User-Defined Gestures with Physical Props in Virtual Reality

by

Marco Aurelio Moran Ledesma

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Applied Science
in
Systems Design Engineering

Waterloo, Ontario, Canada, 2020

Author's Declaration

This thesis consists of materials all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

While the first person singular, “I”, is used throughout this thesis, the contents are collaborative efforts with my supervisors. These include:

- Dr. Mark Hancock
- Dr. Oliver Schneider

Abstract

When building virtual reality (VR) environments, designers use physical props to improve immersion and realism. However, people may want to perform actions that would not be supported by physical objects, for example, duplicating an object in a Computer-Aided Design (CAD) program or darkening the sky in an open-world game. In this thesis, I present an elicitation study where I asked 21 participants to choose from 95 props to perform manipulative gestures for 20 referents (actions), typically found in CAD software or open-world games. I describe the resulting gestures as context-free grammars, capturing the actions taken by our participants, their prop choices, and how the props were used in each gesture. I present agreement scores between gesture choices and prop choices; to accomplish the latter, I developed a generalized agreement score that compares sets of selections rather than a single selection, enabling new types of elicitation studies. I found that props were selected according to their resemblance to virtual objects and the actions they afforded; that gesture and prop agreement depended on the referent, with some referents leading to similar gesture choices, while others led to similar prop choices; and that a small set of carefully chosen props can support a wide variety of gestures.

Acknowledgements

Thank you to my supervisors, Dr. Mark Hancock and Dr. Oliver Schneider, for providing guidance and feedback throughout this project. Thank you to Dr. Catherine Burns and Dr. Daniel Vogel for serving as readers of this thesis. Thanks also to my girlfriend Caroline, for putting up with me being sat in the office for hours on end, and for providing endless guidance and support voluntarily. Thank you to my family and to my friends for keeping me afloat and for encouraging me to fulfill my dreams wherever I wish. Thanks also to the University of Waterloo Biomedical Engineering class of 2022, for being an outstanding group of young people and for helping me discover the career goal I want to pursue in my life. Thanks to my professors Eihab Abdel-Rahman, Rasoul Keshavarzi-Valdani, Orion Bruckman, and Bryan Tripp, for their continuous trust to serve as a teaching assistant and for their irreplaceable mentoring to improve my teaching skills. Finally, thank you Lord for this amazing journey.

This thesis is dedicated to Jesus, Celestino, and Karol, who stand at the window of the Father's House.

Table of Contents

List of Figures **ix**

List of Tables **xiii**

1 Introduction **1**

1.1 Research Questions 4

1.2 Contributions 4

1.3 Thesis Organization 5

2 Background **6**

2.1 Gesture Elicitation Studies in HCI 7

2.1.1 Agreement Assessment Methods 9

2.1.2 Agreement Assessment Issues in Elicitation Studies 10

2.1.3 Gesture-based User Interfaces for CAD Modelling 11

2.2 Haptics Technology in Virtual Reality . . . 12

2.2.1 Active Haptics in Virtual Reality . . . 13

2.2.2 Passive Haptics in Virtual Reality . . . 16

2.2.3 Dynamic Passive Haptics . . . 18

2.3 Chapter Summary . . . 20

3 Study Method **22**

3.1 Methodology . . . 22

3.1.1 Selection of VR Referents . . . 24

3.1.2 Selection of Physical Props . . . 25

3.1.3 Participants . . . 28

3.1.4 Apparatus . . . 29

3.1.5 Procedure . . . 29

3.2 Chapter Summary . . . 31

4 Results **33**

4.1 A Participant-Defined Gesture Set . . . 33

4.2 Gesture Agreement . . . 36

4.3 Prop Agreement . . . 39

4.4 Classification of Gestures with Props . . . 43

4.5 Chance agreement and agreement confidence . . . 45

5 Discussion **51**

6 Conclusion, Limitations, and Future Work **57**

6.1 Limitations 60

6.2 Future Work 62

References **64**

APPENDICES **82**

A PDF Plots From Matlab **83**

A.1 Code for Calculating Prop Agreement Score of Referent C: Bevel 83

A.2 Code for Calculating Confidence Intervals of Prop Agreement Score of Referent C: Bevel 85

B Ethics Approval **87**

C Study Supplementary Materials **90**

List of Figures

3.1 Gesture elicitation study with props that follows a previous methodology (Wobbrock et al., 2009). A participant: a) watches the completion of a referent, b) chooses a set of physical props, and c) performs a gesture with the chosen prop(s) that would complete the referent. 24

3.2 The vocabulary of 95 physical props used for this elicitation study. 50 props were 3D printed, 25 with flexible material and 25 with rigid material. 28 props were LEGO-brand toys, 8 specific models and 20 spare bricks. 13 retail-manufactured objects and 3 handheld gaming controllers from Vive, Oculus and Nintendo Switch. 28

3.3 Two participants performing in the study. On the left, participant 4 picking up a retail-manufactured flute for gesture demonstration. On the right, participant 10 assembling a model with LEGO-brand bricks that was not available. 30

3.4 Study room setting 31

4.1 Three of the thirteen gestures identified for referent C: *Bevel*. Each gesture is articulated as a context-free grammar, showing what physical props can be used with that gesture. 35

4.2 Complete graph that captures gestures with physical props performed by 5 participants for referent C: *Bevel*. Gestures are considered "one-element" sets and similarity is calculated with one another, where a 1 means "same gesture" and 0 means "distinct gesture". No partial similarity is considered. Similarity between two sets is the weight of the edge that joins them. Agreement score is computed by summing the similarity of each pair of sets compared, and dividing by the total number of pairs that can be identical. 37

4.3 Gesture agreement scores of the 20 VR referents sorted in descending order. 40

4.4 Complete graph that captures props sets of 5 participants for referent C: *Bevel*. Similarity of props sets is measured, which is the weight of the edges that joins them. Each measure of set similarity is between 0 and 1, with 0 indicating no shared props and 1 indicating completely identical prop selections. Agreement is computed by summing the similarity of each pair of sets compared, and dividing by the total number of pairs that can be identical. 42

4.5 Venn diagrams showing a pairwise comparison between prop sets of P1 and P4. Three measures of set similarity are proposed to evaluate partial agreement (SIM(P1,P4)): Overlap, Sørensen, and Jaccard. Overlap is the most optimistic, leading to the highest agreement. Jaccard is the most pessimistic, leading to the lowest agreement. Sørensen prevails between the other two (or equal to one of them). In this pairwise comparison, Sørensen and Overlap are the same. 43

4.6 Gesture set with physical props for referents A to J. For each referent, the preferred set of props and the preferred gesture are presented; the initial state of the most common prop-gesture interaction is shown on the left side, and the end state is shown on the right. 44

4.7 Prop agreement scores of the 20 VR referents sorted in descending order. . 45

4.8 Gesture set with props for referents A to J. For each referent, the preferred set of props and the preferred gesture are presented; the initial state of the most common prop-gesture interaction is shown on the left side, and the end state is shown on the right. 48

4.9 Gesture set with props for referents A to J. For each referent, the preferred set of props and the preferred gesture are presented; the initial state of the most common prop-gesture interaction is shown on the left side, and the end state is shown on the right. 49

4.10 Three different categories of gestures were created from qualitative observations: 1) Act, 2) Leverage, and 3)Adapt. Each category groups several gesture possibilities. For a given referent, a participant can: 1a) alter the physical properties of a prop with hands or with more props, 1b) change the position of props in a 3D space, 1c) operates movable parts of a prop, 2a) adhere a prop, or a section of it, to another one, 2b) show a comparison of property of two or more props, 3a) disassemble an initial assembly and re-assemble it differently, and/or 3b) assemble a single model to represent the effect. 50

5.1 Number of gestures from Figure 4.8 and Figure 4.9 for which each prop mentioned the context-free grammar can be used for. 52

List of Tables

3.1 Labels and descriptions of the 20 VR Referents included in this elicitation study. These referents were extracted from prior work (Huang and Rai, 2014; Khan and Tunçer, 2019; Khan et al., 2019; Nanjappan et al., 2018), CAD 3D modelling software, VR CAD modelling applications, and Open-World games. Figure 4.8 and Figure 4.9 provide an illustration of each. 26

4.1 For each VR referent, number of unique gestures recognized from the 21 participants, number of participants that performed each of these unique gestures, and gesture agreement score are reported. 39

Chapter 1

Introduction

Today, much of what we think of Virtual reality (VR) comes from the recent surge and rapid advances in gaming and VR headsets. VR, nonetheless, is more than simply an entertainment platform where users put on a headset and use a pair of handheld controllers to interact with a computer-generated simulation of a three-dimensional space. VR represents an extraordinary change in the way people interact with the digital realm. So far, our experience with computers has been through mice, keyboards, and tactile devices. VR changes all that by surrounding us with fully immersive environments, layered with multiple sets of information that we can interact with. With that being said, many fields have been adopting VR throughout the last decades, such as architecture (Kharvari and Hohl, 2019; Neil, 1996), education (Calderon and Ruiz, 2019; Saunier et al., 2016), manufacturing (Hamid et al., 2014; Liu et al., 2018), and healthcare (Kim and Park, 2009; Qian et al., 2015). For example, with VR, an architect can create the 3D model of an envisioned home, a surgeon can perform on a real body, or a geography class can visit Mars.

In the context of VR, an interaction is often described as the ability to move within a virtual environment and to interact with objects in it (Bostan and Nalbant, 2006). In VR, interaction designers increasingly employ physical objects, or "props", to mimic their real-world counterpart, and research has demonstrated that physical props are effective at improving immersion (Azmandian et al., 2016) as well as enhancing realism (Hoffman et al., 1998; Schulz et al., 2019). In addition, due to the non-physically realistic experience possibilities that VR enables, users may want to perform actions that typically would not be supported by physical props, for example, duplicating an object in a CAD modelling program or darkening the sky in an open-world game. In this thesis, I propose to expand the known input vocabulary of physical props by combining them with gestures, which offer an alternative interaction input that is often described as more natural and intuitive (Jahani and Kavakli, 2018), as well as allowing for easy spatial manipulation (Alaçam, 2014).

Individual gestures and sets of gestures are difficult to design (Long et al., 1999). Instead of arbitrary designs that are technology- or designer-centered, researchers have proposed directly asking people who would use these gestures (Wobbrock et al., 2009). To design gestures, researchers have widely adopted elicitation studies, where people who might use a system provide their input without regard for recognition or technical concerns. Elicitation studies have shown promising outcomes for gesture design in a variety of domains, for example, tabletops (Wobbrock et al., 2009), mobile devices (Ruiz et al., 2011), smartwatches (Arefin Shimon et al., 2016), televisions (Vatavu, 2012), public displays (Kray et al., 2010), virtual reality (Nanjappan et al., 2018), and augmented reality (Piumsomboon et al., 2013).

Additionally, elicited gestures have been shown to be easier to remember and preferred by those without technical expertise (Nacenta et al., 2013).

In this thesis, I conducted an elicitation study for manipulative gestures (Quek et al., 2002) with physical props in VR. I chose referents from CAD modelling software and open world games to ground my work in application areas that offer a rich vocabulary of actions and commands, but might also use physical props. I followed Wobbrock et al.'s approach (2009) by first showing the execution of a referent (action) in VR and then asking participants for their preferred gesture, but also ask them to choose their preferred props(s) to perform the gesture with. The result was a set of twenty user-defined gestures with props, one for each referent.

I developed two representations to handle both gestures and props to analyze the data from this study. 1) I found that context-free grammars were a useful method to capture the various props used in a gesture, their role, and how people used them; they are presented alongside illustrations of our elicited gestures. This language succinctly communicates gestures, could be directly implemented into systems using props for gestures, and can support future, more-involved analysis. 2) While I calculated agreement scores for gestures using methods from previous elicitation studies (Vatavu and Wobbrock, 2015; Wobbrock et al., 2009), these scores were not suitable for the multiple props that participants selected. Thus I introduced a new agreement score based on set similarity metrics, that is used to analyze agreement between both gestures and props. This score is identical to previous agreement scores (Vatavu and Wobbrock, 2015) when used for unimodal selection (e.g., gestures), but also accommodates multimodal selections (e.g., gestures and props).

1.1 Research Questions

Combining manipulative gestures with physical props for interactions in virtual environments raises several research questions. In this thesis, I investigated the following:

1. What gestures would people perform with physical props to complete CAD-like and open-world referents in VR?

2. What physical props would people choose to manipulate objects in VR?

3. How would people leverage physical props to manipulate virtual objects?

1.2 Contributions

The contributions of the research presented in this thesis are:

- A set of gesture-prop combinations for 20 CAD-like and open-world-like referents.

- A language for articulating and implementing gestures with physical props, based on context-free grammars.

- A generalization of the agreement score metric to account for multiple selections in elicitation studies.

- Insights into how people leverage physical props to perform gestures. Results suggest two strategies taken by participants: pick a prop and design a suitable gesture for it, or design a gesture and pick a suitable prop for it.

1.3 Thesis Organization

I organized this thesis into five chapters following the introduction.

- In chapter 2, I present a literature review with relevant background in the fields of gesture elicitation studies and haptics technology for VR. In the first section of this chapter, I analyze several elicitation studies that followed the methodology used for this study. In the second section, I review a large body of haptic feedback devices. I categorize them into active haptics, passive haptics, and dynamic passive haptics. The majority of the work discussed in chapter 2 belongs to the Human Factors in Computing Systems (CHI) and the User Interface Software and Technology (UIST) communities.

- In chapter 3, I present the study methods as well as the design of the elicitation study. I discuss and ground relevant components of the study, such as how I created the vocabulary of props and how I picked the 20 VR referents.

- In chapter 4, I show the results of the elicitation study. In this chapter, I report Gesture and prop agreement scores for each VR referent, I explain context-free grammars to articulate gestures with props, I present the user-defined gesture set for 20 VR referents, and I introduce a classification of gestures with props.

- In chapter 5, I discuss implications of the results for VR systems that use physical props and for future elicitation studies.

- In chapter 6, finally, I provide a general conclusion of my research, highlighting contributions, relevance, limitations, and future work.

Chapter 2

Background

In this chapter, I provide readers with an overview of previous work around gesture elicitation studies for a variety of domains, and haptic feedback to improve immersion and enhance realism in VR. I divide this chapter into two main sections, with two and three subsections, respectively. In section 2.1, I discuss existing gesture elicitation studies in different human-computer interaction fields. In subsection 2.1.3, I discuss two gesture-based user interfaces for CAD modelling, from which I extracted a group of referents for the study I present in this thesis. In section 2.2, I discuss research in the area of haptic technologies in virtual reality experiences. I break down this section into active haptics (2.2.1), passive haptics (2.2.2), and dynamic passive haptics (2.2.3).

2.1 Gesture Elicitation Studies in HCI

Gesture elicitation is a widely used technique in Human-Computer Interaction (HCI) for identifying gesture vocabularies that are self-discoverable (Tsandilas, 2018). In an elicitation study, users are commonly shown the effect of an action (referent) and asked to propose a gesture that would trigger this referent. Vatavu (2012) argues that incorporating users in the design process represents a viable alternative for collecting important data to inform design. Wobbrock et al. (2009) developed a user-defined set of gestures based on the degree of consensus (agreement score Wobbrock et al. (2005)) among participants to complete 27 referents. Agreement analysis can guide the design of gesture sets and help understand how referents or commands naturally map to gestures. Wobbrock et al. (2009) also classified their elicited vocabulary of gestures in a taxonomy for tabletop systems design, which aims to capture the gesture design space in a tabletop environment. Later, the authors compared this gesture set with a gesture set proposed by a designer and found that the user-defined set is easier to remember (Morris et al., 2010).

Since Wobbrock et al.'s work, elicitation studies have become a common practice for determining suitable gestures, but so far most studies have been limited to actions using hands and fingers to interact with specific technology, such as mobile devices (Gomes et al., 2018; Ruiz et al., 2011), televisions (Vatavu, 2012), public displays (Kray et al., 2010), smartwatches (Arefin Shimon et al., 2016), and augmented reality (Piumsomboon et al., 2013). Kray et al. (2010) ran an elicitation study where participants produced natural gestures with their mobile phone to trigger 20 referents. They investigated three different conditions: phone-to-phone, phone-to-tabletop, and phone-to-public displays. Later, Ruiz et al. (2011) presented a study that elicited motion gestures to execute 19 commands on smartphone devices (e.g., answer a call, go to home screen, or ignore call), resulting in a user-defined gesture set. After that, Vatavu (2012) proposed a set of candidate gestures for 12 common TV control referents, such as volume up, help, or open. Then, Piumsomboon et al. (2013) developed a user-defined gesture set for 40 referents in AR. Their vocabulary of referents was divided into categories, for example, transform (rotate, scale, or move), control (play, stop, or reset), and edit (insert, delete, or cut). This study aimed to capture gesture input that users would use to manipulate objects in AR. After that, Shimon et al. (2016) conducted their elicitation study to build a set of non-touchscreen gestures for 31 smartwatches referents. As it can be seen, the elicitation technique has been widely adopted to propose novel interaction techniques with a variety of interfaces.

Other elicitation studies have focused on creating gestures sets while manipulating devices. Valdes et al. (2014) investigated the use of gestural interaction with active tokens for manipulating large information sets on multi-touch and tangible interfaces. The user-generated gesture set includes manipulative gestures such tilting a token or neighboring two tokens. *MagicScroll*, by Gomes et al. (2018), is a rollable table for gestural input. The authors used this device to create a gesture set for action and navigation-based tasks that people perform on mobile platforms. Lastly, Nanjappan et al. (2018) ran an elicitation study to create a set of natural and suitable gestures for 17 VR referents using dual-hand VR controllers.

In this thesis, I adopted the elicitation study methodology to determine a gesture set for using physical props to control virtual reality, and by necessity build on this previous work to be able to incorporate not only the choice of gesture, but also the choice of physical artifact used with the gesture.

2.1.1 Agreement Assessment Methods

To measure agreement among participants' proposals (gestures) in elicitation studies, researchers have used the agreement index A introduced by Wobbrock et al. (2005, 2009), or the agreement rate AR (Vatavu and Wobbrock, 2015) metric, a formalized version of the agreement index A formula. Intuitively, one assumes that: if for a given referent, the number of gestures proposed is equal to the number of participants (i.e., no gesture is repeated), the agreement score would yield zero. Unfortunately, the original agreement index A formula does not capture this. The agreement rate (AR) formula does, since it is defined by Vatavu and Wobbrock as: "the number of pairs of participants in agreement with each other divided by the total number of pairs of participants that could agree". For instance, if not a single pair of participants agree with each other, the score computed is zero.

Ruiz et al., Vatavu, Piumdomboon et al. and Nanjappan et al. used the agreement index formula. Shimon et al., on the other hand, used the agreement rate formula to quantify gesture agreement scores. In this thesis, I used the agreement rate formula to calculate gesture agreement scores. However, I made a generalization in order to introduce a new agreement score based on set similarity metrics that allows one to analyze agreement between both gestures and props.

2.1.2 Agreement Assessment Issues in Elicitation Studies

Tsandilas (2018) identified a series of problems with elicitation studies and agreement assessment methods (Vatavu and Wobbrock, 2015; Wobbrock et al., 2009):

1. Measures such as the agreement index A and agreement rate AR do not consider that agreement between participants can occur by chance, and chance agreement is present in gesture elicitation studies due to bias factors. Tsandilas uses two distributions with properties that reflect how bias affects chance agreement, Zipf-Mandelbrot and discrete half-normal probability distributions, to demonstrate that as bias increases, chance agreement also increases. Prior experience and knowledge with existing interfaces, poor study settings, and low fidelity device prototypes are common sources of bias in elicitation studies (Morris et al., 2014; Ruiz and Vogel, 2015). These and other factors (memorability, acceptability, physical complexity) can influence participants to focus on certain gestures and disregard others. According to Tsandilas, researchers should ensure their agreement measures take bias factors into account.

2. Guidelines for interpreting levels of agreement rate (AR) proposed by Vatavu and Wobbrock (2015) (low, medium, high, and very high) rely on conflicting assumptions that can lead to overoptimistic conclusions about agreement reached by participants. To derive these guidelines, Vatavu and Wobbrock relied on a flawed probability distribution and on a survey of agreement rates from many previous elicitation studies. Tsandilas specifically highlights the need for more research in this field to define universal agreement scores levels. For now, he encourages researchers to avoid reporting

levels to interpret agreement. Instead, as Gwet (2014) and Krippendorff (2004) claim, confidence intervals can be provided to better interpret agreement.

3. The Vrd and Vb significant tests proposed by Vatavu and Wobbrock for comparing agreement scores within participants and between participants, respectively, rely on probabilistic assumptions that yield extremely high Type I error rates.

In this thesis, I followed Tsandilas' recommendations to provide a solution for problem 1. Specifically, I used Scott's (1955) π and Fleiss' (1971) κ_F coefficients to report chance agreement of gestures. This increases transparency and help researchers to better interpret the results. To partially reduce the effect of problem 2 on this work, I adopted the *Jackknifing* re-sampling method (Quenouille, 1949) to report confidence intervals of both gestures and prop agreement. With regard to problem 3, computing statistical significance tests for agreement scores is outside the scope of this work.

2.1.3 Gesture-based User Interfaces for CAD Modelling

In my thesis, I explored the domain of 3D modelling in computer-automated design (CAD), and some work has already explored the use of gestures to perform 3D modelling tasks. Khan et al. (2019) presented a compilation of a set of gestures and speech commands for 3D CAD modelling for conceptual design that were elicited from participants and evaluated by experts individually. In this study, the authors included modelling tasks like "rotate", "scale", and "zoom in", which are some of the referents in this thesis. Huang et al. (2014) also presented a system that recognized hand gestures along with hand position information and converted them into commands for rotating, translating and scaling 3D models. I expanded the vocabulary of commands (referents) of this previous work by incorporating tasks such as “changing colour”, “bending”, “perforating”, “beveling” to my own work.

2.2 Haptics Technology in Virtual Reality

As VR evolves at an accelerated rate, novel and sophisticated technologies surge to bring make virtual environments approximate the physical world. One such technology is haptics technology. VR haptics technology allows people to feel virtual environments via the sense of touch, in addition to visual and auditory perception. Traditionally haptics technology has come in the form of external devices like gloves, shoes, controllers, joysticks, et cetera, through which users receive feedback in the form of vibrations. This feedback provides realistic physical sensations in the hand or other parts of the body, similar to the ones people realize in the real world. Several studies have demonstrated the evident benefits of incorporating haptic feedback in VR experiences, such as fostering embodied interaction, presence, and immersion (Azmandian et al., 2016; Frohner et al., 2019; Schulz et al., 2019).

Nonetheless, VR haptics technology is growing beyond creating vibrations. In recent times, there has been an impressive amount of research in the area of incorporating physical objects or devices and haptic feedback in virtual reality experiences, in the form of active haptics, passive haptics, and dynamic passive haptics.

2.2.1 Active Haptics in Virtual Reality

Active haptic feedback (AHF) technology provides feedback through integrated systems of powered actuators (Zenner and Kruger, 2017) that exert forces on users of VR systems (McClelland et al., 2017) to render virtual content. Researchers in this area have spent a significant amount of time and effort customizing controllers for input in VR, some of which come in the following forms:

- Grounded/ungrounded shape-changing surfaces
- Electro-mechanical actuators
- Pneumatically-actuated interfaces
- Touch-and-texture rendering devices

Grounded/ungrounded Shape-changing Surfaces

Work in this area belongs to the encountered-type haptic systems: active devices that move or change shape such that when users touch a virtual object, they encounter the haptic device (Abtahi and Follmer, 2018). Abtahi and Follmer (2018) designed and evaluated a matrix of actuated pins that physically renders 2.5D content (Leithinger et al., 2015). The authors developed three applications to demonstrate its functionality: a pentagon maze in which the walls of a maze are rendered to provide haptic feedback as the user's finger follows the path, a virtual museum where a sculpted relief piece is reproduced so that the user haptically explore art work, and a ball bouncing on the user's finger. These exploratory applications are limited to single finger interactions. Siu et al. (2018), however, constructed a similar system for manipulating tangibles that allows user interaction using the entire hand. Among the renderings generated with this device, a terrain map and a house can be found. Shape-changing surfaces can provide realistic haptic feedback, but low resolution, pin speed, and surface size limit their usage for manipulating virtual objects the same way it is done in the real world.

Electro-mechanical actuators

The *CLAW* (Choi et al., 2018) VR handheld controller provides articulated movement and force feedback actuation to user's index finger. This device allows for convincing haptic rendering of grasping a ball, pushing a button, touching spaced grooves, and triggering a gun. *CapstanCrunch* (Sinclair et al., 2019) is another handheld controller that reproduces touch and grasp haptic sensations in VR. According to the authors, to successfully render haptics for virtual objects, controllers must be built to produce and endure human-scale forces during interaction. CapstanCrunch possess this feature due to a controlled brake that renders grasp feedback at varying stiffness. This is demonstrated with an experiment where the device renders haptic sensation of a virtual ball being pinched by the user. Other renderings presented in this work were pressing a button and operating scissors.

Touch-and-texture Rendering Devices

Benko et al. (2016) presented a couple of mechanically-actuated hand-held controllers that enables users to feel 3D surfaces, textures, and forces that match visual content: *Normal-*

Touch and *TextureTouch*. The first one provides haptic feedback through an active tiltable and extrudable platform. The second one does it via an array of actuated pins. These devices can render a variety of rigid and deformable 3D objects, such as simple shapes, as well as more complex ones, like cars and animals. On the other hand, Whitmire et al. (2018) built *Haptic Revolver*, a handheld VR controller that renders touch contact, pressure, shear forces, textures and other shapes using a rotating wheel beneath the index finger. This device is capable of rendering demos of a playing card game, painting, sculpting, and typing. It is evident that tiltable platforms, automated pins, rotating wheels or similar devices in this area can render multiple haptic sensations. However, interactivity with virtual environments is achieved only with a finger, which limits their usage for manipulating virtual objects.

Pneumatically-actuated interfaces

Force Jacket (Delazio et al., 2018) is an array of pneumatically-actuated airbags that provide precise force and vibrations to the upper body. Its features allow for rendering of a snowball fight or a snake crawling on the user. *TilePop* (Teng et al., 2019), similarly, reproduces virtual experiences, such as mounting a dinosaur or navigating on a boat with inflatable floor tiles. This work illustrate interesting approaches to rendering virtual experiences, rather than only virtual objects. It is also a promising technique for providing haptic feedback for the whole body.

The four forms of active haptic feedback (AHF) discussed are only a small sample of the extensive literature in this field. Evidently, AHF offers interaction techniques that realistically simulate how people interact with physical objects in the real world, but it heavily relies on systems that are complex, intrusive, expensive, robust, or a combination of these (McClelland et al., 2017).

2.2.2 Passive Haptics in Virtual Reality

In contrast to AHF, passive haptic feedback (PHF) does not use computer-controlled actuators to exert forces (Zenner and Kruger, 2017). Instead, passive haptics uses static physical objects of different materials and building techniques as props in VR. Props are a feasible approach with natural feedback qualities (Arora et al., 2019; McClelland et al., 2017; Muender et al., 2019) to manipulate virtual objects in virtual environments. Manipulating virtual objects with appropriate PHF provides a satisfying sense of presence in VR (Azmandian et al., 2016).

Some authors have presented interesting approaches with single-purpose passive props that can be mapped to virtual objects in different ways. Yang et al. (2018) created *VR Grabbers*, a passive chopsticks-like VR controller for precise virtual object manipulation that works on an *ungrounded haptic retargeting* technique. This device allows grabbing and dropping a variety of virtual objects of different size and geometry. Muender et al. (2019) fabricated a tree, an avatar, and a trampoline using three different fidelities for each: LEGO-brand bricks, 3D-printing, and uniform shapes. The authors evaluated the effect of tangibles with different haptic fidelities on immersion, performance, and intuitive interaction for a 3D scene created in VR. Chang et al. (2017) described *TASC*, a system of tangible objects to create a strong sense of embodiment in a virtual environment for spatial puzzle solving.

This system consisted of a pair of wooden rectangular prisms on a rail that could only be moved along one axis. Strandholt et al. (2020) used a hammer, a screwdriver, and a hand saw, a wooden surface, and a wooden plank to provide hammering, screwing, and sawing haptic experiences. This was done through redirected tool-mediated manipulation.

Since it can be expensive and impractical to map each physical proxy to a virtual object, authors have directed their work to developing multi-purpose or reconfigurable objects. In this way, users enjoy multiple experiences with a single system. In recent times, Arora et al. (2019) presented *VirtualBricks*, a LEGO-based toolkit that enables construction of a variety of controllers and props for VR. Using VirtualBricks, users are able to feel the manipulation of a variety of objects, such as a weapon, a slingshot, a fishing rod, and even a DJ board. In addition, the versatility of this system allows re-implementation of artifacts from past work: TASC (Chang et al., 2017) and Haptobend (McClelland et al., 2017). *HapTwist* (Zhu et al., 2019) is a twistable passive device made of Rubik's Twist to create haptic proxies for distinct hand-graspable VR objects, such as a ping-pong paddle, a steering wheel, a machine gun, and a fishing rod. Araujo et al. (2016) implemented Snake Charmer, an encountered-type motorized robotic arm with a prop (cube) attached. As users reach into a virtual object, the arm spatially aligns with that object's virtual representation and provides a physical surface. This matches one or more of the virtual object's shape, texture, and temperature, for the user to touch and feel. Cheng et al. (2018) developed *iTurk*, a foldable and reconfigurable board to render a suitcase, a fuse cabinet, a railing, and a seat in VR. More recently, Ning et al. (2020) constructed *Haptic-go-around* a surrounding platform that deploys props to provide haptic feedback in any direction in VR. The vocabulary of props includes a ball, a gun, a button, and a helm.

2.2.3 Dynamic Passive Haptics

Zenner et al. (2017) constructed *Shifty*, a weight-shifting mechanism to enhance object perception in VR. The system is an array of mechanical actuators that can change its internal weight distribution to automatically provide mixed active and passive haptic feedback. To reduce lack of generality of passive and active haptics devices, the authors introduced the concept of *dynamic passive haptic feedback*—systems that use actuators to change their passive haptic properties (size, shape, texture, weight, position, etc.) without exerting noticeable active forces on people. Prior work that aligns with this definition can be found in the literature, and it traditionally comes in the form of either *force-feedback devices (exoskeletons and wearables)* or *shape- and/or weight-changing systems*.

Force-feedback devices (exoskeletons and wearables)

Researchers have implemented glove-style exoskeletons to provide haptic feedback for grasping of virtual objects. *Dexmo* (Gu et al., 2016) is an exoskeleton that exerts an opposing force to the user's finger tips. However, rather than applying torque control at each individual joint of the exoskeleton, the system forms a rigid body by using a servomotor that shifts stopping blocks to stop the rotation of all joints. Thus, haptic feedback is provided. Grasping a cube and manipulating an arch are a couple of applications of *Dexmo*. *Wolverine* (Choi and Follmer, 2016) is another wearable haptic device that renders a force directly between the thumb and three fingers to simulate grasps. This device provides the sensation of grasping objects (e.g., a mug) through a brake-based system that resists motion of the thumb and the fingers. More recently, Choi et al. (2017) presented *Grabity*, a gripper-style haptic device that provides various types of force feedback to users, such as touching, grasping, gravity, inertia, and rigid stiffness. Using *Grabity*, users are able to touch an object, grasp it, perceive its size, lift it and feel its weight, as well as recognize if they translated it too fast.

Shape- and/or weight-changing systems

Unlike wearable and exoskeleton devices, researchers have also explored handheld reconfigurable systems that render a variety of haptic sensations in VR. *HaptoBend* (McClelland et al., 2017) is a shape-changing prototype that provides haptic feedback for multiple objects. *HaptoBend* allows users to bend the device into 2D plane-like shapes (e.g., a phone or a tablet) and multi-surface 3D shapes (e.g., a flashlight or a hammer). Strasnick et al. (2018) introduced *Haptic Links*, a set of electro-mechanically actuated physical connections capable of rendering variable stiffness between two handheld virtual reality controllers. In this way, the system creates haptic rendering of diverse two-handed objects, such as a rifle, a bow, a trombone, and a couple of pistols. *Transcalibur* (Shigeyama et al., 2019) is a novel reconfigurable handheld controller that renders a 2D shape by by modifying its mass properties on a 2D planar area. Due to this feature, users can perceive various shapes in a virtual environment, for example, holding a sword, a gun, and a crossbow. Zenner and Krüger (2019) implemented *Drag:on*, a device whose appearance and performance are similar to *Transcalibur*'s. The handheld controller provides haptic feedback based on drag, i.e. air resistance and weight shift. By using *Drag:on*, users are able to experience the haptic sensation of manipulating a ratchet, a shovel, and a wagon.

This large body of work classified into active, passive and dynamic-passive haptic feedback presents many possible opportunities to provide the feeling of manipulating physical objects in VR. The work I present in this thesis builds on this work by focusing on determining a vocabulary of how to leverage physical props to perform actions (i.e., gestures) in a VR system. While I think my research could help inform the design of gestures that incorporate active haptics, the scope is currently limited to the use of props with passive haptics. Nonetheless, I intend to provide both a gesture set that could be used by this other work, and to inform future elicitation studies that incorporate the use of physical props (or the feeling of them) when performing a gesture.

2.3 Chapter Summary

In this chapter, I presented an overview of previous work around gesture elicitation studies, gesture-based user interfaces for CAD modelling, and haptic feedback technology in VR. Since the user-defined gesture set for tabletops presented by Wobbrock et al. (2009), researchers have adopted the elicitation technique to design suitable gestures interaction with a diverse group of technologies, such as public displays (Kray et al., 2010), mobile platforms (Gomes et al., 2018; Ruiz et al., 2011), tangible interfaces (Valdes et al., 2014), smartwatches (Arefin Shimon et al., 2016), televisions (Vatavu, 2012), augmented and virtual reality (Nanjappan et al., 2018; Piumsomboon et al., 2013). To measure agreement among participants' input (e.g., gestures) in elicitation studies, researchers have used the agreement index A (Wobbrock et al., 2009) and the agreement rate AR (Vatavu and Wobbrock, 2015). In this study, I applied the latter to measure gesture agreement for each referent, as it captures it realistically (e.g., if no participants agree with each other, agreement score yields zero). I generalized Vatavu and Wobbrock's 2015 AR formula to calculate prop agreement scores for each referent. Tsandilas (2018) emphasized on a series of limitations of the A and AR metrics. For instance, I adhered to two of his recommendations that reduce to effect of these limitations on my work. 3D modelling referents (tasks) are used in the elicitation study presented in this thesis. The list of tasks includes a few found in past work that already explored the use of gestures to perform 3D modelling tasks (Huang and Rai, 2014; Khan and Tunçer, 2019). Finally, researchers, designers, and hapticians have developed a vast number of devices to render or incorporate physical to virtual environment for interaction with virtual objects. These devices provide active (Abtahi and Follmer, 2018; Benko et al., 2016; Choi et al., 2018; Delazio et al., 2018; Sinclair et al., 2019; Siu et al., 2018; Teng et al., 2019; Whitmire et al., 2018), passive (Araujo et al., 2016; Arora et al., 2019; Chang et al., 2017; Cheng et al., 2018; Huang et al., 2020; Muender et al., 2019; Strandholt et al., 2020; Yan et al., 2018; Zhu et al., 2019), or dynamic passive (Choi and Follmer, 2016; Choi et al., 2017; Gu et al., 2016; McClelland et al., 2017; Shigeyama et al., 2019; Strasnick et al., 2018; Zenner and Kruger, 2017; Zenner and Krüger, 2019) feedback. The work I present in this thesis builds on top of this large body of devices by determining a vocabulary of how to leverage physical props to perform actions (i.e., gestures) in a VR system.

Chapter 3

Study Method

3.1 Methodology

For this study, I chose an elicitation approach (Wobbrock et al., 2009). User experience designers and developers have argued that users should not be made to learn artificial gestural languages to interact with systems or interfaces because it is not a natural approach (Malizia and Bellucci, 2012). Instead, by means of a natural interface, users should be put in a position to employ the same gestures they employ to interact with their environment (e.g. objects), just as evolution and education has taught us. Unfortunately, gestural interfaces often rely on experts and professional designers, instead of considering spontaneity of users (Malizia and Bellucci, 2012) and understanding user preferences (Morris et al., 2010). Consequently, they are forced to adopt an expert-design or technology-driven approach to interaction (e.g. the mouse).

The elicitation technique, which includes the participation of non-technical users in the design of interactive systems, represents a suitable solution for this problem and the one that I adopted for this research. The study I present in this thesis is aligned with Malizia et al.'s goal of gestural interfaces (Malizia and Bellucci, 2012), which is to provide users with a natural and intuitive approach to interact with an eventual VR gestural interface so that, ideally, no significant learning or training for specific gesture-referent mappings is needed. People are knowledgeable of physical props (objects) because interacting with them is an inevitable part of their lives (Alzayat et al., 2014, 2017, 2019). For instance, by incorporating props to this study, users do not need training and they are able to expand the gesture vocabulary that can be achieved with hands and fingers only.

The experimental study I conducted for this research follows Wobbrock et al.'s method (Wobbrock et al., 2009) to design surface gestures for tabletop referents, which has led to gesture sets in other domains (Arefin Shimon et al., 2016; Kray et al., 2010; Nanjappan et al., 2018; Piumsomboon et al., 2013; Ruiz et al., 2011; Vatavu, 2012):

1. The effect of a referent being completed in VR is displayed.
2. Participants are asked to choose a single or multiple of props from a group of 95 (75 props and 20 spare LEGO bricks) randomly arranged and numbered on a grid.
3. Participants are asked to perform a manipulative gesture with the chosen prop(s) that would complete the displayed referent.

Figure 3.1 shows the three stages of the elicitation process I implemented in this study; participants repeated these steps for a total of 20 VR referents.

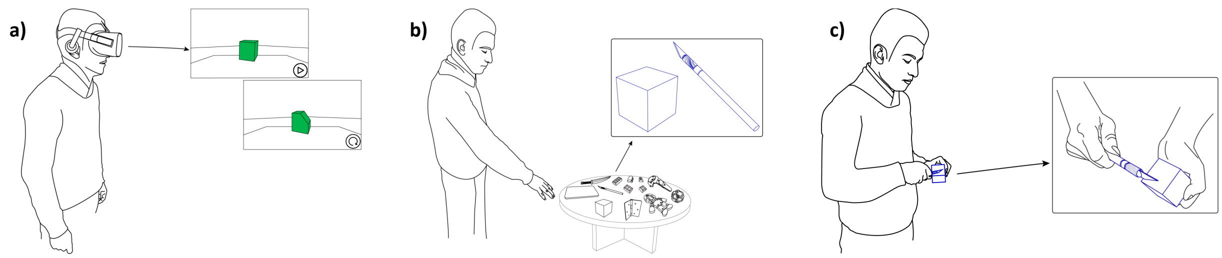

Figure 3.1: Gesture elicitation study with props that follows a previous methodology (Wobbrock et al., 2009). A participant: a) watches the completion of a referent, b) chooses a set of physical props, and c) performs a gesture with the chosen prop(s) that would complete the referent.

3.1.1 Selection of VR Referents

To ground the work I present in this research, and to represent a rich vocabulary of referents, I used CAD-like and open-world game commands. Similar to previous gesture-based user interfaces for CAD 3D modelling (Huang and Rai, 2014; Khan and Tunçer, 2019; Khan et al., 2019; Nanjappan et al., 2018), I included simple tasks like “rotate”, “scale”, “extrude”, “copy-paste” (duplicate), “make a cone”, and “zoom”. I expanded this vocabulary of referents with some more complex ones, like “Twist”, “Bend”, “Bevel”, “Extruded cut” (puncture), “Stretch”, “Deform a model”, “Assemble”, “Colour”, and “Tighten bolt and nut”, which are typically found in standard CAD modelling software and VR CAD modelling applications (e.g., AutoCAD, SolidWorks, Blender, MakeVR Pro, Mesh Maker VR, to name a few). In addition, I surveyed two VR open-world games, like Job Simulator and Rec Room, to incorporate abstract commands like “Turn on a light”, “Open a hinged door”, “Close two sliding doors”, “Change the time”, “Join two spheres”, and “Hang a piece of art on a wall”. Table 3.1 shows a description of the effect of each VR referent included in this elicitation study, and Figure 4.8 and Figure 4.9, in chapter 4, illustrate them in the format of beginning (on the left) and end (on the right). While I assigned labels (names) to the referents, I did not show these to participants in the study to avoid bias (they instead viewed the effect of each of them in a VR headset). The effect of all referents was animated while visualizing them in the VR headset, except for "Duplicate", "Puncture" and "Turn ON", where the effect was presented in a binary way (i.e., go straight from beginning to end).

3.1.2 Selection of Physical Props

Deciding on the physical props to include in this study presented two challenges. What would be the best vocabulary of props? And how large would this vocabulary need to be? Instead of arbitrarily answering these questions, I looked at the sandtray therapy (Homeyer, 2011), a type of psychological therapy of play in which the use of miniatures and figurines in a tray of sand acts as a vehicle for young patients to express feelings and emotions. In other words, patients are given the opportunity to freely associate physical objects with their feelings and emotions, and build a representative scene on the sand. In her manual (Homeyer, 2011), Homeyer suggest categories of objects, such as nature, people, tools, fantasy, and transportation, as well as a helpful guideline for achieving a rich vocabulary of objects. This guideline states that having a limited collection represents a limited vocabulary. On the other hand, having a massive set of items increases cognitive load. Considering these recommendations, I selected a vocabulary of approximately 100 props; 95, to be exact.

Referent		Description
A	Join	a straight line is traced to join two green spheres sized differently
B	Twist	an initially brown straight bar is twisted
C	Bevel	a red cube is beveled from one of its edges
D	Bend	a dark flat rectangular plate is bent
E	Hang	a painting initially on the white ground is hung on a gray virtual wall that includes a second painting
F	Open	a brown hinged door is opened
G	Close	two brown sliding doors are closed
H	Zoom	the view of the camera is zoomed to the model of a virtual house
I	Reshape	a virtual object initially presented as a purple cylinder transforms into a cone
J	Rotate	a palm tree is rotated 180 degrees
K	Stretch	a blue cylinder is stretched from both ends
L	Extrude	a green square cross-sectional area within a blue-coloured cube is extruded
M	Puncture	a circular volume is removed from the center of an initially green solid hexagonal prism
N	Turn ON	a lamp post located in the darkness is turned ON, emitting a reflection on the virtual ground
O	Run time	time is accelerated to make the environment around a virtual house go from daytime to dusk
P	Deform	the top section of a vertically oriented red straight bar is stretched and expanded
Q	Duplicate	a single small staircase is duplicated
R	Assemble	two flat identical blocks contrarily oriented, one green and one blue, are assembled together
S	Colour	an initially blue coloured cube changes to yellow
T	Tighten	a bolt is inserted in a nut

Table 3.1: Labels and descriptions of the 20 VR Referents included in this elicitation study. These referents were extracted from prior work (Huang and Rai, 2014; Khan and Tunçer, 2019; Khan et al., 2019; Nanjappan et al., 2018), CAD 3D modelling software, VR CAD modelling applications, and Open-World games. Figure 4.8 and Figure 4.9 provide an illustration of each.

I picked a first group of props by following categories from sandtray therapy. Therefore, figures like a fox, a flower, a ghost, a wrench, a magnet, a sword, a spaceship, among others, were selected. In a second group, I included props used or explored in prior research (Arora et al., 2019; McClelland et al., 2017; Muender et al., 2019; Weichel et al., 2014; Yan et al., 2018), and later I complemented the vocabulary with abstract props that seemed interesting to explore for CAD modelling and open-world game referents, such as a cube, a plane, and a sphere.

I used an *Eden260v* 3D printer to fabricate 50 of the 95 props, 25 of flexible material and 25 of rigid material. 8 props were LEGO-built specific assemblies, and 20 were spare bricks that would allow participants to construct a desired prop that was not available. The remaining 13 were retail-manufactured objects purchased in museums and hardware stores. I also included 3 handheld controllers, from Vive, Oculus and Nintendo Switch.

I recognized that using free association was beneficial for this elicitation study because it unveils information that matters to people. In the sandtray therapy, inspiration tends to start with associations to the presented repertoire of objects (Hancock et al., 2010), which leads to a scene of what patients feel or think. In this study, free association and inspiration occurred when participants were asked to choose props from a vast vocabulary of props to then perform a gesture with them. Physical props have a relatively fixed purpose. However, the way that people freely do manipulation with them becomes part of the person's means of articulating their own gestures with the picked props. Figure 3.2 shows the vocabulary of objects used in this elicitation study.

Figure 3.2: The vocabulary of 95 physical props used for this elicitation study. 50 props were 3D printed, 25 with flexible material and 25 with rigid material. 28 props were LEGO-brand toys, 8 specific models and 20 spare bricks. 13 retail-manufactured objects and 3 handheld gaming controllers from Vive, Oculus and Nintendo Switch.

3.1.3 Participants

I recruited a total of 21 participants from the university campus, ages 19 to 32 ($M = 25.4$, $SD = 3.7$), 7 identified as women and 14 as men; each was paid \$15. Upon the completion of the study, they completed a validated questionnaire about skills and experience with VR, 3D printing, CAD, and gaming. Five participants in total had previous experience with VR environments and nine had experience operating 3D printers.

3.1.4 Apparatus

Participants watched the execution of the referents using an HTC Vive headset connected to an MSI VR-ready laptop. I used a compact camera facing the participant's seat to record gestures, and a *Fusion* GoPro camera placed near the ceiling of the study room to capture the selection process of the props. Figure 3.3 and Figure 3.4 illustrate the room where the study took place. I modelled the virtual objects using the Blender 3D computer graphics software, and I animated the referents in the Unity 3D game engine.

3.1.5 Procedure

Upon arrival, I compensated participants and asked them to complete a consent form. Then I gave them a detailed verbal description of the experiment. I did not reveal the VR referents, but I showed a sample (not part of the actual referents for the study) on a monitor. For each participant, I randomly arranged the physical props on a table at the center of the room. I pointed out the different fidelities available to the participant: 3D-printed rigid and flexible, LEGO-built figures/bricks, and manufactured retail objects. After that, participants took approximately five minutes to familiarize themselves with the props. I then supervised and, upon request, supported participants for appropriate and comfortable wearing of the VR headset.

The software presented the 20 referents shown in Figure 4.8 and Figure 4.9 in random order. As in Wobbrock et al. (2009), for each referent, participants performed a gesture with physical props. I did not impose restrictions on the number of props to pick, and

Figure 3.3: Two participants performing in the study. On the left, participant 4 picking up a retail-manufactured flute for gesture demonstration. On the right, participant 10 assembling a model with LEGO-brand bricks that was not available.

there was no time limit. The study concluded with a computer-administered demographics questionnaire.

Figure 3.4: Study room setting

3.2 Chapter Summary

I designed a gesture elicitation study using physical props to investigate the research questions stated in section 1.1. By adopting this technique, I aimed to develop a natural and interactive VR interface where people are able to employ gestures with props that they employ to interact with their environment, instead of forcing them to learn artificial gestural languages. The study follows Wobbrock et al.'s (2009) methodology by asking 21 participants from the university campus to choose from 90+ props to perform 3D manipulation for 20 different CAD-like and open-world referents (actions). I extracted referents from previous work (Huang and Rai, 2014; Khan and Tunçer, 2019; Khan et al., 2019; Nanjappan et al., 2018), from 3D modelling software, from two VR modelling applications, and from two open-world games. I created the vocabulary of props by following categories (e.g., nature, people, fantasy) recommended by sandtray therapy (Homeyer, 2011), by extracting some from prior work (Arora et al., 2019; McClelland et al., 2017; Muender et al., 2019; Weichel et al., 2014; Yan et al., 2018), and by incorporating some abstract props (e.g., cube, plane) that seemed interesting to explore. I used four materials: 3D-printed flexible props, 3D-printed rigid props, LEGO-brand figures and bricks, and retail-manufactured objects. In the the following chapter, I provide a detailed description of the results of this elicitation study.

Chapter 4

Results

In this chapter, I propose a participant-defined set of gestures using physical props based on the data collected from my study, descriptions of patterns found in the elicited gestures, and two sets of agreement scores: one for gestures and one for physical props.

4.1 A Participant-Defined Gesture Set

Similar to Piumsomboon et al. (2013) and Nanjappan et al. (2018), I grouped participant responses into a set of gestures for each referent. For each referent, I grouped gestures that exhibited the same pattern together, and I chose the group with the largest consensus to be the representative gesture for the referent. This differs slightly from elicitation studies for 2D systems, such as tabletops (Wobbrock et al., 2009), mobile devices (Ruiz et al., 2011), smartwatches (Arefin Shimon et al., 2016), public displays (Kray et al., 2010), and televisions (Vatavu, 2012), as I loosened constraints from "gestures must be identical" to "gestures must be similar" or "gestures should show the same pattern".

I recognized a total of 197 unique gestures from the 420 gestures elicited from 21 participants for each of the 20 VR referents. I found that the most natural and expressive way of capturing each gesture was as a context-free grammar (CFG), which described the gesture, which props had been chosen for that gesture across all the participants who made that gesture, and how those props fit into the description.

Figure 4.1 illustrates the CFG representation for referent C (Bevel), where a rectangular prism's corner is beveled. The most used gesture (five participants) was *Cut*, where participants took a cube and cut a bevel into it with either a 3D-printed knife or an X-acto knife. Here, the virtual cube is presented as the variable "`PROP`", which represents the object being beveled, and represent the knife and X-acto knife as "`CUTTING_PROP`", which affords the action of cutting:

```
G → Cut section of PROP with CUTTING_PROP
PROP → Cube | Imaginary
CUTTING_PROP → Knife | X-acto Knife
```

Participants doing the *Cut* gesture either held a cube to represent the virtual prism, or held their hand as if they were holding an imaginary prism, then cut the bevel with the `CUTTING_PROP`.

The next most used gesture for referent C (Bevel) was *adhere* (two participants), where participants took either a cube or a sofa, and placed a flat object (either a 3D-printed hinge in its flat position or a 3D-printed plane) on the corner of it:

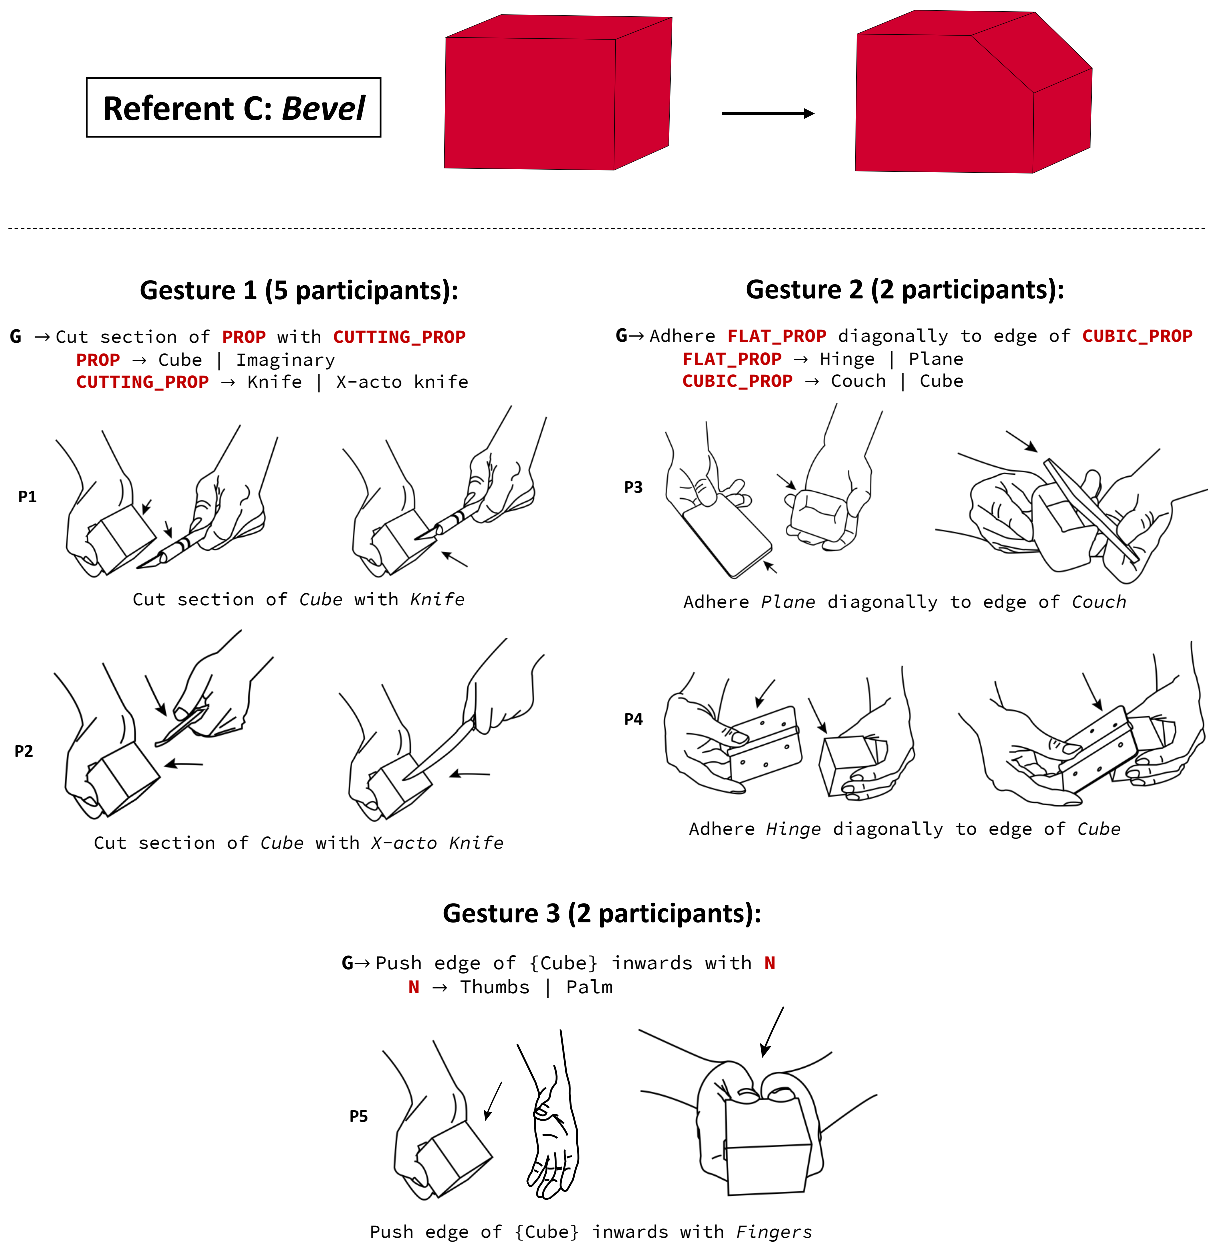

Figure 4.1: Three of the thirteen gestures identified for referent C: *Bevel*. Each gesture is articulated as a context-free grammar, showing what physical props can be used with that gesture.

```
G → Adhere FLAT_PROP diagonally to edge of CUBIC_PROP
FLAT_PROP → Hinge | Plane
CUBIC_PROP → Cube | Couch
```

Another gesture for referent C was *push*, where participants simply attempted to push one of the top edges of a 3D-printed cube with either their two thumbs or their palm:

```
G → Push {Cube} edge of Cube inwards with N
N → Thumbs | Palm
```

For each referent (here, Bevel), I chose the most demonstrated gesture (here, *Cut*) for the proposed gesture set. Figure 4.8 and Figure 4.9 illustrate the complete proposed set of gestures with their corresponding referents.

4.2 Gesture Agreement

With previous gesture studies (Arefin Shimon et al., 2016; Kray et al., 2010; Nanjappan et al., 2018; Piumsomboon et al., 2013; Ruiz et al., 2011; Vatavu, 2012; Wobbrock et al., 2009), participants would suggest one gesture per referent, so agreement was calculated for gestures alone. In this elicitation study, participants both offered a gesture and a selection of props, and may have based their actions on either a gesture or a set of props. Therefore, I analyzed both gesture agreement and prop selection.

For gesture agreement, I used Vatavu and Wobbrock's formula (2015):

$$AR_{gesture}(r) = \frac{\displaystyle\sum_{P_i \subseteq P} \frac{1}{2}|P_i|(|P_i| - 1)}{\frac{1}{2}|P|(|P| - 1)} \tag{4.1}$$

This formula is derived from pairwise comparison of participant gestures. For a given referent, the agreement is calculated by looking at each pair of participants, saying that the two participants agree if they performed the same gesture (similarity of gestures is 1), and do not agree if they performed different gestures (similarity of gestures is 0). For instance, this formula defines agreement as the number of pairs of participants with each other divided by the total number of participants that could agree (Vatavu and Wobbrock, 2015).

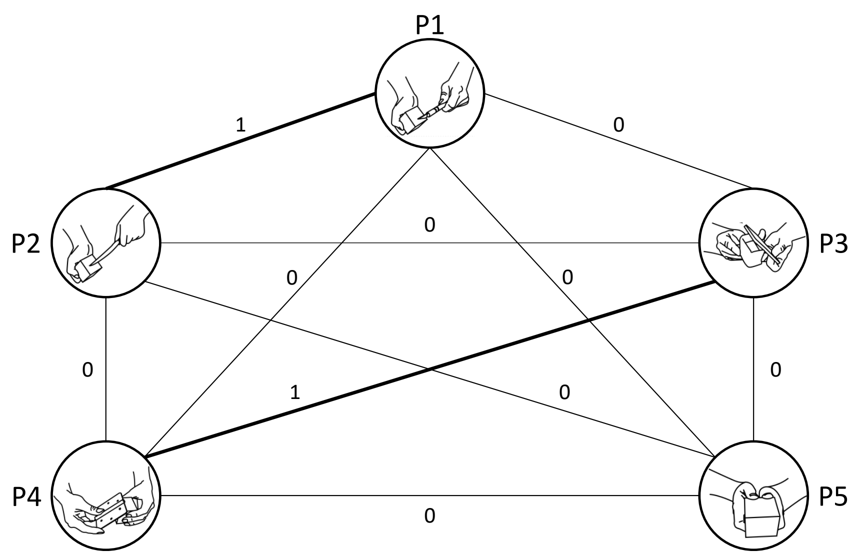

Figure 4.2: Complete graph that captures gestures with physical props performed by 5 participants for referent C: *Bevel*. Gestures are considered "one-element" sets and similarity is calculated with one another, where a 1 means "same gesture" and 0 means "distinct gesture". No partial similarity is considered. Similarity between two sets is the weight of the edge that joins them. Agreement score is computed by summing the similarity of each pair of sets compared, and dividing by the total number of pairs that can be identical.

Figure 4.2 illustrates pairwise gesture agreement for the five participant *Bevel* examples introduced in Figure 4.1. Pairwise agreement can be represented as a complete graph, where nodes represent participants, and edge weights represent the similarity between those participants' gestures, i.e., either 0 (if they are different) or 1 (if they are the same). The agreement for a given referent is the sum of the weights of those edges divided by the total number of edges: $\frac{1}{2}n(n-1)$, where n is the number of participants. This score is thus bounded between 0 and 1, as if everyone agreed and all edges were equal to 1, then the equation would be the number of edges divided by the number of edges. By incorporating the rest of the participants and their gestures for referent C (Bevel) in Figure 4.2, the graph becomes a 20$^{\text{th}}$-grade complete graph. For this referent, a total of 13 different gestures were recognized, whose sizes (i.e. number of participants that performed them) were $5, 2, 2, 2, 2, 1, 1, 1, 1, 1, 1, 1,$ and 1. For instance, using Equation 4.1, the gesture agreement score for *bevel* is:

$$AR_{gesture}(bevel) = \frac{\frac{5\cdot4}{2} + \frac{2\cdot1}{2} + \frac{2\cdot1}{2} + \frac{2\cdot1}{2} + \frac{2\cdot1}{2}}{\frac{21\cdot20}{2}} = \frac{14}{210} = 0.067 \tag{4.2}$$

Table 4.1 shows the number of unique gestures recognized for each referent as well as the sizes of each. This information is used as input to Equation 4.1 to obtain a gesture agreement score for each referent. Figure 4.3 illustrates these scores, in descending order.

While previous formalization of agreement scores (Vatavu and Wobbrock, 2015) did not represent agreement as a graph, I found this representation to be useful when reasoning about agreement for prop selection.

Referent	Number of Gestures	Sizes	Agreement Score
Tighten	5	17, 1, 1, 1, 1	.648
Assemble	6	15, 2, 1, 1, 1, 1	.505
Bend	7	15, 1, 1, 1, 1, 1, 1	.500
Duplicate	7	15, 1, 1, 1, 1, 1, 1	.500
Rotate	7	15, 1, 1, 1, 1, 1, 1	.500
Hang	8	14, 1, 1, 1, 1, 1, 1, 1	.433
Twist	6	12, 5, 1, 1, 1, 1	.362
Run time	10	11, 2, 1, 1, 1, 1, 1, 1, 1, 1	.267
Close	10	10, 2, 2, 1, 1, 1, 1, 1, 1, 1	.224
Join	8	9, 4, 3, 1, 1, 1, 1, 1	.214
Stretch	11	9, 2, 2, 1, 1, 1, 1, 1, 1, 1, 1	.181
Puncture	9	7, 5, 3, 1, 1, 1, 1, 1, 1	.162
Open	9	6, 5, 3, 2, 1, 1, 1, 1, 1	.138
Reshape	9	5, 4, 4, 2, 2, 1, 1, 1, 1	.114
Zoom	13	7, 2, 2, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1	.110
Colour	14	7, 2, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1	.105
Turn ON	15	7, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1	.100
Bevel	13	5, 2, 2, 2, 2, 1, 1, 1, 1, 1, 1, 1, 1	.071
Extrude	14	4, 3, 2, 2, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1	.052
Deform	16	4, 2, 2, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1	.038

Table 4.1: For each VR referent, number of unique gestures recognized from the 21 participants, number of participants that performed each of these unique gestures, and gesture agreement score are reported.

4.3 Prop Agreement

In this study, participants selected props in addition to gestures. I was interested in observing how much prop selection agreed across participants for a given referent. However, prop agreement is more complicated than gesture agreement: participants selected *sets* of physical props, which meant that there could be partial agreement between participants.

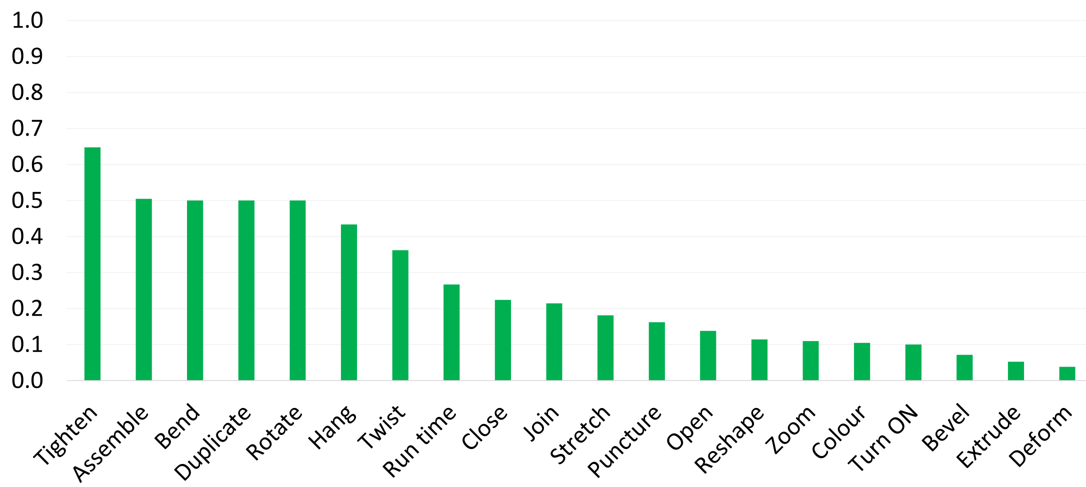

Figure 4.3: Gesture agreement scores of the 20 VR referents sorted in descending order.

To analyze prop agreement, I proposed a generalization of Vatavu and Wobbrock's (2015) agreement rate formula. Instead of using simple equality between pairs, I used different measurements of set similarity (Egghe and Michel, 2002), here denoted by the function SIM:

$$AR_{prop}(r) = \frac{\sum_{i=0}^{P-1} \sum_{j=i+1}^{P} \mathrm{SIM}(P_i, P_j)}{\frac{1}{2}|P|(|P|-1)} \tag{4.3}$$

Equation 4.3 is defined as the sum of the similarity of each pair of sets by the total number of pairs of sets that could be identical, where P is the total number of people. P_i and P_j group all possible pairs of sets for r.

Equation 4.3 is structurally similar to the simple gesture agreement formula (Equation 4.1), but Equation 4.3 offers more flexibility in comparing participants' choices. Thus I constructed a connected graph to represent prop selection similarity, where nodes represent participants, and edge weights represent the similarity between those participants' gestures (a real number between 0 and 1 inclusive).

In the graph representation (Figure 4.4) for the five participant *Bevel* examples I introduced in Figure 4.1, each edge is weighted by SIM, the set similarity between their sets of selected props, rather than just 1 (same) or 0 (different). If SIM is bounded between 0 and 1, then the overall agreement (sum of the weights of all edges divided by the number of edges) is also bounded between 0 and 1. To be analogous to gesture similarity, SIM must also be 0 if there are no props in common, or 1 if they are completely equal (analogous to choosing different gestures or the same gesture).

Figure 4.5 shows the three set similarity metrics that have been identified for SIM: the Jaccard Index (*Jaccard*), the Overlap coefficient (*Overlap*) and the Sørensen-Dice coefficient (*Sørensen*). These three metrics can all satisfy the requirements for a similarity score, and if they are sets of size one (e.g., a single gesture per participant), they collapse into a binary 1 for equality and 0 for inequality. In other words, all three metrics are equivalent to Vatavu et al.'s score when comparing a single agreement.

Figure 4.6 shows the prop agreement score for all three metrics (Jaccard, Sørensen, Overlap) for each referent. The three metrics give similar results for each referent. Jaccard is consistently the lowest agreement, Overlap is the highest, and Sørensen is in between the two. In the analysis made in this thesis, the Jaccard Index is used as the most conservative metric for identifying high agreement.

I wrote a script (section A.1) in the Matlab computing environment in order to solve Equation 4.3 for referent C: Bevel. In the first section, I inputted the number of pairs of

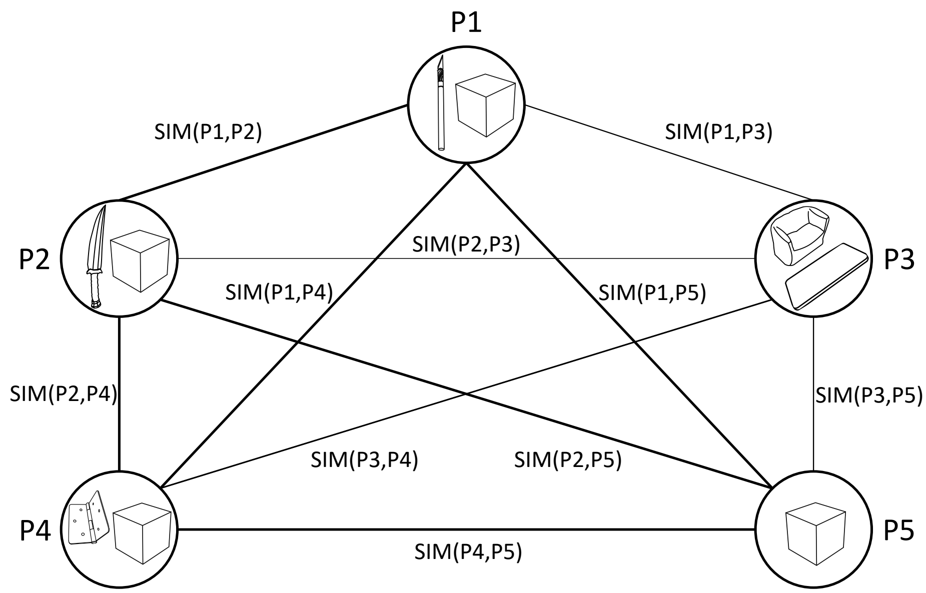

Figure 4.4: Complete graph that captures props sets of 5 participants for referent C: *Bevel*. Similarity of props sets is measured, which is the weight of the edges that joins them. Each measure of set similarity is between 0 and 1, with 0 indicating no shared props and 1 indicating completely identical prop selections. Agreement is computed by summing the similarity of each pair of sets compared, and dividing by the total number of pairs that can be identical.

participants as well as the prop sets each person picked. Thus, in the second section, I calculated the prop agreement score for this referent.

To compute the prop agreement scores for the rest of the referents, I modified the script accordingly. Figure 4.7 reports these agreement scores, sorted in descending order. Both prop and gesture agreement scores can also be seen individually for each referent in Figure 4.8 and Figure 4.9.

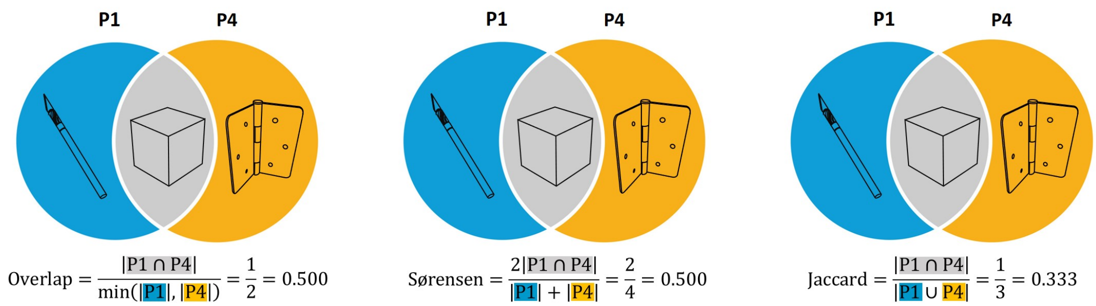

Figure 4.5: Venn diagrams showing a pairwise comparison between prop sets of P1 and P4. Three measures of set similarity are proposed to evaluate partial agreement (SIM(P1,P4)): Overlap, Sørensen, and Jaccard. Overlap is the most optimistic, leading to the highest agreement. Jaccard is the most pessimistic, leading to the lowest agreement. Sørensen prevails between the other two (or equal to one of them). In this pairwise comparison, Sørensen and Overlap are the same.

4.4 Classification of Gestures with Props

I classified the 197 unique gestures elicited in this study into 3 categories: *Act*, *Leverage*, and *Adapt*. From a qualitative analysis with emphasis on recognizing similar or identical patterns in gestures with props, I found that participants performed a gesture from at least one of these categories (i.e., overlap exists). I also found that, within each category, several possibilities can be observed.

In the first category, *Adapt*, I grouped gestures where participants physically acted on props with either their hands or another prop to represent the effect of a referent. Gesture possibilities within this category include a) altering the properties of props (e.g., twist, fold, squeeze), b) changing the position and/or orientation of props in a three-dimensional space (e.g., move up, rotate, throw), or c) operating movable parts of props (e.g., open pliers, fold a hinge).

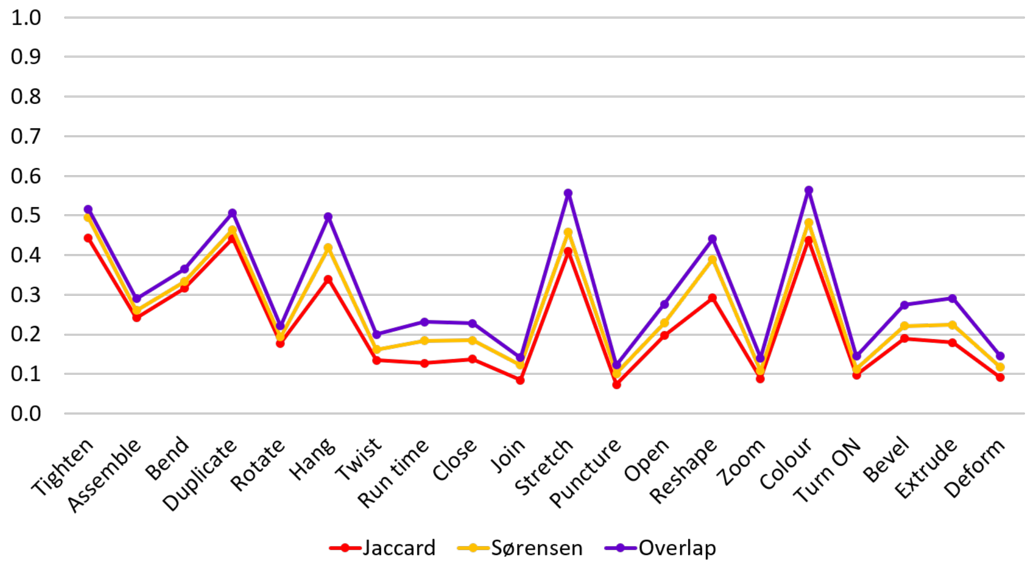

Figure 4.6: Gesture set with physical props for referents A to J. For each referent, the preferred set of props and the preferred gesture are presented; the initial state of the most common prop-gesture interaction is shown on the left side, and the end state is shown on the right.

In the second category, *Leverage*, I clustered gestures where participants leveraged physical properties of props to represent the effect of a referent. This can be done by a) adhering a prop (or a section of it) to another prop (e.g., adhere a vertex section of a prop to the top end of a candle), or b) by simply showing a comparison of a property (e.g., color, shape) of two or more props.

In the last category, *Adapt*, I grouped gestures made with LEGO-brand pieces. This category includes gestures where participants adapt an assembly of bricks to represent the effect of a referent. Gestures in the Adapt category consist of a) disassembling an initial assembly and re-assembling it differently, or b) simply assembling a single model.

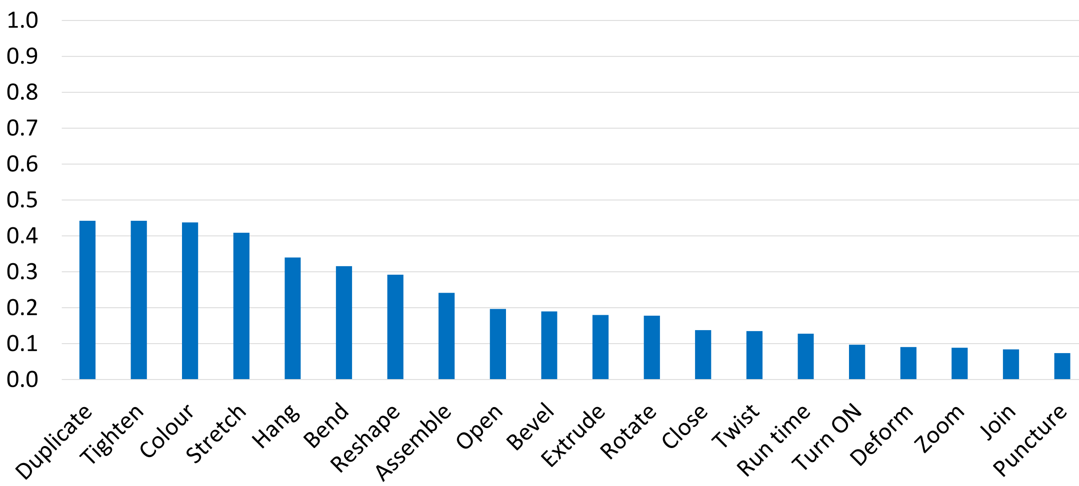

Figure 4.7: Prop agreement scores of the 20 VR referents sorted in descending order.

Figure 4.10 illustrates the three categories of gestures with an example for each possibility that may occur within these categories. The gestures I included in this figure are expressed using a CFG.

4.5 Chance agreement and agreement confidence

Tsandilas (2018) argued that in the presence of any source of bias, chance agreement alters raw agreement scores, thus providing misleading information about how participants agree or disagree. On the other hand, Gwet (2014) and Krippendorff (2004) emphasized that the magnitude of a raw agreement score cannot be interpreted if confidence intervals are not provided. For these reasons, I estimated a common chance agreement for gestures across all referents, and I computed confidence intervals for gesture and prop agreement scores of each referent.

To estimate chance agreement and chance-corrected agreement indices, Tsandilas (2018) suggests using inter-rater reliability measures. In the context of elicitation studies, the terminology of inter-rater reliability can be mapped as follows: study participants become raters, referents become items, and signs (i.e., gestures) become categories. In this study, I adopted the Fleiss' (1971) κ_F coefficient:

$$\kappa_F = \frac{\rho_a - \rho_e}{1 - \rho_e} \tag{4.4}$$

where ρ_a is the proportion of items on which both raters agree (i.e., raw agreement score calculated with Equation 4.1), and ρ_e is the chance agreement (i.e., agreement that would have occurred by chance).

For the chance agreement term ρ_e, Fleiss uses an extended version of Scott's (1955) π for multiple raters:

$$\rho_e = \sum_{k=1}^{q_-} \pi_k^2, \qquad \pi_k = \frac{1}{m} \sum_{i=1}^{m} \frac{n_{ik}}{n_i} \tag{4.5}$$

where m is the total number of items, n_{ik} is the number of ratings for item i having category k, and n_i is the total number of ratings for item i. The term π_k estimates the probability that a rater classifies an item into category k, based on how many times this category has been used across the entire elicitation study. π_k takes bias into account, as it does not assume that each category (gesture) has the same probability to be shown by a rater. However, it assumes that all raters share the same preferences for categories. The overall gesture agreement score (Figure 4.3) of the elicitation study of this thesis was $AR = 0.261$. For the 182 unique gestures recognized I computed $\rho_e = 0.016$, and therefore,

$\kappa_F = \frac{.261-.016}{1-.016} = .249$, reflecting that the overall gesture agreement is affected minimally by chance agreement. In section 6.1, I briefly discuss the reason(s) chance agreement and chance-corrected agreement indices for props are not calculated, as more research is needed on this matter.

Re-sampling methods, such as *jackknifing* (Quenouille, 1949) and *bootstrapping* (Efron, 1992), can be used to calculate variance of agreement scores. Tsandilas (2018) recommends that the jackknife method is a good approach for estimating within-participants agreement differences, while differences between independent groups of participants can be estimated with bootstrapping. Therefore, I used the jackknifing approach to calculate confidence intervals of gesture and prop agreement scores AR for each referent.

$$\upsilon_{jack} = \frac{n-1}{n}\sum_{i=1}^{n}(\hat{AR}_i - \hat{AR})^2, \tag{4.6}$$

where n is the number of participants, $\hat{AR}$ is an estimate based on the full set of observations (i.e., agreement score obtained with Equation 4.1 or Equation 4.3), and $\hat{AR}_i$ is an estimate when leaving out the ith participant. The square root of the variance υ_{jack} gives the standard error: $SE = \sqrt{\upsilon_{jack}}$. Assuming that agreement scores follow a normal distribution, the $(1-\alpha)\%$ confidence interval is $\hat{AR} \pm SE_{jack} \times q$, where $q = q(1-\frac{\alpha}{2}, n-1)$ is the $(1-\frac{\alpha}{2})$-quantile of Student's t-distribution with $n-1$ degrees of freedom. Tsandilas (2018) ran a series of Monte Carlo experiements to examine this assumption.

In Figure 4.8 and Figure 4.9, 95% confidence intervals can be observed for each VR referent. An algorithm written in the Matlab computing environment to obtain confidence intervals for referent C: Bevel can be found in section A.2.

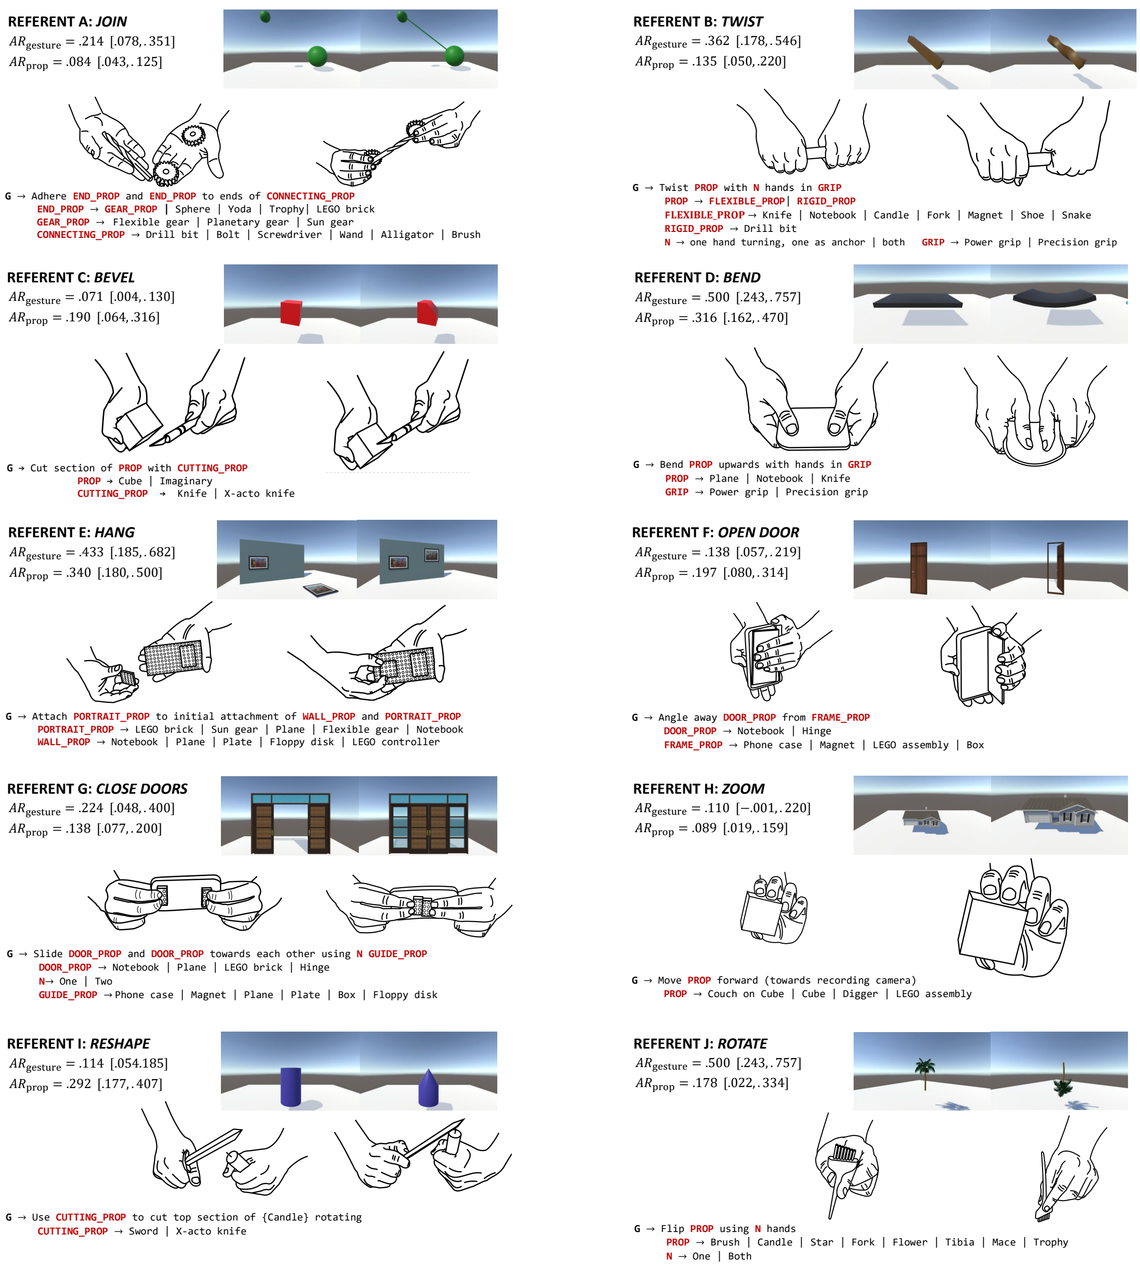

Figure 4.8: Gesture set with props for referents A to J. For each referent, the preferred set of props and the preferred gesture are presented; the initial state of the most common prop-gesture interaction is shown on the left side, and the end state is shown on the right.

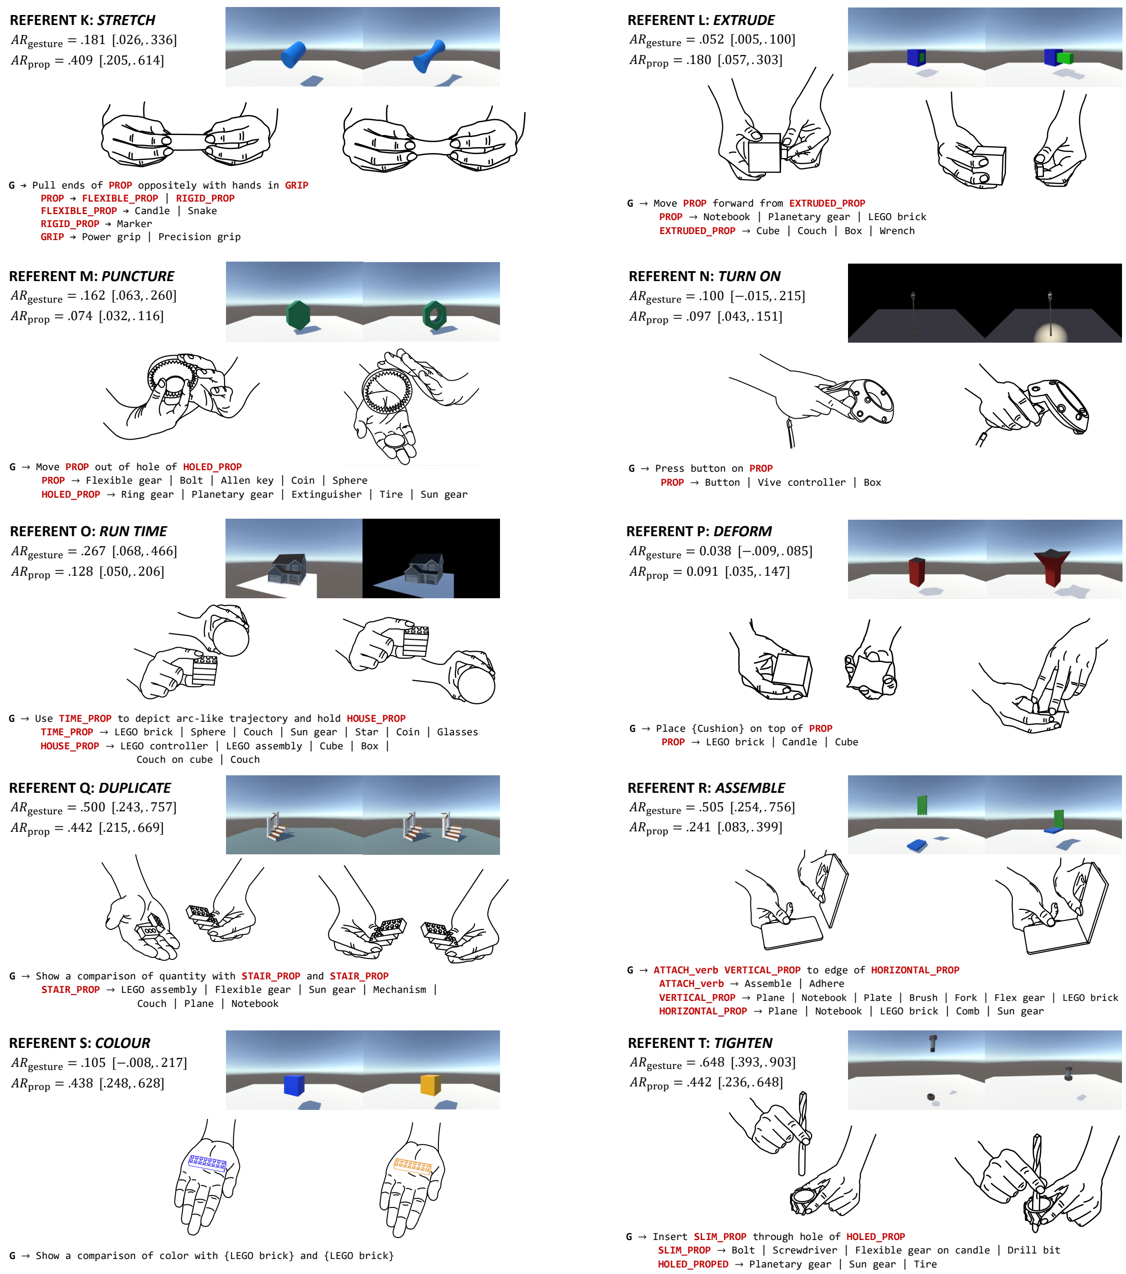

Figure 4.9: Gesture set with props for referents A to J. For each referent, the preferred set of props and the preferred gesture are presented; the initial state of the most common prop-gesture interaction is shown on the left side, and the end state is shown on the right.

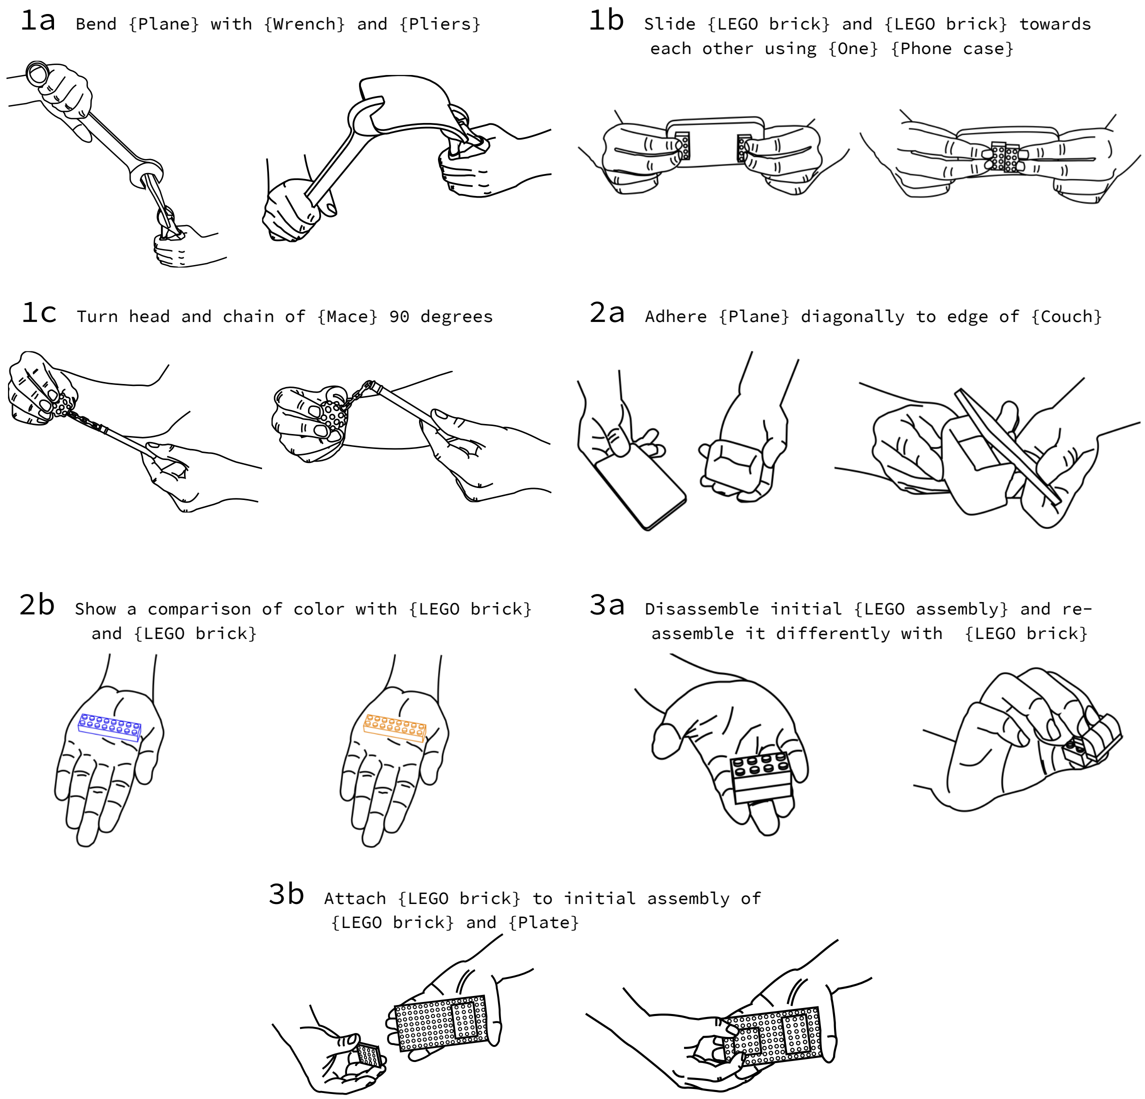

Figure 4.10: Three different categories of gestures were created from qualitative observations: 1) Act, 2) Leverage, and 3)Adapt. Each category groups several gesture possibilities. For a given referent, a participant can: 1a) alter the physical properties of a prop with hands or with more props, 1b) change the position of props in a 3D space, 1c) operates movable parts of a prop, 2a) adhere a prop, or a section of it, to another one, 2b) show a comparison of property of two or more props, 3a) disassemble an initial assembly and re-assemble it differently, and/or 3b) assemble a single model to represent the effect.

Chapter 5

Discussion

The work I present in this thesis has implications for VR systems that use physical props and future elicitation studies.

Props were selected for shape and affordances. When selecting props, participants chose objects that resembled the objects in the virtual world and had affordances (Gibson, 1986; Norman, 2002) for the actions they wanted to take. For example, in referent C (*Bevel*), participants chose a cube as a stand-in for the cube in the virtual world, and a Knife or an X-acto Knife because both afforded cutting the corner away to create the bevel. In referents B (*Twist*) and D (*Bend*), participants picked a single prop that resembled the shape of the virtual object and that was flexible to be twisted or bent; one exception was the drill bit, a rigid object that resembled the action in its spiral shape. In referent E (*Hang*), participants used LEGO-brand pieces because they afforded snapping together.

In some cases, the affordances chosen were metaphorical. In referent N, *Turn On*, participants chose objects that supported pressing (i.e., as a button) to turn on the light. Because objects were chosen for their affordances, and each object had multiple affordances (from its shape, meaning, structure, or material), objects could serve different purposes in different gestures: the knife, for example, afforded cutting in *Bevel* and *Reshape*, but bending in *Twist* and *Bend*.

A small set of props can handle many gestures. A small number of props provided coverage for a large number of gestures. As seen in Figure 3.2, only 48 of the 95 props were chosen; 18 of these were only used for a single gesture. In addition, most prop variables had on average 4.2 props that were selected to fulfill a certain use. This means that a small set of props could support a large set of gestures. In fact, 12 props and a group of spare LEGO bricks would be sufficient to handle all 20 gestures: sun gear, flexible gear, x-acto knife, cube, candle, phone case, notebook, plane, Vive controller, tire, drill bit, and cushion would enable all gestures.

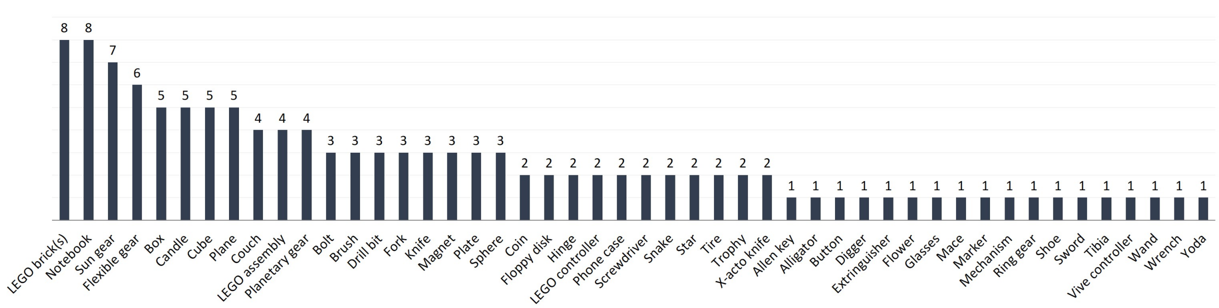

Figure 5.1: Number of gestures from Figure 4.8 and Figure 4.9 for which each prop mentioned the context-free grammar can be used for.

Props were preferred over controllers. Participants rarely selected controllers. While two VR and one Nintendo Switch controllers were available for every referent, only a few participants selected them for seven gestures for referents N (*Turn on Light*), L (*Extrude*), H (*Zoom*), and C (*Bevel*); only one of these gestures was the most-performed for its referent (specifically, for *Turn On Light*). When participants chose the controller, it was not the only prop that could fulfill that purpose. In referent N (*Turn on Light*), the key affordance was that it could be pressed like a button, perhaps as a metaphorical light switch. In referent L (*Extrude*), a participant first pressed the back button of the Vive controller and slowly released it, and in referent H (*Zoom*) a participant meant to use the top orifice of this controller as a lens and placed it near the eye. In referent C (*Bevel*), a participant simply pressed the back button, moved it down a little, and released it. This suggests that gestures with props were more preferred than using a controller.

Props for gestures vs. gestures for props. I presented two separate ways of measuring agreement, which correspond to different strategies by participants: choosing their gesture first and finding props to support that gesture, or picking a desired prop and then figuring out a gesture to perform. As such, I offered two different agreement scores to indicate how consistent people were in their choices.

In other cases, props were likely chosen to fit the gesture. Referents like *twist*, *rotate*, and *assemble* show a high gesture agreement score compared to the prop score, thus suggesting that a lot of props afforded the most used gestures. In these cases, it may be more important to recognize the gesture than the props that were chosen in an implemented system.

I made several attempts at a combined prop-gesture agreement score, but ultimately I decided that this was outside the scope of this thesis. First, developing such a score would need to consider the structure of how props were chosen for different gestures, possibly categorizing the objects by their affordances. Second, it could be that the agreement score should depend more on the prop choice in some cases and more on the gesture in others (e.g., through weights), depending on the referent; more analysis would help understand when each choice was important.

Context-free grammars allow for expressive and programmable gestures. When grouping gestures, I found that a natural way to express them was using context-free grammars (CFGs). This notation captures the variations in props that each gesture had, and groups together the types of props that were used in a certain role in that gesture. For instance, with *Reshape*, a sword and an x-acto knife could be used as "CUTTING_PROP"; presumably other objects could be used to cut as well. Or, with *Close*, a notebook, a plane, a LEGO brick and a hinge in its flat position could be used as "DOOR_PROP". Reasonably other flat and square props could act as "DOOR_PROP".

Using a formalism like a CFG also enables system designers to efficiently describe gestures for a sensing system; if other props not used in our study could fulfill a role, then they could simply be added to the list. Combined with a method for detecting objects and recognizing grammars, then designers could adapt and define gestures to various situations.

Generalized agreement score enables new elicitation studies. To measure agreement between participants' prop selections, I generalized participant agreement scores in order to handle sets of selections. When participants make a single selection, generalized agreement scores are equivalent to the agreement score formula first presented in (Wobbrock et al., 2009) and later formalized in (Vatavu and Wobbrock, 2015). However, with this new agreement score, researchers and designers can study agreement in new situations.

For example, multimodal input methods, such as gesture and speech input, previously had to look at agreement separately (Khan and Tunçer, 2019). If participants are able to choose either a gesture, a voiced utterance, or both, then our generalized agreement score could capture partial overlap between what participants suggest, rather than studying modalities independently. Any of the three similarity metrics we suggested (Jaccard, Sørensen, Overlap) will give a 1 for full equivalence, a 0 for no overlap, and an intermediate value for partial overlap.

The generalized agreement score can also handle elicitation of sets from participants. If a researcher wanted to follow up with a study on preferred props for a referent, they can use this measure to identify which referents have an agreement score. Alternatively, this could work for other modalities: suggesting multiple gestures or multiple commands as options. This technique could also be used for finer-grained comparison of gestures: if two participants provide a gesture for a referent, and they use different sets of fingers, a partial agreement score could capture similarities between gestures.

Chance-corrected agreement relatively close to raw agreement The chance-corrected agreement score was marginally different from the overall gesture agreement score calculated with Vatavu and Wobbrock's (2015) formula. This suggests that there was little to correct, as the chance agreement was relatively close to zero ($\rho_e = .016$). This also suggests that the 21 participants were minimally influenced by bias factors and that they did not focus on a limited number of gestures. Generally, gestures proposed for a given referent were barely proposed for another one. "`Fold {hinge} with hands`" and "`Shake {wand} with hand`" were elicited for up to three referents maximum. None of these gestures, however, resulted to be the representative gesture for any referent in the gesture set. Tsandilas (2018) hypothesized that in elicitation studies that investigated contextual and direct manipulation gestures, such as Wobbrock et al.'s (2009), Piumsomboon et al.'s (2013) and this study, chance agreement could be minimal, compared to studies that focused on abstract gestures for noncontextual operations. The results I reported in this thesis support that hypothesis.

Chapter 6

Conclusion, Limitations, and Future Work

In this thesis, I investigated what gestures people would perform with physical props to complete 20 CAD-like and open-world referents in VR, what physical props people would choose to do these referents, and how they would leverage physical props to manipulate virtual objects. This thesis presents the results of an elicitation study for manipulative gestures with props that aims to answer those questions. Based on 21 participants' multimodal agreement, a user-defined gesture set with props for 20 CAD-like and open-world games referents was developed. For each referent in Figure 4.8 and Figure 4.9, the gesture set presents the most preferred (popular) gesture with its prop and gesture agreement scores calculated from the input of the 21 participants. A new agreement score based on set similarity metrics was used to analyze agreement between both gestures and props. This score is identical to previous agreement scores (Vatavu and Wobbrock, 2015; Wobbrock et al.,

2009) when used for unimodal selection (gestures only), but also accommodates multimodal selections (props). In addition, context-free grammars were found to be a useful approach to capture various props used in a gesture, the role of each, and how participants manipulate or operate them. In this way, I also developed an expressive language for articulating gestures with physical props for each referent in the user-defined gesture set.

To summarize, the contributions of this research are the following:

1. A user-defined set of gesture-prop combinations for 20 CAD-like and open-world-like referents (Figure 4.8 and Figure 4.9).
2. An expressive language for articulating and implementing gestures with physical props, based on context-free grammars (Figure 4.8 and Figure 4.9).
3. A generalization of the agreement score metric to account for multiple selections in elicitation studies (Equation 4.3).
4. Insight into how people leverage physical props to perform gestures in VR (chapter 5).

These contributions are important for the HCI community for several reasons. Firstly, this work confirms that physical objects are widely preferred over hand-held controllers for interaction with virtual environments. However, the gesture set presented in Figure 4.8 and Figure 4.9 reflects that researchers and designers do not require a large vocabulary of props. Instead, focus can be turned towards shapes and affordances of physical props, which allows people to do gesture sets with only few props or to perform several gestures with the same prop(s) (Figure 5.1). Second, the new agreement score metric establishes a path for multimodal elicitation studies that can lead to more robust gesture sets. It is believed that multimodal input communicates more information, exhibits more expressivity, and offers more precision to interact with VR systems. In other words, it increases the scope of what people can do to interact with VR systems. Third, the use of context-free grammars allows to construct an expressive language to articulate gestures with props. This structure could be adopted by future multimodal gesture elicitation studies (e.g., gesture + speech). Computers and systems can also be programmed to understand this language. Lastly, insights into how people leverage props for gestures was gained. Choosing gesture first and then finding props to support it or choosing desired props and then making up a gesture were generally the two strategies that participants followed in the elicitation study.

The elicitation study for this thesis also helped to find an input vocabulary for gestures with physical props in VR that allows users to do things that would not be possible in the real world. Hand-held controllers, like the current state of the art, offer a clear input mapping to execute general commands in VR, such as scrolling down a menu following the move of a joystick. Physical props, on the other hand, have a clear purpose. Phone cases enclose and protect phones, pencils serve as writing utensils, candles are light sources, et cetera. Physical props are not suitable for general commands. The results of this study, nonetheless, suggest that physical props can actually help execute general commands in VR when these are combined with gestures. From Figure 4.8 and Figure 4.9, it can be observed that people are able to darken the sky by depicting an arc-like trajectory with a sphere, or to duplicate an object by snapping two identical models with LEGO-brand bricks. For instance, gestures with props expand the ability of users to do tasks in VR that reality would not support.

6.1 Limitations

A limitation of this elicitation study is that the gesture-prop set I propose requires validation. Similar to the ones proposed by Wobbrock et al. 2009 and Kray et al. (2010), a follow-up study may require showing video footage of the gestures to a new group of participants, to investigate if they can infer the intended command (referent). Usability-wise, it may be needed to ask this new group of subjects to score with one of five responses that range from strongly agree to strongly disagree a series of items, for example, "I found the gestures unnecessarily complex to the intended referents". Memorability-wise, in another study, the same 21 participants could be recruited to obtain a memorability score of the gestures they proposed for each referent.

Another limitation is that the 21 participants recruited for this elicitation study were undergraduate and graduate students from different disciplines across the University of Waterloo campus. Professional designers, architects, and gamers might behave differently when, for example, creating a bevel on a rectangular prism, opening a door in an open-world game, or hanging pieces of art on the wall of an envisioned living room. This foundational work focused on a set of manipulative gestures usable by non-technical people to do CAD 3D modelling in VR or perform in an open-world game, enhancing immersion, intuitiveness, and realism with the incorporation of physical props.

Chance agreement for props remains unknown. So does chance-corrected agreement. Tsandilas (2018) shows how to determine chance agreement ρ_e for gestures in elicitation studies using Scott's (1955) π (section 4.5). Determining chance agreement for props, nonetheless, is more complicated, as it is not clear what the effect of bias is. A quick solution may be to assume that each prop available has the same chance of being chosen by a participant ($\rho_e = 1/q$), where q is the total number of props. However, Tsandilas explains that this assumption is generally not realistic, as it does not account for bias. The Scott's π metric includes bias, but further research is needed to explore whether it can be applied for props. As stated in section 4.5, this metric estimates the probability that a participant classifies a referent into a gesture, based on how many times this gesture has been used across the study. If participants were limited to use one prop per referent, Scott's π could be suitable. However, they form props sets, meaning that both the size of the vocabulary of props and the chance of choosing one prop or another change on the fly. For example, if a participant first picks a wrench, what's the probability of picking a cube afterwards? Also, in section 5, I discussed that participants thought of a gesture and found a group of props suitable for it, or picked a desired group of props and figured out a gesture to perform. Suppose that a participant desires to `stretch` an `object with both hands`. Rigid and hard props would likely not have possibility of being chosen, and the group of props from which this participant could pick would now be reduced. Moreover, does the number of virtual objects in a referent bias participants' selection criteria? Or, should researchers explore chance agreement of shapes and affordances, instead of the actual props? These are some questions I intend to answer in the future in order to estimate chance agreement for props accurately and fairly.

6.2 Future Work

The results and conclusions presented in this thesis direct attention towards several areas that require further research. In the future, I plan to study the relationship between gesture choice and prop choice, formalize the affordances for elicited props, and explore how gestures with physical props could be used in a working system, for example, a VR CAD environment or open-world video game.

It is worth mentioning that the gestures with props shown in Figure 4.8 and Figure Figure 4.9 do not work in VR yet. This may not be easy to implement. VR systems intending to utilize my work would require a method for identifying gestures, detecting props, and recognizing grammars. To identify users' gestures, motion sensing input devices like the Microsoft Kinect could be adopted. With regards to physical props, appearance-based methods may be used. The main challenge, nonetheless, is the simple fact that props may look different depending on lighting conditions, distance, or orientation. Once systems recognize a person has performed a gesture with props, a comparison or association with a context-free grammar database or toolkit should happen. The successful completion of these previous stages would then allow users to perform a task in VR.

This work provides researchers and designers with tools and frameworks to develop more gesture sets. Gestures with props allow people to do things reality does not support, which could help industries striving to adopt VR. In the future, researchers could define prop-gesture vocabularies (sets) to let people explore the solar system, to allow an interior designer to try different layouts of furniture, or to give engineering students the chance of assembling every part of an engine. Moreover, the proposed generalization of the agreement score formula could allow researchers to build more robust and more versatile gesture sets for human interaction with VR and with other technologies.

References

Parastoo Abtahi and Sean Follmer. Visuo-Haptic Illusions for Improving the Perceived Performance of Shape Displays. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18*, pages 1–13, New York, New York, USA, 2018. ACM Press. ISBN 9781450356206. doi: 10.1145/3173574.3173724. URL `http://dl.acm.org/citation.cfm?doid=3173574.3173724`.

Sema Alaçam. The Many Functions of Hand Gestures While Communicating Spatial Ideas An Empirical Case Study. In *Proceedings of the XVIII Conference of the Iberoamerican Society of Digital Graphics - SIGraDi: Design in Freedom*, pages 106–109, São Paulo, dec 2014. Editora Edgard Blücher. doi: 10.5151/despro-sigradi2014-0017. URL `http://www.proceedings.blucher.com.br/article-details/14234`.

Ayman Alzayat, Mark Hancock, and Miguel Nacenta. Quantitative measurement of virtual vs. physical object embodiment through kinesthetic figural after effects. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*, pages 2903–2912, New York, New York, USA, 2014. ACM Press. ISBN

9781450324731. doi: 10.1145/2556288.2557282. URL http://dl.acm.org/citation.cfm?doid=2556288.2557282.

Ayman Alzayat, Mark Hancock, and Miguel A. Nacenta. Measuring Readiness-to-Hand through Differences in Attention to the Task vs. Attention to the Tool. In *Proceedings of the Interactive Surfaces and Spaces on ZZZ - ISS '17*, pages 42–51, New York, New York, USA, 2017. ACM Press. ISBN 9781450346917. doi: 10.1145/3132272.3134114. URL http://dl.acm.org/citation.cfm?doid=3132272.3134114.

Ayman Alzayat, Mark Hancock, and Miguel A. Nacenta. Quantitative Measurement of Tool Embodiment for Virtual Reality Input Alternatives. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, pages 1–11, New York, New York, USA, 2019. ACM Press. ISBN 9781450359702. doi: 10.1145/3290605.3300673. URL http://dl.acm.org/citation.cfm?doid=3290605.3300673.

Bruno Araujo, Ricardo Jota, Varun Perumal, Jia Xian Yao, Karan Singh, and Daniel Wigdor. Snake Charmer: Physically Enabling Virtual Objects. In *Proceedings of the TEI '16: Tenth International Conference on Tangible, Embedded, and Embodied Interaction - TEI '16*, pages 218–226, New York, New York, USA, 2016. ACM Press. ISBN 9781450335829. doi: 10.1145/2839462.2839484. URL http://dl.acm.org/citation.cfm?doid=2839462.2839484.

Shaikh Shawon Arefin Shimon, Courtney Lutton, Zichun Xu, Sarah Morrison-Smith, Christina Boucher, and Jaime Ruiz. Exploring Non-touchscreen Gestures for Smartwatches. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing*

Systems, pages 3822–3833. ACM, may 2016. ISBN 9781450333627. doi: 10.1145/2858036.2858385. URL `https://dl.acm.org/doi/10.1145/2858036.2858385`.

Jatin Arora, Aryan Saini, Nirmita Mehra, Varnit Jain, Shwetank Shrey, and Aman Parnami. VirtualBricks: Exploring a Scalable, Modular Toolkit for Enabling Physical Manipulation in VR. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, pages 1–12, New York, New York, USA, 2019. ACM Press. ISBN 9781450359702. doi: 10.1145/3290605.3300286. URL `http://dl.acm.org/citation.cfm?doid=3290605.3300286`.

Mahdi Azmandian, Mark Hancock, Hrvoje Benko, Eyal Ofek, and Andrew D. Wilson. Haptic Retargeting. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 1968–1979. ACM, may 2016. ISBN 9781450333627. doi: 10.1145/2858036.2858226. URL `https://dl.acm.org/doi/10.1145/2858036.2858226`.

Hrvoje Benko, Christian Holz, Mike Sinclair, and Eyal Ofek. NormalTouch and TextureTouch. *Proceedings of the 29th Annual Symposium on User Interface Software and Technology - UIST '16*, pages 717–728, 2016. doi: 10.1145/2984511.2984526. URL `http://dl.acm.org/citation.cfm?doid=2984511.2984526`.

Barbaros Bostan and Gökhan Nalbant. Interaction in Virtual Reality. *International Symposium of Interactive Medical Design*, 2006. URL `https://www.researchgate.net/publication/229021841`.

Rosalino Rodriguez Calderon and Miguel Angel Garcia Ruiz. Learning physics with virtual environment. In *2019 IEEE World Conference on Engineering Education (EDUNINE)*, pages 1–5. IEEE, mar 2019. ISBN 978-1-7281-1666-2. doi: 10.1109/EDUNINE.2019.8875824. URL `https://ieeexplore.ieee.org/document/8875824/`.

Jack Shen-Kuen Chang, Georgina Yeboah, Alison Doucette, Paul Clifton, Michael Nitsche, Timothy Welsh, and Ali Mazalek. TASC: Combining Virtual Reality with Tangible and Embodied Interactions to Support Spatial Cognition. In *Proceedings of the 2017 Conference on Designing Interactive Systems - DIS '17*, pages 1239–1251. ACM Press, 2017. ISBN 9781450349222. doi: 10.1145/3064663.3064675. URL `http://dl.acm.org/citation.cfm?doid=3064663.3064675`.

Lung-Pan Cheng, Li Chang, Sebastian Marwecki, and Patrick Baudisch. iTurk: Turning Passive Haptics into Active Haptics by Making Users Reconfigure Props in Virtual Reality. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18*, volume 2018-April, pages 1–10. ACM Press, 2018. ISBN 9781450356206. doi: 10.1145/3173574.3173663. URL `http://dl.acm.org/citation.cfm?doid=3173574.3173663`.

Inrak Choi and Sean Follmer. Wolverine. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology - UIST '16 Adjunct*, pages 117–119. ACM Press, 2016. ISBN 9781450345316. doi: 10.1145/2984751.2985725. URL `http://dl.acm.org/citation.cfm?doid=2984751.2985725`.

Inrak Choi, Heather Culbertson, Mark R. Miller, Alex Olwal, and Sean Follmer. Grabity. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*, pages 119–130, New York, NY, USA, oct 2017. ACM. ISBN

9781450349819. doi: 10.1145/3126594.3126599. URL https://dl.acm.org/doi/10.1145/3126594.3126599.

Inrak Choi, Eyal Ofek, Hrvoje Benko, Mike Sinclair, and Christian Holz. CLAW: A Multifunctional Handheld Haptic Controller for Grasping, Touching, and Triggering in Virtual Reality. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18*, pages 1–13. ACM Press, 2018. ISBN 9781450356206. doi: 10.1145/3173574.3174228. URL http://dl.acm.org/citation.cfm?doid=3173574.3174228.

Alexandra Delazio, Ken Nakagaki, Roberta L. Klatzky, Scott E. Hudson, Jill Fain Lehman, and Alanson P. Sample. Force Jacket: Pneumatically-Actuated Jacket for Embodied Haptic Experiences. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18*, pages 1–12. ACM Press, 2018. ISBN 9781450356206. doi: 10.1145/3173574.3173894. URL http://dl.acm.org/citation.cfm?doid=3173574.3173894.

Bradley Efron. Bootstrap Methods: Another Look at the Jackknife. In *Annals of Statistics*, pages 569–593. 1992. doi: 10.1007/978-1-4612-4380-9_41. URL http://link.springer.com/10.1007/978-1-4612-4380-9{_}41.

L. Egghe and C. Michel. Strong similarity measures for ordered sets of documents in information retrieval. *Information Processing & Management*, 38(6):823–848, nov 2002. ISSN 03064573. doi: 10.1016/S0306-4573(01)00051-6. URL https://linkinghub.elsevier.com/retrieve/pii/S0306457301000516.

Joseph L. Fleiss. Measuring nominal scale agreement among many raters. *Psychological Bulletin*, 76(5):378–382, 1971. ISSN 0033-2909. doi: 10.1037/h0031619. URL `http://content.apa.org/journals/bul/76/5/378`.

Jakob Frohner, Gionata Salvietti, Philipp Beckerle, and Domenico Prattichizzo. Can Wearable Haptic Devices Foster the Embodiment of Virtual Limbs? *IEEE Transactions on Haptics*, 12(3):339–349, jul 2019. ISSN 1939-1412. doi: 10.1109/TOH.2018.2889497. URL `https://ieeexplore.ieee.org/document/8587215/`.

James J Gibson. 04-JJ Gibson-Ch8-Affordances. *Chapter Eight The Theory of Affordances*, pages 127–136, 1986. ISSN 10725520. doi: citeulike-article-id:3508530.

Antonio Gomes, Lahiru Lakmal Priyadarshana, Aaron Visser, Juan Pablo Carrascal, and Roel Vertegaal. MagicScroll: A Rollable Display Device with Flexible Screen Real Estate and Gestural Input. In *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services*, pages 1–11, New York, NY, USA, sep 2018. ACM. ISBN 9781450358989. doi: 10.1145/3229434.3229442. URL `https://dl.acm.org/doi/10.1145/3229434.3229442`.

Xiaochi Gu, Yifei Zhang, Weize Sun, Yuanzhe Bian, Dao Zhou, and Per Ola Kristensson. Dexmo: An Inexpensive and Lightweight Mechanical Exoskeleton for Motion Capture and Force Feedback in VR. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 1991–1995. ACM, may 2016. ISBN 9781450333627. doi: 10.1145/2858036.2858487. URL `https://dl.acm.org/doi/10.1145/2858036.2858487`.

K.L. Gwet. *Handbook of Inter-Rater Reliability, 4th Edition: The Definitive Guide to Measuring The Extent of Agreement Among Raters*. Advanced Analytics, LLC, 2014. ISBN 9780970806284. URL `https://books.google.ca/books?id=fac9BQAAQBAJ`.

Nur Suraya Sahol Hamid, Faieza Abdul Aziz, and Amir Azizi. Virtual reality applications in manufacturing system. In *2014 Science and Information Conference*, pages 1034–1037. IEEE, aug 2014. ISBN 978-0-9893193-1-7. doi: 10.1109/SAI.2014.6918317. URL `http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6918317`.

Mark Hancock, Thomas ten Cate, Sheelagh Carpendale, and Tobias Isenberg. Supporting Sandtray Therapy on an Interactive Tabletop. In *Proceedings of the 28th international conference on Human factors in computing systems - CHI '10*, page 2133, New York, New York, USA, 2010. ACM Press. ISBN 9781605589299. doi: 10.1145/1753326.1753651. URL `http://portal.acm.org/citation.cfm?doid=1753326.1753651`.

H. G. Hoffman, A. Hollander, K. Schroder, S. Rousseau, and T. Furness. Physically touching and tasting virtual objects enhances the realism of virtual experiences. *Virtual Reality*, 3(4):226–234, dec 1998. ISSN 1359-4338. doi: 10.1007/BF01408703. URL `http://link.springer.com/10.1007/BF01408703`.

Linda Homeyer. *Sandtray Therapy: A Practical Manual (2nd edition)*. Routledge, 01 2011.

Hsin-yu Huang, Chih-wei Ning, Po-Yao Wang, Jen-Hao Cheng, and Lung-Pan Cheng. Haptic-go-round: A Surrounding Platform for Encounter-type Haptics in Virtual Reality Experiences. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–10, New York, NY, USA, apr 2020. ACM. ISBN

9781450367080. doi: 10.1145/3313831.3376476. URL https://dl.acm.org/doi/10.1145/3313831.3376476.

Jinmiao Huang and Rahul Rai. Hand Gesture Based Intuitive CAD Interface. In *Volume 1A: 34th Computers and Information in Engineering Conference*, page V01AT02A094. ASME, aug 2014. ISBN 978-0-7918-4628-5. doi: 10.1115/DETC2014-34070. URL http://proceedings.asmedigitalcollection.asme.org/proceeding.aspx?doi=10.1115/DETC2014-34070.

Hessam Jahani and Manolya Kavakli. Exploring a user-defined gesture vocabulary for descriptive mid-air interactions. *Cognition, Technology & Work*, 20(1):11–22, feb 2018. ISSN 1435-5558. doi: 10.1007/s10111-017-0444-0. URL http://link.springer.com/10.1007/s10111-017-0444-0.

Sumbul Khan and Bige Tunçer. Automation in Construction Gesture and speech elicitation for 3D CAD modeling in conceptual design. *Automation in Construction*, 106(July 2017):102847, 2019. ISSN 0926-5805. doi: 10.1016/j.autcon.2019.102847. URL https://doi.org/10.1016/j.autcon.2019.102847.

Sumbul Khan, Bige Tunçer, Ramanathan Subramanian, and Lucienne Blessing. 3D CAD modeling using gestures and speech: Investigating CAD legacy and non-legacy procedures. *CAADFutures19*, 2019. URL http://cumincad.scix.net/cgi-bin/works/paper/cf2019{_}042.

Farzam Kharvari and Wolfgang Hohl. The Role of Serious Gaming using Virtual Reality Applications for 3D Architectural Visualization. In *2019 11th International Conference* *on Virtual Worlds and Games for Serious Applications (VS-Games)*, pages 1–2. IEEE, sep 2019. ISBN 978-1-7281-4540-2. doi: 10.1109/VS-Games.2019.8864576. URL `https://ieeexplore.ieee.org/document/8864576/`.

Kimin Kim and Jinah Park. Virtual bone drilling for dental implant surgery training. In *Proceedings of the 16th ACM Symposium on Virtual Reality Software and Technology - VRST '09*, volume C, page 91. ACM Press, 2009. ISBN 9781605588698. doi: 10.1145/1643928.1643950. URL `http://portal.acm.org/citation.cfm?doid=1643928.1643950`.

Christian Kray, Daniel Nesbitt, John Dawson, and Michael Rohs. User-defined gestures for connecting mobile phones, public displays, and tabletops. In *Proceedings of the 12th international conference on Human computer interaction with mobile devices and services - MobileHCI '10*, page 239. ACM Press, 2010. ISBN 9781605588353. doi: 10.1145/1851600.1851640. URL `http://portal.acm.org/citation.cfm?doid=1851600.1851640`.

Klaus Krippendorff. Reliability in Content Analysis: Some Common Misconceptions and Recommendations. *Human Communication Research*, 30(3):411–433, jul 2004. ISSN 0360-3989. doi: 10.1111/j.1468-2958.2004.tb00738.x. URL `https://academic.oup.com/hcr/article/30/3/411-433/4331534`.

Daniel Leithinger, Sean Follmer, Alex Olwal, and Hiroshi Ishii. Shape Displays: Spatial Interaction with Dynamic Physical Form. *IEEE Computer Graphics and Applications*, 35(5):5–11, sep 2015. ISSN 0272-1716. doi: 10.1109/MCG.2015.111. URL `http://ieeexplore.ieee.org/document/7274417/`.

Yuhan Liu, Yong Tang, and Qian Sun. SYVR Virtual Equipment Manufacturing Simulation. In *2018 International Conference on Virtual Reality and Visualization (ICVRV)*, pages 139–139. IEEE, 2018. ISBN 978-1-5386-8497-9. doi: 10.1109/ICVRV.2018.00046. URL https://ieeexplore.ieee.org/document/8711278/.

Allan Christian Long, James A. Landay, and Lawrence A. Rowe. Implications for a gesture design tool. In *Proceedings of the SIGCHI conference on Human factors in computing systems the CHI is the limit - CHI '99*, pages 40–47. ACM Press, 1999. ISBN 0201485591. doi: 10.1145/302979.302985. URL http://portal.acm.org/citation.cfm?doid=302979.302985.

Alessio Malizia and Andrea Bellucci. The artificiality of natural user interfaces. *Communications of the ACM*, 55(3):36–38, mar 2012. ISSN 0001-0782. doi: 10.1145/2093548.2093563. URL https://dl.acm.org/doi/10.1145/2093548.2093563.

John C. McClelland, Robert J. Teather, and Audrey Girouard. Haptobend: Shape-Changing Passive Haptic Feedback in Virtual Reality. In *Proceedings of the 5th Symposium on Spatial User Interaction - SUI '17*, pages 82–90. ACM Press, 2017. ISBN 9781450354868. doi: 10.1145/3131277.3132179. URL http://dl.acm.org/citation.cfm?doid=3131277.3132179.

Meredith Ringel Morris, Andreea Danielescu, Steven Drucker, Danyel Fisher, Bongshin Lee, C. Schraefel, and Jacob O. Wobbrock. Reducing legacy bias in gesture elicitation studies. *interactions*, 21(3):40–45, may 2014. ISSN 10725520. doi: 10.1145/2591689. URL http://dl.acm.org/citation.cfm?doid=2608008.2591689.

MR Morris, JO Wobbrock, and AD Wilson. Understanding Users' Preferences for Surface Gestures. *Proceedings of Graphics Interface*, pages 261–268, 2010. ISSN 13596454. doi: 10.1016/j.actamat.2009.07.058. URL `http://dl.acm.org/citation.cfm?id=1839260`.

Thomas Muender, Anke V Reinschluessel, Sean Drewes, Dirk Wenig, Tanja Döring, and Rainer Malaka. Does It Feel Real? Using Tangibles with Different Fidelities to Build and Explore Scenes in Virtual Reality. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, pages 1–12, New York, New York, USA, 2019. ACM Press. ISBN 9781450359702. doi: 10.1145/3290605.3300903. URL `http://dl.acm.org/citation.cfm?doid=3290605.3300903`.

Miguel A Nacenta, Yemliha Kamber, Yizhou Qiang, and Per Ola Kristensson. Memorability of pre-designed and user-defined gesture sets. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13*, page 1099, New York, New York, USA, 2013. ACM Press. ISBN 9781450318990. doi: 10.1145/2470654.2466142. URL `http://dl.acm.org/citation.cfm?doid=2470654.2466142`.

Vijayakumar Nanjappan, Hai Ning Liang, Feiyu Lu, Konstantinos Papangelis, Yong Yue, and Ka Lok Man. User-elicited dual-hand interactions for manipulating 3D objects in virtual reality environments. *Human-centric Computing and Information Sciences*, 8(1), 2018. ISSN 21921962. doi: 10.1186/s13673-018-0154-5. URL `https://doi.org/10.1186/s13673-018-0154-5`.

Michael J. Neil. Architectural virtual reality applications. *ACM SIGGRAPH Computer Graphics*, 30(4):53–54, nov 1996. ISSN 0097-8930. doi: 10.1145/240806.240816. URL `https://dl.acm.org/doi/10.1145/240806.240816`.

Donald A. Norman. *The Design of Everyday Things.* Basic Books, 2002. ISBN 0465067107. URL http://books.google.com/books?hl=en{&}lr={&}id=w8pM72p{_}dpoC{&}pgis=1.

Thammathip Piumsomboon, Adrian Clark, Mark Billinghurst, and Andy Cockburn. User-defined gestures for augmented reality. In *CHI '13 Extended Abstracts on Human Factors in Computing Systems on - CHI EA '13*, volume 2013-April, page 955. ACM Press, 2013. ISBN 9781450319522. doi: 10.1145/2468356.2468527. URL http://dl.acm.org/citation.cfm?doid=2468356.2468527.

Kun Qian, Junxuan Bai, Xiaosong Yang, Junjun Pan, and Jianjun Zhang. Virtual reality based laparoscopic surgery simulation. In *Proceedings of the 21st ACM Symposium on Virtual Reality Software and Technology - VRST '15*, volume 13-15-Nove, pages 69–78. ACM Press, 2015. ISBN 9781450339902. doi: 10.1145/2821592.2821599. URL http://dl.acm.org/citation.cfm?doid=2821592.2821599.

Francis Quek, David McNeill, Robert Bryll, Susan Duncan, Xin-Feng Ma, Cemil Kirbas, Karl E. McCullough, and Rashid Ansari. Multimodal human discourse: Gesture and speech. *ACM Trans. Comput.-Hum. Interact.*, 9(3):171–193, September 2002. ISSN 1073-0516. doi: 10.1145/568513.568514. URL https://doi.org/10.1145/568513.568514.

M. H. Quenouille. Problems in Plane Sampling. *The Annals of Mathematical Statistics*, 20(3):355–375, sep 1949. ISSN 0003-4851. doi: 10.1214/aoms/1177729989. URL http://projecteuclid.org/euclid.aoms/1177729989.

Jaime Ruiz and Daniel Vogel. Soft-Constraints to Reduce Legacy and Performance Bias to Elicit Whole-body Gestures with Low Arm Fatigue. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15*, volume 2015-April, pages 3347–3350, New York, New York, USA, 2015. ACM Press. ISBN 9781450331456. doi: 10.1145/2702123.2702583. URL `http://dl.acm.org/citation.cfm?doid=2702123.2702583`.

Jaime Ruiz, Yang Li, and Edward Lank. User-Defined Motion Gestures for Mobile Interaction. In *Proceedings of the 2011 annual conference on Human factors in computing systems - CHI '11*, page 197. ACM Press, 2011. ISBN 9781450302289. doi: 10.1145/1978942.1978971. URL `http://dl.acm.org/citation.cfm?doid=1978942.1978971`.

Julien Saunier, Mukesh Barange, Bernard Blandin, Ronan Querrec, and Joanna Taoum. Designing adaptable virtual reality learning environments. In *Proceedings of the 2016 Virtual Reality International Conference on - VRIC '16*, pages 1–4. ACM Press, 2016. ISBN 9781450341806. doi: 10.1145/2927929.2927937. URL `http://dl.acm.org/citation.cfm?doid=2927929.2927937`.

Peter Schulz, Dmitry Alexandrovsky, Felix Putze, Rainer Malaka, and Johannes Schöning. The Role of Physical Props in VR Climbing Environments. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, pages 1–13. ACM Press, 2019. ISBN 9781450359702. doi: 10.1145/3290605.3300413. URL `http://dl.acm.org/citation.cfm?doid=3290605.3300413`.

William A. Scott. Reliability of Content Analysis: The Case of Nominal Scale Coding. *Public Opinion Quarterly*, 19(3):321, 1955. ISSN 0033362X. doi: 10.1086/266577. URL https://academic.oup.com/poq/article-lookup/doi/10.1086/266577.

Jotaro Shigeyama, Takeru Hashimoto, Shigeo Yoshida, Takuji Narumi, Tomohiro Tanikawa, and Michitaka Hirose. Transcalibur: AWeight Shifting Virtual Reality Controller for 2D Shape Rendering based on Computational Perception Model. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, pages 1–11. ACM Press, 2019. ISBN 9781450359702. doi: 10.1145/3290605.3300241. URL http://dl.acm.org/citation.cfm?doid=3290605.3300241.

Mike Sinclair, Eyal Ofek, Mar Gonzalez-Franco, and Christian Holz. CapstanCrunch. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, pages 815–829. ACM, oct 2019. ISBN 9781450368162. doi: 10.1145/3332165.3347891. URL https://dl.acm.org/doi/10.1145/3332165.3347891.

Alexa F. Siu, Eric J. Gonzalez, Shenli Yuan, Jason B. Ginsberg, and Sean Follmer. shapeShift: 2D Spatial Manipulation and Self-Actuation of Tabletop Shape Displays for Tangible and Haptic Interaction. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18*, volume 2018-April, pages 1–13, New York, New York, USA, 2018. ACM Press. ISBN 9781450356206. doi: 10.1145/3173574.3173865. URL http://dl.acm.org/citation.cfm?doid=3173574.3173865.

Patrick L Strandholt, Oana A Dogaru, Niels C Nilsson, Rolf Nordahl, and Stefania Serafin. Knock on Wood: Combining Redirected Touching and Physical Props for Tool-Based

Interaction in Virtual Reality. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–13, New York, NY, USA, apr 2020. ACM. ISBN 9781450367080. doi: 10.1145/3313831.3376303. URL `https://dl.acm.org/doi/10.1145/3313831.3376303`.

Evan Strasnick, Christian Holz, Eyal Ofek, Mike Sinclair, and Hrvoje Benko. Haptic Links: Bimanual Haptics for Virtual Reality Using Variable Stiffness Actuation. *Conference on Human Factors in Computing Systems*, pages 1–12, 2018. doi: 10.1145/3173574.3174218. URL `http://dl.acm.org/citation.cfm?doid=3173574.3174218`.

Shan-yuan Teng, Cheng-Lung Lin, Chi-huan Chiang, Tzu-Sheng Kuo, Liwei Chan, Da-Yuan Huang, and Bing-Yu Chen. TilePop: Tile-type Pop-up Prop for Virtual RealityTilePoP. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, pages 639–649, New York, NY, USA, oct 2019. ACM. ISBN 9781450368162. doi: 10.1145/3332165.3347958. URL `https://dl.acm.org/doi/10.1145/3332165.3347958`.

Theophanis Tsandilas. Fallacies of Agreement: A Critical Review of Consensus Assessment Methods for Gesture Elicitation. *ACM Transactions on Computer-Human Interaction*, 25(3):1–49, jun 2018. ISSN 1073-0516. doi: 10.1145/3182168. URL `https://dl.acm.org/doi/10.1145/3182168`.

Consuelo Valdes, Diana Eastman, Casey Grote, Shantanu Thatte, Orit Shaer, Ali Mazalek, Brygg Ullmer, and Miriam K. Konkel. Exploring the design space of gestural interaction with active tokens through user-defined gestures. In *Proceedings of the 32nd annual*

ACM conference on Human factors in computing systems - CHI '14, pages 4107–4116, New York, New York, USA, 2014. ACM Press. ISBN 9781450324731. doi: 10.1145/2556288.2557373. URL `http://dl.acm.org/citation.cfm?doid=2556288.2557373`.

Radu-Daniel Vatavu. User-defined gestures for free-hand TV control. In *Proceedings of the 10th European conference on Interactive tv and video - EuroiTV '12*, page 45, New York, New York, USA, 2012. ACM Press. ISBN 9781450311076. doi: 10.1145/2325616.2325626. URL `http://dl.acm.org/citation.cfm?doid=2325616.2325626`.

Radu-Daniel Vatavu and Jacob O. Wobbrock. Formalizing Agreement Analysis for Elicitation Studies. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15*, pages 1325–1334, 2015. doi: 10.1145/2702123.2702223. URL `http://dl.acm.org/citation.cfm?doid=2702123.2702223`.

Christian Weichel, Manfred Lau, David Kim, Nicolas Villar, and Hans W. Gellersen. MixFab: A Mixed-Reality Environment for Personal Fabrication. *Proceedings of the Annual ACM Conference on Human Factors in Computing Systems (CHI)*, pages 3855–3864, 2014. ISSN 0160-4953. doi: 10.1145/2556288.2557090. URL `http://dl.acm.org/citation.cfm?doid=2556288.2557090`.

Eric Whitmire, Hrvoje Benko, Christian Holz, Eyal Ofek, and Mike Sinclair. Haptic Revolver. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18*, pages 1–12, New York, New York, USA, mar 2018. ACM Press. ISBN 9781450356206. doi: 10.1145/3173574.3173660. URL `http://dl.acm.org/citation.cfm?doid=3173574.3173660`.

Jacob O. Wobbrock, Htet Htet Aung, Brandon Rothrock, and Brad A. Myers. Maximizing the guessability of symbolic input. In *CHI '05 extended abstracts on Human factors in computing systems - CHI '05*, page 1869. ACM Press, 2005. ISBN 1595930027. doi: 10.1145/1056808.1057043. URL http://portal.acm.org/citation.cfm?doid=1056808.1057043.

Jacob O Wobbrock, Meredith Ringel Morris, and Andrew D Wilson. User-Defined Gestures for Surface Computing. In *Proceedings of the 27th international conference on Human factors in computing systems - CHI 09*, volume 8, page 1083, New York, New York, USA, 2009. ACM Press. ISBN 9781605582467. doi: 10.1145/1518701.1518866. URL http://dl.acm.org/citation.cfm?doid=1518701.1518866.

Yukang Yan, Chun Yu, Xiaojuan Ma, Xin Yi, Ke Sun, and Yuanchun Shi. VirtualGrasp: Leveraging Experience of Interacting with Physical Objects to Facilitate Digital Object Retrieval. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18*, volume 2018-April, pages 1–13. ACM Press, 2018. ISBN 9781450356206. doi: 10.1145/3173574.3173652. URL http://dl.acm.org/citation.cfm?doid=3173574.3173652.

Jackie Yang, Hiroshi Horii, Alexander Thayer, and Rafael Ballagas. VR Grabbers: Ungrounded Haptic Retargeting for Precision Grabbing Tools. *The 31st Annual ACM Symposium on User Interface Software and Technology - UIST '18*, pages 889–899, 2018. doi: 10.1145/3242587.3242643. URL https://doi.org/10.1145/3242587.3242643.

Andre Zenner and Antonio Kruger. Shifty: A Weight-Shifting Dynamic Passive Haptic Proxy to Enhance Object Perception in Virtual Reality. *IEEE Transactions on Visualiza-*

tion and Computer Graphics, 23(4):1285–1294, apr 2017. ISSN 1077-2626. doi: 10.1109/TVCG.2017.2656978. URL https://ieeexplore.ieee.org/document/7833030/.

André Zenner and Antonio Krüger. Drag:on– A Virtual Reality Controller Providing Haptic Feedback Based on Drag and Weight Shift. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, pages 1–12, New York, New York, USA, 2019. ACM Press. ISBN 9781450359702. doi: 10.1145/3290605.3300441. URL http://dl.acm.org/citation.cfm?doid=3290605.3300441.

Kening Zhu, Taizhou Chen, Feng Han, and Yi-Shiun Wu. HapTwist: Creating Interactive Haptic Proxies in Virtual Reality Using Low-cost Twistable Artefacts. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, pages 1–13, New York, New York, USA, 2019. ACM Press. ISBN 9781450359702. doi: 10.1145/3290605.3300923. URL http://dl.acm.org/citation.cfm?doid=3290605.3300923.

Add a title page before the appendices and a line in the Table of Contents

APPENDICES

Appendix A

Matlab Scripts

A.1 Code for Calculating Prop Agreement Score of Referent C: Bevel

```
People = 21                 ;
Pairs = nchoosek(People,2) ;

P1C  = [‘‘Cube’’,‘‘Xacto_knife’’]     ;
P2C  = [‘‘Cube’’,‘‘Flex_knife’’]      ;
P3C  = [‘‘Couch’’,‘‘Notebook_cover’’] ;
P4C  = [‘‘Cube’’,‘‘Hinge’’]           ;
P5C  = [‘‘Cube’’]                     ;
P6C  = [‘‘LEGO1’’,‘‘LEGO2’’]          ;
P7C  = [‘‘Flex_knife’’]               ;
```

```
P8C  = [‘‘LEGO1’’, ‘‘LEGO2’’]          ;

P9C  = [‘‘Abstract_sphere’’]           ;

P10C = [‘‘Cube’’]                      ;

P11C = [‘‘Cube’’,‘‘Flex_knife’’]       ;

P12C = [‘‘Cube’’]                      ;

P13C = [‘‘Cube’’]                      ;

P14C = [‘‘Cube’’]                      ;

P15C = [‘‘Mario_BROS’’]                ;

P16C = [‘‘Cushion’’]                   ;

P17C = [‘‘Cube’’]                      ;

P18C = [‘‘Tire’’]                      ;

P19C = [‘‘Cube’’,‘‘Flex_knife’’]       ;

P20C = [‘‘LEGO1’’, ‘‘LEGO2’’]          ;

P21C = [‘‘Vive_controller’’]           ;


ParticipantsC = {P1C, P2C, P3C, P4C, P5C, P6C, P7C, P8C, P9C, P10C, ..., P21C} ;

C = nchoosek(ParticipantsC,2)                                                  ;

g_C = 0                                                                        ;


for i=1:1:Pairs

 f_C = (numel((intersect(C{i,1},C{i,2})))) / (numel(union(C{i,1},C{i,2}))) ;

 g_C = [(f_C)+(g_C)]                                                       ;

end


AR_bevel = g_C/Pairs ;
```

A.2 Code for Calculating Confidence Intervals of Prop Agreement Score of Referent C: Bevel

```
n = 21                    ;
People = n-1              ;
Pairs = nchoosek(People,2) ;


P1C  = [‘‘Cube’’,‘‘Xacto_knife’’]     ;
P2C  = [‘‘Cube’’,‘‘Flex_knife’’]      ;
P3C  = [‘‘Couch’’,‘‘Notebook_cover’’] ;
P4C  = [‘‘Cube’’,‘‘Hinge’’]           ;
P5C  = [‘‘Cube’’]                     ;
P6C  = [‘‘LEGO1’’,‘‘LEGO2’’]          ;
P7C  = [‘‘Flex_knife’’]               ;
P8C  = [‘‘LEGO1’’, ‘‘LEGO2’’]         ;
P9C  = [‘‘Abstract_sphere’’]          ;
P10C = [‘‘Cube’’]                     ;
P11C = [‘‘Cube’’,‘‘Flex_knife’’]      ;
P12C = [‘‘Cube’’]                     ;
P13C = [‘‘Cube’’]                     ;
P14C = [‘‘Cube’’]                     ;
P15C = [‘‘Mario_BROS’’]               ;
P16C = [‘‘Cushion’’]                  ;
P17C = [‘‘Cube’’]                     ;
P18C = [‘‘Tire’’]                     ;
P19C = [‘‘Cube’’,‘‘Flex_knife’’]      ;
P20C = [‘‘LEGO1’’, ‘‘LEGO2’’]         ;
```

```
P21C = [‘‘Vive_controller’’]          ;

AR_bevel = 0.190;
sumC = 0        ;
q = 1.725       ;

for j=1:1:n
 ParticipantsC = {P1C, P2C, P3C, P4C, P5C, P6C, P7C, P8C, P9C, P10C, ..., P21C} ;
 ParticipantsC(j) = []         ;
 C = nchoosek(ParticipantsC,2) ;
 g_C = 0                       ;

    for i=1:1:Pairs
     f_C = (numel((intersect(C{i,1},C{i,2})))) / (numel(union(C{i,1},C{i,2})));
     g_C = [(f_C)+(g_C)];
    end

 AR_jth_C = g_C/Pairs                       ;
 sumC = [((AR_jth_C - AR_bevel)^2) + sumC] ;
end

Vjack_C = ((n-1)/n)*sumC ;
SE_C = sqrt(Vjack_C)     ;
CIl_C = AR_bevel-(q*SE_C)
CIh_C = AR_bevel+(q*SE_C)
```

Appendix B

Ethics Approval

Dear Mark Hancock and other members of the research team:

Your application has been reviewed by Delegated Reviewers. We are pleased to inform you the **Initial application for 40865 Interacting in Virtual Reality with 3D-Printed Objects** has been given ethics clearance.

This research must be conducted in accordance with the most recent version of the application in the research ethics system and the most recent versions of all supporting materials.

Ethics clearance for this study is valid until Sunday, May 3rd 2020.

The research team is responsible for obtaining any additional institutional approvals that might be required to complete this Expedited study.

University of Waterloo Research Ethics Committees operate in compliance with the institution's guidelines for research with human participants, the Tri-Council Policy Statement for the Ethical Conduct for Research Involving Humans (TCPS, 2nd edition), Internalization Conference on Harmonization: Good Clinical Practice (ICH-GCP), the Ontario Personal Health Information Protection Act (PHIPA), and the applicable laws and regulations of the province of Ontario. Both Committees are registered with the U.S. Department of Health and Human Services under the Federal Wide Assurance, FWA00021410, and IRB registration number IRB00002419 (Human Research Ethics Committee) and IRB00007409 (Clinical Research Ethics Committee).

Renewal: Multi-year research must be renewed at least once every 12 months unless a more frequent review has been specified on the notification of ethics clearance. This is a requirement as outlined in Article 6.14 of the Tri-Council Policy Statement for the Ethical Conduct for Research Involving Humans (TCPS2, 2014). The annual renewal report/application must receive ethics clearance before Friday, April 10th 2020. Failure to receive ethics clearance for a study renewal will result in suspension of ethics clearance and the researchers must cease conducting the study. Research Finance will be notified ethics clearance is no longer valid.

Amendment: Changes to this study are to be submitted by initiating the amendment procedure in the research ethics system and may only be implemented once the proposed changes have received ethics clearance.

Adverse event: Events that adversely affect a study participant must be reported as soon as possible, but no later than 24 hours following the event, by contacting the Director, Research Ethics. Submission of an adverse event form is to follow the next business day.

Deviation: Unanticipated deviations from the approved study protocol or approved documentation or procedures are to be reported within 7 days of the occurrence using a protocol deviation form.

Incidental finding: Anticipated or unanticipated incidental findings are to be reported as soon as possible by contacting the Director, Research Ethics. Submission of the incidental findings form is to follow within 3 days of learning of the finding. Participants may not be contacted regarding incidental findings until after clearance has been received from a Research Ethics Committee to contact participants to disclose these findings.

Study closure: Report the end of this study by submitting a study closure report through the research ethics system.

Coordinated Reviews: If your application was reviewed in conjunction with Wilfrid Laurier University, Conestoga College, Western University or the Tri-Hospital Research Ethics Board, note the following: 1) Amendments must receive prior ethics clearance through both REBs before the changes are put in place, 2) PI must submit the required annual renewal report to both REBs and failure to complete the necessary annual reporting requirements may result in Research Finance being notified at both institutions, 3) In the event that there is an unanticipated event involving a participant that adversely affects them, the PI must report this to both REBs within 24 hours of the event taking place and any unanticipated or unintentional changes which may impact the research protocol shall be reported within seven days of the deviation to both REBs.

Initial application ethics clearance notification: Your clearance notification will be added to the record within 24 hours. Go to “View Admin Attachments” in the research ethics system (right-hand side) to print a copy of the initial application ethics clearance notification.

****

Best wishes for success with this study.

If you have any questions concerning this notification, please contact the Research Ethics Office or email researchethics@uwaterloo.ca.

Appendix C

Study Supplementary Materials

INFORMATION LETTER

Date: March 25, 2019
Project: Interacting in Virtual Reality with 3D-Printed Objects

Principal Investigator
Dr. Mark Hancock
Management Sciences
University of Waterloo, Canada
mark.hancock@uwaterloo.ca

Student Investigator
Marco Moran-Ledesma
Systems Design Engineering
University of Waterloo, Canada
mamoranl@uwaterloo.ca

Study Overview

You are invited to participate in a research study that aims to explore new interaction techniques in Virtual Reality (VR) using 3D-printed objects. Our project will focus on providing a design vocabulary of 3D-printed objects for use in VR. The goal is to obtain guidelines for how to better design interaction in VR.

What you will be asked to do

As a participant, you will be asked to watch short videos within a VR environment through the *HTC Vive* headset. Each video will present a task, e.g. stacking cubes, parking a car, decorating a living room, etc. Following this, you will be asked to pick one 3D-printed object of your choice from a bucket and show the researcher how you would use it to complete the task you just saw, which will be video taped. Furthermore, you will be asked to complete a computer-administered survey and a paper-pencil questionnaire. The computer-administered survey will be operated by Qualtrics. When information is transmitted or stored on the internet privacy cannot be guaranteed. There is always a risk your responses may be intercepted by a third party (e.g., government agencies, hackers). If there are any questions that you do not wish to answer you can skip them.

Here are pictures of the HTC Vive headset and some of the 3D printed objects that you will have the opportunity to choose:





Participation and Remuneration

In order to participate in the study, you must not have visual, audio and/or balance impairments.
Participation in this study is completely voluntary and will take approximately 90 minutes of your time. You may decline to answer any questions presented during the study if you so wish. Furthermore, you may decide to withdraw from this study by advising the researcher and may do so without any penalty. $15 in cash will be given to you for your time invested in this important step of this research. The amount received is taxable. It is your responsibility to report this amount for income tax purposes.

Benefits of the Study

Since the current Virtual Reality devices found in the market offer limited interactivity through conventional hand-held controllers, users can't interact with virtual objects in the same way as real-world objects. Being able to feel the shape of an object when grasping it in VR enhances a sense of presence and allows object manipulation, and although prior work focuses on stimulating fingers, grasping a 3D object requires the sensation of touch using the entire hand. The goal of this research is to create a system with the help of 3D-printed objects that enables interactions in virtual environments in a broader and more-realistic way. As a result, people will be able to manipulate physical objects with their hands and observe this manipulation in the virtual environment, resulting in a much more immersive experience.

Risks to Participation in the Study

Some individuals may experience dizziness, headaches, or nausea as a result of wearing head mounted display devices. There is a very rare chance (1 in 4000) of severe dizziness, seizures, eye or muscle twitching or blackouts with use of VR, with this risk being higher in children and people under 20. These symptoms can be more likely if you are fatigued, under stress or anxiety, or suffering from cold/flu/headaches. The symptoms from VR exposure can persist and become more apparent hours after use. In addition, some individuals may experience a loss of balance due to using the head mounted device, and this may result in a fall.

To reduce the likeliness of any discomfort caused by the head mounted display device, you will have breaks between conditions and may request a break at any time. If need be, we will also seek out the appropriate medical assistance, however it is unlikely that this will be necessary as the risks listed above are quite rare. As there is a possibility for the symptoms from VR exposure to persist and become more apparent hours after use, you should remain alert for any symptoms in the next few hours. If at any time you have strong adverse reactions associated with motion sickness, as a safeguard, testing will end. Please note that the VR head mounted display will be cleaned with alcohol sanitizers after each participant's use.

To safeguard the possible symptoms of headaches, nausea, and potential falls, a comfortable chair where you can sit and a cot where you can lay down if necessary are available along with water bottles. In the case of a fall, foam mats are located on the floor and a spotter assigned to prevent actual falls. If symptoms persist after the study, the participant will be advised to remain in the testing room until feeling well. If the symptoms return after leaving the testing space, the participant is advised to consult a doctor and inform researchers, and to avoid driving, operating machinery, engaging in demanding activities that have potentially serious consequences, or other activities that require unimpaired balance and hand-eye coordination (such as playing sports or riding a bicycle, etc.) until fully recovered from any symptoms.

Confidentiality and Security of Data

Your identity will be confidential. Your name will not be included in any thesis or report resulting from this study. Furthermore, because the interest of this study is in the average responses of the entire group of participants, you will not be identified individually in any way in any written reports of this research. Paper records of data collected during this study will be retained in a locked office at the University of Waterloo for a minimum of 7 years, to which only researchers associated with this study have access. Electronic data and audio-video recordings will be kept for a minimum of 7 years on a password-protected computer in a locked office at the University of Waterloo, to which only researchers associated with this study have access to. All identifying information will be removed from the records prior to storage.

This study has been reviewed and received ethics clearance through a University of Waterloo Research Ethics Committee (ORE #40865). If you have questions for the Committee contact the Office of Research Ethics, at 1-519-888-4567 ext. 36005. For all other questions regarding this research study, please feel free to contact any member of the research team listed on this form.

Interacting in Virtual Reality with 3D-Printed Objects

Participant Consent Form

By signing this consent form, you are not waiving your legal rights or releasing the investigator(s) or involved institution(s) from their legal and professional responsibilities. I have read the Information Letter regarding the study being conducted by the M.A.Sc. Student Marco Moran-Ledesma in Systems Design Engineering at the University of Waterloo. I have had the opportunity to ask questions related to this study, to receive satisfactory answers to my questions, and any additional details I wanted.

I am aware that I have the option of allowing my performance in the study and my interview to be audio and video recorded to ensure an accurate acquisition of my responses. When video recordings are shown, faces will be blurred. I am also aware that excerpts from the performance and the interview may be included in the thesis and/or publications to come from this research, with the understanding that the quotations will be anonymous.

I was informed that I may withdraw my consent at any time during the study without penalty by advising the researcher.

Only researchers associated with this study will have access to the password-protected study records. Paper files will be locked securely. We will keep your data for a minimum of **7** years. You can withdraw your consent to participate and ask that your data be destroyed by contacting one of the researchers within this time. It is not possible to withdraw your once papers and publications have been submitted to publishers. All data will be destroyed according to University of Waterloo policy.

This study has been reviewed and received ethics clearance through a University of Waterloo Research Ethics Committee (ORE #40865). If you have questions for the Committee contact the Office of Research Ethics, at 1-519-888-4567 ext. 36005 or ore-ceo@uwaterloo.ca. If you have any questions regarding this study, you can also contact the principal investigator, Dr. Mark Hancock at mark.hancock@uwaterloo.ca or 519-888-4567 ext. 36587.

With full knowledge of all foregoing, I agree, of my own free will, to participate in this study.

☐YES ☐NO

I agree to have both my performance and interview audio and video recorded.

☐YES ☐NO

I agree to the use of anonymous quotations in any thesis or publication that comes of this research.

☐YES ☐NO

Participant Name: ______________________________ (Please print)

UNIVERSITY OF WATERLOO

Participant Signature: ______________________________

Witness Name: _________________________________ (Please print)

Witness Signature: ________________________________ Date: _YYYY_/_MM_/_DD_