Utrecht University

Department of Information and Computing Sciences

Game and Media Technology - Master Thesis

Immersive Multimodal Virtual Reality Experiences

Using Visual and Auditory Stimuli to Improve Tactile Experiences

Author:
Nina Rosa

Supervisors:
Dr. Wolfgang Hürst
Prof. Peter Werkhoven

ICA-3654257

August 10, 2015

Abstract

Haptic feedback, such as the sensation of 'being touched', is an essential part of how we experience our environment. Yet, it is often disregarded in current virtual reality (VR) systems. In addition to the technical challenge of creating such tactile experiences, there are also human aspects that are not fully understood, especially with respect to how humans integrate multimodal stimuli. In this research, we proved that the visual stimuli in a VR setting can influence how vibrotactile stimuli are perceived. In particular, we identified how visual cues that are generally associated with the characteristic of weight, namely size and falling speed, influence tactile perception, whereas a similar effect could not be achieved for a temperature-related visual cue, namely color. Such an effect was also not achieved for impact in a realistic precipitation scenario with other typical representative visual and auditory indications, however the results did demonstrate that fundamentally consistent cues, as were used in the case of weight, are necessary to trigger such a response rather than only real-world knowledge. Our results have technical implications – for example, suggesting that a rather simple vibration motor may be sufficient to create a complex tactile experience such as perceiving weight when correctly presented with multimodal stimuli – and relevance for practical implementations – for example, indicating that vibration intensities need to be 'exaggerated' to achieve certain effects.

Acknowledgements

I would like to express my deepest gratitude to my first supervisor Dr. Wolfgang Hürst for his engagement, encouragement and meticulous feedback throughout this master thesis, and also to my second supervisor Prof. Peter Werkhoven for his guidance and assistance. They have both supported me in pursuing research after my master studies, and for that, I am extremely grateful. Special thanks goes out to the staff at Elitac B.V. for supplying the necessary hardware and their valuable insights and advice during this research. Also, I would like to thank the participants in my experiments, who have willingly shared their precious time for my tests. Lastly, I would like to thank my family and friends, who have supported me throughout the entire process and kept me sane during this stressful period.

Table of Contents

Preface

This thesis, entitled *Immersive Multimodal Virtual Reality Experiences - Using Visual and Auditory Stimuli to Improve Tactile Experiences*, investigates how visual and auditory stimuli influence vibrotactile perception in passive virtual reality experiences. Instead of a traditional thesis report, the goal was to deliver a scientific paper aimed at a journal publication. Therefore, the major deliverables of this thesis are:

- A scientific paper aimed for publication in a journal. Because of the successful outcome, results from the first experiment have been summarized in a scientific conference paper and submitted to the International Conference on Multimodal Interaction (ICMI 2015), one of the leading events in computer science explicitly addressing the area of multimodal interaction. Co-authors include my two supervisors, Dr. W. Hürst and Prof. P. Werkhoven, and W. Vos, an external partner from Elitac B.V., the company providing the tactile hardware required for this project. The journal paper will be submitted to a suitable journal once the outcome of the conference paper review has been received. It can be found in Chapter 1 of this document.
- An annotated appendix to complement the scientific paper (Chapter 2 of this document). It provides further information about relevant parts of the thesis project not covered in the papers, most importantly:
  - A detailed initial literature study identifying the most promising and relevant areas to investigate in this project.
  - Detailed information on the research done in this thesis, but not covered or included in the scientific paper.
  - Information about further thesis contributions; most importantly a demonstration game illustrating the scientific results of the thesis in a practical example.

Further contributions and deliverables of this thesis include:

- Related source code and programs for the virtual environments and testing infrastructure used in the scientific experiments summarized in the two papers and the final demo game. Information about the implementation and development work can be found in Sections 3.2, 4.2, and 5.2 of the annotated appendix in Chapter 2.
- A website containing a description, media (video and images from the paper and demo), and results, targeted at a broad audience: `www.jelmerdejongonline.nl/ninammvr/`

Chapter 1

Scientific Paper

In this section, the scientific paper is introduced, which outlines the results and findings from this thesis. The experiments described in Section 3 were the basis of the conference paper submitted to the International Conference on Multimodal Interaction (ICMI 2015).

Altering Tactile Perception through Abstract and Realistic Multimodal Stimuli in Virtual Reality

Rosa, Nina
Utrecht University
Information and Computing Sciences
Utrecht, Netherlands

ABSTRACT

Haptic feedback, such as the sensation of 'being touched', is an essential part of how we experience our environment. Yet, it is often disregarded in current virtual reality (VR) systems. In addition to the technical challenge of creating such tactile experiences, there are also human aspects that are not fully understood, especially with respect to how humans integrate multimodal stimuli. In this research, we proved that the visual stimuli in a VR setting can influence how vibrotactile stimuli are perceived. In particular, we identified how visual cues that are generally associated with the characteristic of weight, namely size and falling speed, influence tactile perception, whereas a similar effect could not be achieved for a temperature-related visual cue, namely color. Such an effect was also not achieved for impact in a realistic precipitation scenario with other typical representative visual and auditory indications, however the results did demonstrate that fundamentally consistent cues, as were used in the case of weight, are necessary to trigger such a response rather than only real-world knowledge. Our results have technical implications – for example, suggesting that a rather simple vibration motor may be sufficient to create a complex tactile experience such as perceiving weight when correctly presented with multimodal stimuli – and relevance for practical implementations – for example, indicating that vibration intensities need to be 'exaggerated' to achieve certain effects.

Categories and Subject Descriptors

H.1.2 [**Models and Principles**]: User/Machine Systems; H.5.1 [**Information Interfaces and Presentation**]: Multimedia Information Systems—*Artificial, augmented, and virtual realities*

General Terms

Human Factors

Keywords

Virtual reality, passive touch, multimodal experiences

1. INTRODUCTION

Virtual reality (VR) set-ups generally aim at creating a high level of realism. Yet, most current systems only support vision and audio, omitting other senses, such as haptic experiences. While adding, for example, the sensation of *passive touch*, (e.g., being touched by a character or object) would likely increase realism, implementing this is very difficult. First, there is the technical challenge: tactile devices are only able to simulate one single property of touch, such as weight or temperature of an object. In addition, it is not straightforward how to integrate the haptic sensation with other senses. Past research studied, for example, how the haptic, visual and auditory sense work together, and how they can change the perception of an object's texture [23, 28, 10, 21], hardness and stiffness [12, 3], temperature [16, 13], and weight [8, 31, 11, 1].

Inspired by these studies, many of which used real haptic objects, our general aim is to verify if similar experiences exist in pure virtual environments. To accomplish this, an initial exploratory experiment was conducted to investigate if visual signals that are commonly associated with the two characteristics *weight* and *temperature* have an influence on vibrotactile perceptions, i.e. tactile feedback created by vibration motors placed on, for example, your arm. Ideally, we hope that the richer visual stimuli can influence this rather simple tactile perception in a way that creates a better, more realistic VR experience. We show that this is indeed the case for weight and also identify that vibration intensity should be adjusted to make differences in weight appear more realistic. Using visual cues indicating temperatures does however not show similar effects.

In a second experiment we examine these findings further, by investigating whether more complex visual and auditory cues associated with precipitation impact, learned by real world knowledge, can influence the perception of vibrotactile sensations. We show that in this case tactile perception depends on both the presented modalities and the precipitation intensity; the presented cues must be fundamentally consistent, as in the case of weight, in order to influence tactile perception in virtual reality, and relating these cues to real world knowledge is insufficient to evoke such influence.

In addition to these concrete findings, we also identify potential issues concerning general psychophysical experiments and the body transfer illusion, which may lead to interesting research problems worth further investigation.

The structure of this paper is as follows. In Section 2, background information on the topics VR, presence, and multimodal integration is given. The experiments are described and discussed in Sections 3 and 4. Section 5 closes with a conclusion and directions for future research.

2. BACKGROUND AND RELATED WORK

2.1 Virtual Reality and Tactile Sensations

Current consumer VR devices generally address the visual and auditory component but lack in modalities associated with tactile, gustatory, and olfactory senses. While there are tools for motion tracking that also support kinesthetic senses, we are still far from being able to create a full, realistic tactile VR experience. Good overviews on how to simulate haptic feedback can be found in [29] (addressing classifications and techniques) and chapter 3.3 of [6] (focusing on devices). Newer techniques not covered in these references include, for example, electrovibrations adding tactile feedback to touch screens [17]. Vibrotactile displays, i.e. vibration motors or *tactors* placed on your skin, can be used to create the illusion of touch through vibration.

Although these approaches have proven their feasibility, not all are equally applicable. Especially for consumer use, common aims are inexpensiveness and expressiveness, to impose limited cumber upon the user, that the system should be easily scalable and reconfigurable, and that the mapping should be straightforward [18]. The vibrotactile displays used in our work to create the illusion of passive touch satisfy all of these requirements besides the last. The *general aim of our research* is therefore to provide better insight about the integration of such devices in the overall interaction experience and increase knowledge about mutual dependencies between stimuli of different modalities.

2.2 Presence and Passive Experience in VR

According to the model by Steuer [30], presence in VR is a human experience and a consequence of immersive technologies. It has two determining dimensions: vividness with the two contributing factors breadth and depth of included modalities; and interactivity with the three contributing factors speed, range, and mapping. Many studies target increasing interactivity by focusing on improving task performance [2, 24, 5]. Yet, when the goal is to improve experiences that users can take part in passively, vividness becomes a vital part of the system, even more urgent than interactivity. Here, we focus on this less studied aspect by specifying the *sub-goal of our research* as investigating the experience of passive touch under varying visual and auditory conditions.

A phenomenon closely related to presence is the body transfer illusion. Such an illusion causes users to perceive part of or an entire artificial body as their own. For this reason, it is believed that it can increase presence in a virtual environment [25]. Factors that enhance this illusion are first person perspective over a fake humanoid body and congruent visuotactile cues [19, 22]. It has been shown that a first person perspective of a life-sized virtual human body alone is sufficient to generate a body transfer illusion [27]. In our experiments we thus apply this methodology focusing on first person body experiences.

2.3 Multimodal Integration

Research from cognitive science identified important illusions with respect to multimodal integration, including the McGurk effect [20], ventriloquist effect [15], double-flash illusion [26], and the rubber hand illusion [4] (an example of the body transfer illusion). Other illusions that specifically concern the tactile sense are changing surface texture and roughness [23, 28, 10, 21], changing object hardness and stiffness [12, 3], changing object temperature [16, 13], and changing object weight and collision force [8, 31, 11, 1]. Many of these studies used the actual tactile property to test the effect of adding visual or auditory cues. However, simulating all of these properties is not feasible in a VR system, thus a substitution is needed. The *concrete research problem addressed in this paper* is therefore to investigate whether the findings above would still apply when using vibrotactile stimuli in a passive situation, rather than an active task such as in [11].

3. EXPERIMENT 1: ABSTRACT WEIGHT AND TEMPERATURE

In the following, we outline the objective of the experiments (3.1), specify the general setup and methodology (3.2 and 3.3). Also, the unique characteristics and results of the weight and temperature tests are described (3.4 and 3.5) and discussed (3.6).

3.1 Objective

Resulting from the general aim, sub-goal, and related research problem introduced in the preceding section, we phrase our first *research question* as:

> *Is it possible to create the illusion of experiencing different intensities of a certain property (weight or temperature) using a rather simple and unrelated type of touch (vibrations) together with compelling, type-related visuals (speed/size and color, respectively)?*

We address this via an exploratory investigation where the two properties weight and temperature of a 'touching' object are studied. Intensity of each property will be varied: tactile intensity is mapped to a vibration intensity; and visual intensity is controlled by different cues motivated by related work (cf. 2.3), cf. Table 1.

The *objective* is therefore to measure the perceived weight and temperature intensity with a range of visual intensities. To do this, we use two psychophysical experiments in the form of a matching task (cf. 3.3). We motivate this methodology by the fact that experience is a qualitative characteristic and thus difficult to verify objectively. Although the used matching task is technically a performance verification, it verifies how humans perceive the intensity of varying multimodal stimuli, thus also allowing us to draw conclusions about their expected experience.

Property	Cue	Levels (low - med - high)		
Temperature	Color	Blue	- Gray	- Red
Weight	Size & Speed	Small & Slow	- Med & Normal	- Large & Fast

Table 1: Tactile properties with corresponding visual cues and intensity levels used in the first experiment.

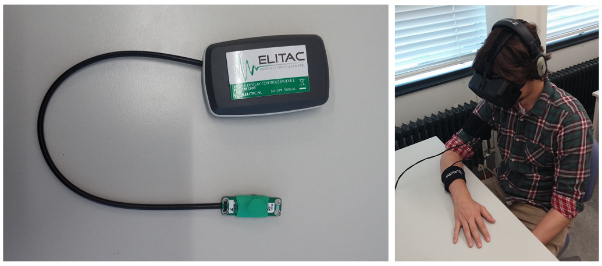

Figure 1: The measurement setup; (left) the Elitac control module and a single tactor, (right) a participant sitting in the correct position wearing the HMD, tactile sleeve and headphones.

3.2 Framework and Material

The same framework is used for the following two experiments. A plain indoor room was created with Unity Pro version 4.6.3f1, and scripts were written in C# using Microsoft Visual Studio 2013. Assets used to create the environment were free downloads from the Unity Asset Store. An Oculus Rift Development Kit 2 is used as head-mounted display (HMD) to create the visual stimuli and VR experience. It has a resolution of 960×1080 pixels per eye, and a nominal field of view of 100 degrees. Tactile stimuli are supplied via vibrations through an Elitac tactile display, with one tactor (vibration motor) and a control module attached to the participant's arm using an elastic band with Velcro, cf. Figure 1. This tactile display has 16 intensity levels on a logarithmic vibration power scale corresponding to a linear perceived intensity scale with fundamental frequency 158 ± 2.4 Hz at maximum vibration strength. The root mean square acceleration at maximum vibration strength is 55.5 ± 9.5 m/s^2. Sennheiser HD201 headphones are used to produce pink noise in the background to mask the sound of the tactor and eliminate all other acoustic influences.

3.3 Method and Procedure

Before starting any experiment, participants were seated in a neutral room, signed a consent form and filled out a general information form. All subjects volunteered and were not reimbursed for their time. The experimenter gave instructions on how they should be positioned, and aided them while putting on the HMD and tactile display. Participants familiarized themselves with the virtual environment by looking around. In the setup, the avatar was sitting on a chair with the left arm placed on their lap under the table (out of sight), and the right arm on the table such that the lower arm was resting horizontally in front of them. The avatar was a humanoid figure (its gender matched the participant's) placed such that the participant had a first-person perspective; cf. Figure 2. Before each condition, a short training session took place, to make sure the participant understood the procedure and to induce a body transfer illusion. Because of the first person perspective, we can assume that this body transfer illusion occurs for all subjects [27].

Each subject tested two conditions: the control (no indicative visual cues), and either the weight or temperature conditions. The order of the conditions was counterbalanced over all participants. The conditions were made up of trials, each trial consisting of a matching task. A matching task uses a method of adjustment described in [9] (Ch. 3). In a general matching task, participants are presented with a stimulus that must be adjusted such that it is above or below a certain threshold. In our experiments, a variation was used in which the participant had to adjust the tactile comparison stimulus until it matched (was perceived as) the tactile reference stimulus. The task is further elaborated in sections 3.4.2 and 3.5.2. The participant's final choice of the matching comparison stimulus was logged. With such a task, we are able to obtain objective measurements about the perception of stimuli.

There were three tactile reference intensity levels (5, 8, 11), three visual intensity levels (0, 1, 2), two starting tactile comparison intensity levels (2 levels lower and 2 levels higher; i.e. 3 and 7 for reference 5), and four repetitions for each starting level, resulting in 72 trials in total. The reference intensity levels were chosen from a practical point of view: the levels are approximately in the center of the allowed range, such that the participants had room for adjustment in both directions. The control condition had no visual intensity levels, thus leading to 24 trials; resulting in a total of 96 trials per participant. A break of about three minutes was taken after every 24 trials (so three breaks in total), where the participant was allowed to take the HMD off. After taking part in the experiment, each participant was verbally asked two questions:

1. Did you have the feeling the virtual arm was your own arm?
2. Did you use the visual weight/temperature of the ball to deduce your answer or did you use the intensity of the sensation on your arm?

While these subjective measures were not the main focus of this study, they can provide additional insight especially when interpreting the relevance of the measured quantitative data. Duration of the experiment for one subject was approximately one hour.

3.4 Experiment 1A: Weight

3.4.1 Participants

Twelve subjects took part in the weight experiment. Ages ranged from 21-27 (average 22.8). Eleven male and one female, eleven were right-handed, and one was mixed-handed. Four had no prior experience with HMDs, seven had taken part in one or a few demonstrations, and one owned one personally. Due to the basic characteristic of the task and related measures, we do not expect that these differences have any influence on the outcome nor did we observe any related indications – neither during the tests nor in the related data analysis. Likewise, we verified visual clarity during the training session in order to reduce any related influence due to colorblindness (two subjects) and prescription glasses (six subjects), and checked that none had any restrictions concerning touch sensations on the skin. Data of an additional thirteenth participant had to be discarded due to a system error that appeared during the test.

3.4.2 Matching Task

Each experiment consisted of the actual weight test, and a pure tactile control condition (without any visuals). At the

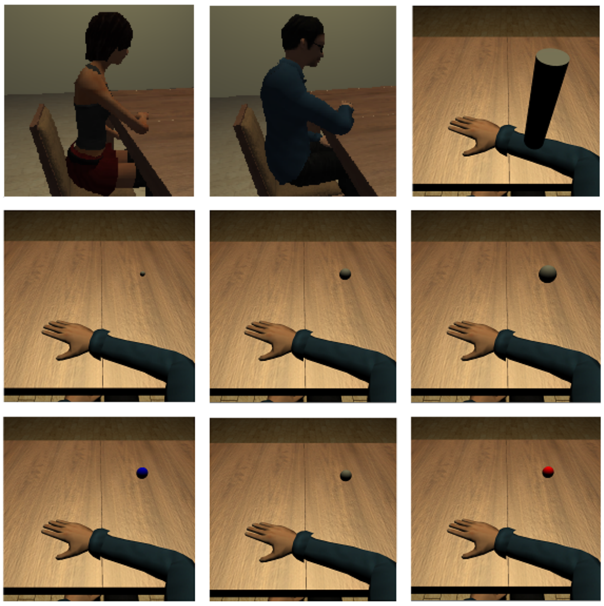

Figure 2: The virtual environment: (top-left) the female avatar, (top-center) the male avatar, (top-right) the occluder for the control condition, (center row) weight visuals where the steps in speed were exponential, (bottom row) temperature visuals.

beginning of the weight experiment, the participant received the following instructions: 'In this experiment, you will feel a virtual object with a certain weight falling on your right arm. In each round, the object will first be blocked from your view, and the second time you will be able to see it. These two sensations on your arm may feel different. After feeling both, you must indicate whether the second ball felt as heavy as the first, or that the second felt lighter or heavier than the first. The goal is to adjust the tactile sensation such that you experience them the same.'

In each trial, the participant was presented with a reference and a comparison stimulus. In the weight condition, the reference stimulus was a tactile-only stimulus, while the comparison stimulus was a combination of a tactile and a visual stimulus: a sphere that was small and slow, medium and normal, or large and fast. For the reference stimulus, an occluder in the form of a cylinder would appear, covering the vertical trajectory the sphere would cover; cf. Figure 2. Then two audio beeps with a 700 ms interval would occur, and after another 700 ms the tactor on the arm was activated at intensity level x for 200 ms so the participant felt a vibration. This was followed by the comparison condition, where the occluder disappeared. After one beep a gray sphere with a certain size appeared in mid-air at eye-level and remained in this position for 700 ms. Then another beep occurred, after which the sphere fell in a vertical trajectory onto the virtual arm with a certain speed. Upon contact, the tactor was activated at intensity y for 200 ms and the sphere disappeared. At this stage the participant replied whether the sensations were the same, or whether the second was lighter or heavier than the first. If they answered that they were the same, the trial was over, and a new reference stimulus was presented. Otherwise, the tactile comparison intensity was adjusted to $y + 1$ if the participant answered 'lighter', or to $y - 1$ if they said 'heavier'. The reference and comparison stimuli were then presented as before, but now the tactile comparison intensity was at the adjusted level. This presenting and adjusting continued until the participant felt that the sensations matched. Table 2 outlines a single trial. All *tactile reference intensity - visual comparison intensity - final tactile comparison intensity* combinations were logged separately for each participant.

Step	Experiment
1	Tactile Reference α
	Tactile Comparison μ + Visual x
2	Tactile Reference α
	Tactile Comparison ν + Visual x
$\vdots$	$\downarrow$ adjust tactile
m	Tactile Reference α
	Tactile Comparison π + Visual x

Table 2: Outline of a single trial in the weight and temperature conditions. Gray cells indicate the presence of the occluder.

For the control condition, the participant was given no verbal or visual indication of the tactile property of the sphere. Terminology was changed to 'less intense'/'more intense' rather than 'lighter'/'heavier'. Also, the comparison stimulus was identical to the reference stimulus. That is, the spheres were always visually blocked by the occluder. Table 3 outlines a single trial in the control condition. All *tactile reference intensity - final tactile comparison intensity* combinations were logged separately for each participant.

Step	Experiment
1	Tactile Reference α
	Tactile Comparison μ
2	Tactile Reference α
	Tactile Comparison ν
$\vdots$	$\downarrow$ adjust tactile
m	Tactile Reference α
	Tactile Comparison π

Table 3: Outline of a single trial in the control condition. Gray cells indicate the presence of the occluder.

3.4.3 Results

First, we discuss the method for handling the results of the weight experiment. Consider one particular participant. The mean of all answers (eight values) is determined for every combination of tactile reference intensity and visual comparison intensity. This results in twelve values per participant. These values are then normalized by subtracting 5, 8 or 11, depending on the tactile reference intensity. The normalized means are given in Table 4 and are visualized in Figure 3. Normality was tested using Shapiro-Wilk tests; four cases were not normally distributed: *Ref5-Control* ($W(12) = 0.833; p = 0.023$), *Ref5-Visual1* ($W(12) = 0.752; p = 0.003$), *Ref5-Visual2* ($W(12) = 0.643; p < 0.001$), and *Ref8-Visual2* ($W(12) = 0.814; p = 0.014$).

	Ref5	Ref8	Ref11
Control	0.292	-0.094	-0.385
Visual0	0.177	-0.198	-0.708
Visual1	0.271	-0.073	-0.573
Visual2	0.510	0.167	-0.333

Table 4: Normalized means over all conditions of the weight experiment.

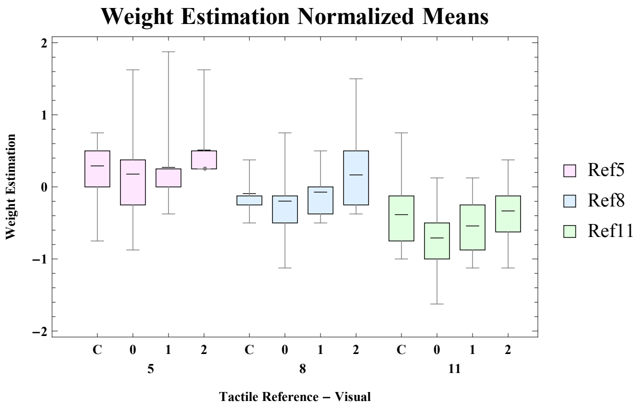

Figure 3: Boxplots of the normalized results of the weight experiment with mean markers. The labels *5*, *8*, *11* refer to the tactile reference intensity, *C*, *0*, *1*, *2* to the visual case.

The data was analyzed with a two-way repeated measures ANOVA, with factors visual comparison intensity (three levels) and reference intensity (three levels). Mauchly's Test of Sphericity indicated that the assumption of sphericity had been violated in the reference intensity factor ($\chi^2(2) = 12.530, p = 0.002$), and a Greenhouse-Geisser correction was used for this factor. The analysis gave the following results: there was a significant main effect over reference intensity ($F(1.167, 12.833) = 45.438; p < 0.00001$), a significant main effect over visual intensity ($F(2, 22) = 8.855; p = 0.002$), but the interaction effect was not significant ($F(4, 44) = 0.028; p = 0.998$). Pairwise comparisons with a Bonferroni adjustment showed that each reference intensity differed significantly from the others (all $p < 0.001$), and that visual intensities 0 and 2 ($p = 0.005$) and 1 and 2 ($p = 0.013$) differed significantly. Nine paired t-tests were run between control results and each visual intensity within each reference intensity. These showed that the results of none of the visual conditions differed significantly from those of the corresponding control.

Qualitative results of the post-experiment questions are summarized in Table 5 in column *Weight*. Three participants answered that they had a strong feeling that the avatar's arm was their own, five had a less strong feeling, and four responded to not feeling any connection at all. Regarding using visual weight or intensity, two participants responded that they used the size and the speed of the spheres, and the other ten used the vibration intensity, of which four stated they tried to use the size and weight, but found it unreliable after a few rounds.

Question	Answer	Weight	Temperature
1. Virtual arm your own	strong	3	2
	weak	5	2
	none	4	6
2. Visual or tactile	visual	2	0
	tactile	10 (4)	10 (4)

Table 5: Answers to the post-experiment questions by experiment. The number in brackets indicates the number of participants who initially used visual cues.

	Ref5	Ref8	Ref11
Control	-0.113	-0.150	-0.463
Visual0	0.075	-0.213	-0.250
Visual1	0.050	-0.338	-0.475
Visual2	-0.188	-0.188	-0.425

Table 6: Normalized means over all conditions of the temperature experiment.

3.5 Experiment 1B: Temperature

3.5.1 Participants

Ten participants took part in the temperature experiment – eight male, two female, ages from 22-24 (average 23.2), nine right-handed, one ambidextrous. Four had no prior experience with HMDs, four had taken part in one or a few demonstrations, and two had worked on projects using one. None were colorblind, two wore prescription glasses, and none had any restrictions concerning touch sensations on the skin. Again, we do not see any influences of these parameters on the evaluation results. Data of three additionally tested subjects had to be discarded due to one system error and two premature terminations do to the subjects being unfit as a result of strain on their eyes.

3.5.2 Matching Task

The method used during this experiment is equivalent to that described in section 3.4.2, now using the terminology 'temperature' instead of 'weight', and 'colder'/'warmer' instead of 'lighter'/'heavier'. In addition, all spheres have the same size and falling speed, but different colors: blue, gray and red, which are commonly associated with cold, neutral, and warm temperature, respectively; cf. Figure 2.

3.5.3 Results

The results were normalized as in the previous experiment, see section 3.4.3. The normalized means are given in Table 6 and visualized in Figure 4. Normality was tested using Shapiro-Wilk tests, which showed that three cases were not normally distributed: *Ref5-Control* ($W(10) = 0.843; p = 0.048$), *Ref5-Visual0* ($W(10) = 0.840; p = 0.044$), and *Ref8-Visual1* ($W(10) = 0.810; p = 0.019$).

Mauchly's Test of Sphericity indicated that the sphericity assumption had been violated in the interaction effect ($\chi^2(9) = 18.643; p = 0.032$), so a Greenhouse-Geisser correction was used for this effect. A two-way repeated measures ANOVA showed that there was no significant main effect over reference intensity ($F(2, 18) = 2.028; p = 0.161$), no significant main effect over visual intensity ($F(2, 18) = 0.946; p =$

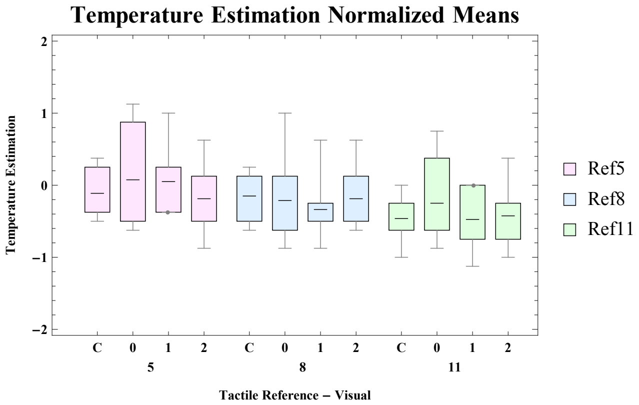

Figure 4: Boxplots of the normalized results of the temperature experiment with mean markers. The labels *5*, *8*, *11* refer to the tactile reference intensity, *C*, *0*, *1*, *2* to the visual case.

0.407) and there was also no significant interaction effect ($F(2.037, 18.332) = 1.817; p = 0.190$).

Although there was no significant main effect over reference intensity in the results of the temperature experiment, the results of both experiments are combined to investigate this effect further. Three independent samples t-tests were run for each reference intensity, which showed that only the results for reference intensity 5 differed significantly between groups ($t(86) = 3.293; p = 0.001$). In the same way, three independent sample t-tests were run for the control results, which showed that again only the results of reference intensity 5 differed significantly between groups ($t(20) = 2.479; p = 0.022$). Normality was tested for each reference intensity with Shapiro-Wilk tests, which showed that reference intensities 5 and 8 were not normal ($W(88) = 0.959; p = 0.007$ and $W(88) = 0.950; p = 0.002$). Sphericity was violated according to Mauchly's Test ($\chi^2(2) = 28.988; p = 0.000001$) so a Huynh-Feldt correction was used. A one-way repeated measured ANOVA with factor reference intensity (three levels) showed that there was a significant main effect ($F(1.578, 137.298) = 58.238; p < 0.0000001$), and pairwise comparisons with Bonferroni correction showed that each comparison was significant ($p < 0.000001$).

Qualitative results of the questions are summarized in column *Temperature* of Table 5. Two participants answered that they had a strong feeling the avatar's arm was their own, two had a less strong feeling, and six responded to not feeling any connection at all. All participants used the vibration intensity to come to an answer, of which four stated they thought they may have been influenced by the color of the sphere.

3.6 Discussion

In these experiments the research question if it is possible to create the illusion of experiencing different intensities of a certain property (weight or temperature) using a rather simple and unrelated type of touch (vibrations of different intensities) together with compelling, type-related visuals (speed/size and color, respectively) was investigated. Using a matching task, we speculated that related quantitative measures enable us to identify if people are likely to associate such tactile experiences with these properties and if or how concrete intensity levels have to be set when visual stimuli change. We hoped that the richer visual stimuli can be used to change a user's tactile perception – ideally in a way that enables us to simulate more complex tactile experiences.

This conjecture was verified for the weight property, but not for the temperature property. In the weight experiment, the consistent increase in equally perceived tactile intensities when visual intensity was increased means the sensation felt on the skin was estimated 'lighter' when the visual sphere looked heavier, even though the true tactile vibration intensities were held constant. This is consistent with the speed-force illusion [1] and the size-weight illusion [31]. However, after analysis of the post-experiment responses, it was clear that almost all participants eventually used the intensity of the vibrations in order to respond in each trial. This means that the visual cues caused participants to subconsciously generate an expectation which changed not only the intensity of the felt sensation, but possibly also the type. However, consciously no association to weight was made.

While we can thus *not* conclude that people will associate such related vibrations with the characteristic of weight, these observations are essential for system design when including tactile feedback into such a VR setting. Roughly speaking, if a visual stimuli suggests a higher weight, but the matching tactile feedback is experienced as less intense, there is a high likelihood that the illusion of weight is not only not apparent but even 'destroyed'. The added modality would then not increase the experience but make it even worse, whereas an adjustment of intensity as suggested by our results will guarantee a 'perceived' match of both modalities.

In the temperature experiment, no significant differences were found over visual intensity. The question that arises is why this occurred for the property weight, but not for temperature. The participants could either not associate vibrations with temperature, or color with temperature, or both.

An explanation concerning the first scenario is that impact induced by the weight of an object can either be felt through a vibration, or is easier to translate from a vibration than the temperature of an object – a reasoning that can be supported by examining different receptors found in human skin. Mechanoreceptors respond to mechanical pressure and distortion, and there exist different types: Pacinian corpuscles, Meissner corpuscles, Merkel receptors, Ruffini end organs, receptors in hair follicles. Each one responds to vibrations in a certain frequency range. Since impact and vibration are of the same nature, i.e. they both cause skin distortion, it may be the case that both sensations can be felt, and which is finally chosen is caused by surrounding/congruent factors. In the case of the change in weight, the participants were visually expecting a sense of impact to occur, and not a vibration. Changes in temperature, on the other hand, are sensed through thermoreceptors, causing the perceptual 'gap' between vibration and temperature to be too large to bridge despite the visuals. The second scenario, i.e. that participants could not associate temperature with color, was disproved in a study by Ho et al. [14]. By objective performance measurements, it was shown that humans make a cross-modal correspondence between color and temperature, specifically in the direction color-temperature, and not temperature-color.

An interesting observation can be made when looking at the results with respect to the order of presentation. Because we used a method of adjustment, we had to present the reference stimulus first, followed by the comparison that subjects had to adjust. This is in contrast to references (e.g., [9]) suggesting a counterbalanced representation in order to avoid a shift in results; generally, for psychophysical experiments, it is assumed that the second stimulus is often judged as 'greater' than the true equality. While in our case, 'greater' does not necessarily correspond to a higher vibration intensity, it does however suggest that a shift in one direction may occur. A significant main effect was found over tactile reference intensity in the weight experiment, but not in the temperature experiment. The combined results, however, show this effect as well. This means that low vibrotactile intensities are experienced as weaker than they really are, and high vibrotactile intensities as stronger. Because this shift is not consistently in the same direction, it may indicate an interesting aspect with respect to the design of psychophysical experiments worth investigating in future research.

Another noteworthy aspect appears when examining the post-experiment answers with respect to the so-called rubber hand illusion [4] – a body illusion experiment where subjects 'experience' a rubber hand as being their own due to a visual stimulus on the artificial hand (e.g., being stroked by the object) that is congruent to a real tactile sensation on their (visually hidden) real hand. The results from Table 5 suggest that participants in the weight experiment were able to generate a stronger rubber hand illusion than those in the temperature experiment. This in turn suggests either that when visual and haptic information are 'more congruent' then a strong rubber hand illusion occurs, or that when a strong rubber hand illusion occurs then visual and haptic information are 'more congruent'. The findings of [22] are in line with the first suggestion; correct visual perspective together with correlated multisensory information trigger a strong illusion. However, in [19] it was observed that incongruent cues are not experienced as incorrect when the illusion was strong, which is in line with the second suggestion. It is clear that these elements are strongly related, however at this stage it is unclear how.

Two rather curious artifacts worth mentioning are first, that the results from the control condition in the case of tactile reference intensity 5 differed significantly between experiment groups, and not in the case of intensities 8 and 11. However, the control conditions were identical in both experiments and thus this difference should only be caused by chance. Second, from the answers to the post-experiment questions it was clear that the body transfer illusion was not very strong for most participants. This is in contrast to our assumption that a first person view will be sufficient to achieve this (which in turn was based on related research; cf. section 2.3). Because a first person experience is not essential for the tests presented in this paper, we do not suspect any related impact on our results. The observed effect does however raise an interesting question for future research. Is this lack of body transfer just due to the used setup (e.g., several participants noted that finger movements in the real world were not directly mapped to the virtual world) or does the introduced combination of modalities, i.e., tactile and visuals have an effect as well?

4. EXPERIMENT 2: PRECIPITATION IMPACT

In this section, the objective (4.1), and setup and methodology for the follow-up experiment are described (4.2 - 4.5) followed by the results and discussion (4.6 and 4.7).

4.1 Objective

Experiment 1 showed that indications of weight in the form of size and falling speed altered tactile perception. Yet, the scenario was abstract in order to eliminate other factors; a more likely setting is one which resembles the real world. In such a scenario, the user can deduct properties from other cues derived from real world knowledge too, both visually and audibly.

This conjecture is tested in a next experiment where tactile impact is investigated in combination with both visual and auditory cues in a realistic scenario, i.e. the cues are learned from real world knowledge rather than exponentially growing type-related scales as in the weight experiment. Terminology was changed to impact rather than weight, focusing more on the sensation directly related to the collision and not the underlying reason. Both visual and auditory cues were applied in order to extensively investigate a multimodal effect. Precipitation was chosen as the feature with varying intensity since this consists of objects that cause a distinct degree of impact with each material phase. A potential drawback was that wetness could not be simulated, but since we are specifically testing impact, we do not expect this to have a negative effect on the quantitative results. This leads to the second *research question*:

> *Is it possible to create the illusion of experiencing different intensities of impact using an unrelated type of touch (vibrations) by realistic visual or auditory stimuli that are learned from real-world knowledge?*

As before, tactile intensity is mapped to a vibration intensity; and visual and auditory intensity is controlled by different cues commonly assumed to represent the related characteristic, cf. Table 7. The *objective* is to measure the perceived impact intensity with a range of visual and auditory intensities. Again, a matching task is used to accomplish this.

Property	Cue	Levels (low - med - high)
Impact	Precipitation	Snow - Rain - Hail

Table 7: Tactile property with corresponding cues and intensity levels used in the second experiment.

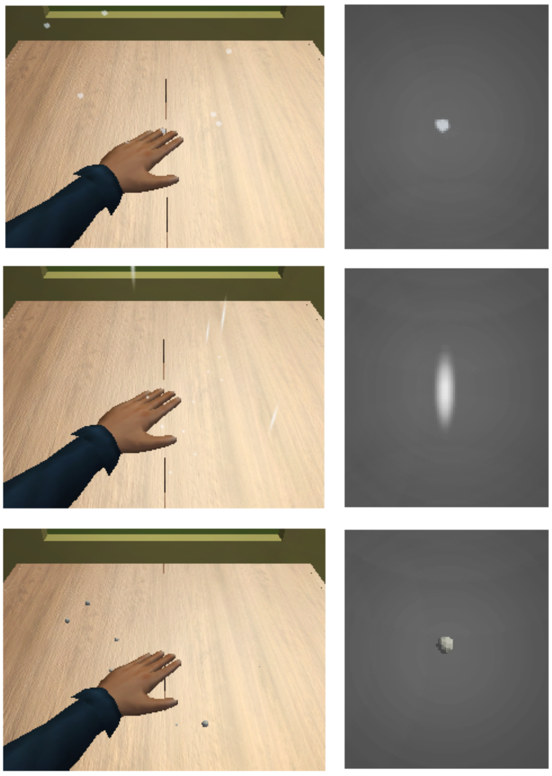

Figure 5: The precipitation intensities as seen by a male participant: (top row) snow, (center row) rain, (bottom row) hail.

4.2 Framework and Material

A similar framework was used as in the previous experiments. Precipitation was visualized with the *Simple Weather Pack* asset by *Caffeine Powered Technologies*, of which the parameters were adjusted to make them visually appealing, and as distinct as possible from each other; cf. Figure 5. Audio recordings of a blizzard, rain and a hail storm from YouTube were used:

- Meditation To The Sound Of A Blizzard (HD)
  Cryptic Mind
  `https://www.youtube.com/watch?v=GQAiB9Vlbh0`
- "Rain Sounds" with no Music 90 mins "Sleep Sound"
  TexasHighDef
  `https://www.youtube.com/watch?v=yhx3JIF7bGU`
- Hail Storm 60 minutes background Sound Effect
  Audio Productions
  `https://www.youtube.com/watch?v=alaf41iMn34`

These were chosen such that the various levels were as distinct as possible from each other, however all with equal volume. The Oculus Rift Development Kit 2, the Elitac tactile display and Philips SHC8535 wireless headphones were used, which constantly played pink noise in the background to mask the sound of the tactors. Since precipitation falls over a certain area rather than one specific location, the tactile display was reconfigured to have three tactors, one for the backside of the hand and two along the back of the forearm, cf. Figure 6.

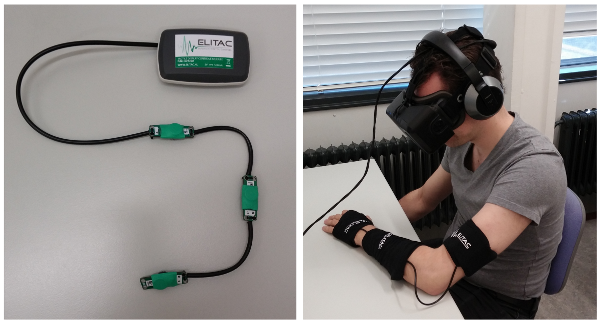

Figure 6: The measurement setup; (left) the Elitac control module and three tactors, (right) a participant sitting in the correct position wearing the HMD, wireless headphones and tactile sleeve.

4.3 Method and Procedure

As before, participants were seated in a neutral room, signed a consent form and filled out a general information form. All subjects volunteered and were not reimbursed for their time. Participants were positioned identical to the virtual environment, with their left hand placed in front of them and their left arm resting on the table. Their right hand was placed under the table out of sight, where a keyboard was placed such that they could use the arrow keys, cf. Figure 6. The avatar was a humanoid figure (its gender matched the participant's) placed such that the participant had a first-person perspective; cf. Figure 5. The virtual environment was a one-roomed house with windows and a hole in the roof, placed in a valley without any further features, cf. Figure 7. The participants were allowed to familiarize themselves with the environment by looking around, and experienced a visual-audio demonstration (without tactile feedback) of each form of precipitation.

Each subject was tested in four modality conditions: control (C), visual (V), audio (A), visual-audio (VA). Each of these had tactile feedback, and the control condition had no visual and no auditory feedback. Before each condition, a short training session took place. The conditions were pairwise counterbalanced, i.e. conditions with audio were always performed one after the other, and either preceded or followed by the conditions without audio, e.g. VA-A-C-V and C-V-A-VA. This was chosen over normal counterbalancing in order to split the experiment over audio-inclusive and audio-exclusive cases. As before, every trial consisted of a matching task, equivalent to those described in sections 3.4.2 and 3.5.2; the task for this experiment is further elaborated in section 4.5.

The previous experiment showed no need for multiple tactile reference intensity levels, thus only one level (7) was used as the reference intensity, with four different starting tactile comparison intensity levels (3, 4, 10, 11), with two repetitions for each starting level. Four was chosen rather than two as in the previous experiments to create a greater sense of randomness within the trials for the participants. Corresponding to the three precipitation intensity levels (0, 1, 2), there were three visual intensity levels, and three audio intensity levels; these were not used in the control condition.

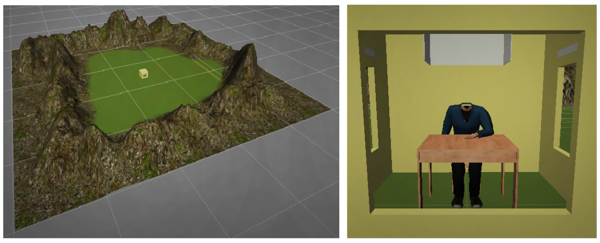

Figure 7: The virtual environment: (left & center) the valley and house, (right) the inside of the house.

This gave a total of 80 trials per participant. After every condition, the participant was allowed to take a break of approximately three minutes during which they would take off the HMD. After the experiment, each participant filled out a questionnaire with questions regarding their experience:

1. Did you have the feeling the virtual arm was your own arm?
2. How convincing did you find the visuals?
3. How convincing did you find the audio fragments?
4. How convincing did you find the tactile sensations?
5. Did you feel like things were actually falling on your arm?
6. Did these falling things feel more real when visuals and/or audio were added?
7. Did you use the visual, audio and/or tactile impact of the precipitation to deduce your answer?
8. What was your overall opinion of the experiment?

The first seven questions were answered using seven-point Likert scales. The entire procedure took approximately one hour per participant.

4.4 Participants

Sixteen participants were recruited for this experiment, none of which took part in the previous experiments. Ages ranged from 19-27 (mean 23). There were fourteen males and two females, and due to the basic characteristic of the task we do not suspect any gender influence on the results. Fifteen were right-handed and one mixed-handed. Four had no previous experience with visual VR (e.g. with an HMD), three had experienced one or a few demos, seven had taken part in another unrelated VR experiment, and two had previously designed and performed their own experiment. Seven had no previous experience with tactile VR (e.g. with a tactile suit), seven had previously taken part in another unrelated tactile experiment, and two had previously designed and performed their own experiment. Ten were near-sighted, of which four never used correction, two had glasses which they took off, and four wore contacts. Visual clarity was confirmed during the visual-audio demonstration. No participants indicated having any restrictions to their touch sense.

4.5 Matching Task

At the beginning of the experiment, the participant received the following instructions: 'In every round you will feel two sensations, each of which is preceded by a beep. The goal is to make these sensations match: after feeling both you must decide whether the second sensation needs to be more or less intense in order to match the intensity of the first. You will make these adjustments yourself using the up and down arrow keys on the keyboard placed by your right hand. Once adjusted, the sensations will be presented again according to your adjustments. If you would like to repeat the sensations you just felt, you can press the left arrow key. Once you feel they match, you can press the right arrow key, after which a new pair of sensations will be presented.'

The outline of a trial was identical to that of the previous experiments, cf. Tables 2 and 3. The reference stimulus was tactile-only, and the comparison stimulus was either tactile-only (in C), a combination of tactile and visual (in V), a combination of tactile and audio (in A), or a combination of tactile, visual and audio (in VA), where the levels of visual and auditory stimuli corresponded to precipitation intensities: snow, rain, and hail. At the start of the trial, a beep occurred and a curtain went down to block the view of the hand and arm. The reference stimulus was presented, which the participants could not see but could feel at intensity level 7, and stopped after two seconds. Then, for C and A the curtains stayed down and for V and VA the curtains went up, a beep occurred, and the comparison stimulus with possibly different intensity was presented, also for two seconds. After a final beep, the participant could either make an adjustment to the intensity of the comparison stimulus, repeat the last trial, or accept the intensities as equal. In the first two cases the same possibly adjusted pair was presented again, and in the last case a new pair was presented.

During the two seconds of stimulus presentation, either one, two or three tactile sensations were presented, each for 100 ms, with an interval of at least 12 ms between the end of one stimulus and the beginning of the next to avoid possible crashes by the tactile display. The virtual left arm consisted of three collision sections: the hand, the lower forearm, and the upper forearm. Corresponding to these sections, the sensations were mostly stimulated on the center of back of the hand, sometimes on the center of the lower forearm, and rarely on the center of the upper forearm.

4.6 Results

The data was handled as in the previous experiments (see section 3.4.3 for a description), which gave nine normalized means per participant; the results are illustrated in Figure 8, and the means are given in Table 8. Ten Shapiro-Wilk tests showed that only the results of the control condition were not normally distributed ($W(16) = 0.861; p = 0.020$). A two-way repeated measures ANOVA with factors precipitation intensity (three levels) and modality condition (three levels) was run. This showed no significant main effect over modality condition ($F(2, 30) = 1.632; p = 0.212$), no significant main effect over precipitation intensity ($F(2, 30) = 1.938; p = 0.162$), and no significant interaction effect ($F(4, 60) = 2.364; p = 0.063$).

Since the observed $p = 0.063$ for the interaction effect is close to the desired $p = 0.050$, the simple main effects are investigated; the corresponding p-values of the pairwise comparisons are listed in Tables 9 and 10; the first compares modality conditions for each precipitation intensity, and the second compares precipitation intensities for each modality condition. In these tables we see that three significant dif-

	C	V	A	VA
Snow	-0.148	-0.391	-0.382	-0.367
Rain		-0.094	-0.453	-0.227
Hail		-0.211	-0.242	-0.383

Table 8: Normalized means over all modality conditions and precipitation intensities.

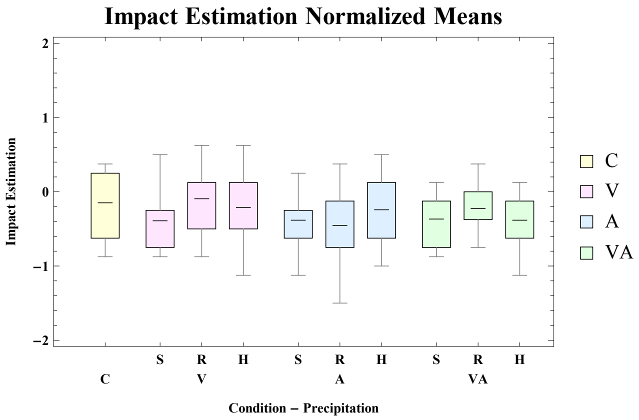

Figure 8: Boxplots of the normalized results of the precipitation impact experiment, grouped by modality condition; *C*, *V*, *A*, *VA* refer to the modality conditions, and *S*, *R*, *H* to the precipitation intensities.

ferences were detected, namely for *Rain* between *V* and *A*, *V* between *Snow* and *Rain*, and *A* between *Rain* and *Hail*.

Nine t-tests revealed that cases *A-Rain* and *VA-Hail* differed significantly from condition *C* ($t(15) = 2.298; p = 0.036$ and $t(15) = 2.570; p = 0.021$, respectively), and *V-Snow* and *VA-Snow* differed nearly significantly from *C* ($t(15) = 1.873, p = 0.081$ and $t(15) = 1.840; p = 0.086$, respectively).

Three independent-samples t-tests were performed to compare the normalized results of the control conditions of experiment 1 to that of this experiment; the results of all controls (combining those of the weight and temperature experiment) are illustrated in Figure 9. Note that although the exact circumstances between the experiments were not identical, the underlying test was. There were near significant differences between tactile reference intensities 5 and 7 ($t(36) = 1.839; p = 0.074$), and between 11 and 7 ($t(36) = -1.998; p = 0.053$), but not 8 and 7 ($t(36) = 0.254; p = 0.801$).

Qualitative results of the post-experiment questions are summarized in Table 11 and Figure 10. There are few points of interest here. Firstly, a fairly strong body transfer illusion occurred compared to the weight-experiment. Secondly, the convincingness of the stimuli for each modality were moderate and did not differ significantly according to a one-way ANOVA ($F(2, 30) = 0.638; p = 0.535$). Thirdly, the participants felt to a certain extent that objects were actually falling on their arm, and agreed that these felt more real when visual and auditory stimuli were added. Lastly, rather than only using the tactile stimuli to execute the matching task, some participants also used the visual and auditory stimuli; a t-test showed the difference between the answers regarding the visual and auditory stimuli was not significant ($t(15) = 1.321; p = 0.206$).

	Snow	Rain	Hail	AllPrec
V-A	0.943	0.007*	0.776	0.132
V-VA	0.853	0.217	0.182	0.102
A-VA	0.880	0.065	0.325	0.686

Table 9: P-values for all pairwise comparisons of the normalized results between modality conditions for each precipitation intensities.

	V	A	VA	AllCond
Snow-Rain	0.045*	0.557	0.153	0.115
Snow-Hail	0.156	0.193	0.903	0.130
Rain-Hail	0.354	0.043*	0.068	0.743

Table 10: P-values for all pairwise comparisons of the normalized results between precipitation intensities for each modality condition.

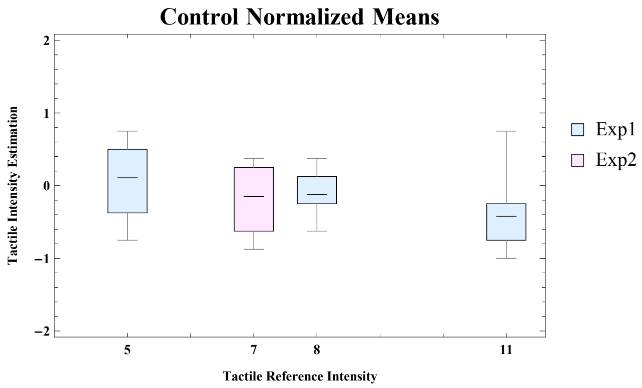

Figure 9: Boxplots of the normalized results of the control conditions of experiments 1 and 2; *5*, *7*, *8*, *11* refer to the tactile reference intensity.

Question	Mean	St. Dev.
1. Virtual arm your own	5.1	1.48
2. Convincing visuals	4.3	1.35
3. Convincing audio	4.8	1.56
4. Convincing tactile	4.3	1.30
5. Things actually falling	5.0	1.63
6. Things more real with V&A	5.3	1.85
	V 3.2	1.97
7. Deduce answer	A 2.6	1.46
	T 6.6	0.72

Table 11: Answers to the post-impact-experiment questions on a 7-point Likert scale.

4.7 Discussion

After demonstrating the potential of vibrotactile feedback in the previous experiment for the experience of weight, this potential was further investigated for the experience of precipitation impact. Now, the cues were intentionally chosen to be related to a realistic case, that is, extracted through

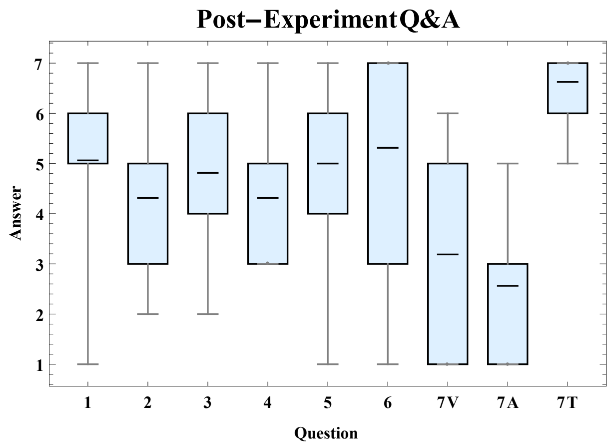

Figure 10: Boxplots of the answers to the post-impact-experiment questions on a 7-point Likert scale.

real world knowledge rather than type-related for example through exponentially increasing size and speed. Also, both visual and auditory cues were used, in order to more extensively investigate a multimodal effect. It was expected that in this case a similar effect would occur over precipitation intensity as in the weight experiment.

By evaluation of the t-tests of the normalized results, and examination of Figure 8 and Table 8, it is likely that an effect has occurred as a consequence of multimodality: when visual and/or auditory feedback were added to the scene, overall the tactile stimuli were perceived as stronger. Because of the significant differences found between various cases and condition C (which had no visual or auditory cues), it is reasonable that this is not only due to the unbalanced presentation of reference and comparison stimuli as described in section 3.6. This multimodal effect is further elaborated, followed by comparisons with the previous experiment concerning the post-experiment questions and control conditions.

The two-way repeated measures ANOVA showed no significant main effects, nor a significant interaction effect. Yet, since the observed $p = 0.063$ was very close to the desired $p = 0.05$, there is reason to believe that the effect of precipitation intensity on perceived tactile intensity may depend on the multimodal condition (Figure 8 and Table 10).

To examine this dependency, each condition is discussed separately. In each case, as suggested by the previous weight experiment, it was expected that stronger visual and/or auditory indication of precipitation impact, e.g. by hail, would lead to a weaker perceived tactile sensation. Likewise, a weaker indication, e.g. by snow, would lead to a stronger sensation; cf. section 3.6. With this, the results should have been strongest - weaker - weakest for the results of *Snow* - *Rain* - *Hail*, which would correspond to an increasing trend for each modality condition in Figure 8. Since there was no main effect over precipitation intensity, it is clear this did not occur. This would mean that in this experiment the steps in impact intensity as indicated by the visual and auditory stimuli were either not convincing enough or too fundamentally inconsistent. According to the post-experiment answers, the stimuli of each modality were moderately to highly compelling, thus fundamental inconsistency is further investigated.

In the condition V, there was only one significant difference in tactile perception between the precipitation intensities, specifically the tactile percept in *Snow* was stronger than that in *Rain*. These results can be partially explained by relating Figure 5 to the size-weight and speed-force illusions. Snow and rain particles are clearly visually different; the shape of a snow particle is round, while the rain particles look more like cylinders due to their malleable liquid form. This means that, indeed, a rain particle is larger in size than a snow particle. Hail particles were the same size as the snow particles; the difference between these two lies in the falling speed and collision effect: snow falls relatively slow, and melts/disappears on contact, while hail falls relatively fast and bounces off the collided surface. Therefore, the results suggest that size plays a more important role in the deduction of impact than speed. A follow-up study is necessary to confirm this conjecture.

For condition A, there was also only one significant differences in tactile perception between the precipitation intensities, specifically the tactile percept in *Hail* was weaker than that in *Rain*. Further investigation of the sound clips provides a possible explanation. Participants likely deduce characteristics of precipitation by the following sound characteristics: collision frequency, and collision type; the clips did not differ in any other major characteristics. Collision frequency corresponds to the density of the precipitation, which in this experiment was chosen such that the snow sound had zero collisions (since the sound caused by a snow particle collision is negligible), and the rain and hail sounds had fairly equal collision frequency. The type of collision, on the other hand, was differentiable between rain and hail: while the rain particles sounded like they were landing on water, the hail particles sounded like they were landing on solid objects. A possible explanation for the results therefore is that a difference in collision type is crucial to make a distinction between rain and hail, while either a difference in collision frequency between snow and rain/hail does not effect perception, or the absence of collision sounds for snow 'broke' the illusion of collision as a whole.

For condition VA, the absence of significant differences in tactile perception results between precipitation intensities is reasonable considering the inconsistent differences in V and A.

The qualitative results concerning the strength of the body transfer illusion in the weight experiment appeared higher than in the temperature experiment. Comparing the qualitative results gained in the impact experiment with these previous results, we see a seemingly higher result than before, cf. Table 5 and Figure 10. In section 3.6, the implication of a correlation between strength of body transfer illusion and congruence of stimuli was discussed and is further elaborated here with respect to the results of this experiment. The stronger body transfer illusion in this experiment compared to the previous experiments indicates that the addition of sound was crucial, and that replicating the weight and temperature experiments with the addition of sound would lead to different qualitative results. These findings, although so far only speculation, are similar to those of [7], where the sense of presence (which is clearly related to the strength of the body transfer illusion [25]) and memory increased by the number of presented modalities. Moreover, the variation of fidelity of the visual stimuli did not influence the sense of presence.

When comparing the responses to the question on how participants deduced their answer in the matching task (question 7), we see that the responses correspond more to those of the weight experiment than to those of the temperature experiment. Specifically, the primary deduction sensation was tactile, however other cues were not entirely excluded. This indicates that the visual and auditory cues affected the generated expectation of impact felt through the vibrations. Unfortunately, as discussed above, the stimuli were too inconsistent to find a main effect over precipitation intensity, thus this conjecture cannot be confirmed with this experiment.

Lastly, the results of the control condition in the impact experiment are consistent with those of the weight and temperature experiments. Although in the impact experiment a tactile reference intensity lower than 5 or higher than 11 were not tested, the results for tactile reference intensity 7 support the previous discussion in section 3.6.

5. CONCLUSION AND FUTURE WORK

The work presented in this paper was motivated by the goal to gain a better understanding of passive touch perception and to investigate if a simple vibration feedback has the potential to improve VR experience when used in combination with a richer visual and auditory stimuli. In a series of experiments, we have shown that visual cues that give an indication of weight, specifically size and falling speed, can change the perceived intensity of a vibration felt on the skin upon collision. Precisely, the sensation felt on the skin was estimated 'lighter' when the visual ball looked heavier. This means that future VR systems with the goal of creating different weight intensities would need to substantially exaggerate the weight by vibration in order for it to seem realistic for humans. This was not the case when using vibrations to simulate a temperature in combination with color as a visual cue. We conclude that it is not possible to let humans perceive different temperatures through vibrations and accompanying visual color cues; while other temperature-indicating cues may produce other results, our study suggests that such an investigation may not be worthwhile. In the case of impact, tactile perception depended on both presented modality and precipitation intensity, suggesting that visual and auditory cues related to the type of precipitation that were learned through real world knowledge were insufficient to evoke a tactile perception effect similar to the results of the weight experiment. This indicates that such an effect can only occur when the presented cues are fundamentally consistent, e.g. on exponentially growing type-related scales of size and speed, which is an important aspect for multimodal VR design.

The results of this research also identified two more general open problems. The first was of fundamental ground regarding general psychophysical experiments, namely what the consequences are of presenting reference and comparison stimuli successively and in the same order. Specifically, our results suggest that in the case of vibrotactile stimuli, the bias in judgement due to unbalanced presentation is not consistently in one direction as was previously thought. The second problem concerned body transfer illusion, and questioned whether the body transfer illusion was a consequence or a cause of accepting incongruent stimuli as plausible. This is an important practical observation, because both possibilities show that only one is needed in the design of a compelling multimodal experience.

This study made a first step in testing various aspects of passive tactile perception through multimodal integration, and the results point out some interesting and relevant follow-up research. An initial idea is to broaden the range of acceptable touch; this study was only able to add the concept of weight to this range using visual speed and size, but it is not misplaced to insist that this could happen for other properties and senses. Further, we remark that many findings in related studies, as in the weight experiment, only concern the influence of visual cues on tactile perception, and very few the influence of auditory cues, causing the relationship between the three to be less understood. This research made a step in that direction. The results are not conclusive but suggest that this is indeed an important direction for further evaluation; more research is necessary to properly combine these senses to create even better, more immersive experiences.

6. REFERENCES

[1] K. Arai and K. Okajima. Tactile force perception depends on the visual speed of the collision object. *Journal of Vision*, 9(11):19, 2009.

[2] K. W. Arthur, K. S. Booth, and C. Ware. Evaluating 3D task performance for fish tank virtual worlds. *ACM Transactions on Information Systems (TOIS)*, 11(3):239–265, 1993.

[3] F. Avanzini and P. Crosato. Integrating physically based sound models in a multimodal rendering architecture. *Computer Animation and Virtual Worlds*, 17:411–419, 2006.

[4] M. Botvinick and J. Cohen. Rubber hands 'feel' touch that eyes see. *Nature*, 391:756–756, 1998.

[5] D. Bowman, D. Johnson, and L. Hodges. Testbed evaluation of virtual environment interaction techniques. *Presence*, 10(1):75–95, 2001.

[6] G. Burdea and P. Coiffet. *Virtual Reality Technology*. John Wiley & Sons, Inc., Hoboken, New Jersey, second edition, 2003.

[7] H. Q. Dinh, N. Walker, L. F. Hodges, C. Song, and A. Kobayashi. Evaluating the importance of multi-sensory input on memory and the sense of presence in virtual environments. In *Virtual Reality, 1999. Proceedings, IEEE*, pages 222–228.

[8] R. R. Ellis and S. J. Lederman. The role of haptic versus visual volume cues in the size-weight illusion. *Perception & Psychophysics*, 53(3):315–324, 1993.

[9] G. A. Gescheider. *Psychophysics: the Fundamentals*. Lawrence Erlbaum Associated, Inc., 1997.

[10] S. Guest, C. Catmur, D. Lloyd, and C. Spence.

Audiotactile interactions in roughness perception. *Experimental Brain Research*, 146(2):161–171, 2002.

[11] E. Heineken and F. P. Schulte. Seeing size and feeling weight: The size-weight illusion in natural and virtual reality. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 49(1):136–144, 2007.

[12] Y. Hirano, A. Kimura, F. Shibata, and H. Tamura. Psychophysical influence of mixed-reality visual stimulation on sense of hardness. In *Virtual Reality Conference (VR), 2011 IEEE*, pages 51–54. IEEE, 2011.

[13] H.-N. Ho, D. Iwai, Y. Yoshikawa, J. Watanabe, and S. Nishida. Combining colour and temperature: A blue object is more likely to be judged as warm than a red object. *Scientific Reports*, 4(5527).

[14] H.-N. Ho, G. H. Van Doorn, T. Kawabe, J. Watanabe, and C. Spence. Colour-temperature correspondences: When reactions to thermal stimuli are influenced by colour. *PloS one*, 9(3):e91854, 2014.

[15] I. P. Howard and W. B. Templeton. Human spatial orientation. 1966.

[16] S. Kanaya, Y. Matsushima, and K. Yokosawa. Does seeing ice really feel cold? Visual-thermal interaction under an illusory body-ownership. *PloS one*, 7(11):e47293, 2012.

[17] S.-C. Kim, A. Israr, and I. Poupyrev. Tactile rendering of 3D features on touch surfaces. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology*, pages 531–538. ACM, 2013.

[18] R. W. Lindeman, Y. Yanagida, H. Noma, and K. Hosaka. Wearable vibrotactile systems for virtual contact and information display. *Virtual Reality*, 9:203–213, 2006.

[19] A. Maselli and M. Slater. The building blocks of the full body ownership illusion. *Frontiers in Human Neuroscience*, 7(83), 2013.

[20] H. McGurk and J. MacDonald. Hearing lips and seeing voices. 1976.

[21] R. Nordahl, A. Berrezag, S. Dimitrov, L. Turchet, V. Hayward, and S. Serafin. Preliminary experiment combining virtual reality haptic shoes and audio synthesis. In *Haptics: Generating and Perceiving Tangible Sensations*, pages 123–129. Springer, 2010.

[22] V. I. Petkova and H. H. Ehrsson. If I were you: perceptual illusion of body swapping. *PloS one*, 3(12):e3832, 2008.

[23] G. L. Poling, J. M. Weisenberger, and T. Kerwin. The role of multisensory feedback in haptic surface perception. In *11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2003. Proceedings*, pages 187–194. IEEE, 2003.

[24] P. Richard, G. Burdea, G. Birebent, D. Gomez, N. Langrana, and P. Coiffet. Effect of frame rate and force feedback on virtual objects manipulation. *Presence-Teleoperators and Virtual Environments (MIT Press)*, 15:95–108, 1996.

[25] M. V. Sanchez-Vives and M. Slater. From presence to consciousness through virtual reality. *Nature Reviews Neuroscience*, 6(4):332–339, 2005.

[26] L. Shams, Y. Kamitani, and S. Shimojo. Illusions: What you see is what you hear. *Nature*, 408(6814):788–788, 2000.

[27] M. Slater, B. Spanlang, M. V. Sanchez-Vives, and O. Blanke. First person experience of body transfer in virtual reality. *PloS one*, 5(5):e10564, 2010.

[28] A. Somada, A. Kimura, F. Shibata, H. Tamura, et al. Psychophysical influence on tactual impression by mixed-reality visual stimulation. In *Virtual Reality Conference, 2008. VR'08. IEEE*, pages 265–266. IEEE, 2008.

[29] M. A. Srinivasan and C. Basdogan. Haptics in virtual environments: Taxonomy, research status, and challenges. *Computers & Graphics*, 21(4):393–404, 1997.

[30] J. Steuer. Defining virtual reality: Dimensions determining telepresence. *Journal of Communication*, 42(4):73–93, 1992.

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

Chapter 2

Annotated Appendix

1 Introduction

Virtual reality (VR) has been defined in many different ways since its birth in the 1960s. One way to define it is hardware based, stating that virtual reality is a technological system through which virtual environments (VEs) can be sensed [11, 46]. Other definitions are based on the resulting human experience, stating that it is a medium through which the user experiences presence [53]. This latter definition considers two determining dimensions to the quality of experienced presence: *vividness* and *interactivity*. The first in turn has dimensions *breadth* (number of presented modalities) and *depth* (resolution of each presented modality), and the second *speed* (system response time), *range* (number of attributes with which interaction is possible), and *mapping* (how human actions are mapped to system actions); see Figure 1. This framework suggests that a highly *vivid*, *interactive* multimodal system allows a high sense of presence in a VE.

If we want to create such a multimodal VR system, we need to fully understand how humans experience multimodality in the real world. Unfortunately, this field of research is very complex and in part still unexplored. There is evidence for both plasticity (where the brain processes stimuli of one modality with neurons from a different modality) and interaction (where one modality alters the perception of another) among sensory modalities [49]. This latter aspect is explained by various theories on combination and integration strategies [12]. An example of such a theory is *visual capture*, stating that the visual modality dominates over other presented modalities when creating a percept. Another widely known theory is the *maximum likelihood estimation* model, where humans integrate signals such that the variance is lowest, thus reliability of the formed percept is highest. Lastly, *Bayesian integration* also focuses on increasing reliability, but through a form of Bayesian inference in sensorimotor learning.

Currently, most VR systems achieve multimodality only through visual and auditory system output; haptic, gustatory, and olfactory output are hardly supported. This may be due to the complexity of these modalities. For example, the haptic sense includes all sensations corresponding to movement (kinesthetic) and contact (tactile), and in turn tactile sensations consist of many different properties such as hardness, temperature, wetness, texture, and more. Furthermore, while visual stimuli are limited to perception by the eyes and auditory stimuli by the ears, tactile stimuli are perceived through skin over the entire body.

Focusing on the human experience, tactile feedback can be categorized into passive and active touch [16], which are related to the dimensions vividness and interactivity from [53], respectively. To create a high sense of vividness in a system including tactile feedback, one would need to simulate every single type of touch known to man. This is not feasible from a technical or practical perspective. Although simulation may not be possible, however, perception may indeed be achievable using convincing signals from other modalities; that is, it may be possible to reach a high degree of *vividness* with high *breadth* and low *depth*, for example by taking advantage of the visual modality, following the theory of *visual capture*. From a practical aspect, for this goal vibrotactile feedback devices are promising, since they are inexpensive, light-weight and flexible in arrangement.

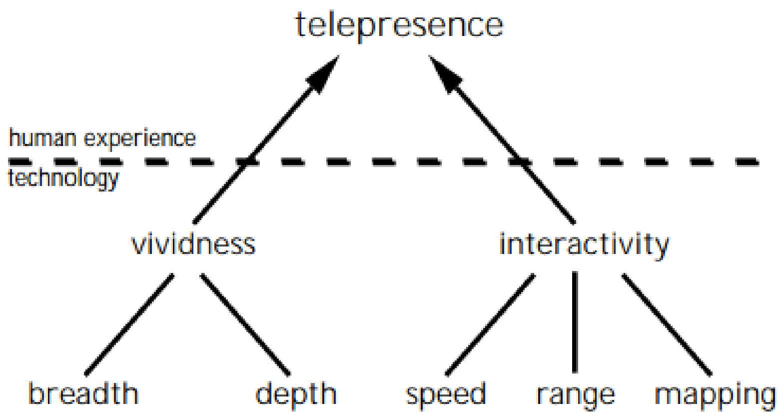

Figure 1: Determining dimensions of telepresence, i.e. presence through a medium; copied from [53].

This opportunity was the starting point for this thesis. Specifically, given the theoretical evidence that multimodality might lead to a better VR experience on the one hand, combined with the technologically feasible opportunities (but also limitations) provided by vibrotactile feedback devices on the other hand, the ultimate aim was to investigate related multimodal interaction effects in VR. Related insight about the perception of tactile sensations simulated by vibrotactile sensations combined with visual and auditory stimuli would allow the creation of richer, better, and even completely new VR experiences.

In order to identify the areas with the most promising potential, and to specify the concrete goal of this thesis, an initial literature study investigating various research opportunities in tactile VR was performed. This study is summarized in Section 2. As a result, the field of visuotactile illusions was identified as most promising and feasible for this thesis project. This lead to the investigation of the characteristics of weight, temperature, and impact for passive touch, and if and how they can be simulated using vibrotactile feedback in combination with visual and (partly) auditory stimuli.

The identified research questions were investigated in three experiments summarized in a scientific paper aimed for publication in a related journal. This annotated appendix contains additional information and background data relevant for the thesis, but not contained in the paper. Section 3 summarizes the results of the weight and temperature experiments (which in turn are the core of the submitted ICMI paper); Section 4 contains background information on the impact experiment and related results.

Section 5 introduces and describes a final demonstration program that was created to illustrate the achieved results and their relevance for a practical context. Section 6 summarizes the thesis results, and provides further comments on possible future work and research directions.

2 Research Opportunities in Tactile VR - Literature Study

This section encompasses the different fields investigated during the preparation phase of this thesis project, in order to formulate research questions. These specific fields were selected due to the potential benefits tactile VR might offer concerning performance and experience. Section 2.1 discusses different approaches to VR menu display and interaction, Section 2.2 describes the apparent motion phenomenon and opportunities through VR, and Section 2.3 demonstrates experiences that are possible through a body transfer illusion in VR. Section 2.4 states the topic that was finally chosen for the experimentation phase of the thesis. In the following, when discussing virtual reality, immersive virtual reality is implied, e.g. by means of a head-mounted display, and not non-immersive virtual reality, e.g. by means of a desktop. Note that the corpus used here is not exhaustive, as this research was only exploratory.

2.1 Menus in VR

To date, there is no standard for VR menu display or interaction. This issue is strongly related to general interaction techniques in VR, for which there are three factors contributing to its complexity: the opportunity to use non-standard input (e.g. gestures instead of keyboard input), the shift from object-oriented to noncommand-oriented interaction [39], and the addition of a new dimension to the virtual environment. This section discusses various approaches of menu systems that can be applied in VR; an extensive overview of menu approaches for desktop VR, immersive VR and augmented reality can be found in [10], together with various classification systems.

Jacoby and Ellis [23] laid a groundwork for this research field in 1992. By investigating characteristics of both conventional menus and VEs, they provided a frame of reference for the design of virtual menus. In [6], the question of whether common 2D graphical user interfaces (GUIs) are appropriate for VEs is addressed, at the same time disagreeing with the mainstream view that all VE interaction should be ‘natural’ [39]: system control interfaces must ensure high usability and performance, and therefore may be application specific. Five requirements of a menu system using pinch gloves are formulated (note that they are applicable to general VR menus as well):

- the new system needs to be at least as efficient and precise as other menu types
- its use should not cause the user significant discomfort
- it should not occlude the environment
- the menu system should be appropriate for both novice and expert users
- expert users should be able to do eyes-off interaction with the menu

Three menu systems were evaluated: a floating menu, a pen and tablet menu, and the TULIP menu, where the latter was a novel implementation of a body-centered menu using pinch gloves. The evaluation of each method showed that the ‘simple’ pen and tablet method was faster, however user preference was highest and discomfort lowest for TULIP.

An approach that was common in the years that followed was using wrist orientation [15, 9]; this switch may have been caused by the trend towards gestural commands after the release of the Nintendo Wii and Microsoft Kinect. An example of such an approach is the rapMenu: a menu system using approximate wrist orientation and pinching on a circular menu [38]. This system was evaluated in [37] through a comparative analysis with a system not using pinch, thus completely relying on precise tilting, where the rapMenu outperformed the tilt menu in many aspects, and even eyes-free selection was shown to be promising. A more extensive study was conducted in [13], where the usability of circular menus with a 6 degree-of-freedom controller was investigated by testing various selection methods, hierarchical layouts, abort methods and dead zone (center of menu) sizes. Results showed that performance and preference for selection through ray-casting was in fact higher than for wrist orientation. Moreover, hierarchy through depth-overlay was fastest to use, and aborting by dead zone was preferred; the optimal size of this zone was not specified. A pilot study was performed to extend functionality of these menus by elements that are common for two-dimensional GUIs, such as sliders and buttons. The results of an expert review indicated high performance and efficient design.

Approaches that step away from this orientation-based selection are presented in [32]. A finger-count menu selection approach, using enumerated menu items, is compared against hand-n-hold selection, thumbs-up selection and 3D marking menus, which all rely on the hand to operate as a 2D cursor. The finger-count method resulted in the lowest completion times and highest user preference, and moderate accuracy. Although not tested in immersive VR, this still proves to be a viable option for 3D menu selection tasks. An example of a completely different approach that does not use gesture input is the earPod [56]: an eyes-free menu using touch input on a device with a touch pad and audio feedback. The touch input works in the same way as an iPod-like list menu, and a comparative evaluation between these two menus was performed. This comparison showed that overall there was no difference in performance between the visual and audio menu systems, however there was a high learning curve for the audio menu. The preferences of the users were spread. Although not actually designed for a VE, it has great potential as a VR menu.

Conclusion

A point that stands out is the lack of tactile feedback in the discussed systems, an aspect that is typical for interaction in non-immersive VEs such as desktops and smartphones. A research question that arises is: *Can menu performance be increased through vibrotactile feedback during item selection?* The reason this research topic was not chosen was that, although each of these discussed menu systems gives an indication of where potential lies, the findings may be indeed too application specific, making it unreasonable to formulate concrete recommendations concerning general menu design and interaction.

2.2 Apparent Motion

Tactile sensations can be simulated using tactile feedback devices. Currently, the delivered sensations are discrete: they have a certain fixed type, duration, intensity, onset and location. Using these however, we are able to create phantom sensations with other characteristics. In particular, using two vibrotactile stimuli with distinct locations, we are able to simulate a sensation somewhere in between these locations. There are two main methods to accomplish this: funneling and saltation. The first method uses difference in intensities to determine the final location, whereas the second uses variation in inter-stimulus onset interval. These methods have been shown to not be limited to phantom locations that are on the skin: it is possible to elicit exosomatic sensations [3, 35, 26].

Through funneling and saltation, we are able to create sensations where there are no actuators, thus making it possible to produce apparent motion, i.e. a moving sensation on the skin. Research into the the requirements and perception of this phenomenon, for both vibrations and other tactile stimuli, has been ongoing for over a century. In these studies it was common to use saltation to induce the illusion, which caused participants to feel 'hopping' sensations on the skin, after which the phenomenon became known as the cutaneous rabbit effect [14]. The main contributing factors to the illusion were found to be inter-stimulus onset interval and stimulus duration [48, 27], number of actuators [4, 28, 29], and shape of motion [30].

In [31], the application of saltation for information display, namely different types of messages with varying urgency, was investigated. Recognition of pattern and speed were tested using three rings, each with five tactors, around the upper arm. Pattern recognition rate was nearly 100%, however absolute speed recognition was around 80%. Since the goal was to ultimately map these patterns to abstract concepts such as 'urgent personal call' vs. 'work call', the application would have an extra level of complexity, thus either more training would be necessary, or different saltation parameters would be needed. In [40], this application was further tested; the objectives were to find the shortest time interval with which humans can differentiate between single directional motion and back-and-forth motion along the length of the upper arm, and the minimum number of tactors necessary to differentiate between circular directions around the upper arm. The results showed that a time interval of 400ms was appropriate, where this interval can be smaller if participants correctly sense the starting point of the sequence, and four tactors were sufficient to display circular motion properly.

The funneling illusion was investigated in a pilot study in [7], where the goal was to define optimum values for velocity and distance of the perceived motion such that the motion sensation was perceived as continuous, specifically to the dorsal forearm. The results showed that the highest subjective scores were for a distance of approximately 60mm and a velocity of approximately 60mm/s. This project was continued in [45], where the influence of intensity variation (linear/logarithmic), body site, limb axis, sensation duration and gender on the quality of the perceived continuous sensation was tested. It was shown that linear intensity variation was preferred, optimal stimulus duration was 110ms, males preferred logarithmic intensity variation while females preferred linear, and linear intensity variation is favoured for the longitudinal orientation. Lastly, whereas the previous two studies focused on the quality of the illusion, [2] focused on spatial resolution on the forearm, where participants experienced a motion illusion and had to determine the starting and ending points of the stroke. The results showed that localization was higher near the real tactors than in between, and highest around the elbow.

For the application of a ‘tactile chair’ to enhance entertainment experiences, Disney Research Pittsburgh applied funneling using a 2D array of tactors placed on the backrest of a chair [20]. To accomplish this, the ‘Tactile Brush’ Algorithm was developed [22], where the control space for straight-line apparent motion was determined. This control space was further investigated in [21] in three experiments, where the stimulus onset asynchrony thresholds were determined for (1) various frequencies, intensities and durations of stimulation and body sites; (2) different motion directions and actuation point spacings; (3) single axis (four tactors) and two-dimensional apparent motion (four-by-three tactors). The findings were combined in the design of the ‘Surround Haptics’ system [19]. Although this system focuses on enhancing the experience of so far only visual and auditory stimuli, the combination of the continuous sensations with the corresponding visual and auditory stimuli is not tested. A similar aspect was tested in [8], where crossmodal cueing instead of intramodal cueing proved to shift the perception of Ternus apparent motion from element motion to group motion.

Conclusion

As highlighted above, more research is necessary into the multimodal integration effects on perception of apparent motion, resulting in the research question: *Is it possible to change properties such as quality, direction and intensity of apparent motion by multimodal stimuli?* Pilot tests with the Elitac tactile display showed that the requirements for apparent motion using vibrotactile stimuli could not be met without frequent system crashes, thus this topic was finally not chosen.

2.3 Body Ownership

When sensory stimulation in VR correlates correctly with the real world experience, a body transfer illusion can occur: the illusion of owning (part of) a body that is not one’s own. A well known example of this illusion is the rubber hand illusion [5], where a rubber hand ‘replaces’ the participant’s own hand, through believable hand position (visually and proprioceptively) and synchronous tactile stimulation. While originally discovered in the real world, this illusion was proven to occur in VR with similar methodology [18, 50]. Later it was also shown that visuomotor correlation is sufficient to evoke the illusion in VR [47]. Similar findings were made for body transfer illusions [44, 51, 36].

Initially it was common practice to make the fake body look as real as possible, in order to evoke a stronger illusion. However, this illusion can occur when the bodies have very clear differences; in [43], it was shown that gender and differences in the precise shape of the body were not important factors for the body transfer illusion, although the fake body did still need to resemble an operable body: objects such as a rectangular box did not evoke an illusion. Other studies have shown that humans accept larger belly sizes [41] and longer arms as their own body parts [25]. When using such a different body, behavioral and operational phenomena can occur. A behavioral consequence is the so called Proteus effect [55]. This effect indicates that an individual’s behavior can change according to their self-representation. This was investigated for different appearances, such as attractiveness and body height [55], age [1], and skin color [24, 42].

An example of an operational consequence is the concept of homuncular flexibility: the ability of users to learn to control different bodies by changing the relationship between tracked and rendered motion [33]. Note that VR offers the opportunity to test these scenarios, which is otherwise hardly possible in the real world. Steptoe et al. investigated the capabilities of humans with a tail-like appendage, which was either controlled by lateral hip movement or moved at random [52]. To examine their actions, a game was designed in which users had to block beams that where in such a formation that each beam could logically be blocked by either the hands, feet or tail. After the game, a threat to first the tail and then the body would occur to measure the extent of body ownership through anxiety responses. An outcome was that controlled movement of the tail had a positive influence on the ownership of the tail, resulting in anxiety towards the threat on the tail. Furthermore, performance of the tail in the game for the controlled movement condition was over 30% higher than for the random movement condition. Similarly, in [54] remapping of movements was tested in two experiments: alteration of and addition to the body schema. In the first, a balloon-popping task was used to show that performance in a condition with altered arm and leg movement strength was just as high as in a normal movement scenario. In the second experiment, a long third arm originated at the center of the chest was added to the avatar, and could be controlled by rotation of both real wrists. In a block-selecting task, where the long arm could reach the farthest blocks without the participant having to step forward, long-armed participants completed more trials than normally-armed participants.

Conclusion

The possibilities for projected bodies seems endless. The question remains how far we can take this deformation: at what point does the body transfer illusion no longer occur, and at what point are we no longer able to control other bodies? Relating this to tactile feedback, a research question would be: *Can more radical deformations be achieved when visuotactile stimulation is applied on top of visuomotor congruence?* This topic was not chosen due to the initially seemingly small computer scientific relevance.

2.4 Visuotactile Illusions in VR

The topic that was finally chosen over the other three was visuotactile illusions in VR. This topic had been introduced in a previous research project (Small Project, Game and Media Technology, INFOMSPGMT), which focused on spatial localization on the skin through multimodal integration in VR. This thesis took on a new direction with completely different properties of tactile sensations, namely weight, temperature and impact. Practical aspects such as availability of resources and expected feasibility given the existing time frame had been considered in this decision as well. For the literature study of this topic, the reader is referred to the accompanying scientific papers of this project.

3 Experiment 1: Abstract Weight and Temperature

3.1 Motivation

During design of this experiment, it was first contemplated which tactile properties seemed feasible. The overview on 'knowledge about objects' defined in [34] (see Table 1) provided a starting point for choosing properties. Since this project focuses on passive tactile sensations, the chosen properties were temperature and weight.

Property type	Knowledge about object	Exploratory procedure	Active/Passive
Substances	Texture	Lateral motion	A
	Hardness	Pressure	A
	Temperature	Static contact	P
	Weight	Unsupported holding	P
Structure	Weight	Unsupported holding	P
	Volume	Enclosure, Contour following	A
	Global shape	Enclosure	A
	Exact shape	Contour following	A
Functional	Part motion	Part motion test	A
	Specific function	Function test	A

Table 1: Postulated links between knowledge about objects and exploratory procedures; first three columns copied from [34].

3.2 Implementation

3.2.1 Hardware

The VR hardware used during this experiment consisted of an Oculus Rift Development Kit 2 and an Elitac tactile display. Initially, a *Leap Motion Controller* was used to incorporate motion tracking to evoke a strong body transfer illusion. This was implemented using the Unity package available on the Leap Motion website. Although the program was decent at tracking general motions, jitter occurred when the user stayed still, which was a large part of the experiment. Motion would moreover encourage active exploration of properties that are not stimulated in this experiment. Furthermore, the provided arm assets resulted in two floating lower arms without the rest of the body, in turn diminishing the illusion. Although it may have been possible to resolve the latter issue, the first issue was hardware dependent. Thus in the interest of time and since it was not the focus of the experiment, it was decided to not incorporate motion tracking to the program.

The Elitac tactile display consisted of one pancake motor that was placed on the center of the forearm and a control module. This module communicated with the program through Bluetooth; this wireless set up was chosen over a USB-cable connection to limit cumber. This however did cause a lag of approximately 50 ms (depending on the quality of the Bluetooth connection), while a USB connection only caused a 2 ms lag. This was so small that it went unnoticed during the experiment. Using a user diagram protocol (UDP) client and the Bluetooth connection, commands in the form of messages were sent to the module over the UDP channel. In such a message, the tactor identification number, intensity, duration and onset must be specified.

The Oculus Rift SDK and Runtime for Windows, both version 0.4.4, and Unity 4 Integration were used. The Oculus was used in Extended Mode, where the experiment was running on the HMD, and the scene and console on the other screen, thus the events could be followed and controlled by the experimenter.

3.2.2 Software

The experiment environment was created in Unity, an engine to create 2D and 3D games. The C# scripts were written using Visual Studio 2013; the UML diagram of the written scripts is shown in Figure 2. The program constantly checks for key input: the first key input specifies whether a training or experiment session is to start, and whether weight, temperature or the control condition is being tested. Once this is specified the total list of rounds was constructed, a random round is removed and presented. This presentation consists of eight phases: block view, beep, beep, tactile stimulus, unblock view, beep and sphere appears, beep and sphere falls, tactile stimulus. The first four phases belong to the reference stimulus and the last four to the comparison stimulus. In the case of the control condition, the comparison was identical to the reference stimulus, thus the first four phases occurred twice. Directly after such a presentation, the program awaits key input indicating the necessary adjustment, requesting a repetition or agreeing the weights/temperatures were equal, after which a new round was presented. After 24 trials the program would indicate it was time for a break.

Besides the reported answers of the participants, two other aspects were recorded: the trail of adjustment and whether the participants hit one of the two extremes on the range of possible tactile intensities. Both pieces of information were eventually not used for analysis, thus further experiments would not include this information.

In case of errors and other events causing the experiment to stop temporarily, a crash-reader was implemented. In the case of a crash, the crash-reader would read the existing results of the participant, remove the already executed rounds from the total list of rounds, and continue with a new round. Although the possibility of crashes had been foreseen, still two participant's results were discarded due to crashes. In the case of the first, the crash reader still had a few bugs that were detected during this crash. These bugs were later fixed and the crash-reader worked as expected. In the case of the second, the computer froze and a restart for the computer, Unity and the tactile display would require too long and the participant had to leave during this time. Discarding these results were not crucial because of the moderate total number of participants.

The idea for specifying the training session was added later on in the programming process, however could no longer easily be integrated into the system. Therefore, many parts of the program contain duplicate code. In the same way, since the control-rounds differed from the other rounds and was added later on, this also resulted in duplicate code. For example, although not shown in Figure 2, there were four *RoundsList* instances, four *allRounds* lists, and four *currentRounds*. This disorganization was a point of improvement for the second experiment.

Statistical analysis was performed in IBM SPSS Statistics 21. Results were preprocessed in two programs, one that regroups the data written in C# in Visual Studio 2013, and another written in Wolfram Mathematica 9 that processes the data such that it is readable by SPSS.

3.3 Issues

Technical issues occurred that were not foreseen during implementation. Firstly, there was one position of the head that caused the body to disappear. While this was avoided by asking the participants to simply not place their head in that position, it was still bothersome in some cases. It was thought that this was due to the distances of the clipping fields, however it was not clear how to change these when this issue was discovered. Secondly, since the user had no further interaction with the program, many thought the experience was boring. This may have also been related to their expectations to new technology. Both issues were points of improvement for the second experiment.

3.4 Data

The results for the weight and temperature experiments are discussed in the scientific paper. Here, a more detailed, complete overview of the concrete results is provided: for each participant the mean and standard deviation of the eight recorded values are given for each tactile reference intensity and each visual intensity. The results for the weight experiment can be found in Table 2 on page 26, and for the temperature experiment in Table 3 on page 27.

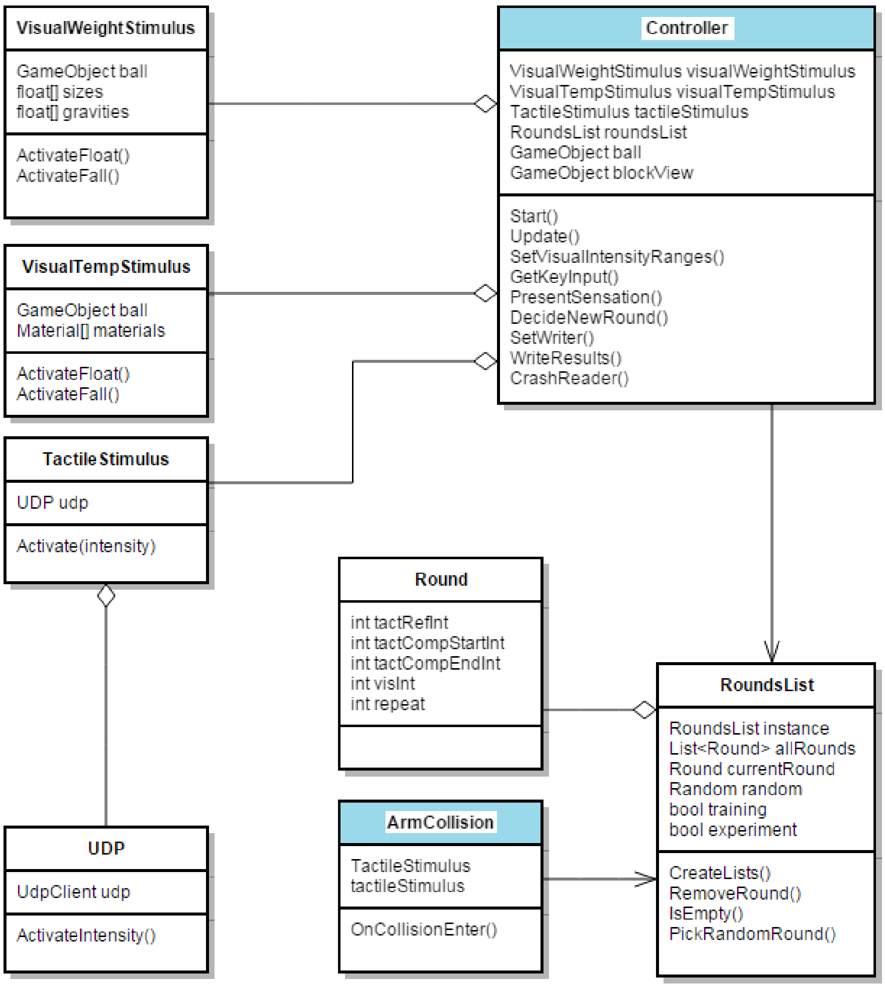

Figure 2: Simplified UML diagrams of the program for experiments 1A and 1B. Blue indicates that the class derives from MonoBehaviour in Unity.

	5		8		11	
	M	SD	M	SD	M	SD
C	5.000	1.069	7.875	1.126	10.250	0.886
0	5.000	1.195	7.375	1.061	10.000	1.690
1	5.125	1.126	7.875	1.959	10.000	1.512
2	5.250	1.165	7.750	1.488	10.500	1.773

(a) Participant 1

	5		8		11	
	M	SD	M	SD	M	SD
C	5.5	1.195	7.750	1.282	10.375	1.061
0	4.625	0.916	7.500	1.512	9.750	0.886
1	4.875	1.246	7.500	1.069	10.125	0.991
2	5.625	1.506	7.625	0.916	9.875	0.991

(b) Participant 2

	5		8		11	
	M	SD	M	SD	M	SD
C	5.5	0.926	7.875	0.991	10.125	1.458
0	6.625	0.916	8.750	1.035	10.500	1.195
1	6.875	1.246	8.5	0.926	11.125	0.641
2	6.625	0.916	9.5	0.926	11.125	0.991

(c) Participant 3

	5		8		11	
	M	SD	M	SD	M	SD
C	5.750	1.389	8.375	1.188	11.75	1.035
0	4.750	1.035	7.500	1.309	10.375	1.302
1	5.000	1.512	8.000	1.309	10.75	1.669
2	5.375	1.302	7.875	1.642	10.625	1.768

(d) Participant 4

	5		8		11	
	M	SD	M	SD	M	SD
C	5.500	0.535	7.500	0.535	10.000	0.926
0	5.125	0.641	8.375	0.916	10.375	1.302
1	5.750	0.707	8.250	1.035	10.625	0.916
2	5.500	1.414	8.000	0.535	10.875	0.835

(e) Participant 5

	5		8		11	
	M	SD	M	SD	M	SD
C	5.250	0.707	7.750	1.165	10.875	1.356
0	5.375	1.188	7.750	1.282	10.000	1.195
1	5.000	1.195	8.000	1.069	10.375	1.188
2	5.250	1.389	8.500	1.309	10.750	1.282

(f) Participant 6

	5		8		11	
	M	SD	M	SD	M	SD
C	5.375	1.768	8.375	1.598	10.875	1.808
0	5.000	1.414	7.625	1.768	10.375	1.506
1	5.250	1.909	7.875	1.553	10.125	1.246
2	5.625	1.302	7.875	1.808	10.625	1.598

(g) Participant 7

	5		8		11	
	M	SD	M	SD	M	SD
C	4.250	0.707	7.750	0.886	10.375	1.685
0	5.000	1.309	7.625	1.061	11.000	1.309
1	4.625	1.188	7.875	0.991	10.500	1.309
2	5.250	1.753	8.750	1.035	11.375	1.302

(h) Participant 8

	5		8		11	
	M	SD	M	SD	M	SD
C	5.375	1.302	8.125	1.553	11.125	2.03
0	5.125	1.356	7.875	1.727	10.125	1.246
1	5.125	1.356	7.625	1.768	9.875	1.126
2	5.375	2.066	7.875	1.458	10.625	1.768

(i) Participant 9

	5		8		11	
	M	SD	M	SD	M	SD
C	5.500	0.926	7.875	0.991	10.625	1.598
0	5.875	1.356	8.500	1.414	11.125	1.642
1	5.375	0.916	8.000	1.414	10.750	1.389
2	5.375	0.744	8.625	0.916	11.125	1.356

(j) Participant 10

	5		8		11	
	M	SD	M	SD	M	SD
C	5.625	1.061	7.750	1.165	10.750	1.581
0	4.125	0.991	6.875	0.835	9.375	0.916
1	5.125	1.126	7.500	1.604	10.750	1.282
2	5.375	1.768	7.875	1.126	10.125	1.356

(k) Participant 11

	5		8		11	
	M	SD	M	SD	M	SD
C	4.875	0.991	7.875	1.356	10.250	1.282
0	5.500	1.195	7.875	0.991	10.500	1.195
1	5.125	1.885	8.125	1.808	10.500	1.414
2	5.500	1.414	7.750	1.282	10.375	1.302

(l) Participant 12

Table 2: Results of Experiment 1A (Weight). Each table represents the means (M) and standard deviations (SD) of a single participant: the rows represent the visual intensity (control, 0, 1, 2) and the columns the tactile reference intensity (5, 8, 11).

	5		8		11	
	M	SD	M	SD	M	SD
C	4.625	0.744	7.375	1.188	10.625	1.598
0	4.375	1.506	7.625	0.916	11.375	1.847
1	4.625	1.744	7.500	1.069	11.000	1.309
2	4.500	1.069	7.625	1.061	11.375	1.685

(a) Participant 1

	5		8		11	
	M	SD	M	SD	M	SD
C	5.375	1.408	7.875	0.641	10.625	1.188
0	4.625	1.302	7.250	0.886	10.125	1.356
1	4.875	1.246	7.500	1.195	10.125	1.356
2	4.500	1.309	7.375	1.061	10.000	1.069

(b) Participant 2

	5		8		11	
	M	SD	M	SD	M	SD
C	5.250	1.389	8.125	1.356	11.000	1.512
0	6.125	0.835	7.875	1.553	10.375	1.598
1	6.000	0.926	8.25	1.165	11.000	1.309
2	5.125	1.356	8.625	1.061	10.625	0.744

(c) Participant 3

	5		8		11	
	M	SD	M	SD	M	SD
C	5.375	1.847	8.000	1.309	10.250	1.065
0	4.750	1.982	9.000	3.117	11.75	3.012
1	4.625	2.560	7.750	2.188	10.25	1.909
2	4.125	2.295	7.625	1.996	10.375	1.996

(d) Participant 4

	5		8		11	
	M	SD	M	SD	M	SD
C	4.500	1.309	8.125	1.959	10.75	1.753
0	5.875	1.356	8.25	1.488	10.75	0.886
1	5.500	1.690	8.625	1.061	10.75	1.753
2	5.250	1.389	8.250	1.282	10.500	1.195

(e) Participant 5

	5		8		11	
	M	SD	M	SD	M	SD
C	4.625	1.506	7.875	0.641	10.375	0.744
0	4.750	1.581	7.375	1.061	10.250	1.035
1	5.250	1.282	7.125	1.246	9.875	1.126
2	4.875	1.642	7.500	1.069	10.500	1.195

(f) Participant 6

	5		8		11	
	M	SD	M	SD	M	SD
C	4.625	1.188	8.000	0.756	10.375	1.506
0	5.375	1.506	7.625	1.061	10.750	1.389
1	5.250	1.389	7.500	1.195	10.500	1.309
2	5.125	0.991	7.500	1.195	10.250	1.488

(g) Participant 7

	5		8		11	
	M	SD	M	SD	M	SD
C	4.625	0.744	7.375	1.188	10.625	1.598
0	4.375	1.506	7.625	0.916	11.375	1.847
1	4.625	0.744	7.500	1.069	11.000	1.310
2	4.500	1.069	7.625	1.061	11.375	1.685

(h) Participant 8

	5		8		11	
	M	SD	M	SD	M	SD
C	5.000	1.195	7.500	1.309	10.000	1.195
0	4.500	1.414	7.125	0.991	10.375	1.302
1	4.625	1.188	7.500	1.195	10.375	1.685
2	4.500	1.414	8.125	1.125	10.750	1.282

(i) Participant 9

	5		8		11	
	M	SD	M	SD	M	SD
C	4.875	0.641	8.250	1.035	10.750	0.886
0	6.000	1.195	8.125	0.991	10.375	1.685
1	5.125	0.991	7.375	0.916	10.375	1.506
2	5.625	1.506	7.875	0.991	10.000	1.195

(j) Participant 10

Table 3: Results of Experiment 1B (Temperature). Each table represents the means (M) and standard deviations (SD) of a single participant: the rows represent the visual intensity (control, 0, 1, 2) and the columns the tactile reference intensity (5, 8, 11).

4 Experiment 2: Precipitation Impact

4.1 Motivation

In the previous experiment, correlations were used that have been fundamentally tested within cognitive science, namely size and speed with weight and color with temperature. However, in real life the differences between objects may not be so 'exponentially' obvious as tested here. Therefore, the goal for the next experiment was to test whether similar results would occur when not using these known correlations, but correlations humans make through experiences in realistic events. Also, rather than weight, terminology was changed to impact, focusing more on the sensation directly related to the collision and not the underlying reason. To test various levels of impact intensity, the goal was to find a single object that could vary in impact. For this reason precipitation was chosen in the form of snow, rain and hail.

4.2 Implementation

The implementation of this program was very similar to that of the previous experiment. The tactile display was reconfigured to have three tactors rather than one, and different wireless headphones were used to relieve the user of a cable connection. The main differences in the software were the addition of *AudioStimulus*, the change in functioning of *VisualStimulus* through particle systems, and the simplification of training, experimental and control sessions, resulting in the removal of a large amount of duplicate code. See Figure 3 for the UML diagram.

4.3 Issues

As described in section 3.3, the solution to the disappearing arm was thought to lie in the adjustment of the clipping fields. The values of these were changed, however, this adaption did not fix the problem. Again, the workaround was to ask participants to not sit in this position. Another issue that occurred, but had been foreseen, was that the position of the visual collisions did not exactly match the position of the tactile collision. The positions of the visual collisions was left random to keep the the precipitation as realistic as possible, however because of this not all positions could be tactually stimulated. This is not a problem for the forearm due to its low tactile resolution, however tactile resolution on the back of the hand is higher than on the forearm, and even more so on the fingers [17] (Ch. 14). Rather than stimulating the whole hand, it was chosen to only stimulate the center of the back of the hand. This was definitely noticed by the participants, as was clear from the post-experiment remarks. Some also remarked that this decreased the strength of the body transfer illusion.

4.4 Data

The results for the impact experiments are discussed in the scientific journal paper. A complete overview of the concrete results is provided here: for each participant the mean and standard deviation of the eight recorded values are given for each precipitation intensity and each visual intensity. The data for the impact experiment can be found in Table 4 on page 30.

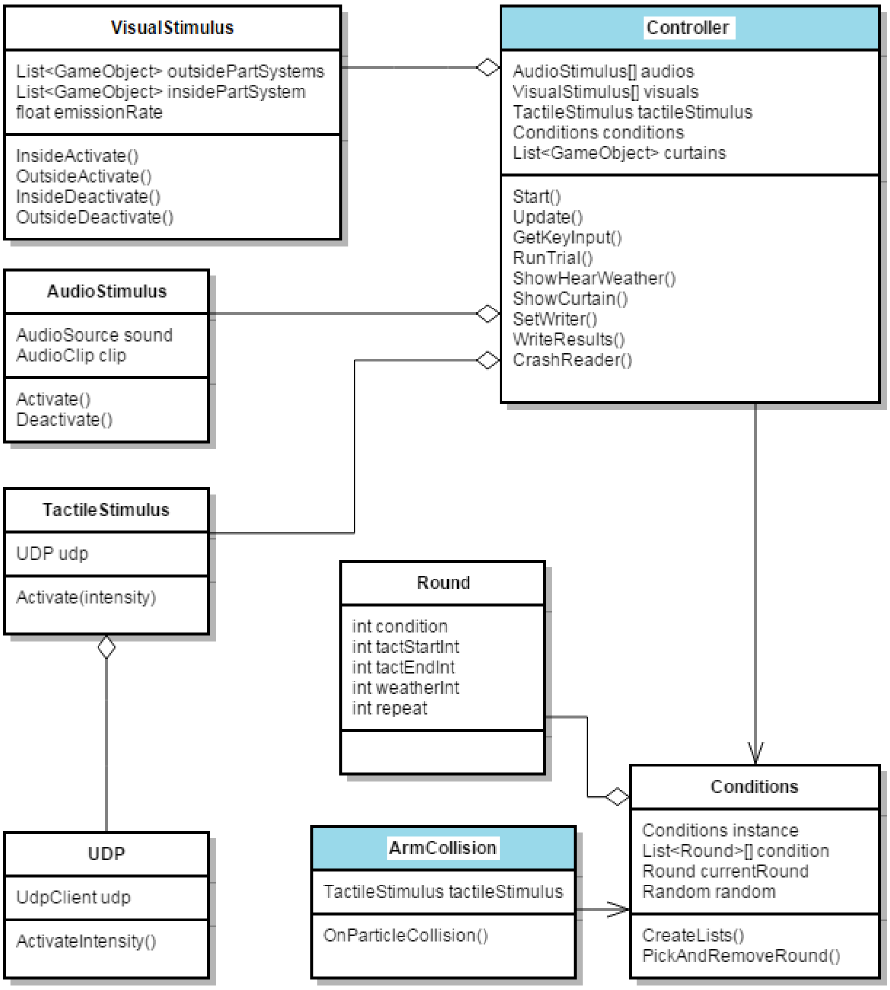

Figure 3: Simplified UML diagrams of the program for experiment 2. Blue indicates that the class derives from MonoBehaviour in Unity.

	C		V		A		VA	
	M	SD	M	SD	M	SD	M	SD
0	7.250	2.278	6.250	2.586	5.870	1.965	6.125	2.315
1			7.000	2.550	6.375	1.654	6.750	1.854
2			6.250	2.586	6.375	2.176	6.625	2.288

(a) Participant 1

	C		V		A		VA	
	M	SD	M	SD	M	SD	M	SD
0	6.750	0.661	7.125	1.536	6.625	1.932	6.625	2.578
1			6.875	1.833	6.250	2.278	7.000	1.658
2			6.875	1.833	6.750	1.984	6.875	2.147

(b) Participant 2

	C		V		A		VA	
	M	SD	M	SD	M	SD	M	SD
0	6.500	2.236	7.500	2.179	6.750	2.817	6.500	2.656
1			6.750	2.278	6.375	2.497	7.125	2.522
2			7.375	1.996	7.000	2.616	6.750	2.681

(c) Participant 3

	C		V		A		VA	
	M	SD	M	SD	M	SD	M	SD
0	7.250	2.537	6.500	2.784	6.750	2.586	6.500	2.121
1			6.875	2.666	6.875	2.713	7.000	2.121
2			6.500	2.784	6.625	2.546	6.875	2.976

(d) Participant 4

	C		V		A		VA	
	M	SD	M	SD	M	SD	M	SD
0	7.250	2.385	6.750	2.107	7.00	2.500	6.875	2.713
1			7.125	2.260	6.625	2.233	6.500	1.871
2			7.250	2.222	6.750	2.586	6.375	2.233

(e) Participant 5

	C		V		A		VA	
	M	SD	M	SD	M	SD	M	SD
0	7.250	0.968	6.500	1.118	6.625	1.218	6.875	1.536
1			7.125	1.533	6.875	1.452	6.875	1.269
2			7.500	1.658	7.500	1.732	6.500	2.616

(f) Participant 6

	C		V		A		VA	
	M	SD	M	SD	M	SD	M	SD
0	6.625	2.058	6.875	1.900	6.750	2.437	6.250	2.385
1			6.125	2.472	6.500	2.550	6.625	1.932
2			5.875	2.204	6.125	2.204	6.625	2.233

(g) Participant 7

	C		V		A		VA	
	M	SD	M	SD	M	SD	M	SD
0	6.125	1.763	6.125	2.204	7.125	2.421	6.750	2.463
1			6.750	2.332	6.500	2.550	6.750	1.639
2			6.500	2.121	6.750	2.385	6.000	1.803

(h) Participant 8

	C		V		A		VA	
	M	SD	M	SD	M	SD	M	SD
0	6.375	1.576	6.125	1.536	5.875	1.763	6.875	1.833
1			6.250	0.433	5.500	1.118	6.250	2.463
2			6.750	2.165	6.000	1.225	6.375	1.409

(i) Participant 9

	C		V		A		VA	
	M	SD	M	SD	M	SD	M	SD
0	6.625	1.576	6.375	1.409	6.500	1.803	6.125	1.536
1			7.000	1.658	7.375	1.495	6.375	1.576
2			6.500	1.871	7.125	1.166	6.500	1.581

(j) Participant 10

	C		V		A		VA	
	M	SD	M	SD	M	SD	M	SD
0	7.250	1.299	6.750	1.854	7.250	1.299	6.750	1.392
1			7.500	1.323	6.250	0.968	7.000	1.732
2			6.750	1.561	6.875	1.763	7.000	1.414

(k) Participant 11

	C		V		A		VA	
	M	SD	M	SD	M	SD	M	SD
0	6.375	0.857	6.625	0.696	6.375	0.696	6.250	1.897
1			7.000	1.225	6.875	1.452	6.750	1.392
2			6.500	2.000	6.375	0.857	6.750	1.299

(l) Participant 12

	C		V		A		VA	
	M	SD	M	SD	M	SD	M	SD
0	7.250	1.479	6.625	1.576	6.625	2.288	6.875	1.691
1			7.625	2.233	6.625	2.118	6.625	2.118
2			6.875	1.691	7.125	2.272	6.625	2.118

(m) Participant 13

	C		V		A		VA	
	M	SD	M	SD	M	SD	M	SD
0	6.375	0.696	6.250	0.829	6.625	0.992	7.000	0.866
1			6.500	0.866	6.625	0.857	6.750	0.661
2			6.375	1.218	7.125	0.599	5.875	1.533

(n) Participant 14

	C		V		A		VA	
	M	SD	M	SD	M	SD	M	SD
0	7.000	0.500	6.625	0.696	6.250	1.897	6.625	1.317
1			6.375	0.857	6.250	1.897	6.625	0.696
2			7.125	0.599	6.500	1.000	7.000	1.000

(o) Participant 15

	C		V		A		VA	
	M	SD	M	SD	M	SD	M	SD
0	7.375	1.867	6.750	1.854	6.875	1.965	7.125	1.533
1			7.625	1.111	5.875	1.452	7.375	1.728
2			7.625	1.218	7.125	1.269	7.125	1.615

(p) Participant 16

Table 4: Results of Experiment 2 (Precipitation Impact). Each table represents the means (M) and standard deviations (SD) of a single participant: the rows represent the precipitation intensity (0, 1, 2) and the columns the modality condition (C, V, A, VA).

5 Demonstration

The purpose of the demonstration was to apply the findings of this thesis project in a practical setting and to illustrate the research result with a concrete example. Since the precipitation impact experiment yielded results concerning interaction effects that were likely due to specific characteristics of the used stimuli, it was decided to apply the more apparent findings of the weight experiment.

5.1 Game

A game was created that takes place in a sorting center, where objects need to be sorted by weight. Unfortunately, there are no scales, thus sorting relies completely on the player. Objects fall on their arm and they must place the objects in the correct tubes. Before playing, the players are given instructions, cf. Figure 4, and are trained by experiencing the five possible weights (each with its own corresponding tube), cf. Figure 5. During the game however, only the first, third and fifth weights are used and the sizes of the objects vary (the falling speeds are constant), cf. Figure 6. Twelve objects must be sorted, after which the game is over and a score form appears, cf. Figure 7. Every correctly placed object is worth two points, every object placed one tube next to the correct one is worth one point, and every other possibility is worth zero points.

Figure 4: The instruction screens the player sees before the game starts.

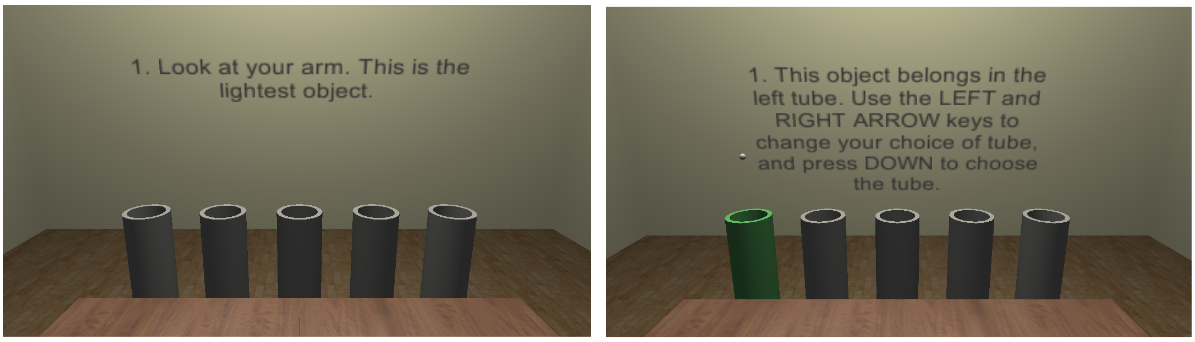

Figure 5: The first of five training rounds; (left) the first ball is hovering over the arm (outside of figure), (right) the green tube lights up indicating the correct tube the ball should be placed in.

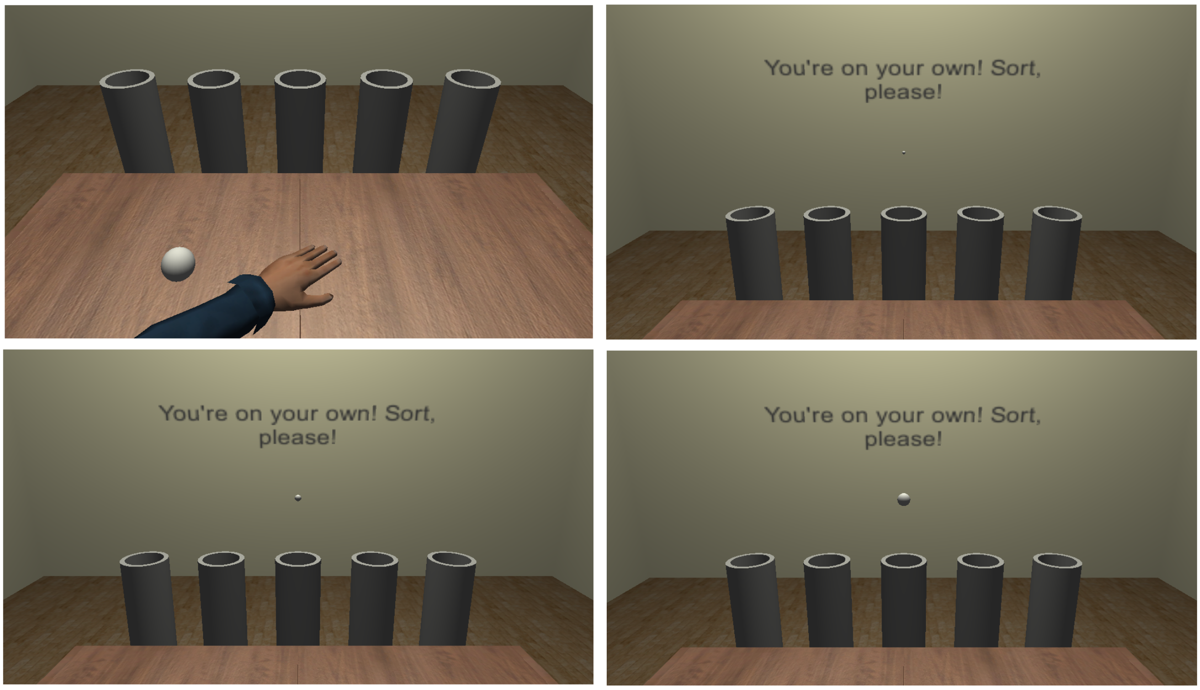

Figure 6: The game; (top-left) the arm of the participant and a ball hovering over it, (top-right) smallest ball, (bottom-left) medium ball, (bottom-right) largest ball.

SIZE	WEIGHT	ANSWER	POINTS
S	5	5	2
M	5	6	1
L	5	6	1
S	7	8	1
M	7	7	2
L	7	7	2
S	7	8	1
M	7	6	1
L	7	5	0
S	9	8	1
M	9	7	0
L	9	9	2

TOTAL SCORE: 14/24

Press the DOWN ARROW key to quit.

Figure 7: The score form that appears at the end of the game.

5.2 Implementation

To implement this, the same Unity project for the weight experiment was extended with new scenes, cf. Figure 8. The UML diagram can be found in Figure 9. Compared to the weight experiment, the code for the demonstration is simpler and the program focuses more on usability. In particular, only one characteristic of weight is used, namely size, and this characteristic is only presented visually (not audibly). This eliminates the need to distinguish between many different types of stimuli. Furthermore, the game should be playable without supervision, so the actions required from the player are explained in the game at each step. A peculiarity that occurred was that, here, the arm did not disappear while the exact same configuration settings were used throughout the thesis.

5.3 Relevance

According to the results of the weight experiment, it is highly unlikely that a participant will sort each object correctly. Every time the second and fourth tubes are used reflects the added level of difficulty of this task due to the effect of the visual stimuli. After playing, the participant is explained that, as is clear from the score sheet, not all tactile intensities from the training session were actually used and that the sizes were the cause of their incorrect answers. This game, although in a very simple form, clearly demonstrates the opportunity to create a larger range of tactile sensations by taking advantage of the visual modality.

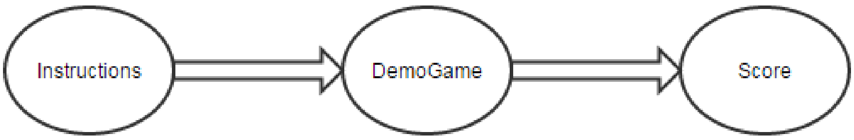

Figure 8: Outline of scenes used in the demonstration program.

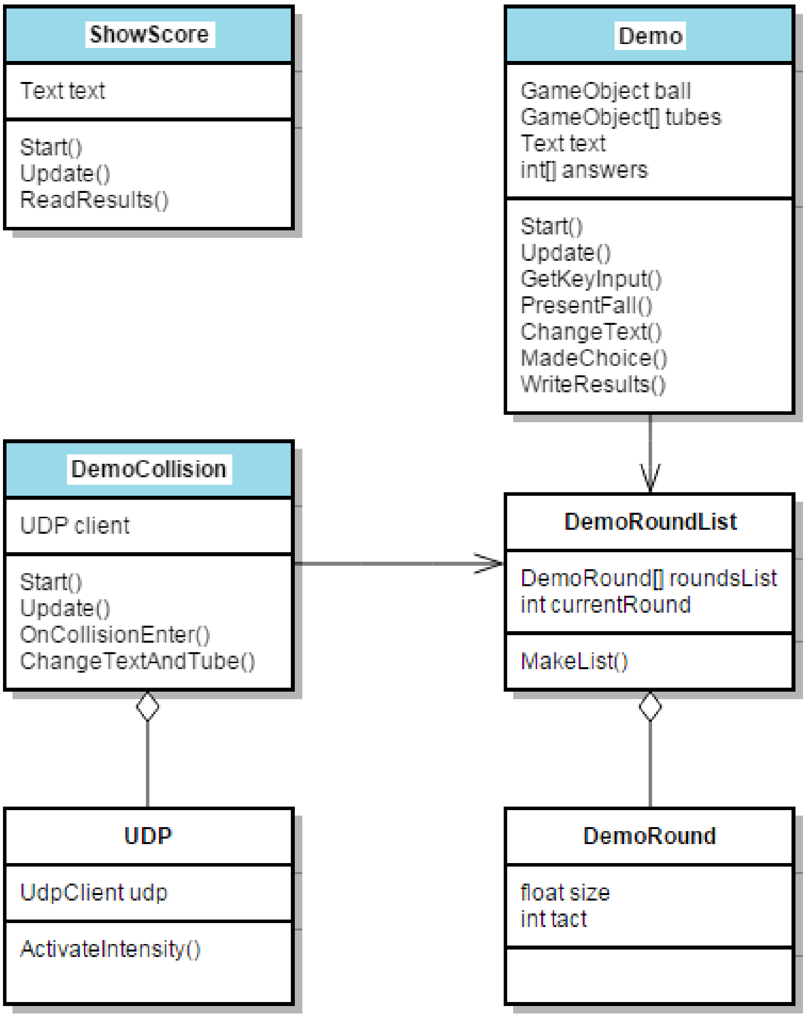

Figure 9: Simplified UML diagrams of the demonstration program. Blue indicates that the class derives from MonoBehaviour in Unity.

6 Conclusion and Future Work

Multimodal VR is a complex field with many research opportunities. In this thesis, the aim was to investigate multimodal interaction effects in VR, specifically the effect of multimodal presentation on tactile perception. A literature study identified four potential research questions:

- Can menu performance be increased through vibrotactile feedback during item selection?
- Is it possible to change properties such as quality, direction and intensity of apparent motion by multimodal stimuli?
- Can more radical deformations be achieved when visuotactile stimulation is applied on top of visuo-motor congruence?
- Is it possible to create the illusion of experiencing different intensities of a certain property using a rather simple and unrelated type of touch together with compelling multimodal stimuli?

The last question was further investigated through several experiments. The weight and temperature experiments investigated tactile perception through fundamentally consistent visual cues, and the impact experiment through visual and auditory cues derived from real world knowledge. Through these experiments, recommendations concerning design of multimodal VR experiences are formulated. It was concluded that weight-indicative visual cues in the form of size and falling speed can change the intensity of a vibration felt on the skin upon collision; this change was consistent with the size-weight and speed-force illusions. Similar results did not occur for temperature-indicative visual cues in the form of color, and although only one type of temperature-related cue was used, the discussion suggests that other cues may not be worthwhile. Despite having found no concrete outcome in the impact experiment, the results strongly support the results of the weight experiment and also highlight the need to investigate the effect or auditory cues on tactile perception. Two more open problems were identified regarding successive presentation in psychophysical experiments and the relationship between the body transfer illusion and perceiving multimodal stimuli as congruent.

This study identified various future research areas. Several interesting aspects were excluded from this project and are certainly worth further investigation. For example, this investigation only added the concept of weight to the range of acceptable touch. This study only focused on passive touch, however many directions are possible within actively explored properties, and likely even properties of apparent motion. Also, motion tracking was not investigated due to practical issues. The addition of this may have an impact on the strength of the body transfer illusion, which in turn may affect the influence of visual and auditory stimuli on tactile perception. The trail of adjustment was also not further investigated, since this concerns the mechanism behind tactile perception rather than the final percept. Such examination may lead to interesting findings regarding the processes humans use during multimodal integration in VR. Another promising future research direction resulted from the achievements of this thesis, specifically, the impact experiment demonstrated the new level of complexity when using not only visual and tactile stimuli, but also auditory stimuli. Tri-modal VR design must be further investigated in order to create effective immersive experiences. Lastly, new intriguing ideas evolved from this thesis, namely whether the effects established in this thesis also occur in augmented reality: how does mixing of real and virtual stimuli from different modalities affect multimodal integration?

7 References

[1] Banakou, D., Groten, R., and Slater, M. Illusory ownership of a virtual child body causes overestimation of object sizes and implicit attitude changes. *Proceedings of the National Academy of Sciences 110*, 31 (2013), 12846–12851.

[2] Barghout, A., Cha, J., El Saddik, A., Kammerl, J., and Steinbach, E. Spatial resolution of vibrotactile perception on the human forearm when exploiting funneling illusion. In *IEEE International Workshop on Haptic Audio Visual Environments and Games, HAVE 2009* (2009), IEEE, pp. 19–23.

[3] Bekesy, G. *Sensory inhibition.* Princeton University Press, 1967.

[4] Bice, R. C. Apparent movement in vibrotactile displays. *Perceptual and Motor Skills 29* (1969), 575–578.

[5] Botvinick, M., and Cohen, J. Rubber hands 'feel' touch that eyes see. *Nature 391* (1998), 756–756.

[6] Bowman, D., and Wingrave, C. Design and evaluation of menu systems for immersive virtual environments. In *Proceedings in the IEEE Virtual Reality* (2001), IEEE, pp. 149–156.

[7] Cha, J., Rahal, L., and El Saddik, A. A pilot study on simulating continuous sensation with two vibrating motors. In *IEEE International Workshop on Haptic Audio visual Environments and Games, HAVE 2008* (2008), IEEE, pp. 143–147.

[8] Chen, L., Shi, Z., and Müller, H. J. Influences of intra-and crossmodal grouping on visual and tactile ternus apparent motion. *Brain Research 1354* (2010), 152–162.

[9] Chertoff, D., Byers, R., and LaViola Jr, J. An exploration of menu techniques using a 3D game input device. In *Proceedings of the 4th International Conference on Foundations of Digital Games* (2009), ACM, pp. 256–262.

[10] Dachselt, R., and Hübner, A. Three-dimensional menus: A survey and taxonomy. *Computers & Graphics 31*, 1 (2007), 53–65.

[11] Ellis, S. R. Nature and origins of virtual environments: A bibliographical essay. *Computing Systems in Engineering 2*, 4 (1991), 321–347.

[12] Ernst, M. O., and Bülthoff, H. H. Merging the senses into a robust percept. *Trends in Cognitive Sciences 8*, 4 (2004), 162–169.

[13] Gebhardt, S., Pick, S., Leithold, F., Hentschel, B., and Kuhlen, T. Extended pie menus for immersive virtual environments. *IEEE Transactions on Visualization and Computer Graphics 19*, 4 (2013), 644–651.

[14] Geldard, F. A., and Sherrick, C. The cutaneous "rabbit": A perceptual illusion. *Science 178* (1972), 178–179.

[15] Gerber, D., and Bechmann, D. The spin menu: a menu system for virtual environments. In *Proceedings of the IEEE Virtual Reality 2005* (2005), IEEE, pp. 271–272.

[16] Gibson, J. J. Observations on active touch. *Psychological Review 69*, 6 (1962), 477–491.

[17] GoldStein, E. B. *Sensation and Perception.* Wadsworth, Cengage Learning, 2007.

[18] IJsselsteijn, W., de Kort, Y., and Haans, A. Is this my hand I see before me? The rubber hand illusion in reality, virtual reality, and mixed reality. *Presence 15*, 4 (2006), 455–464.

[19] Israr, A., Kim, S.-C., Stec, J., and Poupyrev, I. Surround haptics: tactile feedback for immersive gaming experiences. In *CHI'12 Extended Abstracts on Human Factors in Computing Systems* (2012), ACM, pp. 1087–1090.

[20] Israr, A., and Poupyrev, I. Exploring surround haptics displays. In *CHI'10 Extended Abstracts on Human Factors in Computing Systems* (2010), ACM, pp. 4171–4176.

[21] Israr, A., and Poupyrev, I. Control space of apparent haptic motion. In *World Haptics Conference (WHC), 2011 IEEE* (2011), IEEE, pp. 457–462.

[22] Israr, A., and Poupyrev, I. Tactile brush: drawing on skin with a tactile grid display. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2011), ACM, pp. 2019–2028.

[23] Jacoby, R. H., and Ellis, S. R. Using virtual menus in a virtual environment. In *SPIE/IS&T 1992 Symposium on Electronic Imaging: Science and Technology* (1992), International Society for Optics and Photonics, pp. 39–48.

[24] Kilteni, K., Bergstrom, I., and Slater, M. Drumming in immersive virtual reality: the body shapes the way we play. *IEEE Transactions on Visualization and Computer Graphics 19*, 4 (2013), 597–605.

[25] Kilteni, K., Normand, J.-M., Sanchez-Vives, M. V., and Slater, M. Extending body space in immersive virtual reality: a very long arm illusion. *PLOS ONE 7*, 7 (2012).

[26] Kim, Y., Lee, J., and Kim, G. Extending out of the body saltation to 2D mobile tactile interaction. In *Proceedings of the 10th Asia Pacific Conference on Computer Human Interaction* (2012), ACM, pp. 67–74.

[27] Kirman, J. H. Tactile apparent movement: the effects of interstimulus onset interval and stimulus duration. *Perception & Psychophysics 15*, 1 (1974), 1–6.

[28] Kirman, J. H. Tactile apparent movement: the effects of number of stimulators. *Journal of Experimental Psychology 103*, 6 (1974), 1175.

[29] Kirman, J. H. The effect of number of stimulators on the optimal interstimulus onset interval in tactile apparent movement. *Perception & Psychophysics 17*, 3 (1975), 263–267.

[30] Kirman, J. H. Tactile apparent movement: the effects of shape and type of motion. *Perception & Psychophysics 34*, 1 (1983), 96–102.

[31] Kohli, L., Niwa, M., Noma, H., Susami, K., Yanagida, Y., Lindeman, R. W., Hosaka, K., and Kume, Y. Towards effective information display using vibrotactile apparent motion. In *14th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2006* (2006), IEEE, pp. 445–451.

[32] Kulshreshth, A., and LaViola Jr, J. J. Exploring the usefulness of finger-based 3D gesture menu selection. In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems* (2014), ACM, pp. 1093–1102.

[33] Lanier, J. Homuncular flexibility. https://edge.org/response-detail/11182. Extracted 19-06-2015.

[34] Lederman, S. J., and Klatzky, R. L. Hand movements: A window into haptic object recognition. *Cognitive Psychology 19*, 3 (1987), 342–368.

[35] Lee, J., Kim, Y., and Kim, G. Funneling and saltation effects for tactile interaction with virtual objects. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2012), ACM, pp. 3141–3148.

[36] Maselli, A., and Slater, M. The building blocks of the full body ownership illusion. *Frontiers in Human Neuroscience 7*, 83 (2013).

[37] Ni, T., Bowman, D. A., North, C., and McMahan, R. P. Design and evaluation of freehand menu selection interfaces using tilt and pinch gestures. *International Journal of Human-Computer Studies 69*, 9 (2011), 551–562.

[38] Ni, T., McMahan, R. P., and Bowman, D. A. Tech-note: rapmenu: Remote menu selection using freehand gestural input. In *IEEE Symposium on 3D User Interfaces, 3DUI 2008* (2008), IEEE, pp. 55–58.

[39] Nielsen, J. Noncommand user interfaces. *Communications of the ACM 36*, 4 (1993), 83–99.

[40] Niwa, M., Lindeman, R. W., Itoh, Y., and Kishino, F. Determining appropriate parameters to elicit linear and circular apparent motion using vibrotactile cues. In *EuroHaptics conference, 2009 and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. World Haptics 2009. Third Joint* (2009), IEEE, pp. 75–78.

[41] Normand, J.-M., Giannopoulos, E., Spanlang, B., and Slater, M. Multisensory stimulation can induce an illusion of larger belly size in immersive virtual reality. *PLOS ONE 6*, 1 (2011).

[42] Peck, T. C., Seinfeld, S., Aglioti, S. M., and Slater, M. Putting yourself in the skin of a black avatar reduces implicit racial bias. *Consciousness and Cognition 22*, 3 (2013), 779–787.

[43] Petkova, V. I., and Ehrsson, H. H. If I were you: perceptual illusion of body swapping. *PLOS ONE 3*, 12 (2008).

[44] Petkova, V. I., Khoshnevis, M., and Ehrsson, H. H. The perspective matters! Multisensory integration in egocentric reference frames determines full-body ownership. *Frontiers in Psychology 2* (2011), 35.

[45] Rahal, L., Cha, J., and El Saddik, A. Continuous tactile perception for vibrotactile displays. In *IEEE International Workshop on Robotic and Sensors Environments, ROSE 2009* (2009), IEEE, pp. 86–91.

[46] Sanchez-Vives, M. V., and Slater, M. From presence to consciousness through virtual reality. *Nature Reviews Neuroscience 6*, 4 (2005), 332–339.

[47] Sanchez-Vives, M. V., Spanlang, B., Frisoli, A., Bergamasco, M., and Slater, M. Virtual hand illusion induced by visuomotor correlations. *PLOS ONE 5*, 4 (2010).

[48] Sherrick, C. E., and Rogers, R. Apparent haptic movement. *Perception & Psychophysics 1*, 3 (1966), 175–180.

[49] Shimojo, S., and Shams, L. Sensory modalities are not separate modalities: plasticity and interactions. *Current Opinion in Neurobiology 11*, 4 (2001), 505–509.

[50] Slater, M., Perez-Marcos, D., Ehrsson, H. H., and Sanchez-Vives, M. V. Towards a digital body: the virtual arm illusion. *Frontiers in Human Neuroscience 2* (2008).

[51] Slater, M., Spanlang, B., Sanchez-Vives, M. V., and Blanke, O. First person experience of body transfer in virtual reality. *PLOS ONE 5*, 5 (2010).

[52] Steptoe, W., Steed, A., and Slater, M. Human tails: ownership and control of extended humanoid avatars. *IEEE Transactions on Visualization and Computer Graphics 19*, 4 (2013), 583–590.

[53] Steuer, J. Defining virtual reality: Dimensions determining telepresence. *Journal of communication 42*, 4 (1992), 73–93.

[54] Won, A. S., Bailenson, J., Lee, J., and Lanier, J. Homuncular flexibility in virtual reality. *Journal of Computer-Mediated Communication* (2015).

[55] Yee, N., and Bailenson, J. The Proteus effect: The effect of transformed self-representation on behavior. *Human Communication Research 33*, 3 (2007), 271–290.

[56] Zhao, S., Dragicevic, P., Chignell, M., Balakrishnan, R., and Baudisch, P. Earpod: eyes-free menu selection using touch input and reactive audio feedback. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2007), ACM, pp. 1395–1404.