{
    ...
    WoodPlank.GetComponent<Rigidbody>().isKinematic = false;
    ...
}
```

CodeSnippet 4. Downwards movement added and drop enacted.

A visual marked line is made to show the progress of the cut and it follows the virtual saw blade position. An audio sample on loop will play when the interaction is ongoing (i.e., the saw is colliding with the board and being moved). These are still used for the non-surface conditions, but the redirection will not be used. The collision areas will not move either, so participants will be able to keep sawing at the edge and see the line move. It is up to them to follow the line if they wish.

3.3 ENVIRONMENT

The tools and their required objects for interaction have already been mentioned. The table, nails, screws, and board assets have all been chosen based on their visuals that are grounded in reality (see Figures 13 and 21). The table and board were calibrated to the real table and board. For the table, this was done by marking the height and four corners, to have the virtual table change its transform after these values. The board was done the other way around. Requiring more precise values for where to cut (i.e., angles and specific positions), the simplest and best solution was to place the real board in the predefined virtual board's position. The environment itself is built to look like a workshop, with enough details to give the feel, but not enough to draw away attention from the place of interaction (i.e., the table with the board, nails, and screws). The room has been built to fit the layout set up in the laboratory for the experiment, which is described in greater details in the *4 Methods* chapter.
A virtual screen is placed on the wall opposite of the table (see Figure 21). This is used to display the questionnaire between conditions, to allow participants to keep on the HMD.

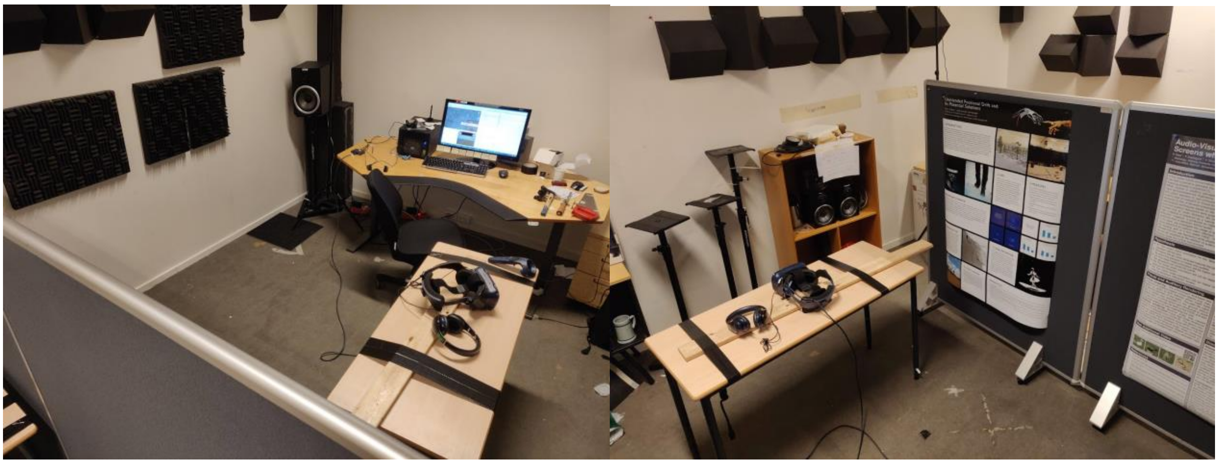

Figure 20. Picture of the test environment setup.

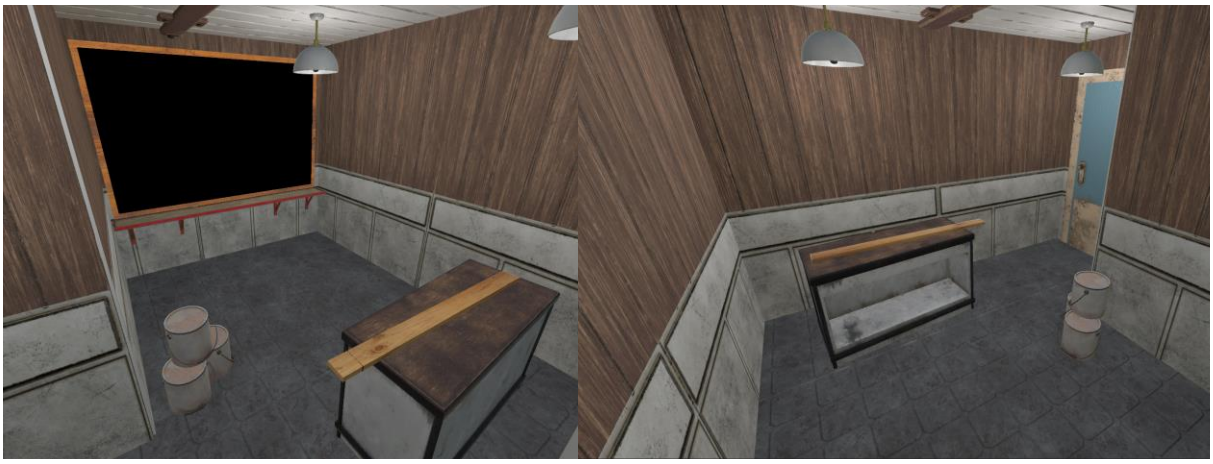

Figure 21. Screenshot of the virtual environment representation.

An environment handler was created so that everything could be controlled during runtime. Displacement could be turned on and off when appropriate, as well as the current tool could be switched. The questionnaire screen would show questions related to the current tool.

3.4 PERFORMANCE MEASUREMENT

There are two kinds of data being captured in regards to the performance measurement: overshooting and precision. Both measures the transform of the virtual objects in certain ways. All of the conditions measure this data a bit differently, but all of them log the data the same way. A call is made to a method holding the transform data and a string with information about the condition, whether it is a controller or proxy, and whether surface is in use or not. All measurement logs capture at 100 ms intervals. This was enough to gather data and also spare the system for writing too much data at once.

Overshooting

- For the hammer condition the data saved is the distance between the nail head and the tip of the virtual hammer head. On collision the current distance is saved as well as for the next 500 ms at 100 ms intervals, making for five entries. Informal testing showed these times to capture a reasonable overshoot, before the hand was being moved up again. The intention is to capture how far below the hand moves after impact.
- For the screwdriver it is also the distance between the screw head and tip of the virtual screwdriver that is saved. However, the data here will be captured for as long as the screwdriver is in contact with the screw. Once again, this happens at 100 ms intervals. The intention how precisely the screwdriver tip matches the screw during interaction.
- For the saw the data is captured in the same way as the screwdriver. The height - here the y-axis - between a midpoint on the blade of the saw and the marked line, is calculated at 100 ms intervals from when the interaction begins. This will indicate if they are sawing with the tip of the saw or closer to the handle, as well as to indicate how far they move their hand.

Precision

- For the hammer, a distance vector on the horizontal plane centered on the nail head is calculated and the x and y coordinates of the hammer in relation to this point - centered on the hammer head - are saved along with the magnitude of the vector. This shows at what position around the nail the hammer hit, as well as the distance itself. This information is saved on impact.
- For the screwdriver this is done similarly, except that once more, the screwdriver saves the data at every 100 ms interval until the tip is out of the collider.
- For the saw, the distance between the marked line and the virtual saw on the x-axis is saved at every 100 ms interval. If the saw is at the left of the marked line, it will result in a positive value; and a negative value if it is at the right.

On every collision, the name and current time is saved before the impact data. This is done to see how often the interaction would start over (e.g., the hammer is lifted for a second hit or the screwdriver misses the screw and must readjust).

It was considered for the system also to record location and rotations at similar intervals, but the system ran noticeably slower at times while this was recording. This goes against the interaction fidelity requirement on latency put forth in *2.7 Design Requirements* and was therefore decided against.

4 METHODS

This chapter will describe the set of methods and strategies for answering the Final Problem Statement.

4.1 EXPERIMENT OBJECTIVE

As mentioned previously, the main question lies in regards to the enhancement of the virtual tool interaction through the passive haptics and perceptual illusions, and the extent to which the experience differs in terms of realism, when using standard controllers versus different amount of proxies. The choice of the specific tools, as explained in the subchapter *3.1 Chosen Tools*, was made based on the choice criterias seen in Table 2. When focusing on the realism of the interaction in itself, these tools differ greatly in terms of movements needed to perform the tasks with them, the level of resistance encountered and the type of force.

4.2 EXPERIMENTAL DESIGN AND SETUP

The test will use the within-group experimental design. Each tool will be tested in a block of three comparison conditions. Each block is considered a separate test as no direct comparison is made between the tools themselves. In order to preserve validity, the order of the tool blocks, as well as the order of the conditions within the blocks, will be completely randomized for each participant.

As previously mentioned, an implementation of each tool interaction is made, using the techniques of redirection, translation and rotation, in order to create credible perceptual illusions of the interaction taking place with minimal use of proxies. Therefore, a three level comparison will be conducted to test three conditions for each tool interaction as follows:

- Controller: Standard controller serving as the proxy for the tool, with no surface (in air)
- Haptic-Air: Real proxy serving as the tool (in air)
- Haptic-Surface: Real proxy serving as the tool, with real proxy serving as a surface for the interaction

The reason for not including a condition for using the controller with a surface proxy is to avoid damage to the controller, since it was not designed for hard impacts or repeated friction against surfaces.

4.2.1 SAMPLE MANAGEMENT

Participants will be invited to the test through a preliminary questionnaire that will be posted on social outlets and network groups for technical students and graduates. The questionnaire only inquires about demographic information and their level of experience within VR. Additionally, more will be chosen through convenience sampling by asking students and staff on campus. It is not necessarily expected that the participants would have experience with VR.

4.2.2 PROCEDURE

The test will begin with a participant filling out an initial questionnaire containing basic demographic questions and information regarding their previous experience in VR as well as usage of the three tools implemented.
The testing will be done in three major blocks, each block representing one of the tools. The sequence of the tasks in a block is represented in Figure 22. A more detailed description follows below.

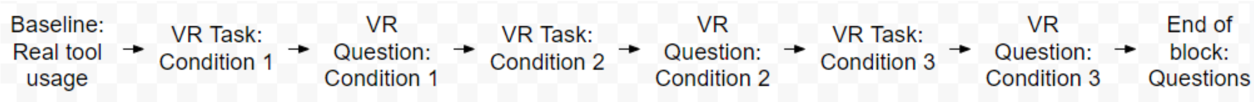

Figure 22. Sequence of the tasks and condition blocks for the experiment.

A baseline of real tool usage will be used, therefore each participant will start out by performing a simple task, depending on the tool in question, as follows:

- Hammer: Hammer down a nail into a wooden board
- Screwdriver: Screw down a screw into a wooden board
- Handsaw: Cut approximately 5 cm into a wooden board

Figure 23. The baseline area setup and tools used.

The reason for this is for them to get a reasonable baseline of real world experience in using these tools, regardless of their previous level of experience. This is done in order to have a fresh recollection of the forces and movements felt in the interaction, in order to allow the participants to give informed answers and make an accurate comparison with the virtual experience.
The next stage of the test is the first block. The participants are handed the VR headset and instructed in their task. Whilst in the VE they are also handed the tool for the first condition. They should perform the task related to the specific tool of the block.

The three virtual reality tasks are as follows:

- Hammer: Hammer down three virtual nails into the table surface
- Screwdriver: Screw down two screws into the virtual wooden board
- Handsaw: Cut along a marked line all the way through the virtual wooden board

After the task is finished, they will answer a set of questions, displayed on a wall in the virtual world (seen in Figure 24). The questions are in regards to the interaction they performed and rating of the realism. Table 3 in the subchapter *4.3 Measurements* shows the exact questions. The same sequence of task and questioning repeats twice more for the other two conditions of a block. At the end of a block, participants have to take off the VR headset and answer the after-block questions.

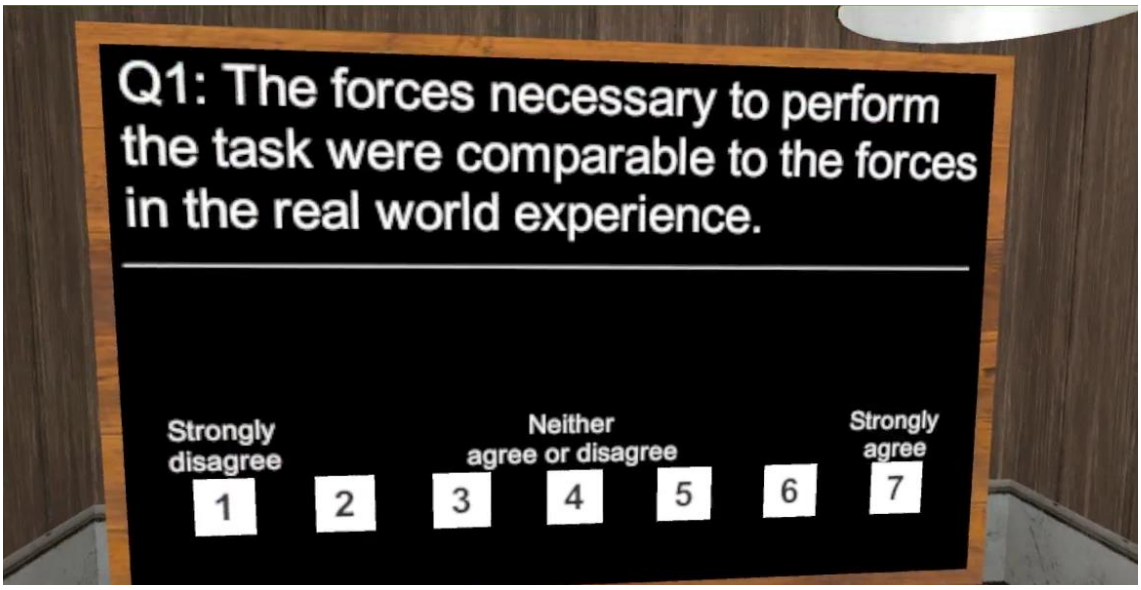

Figure 24. Virtual board showing the questions and answer scale.

After each block, there will be another set of questions referring to the participant's preference of the conditions, those will be answered outside VR. At the end of all three blocks, there will be further questioning in case clarification from the participant is needed based on the observations made of their actions.

4.2.3 SETUP

The experiment will take place in a controlled environment: a private isolated room. Because the participants will not be allowed to see the part of the setup with the tool and surface props, the room will be divided in two sides, one for the real tool tryout and the other serving as the virtual area space. A layout of the experiment room can be seen in Figure 25. In the virtual area space, there will be a table that acts as the interactive surface proxy, which will be moved away by the researchers when the air conditions are performed. For safety reasons, the participants will be told when the surface is removed/introduced.

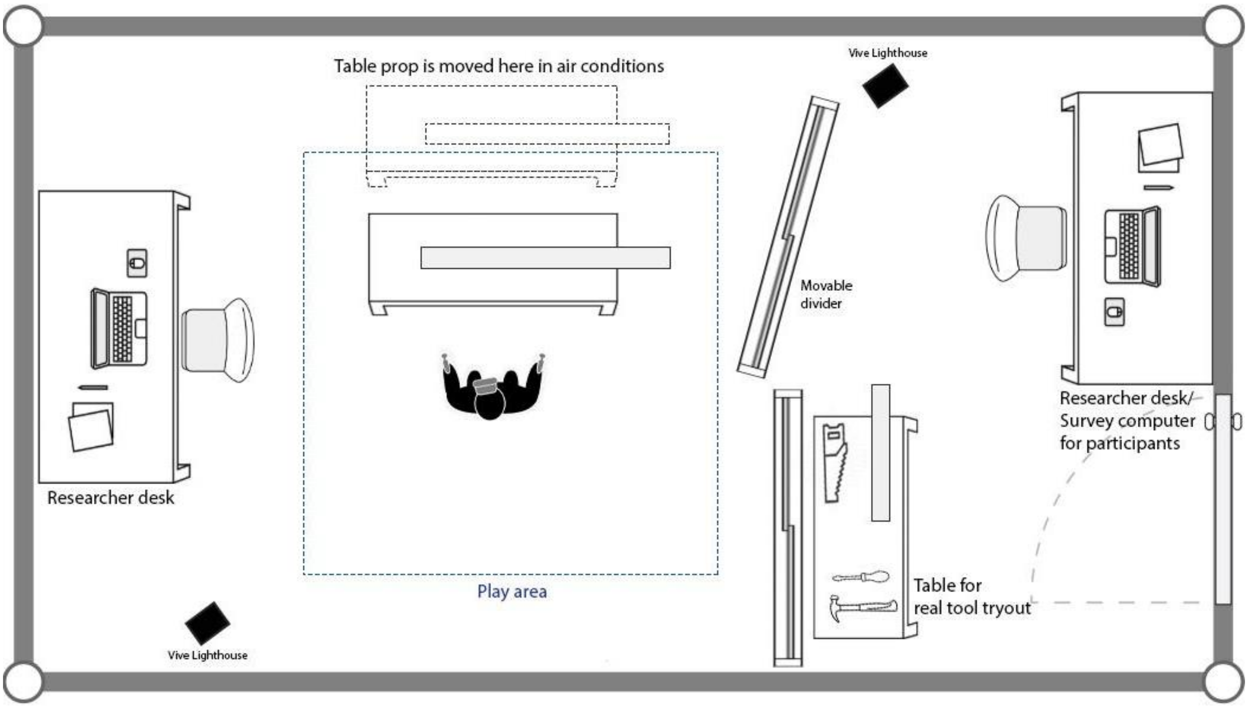

Figure 25. Experiment room layout.

4.3 MEASUREMENTS

This subchapter will present how the participant data will be measured and collected.

4.3.1 PERFORMANCE DATA

A set of data from the participants' interaction will be logged, such as their precision, overshooting amount and completion time. The precision for each tool is regarded in different ways, as follows:

- Hammer: how close to the nail head the impact happens (mean distance during each trail)
- Screwdriver: how much the tip of the screwdriver wanders from the screw point during rotations (mean distance during each trail)
- Handsaw: how far the blade wanders from the marked line during the sawing motion (mean distance during each trail)

As well as overshooting:

- Hammer: how far it travels after impact (mean distance during each trial)
- Screwdriver: how much it wanders on the y-axis (mean distance during each trial)
- Handsaw: indicating how it is used (mean distance during each trial)

Technical details can be found in *3.4 Performance Measurement*. Means are taken for the precision to see if resting the tool on a surface aids in the interaction. Having no surface would mean more imprecision because of the missing support; the interaction is less stable. The overshooting indicated different things for each tool: for the hammer it can tell it was harder to stop after an impact; for the screwdriver if keeping it on point was difficult - like precision, but for the y-axis; for the saw, the measurement shows the height at which the saw was used, as well as how far the saw moved (i.e., indicated from standard deviation).

4.3.2 QUESTION SETS

Additionally, as previously mentioned, sets of questions will be used to collect information on the participant's experience, as well as free observation by the researchers. These observations of the participants will be done in terms of:

- Movements performed during the tasks (how closely to they resemble the movements done during the baseline, how precise they are)
- Parts of the body used to apply force (compared to the baseline, do they put the same amount of effort)
- Notable behaviour (tries to feel the tools, performs the task playfully like in a game etc.)
- Notable reaction (confusion, frustration, excitement)
- Talking aloud (what do they say, do they exclaim things, do they ask questions)
- Difficulties they might encounter (task is hard to perform, not knowing how to control the tools)
- Issues with the system (any issue caused by the hardware or software)

Follow-up questions might be asked at the very end if clarifications are needed based on the observations.

The main part of the questionnaire regarding to the measurements of realism after each tool trail, was developed from the *“Reality Judgement and Presence Questionnaire”* (Baños et al., 2000), and adapted specifically for the situation in this experiment. The reasoning for this choice of questionnaire was our focus on differentiating between the concepts of realism and presence, and to consider them as separate, focusing on realism. Therefore, five questions were adapted and constructed, pertaining to realism and more specifically, the realism of an interaction. Table 3 contains a showcase of all question sets and their measurement scale.

Question	Response options
After each condition	
1. The forces necessary to perform the task were comparable to the forces in the real world experience.	1 Strongly Disagree -- 7 Strongly Agree
2. The experience of hammering/screwing/sawing in the virtual world was realistic.	1 Strongly Disagree -- 7 Strongly Agree
3. My experience of hammering/screwing/sawing in the virtual world was comparable to the experience of hammering/screwing/sawing in real life.	1 Strongly Disagree -- 7 Strongly Agree
4. I felt like the virtual nail/screw/board responded to my actions.	1 Strongly Disagree -- 7 Strongly Agree
5. I felt like I was actually interacting with a nail/screw/board.	1 Strongly Disagree -- 7 Strongly Agree
After each block	
1. Which condition did you find to be the most realistic?	First/Second/Third/ None
2. Which of the three conditions did you prefer the most?	First/Second/Third/ None
3. Why?	(Open answer)
At the end (after all blocks)	
1. To what extent was the interaction in the virtual world consistent with the interaction in the real world?	1 Not at all consistent -- 5 Extremely consistent

Table 3. Question sets used in the experiment.

4.4 PARTICIPANTS

Participants for the experiment were 20 individuals: 15 males and 5 females, all aged between 18-44. They participated voluntarily, no reward or gift has been offered for participation. Their prior experience with VR and tool usage was as follows:

<u>Previous VR experience</u>

- 18 have experienced virtual reality before to various degrees, while 2 have never experienced virtual reality at all

- 16 out of the 18 with previous VR experience, have tried the HTC Vive system before and were accustomed to the controls and headset, while the remaining 2 have tried a similar system, the Oculus Rift

Previous tool experience

- Hammer: All 20 participants have used a hammer before
- Handsaw: 19 participants have used a handsaw before
- Screwdriver: All 20 participants have used a screwdriver before

The overall time period for each test participant were 25-30 minutes. In the following chapter, the quantitative and qualitative results obtained from the test will be presented. Each subchapter will have separate sections for each of the tools.

5 RESULTS

This chapter will present the results of the experiment run to test the difference in realism between three different conditions, for three tools. The test took place at the Multisensory Experience Lab (ME-Lab) located at Aalborg University Copenhagen. Prior, the implementation had been calibrated and a pilot test was conducted to adjust parameters for the algorithm and test the progress of the experiment to refine it.

5.1 QUANTITATIVE RESULTS

In this subchapter the quantitative results will be presented (i.e., the questionnaire and performance measurements). Data obtained from the Likert scales in the test questionnaire were treated as ordinal and a Friedman test was run to determine if there were differences in the distributions of scores between the haptic conditions. Performance data is treated by a one-way repeated measures ANOVA. All values for performance data are rounded to one decimal place. This leaves one tenth of a millimeter. Observation of the updated position in real time indicate that values below this begin to become unreliable. Each condition will be handled as their own test and the following parts will present the results for each of the tools respectively.

5.1.1 HAMMER

In this section, the hammer's results from the test will be presented. The results will be further discussed in the *6 Discussion* chapter. A Friedman's significance test was run to determine if there were differences in scores for each question asked for the hammer, between the three conditions: *controller, haptic-air, and haptic-surface*. The significance level were $p < .05$. All scores were found to be statistically significantly different (see Table 4). Subsequently, a post hoc pairwise comparison was performed with a Bonferroni correction for multiple comparisons, for each question respectively.

Hammer	
Question 1	$\chi^2(2)$= 24.261, p < .0005
Question 2	$\chi^2(2)$= 19.972, p < .0005
Question 3	$\chi^2(2)$= 22.143, p < .0005
Question 4	$\chi^2(2)$= 11.878, p = .003
Question 5	$\chi^2(2)$= 19.972, p = .007

Table 4. Statistical scores for each question.

Q1: The forces necessary to perform the task were comparable to the forces in the real world experience

Comparisons	p-values
Controller - Haptic-air	$p = .001$
Controller - Haptic-surface	$p < .0005$
Haptic-air - Haptic-surface	$p = 1$

Table 5. Pairwise comparison for Question 1.

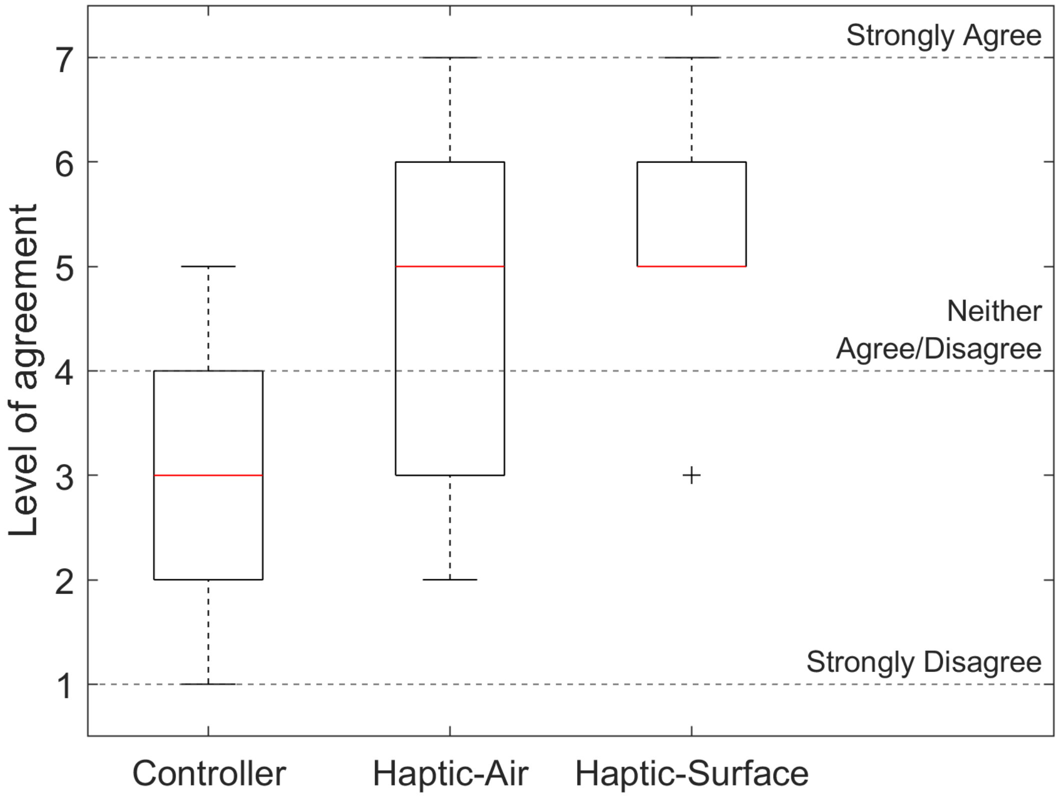

Figure 26. Box plot illustrated for Question 1.

Q2: The experience of hammering in the virtual world was realistic

Comparisons	p-values
Controller - Haptic-air	p = .043
Controller - Haptic-surface	p < .0005
Haptic-air - Haptic-surface	p = .246

Table 6. Pairwise comparison for Question 2.

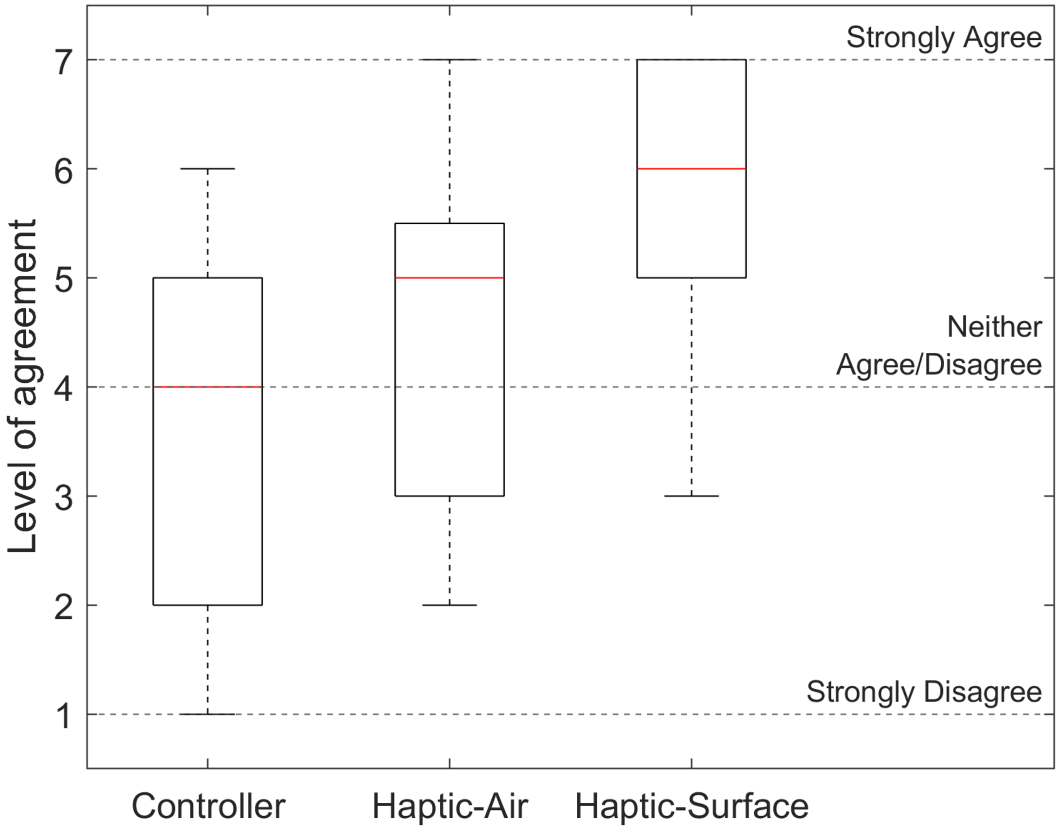

Figure 27. Box plot illustrated for Question 2.

Q3: My experience of hammering in the virtual world was comparable to the experience of hammering in real life

Comparisons	p-values
Controller - Haptic-air	p = .342
Controller - Haptic-surface	p < .0005
Haptic-air - Haptic-surface	p = .017

Table 7. Pairwise comparison for Question 3.

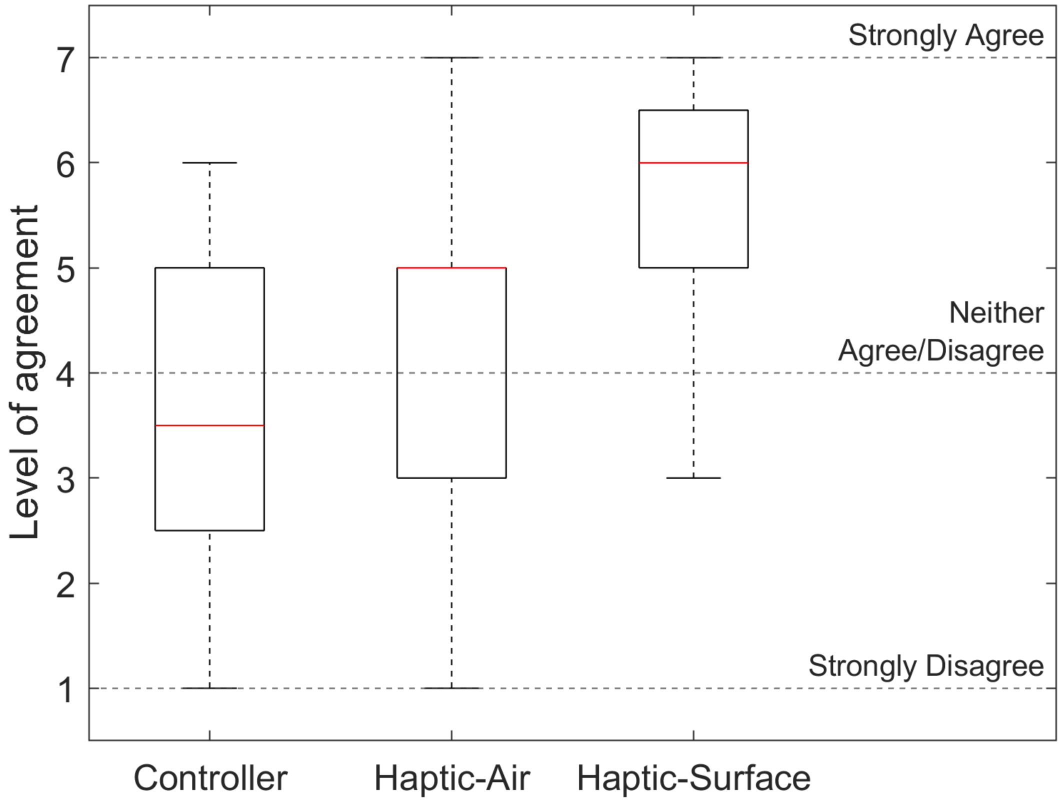

Figure 28. Box plot illustrated for Question 3.

Q4: I felt like the virtual nail responded to my actions

Comparisons	p-values
Controller - Haptic-air	p = 1
Controller - Haptic-surface	p = .027
Haptic-air - Haptic-surface	p = .173

Table 8. Pairwise comparison for Question 4.

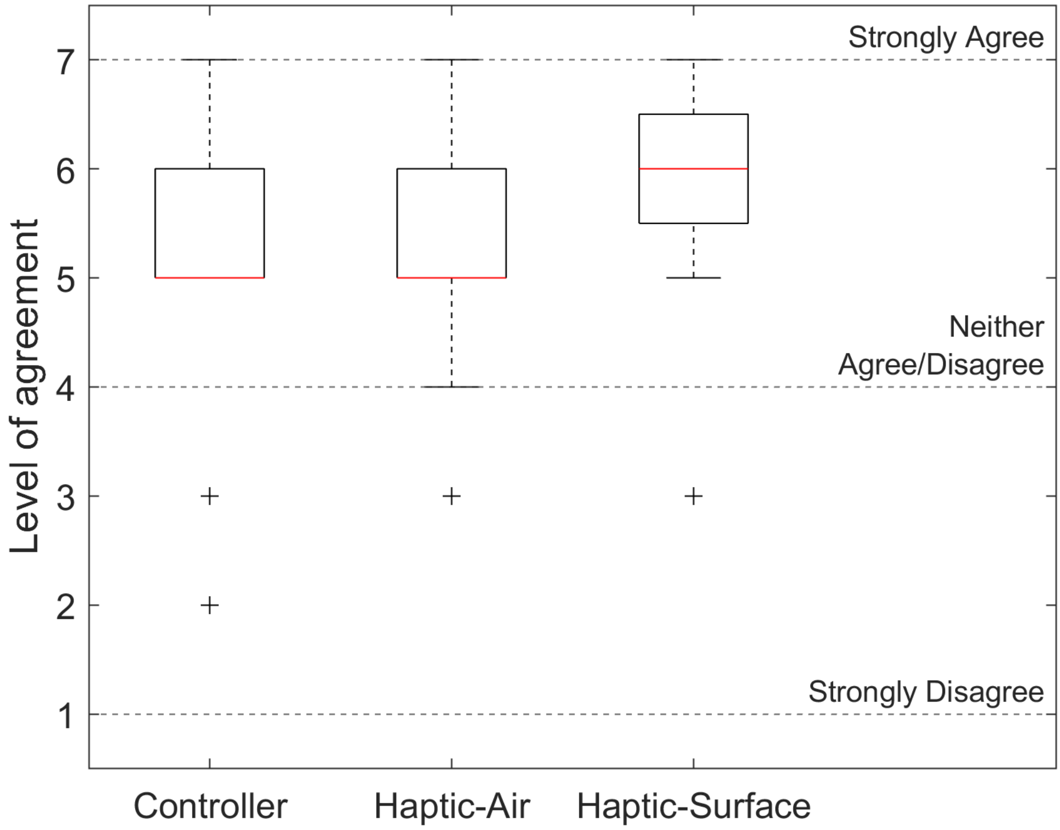

Figure 29. Box plot illustrated for Question 4.

Q5: I felt like I was actually interacting with a nail

Comparisons	p-values
Controller - Haptic-air	p = .618
Controller - Haptic-surface	p = .034
Haptic-air - Haptic-surface	p = .618

Table 9. Pairwise comparison for Question 5.

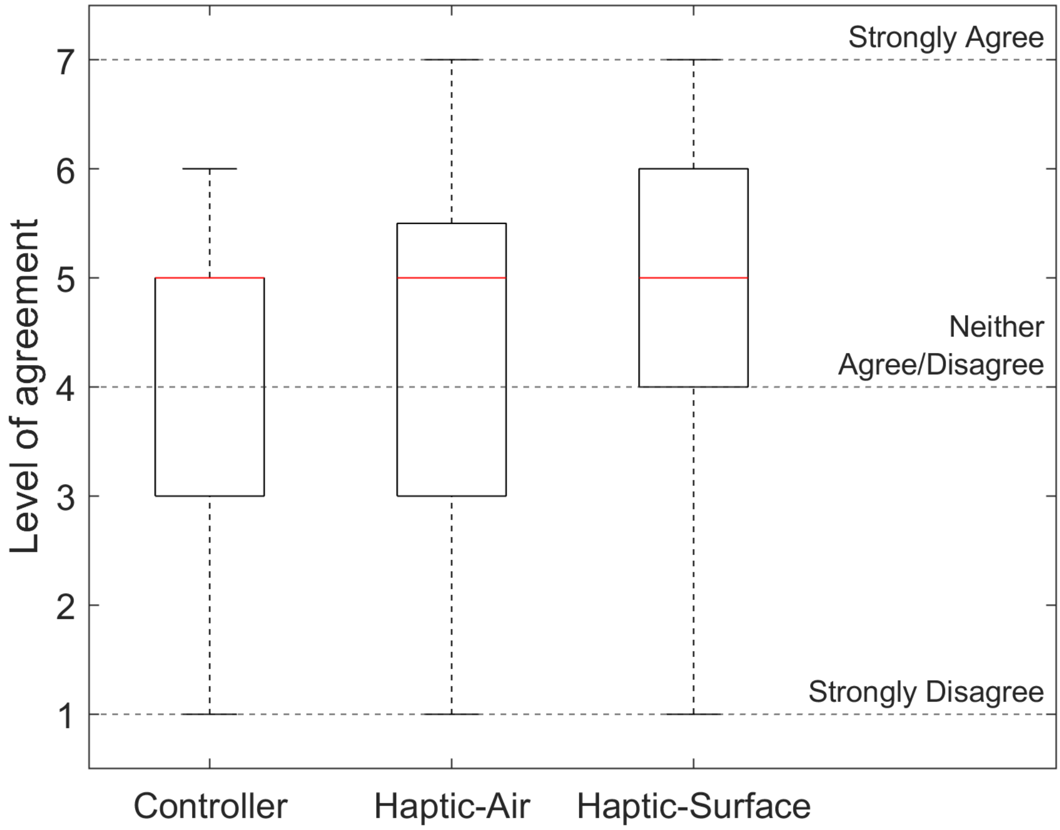

Figure 30. Box plot illustrated for Question 5.

Questions Summarization

The post hoc pairwise comparisons shows a significant difference in the rating for all questions, between the controller and the haptic-surface conditions. However, there is not always a significant difference between neither controller and haptic-air, nor haptic-air and haptic-surface. The box plots illustrated in Figures 26-30 shows the controller version to generally trend the lowest and haptic-surface the highest, with haptic-air ranging in-between them. The median for haptic-surface never falls below the middle score of 4, though the same does also hold true for haptic-air. Outliers were found in Question 1 and Question 4. As seen in Figure 29 outliers were found for all conditions. Observation showed no issues in how the nail responded to the interaction, so the answer here may be due to a different understanding of what the question asked. This speculated is discussed further in the *6 Discussion*.

Performance Measurements

Performance data was captured during interaction with the hammer. It measured the overshoot amount, meaning how much the hammer travelled for 500 ms after virtual impact with the nail. The total means for each condition was calculated for all participants and a one-way repeated measures ANOVA was conducted to determine whether there were statistically significant differences in the amount of overshoot when hammering down with the different conditions. Existence of outliers were assessed by box plot and participant 17 was found as an outlier in the haptic-air condition. Observational data showed the high amount of overshooting is due to excessive movement rather than equipment malfunctioning. It was therefore decided to continue the test regardless. The data was normally distributed in all conditions as assessed by Shapiro-Wilk test ($p > .05$). The assumption of sphericity was violated, as assessed by Mauchly's test of sphericity, $\chi^2(2)$ = 10.873, $p = .004$. Therefore, a Greenhouse-Geisser correction was applied (ε = 0.688). The amount of overshooting was significantly different between the three conditions, $F(1.1376, 26.145) = 7.728$, $p = .005$, partial $\eta^2 = .289$, with the amount of overshooting seen for controller (M = 90.5, SD = 31.6 mm), haptic-air (M = 95.6, SD = 58.2 mm), and haptic-surface (M = 60.1, SD = 30.3 mm). Post hoc analysis with a Bonferroni adjustment revealed that the amount of overshooting was not significantly different between controller and haptic-air (M = 5.2 mm, 95% CI [-33.8, 23.5], p = 1), but there were a significant decrease from controller to haptic-surface (M = 30.4 mm, 95% CI [15.6, 45.2], $p < .0005$), and haptic-air to haptic-surface (M = 35.6 mm, 95% CI [5, 66.1], $p = .02$).

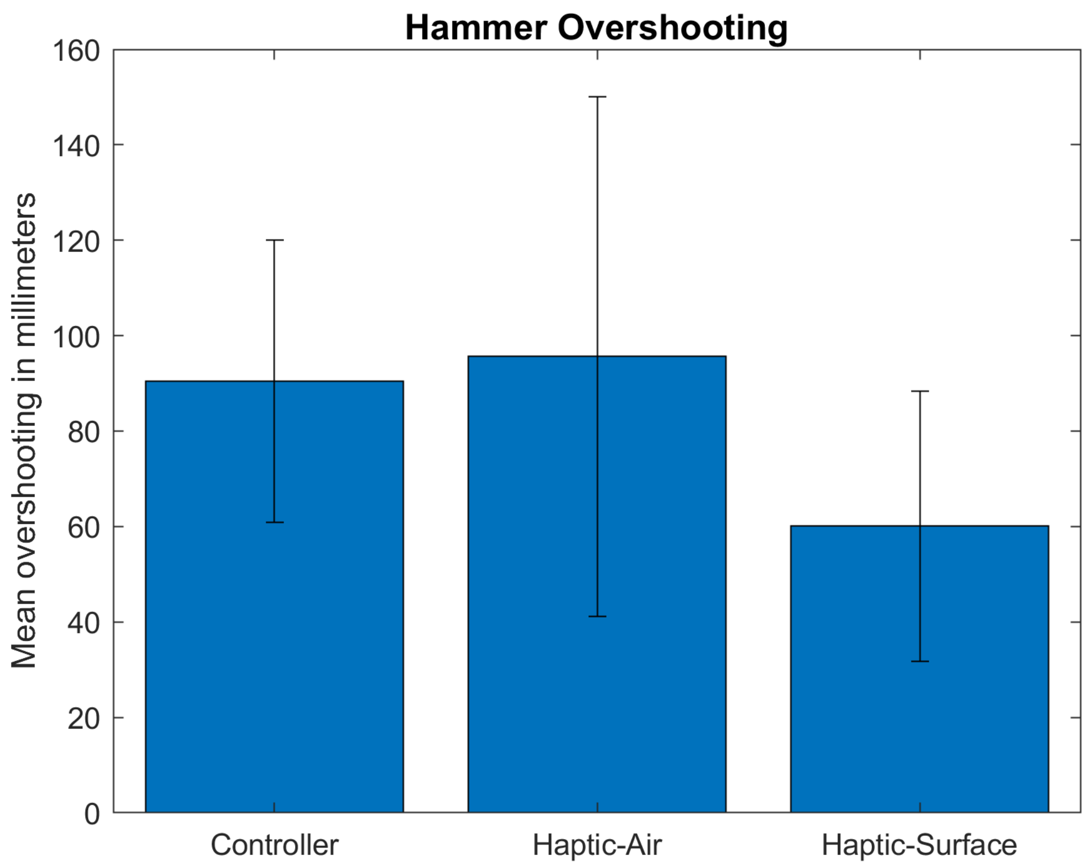

Figure 31. Showing the mean amount of overshoot in millimeters for each condition.

Performance data was also captured in regards to the precision of the hammer, as described in the *3.4 Performance Measurement*. The average distances of all hits throughout a session were calculated, for each participant and condition. A one-way repeated measures ANOVA was conducted to determine whether there were statistically significant differences in the precision between conditions. Existence of outliers were assessed by box plot and participant 4 was found as an extreme outlier in the controller condition. Further analysis of the data showed that many of the participant's hits had a distance above a reasonable amount - reasonable here meaning that the hammer wouldn't be able to collide with the nail at this distance. As such these data points must be due to measurement error. We are unable to say what the actual distance for these collisions are and will therefore treat any outlier as measurement error in the system, meaning they will be excluded in the test. Subsequently participant 4 and 9 were excluded. The data was normally distributed in all conditions as assessed by Shapiro-Wilk test ($p > .05$). The assumption of sphericity was met as assessed by Mauchly's test of sphericity, $\chi^2(2) = .143$, $p = .931$. Distances were significantly different between the three conditions, $F(2, 34) = 9.146$, $p = .001$, partial $\eta^2 = .35$, with the distances seen for controller ($M = 13.4$, $SD = 4.5$ mm), haptic-air ($M = 12.7$, $SD = 3.8$ mm), and haptic-controller ($M = 9.6$, SD 2.8 mm). Post hoc analysis with a

Bonferroni adjustment revealed that there was no significant difference between controller and haptic-air (M = 0.61 mm, 95% CI [-1.82, 3.05], p = 1), but there were a significant decrease between controller and haptic-surface (M = 3.79 mm, 95% CI [1.16, 6.43], p = .004), as well as haptic-air and haptic-surface (M = 3.18 mm, 95% [0.68, 5.69], p = .011).

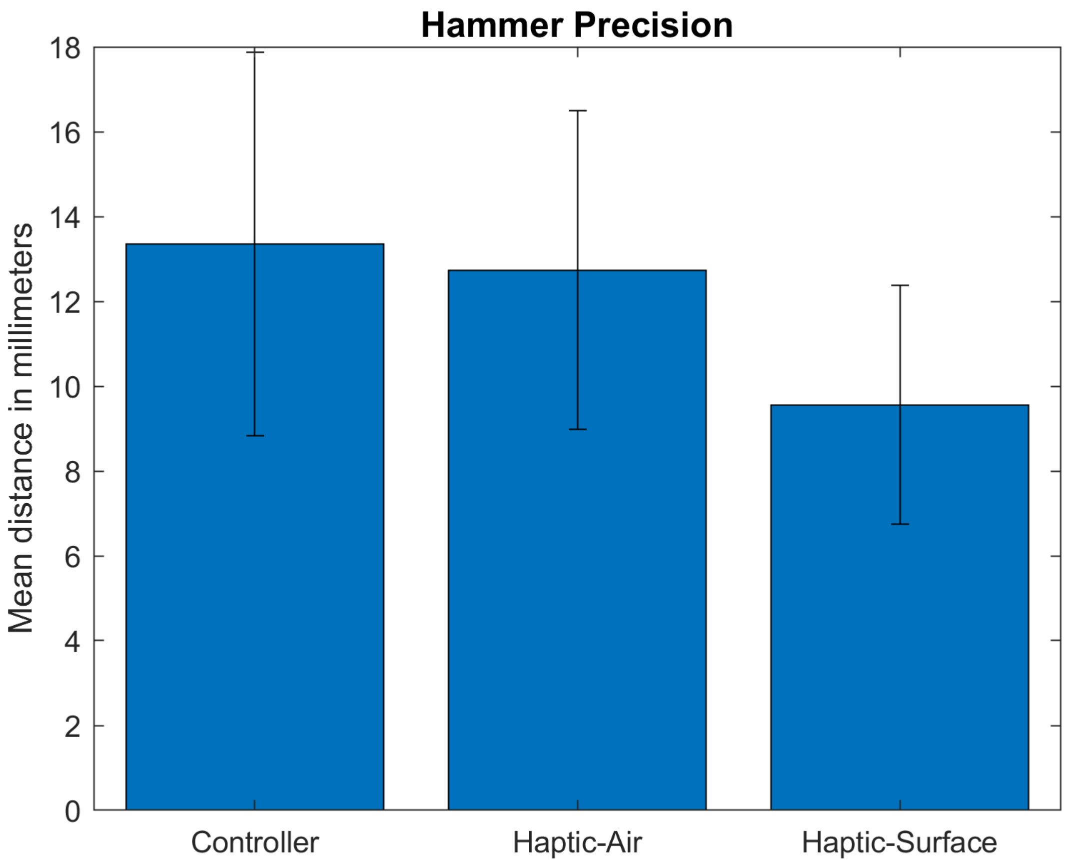

Figure 32. Showing the mean distance in millimeters for each condition.

Preference

In regards to most preferred condition in the case of the hammer, the data can be seen in Table 10. All participants have chosen a condition. While 0 participants have chosen the controller as their preferred, the other two choices are split between 65% choosing the Haptic-surface condition, with the rest 35% for the Haptic-air.

Condition	Realism Frequency	Preference Frequency
Controller	0	0
Haptic-air	5	7
Haptic-surface	14	13
None	0	0

Table 10. Hammer data for condition considered the most realistic and most preferred.

When choosing the most realistic condition, 1 participant chose the controller. However, based on their open answer, it was concluded to be a mistake. As can be seen in Table 3, (*section 4.3.2 Question sets*) , the response options for this question were phrased as "First/Second/Third", referring to the order of the conditions they tried. The participant responded with "First", which meant the Controller. However, their reasoning of the choice was that they could do the task quicker compared to the first condition in which the virtual hammer had some tracking problems and jumped around for a while. Therefore we consider this choice invalid. Their choice of preferred condition was the haptic-air, however we cannot conclude that they also considered this condition the most realistic. Since we cannot deduce from their open answer which choice they wanted to make, this entry will be disregarded, leaving 19 valid participants for this case. The distribution is as follows: 26.31% chose haptic-air, while the rest 73.69% went for the haptic-surface. In terms of matching answers, 16/20 participants (80%) have chosen the same answer as both their preferred condition and the one they found most realistic.

5.1.2 SCREWDRIVER

In this section, results from the test will be presented for the screwdriver conditions. The results will be further discussed in the *6 Discussion* chapter. A Friedman's significance test was run to determine if there were differences in scores on each questions asked for the screwdriver condition, between the three conditions: *controller, haptic-air, and haptic-surface*. The significance level were $p < .05$. All scores were found to be statistically significantly different (see

Table 11). Subsequently, a post hoc pairwise comparison was performed with a Bonferroni correction for multiple comparisons, for each question respectively.

Screwdriver	
Question 1	$\chi^2(2)$ = 18.375, $p < .0005$
Question 2	$\chi^2(2)$ = 18.771, $p < .0005$
Question 3	$\chi^2(2)$ = 17.738, $p < .0005$
Question 4	$\chi^2(2)$ = 11.789, $p = .003$
Question 5	$\chi^2(2)$= 8.542, $p = .014$

Table 11. Statistical scores for each question.

Q1: The forces necessary to perform the task were comparable to the forces in the real world experience

Comparisons	**p-values**
Controller - Haptic-air	p = .464
Controller - Haptic-surface	p < .0005
Haptic-air - Haptic-surface	p = .053

Table 12. Pairwise comparison for Question 1.

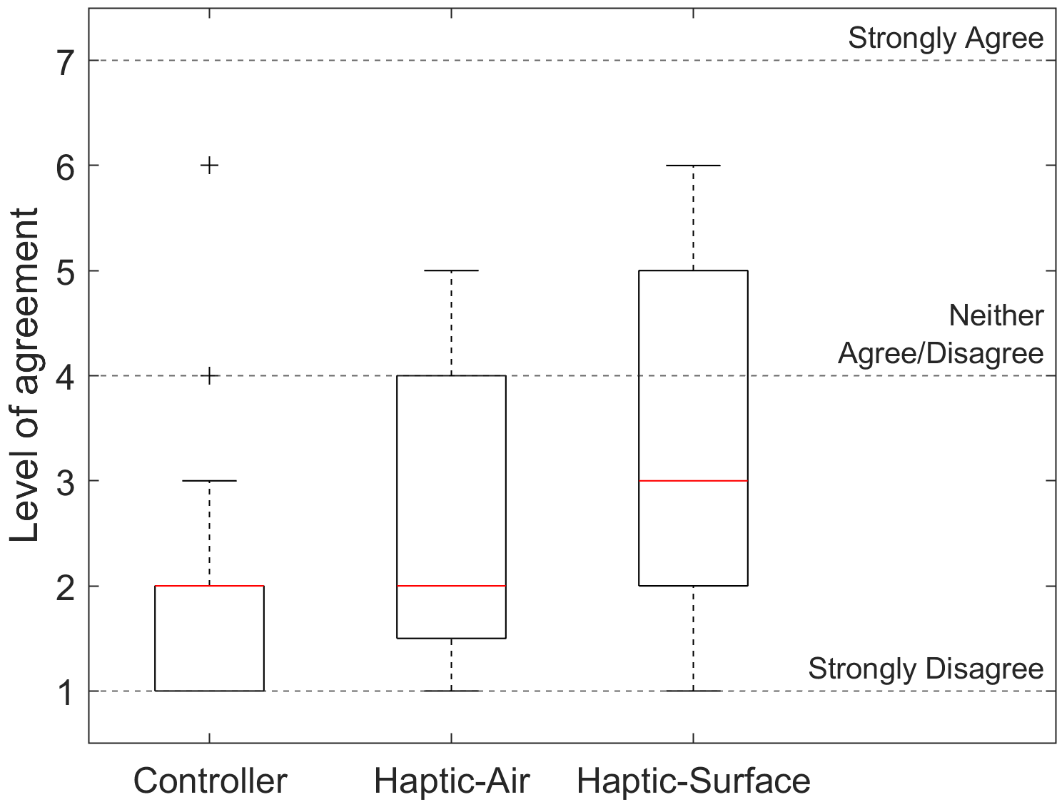

Figure 33. Box plot illustrated for Question 1.

Q2: The experience of screwing in the virtual world was realistic

Comparisons	p-values
Controller - Haptic-air	p = .053
Controller - Haptic-surface	p < .0005
Haptic-air - Haptic-surface	p = .291

Table 13. Pairwise comparison for Question 2.

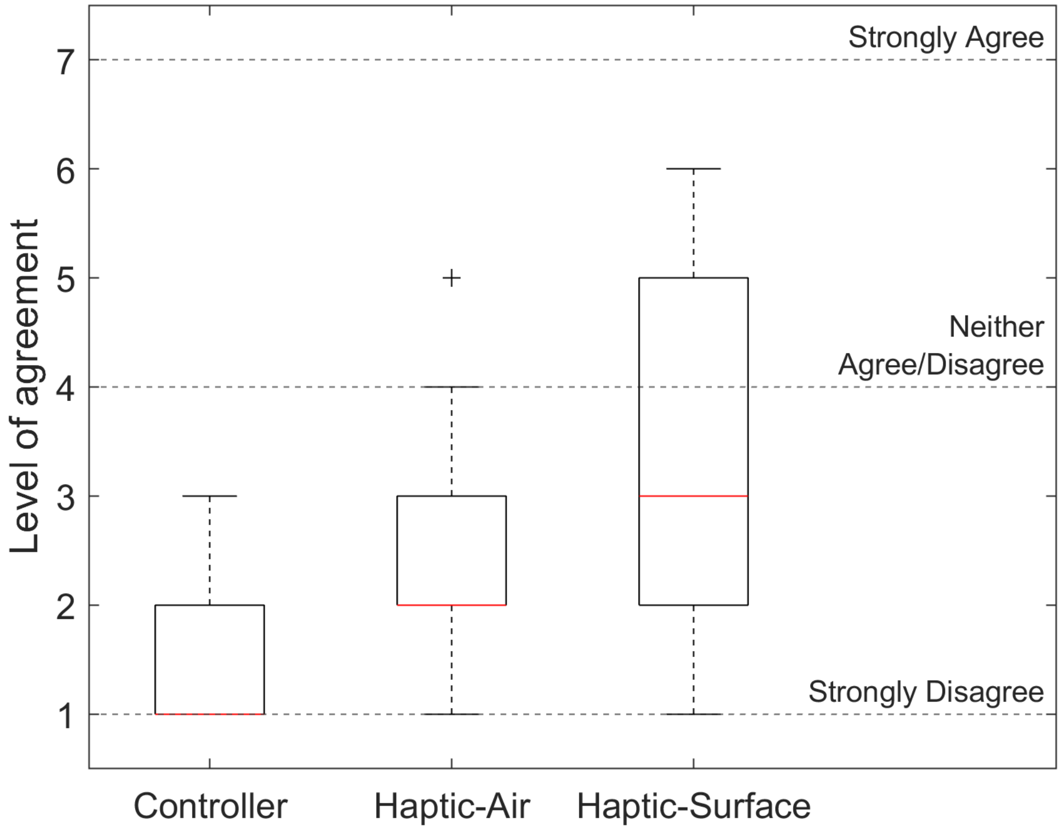

Figure 34. Box plot illustrated for Question 2.

Q3: My experience of screwing in the virtual world was comparable to the experience of screwing in real life

Comparisons	p-values
Controller - Haptic-air	p = .537
Controller - Haptic-surface	p = .001
Haptic-air - Haptic-surface	p = .066

Table 14. Pairwise comparison for Question 3.

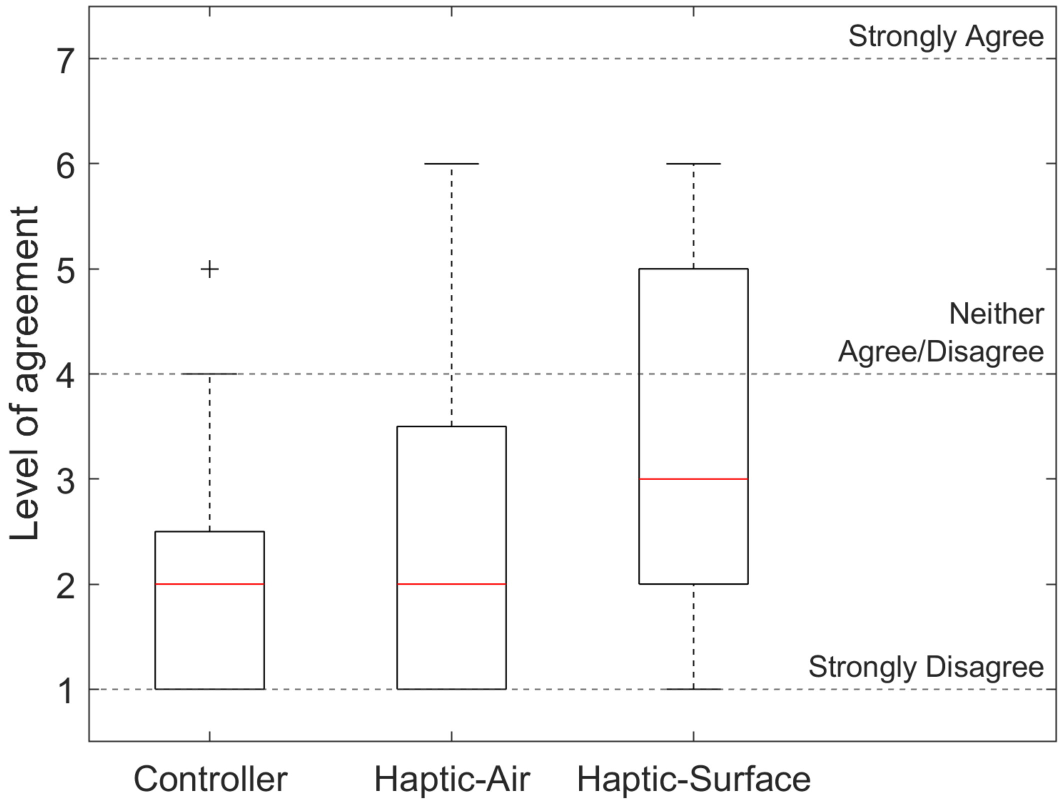

Figure 35. Box plot illustrated for Question 3.

Q4: I felt like the virtual screw responded to my actions

Comparisons	p-values
Controller - Haptic-air	p = 1
Controller - Haptic-surface	p = .013
Haptic-air - Haptic-surface	p = .173

Table 15. Pairwise comparison for Question 4.

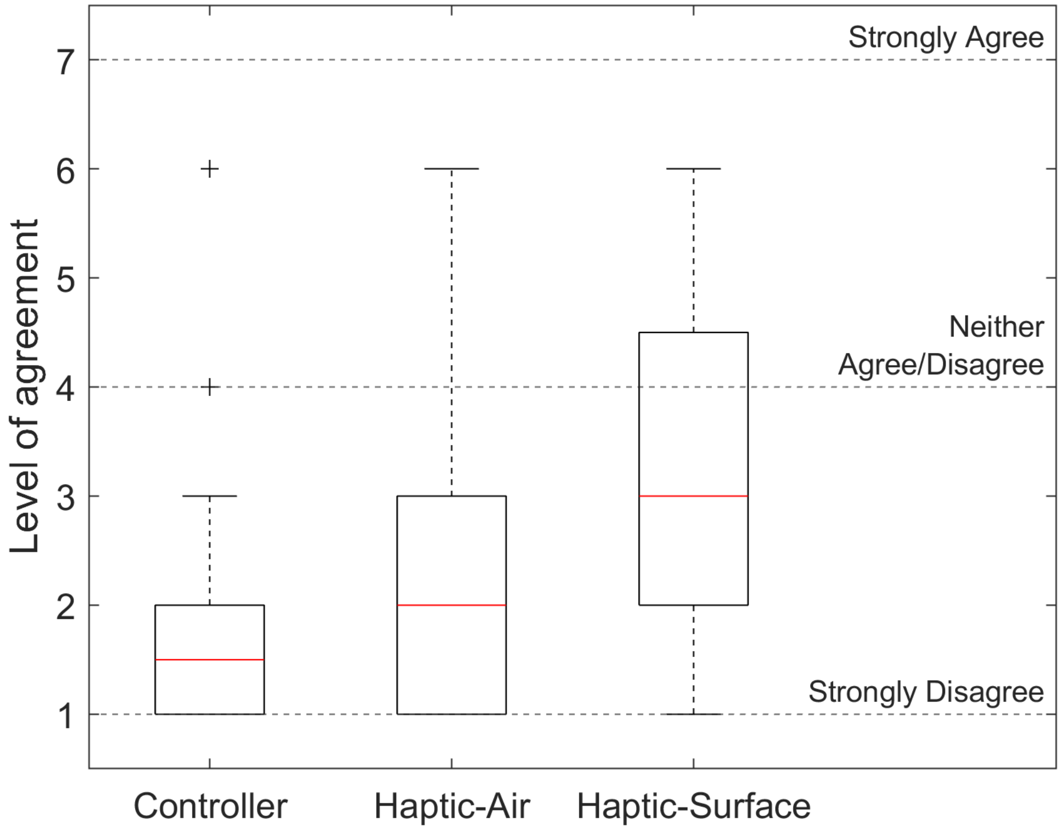

Figure 36. Box plot illustrated for Question 4.

Q5: I felt like I was actually interacting with a screw

Comparisons	**p-value**
Controller - Haptic-air	p = 1
Controller - Haptic-surface	p = .053
Haptic-air - Haptic-surface	p = .173

Table 16. Pairwise comparison for Question 5.

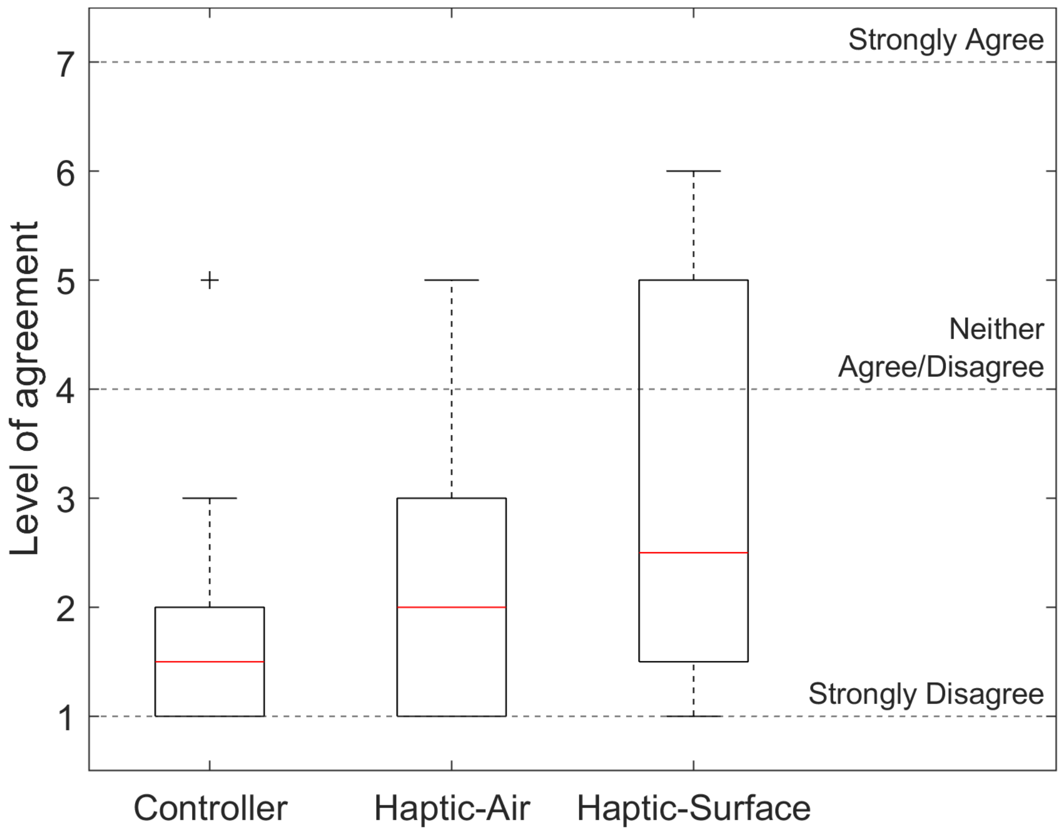

Figure 37. Box plot illustrated for Question 5.

Questions Summarization

All scores are siding towards strongly disagree, the median score given never going above 3. However, both haptic-air and haptic-surface fare better than the controller - though haptic-air does not always significantly. Participants gave very similar scores overall for the haptic-surface condition when comparing all questions (Figures 33-37). All sets had outliers and no score of 7 was given for any condition. There were a significant difference between controller and haptic-surface the first four questions. In Question 5, when participants were asked to state the extent of which they agreed to *"It felt like I was actually interacting with a screw"*, no significant difference was found.

Performance Measurements

Whilst overshooting data was also collected for the screwdriver, several issues with the implementation came to light during the experiment. Follow up questions also had participants note that the screwdriver did not handle as expected or did not respond correctly. Observation of the system during runtime found the displacement of the screwdriver to not work optimally in the haptic-surface condition. The displaced virtual tool would jump below the screw - likely caused by incorrect calibrations of the height of the screw and size of the screwdriver that causes the displacement to be applied in the opposite direction. An assessment made with a box plot found several outliers in all conditions. Those in the haptic-surface condition all had much lower values, less than zero, which confirms the suspicion above. Therefore a statistical test for the overshooting was deemed unwise to run. Issues will be further discussed in *6.3 System Issues*.

Performance data was also captured for precision with screwdriver (i.e., the distance between screwdriver tip and screw). Participant 1 was missing a dataset for one of the conditions and was therefore excluded in the test. Means for all condition sessions where calculated for each participant and a one-way repeated measures ANOVA was conducted to determine whether the distances of each condition were statistically significantly different. There were no outliers and the data was normally distributed for each condition, as assessed by box plot and Shapiro-Wilk test ($p > .05$). The assumption of sphericity was met, as assessed by Mauchly's test of sphericity, $\chi^2(2)$ = .617, $p = .735$. Distances between the conditions were found to be significantly different, $F(2, 36) = 19.671$, $p < .0005$, partial $\eta^2 = .372$, with the distances seen for controller (M = 10, SD 2.5 mm), haptic-air (M = 10.6, SD = 2.5 mm), and haptic-surface (M = 8.7, SD = 1.8 mm). Post hoc analysis with a Bonferroni adjustment revealed that there were no significant differences between controller and haptic-air (M = 0.6 mm, 95% CI [-1.72, 0.4], p = .357). However, a significant decrease was found between the controller and haptic-surface (M = 1.2 mm, 95% CI [0.04, 2.4], p = .042), and between haptic-air and haptic-surface (M = 1.9 mm, 95% CI [0.87, 2.89], p < .0005).

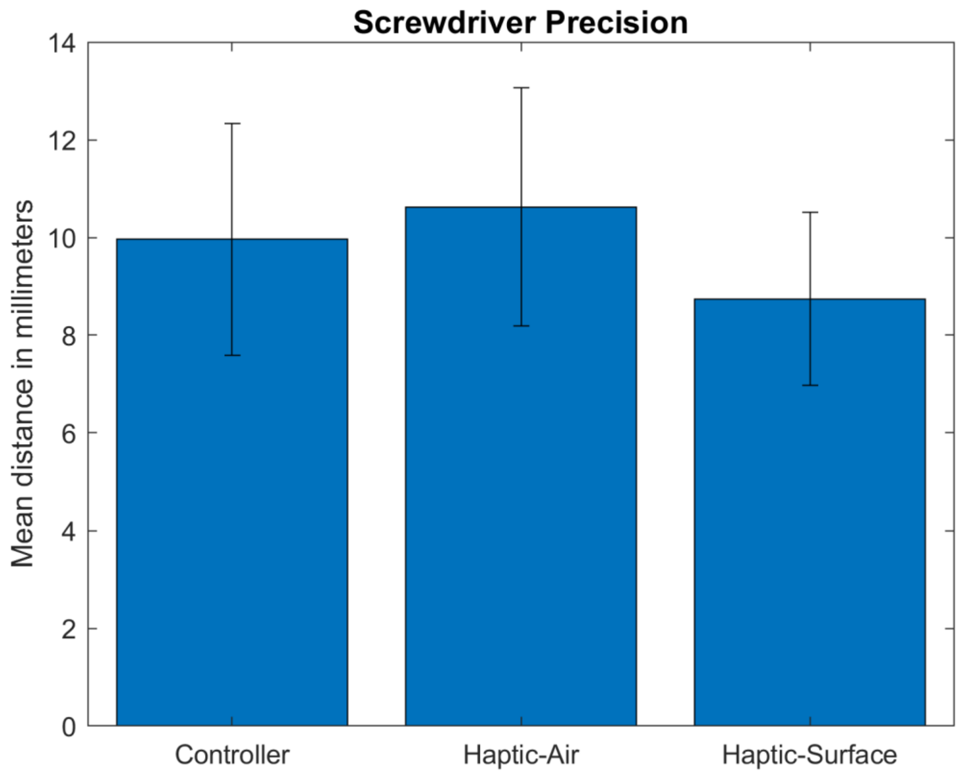

Figure 38. Showing the mean distance in millimeters for each condition.

Preference
In regards to most preferred condition in the case of the screwdriver, the data can be seen in Table 17. 19 participants have selected a choice, while 1 has chosen the *"none"* option. No participants chose the controller. The other two choices are split between 75% choosing the haptic-surface condition, with the rest 20% for the haptic-air.

Condition	Realism Frequency	Preference Frequency
Controller	0	0
Haptic-air	4	4
Haptic-surface	15	15
None	1	1

Table 17. Screwdriver data for condition considered the most realistic and most preferred.

When choosing the most realistic condition, the distribution was exactly the same as for preference choices. In terms of matching answers, 20/20 participants (100%) have chosen the same answer as both their preferred condition and the one they found most realistic.

5.1.3 SAW

In this section, results from the test will be presented for the saw condition. The results will be further discussed in the *6 Discussion* chapter. A Friedman's significance test was run to determine if there were differences in scores on each questions asked for the saw condition, between the three conditions: *controller, haptic-air, and haptic-surface*. Significance level was $p < .05$. All scores were found to be statistically significantly different (see Table 18). Subsequently, a post hoc pairwise comparisons were performed with a Bonferroni correction for multiple comparisons, for each question respectively.

Saw	
Question 1	$\chi^2(2)$ = 31.662, p < .0005
Question 2	$\chi^2(2)$ = 28.794, p < .0005
Question 3	$\chi^2(2)$= 21.377, p < .0005
Question 4	$\chi^2(2)$ = 16.478, p < .0005
Question 5	$\chi^2(2)$ = 24.738, p < .0005

Table 18. Statistical scores for each question.

Q1: The forces necessary to perform the task were comparable to the forces in the real world experience

Comparisons	p-values
Controller - Haptic-air	p = .291
Controller - Haptic-surface	p < .0005
Haptic-air - Haptic-surface	p = .003

Table 19. Pairwise comparison for Question 1.

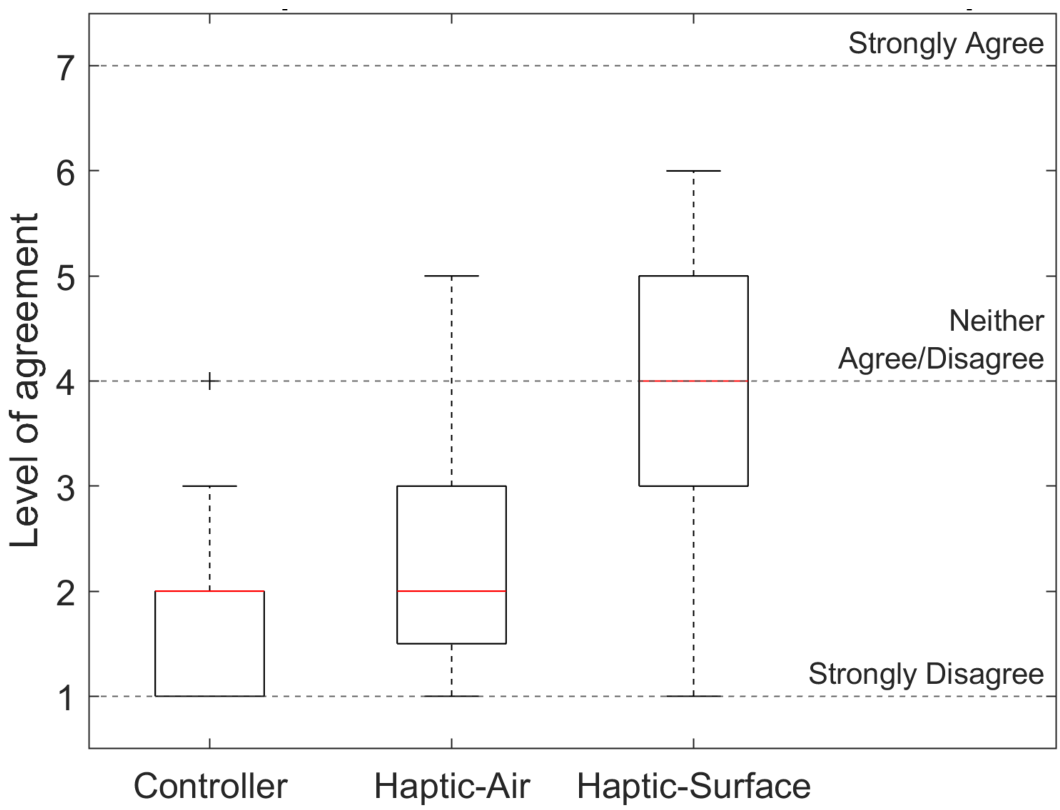

Figure 39. Box plot illustrated for Question 1.

Q2: The experience of sawing in the virtual world was realistic

Comparisons	p-values
Controller - Haptic-air	p = .399
Controller - Haptic-surface	p < .0005
Haptic-air - Haptic-surface	p = .005

Table 20. Pairwise comparison for Question 2.

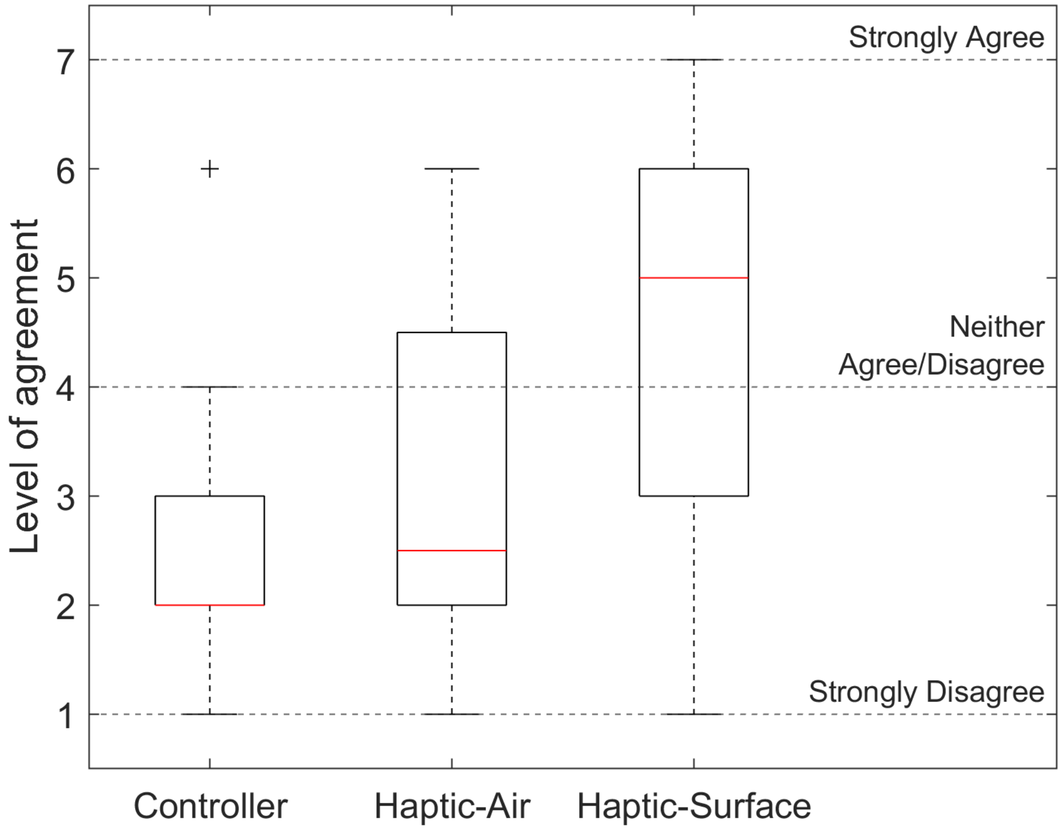

Figure 40. Box plot illustrated for Question 2.

Q3: My experience of sawing in the virtual world was comparable to the experience of sawing in real life

Comparisons	p-values
Controller - Haptic-air	p = .618
Controller - Haptic-surface	p < .0005
Haptic-air - Haptic-surface	p = .022

Table 21. Pairwise comparison for Question 3.

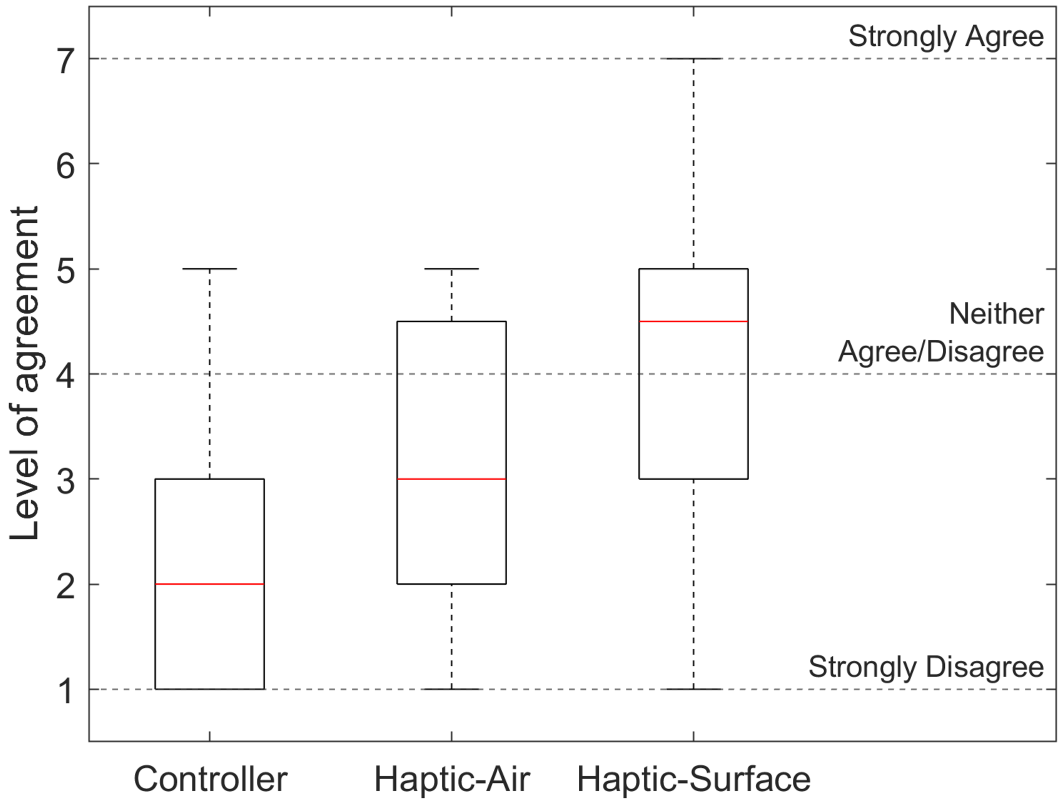

Figure 41. Box plot illustrated for Question 3.

Q4: I felt like the virtual board responded to my actions

Comparisons	p-values
Controller - Haptic-air	p = 1
Controller - Haptic-surface	p = .01
Haptic-air - Haptic-surface	p = .066

Table 22. Pairwise comparison for Question 4.

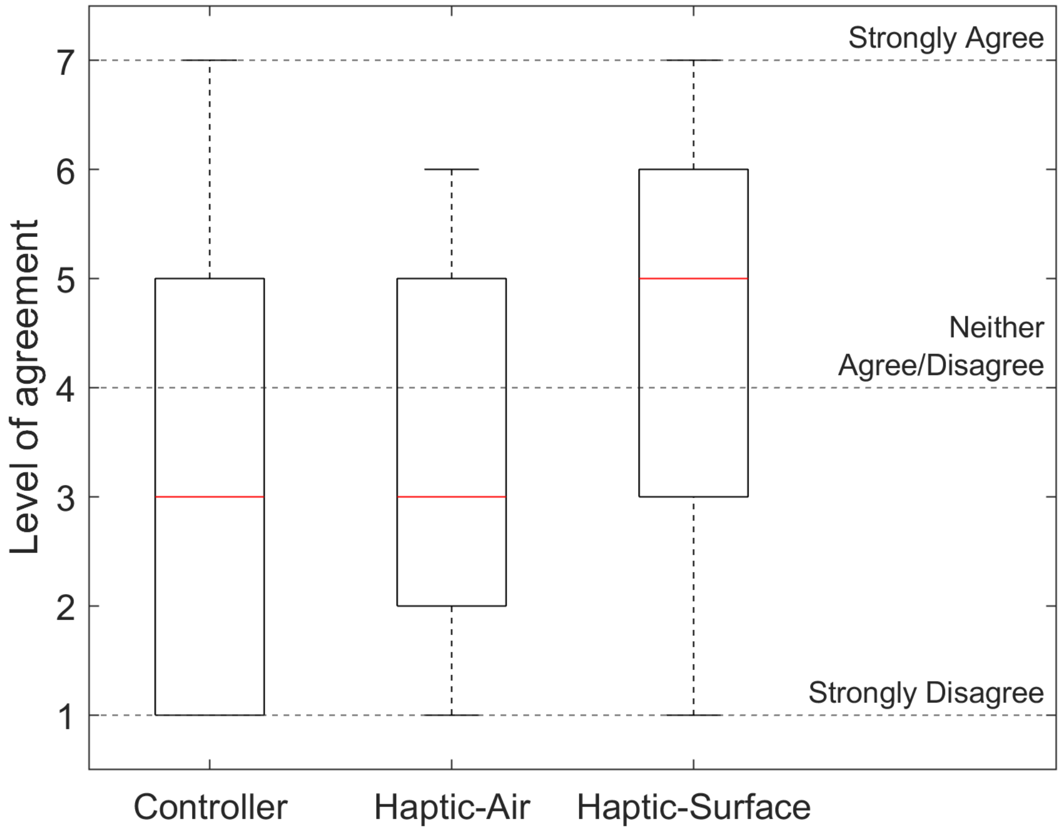

Figure 42. Box plot illustrated for Question 4.

Q5: I felt like I was actually interacting with a board

Comparisons	**p-values**
Controller - Haptic-air	p = 1
Controller - Haptic-surface	p < .0005
Haptic-air - Haptic-surface	p = .003

Table 23. Pairwise comparison for Question 5.

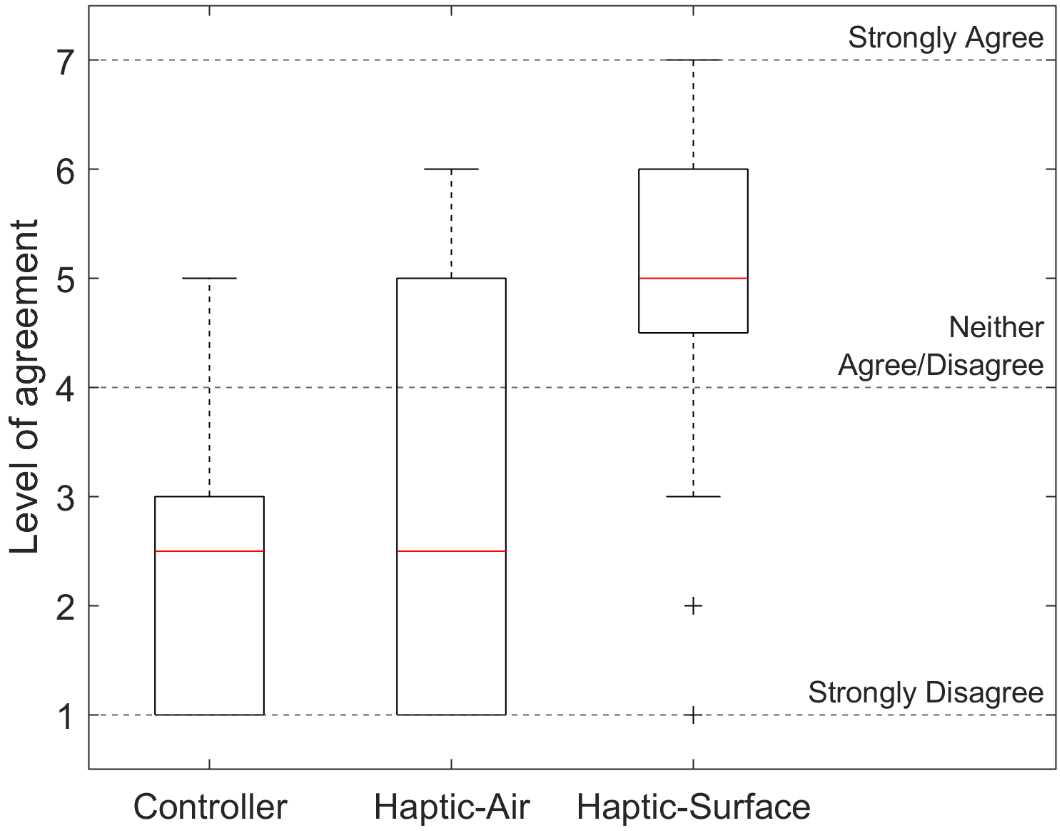

Figure 43. Box plot illustrated for Question 5.

Questions Summarization

The pairwise comparisons of all the questions (Tables 19-23) only shows significant differences in the comparisons with haptic-surface, except for Question 4 in the haptic-air/haptic-surface comparison. None is shown in the controller and haptic-air comparisons for any question. The box plot visualized in Figures 39-43 show the haptic-surface version to be ranged higher than controller and haptic-air. The median answer is either 4 or above. Outliers were found in Question 1, Question 2, and Question 5. The haptic-surface condition saw answers ranging from 1 to 7 in all but the last question. Haptic-air mostly ranged from 1 to 6 and controller ranged for 1 to 5, with an exception in Question 4.

Performance Measurements

As noted in *3.4 Performance Measurement*, overshooting and precision were gathered for the saw condition. For overshooting, an average for every condition session was calculated and a one-way repeated measures ANOVA was conducted to determine whether there were statistically significant differences in the amount of overshoot between all conditions. There were no outliers and the data was normally distributed for each condition, as assessed by box plot and Shapiro-Wilk test ($p > .05$) respectively. The assumption of sphericity was met as assessed by Mauchly's test of sphericity, $\chi^2(2) = 1.768$, $p = .413$. The overshoot amount were found to be significantly different between the different conditions, $F(2, 38) = 6.762$, $p = .003$, partial $\eta^2 = .262$, with the amounts seen for controller (M = 189.1, SD 36.7 mm), haptic-air (M = 160.2, SD 70 mm), and haptic-surface (M = 124.3, SD 59.4 mm). Post hoc analysis with a Bonferroni adjustment revealed that in regards to the overshoot amount haptic-air was not found to be significantly lower than controller (M = 28.9 mm, 95% CI [-18.1, 75.9], p = .368). Neither was the different amount between haptic-surface and haptic-air (M = 35.9 mm, 95% CI [-15.9, 87.7], p = .255). However, overshoot amount was found significantly lower for the haptic-surface compared to controller (M = 64.8 mm, 95% CI [25.4, 104.2], p = .001).

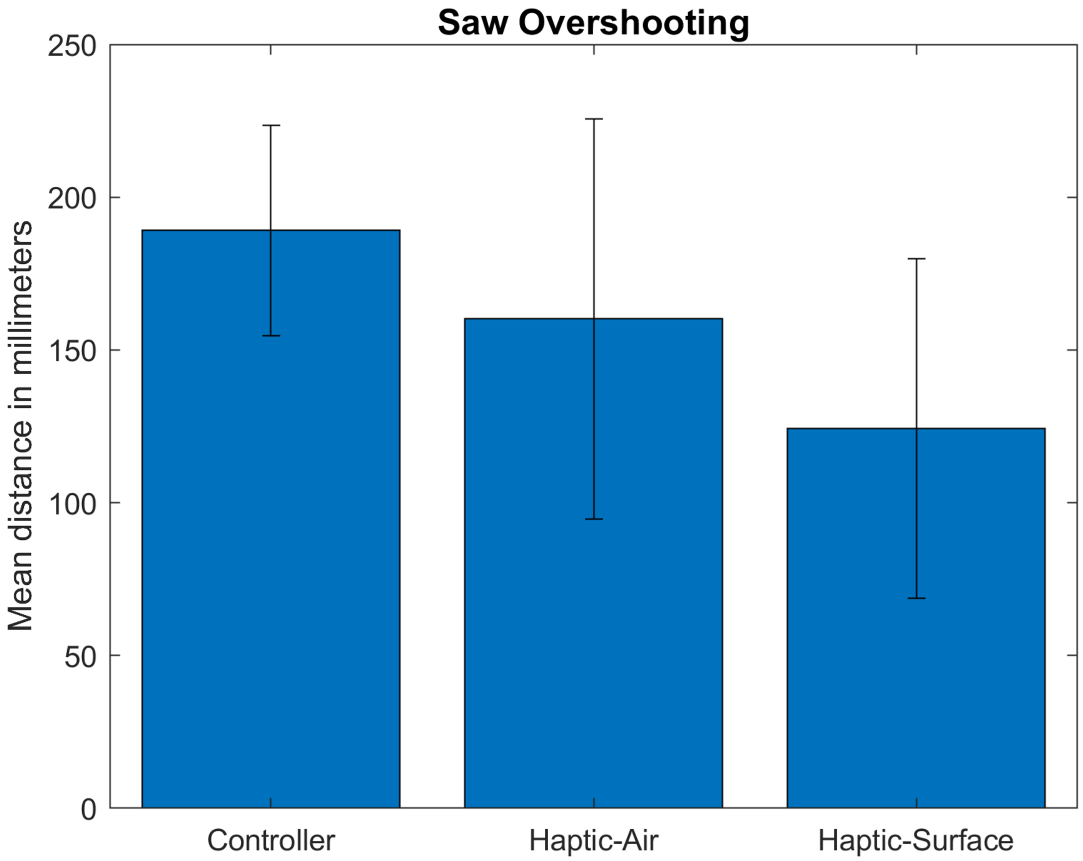

Figure 44. Showing the mean amount of overshoot in millimeters for each condition.

For precision of the saw, the average was calculated for each participant of all their sessions and a one-way repeated measures ANOVA was conducted to determine whether there were statistically significant differences in the distances between all conditions. Outliers were assessed by box plot and participant 4 was found as one in the haptic-surface condition. Analysis of the data found it to be unusual, but no errors seemed to exist in the data entries. Therefore it was kept as part of the test. The data was normally distributed in all conditions as assessed by Shapiro-Wilk test ($p > .05$). The assumption of sphericity was met as assessed by Mauchly's test of sphericity, $\chi^2(2) = 1.120$, $p = .571$. Distances were found to be significantly different between the different conditions, $F(2, 38) = 11.883$, $p < .0005$, partial $\eta^2 = .385$, with the distances seen for controller (M = -2.3, SD 8.4 mm), haptic-air (M = -4.2, SD = 6.8 mm), and haptic-surface (M = 7, SD = 9.6 mm). Post hoc analysis with a Bonferroni adjustment revealed that there were no significant difference between controller and haptic-air (M = 1.9 mm, 95% CI [-3.71, 7.53], p = 1). However, a significant increase was found between controller and haptic-surface (M = 9.3 mm, 95% CI [2.28, 16.23], p = .007), and between haptic-air and haptic-surface (M = 11.2 mm, 95% CI [4.55, 17.79], p = .001).

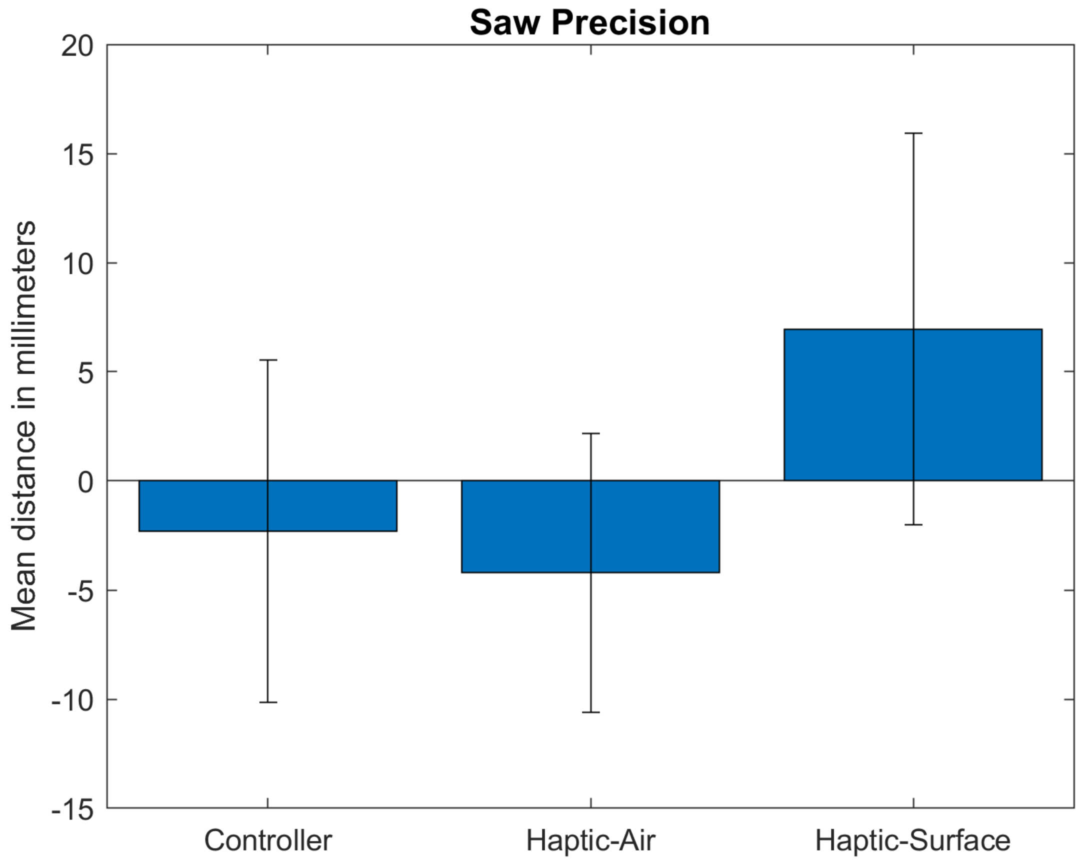

Figure 45. Showing the mean distance in millimeters for each condition, positive values indication distance to the left of the marked line and negative to the right.

It is perhaps not surprising that the haptic-surface condition tends towards the left side, as the redirection is being added to this direction during interaction. Observation and follow-up question showed that no one actively noticed that their hand were being moved. So even though the action was less precise, no one noted themselves sawing any differently.

Preference

In regards to most preferred condition in the case of the saw, the data can be seen in Table 24. All participants have made a choice. No participants chose the controller. The other two choices are split between 90% for the Haptic-surface condition, with the other 10% choosing Haptic-air.

Condition	Realism Frequency	Preference Frequency
Controller	0	0
Haptic-air	0	2
Haptic-surface	20	18
None	0	0

Table 24. Saw data for condition considered the most realistic and most preferred.

When choosing the most realistic condition, all participants chose the haptic-surface condition. In terms of matching answers, 18/20 participants (90%) have chosen the same answer as both their preferred condition and the one they found most realistic.

5.1.4 OVERALL RESULTS

In terms of the extent to which the interaction in the virtual world was consistent to the interaction in the real world, participants were asked to rank the overall interaction on a 5-point scale. The results show that they found the interaction to be mainly *"somewhat consistent"* (M = 3.00, SD = 0.725, Mdn = 3.00), with the percentages shown in Table 25 below.

Consistency level	Percentage
1 - Not at all consistent	5%
2 - Not so consistent	10%
3 - Somewhat consistent	65%
2 - Very consistent	20%
5 - Extremely consistent	0%

Table 25. Results showing frequency spread for level of consistency between virtual world and real world interactions (overall).

5.2 QUALITATIVE RESULTS

In this section, the qualitative data will be presented. This data was gathered from self-reports of the questions pertaining to participants' preference of condition as well as the condition regarded as most realistic. Additionally, notable trends from the observation sheet will also be mentioned.

5.2.1 HAMMER

When describing their reason for the choice of preferred condition, approximately half of the participants have made a positive note on the haptic feedback provided by the presence of the surface proxy. Some of them have noted that having a place of impact for the hammer has increased the realism of the interaction, while others have mentioned how simply having some form of feedback when hitting the top of the nail has kept the impression of reality. Six participants made a comment on the weight of the tool proxy, generally stating that it was an important factor affecting the realism, while some mentioned that it has helped provide the needed force to perform the interaction in a realistic manner. Six different participants mentioned how their choice of preference was influenced by what they also regarded as most realistic.

From the observation data it can be seen that approximately a quarter of the participants use their free arm for support in the haptic-surface condition, placing their hand on the surface and leaning. Both the controller and haptic-air conditions have participants perform similar movements, using mostly their forearm and wrist for the interaction. A majority of the participants have been observed to restrain themselves from using the same force and speed when hitting the hammer on the surface, as when compared to the real tool tryouts.

5.2.2 SCREWDRIVER

Half of the participants have mentioned the positive impact of the haptic feedback provided by the surface-proxy. Most of them have stated how the presence of a surface has helped with precision in the task of using the screwdriver, because they could rest the tip of it and support their rotations more accurately when compared to the other two conditions. Three participants have made a comment on the weight of the tool having an importance in controlling it. Five participants confirmed that their preferred choice of tool coincided with what they found to be most realistic. Issues reported by the participants are covered in *6.3 System Issues*.

Observation data shows that three quarters of the participants have used both hands in using the screwdriver prop, for both haptic-air and haptic-surface conditions, when not nearly the same percentage used the same for the real tool tryouts. Nearly half the participants had discovered a different unintended way of doing the screwdriver interaction instead of rotating it, and stated it worked out better that way. This alternative usage will be addressed in *6.3 System Issues*. Most participants had issues with keeping the screwdriver stable while rotating in both air conditions. In terms of effort, none of the participants used the same rotational force as in the real world interaction.

5.2.3 SAW

A big percentage of participants mentioned that haptic feedback of the surface proxy when sawing has aided the realism of the interaction. Comments regarding that were made, stating it helped put in the necessary force of the interaction when they were able to support themselves with the free hand and lean over the table. Five participants noted how the weight provided by the tool proxy was also a factor increasing the realism. Nearly half the participants' choice of preference in this case was tied to what they also found to be more realistic or closer to their real life expectancy of how the interaction should feel.

The observation data shows a considerable amount of participants using very different body posture and movement for the sawing in both air conditions, mainly only using their arm to perform back and forth movements in air. When the surface was introduced their posture changed and they used their free hand to lean over the table and hold the board down. The sawing motion used in this case was much closer to their real world interaction, even though they did not encounter the same strength of resistance from the blade of the saw.

6 DISCUSSION

This chapter aims to examine and discuss the results obtained from the experiment, both the quantitative and qualitative responses, in order to seek what worked and what needs improvement.

6.1 PERCEIVED REALISM

This part of the discussion will go through the results for each of the tools, comparing the conditions in regards to realism.

The first look at the results, will be for the hammer. Participants found little difference in Question 1 between haptic-air and haptic-surface, but both were significantly higher than the controller. The similarity could be based solely on the introduction of the tool proxy and the participants relating the word *forces* to weight. In the follow-up interview, some participants mentioned that the weight of the tool proxy helped provide the needed force to perform the interaction in a realistic manner. The answers for Question 2 also shows both haptic conditions significantly higher than the controller. Haptic-surface was rated higher than haptic-air, though not significantly, with participants stating that having a place of impact for the hammer increased realism. Some specifically mentioned having the feedback when hitting on top of the nail is what kept the impression of realism. This points to the fact that it is not solely the addition of another proxy, but the perceptual illusion added with it that caused the heightened realism. In Question 3, when asked whether the experience was comparable to real life, participants found haptic-surface significantly more comparable than the other conditions. Question 4 asked if the nail responded to their actions and there was some extent of agreement for all conditions, though haptic-surface was significantly higher than the controller. Having the impact really happen might give in to the illusion that it responds more. It is very possible that the scores for controller and haptic-air would not have been this high, had it not been for the sound. However, it would require a different study to test audio feedback against passive haptic feedback. Several outliers were found for this question, the reasoning likely being a different understanding of the question as no issues were noted during observation. Perhaps they could have understood the question as the nail acting as in real life, which it was not modelled to do. The answers were similarly spread out for Question 5, but none were considered outliers here. Participants still found the haptic-surface to feel most like they were interacting with a nail, significantly higher than the controller, but not the haptic-air condition. More than 70% of the participants specifically stated the haptic-surface as the most realistic, the rest choosing haptic-air. Around 85% of the people choosing the haptic-surface as most realistic also chose this condition as their most preferred. Observation of the motion and posture of participants, saw them using their free arm to support themselves on the surface proxy, to get a more steady aim - which maybe accounts for the higher precision (Figure 32). The motion of the hit in all conditions was similar to that of the real life tryout - covered more in the next subchapter. It is clear from the results that haptic-surface is the most realistic of the conditions for the hammer and that there is a fair amount of realism to it. Haptic-air is a close second in some cases as no significant difference is found, making it somewhat realistic. Controller is significantly lower, showing that this should not be a choice if the decision is to make a realistic experience.

The screwdriver can definitely not be considered realistic in this case, as the median for all answers fell below 4 (neither agree/disagree). This is most likely due to the issues found with the implementation as will be discussed further below. However, it can be said that the haptic-surface did better than the other conditions. Half of the participants mentioned that having the surface proxy for support made a positive impact, mainly because it aided in the precision of the task which does show in the performance measurement of precision, in which haptic-surface was significantly lower in the amount it moved away from the screw. Question 4 also showed that participant felt more like the screw was responding to actions, in the haptic-surface condition. Again, likely due to the added stability aiding in precision. It did not feel like they were interacting with a screw as seen in Question 5, but there is still a ranking with haptic-surface first and controller last, though not significantly. Looking at the forces involved in Question 1, participants did not find them similar to real life in any of the conditions, but haptic-surface was still found to be the most comparable. This is likely due to simply adding another proxy. The reason for this is that the perceptual illusion only worked at times and the interaction with the screw did not work at all, these are unlikely to have been a positive factor. In Question 2 haptic-surface felt more realistic than the other conditions, though not significantly more than haptic-air. Both haptic conditions felt significantly better than the controller. Which again is likely due to the proxies. Comparability to the real world in Question 3 found the same pattern. Haptic-surface is likely the highest due to having the surface proxy for support, as mentioned prior. 75% of participants did choose the haptic-surface condition as the most realistic when asked, though the experience overall was not realistic. The same participants also chose this as their preferred condition. Observational data, however, did show some participants using the screwdriver in a manner similar to the real tool tryout, though more participants did use two hands for the proxy. This is likely due to the size difference. This experiment clearly shows that the screwdriver as implemented is not realistic. The order of realism in descending order is haptic-surface, haptic-air, and controller. It is highly likely that the poor implementation is what has affected the overall results and the differences we are seeing are simply due to proxy differences. Though further studies would have to prove this.

The saw can be considered realistic, at least to some extent for the haptic-surface. For Question 1 the median for haptic-surface was on a 4, but it was significantly higher than both of the other conditions. Participants commented that the surface proxy helped with the necessary force of the interaction, allowing them to lean with one hand on the table. Question 2 once more found haptic-surface significantly higher than both of the other conditions. The median was 5, meaning it was somewhat realistic to saw in the virtual world. A large part of the participants did mention the haptic feedback of the surface proxy was what aided the realism of the interaction. Here the weight of the tool proxy was also mentioned as adding to the realism. Question 3 once more significantly higher, but the median was closer to 4. It was the most comparable experience to sawing in real life, but not truly comparable. For Question 4 haptic-surface is still highest, but less significantly. Even though observation showed the board responding as it should for all conditions - though with some issues in haptic-surface - the scores were not high. This could be due to expectation how it should have responded: move slower, faster, the cut should follow the saw blade, etc. Or perhaps it is due to system issues, which is covered further below. Question 5 show haptic-surface significantly higher. The scores are also higher than in Question 4. That participants answered this way mostly shows that the addition of a proxy for the surface added a lot to the experience, as mentioned prior. Though it should be said that the board did not always drop correctly. Even when it did, the affected participants did not seem to notice that much had happened. Sometimes a small exclamation would be made. When asked about what had happened during the interaction, some participants did find it really interesting when they began thinking about it. Though it is difficult to say, the reason for this could be that it felt so expected that they did not second guess what had happened. This speaks for the strength of this perceptual illusion, but still asks the question why the scores were not higher. In general answers ranged wildly for all questions, in some cases with outliers, though groupings of answers are clear. System issues could be the culprit or perhaps a different understanding in what they are answering. Regardless, all participants found the haptic-surface condition most realistic, 90% preferring this condition. Through observation a considerable amount of participants were seen using postures and movements unlike that seen at the real tool tryout in the air conditions - what we would classify as not realistic. For the haptic-surface condition this posture and movement changed, and became more alike. Their free hand was used to lean and support themself on the surface proxy, and the motion was much closer to the real world interaction. This is especially interesting as neither posture nor movement is required as the same level of resistance is encountered during the cutting motion. In all saw conditions, there were a few remarks made about the rhythm of the audio feedback. Because in the situation of this tool, the contact with the surface and motion was continuous, the sound being played was also continuous and it had a very constant rhythm, leading some people to try to follow it. This has caused them to not saw at their own pace, and pay more attention to the audio feedback than in the case of the other tools. Though one participant did mention that it followed his rhythm. For the saw it is also difficult to say if amount of realism is added due to the passive haptics or the perceptual illusion. More testing with less issues would of course be required. Further discussion on the passive haptics’ impact in the experiment is covered in the next subchapter.

Interesting to note that for the choices of most realistic and most preferred, a majority of the choices coincided, meaning that participants have actively chosen their preferred condition to be the one they also considered the most realistic. Comments on the subject do reveal that to be the case. From the results it is reasonable to assume that people prefer higher realism for their interactions in tool simulations - whether this is a matter of passive haptics, perceptual illusions, or a mix is difficult to say. The next subchapter will further discuss the use of passive haptics.

6.2 PASSIVE HAPTICS

This subchapter will discuss the effect of the passive haptics and how they were interpreted and felt by the participants.

A big number of participants were observed to restrain the speed they hit with the hammer, during the surface condition. They seemed to have kept back from fully impacting the table, unlike they did in the baseline tryouts. When asked about it post-experiment, some have mentioned that they held back because they were afraid they would damage the props. Others have reasoned that they were still aware they were in VR, therefore they did not feel compelled to use the same amount of force as in real life, even though they were told they could. Perhaps it is the knowledge of being in a VE that affects how you do things, as you know you will probably not have to exert the same forces to get the same outcome. If the perceptual illusions were made to mimic the force needed for a real nail, then the interaction may have been much more similar. Of course this does nothing for those who held back because of fear of damaging the system. Contrary to this reaction, a few other participants did use full force just like in the real world, and when questioned they stated it felt good to be able to do that in a virtual environment. However, this behaviour does speak against the realism of the hammer.

A number of participants made out loud remarks when handed the props for the first time. It seemed they were surprised, as they mainly did not expect the heft of the tools being handed to them. Some have tried touching the tools and swung them around, and it was revealed upon questioning that they wanted to figure out what they were holding and how accurate it is portrayed in VR. Overall the passive haptics were very noticeable, as it was shown with the results. The observational data has shown that participants made a lot of remarks when first handed the prop, as well as first trying a task on the surface.

In terms of resistance, it was clear that the implementation did not emulate this quite high enough for some people to consider the experience realistic. Several participants remarked how they would have given overall higher scores if only there was more resistance in the case of the screwdriver and occasionally the saw as well. Even though most of the issues with the screwdriver was because of the rotational issues from the implementation (mentioned in the next subchapter), during surface conditions some remarks have been made that compared to the baseline, the interaction did not offer the same amount of resistance when trying to screw down the screws or cut the board with the saw. These remarks indicate the speculation that differences in screwdriver is due to the passive haptics rather than the perceptual illusions. It also confirms that perceptual illusion can play a part in realism if implemented without issues and a higher complexity.

As noted in the *5.2 Qualitative Results* subchapter, many participants had something to say about the weight of the tools having importance in the realism.

“I felt the same heaviness and force as with real hammer, plus it was not passing through the board but was touching the surface keeping the impression of reality.’ - Participant 2 / Hammer

“The weight was very important, and the fact that I felt something physical was a major bonus to the feeling and experience” - Participant 7 / Screwdriver

“It felt most realistic because the tool was heavy like in real life and i could feel something with my left hand that was "on the board" - Participant 20 / Saw

Although the weight of the props was not a focus of the implementation, it is interesting to note how weight was one of the things that gathered the most attention and seemed to have influenced the realism of the experience. Though it is not a huge surprise as the research discussed in the *2 Analysis* chapter has shown weight to be important for passive haptics. However, there was one notable point made about the weight of the tools. One participant said how although having heavy tools adds to the realism, it is quite difficult to be precise and stop the motions in the air conditions. Giving the example of the hammer, without the presence of surface the participant had to stop the motion of the hammer just from his wrist, which he noted was not very comfortable and could lead to strain. The specific comments about passive haptics show that differences in results are to a large extent due to passive haptics.

6.3 SYSTEM ISSUES

System issues was found working with the implementation and running the experiment. Some of which affected the experiment and therefore the results. The biggest hurdle was likely the calibration of the system. A part of it being that the virtual room of the VIVE setup would drift over time or simply have moved on a restart of the Unity scene. The system would also at times present a tilt of the floor that was noticeable enough that one end of the table would clip with objects placed on it, but the other would float above. This drift would move the virtual scene by up to 20 centimeters and rotate it by a few degrees. It was rather easily solved by knowing about it and making sure to run each tool session, without shutting down the scene for the rest of the duration of the experiment. Between scenes a quick check would be made to confirm that nothing had moved. For the tilt, it was found that holding the HMD looking upwards when starting SteamVR could solve the issue. These issues are unlikely to have affected the test as precautions were taken. Second significant problem is the calibration of the objects in the implementation. Table, board, nails, and screws all required precise placements for them to function properly. Due to the issues of the VIVE system, the table would have to be calibrated each day and sometimes several times a day. This hassle could like have been avoided by tracking the table in real time. For the nails and screws, incorrect calibration would mean that the real tool and virtual tool may not hit the table and objects’ tops at the same time. This makes for an incomplete experience and could very well have affected participants’ perception in an uncontrolled way. The wooden board proxy also needed precise placement, otherwise the feeling of the board piece being cut off would not occur. Calibration is a large part of making such a VR experience happen, like the one presented in this thesis. It is important to keep in mind when designing for it.

Of issues that affected the results the most noticeable was without a doubt the screwdriver. Many participants felt like the screwdriver had erratic responsiveness, and chose the preferred condition based on when the screws responded best to their actions.

"I felt that the screw was responding to my movements more" - Participant 2 / Screwdriver Haptic-Surface

"I could rest the screwdriver on the board which helped. I also learned how the screws responded best" - Participant 3 / Screwdriver Haptic-Surface

"The interaction was hard to make work in the first 2 instances. The interaction felt more realistic in the third." - Participant 18 / Screwdriver Haptic-Surface

Participants also mentioned this as a difficulty, and further observation easily showed this to be true. The screwdriver would pop below the table at times during the haptic-surface condition. This was most likely the height of the table being calibrated incorrectly, such that the tip of the screwdriver would be below the point and thus add the offset in the wrong direction. This mostly happened on the first test day and disappeared after the setup was calibrated for the next test days. This was also mentioned in the results for overshooting performance measurement and why it was excluded. On top of that, participants had difficulties getting the screws to screw down correctly, which is clearly a fault in the implementation. Screws would follow the rotation of the screwdriver, but not go down reliably, even going back up sometimes. Some participants found a method of spiralling the screwdriver around the edge of the screw to work, but of course this is far from intended, and does not mimic a screwdriver usage in real life.
The screwdriver offset issue only happened in the haptic-surface condition, whilst the screw issue happened in all of them. It is extremely likely that these issues severely affected the scores given for the different questions in each condition. A possible addition to this interaction would be to snap the screwdriver to the screw, assisting the interaction and allowing for it to feel more precise - as well as aligning with the groove.

Saw had issues with the calibration as well, causing the board to drop before the saw did. This was a mixture of calibration issues and the implementation not being as good as it could be. The implementation looked for the progress of the saw, from a point that existed around the real saw's handle. As people move their hand horizontal as well when they saw, it caused the hand to move over the edge - dropping the board - whilst the blade was still on the board. The more optimal solution would have been to find the intersection point of the real saw with the board. This point changes over time, but when that leaves the edge, it will be certain that the entire saw is moving over the edge. Another issue was a missing weight to the peak detection. This was found too late to change and meant that even holding the saw still would make the system think it was being used to saw. This is due to the tracker always moving slightly because of imprecision, meaning the saw will think it has moved from its last position.

Hammer appeared to have the least issues, but some were found. Looking at Figure 31 for the overshooting amounts, the haptic-surface condition is much higher than one would expect. The hammer is supposed to stop when it hits the table and thus the amount should not be over zero. Regardless, a distance was measured which is likely due to the tracker continuing after the impact. It was observed that on a hit, the virtual hammer would move downwards for a brief moment before going back up again. This could be due to the accelerometer not faring well with impacts. This distance was likely made bigger due to the same offset issue as the screwdriver.

6.4 TEST VALIDITY AND BIAS

Overall the experiment consisted of 20 participants. While this is not representative of a larger population, it still resulted in interesting findings. However, for the experiment to have higher validity and reliability, more participants would be needed.

There is always the factor of human error and misinterpretation with self-reported measures. While some participants did ask clarifications when needed, it is never a guarantee that all participants have the same interpretation of a question, or they simply might have misunderstood it.
Additionally, participants might have also not had a clear enough recollection of the baseline to compare when asked questions relating to that. It was noted that in a few cases, for the first question *"The forces necessary to perform the task were comparable to the forces in the real world experience"* participants were either not quite sure what the question meant or they did not actively think about the forces at play during the real tool tryouts.
Furthermore, there is also the factor of fatigue to be taken into account. Because the test was divided into three tool blocks, with participants having to go in and out of VR three times, it is possible that towards the end some might have tried to end it quicker, by rushing the task or answering the question with not as much thought. Of course there is no way to know for sure if any of this was the case, but we have asked almost all participants to estimate how much time they think they have spend during the whole experiment since entering the room and answering the last question, and surprisingly all of them have estimated less time than they actually spent. When questioned further about it, they have noted that because of the constant switching between real life and VR, it felt like time passed quicker. The randomized order of the blocks should also help negate the possible bias because of fatigue in the end.

From some of the answers and follow-up questions, it was also revealed that a few participants regarded this more as a game. This was also confirmed through observation in several cases, when participants were rather playful with the tools. This could have possibly lead to them not performing the tasks as we expected (i.e., as they would perform them in the real world). This issue could be solved in a future test by having clearer and more concrete instructions for the participants, where it is emphasized that the experience is not a game.

6.5 FUTURE WORK

This subchapter will point to future areas of research and improvements in regards to the subject of this thesis.

Maybe not surprisingly our experiment has shown that realism is higher in the haptic conditions, compared to the controller condition. To gain more knowledge about realistic tool interaction and the amount it is affected by passive haptics and perceptual illusions, it would be valid to test different algorithms and parameters of these. As presented in the *2 Analysis* chapter, there are several techniques and variations that can be applied to different degrees of success depending on the intended purpose. In the case of the tools selected for this thesis, some variations of the redirected touching and haptic retargeting methods can be implemented to test which would be best suited for the interaction at hand. More complexity could also be added to the interaction in form of resistance (e.g., making it feel like the screw tightens or the saw is digging into the wood).

In a future implementation, the system can be enhanced with hand tracking technology, so users can get a clearer idea of their position and their limbs when handling the tools. Few notes have been made by participants that they actively noticed the lack of body representation, therefore a future study can include real time tracking of user's hands. Even though the realism of the interaction itself was the focus in the current thesis, it would be interesting to see if the scores would change in the presence of virtual limbs.

Audio feedback is very likely to have an impact, but it would also be relevant to see how important this is. Rather than a sample on a loop, as used for the screwdriver and saw conditions, one could explore better simulated sounds and see how they impact. With or without perceptual illusions.

Observations taken during our experiment gives some indication of how similar/different participants' movements and postures were between virtual and real interaction with the tools. However, it does not give a clear picture. Future studies could utilise motion capture to see exactly how similar these movements are.

Future studies should thoroughly consider how the calibration is done and the VR system in use. Proxies that replicate the world needs to be precise in their placement, otherwise the experience can easily fall apart - as seen with the issues found with the screwdriver.

7 CONCLUSION

The Final Problem Statement was as follows:

> *“When comparing the three conditions of controller, haptic-air, and haptic-surface in a virtual reality tool simulation, to what degree can passive haptics and perceptual illusions enhance the interaction in terms of realism?”*

When it comes to passive haptics impacting the realism of the interaction, the results have shown that it has done so. The controller condition for all tools has significantly lower realism scores, as expected. Adding at least one prop to the experience did enhance the realism of it, however, not always to a statistically significant extent. The prop added to serve as a surface for contact seemed to have the highest impact, as it allowed users to have more control over their actions. This has been confirmed by both the self-reported measures, observational data, and performance data as multiple people have used the surface to their aid, and their general movements were closer to how they performed in the real world tasks.

When including the perceptual illusions in the experience, it is hard to tell what the extent of their influence are over the results. However, the majority of participants did not actively notice the displacement occurring during their actions, pointing to their perception being successfully altered as to believe to some extent that what they were seeing was actually occurring.

Future work with an improved implementation and a more consistently working algorithm, as well as more focused and clear questions on perceived realism, would be able to give an answer about the extent to which each illusion and prop has affected their realism in the experience.

Overall, the subject of tool interactions in virtual reality is an interesting specific area of the field that we believe is worth looking into even further, as it could provide insight into more realistic and believable virtual interactions in the future, from both a perceptual aspect and a behavioural one.

8 BIBLIOGRAPHY

Aguerreche, L., Duval, T., & Lécuyer, A. (2010, November). Reconfigurable tangible devices for 3D virtual object manipulation by single or multiple users. In Proceedings of the 17th ACM Symposium on Virtual Reality Software and Technology (pp. 227-230). ACM.

Azmandian, M., Hancock, M., Benko, H., Ofek, E., Wilson, A.D. (2016). Haptic Retargeting: Dynamic Repurposing of Passive Haptics for Enhanced Virtual Reality Experiences. In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16). ACM, New York, NY, USA, 1968-1979. DOI: https://doi.org/10.1145/2858036.2858226

Baños, R. M., Botella, C., Garcia-Palacios, A., Villa, H., Perpiñá, C., & Alcaniz, M. (2000). Presence and reality judgment in virtual environments: a unitary construct?. CyberPsychology & Behavior, 3(3), 327-335.

Barfield, W., Zeltzer, D., Sheridan, T., & Slater, M. (1995). Presence and performance within virtual environments. Virtual environments and advanced interface design, 473-513.

Chittaro, L., & Buttussi, F. (2015). Assessing knowledge retention of an immersive serious game vs. a traditional education method in aviation safety. IEEE transactions on visualization and computer graphics, 21(4), 529-538.

Dominjon, L., Lecuyer, A., Burkhardt, J., Richard, P., & Richir, S. (2005). Influence of control/display ratio on the perception of mass of manipulated objects in virtual environments. IEEE Proceedings. VR 2005. Virtual Reality, 2005. (pp. 19-25). doi:10.1109/vr.2005.1492749

Fujinawa, E., Yoshida, S., Koyama, Y., Narumi, T., Tanikawa, T., & Hirose, M. (2017, November). Computational design of hand-held VR controllers using haptic shape illusion. In Proceedings of the 23rd ACM Symposium on Virtual Reality Software and Technology (p. 28). ACM.

Gonzalez-Franco, M., & Lanier, J. (2017). Model of illusions and virtual reality. Frontiers in psychology, 8, 1125. doi: 10.3389/fpsyg.2017.01125

Han, D.T., Suhail, M., Ragan, E.D. (2018). Evaluating Remapped Physical Reach for Hand Interactions with Passive Haptics in Virtual Reality, Visualization and Computer Graphics IEEE Transactions on, vol. 24, no. 4, pp. 1467-1476, 2018

Hoffman, H. G. (1998, March). Physically touching virtual objects using tactile augmentation enhances the realism of virtual environments. In Proceedings. IEEE 1998 Virtual Reality Annual International Symposium (Cat. No. 98CB36180) (pp. 59-63). IEEE.

Hvass, J., Larsen, O., Vendelbo, K., Nilsson, N., Nordahl, R., & Serafin, S. (2017, June). Visual realism and presence in a virtual reality game. In 2017 3DTV Conference: The True Vision-Capture, Transmission and Display of 3D Video (3DTV-CON) (pp. 1-4). IEEE.

Insko, B. E., Meehan, M., Whitton, M., & Brooks, F. (2001). Passive haptics significantly enhances virtual environments (Doctoral dissertation, University of North Carolina at Chapel Hill).

Kohli, L. (2013). Redirected touching (Doctoral dissertation, University of North Carolina at Chapel Hill).

Leder, J., Horlitz, T., Puschmann, P., Wittstock, V., & Schütz, A. (2019). Comparing immersive virtual reality and powerpoint as methods for delivering safety training: impacts on risk perception, learning, and decision making. Safety science, 111, 271-286.

Lohse, A., Kjaer, C., Hamulic, E., Lima, I., Jensen, T., Bruni, L., & Nilsson, N. (2019). Leveraging Change Blindness for Haptic Remapping in Virtual Environments. *In IEEE 5th Workshop on Everyday Virtual Reality* (WEVR).

Matsumoto, K., Hashimoto, T., Mizutani, J., Yonahara, H., Nagao, R., Narumi, T., Tanikawa, T., Hirose, M. (2017). Magic table: deformable props using visuo haptic redirection. In SIGGRAPH Asia 2017 Emerging Technologies (SA '17). ACM, New York, NY, USA, Article 9, 2 pages. DOI: https://doi.org/10.1145/3132818.3132821

McClelland, J. C., Teather, R. J., & Girouard, A. (2017, October). Haptobend: shape-changing passive haptic feedback in virtual reality. In Proceedings of the 5th Symposium on Spatial User Interaction (pp. 82-90). ACM.

McMahan, R. P. (2011). Exploring the effects of higher-fidelity display and interaction for virtual reality games (Doctoral dissertation, Virginia Tech).

Meehan, M., Insko, B., Whitton, M., & Brooks Jr, F. P. (2002, July). Physiological measures of presence in stressful virtual environments. In Acm transactions on graphics (tog) (Vol. 21, No. 3, pp. 645-652). ACM.

Nilsson, N.C., Nordahl, R., Serafin, S. (2017). Waiting for the ultimate display: can decreased fidelity positively influence perceived realism?, in IEEE 3rd Workshop on Everyday Virtual Reality (WEVR), Los Angeles, CA, 2017 pp. 1-5. doi: 10.1109/WEVR.2017.7957710

Razzaque, S., Kohn, Z., & Whitton, M. C. (2005). Redirected walking (pp. 4914-4914). University of North Carolina at Chapel Hill.

Rietzler, M., Geiselhart, F., Frommel, J., & Rukzio, E. (2018, April). Conveying the perception of kinesthetic feedback in virtual reality using state-of-the-art hardware. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (p. 460). ACM.

Sallnäs, E. L., Rassmus-Gröhn, K., & Sjöström, C. (2000). Supporting presence in collaborative environments by haptic force feedback. ACM Transactions on Computer-Human Interaction (TOCHI), 7(4), 461-476.

Sanchez-Vives, M. V., and Slater, M. (2005). From presence to consciousness through virtual reality. Nat. Rev. Neurosci. 6, 332–339. doi: 10.1038/nrn1651

Simeone, A. L., Velloso, E., & Gellersen, H. (2015, April). Substitutional reality: Using the physical environment to design virtual reality experiences. In Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (pp. 3307-3316). ACM.

Slater, M., Steed, A., & Chrysanthou, Y. (2002). Introduction: a phantom world of projections. In *Computer graphics and virtual environments* (pp. 1-47). Harlow: Addison Wesley.

Slater, M. (2009). Place illusion and plausibility can lead to realistic behaviour in immersive virtual environments. Philosophical Transactions of the Royal Society B: Biological Sciences, 364(1535), 3549-3557.

Undervisningsministeriet. (2018). Fra folkeskole til faglært - erhvervsuddannelser til fremtiden (Isbn 978-87-93635-83-8). Retrieved from: https://www.regeringen.dk/publikationer-og-aftaletekster/fra-folkeskole-til-faglaert/

Usoh, M., Catena, E., Arman, S., & Slater, M. (2000). Using presence questionnaires in reality. Presence: Teleoperators & Virtual Environments, 9(5), 497-503.

Viciana-Abad, R., Lecuona, A. R., & Poyade, M. (2010). The influence of passive haptic feedback and difference interaction metaphors on presence and task performance. Presence: Teleoperators and Virtual Environments, 19(3), 197-212.

Witmer, B. G., & Singer, M. J. (1998). Measuring presence in virtual environments: A presence questionnaire. Presence, 7(3), 225-240.

Zenner, A., Krüger, A., (2017). Shifty: A Weight-Shifting Dynamic Passive Haptic Proxy to Enhance Object Perception in Virtual Reality. In *IEEE Transactions on Visualization and Computer Graphics*, vol. 23, no. 4, pp. 1285-1294. doi: 10.1109/TVCG.2017.2656978

9 ACKNOWLEDGEMENT

We would like to thank our supervisors Dr. Rolf Nordahl and Niels C. Nilsson, for the immense help that you have given to us, making this piece of research an interesting and pleasant experience.
We would also like to extend our thanks to the team behind the Multisensory Experience Lab at Aalborg University in Copenhagen.
Special thanks to Kenneth Jensen for his assistance.
At last, we would like to thank our families and loved ones for their support during the heavy hours.