Haptic Feedback in Virtual Reality with Deformation and Shape-Change

by

John McClelland

A thesis submitted to the Faculty of Graduate and Postdoctoral Affairs in

partial fulfillment of the requirements for the degree of

Masters of Information Technology

in

Digital Media

Carleton University

Ottawa, Ontario

© 2018, John McClelland

Abstract

Past Virtual Reality (VR) research shows that haptic feedback increases presence and improves users' task performance. However, providing haptic feedback for multiple virtual objects usually requires complex, immobile systems, or multiple haptic props. We present a new approach that applies deformable, shape-changing devices to VR haptics, leveraging the dominance of human vision in VR to provide realistic haptic feedback with physical shape approximations. Our first study evaluates our HaptoBend prototype through an elicitation study. Results support the use of physical shape approximations and reveal important user preferences. We translate these results and past work into a Design Criteria to inform our second prototype, Adaptic. In our second study, we compare docking performance and adherence to our Design Criteria with Adaptic, a Razor Hydra Controller, and haptic props. We found Adaptic did well in satisfying our Design Criteria and had little difference in performance compared to the other haptic approaches.

Acknowledgments

Over my time working on this thesis I have had the privilege of receiving support from my amazing family, friends, and colleagues. First I would like to thank my parents for their unconditional support in every possible way. From understanding when I needed space to work on a deadline to snapping into action when I needed their help, they have truly been there from me every step of the way. The advice and support of my brother, sister, and brother-in-law have also done a lot to help guide me through this experience.

I would also like to thank my supervisors Audrey Girouard and Robert J. Teather. I couldn't have asked for better people to oversee my work. I can't see any way that this thesis could have been completed without their constant motivation and support. Their trust in my abilities has done a lot for my own self confidence and made me a better person within and outside of academics. They also introduced me to community of supportive colleagues including those at the Creative Interactions Lab. Each person I have met through them has contributed something to this thesis and I am thankful to all of them.

Finally, my friends deserve a huge amount of praise for dealing with my canceled plans, late arrivals, and lapses in communication that resulted from focusing on this thesis. I especially want to thank Andrew Parks, Nader Kaddour, Stephanie Bourassa, Thomas Corkum, and especially my partner Chloe Barker for their love, support, and empathy over the last two years.

Table of Contents

List of Figures

List of Tables

Chapter 1: Introduction

1.1 Motivation

Haptics facilitate richer user interactions by adding a dimension of physical feedback to digital interfaces. Examples such as vibration from a mobile device and detailed flight simulator hardware used for training shows that a wide variety of applications and complexity are possible for haptics. Haptic feedback is especially valuable in virtual reality (VR), without it the disparity between visual and physical experiences negatively affects presence during direct contact with a virtual object [71]. Likewise, the absence of tangible interaction removes an important reference introducing constraints to 3D manipulation tasks. Even with the importance of VR haptics, no solution for easily accessible general purpose haptic feedback exists.

Much of the development in VR thus far focuses on advancements in the visual aspects of a virtual environment (VE), while haptic feedback in VR shows much slower progress [4,8]. As a result, the popularity of today's commercial head-mounted displays illudes existing commercial and research haptic devices for VR [12,51]. Convention separates most approaches to VR haptics into two main categories: passive haptic feedback (PHF) and active haptic feedback (AHF). AHF devices [51,60,69] generalize easily to virtual objects by actuating components to limit a user's movement and providing force feedback. PHF provides a more accessible approach to haptics by using physical props, similar to a corresponding virtual object [5,32,74].

Past research reveals several persistent issues found within AHF and PHF. AHF relies on expensive systems that are complex, intrusive and lack physically robust feedback [84]. In contrast, PHF suffers from the complexity of switching props, and the excessive number of props needed for a general haptic system [4]. We classify these prevalent issues in VR haptics into three categories:

1. **Complexity:** Intimidatingly intricate AHF systems, and an excessive number of props both result in an undesirable level of complexity for users.
2. **Limited Interactions in VR:** The need to switch between props for different virtual objects, and intrusive hardware from AHF systems limit the interactions users can participate in and create breaks in presence.
3. **Inadequate Haptic Feedback:** Oversimplified PHF, such as universal, wand-style controllers, and AHF that lack physically robust feedback disrupt immersion and presence in VR.

Our solution to these issues applies a combination of deformable[1], shape-changing[2] devices, and the dominance of human vision to VR haptics. To our knowledge, no other research has explored the application of deformable, shape-changing devices to VR haptics. Therefore, our research targets this gap to assess the usability of deformable VR haptic devices.

Similar to conventional VR controllers, past examples [25,58,65] show the form factor of deformable and shape-changing devices accommodate handheld interactions. By

---

[1] **Throughout this thesis, we will refer to deformation as changes in shape through user manipulation.**

[2] **Throughout this thesis, we will refer to shape-change as changes in shape through self-actuation.**

transitioning into different shapes to create physical approximations of virtual objects these devices also address issues of over-complexity by reducing the need for multiple PHF props. While physical approximations could also integrate other factors, such as size and weight, we focus on shape approximation as it aligns best with the affordances of deformable and shape-changing devices. The dominance of vision over other senses in VR improves the realism of shape approximations allowing them to serve as adequate haptic feedback [4,18,41,70,83,84]. Therefore, we believe our approach of applying deformable, shape-changing devices to VR haptics is a valid way to mitigate the issues of complexity, limited interactions in VR, and inadequate haptic feedback found in PHF and AHF approaches today.

1.2 Research Goals and Questions

Our research aims to answer the question*: Do deformable, shape-changing devices offer a positive alternative to conventional haptic feedback approaches for VR*? To simplify this question, we define a goal for testing each of our user studies with corresponding research questions meant to facilitate reaching each goal. We also developed two prototypes, HaptoBend (Figure 1) and Adaptic (Figure 2), in the interest of completing each of the user studies.

Goal 1 (HaptoBend): Investigate user preferences for a device that can change shape to create haptic feedback for different objects.

1.1. User Impressions: In what contexts do users feel comfortable using a deformable haptic device meant for VR?

1.2. **Shape Approximation:** Do users prefer physical shape approximations of corresponding virtual objects for haptic feedback with a deformable haptic device?

1.3. **Contributing Factors:** What factors influence how users prefer to use a deformable haptic device in VR?

Goal 2 (Adaptic): Compare a deformable, shape-changing device with other approaches for VR haptics.

2.1. **Design Criteria:** Does Adaptic satisfy the Design Criteria we outlined for deformable, shape-changing VR haptic devices?

2.2. **Performance:** How well does a deformable haptic device perform compared to other haptic approaches in VR?

2.3. **Preference:** Do users prefer a deformable device over alternative approaches for VR haptics?

1.3 Contribution

As a solution to the prevalent issues in VR haptics we developed two prototypes, HaptoBend (Figure 1) and Adaptic (Figure 2) and evaluated each with an informative user study. Both prototypes represent original contributions as the first attempts to combine the affordances of deformable and shape-changing devices with VR haptics. With HaptoBend we performed a user study resulting in the following contributions:

- The design for HaptoBend, a deformable PHF device compatible with a range of virtual objects.

- The first elicitation study outlining preferred PHF shapes for virtual objects using a deformable device.

- Further evidence that physical shape approximations are sufficient when providing PHF for 2D and 3D virtual objects.

Using the findings from our first study, we improved on HaptoBend to create another prototype, Adaptic. Running a user study to evaluate Adaptic resulted in these contributions:

- The design for Adaptic, a VR haptic device compatible with a range of virtual objects that allows both deformability and actuated shape-change.

- The first comparative docking task study exhibiting equivalent performance of a deformable VR haptic device to when compared to other haptic approaches.

- Our Design Criteria for implementing a deformable, shape-changing VR haptic devices.

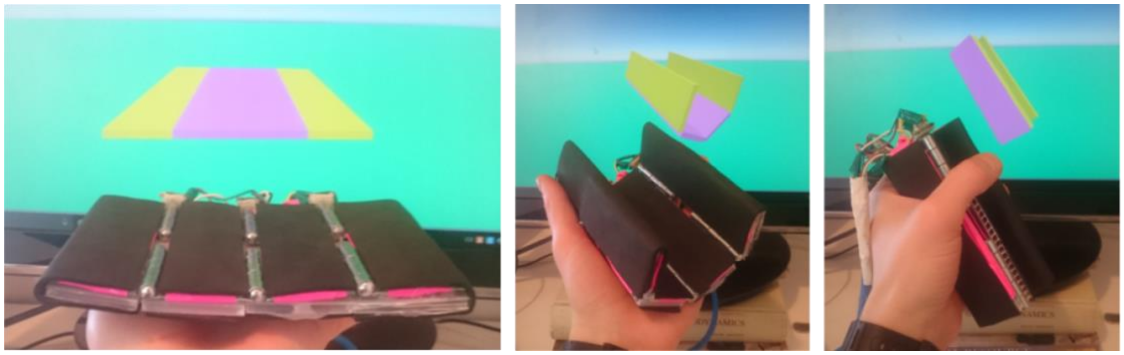

Figure 1. HaptoBend and its real-time digital reconstruction.

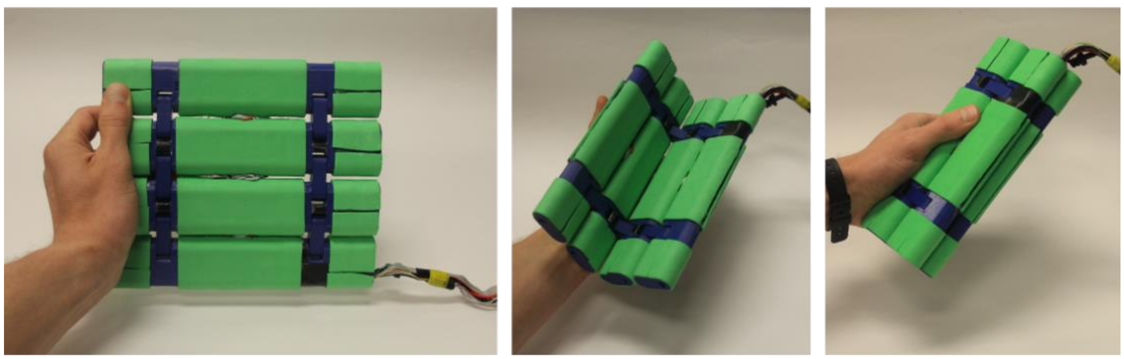

Figure 2. Our Adaptic prototype.

1.4 Thesis Outline

After introducing our topic of interest in Chapter 1, we discuss the main areas of research that contribute to our work in a literature review found in Chapter 2. The literature review covers past approaches in VR haptics and identifies persistent gaps. It also covers the research behind the new approach we propose by discussing deformable and shape-changing devices, visual dominance in VR, and the methods we apply in our user studies.

Chapters 3 covers our first user study with HaptoBend, a deformable prototype for VR haptics. The chapter introduces the reader to HaptoBend and explains our prototyping process. We also report our research methodology, and results from running an elicitation user study with HaptoBend including user preferences for using the device. Chapter 4 introduces Adaptic, our second prototype and outlines our process for developing this iteration by creating our Design Criteria which incorporates design feedback from our first user study and past requirements for similar devices. Chapter 5 explains our second user study where we compare Adaptic's performance to other approaches for VR haptics using docking tasks and discuss our results.

Chapter 6 wraps up our research with a summary of our process and contributions. We then conclude by giving recommendations for future work that applies deformable, shape-changing devices to VR haptics.

1.5 Publications and Presentations

We presented HaptoBend and the first user study in this document, which explores the shape preferences of our prototype, at the 2017 ACM Symposium on Spatial User Interaction (SUI).

> John C McClelland, Robert J Teather, and Audrey Girouard. 2017. HaptoBend: shape-changing passive haptic feedback in virtual reality. In *Proceedings of the ACM Symposium Spatial User Interaction*, 9p.

HaptoBend also appeared at SUI 2017 as a demo which received the "Best Demo" award.

John C McClelland, Robert J Teather, and Audrey Girouard. 2017. Haptic Feedback with HaptoBend: utilizing shape-change to enhance virtual reality. In *Proceedings of the ACM Symposium Spatial User Interaction*, 150p.

Chapter 2: Literature Review

Past research in VR shows benefits of haptic feedback in both spatial awareness and presence [36]. We look at the two main categories of haptic feedback, PHF and AHF, to identify their benefits and shortcomings. Next, to differentiate our approach from past work, we review several of the devices that resulted from acknowledging issues with traditional PHF and AHF. We then draw on past research to show the benefits of the three components that make up our approach: visual dominance in VR, plane-like deformable, and shape-changing devices. Finally, we provide a review of our main assessment procedures by examining the use of elicitation studies and docking task studies in related contexts.

2.1 Passive Haptics

PHF applies physical proxies that are similar to corresponding virtual objects to increase realism. The simplicity of this approach can make it accessible by tracking existing objects or fabricating them through processes like 3D printing [84], or even constructing them out of paper [37]. While PHF props demonstrate benefits, research also points to the drawback of requiring separate physical objects for each virtual object [84]. Most notably these include the need for multiple PHF props and complexity of prop switching, both of which increase in difficulty as the variety of virtual objects grow.

A long history of research in VR supports the benefits of PHF. Hinckley et al. [32] contributed some of the first research exploring the benefits of PHF in 1994. Their study assessed the use of a tracked doll head as PHF for 3D brain models, along with a plane-cutting prop for viewing inside a brain model. While they use a monitor with 3D graphics instead of a VR headset, the work still laid some of the groundwork for PHF use in VR. Their research compares the PHF interface with a direct manipulation interface by asking neurosurgeons to explore the brain model with both. Results from this show a higher preference for the PHF approach due to its metaphor creating a better connection with the neurosurgeons' mental models.

The use of 3D props as physical proxies persists as one of the main approaches for PHF. Phobia treatment with props is another example that demonstrates its potential for high levels of realism. Using physical replicas of spiders for PHF, Carlin et al. [11] and Garcia et al. [22] performed studies which show VR as effective in treating arachnophobia. Carin et al. [11] used a spider toy as realistic PHF in a VE developed to treat a single patient's arachnophobia with a visual and haptic experiences. Researchers gathered data on fear during weekly sessions over 3 months and found a drastic reduction over that time. In a later study, Garcia et al. [22] executed a similar study, comparing 12 participants that received treatment with VR, including a toy spider for PHF, to 11 participants who did not. Of the 12 participants who received treatment 83% had significant improvements to their phobia while 0% of the untreated group improved.

Hoffman et al. [33] was the first to provide empirical evidence to support the benefits of PHF by examining its effect on realism in a virtual kitchen. Participants picked up and manipulated a plate with one of two conditions, either using a wand controller or a tracked real world plate. Through a series of questions the researchers found PHF led to higher levels of realism through participants basing their perception of the VE on their haptic experience. Similar studies continue today, the recent work by Besançon et al. [5] shows the benefits of 3D props over touchscreen and mouse-based interactions. Their study compared these three input modalities with a docking task study to collect data on speed, accuracy, workload and user preference. The results show using a PHF prop led to faster task completion at an equal level of accuracy, participants also preferred using the tangible prop which led to lower workload levels.

Another common PHF approach is the use of 2D plane-like surfaces for a variety of 2D interactions. One of the earliest examples is seen in the work of Stoakley et al. [74], where a clipboard served as PHF for a worlds-in-miniature metaphor (WIM) to facilitate selection and navigation. They used a 3D tracked clip board representing a handheld miniature model of the VE, allowing users to observe the entire VE at a smaller scale and manipulate it with 6 degrees of freedom (DOF). A 3D tracked tennis ball also provided PHF for selection and manipulation of virtual objects in the full-scale VE and the WIM. Informal observations showed the interface to be fairly intuitive and effective for 3D selection and manipulation.

Lindeman et al. [46] contributed early findings on the benefits of handheld devices over fixed devices. They developed a Haptic Augmented Reality Paddle (HARP) that provides PHF for 2D interfaces in a VE. The authors assessed the impact of PHF by running participants through the same set of tasks with a HARP present or absent. When using a HARP it was either in a fixed position or held in the hand of a participant. Their results showed PHF increased speed and accuracy in selection and docking tasks, and overall the handheld approach provided more benefits. Teather et al. [77] and Joyce & Robson [38] echo the positive effect of using a flat plane for PHF in 2D interfaces through their research with Fitt's Law tests in VR.

2.2 Active Haptics

AHF devices [51,60,69] can represent the shapes of different virtual objects by actuating to limit a user's joint movement and providing force feedback. Through our review of AHF we see their main drawbacks stem from using an apparatus that is either large, complex, intrusive, expensive or a combination of these. In addition, active haptics can lack the robust physical feedback of solid PHF props, depending on the strength of their actuators.

Some of the earliest research on AHF by is seen in the work of Ouh-Young et al. [57]. Their work compared user performance in a docking task with two conditions: only visual feedback and only AHF. They used a very large, ceiling mounted force-reflecting joystick known as the Argonne E-3 Remote Manipulator (ARM) for AHF and presented visual feedback in 3D stereo vision with alternating polarization plate glasses. After running the study, results showed that AHF led to much higher performance, reporting that participants could complete the task more than twice as fast. Work with large, mounted AHF systems like the ARM continues to develop to this day. However, devices like the EXO-UL3 [60], a wall mounted, full-arm exoskeleton, provide an example that shows much of this work focuses on technical assessments rather than user interactions.

String-based systems are another approach to AHF. To create force feedback these systems connect several strings to a contact point, such as a finger, and increase or reduce tension on the strings accordingly to emulate the physical feedback of objects. SPIDAR [69] demonstrates several approaches to string-based AHF. Their research documents an iterative process that progresses from providing force feedback to one finger, to four fingers (including the thumb) on each hand for applications like surgery practice in a VE.

While effective, all our previous examples are large, complex and intrusive. One example that avoids these issues is the PHANTOM [51]. Developed in 1994, it is a much more manageable in size using a single stylus connected to a base that provides force feedback. While the PHANTOM is relatively simple compared to other examples, Achibet et al. [1] showed combining it with other haptic devices can yield positive results for more detailed interactions. However, this approach occupies two hands, one for each haptic device, to control interactions with one virtual hand. Another smaller scale AHF approach focuses directly on hands with glove-like exoskeleton systems as shown by CyberGrasp [78] and the Rutgers Master II [7].

2.3 Diverse Haptic Systems for VR

Seeking to address issues found in traditional AHF and PHF systems like the EXO-UL3 [60] and the work of Hinckley et al. [32], more recent research explores the possibilities for simpler and more diverse haptic systems. One approach for general haptics creates hybrid devices by combining aspects of AHF and PHF.

Zhao et al. [85] applied Zooid [23] swarm robots to construct PHF props out of a set of building blocks. They performed a technical evaluation to measure the speed and accuracy of their approach. While their approach allows for a diverse set of props from a limited amount of materials, the process appears to limit interactions through slow speed and limited accuracy of prop construction.

Several other devices aimed at general haptics use a linear sliding mechanism that restricts movement by locking at specified points. Aguerreche [2] introduced the concept of Reconfigurable Tangible Devices (RTD), with two prototypes. They are assembled of telescopic rods joined at pivoting connections that create the outline of a triangle (RTD-3) and a square (RTD-4) respectively. Both allow manipulation in different shapes that users can maintain by locking the rods and pivot points in place. Using a collaborative "pick and place" task in a user study, they gathered data on RTD-3 that suggested participants prefer RTD over non-tangible techniques in the context of collaborative manipulation.

Emulating AHF gloves mentioned in the previous section, Wolverine [16] and Grabity [15] use a similar break-sliding mechanism. Wolverine [16] is a small hand-based device that restricts movement between users' fingers and thumb with break-based locking sliders that move along rods. A technical evaluation of the device shows it can deliver over 100N of force between each finger and the thumb, and it can render haptic feedback close to human perception. However, Wolverine does lack variable stiffness, meaning it is best suited for rigid virtual objects. Instead of a sliding mechanism for each finger, Grabity [15] uses one for the entire hand but also incorporates voice coil actuators that stretch a user’s finger pad skin to mimic weight and grasping. Researchers evaluated Grabity with two users studies that asked participants to evaluate simulated object weights. Results showed participants could differentiate between weights and determine how heavy they were relative to one another. Both Grabity [15] and Wolverine [16] are simple and provide adequate haptic feedback, but their approach of attaching hardware to a user’s hands can limit interactions available in VR, especially those requiring both hands to interact with each other.

Another approach to general haptics removes physical interactions with objects altogether. One method uses electrical muscle stimulation (EMS) to mimic the feel of interacting with virtual objects [49,50]. A series of user studies showed the use of EMS for haptics contributes to higher realism, when compared to no haptic feedback [49] and vibrotactile feedback [50]. Ultrahaptics [12,76], provides physical stimulation with a matrix of ultrasonic speakers that create converging frequencies at points in 3D space. Carter et al. [12] tested Ultrahaptics in the absence of a visual display, using a Leap Motion to track hand position. Their findings showed participants could distinguish between points of haptic feedback with different intensities placed close to one another, and differentiate between changes in vibration frequency. While these examples of discarding physical elements do aid simplicity and freedom of movement in VR, the lack of tangible objects limits their robustness leading to inadequate haptic feedback for some situations.

2.4 Visual Dominance

By concentrating more on the user instead of hardware, past research shows visual dominance over other senses in VR allows the use of physical approximations to work as immersive haptic feedback for users [4,18,20,41,70,84]. Simeone et al. [70] found 3D PHF props can still be effective when their shape is not an exact match to corresponding virtual objects. They performed two studies, the first used a mug as PHF for different virtual objects including mugs that varied in parameters such as shape, size and texture. The results showed that many of the virtual objects did not lead to significant differences in believability when compared with a baseline virtual object that was identical to the PHF mug. The second study used only one virtual object, a lightsaber, and tested a replica lightsaber, flashlight, and umbrella as PHF. Their findings show that the flashlight performed best, with few statistical differences from the replica.

Redirected touching is another benefit of visual dominance that enables the use of one haptic device for multiple virtual objects [4,40]. The process involves either redirecting a user's hand, or warping the perceived location of virtual objects relative to the haptic device. Using these approaches, Azmandian et al. [4] developed several techniques for reusing PHF devices through warping the user's visual environment. The techniques include: Body Warping which manipulates a user's tracked body parts, World Warping which changes the world position of virtual objects and Hybrid Warping, a combination of both techniques. Comparing these techniques against a wand-based interface showed they offer higher levels of presence and realism, with the hybrid technique performing the best.

Navigation in VR provides further support for visual dominance by showing that warping a VE can remain undetected through redirected walking [64,73]. Razzaque et al. [64] allowed users to walk around VEs larger than their physical environment using redirected walking. Their technique subtly rotates the VE so that areas of the physical space are reused, relying on visual dominance for the rotations to go unnoticed. With a user study the researchers showed this technique was effective for navigating in a physical room roughly half the size of the VE.

2.5 Deformable Plane-Like Devices

Other than our own prototypes, we are not aware of any deformable haptic devices for VR. Therefore, our approach focuses on applying the observed benefits of deformables in the real world to haptic experiences in VR.

Flexible devices are a commonly pursued research area in HCI, with a large number prototypes utilizing a flat, plane-like form factor [25,42,62,72]. Studies of these devices show deformability allows users to create shapes that complement the context of their use for both fully flexible devices, and those with rigid displays connected by hinges [10,14,25,31,62]. Gomes & Vertegaal [25] explored display interactions based on paper metaphors by developing PaperFold, a prototype made of three displays with flexible, detachable connections. With PaperFold they used participatory design to study user preferences for transitioning between shapes. Results aligned with our prototyping approach by showing a preference towards using a single device that allows transitioning from a flat panel to a 3D triangular prism, and that users associate folding PaperFold into 3D shapes with beginning 3D interactions.

The ability to collect rich input through detailed shape tracking is another benefit of deformables. Richer tracking can facilitate more detailed display interactions [21,65,72], capture emotional states [75], monitor posture [29], and even guide users through origami folding patterns [35]. FlexSense, developed by Rendl et al. [65], offers an effective approach for enabling diverse deformable interactions with rich continuous input. The authors took a novel approach with this device by printing their own piezoelectric bend sensors on plastic film to capture bending, rolling and flexing motions. They also developed two algorithms using computer vision and the physical properties of the plastic film for digital reconstruction of the input device. The resulting prototype allows accurate shape reconstruction and tangible feedback for 2.5D interactions when lying flat and 3D interactions during in-air use, showing potential for diverse PHF in VR.

Deformable devices exhibit important advancements in tracking needed to create digital recreations of their shape while avoiding occlusion. Some examples include internally based tracking methods with IMUs [29], piezoelectric bend sensors [48,65,75] and hinges with integrated potentiometers [35]. Hermanis et al. [29] created a prototype using a rectangle of thick fabric that tracks deformation using a grid accelerometer/magnetometer sensors. Using the position of each sensor allowed them to create real-time digital reconstructions to capture the device's shape. When compared to external sensing of the device using a Kinect V2 the device was only slightly less accurate, while also avoiding shortcomings of external sensing, such as occlusion and limited mobility.

2.6 Actuated Shape-Changing Devices

While unexplored in VR, we believe shape-change will allow richer haptic feedback through more expressive physical and visual interactions. Common approaches for achieving actuated shape-change include servo motors [28,47,53,63,68], shape memory alloys (SMAs) [24,27,58,68], particle jamming [19] and linear actuators [56].

Past research shows shape-change as an effective means to visually communicate a variety of emotional states through changes in their physical form [24,28,43,47,59,75]. MorePhone [24], a phone-sized e-ink display, is one example of successful visual communication through shape-change. Using shape memory alloy wire to actuate, the prototype can bend itself at each individual corner and along its horizontal center line. Researchers ran two user studies to assessing MorePhone's ability to communicate notification types and urgency through visual and tactile feedback. Their results show the device is effective at visually communicating notification type and urgency, however tactile input appeared to have an understated impact. Lindlbauer et al. [47] created a display integrating spatial augmented reality (AR) through projection mapping with a shape-changing interface by actuating folds in a piece of paper with servo motors. Projecting onto the shape-changing paper complimented 3D graphics through depth cues and showed increases in "realism" and immersion during an informal study.

In addition to visual communication, changes in shape facilitate different tangible experiences. Physical interactions with users through shape-change can add an expressive dimension through force feedback [24,54,58]. Complimentary shapes also have the power to enhance different digital interfaces [19,27,53,58,68], or facilitate exploration for new interactions with digital devices [53,56]. Park et al. [58] exhibited impactful communication with both visual and tactile communication through a single device. First, they developed Bendi, an I/O device allowing input though a low-profile joystick and shape-changing output with SMAs integrated into a rubber honeycomb structure. The researchers assessed Bendi with several romantic couples who used the device for mobile communication. Results show that participants felt comfortable communicating through Bendi and it was common for couples to construct their own vocabulary using tactile and visual communication.

2.7 Gesture Elicitation Studies

Wobbrock et al. [81,82] provided the initial structure for gesture elicitation studies. In their method participants are shown an action and asked to map a multi-touch gesture to it. Afterward, the researchers calculate an agreement score for each action, showing the level of consensus for the most preferred corresponding gesture. Further work by Vatavu & Wobbrock [79] improved the method for calculating agreement scores by increasing accuracy. Since its initial use of measuring consensus in multi-touch gestures, elicitation studies are now common throughout the field of human computer interactions. Lee et al. [45] and PaperPhone [42] show the adoption of gesture elicitation studies for tangible UIs allowing deformation. Researchers also used gesture elicitation studies to explore the manipulation of digital 3D objects in screen-based 3D UIs [8] and augmented reality [44,61]. Our study applied gesture elicitation studies to assess shape preferences for PHF enabled by HaptoBend.

2.8 Docking Task Studies

Researchers regularly use docking tasks to assess the performance of input modalities for 3D user interfaces including VR. A docking task will start by outlining a target location defined by a specified location, rotation, or both. To complete the task, a participant must move a virtual object to the target using the specified input modality, "docking" it at that location [5]. It is common to collect data on speed and accuracy during this process, allowing a comparison of different input modalities or other independent variables.

Work by Chen et al. [13] provides an example of early docking task studies. Their research compared four different 2D user interfaces for rotating 3D models. In a user study, they asked participants to complete docking tasks by rotating a 3D model of a house so that it matches the rotation seen in a copy of the model. Participants repeated a series of 9 dockings with each interface allowing the collection of data on speed and accuracy. For each interface, three target rotations were "simple" (involving only the x, y, or z axis), while the other 6 were complex (involving each axis). Their findings showed that the used of sliders to control rotation is more suitable for high precision with a simple task, while direct manipulation is faster and just as accurate for complex tasks.

Over time docking tasks have continued their popularity as an assessment for 3D manipulation devices in VR and other 3D user interfaces. Besançon et al. [5]'s recent study provided inspiration for our methodology seen in Chapter 5. They used a docking task study to compare a mouse, touchscreen, and PHF prop as input for moving and rotating a virtual object. In a comparative user study, participants used each condition to dock a 3D model of a teapot at a target location and rotation defined by a semitransparent copy of the teapot. During the study, they collected data on speed, accuracy, workload, fatigue and personal preference. Results suggest equal accuracy for each condition, however, the PHF prop took the least amount of time while the mouse took the most. PHF also resulted in the lowest workload, but they found no difference in fatigue between each condition.

2.9 Summary

Our review of related work reveals a solid base supporting our research. Past work in PHF, AHF, and novel approaches to haptics build a strong case for the benefits of haptic feedback in VR. PHF is an accessible way to enhance realism [11,22,33] while improving selection and manipulation in VR [5,32,38,46,74,77]. However, the use of PHF leads to an inconvenient amount of hardware for a variety of virtual objects and the need to switch between haptic props. Research in AHF produces impressively engineered devices that generalize to a variety of virtual objects and improve user performance in VR [1,7,51,57,60,69,78]. Unfortunately, AHF devices are usually too large, complex, expensive and/or weak for use outside of research. Through the acknowledgment of issues in PHF and AHF researchers continue to advance VR haptics by combining PHF and AHF [15,16,23], and introducing entirely new approaches [2,12,49,50,76]. While these approaches offer some noticeable benefits over traditional VR haptics issues like complexity, interaction limitations and weak feedback persist.

Informed by past work in VR haptics, our approach applies the benefits of deformable and shape-changing devices with the affordances of visual dominance in VR. Past work in visual dominance shows users are tolerant to differences between the physical and virtual world including mismatch between haptic feedback [70], manipulations of their rotation to the VE [64,73], and adjustments to their manipulation of virtual objects [4,40]. The research we reviewed on deformable plane-like devices exhibits their potential for VR by allowing a variety tangible experiences with one device [25] and collecting rich input for detailed interactions and tracking [29,65]. Past research shows actuated shape-changing devices map well to VR haptics by creating expressive tangible experiences [24,58] through physical and visual communication. To assess the benefits of applying deformation, shape-change and visual dominance to VR haptics, we drew upon the well-documented use of elicitation studies for user preferences, and docking task studies for device performance.

Chapter 3: Shape Elicitation User Study

3.1 Introduction

Our first study marks the first exploration into applying a deformable device to VR haptics, and begins our investigation into whether deformable, shape-changing devices offer a positive alternative to conventional haptic feedback approaches. Because of the original nature of our research, this study's goal was to establish a baseline of user preferences by assessing research questions in three main areas of interest:

1.1. **User Impressions:** In what contexts do users feel comfortable using a deformable haptic device meant for VR?

1.2. **Shape Approximation:** Do users prefer physical shape approximations of corresponding virtual objects for haptic feedback with a deformable haptic device?

1.3. **Contributing Factors:** What factors influence how users prefer to use a deformable haptic device in VR?

We use our HaptoBend prototype to investigate these questions with a device that provides haptic feedback for different objects by changing shape through deformation. HaptoBend leverages a simple design with a row of four rigid sections connected by hinges. The handheld device can transition from a flat plane to a variety of different shapes depending on how each section is bent. This mitigates prop-switching by letting users deform the device into a variety of physical shape approximations for haptic feedback that mimics different virtual object shapes. Past research observes visual dominance in VR allows these physical shape approximations to serve as realistic haptic feedback [4,18,41,70,84]. The design we chose takes inspiration from similar flexible plane devices [25,29,65], but maintains originality through a corresponding real-time digital reconstruction of HaptoBend meant for use in VR.

To address our first research question, we encouraged users to interact with HaptoBend in an exploratory manner, allowing us to gather design feedback and ideas on relevant applications for our approach. We gained insight on the second and third questions by performing an experiment modelled after Wobbrock et al.'s [82] gesture elicitation studies. To the best of our knowledge this marks the first elicitation study for PHF in VR. Participants selected their preferred shape for HaptoBend to map to different 2D, plane-like virtual objects, and multi-surface 3D virtual objects, allowing us to find overall preferences, satisfaction and influential factors that contributed to the haptic experience of a deformable haptic device.

3.2 HaptoBend Prototype

The construction of HaptoBend relies on a simple design that integrates four 1.5" x 5" rigid sections with hinged connections. Together, the panels create a bendable plane measuring 6" x 5" when lying flat and weighing 358.8 grams. Steel hinges used to construct HaptoBend give it a sturdy build, all of which we covered with a thin layer of black foam to create a comfortable surface texture. Our overall design draws from past flexible plane devices [25,29,65], however, none combine the same construction and sensing methods to create a digital reconstruction of the device in VR.

To represent HaptoBend (Figure 1) digitally in VR it incorporates three twist potentiometers located at each hinge axis to sense the bend angle of each panel and an Adafruit BNO055 IMU to sense yaw, roll and pitch of the entire device. The sensors all provided input to an Arduino Uno feeding serial data to a PC that integrates all the hardware together in Unity 5.5. A C# script utilized the SerialPort class to capture incoming serial data from the Arduino. Additional scripts written in C# for Unity apply the sensor data to a simplified 3D model of HaptoBend that creates a real-time virtual representation of the device. An Oculus CV1 VR headset displays this to users, integrated through the Oculus SDK for Unity.

3.3 Methodology

During the user study, we collected qualitative data on participants' first impressions by asking them to think aloud while familiarizing themselves with HaptoBend. Next, we guided participants through an exercise based on gesture elicitation to test which PHF shapes users preferred for a variety of virtual objects. Participants rated each preferred HaptoBend shape in terms of goodness and ease, as defined by Wobbrock et al. [81,82] to gain insight into the quality of their interactions. Participants then shared their final thoughts on the pros and cons of the device.

3.3.1 Participants

We recruited 20 participants, aged 21 to 38 years (μ = 27.8 years). Twelve participants were male, 7 were female and 1 answered “other”. The majority used VR, played video games and used 3D modelling software at least once a month. Four used VR daily, while 3 had never experienced VR before.

3.3.2 Prototype Setup

The Arduino Uno collecting data from HaptoBend’s sensors fed serial data to PC running Windows 10 (64 bit) with a 3.2GHz CPU, 8 GB of RAM, with a NVidia GeForce GTX1060 3 GB GPU. Participants viewed the real-time virtual representation of HaptoBend through an Oculus Rift CV1 head-mounted display. The setup allowed participants to interact with HaptoBend physically using the prototype and visually through their headset.

The VE depicted a simple scene consisting of a flat plane, a horizon, a simplified 3D model of HaptoBend (Figure 1), and, during the elicitation section, one of 6 virtual objects (Figure 3). The 3D model of HaptoBend reflected the bend angle of each panel and the device’s overall rotation in real time (Figure 1). We simultaneously displayed the “target” virtual object and HaptoBend’s 3D model when we asked participants to perform interactions between the two. The 6 virtual objects are depicted in Figure 3. We selected the virtual objects to provide three objects commonly used for 2D interactions, and three objects commonly used for 3D interactions, across three size categories: one roughly the same size as HaptoBend (called “medium”), one smaller, and one larger. The 2D models were a smartphone, a notebook and a large tablet, while 3D objects consisted of a pen, a flashlight and a sledge hammer. We used royalty-free 3D models for the virtual objects.

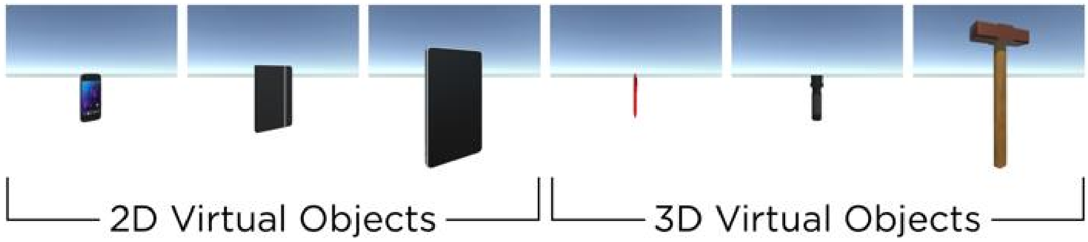

Figure 3. Virtual objects used in shape elicitation. Left to right: smartphone, notebook, tablet, pen, flashlight and sledge hammer.

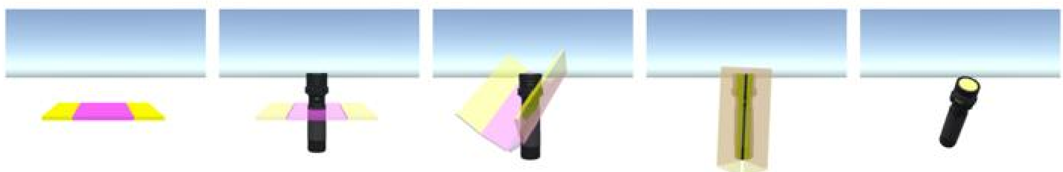

Figure 4. Example of HaptoBend mapping process to a virtual object: the virtual object appears (left); the participant determines their preferred shape (middle images), before notifying the experimenter that the shape is ready; the virtual model of HaptoBend disappears and the participant controls the rotation of the virtual object with the prototype (right).

3.3.3 Procedure

After participants completed a consent form and demographic questionnaire, they received a detailed description of HaptoBend as a flat plane with the ability to bend at its panel connections to create 3D shapes. The experimenter assisted each participant with fitting the CV1 correctly and ensured proper use of all devices during the study.

3.3.3.1 Think Aloud

We address our first research question through this section of the study by following the think aloud assessment presented by Ahmaniemi et al. [3]. The experimenter asked each participant to familiarize themselves with HaptoBend by contorting it into different shapes and brainstorming applications for the device in VR. During this time, the researcher gathered general qualitative data.

3.3.3.2 Shape Elicitation

The elicitation phase draws on the work of Wobbrock et al. [81,82] to inform our second and third research questions by assessing if HaptoBend's physical shape approximations of virtual objects create satisfactory PHF. Research by Gomes et al. [25] inspired us to examine if participants prefer using HaptoBend as a 2D shape for 2D virtual objects and 3D shapes for 3D virtual objects.

Upon starting the shape elicitation phase, target object models appeared one at a time co-located with the 3D model of HaptoBend (Figure 4). Following our instructions, participants held HaptoBend in the shape and orientation that they felt was most preferable for controlling the virtual object. Participants could choose any shape they wanted subject to the physical limitations of HaptoBend. They performed this task with the intent of using the target object as they would in the real world. Upon completing the task, participants notified the experimenter who pressed a key causing the HaptoBend model to disappear and applied its rotation to the virtual object. Figure 4 illustrates an example of the mapping process.

Similar to Wobbrock et al. [82], after selecting each mapping, participants verbally rated it in terms of goodness and ease on a 7-point Likert scale (1 = strongly disagree, 7 = strongly agree). For goodness participants rated the statement, "The shape I picked is a good match for its intended purpose". The "purpose" in this question referred to providing haptic feedback while using the virtual object as they would in the real world. Participants rated the statement, “The shape I chose was easy to perform” for ease. “Perform” in this question refers to performing the task of deforming HaptoBend into the participant’s most preferred haptic shape and mapping that shape to the corresponding virtual object. After completing both the think aloud and shape elicitation phases, participants completed a post-questionnaire, which asked them to record what they liked and disliked about HaptoBend.

3.3.4 Design

The shape elicitation phase employed a 2 × 3 within-subjects design with the following independent variables and levels:

- **Object type:** 2D flat objects, 3D multi-surface objects
- **Object size:** small, medium, large

These independent variables yielded the 6 different virtual object combinations (Figure 3). Each participant mapped HaptoBend to each of the 6 virtual objects once. Across all 20 participants, this resulted in 120 trials. To counterbalance fatigue and training effects, we randomized the order of the virtual objects for each participant.

We recorded three dependent variables during the elicitation phase: shape (the shape users deformed HaptoBend into), goodness (collected through goodness ratings described in 3.3.3.2), and ease (collected through ease ratings described in 3.3.3.2). We also calculated agreement scores using the process outlined by Vatavu & Wobbrock [79], as described in Section 3.4.3.

3.4 Results

We first present participants' impressions of HaptoBend gathered from the think aloud phase and post-questionnaire. Next, we report shape elicitation results.

3.4.1 Shapes

We allowed participants to reuse shapes for different virtual objects, as in Wobbrock et al. [82], which led to a total of 8 original shapes, illustrated in Figure 5. We classified four as 2D shapes and four as 3D shapes. Shapes received the designation "2D" if the intent of the shape was to create a single flat plane, while we classified shapes that utilized multiple intersecting planes as "3D".

Figure 6 shows shape-use frequency for each virtual object. Participants used Shape E the most for 3D shapes and overall with 28 uses. Totals for the rest of the 3D shapes amount to 24 for F, 1 for G, and 1 for H. Shape A showed the highest use of 2D objects with 21 uses followed by shapes B and C with 19, and 7 for D.

As expected, frequency of use changed to match the virtual object encountered. The most common shapes for the 2D virtual objects were Shape B for the smartphone, Shape C for the notebook and Shape A for the tablet. The most common shapes for the 3D virtual objects were Shape F for the pen, Shape E for the flashlight, and Shape E for the sledge hammer.

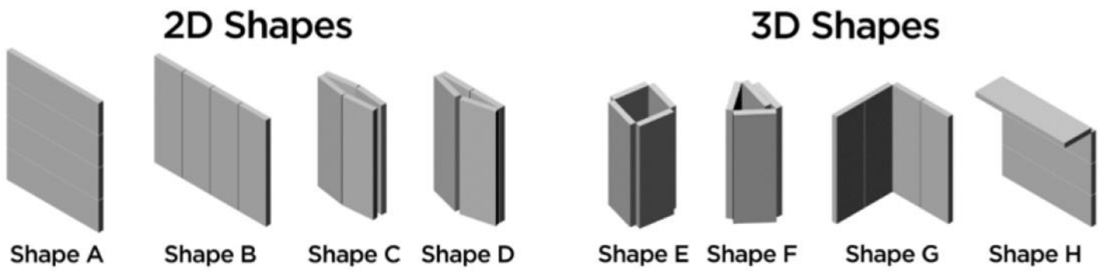

Figure 5. Shapes produced by participants during the elicitation section.

3.4.2 Agreement Scores

Agreement scores represent participant consensus in the shapes mapped to each virtual object. An agreement score of 1 means all participants chose the same shape for a given virtual object. Lower agreement scores indicate a greater variety in the shapes chosen for a given virtual object. Most past elicitation studies relied on the method proposed by Wobbrock et al. [81,82]. We employ an updated equation outlined by Vatavu & Wobbrock [79] which, unlike the previous approach, puts scores on a true 0-to-1 scale. As a result, agreement scores are lower, however, they are more accurate, allow calculation of coagreement scores, and enable statistical significance tests. Following Vatavu & Wobbrock [79], we can calculate agreement score (*AR*) with equation (1).

$$AR(r) = \frac{|P|}{|P|-1} \sum_{P_i \subseteq P} \left(\frac{|P_i|}{|P|}\right)^2 - \frac{1}{|P|-1} \tag{1}$$

For virtual object r, P is the total number of shapes participants used in the elicitation exercise and P_i is a set of identical shapes within P. Equation (2) shows an example of this equation in use to calculate the agreement score for the virtual sledge hammer, where participants selected three different shapes.

$$A_{hammer} = \frac{|20|}{|20|-1}\left(\left(\frac{10}{20}\right)^2 + \left(\frac{9}{20}\right)^2 + \left(\frac{1}{20}\right)^2\right) - \frac{1}{|20|-1} = 0.426 \qquad (2)$$

We calculated agreement scores for each virtual object as seen in Figure 7. Scores range from 0.216 to 0.489, with the smartphone receiving the lowest score, and the highest achieved by the flashlight.

To compare agreement scores we used Cochran's Q test as outlined by Vatavu & Wobbrock [79], which yielded 7 significantly different pairs of conditions. The smartphone's agreement score was significantly lower than the notebook ($V_{rd\,(1, N=40)} = 13.47$, $p < 0.001$), the tablet ($V_{rd\,(1, N=40)} = 8.12$, $p < 0.01$), the flashlight ($V_{rd\,(1, N=40)} = 31.44$, $p < 0.001$), and the sledge hammer ($V_{rd\,(1, N=40)} = 18.61$, $p < 0.001$) virtual objects. The agreement score for the pen was significantly lower than the flashlight ($V_{rd\,(1, N=40)} = 21.83$, $p < 0.001$) and the sledge hammer ($V_{rd\,(1, N=40)} = 10.92$, $p < 0.001$). Finally, the flashlight had a significantly higher agreement score than the tablet ($V_{rd\,(1, N=40)} = 7.84$, $p < 0.01$).

In addition, we calculated agreement scores for the broader categories of 2D shapes and 3D shapes for each virtual object. This facilitated an assessment of whether participants prefer 2D or 3D PHF shapes when mapping to virtual objects intended for 2D vs. 3D interactions. All objects received an agreement score of 1 except for the pen (score of 0.605) and the flashlight (score of 0.900). These results give insight into our second and third research questions by showing high consensus for mapping 2D virtual interactions to 2D shapes and 3D virtual interactions to 3D shapes.

3.4.3 Goodness and Ease Ratings

We used both goodness and ease ratings to examine our second research question concerning the level of satisfaction users experienced with haptic feedback from HaptoBend. Participants rated the shapes produced for each virtual object in terms of **goodness**, i.e., quality of the mapping. Overall ratings for all the shapes were positive, except for Shape G. Goodness scores are seen in Figure 8. These ratings show positive ratings for all objects, with the flashlight and sledge hammer receiving only positive ratings.

To compare overall goodness ratings between virtual objects we used the Mann-Whitney U test. The goodness ratings for the pen were significantly lower than all the other objects: smartphone ($U = 94.0$, $p < 0.05$) notebook ($U = 94.0$, $p < 0.05$), tablet ($U = 88.5$, $p < 0.005$), flashlight ($U = 61.0$, $p < 0.001$), and sledge hammer ($U = 84.5$, $p < 0.001$),

Ease ratings allow an assessment of HaptoBend's ability to deform into a participant's desired shape. We also summed ease rating results for each shape-object mapping (Figure 9). The only virtual objects showing a negative ease rating are Shape G mapped to the pen and Shape A mapped to the tablet.

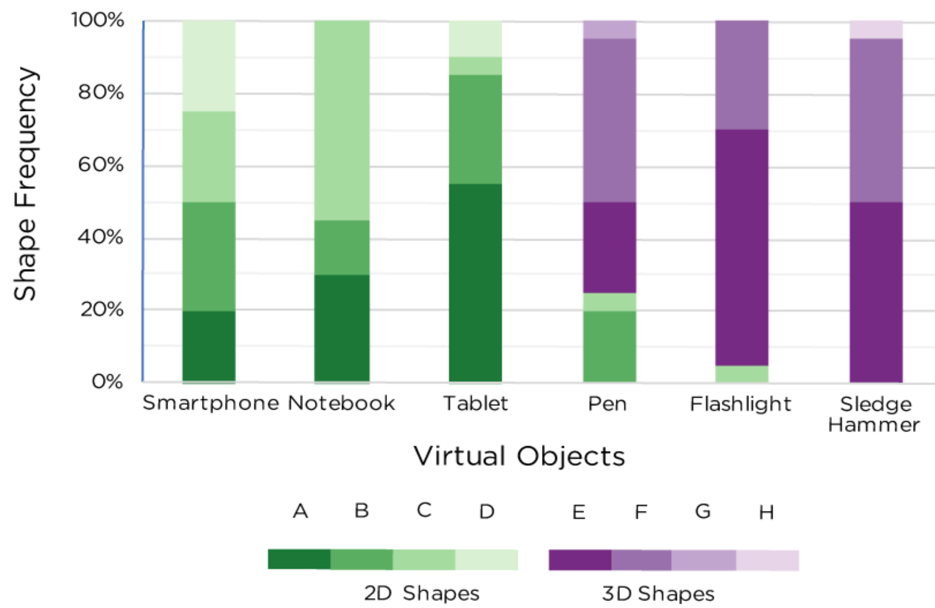

Figure 6. Frequency of use for each shape as mapped to each object type and size.

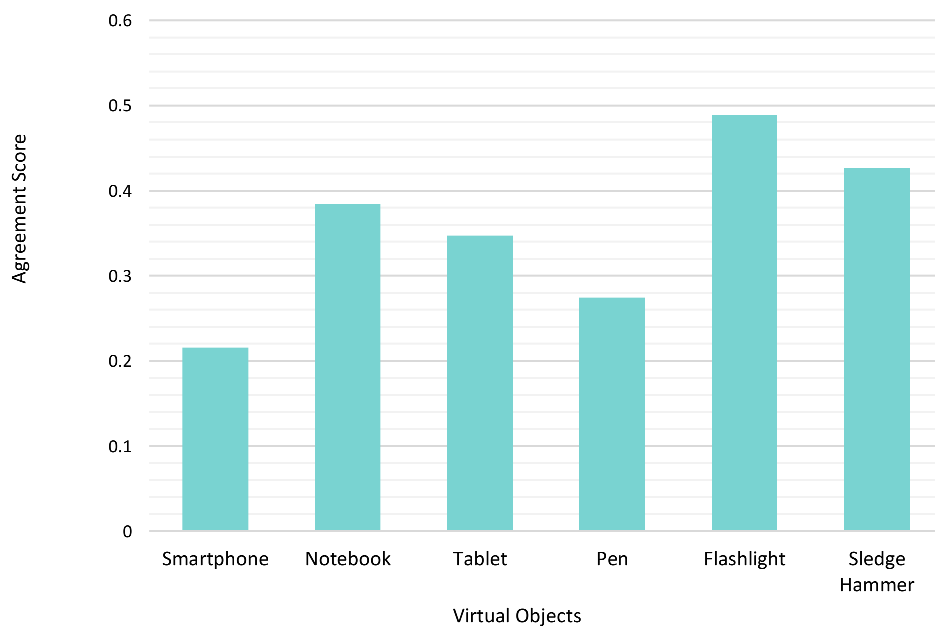

Figure 7. Agreement scores for each virtual object.

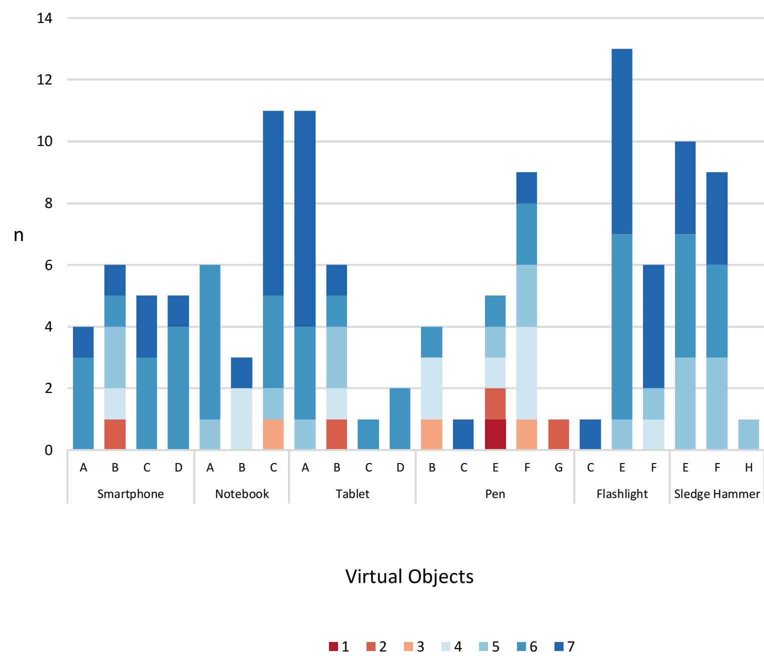

Figure 8. Goodness rating summed for each shape within each virtual object presented to participants.

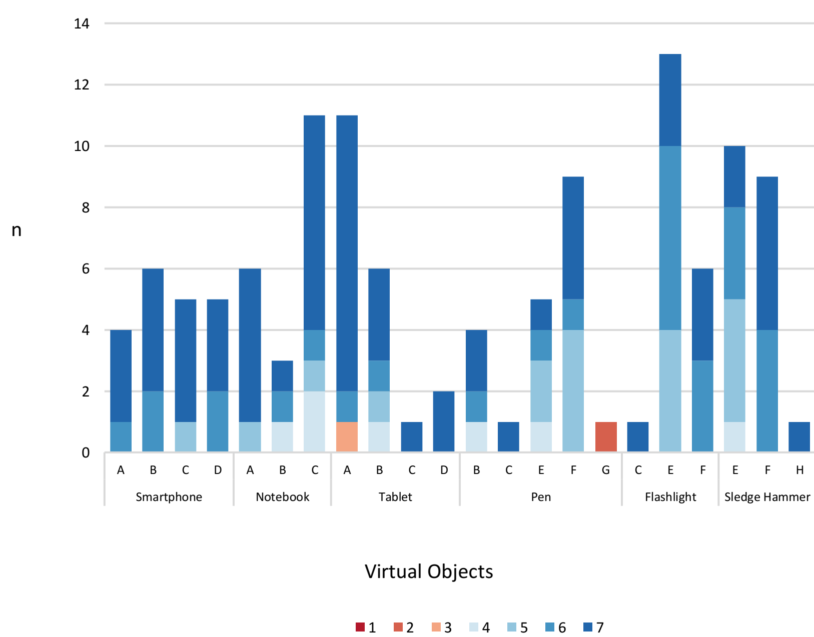

Figure 9. Ease rating responses for ease summed over the entire elicitation phase for each shape.

3.4.4 Final Thoughts Questionnaire

Upon completing the experiment, participants shared their overall thoughts on HaptoBend. Positive feedback included 10 participants praising HaptoBend's responsiveness, followed by 9 enjoying the foam texture, 8 valuing its ability to bend, and 8 appreciating the digital model of the device. Six participants also saw its diversity of application as a plus. In terms of negative feedback, 6 participants noted the limits of including only three hinged areas, 5 expressed dissatisfaction with the inability to fold the device completely flat, and 5 saw the size difference between some virtual models and the device as a negative.

3.5 Discussion

Our study revealed a large variety of applications that show promise for a deformable haptic device. High goodness and ease ratings show participants generally felt positive about using physical shape approximations of virtual objects as their preferred PHF shapes. Other than shape, we observed several other factors, such as perceived functionality and weight, that had notable impacts on participant preferences. These findings support HaptoBend as a simple, mobile and more accessible alternative to large, complex and costly general purpose haptic systems.

3.5.1 Question 1: User Impressions

Participants expressed excitement during the Think Aloud portion of the study, describing HaptoBend as enjoyable and easy to use. When asked to think of possible applications for HaptoBend, most were already coming up with virtual objects it could physically represent in VR, the most popular being a book. We classified the proposed applications for HaptoBend into 5 broad categories: 3D modelling, 3D manipulation, games, education, and haptic stand-in.

The variety of suggestions shows a high level of optimism for the device's usefulness. Suggestions for 3D modelling included creating 3D primitives with HaptoBend to quickly assemble models in CAD or other modelling software, and allowing collaboration with several people each using their own HaptoBend. 3D manipulation ideas incorporated both manipulating entire virtual objects to view each side and manipulating joints within a model for animation purposes. Game ideas included using HaptoBend for multiple handheld tools within a game, such as transitioning between a weapon and a book or for a single static object used throughout a game. Educational suggestions for HaptoBend referenced using it to teach geometry, and leveraging it in the context of hand rehabilitation exercises allowing the sensors to give important feedback to users. The haptic stand-in category encompasses suggestions for using HaptoBend as a physical proxy for specific virtual objects. Many of these suggestions leveraged the dynamic nature of the hinges with objects like a book and a laptop.

3.5.2 Question 2: Shape Approximation

3.5.2.1 2D vs 3D Shape Mapping

In our assessment of haptic shape preferences for different virtual objects, participants deformed HaptoBend into their preferred shape to map to different 2D, plane-like virtual objects, and multi-surface 3D virtual objects. The division between these categories allows for insight into whether our participants preferred physical shape approximations of virtual objects. As one might expect, our results across all measures strongly suggest that participants did prefer 2D PHF shapes for 2D virtual objects and 3D PHF shapes for 3D virtual objects. These findings are especially prevalent in the frequency of shapes used for each virtual model and the agreement scores comparing 2D and 3D shapes.

There is also support showing the use of 2D shapes for 3D objects may have a negative effect. In general, participants mapped a strong majority of the virtual objects to shapes of the same dimension. The two exceptions to this were the pen and the flashlight, with 25% and 5% of participants mapping them to 2D shapes respectively (Figure 6). The resulting goodness and ease ratings for the pen are the lowest of all the objects. Observing this type of behavior points to the importance of similar physical shapes for haptic feedback in VR allowed by deformable PHF devices, like HaptoBend.

3.5.2.2 Approximating Shapes

We used a relatively simple and inexpensive design for HaptoBend, based on Simeone et al. [70] who showed physical approximations of virtual objects produce satisfactory PHF. The results from our elicitation study support Simeone's findings [70]. Participants consistently rated approximate shapes made with HaptoBend as good PHF for more detailed virtual objects. These results also align with findings from Aguerreche et al. [2], which supports the use of physical shape approximations for PHF. One participant summed the effects of shape approximation best when mapping to the pen stating: "Even the lack of roundness doesn't really matter. What matters more is that it feels like I'm holding something sort of elongated barrel shape in my hand." The quote is especially significant when one considers that the pen yielded our weakest overall results.

The flashlight performed especially well with the highest agreement score and the high goodness ratings. Seven participants even mentioned no noticeable difference between HaptoBend's angular shapes and the cylindrical flashlight. One went as far as saying, "I don't think you could get any closer to the shape (of the virtual object)" and another stated, "this feels like a flashlight." The sledge hammer, and notebook received less pronounced but similar results, with agreement scores and goodness ratings that were not significantly different from the flashlight.

Some of our virtual objects illustrate possible limitations to how strongly vision dominates touch, as reported by Simeone et al. [70]. The tablet, smartphone and pen all received significantly lower agreement scores than the flashlight. Even with less impressive results these virtual objects still received high goodness ratings and encouraging comments from participants. Enabling shapes that more closely matching the size and shape of these virtual objects might lead to improvements here.

3.5.3 Question 3: Contributing Factors

While our approach focuses on creating shape approximations of virtual objects, we acknowledge that other physical characteristics may impact user preferences for haptic devices. Looking at our results, we found a participant's perception of a virtual object's function, size and weight also had potential to influence their use of HaptoBend.

3.5.3.1 Perceived Function

We observed that the perceived function of different virtual objects influenced users' most preferred haptic shaped. One example of this appeared from investigating why the smartphone resulted in the lowest agreement score. HaptoBend's physical constraints appear to be an influence here as they limit the device's hinges from rotating a full 180°. As a result, the Shapes C and D, the closest in size to the smartphone, could not fold completely flat. Nine participants mentioned this physical constraint as a problem. Four commented they would have chosen shapes C or D, but instead chose larger Shapes A or B to achieve a completely flat shape. Since the functionality of a smartphone is dependent on using a flat touchscreen, participants had to choose between a shape similar in size, or a shape perceived to more closely fit the function of this virtual object.

Functionality factored into the participants' opinions of the tablet as well. Using shape A, the closest in shape to a tablet, positioned HaptoBend's wire connections in a conflicting location for conventional tablet grip positions [80]. As a result, some participants avoided this shape.

3.5.3.2 Size

Size may also have been a factor for the tablet as all shapes enabled by HaptoBend were smaller than it. Of all the virtual objects, size impacted the pen most. Eleven participants described HaptoBend as too large to map well to it, which yielded significantly lower goodness ratings compared to all other virtual objects.

3.5.3.3 Weight

While HaptoBend's design allowed changes in shape and size, it does not support changes in weight. As described earlier, Zenner et al. [84] showed the importance of PHF objects approximating a virtual object's expected weight. Our results suggest that HaptoBend is still able to provide PHF for a variety of virtual objects, even without dynamic weight distribution like Zenner et al.'s Shifty [84]. Earlier work by Zenner [83] provides insight into these observations by describing some level of tolerance to weight differences for PHF objects. While we observed some level of disparity, HaptoBend does not differ drastically from the real-world weights of any 2D virtual objects we used. Participants also seemed to tolerate these weight differences well as they rarely mentioned them during the elicitation study.

The range of 3D objects had a larger weight disparity with HaptoBend than the 2D objects. The weight difference with the pen and sledge hammer are particularly pronounced. In general participants felt the weight of HaptoBend and the flashlight was similar, leading to mostly positive comments. The sledge hammer would be far heavier than HaptoBend, however, participants had mixed opinions on this: 2 mentioned weight positively, and 2 mentioned weight negatively. The pen was the only virtual object where participants noticed a pronounced difference in weight. Six participants mentioned HaptoBend was too heavy for this virtual object, contributing its lower goodness ratings.

3.5.4 Future Improvements

The high goodness and ease ratings achieved by HaptoBend point to a high potential for deformable devices to provide PHF in VEs. Overall, mimicking a virtual object's shape appeared effective in emulating users' expected haptic feedback. These results align with past work from Ninja Track [39] and Aguerreche et al. [2] who took similar approaches by emulating the shape of different real-world objects for digital interactions. Future work should test these findings further with virtual objects that have a larger variety in size and shape. A greater variety of virtual objects would allow richer insight into the ability of HaptoBend to produce realistic PHF through approximations of shapes and assess the importance of differences in function, size, weight and other physical properties.

Poor performance in agreement for the small virtual objects points to a need for higher resolution shapes by dividing HaptoBend into more panels, or replacing them with a flexible material. Higher resolution would especially improve PHF for smaller and more intricate objects. Another important factor for resolution is hinges that allow 180°, or even 360° rotation, for fully flat bends. In combination, these two improvements would alleviate much of the negative shape feedback HaptoBend received. Six participants recommended adding a feature that locks HaptoBend's panels to prevent the devices from changing shape once it is mapped. A shape-locking feature would also increase functionality by creating physical consistency for interactions.

At points in the study where participants used HaptoBend to control virtual objects, they were eager to use those objects for their expected functions. However, we note that to support such functions, we would have to add additional sensors. For example, capacitive touch sensors would enable (simulated) touch screen interactions. Adding a 3D position tracker would also facilitate richer spatial interaction. Applying flexibility to normally ridged objects to increase interactions through bend gestures also gained support by the suggestion of 6 participants.

Functionality also seemed to suffer from the wires connecting HaptoBend's sensors to the Arduino Uno. A future version of the device could use wireless data transmission (e.g., via Bluetooth) to eliminate this problem, and may yield a better experience. These modifications could lead to mobile version of HaptoBend with the potential for augmented reality applications.

3.6 Summary

Our first study examined HaptoBend, a deformable haptic feedback device for VR, to investigate into user preferences for this new approach to VR haptics. HaptoBend used a simple design to create physical shape approximations of virtual objects that draws on visual dominance to provide satisfactory haptic feedback in VR. Due to the originality of our approach we performed a user study focused on establishing a baseline of user preferences with questions addressing three main areas of interest: *user impressions*, *shape approximation*, and *contributing factors*.

To answer these questions, we incorporated a methodology that gathered qualitative feedback through exploring HaptoBend and user preferences with an elicitation task. We found that user impressions of HaptoBend were positive and users expressed interest in using the device for a large variety of applications. High goodness and ease ratings along with user feedback suggests HaptoBend's physical shape approximations provided satisfactory haptic feedback for the majority of virtual objects we tested. Besides shape, our participants also revealed factors other than shape contribute to their most preferred shapes when using HaptoBend. The most prominent of these were the perceived function of an object, weight and size. Our findings also build evidence for the use of physical shape approximations with deformable devices as a legitimate solution to current issues with traditional PHF and AHF. The positive performance of HaptoBend points to a bright future for deformable PHF devices in VR, with many areas open for further research.

Chapter 4: Adaptic Prototype

4.1 Introduction

To further explore whether shape-changing devices offer a positive alternative to conventional haptic feedback approaches, we designed a new prototype called Adaptic that *adapts* to different *haptic* needs (see Figure 2 and Figure 10). After our elicitation user study with HaptoBend, we combined our findings with past research to create a Design Criteria that guided our development of Adaptic. Like HaptoBend, Adaptic addresses over complexity through a simple design that focuses on creating physical shape approximations to mitigate the need for multiple props. Adaptic also allows increased functionality including self-actuated shape-change, translation tracking, and double hinge connections that allow more shape options.

Figure 10. Adaptic in several shapes showing the affordances of its hinged connections.

4.2 Design Criteria

The Design Criteria we constructed outlines a guide to aid further development of deformable, shape-changing devices for VR haptics. As a base our Design Criteria uses foundational research from Holman & Vertegaal [34] that describes best practice requirements for organic user interfaces (OUIs), which encompasses deformable and shape-changing devices. While well recognized, their research does not speak specifically to the original context of VR haptics. As a result, their requirements do not account for issues of *complexity*, *limited interactions in VR*, and *inadequate haptic feedback* that we noted in Chapter 1. To account for VR-specific issues, we integrated design feedback from HaptoBend into our Design Criteria, allowing us to focus directly on the application of deformable, shape-changing devices to VR haptics.

Before creating our Design Criteria, the only research exploring deformable haptic devices in VR was our evaluation of HaptoBend, and no research had explored shape-changing haptic devices for VR. As a result, we do not evaluate our design criteria with past research. Instead we hope it will act as a first step towards formalizing best practices for deformable and shape-changing VR haptic devices that future research in this area will continue to refine.

The base of our design criteria originates from Holman & Vertegaal's [34] work, which defines best practices for OUIs. Their work describes OUIs as "promoting well-being through diversity of posture and ergonomic fit," further defined by 3 requirements:

- **Input Equals Output:** To achieve physical immersion, OUIs sense manipulations of their shape and connect that to complimentary visual, auditory and haptic output.

- **Function Equals Form:** OUIs leverage physical metaphors to form intuitive digital interfaces.

- **Form Follows Flow:** OUIs allow changes in shape to compliment different contexts of use.

To ensure our Design Criteria focused on the context of VR haptics for deformable, shape-changing devices, we looked to design feedback from our elicitation study with HaptoBend. Participant feedback from this study identified several areas of improvement that we observed repeatedly throughout the study. We classify these into 4 categories for our Design Criteria to account for:

- **Shape-Locking:** Some haptic shapes were hard to maintain due to hinge movement. As a result, participants asked for a shape-locking option that would keep hinges in place when necessary.

- **Shape Possibilities:** HaptoBend's physical constraints limited its hinges to rotating only 150° from a flat, neutral state. In response, participants asked for more bending options to expand shape possibilities.

- **Minimize Conflicting Hardware:** The external location of HaptoBend's potentiometers and their associated wiring conflicted with some participants' preferred hand positioning.

- **Richer sensing:** Participants listed several additional input modalities for improving HaptoBend, most popular of these was 3D translation tracking.

By combining the work of Holman & Vertegaal [34] with design feedback from HaptoBend, we constructed our Design Criteria for developing deformable VR Haptic devices. Our result outlines 4 requirements:

- **Comfort:** In the interest of ergonomics and freedom of interaction, a user should be able to comfortably hold the device in their hands using a variety shapes and orientations.

- **Simplicity:** The device's design should remain mechanically simple enough that it is inviting to the user and should not limit users from performing other interactions in VR.

- **Shape diversity:** For adequate haptic feedback in VR, shape approximations are sufficient, but a variety of them need to be possible for a diversity of virtual object shapes.

- **Responsiveness:** Tracking of the shape and orientation should be accurate and responsive enough to maintain a high level of immersion.

4.3 Developing the Prototype

Using our Design Criteria, we developed Adaptic to address the three categories of issues in VR haptics.

4.3.1 Form Factor

To adhere to the comfort and simplicity design considerations, we maintained a form factor similar HaptoBend with four rigid elliptic flattened cylinders measuring 41mm x 20mm x 206mm, each covered by a layer of foam. To connect one rigid panel to another we used double-hinged connections (see Figure 10) that allow full 360° rotation. Enabling this range of rotation means that each section can be folded perfectly flat on top of its neighbor, resulting in more complex shape options when compared to other hinged devices [25,52]. For example, the single hinged design used in HaptoBend did not allow fully flat bends between sections [52]. We designed the panels to be modular and 3D printed them with PLA filament for assembly with metal screws. Together, the panels create a bendable plane measuring 200mm x 173mm when lying flat and weighing 494.8 grams.

Each hinged connection contains a 10k potentiometer to measure its current position. It can also lock in place and actuate to a desired angle using two Tower Pro MG92B 14g micro servos which are rated to travel 60° in 0.08 sec. and supply 3.5kg of torque. The result allows each hinge to function in three different modes (Figure 10):

- **Deformable:** allows the user to freely manipulate the device's shape.

- **Shape-locking:** prevents bending along specified hinges to mimic the physical attributes of a virtual object.

- **Shape-changing:** actively bends the device to a specified shape and provides animated haptic feedback. Actuation from flat to a compact wand-like shape takes approximately 2 seconds.

4.3.2 Tracking

Following Holman & Vertegaal's [34] "input equals output" recommendation, Adaptic tracks its overall rotation and each hinge angle. Our tracking method allows for a real time digital model of the device's current shape, which users can interact with in VR. Two different methods are available for tracking Adaptic in 3D space. The first allows both rotation and translation tracking through a Razor Hydra's magnetic tracking system. To use the Hydra's system, we disassembled one of the Hydra controllers to extract its tracking module. We then put the tracking module into one of Adaptic's panels to track it. Unfortunately, engaging Adaptic's servo motors interferes with the magnetic tracking system, meaning shape-changing and shape-locking can't coincide with 3D translation tracking. The second method uses a BNO055 IMU that includes an accelerometer, gyroscope, and magnetometer. The IMU is limited to 3D rotational tracking, but Adaptic's servo motors do not interfere with it, therefore shape-changing and shape-locking are available.

4.3.3 Pipeline

A Teensy 3.5 handles data collection from each sensor and communicates with each servo. The Teensy communicates with a PC running Unity via a custom C# script utilizing the SerialPorts Class to import sensor data and control the servos. To integrate the Hydra tracking system into Unity we used the SixenseUnityPlugin SDK. In addition, we wrote C# scripts for Unity to translate sensor data into a real-time digital representation of Adaptic and control each servo individually. We can also map all the sensor data captured from Adaptic to other virtual objects, allowing users to manipulate them in 3D space with Adaptic as haptic feedback.

4.4 Connection to Design Criteria

We used our Design Criteria to inform our development of Adaptic, with the goal of improving our approach for haptic feedback. To better explain our strategy for satisfying the Design Criteria, we describe it below for each requirement.

Comfort: The size of Adaptic fits in users' hands in a variety of shapes (Figure 10). However, accommodating the internal servos meant that Adaptic was slightly larger than HaptoBend, creating limits on the range of virtual objects it can provide haptic feedback for.

Simplicity: We used a simple and unintimidating design for Adaptic by we embedding all sensors and actuators inside Adaptic and minimized the presence of external wires. Deforming the device is easy to understand in VR given the virtual representation of Adaptic paired with real-time tracking.

Shape diversity: Users can achieve a variety of shape approximations with Adaptic. Features like the double-hinged connections expand its range of shapes over other hinged devices like HaptoBend and PaperFold [25]. Shape locking also allows Adaptic to provide more consistent and robust haptic feedback in comparison to HaptoBend [52].

Responsiveness: We find the tracking of Adaptic to be both responsive and accurate. Constructing a real-time digital representation of Adaptic allows a pleasant connection between the haptic experience and what users visually observe in a VE. Three-dimensional rotation and translation tracking also allow users to interact with Adaptic in VR as they would with real-world objects.

4.5 Conclusion

In this chapter, we outlined two main contributions in our work towards mitigating these issues. The first contribution is our Design Criteria, which provides a framework for continuing to integrate shape-changing, deformable devices in VR. To develop our Design Criteria we incorporated design feedback from our study with HaptoBend and foundational research in the area of OUIs [34] for an informed result. Our second contribution comes from applying our Design Criteria to create Adaptic, an original approach to mitigating issues in VR haptics by applying the affordances of deformability and shape-change. When combined with visual dominance in VR [4,18,41,70,84], these affordances allow Adaptic to create physical shape approximations for a variety virtual objects.

Chapter 5: Docking Task User Study

5.1 Introduction

In the second user study, we furthered our understanding of deformable, shape-changing devices as a positive alternative to conventional haptic feedback approaches. We employed our Adaptic prototype to expand on our past research by comparing our use of a deformable, shape-changing device, with other approaches to VR haptics. While Adaptic has shape-changing capabilities, it is important to note that this study does not specifically evaluate shape-change, which means we cannot apply our findings directly that approach. However, we do use Adaptic in a variety of static shapes that it can actuate into, allowing us to gain some idea of how a shape-changing device might perform by proxy. To compare performance between haptic approaches, we rely on three main research questions:

2.1. **Design Criteria:** Does Adaptic satisfy the Design Criteria we outlined for deformable, shape-changing VR haptic devices?

2.2. **Performance:** How well does a deformable haptic device perform compared to other haptic approaches in VR?

2.3. **Preference:** Do users prefer a deformable device over alternative approaches for VR haptics?

The user study we carried out to address these questions marks the first docking task study to explore the performance of a deformable device for VR haptics. We test each approach against our **Design Criteria** by asking targeted questions on each point of the criteria for each haptic approach. To measure **performance,** we collected data on the time and accuracy of each docking, and asked questions about participants' experience with each approach. **Preference** measurements consisted of asking participants' to rank each of the approaches at the end of the study. Combined, these measures give a rich picture of the performance of Adaptic compared to other haptic approaches and continue to further our understanding of deformable devices in the context of VR haptics.

5.2 Methodology

To compare Adaptic to other haptic approaches, we tested its performance along with two other approaches using a docking task based on the work of Besançon et al. [5]. Our study focuses on the impact of shape disparity between a virtual object and its corresponding haptic device. Therefore, the other haptic approaches we tested alongside Adaptic were *static controller*, an off the shelf, conventional VR controller that assumed the same shape for each object, and the *haptic props* approach, which consisted of using a similar real-world object for each corresponding virtual object. These approaches cover a full spectrum with a single shape for all virtual objects using the *static controller* approach, Adaptic's shape approximations used for the *deformable* approach, and near-identical shapes used for the *haptic props* approach (see Figure 11).

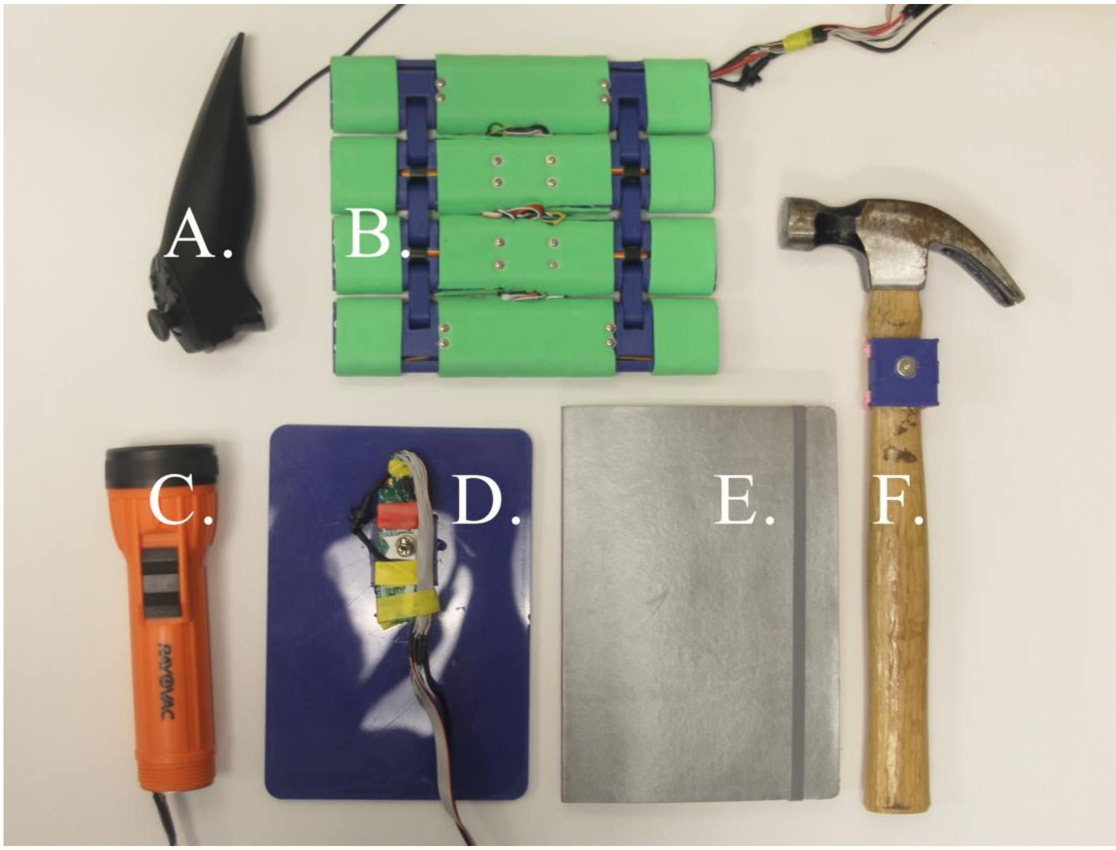

Figure 11. All haptic devices used during the user study. The Razer Hydra controller (A.) fulfills the static controller condition. The deformable condition uses Adaptic (B.). The haptic props condition uses a real flashlight (C.), a 3D printed tablet model (D.), a real notebook (E.), and a real hammer (F.). The tracking component extracted from the Hydra controller is seen on the 3D printed tablet (D.).

5.2.1 Participants

We recruited 23 participants with an average age of 23.43 years ($\sigma = 4.59$), 13 were male and 10 were female. Each received $15 for participating in the study. The majority were students, with 16 responding as undergraduate students, and 4 answering that they were graduate students. Most had little to no experience using VR in the past with 14 experiencing VR approximately once a year and 3 having never experienced VR. To capture an understanding of their experience with 3D interfaces we also asked how often they play video games and use 3D modelling software. Experience with video games varied fairly evenly for participants: 2 played more than once a day, 4 played once a day, 5 played once a week, 5 played once a month, 4 played once a year, and 3 had never played video games. Overall, participants were not very experienced with 3D modelling software: 7 used them once a month, 9 used them once a year and 9 having never used them.

5.2.2 Hardware Setup

Each haptic approach required a different combination of hardware. For the *static controller* approach, we used a single Razor Hydra controller, for the *deformable* approach our Adaptic prototype served as haptic feedback, and the *haptic props* approach incorporated four real-world objects corresponding to the four virtual objects tested. Our tracking system permitted us to co-locate each haptic device with any of the virtual objects, allowing participants to move and rotate an object with a one-to-one mapping.

To maintain a consistent experience, we used the Razor Hydra's magnetic tracking system (Figure 11) to track the rotation and translation of all the haptic approaches. Using the Hydra's tracking system for Adaptic and the haptic props required disassembling several Hydra controllers to remove their tracking modules. The small form factor of the tracking module (Figure 11) allowed us to fit it discretely inside of Adaptic and some of the haptic props, for the other haptic props we placed the module on the surface of the object in a location that would not interfere with hand positioning. The use of magnetic tracking also avoided issues of occlusion found in other tracking systems, which were of especially high concern for Adaptic given different areas become occluded depending on what shape it is in.

There were two issues that did arise from using magnetic tracking. The first involved using an actual tablet as a haptic prop. The tablet's hardware interfered with the magnetic tracking creating noticeable inaccuracies. To fix this problem we 3D printed a copy of the tablet (Figure 11) to maintain an accurate shape while avoiding the metal components that affected the magnetic tracking system. To print the tablet, we used a 3D model that was slightly different than the tablet used with HaptoBend. The other issue we encountered was interference from powering Adaptic's servo motors. Engaging the servos rendered the tracking unusable, as a workaround to this we decided not to power the servos during the docking task. The stiffness of the servos allowed to Adaptic to easily stay in all the desired haptic shapes except for the completely flat shape used for the tablet, which required some intentional support by the participants.

To integrate the Razor Hydra tracking system into our VE we used the SixenseUnityPlugin SDK along with custom C# scripts in Unity 2017.1.0f3. We ran Unity on a PC running Windows 10 (64 bit) with a 4.20GHz CPU, 32 GB of RAM, and a NVidia GeForce GTX 1080 8 GB GPU. To display the VE participants wore an Oculus CV1 head-mounted display. Figure 12 shows our setup in use.

5.2.3 Procedure

After participants completed a consent form and demographics questionnaire, they received a detailed description of the docking task, and each of the haptic approaches we asked them to use. We described Adaptic as a flat plane with the ability to bend at into 3D shapes and create physical shape approximations for haptic feedback. Each participant then received assistance from a researcher to achieve a satisfactory fit with the CV1.

For each of the three haptic approaches a participant completed a docking at 9 different docking positions (see Figure 14) with four different virtual objects for a total of 108 dockings (3 haptic approaches * 9 docking positions * 4 virtual objects). The study began with one haptic approach for which a participant completed 9 dockings in succession for each of the virtual objects. They then repeated same process with the second haptic approach chosen, followed by the third. We randomized the order of haptic approaches, docking positions and virtual objects to counterbalance fatigue and training effects.

The virtual objects were a subset of those used in the elicitation study with HaptoBend consisting of a hammer, a flashlight, a tablet and a notebook (see Figure 11). We made some adjustments to the size of the virtual objects with HaptoBend so that they better matched our haptic props. For the *deformable* approach a researcher manipulated Adaptic into a different shape for each virtual object (see Figure 13), which corresponded with the most preferred shape for that object in the elicitation study.

The subset of virtual objects (Figure 13) did not include the pen and smartphone virtual objects. During the elicitation study participants complained about the size disparity between the pen and HaptoBend, which was even more pronounced with Adaptic. The smartphone also received complaints due to HaptoBend's hinge rotation limitations. Both issues affected participants' most preferred shapes for the pen and smartphone, and led them to produce the lowest and second-lowest agreement scores respectively. With these low levels of agreement, we felt uncomfortable assigning a single haptic shape to either of these objects for this study, causing use to exclude them.

5.2.3.1 Docking

For one “docking” a participant used a haptic device to align a virtual object with a specified docking position defined by a target location and target rotation. To show the docking position we placed a semi-opaque copy of the current virtual object for participants to align with, meaning a perfect overlap with the copy would result in 100% accuracy. Because some of the virtual objects had rotational or near-rotational symmetry we also augmented them and the targets with color-coded XYZ coordinate arrows as an orientation reference.

Each docking task began with a seated participant holding one of the haptic devices in front of them at a specified starting position. We told participants to inform us when they felt they had achieved a satisfactory level of accuracy for the docking location, with time and accuracy weighted evenly. A researcher began each docking with the verbal command “go”, allowing the participants to begin aligning the virtual object with the docking location. Once satisfied, a participant would verbally indicate it to the researcher allowing them to stop the timer and record the virtual object’s location and rotation.

5.2.3.2 Questionnaires

Similarly to elicitation studies [82], each time a participant completed all 9 dockings for a virtual object a researcher asked them to rate match goodness and accuracy ease, both on a 7-point Likert scales (1 = strongly disagree, 7 = strongly agree). Match goodness asked participants to rate the statement, “The controller I used for this virtual object was a good match for this task." For accuracy ease, participants rated the statement, "It was easy to achieve the level of accuracy I wanted."

Once a participant completed all the dockings for a haptic approach, we asked them to complete 2 questionnaires on that specific approach. The first questionnaire was the NASA Raw Task Load Index (RTLX), to measure workload with a series of 20-point Likert scale questions (1 = very low, 20 = very high) on mental demand, physical demand, temporal demand, performance, effort, and frustration. Following the process of Besançon et al. [5], we opted for the RTLX instead of the full NASA TLX as they found the additional section to be lengthy and confusing. They also cite Hart's survey [26], who indicates that the RTLX is similar in effectiveness when compared to the full TLX. A Heuristics Questionnaire followed the NASA RTLX, with the purpose of testing how well each haptic approach fit our Design Criteria. The questionnaire targeted each point of criteria with a statement:

1. **Comfort:** I found the haptic approach comfortable to hold for all virtual objects.

2. **Simplicity:** The design of this haptic approach felt inviting to use and would not limit my interactions in virtual reality.

3. **Shape diversity:** I felt this haptic approach provided good shape for each of the virtual objects in the study.

4. **Responsiveness:** The responsiveness of this haptic approach allowed me to feel fully immersed in the virtual environment.

Participants rated each on a 7-point Likert scale scales (1 = strongly disagree, 7 = strongly agree).

5.2.3.3 Ranking

Again following Besançon et al. [5], after completing the docking tasks with all haptic approaches, participants completed a questionnaire to determine which approach they preferred most. The researcher informed each participant they need to complete 15 more docking tasks and they needed to choose their most preferred haptic approach to use, as well as their second and third most preferred approach. After recording their ranking, we revealed that no more docking tasks were necessary. According to Besançon et al. [5], this approach allows a better picture into which approach users preferred in the context of actual use and filters out other contributing factors such as the novelty of a device.

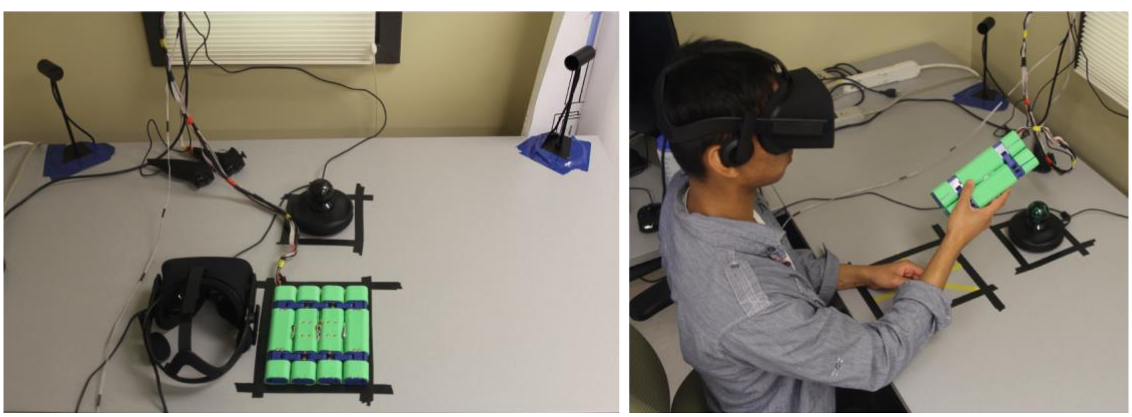

Figure 12 Experiment setup with Adaptic including the Oculus headset and its tracking system, and the Razer Hydra tracking system.

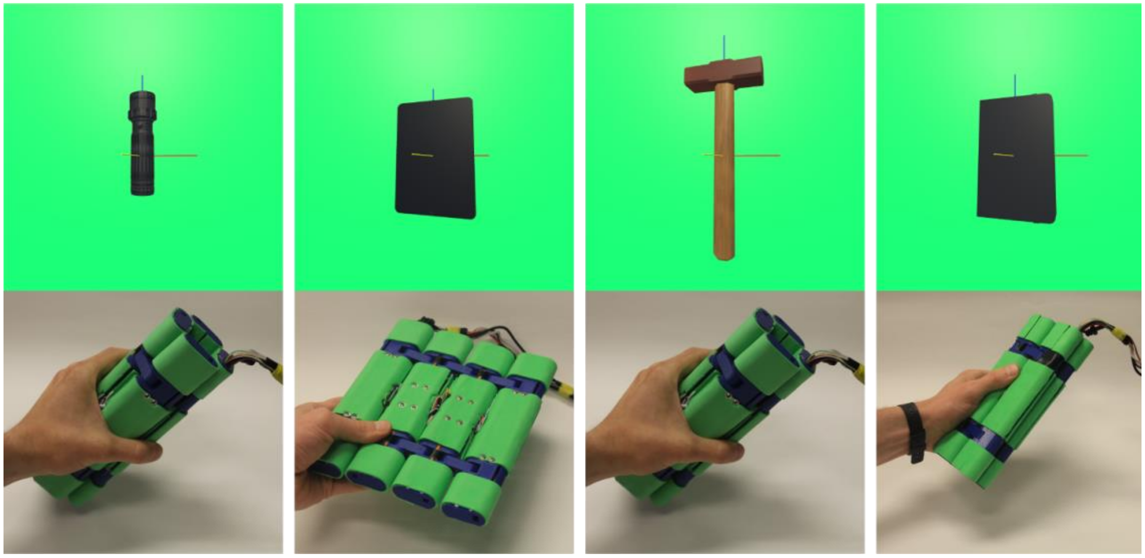

Figure 13. Virtual objects used in the user study with the Adaptic in its corresponding haptic shapes for each. From left to right they show the flashlight, the tablet, the hammer and the notebook.

5.2.4 Design

The docking user study draws from a 3 x 4 x 3 x 3 within-subjects design by incorporating the *independent variables*:

- **Haptic approach**: static controller, deformable, haptic prop
- **Virtual object**: hammer, flashlight, tablet, notebook
- **Location**[3] (cm): (-24.00, 17.53, 19.80), (0.00, 29.58, 19.80), (24.00, 17.53, 19.80)
- **Rotation**[4]: (300°, 330°, 345°), (330°, 345°, 300°), (345°, 300°, 330°)

[3] **Referred to as "left target location", "center target location", and "right target location" respectively.**
[4] **Referred to as "rotation 1", "rotation 2", and "rotation 3" respectively.**

As a result, each participant experienced an original combination of each independent variable for all the 108 dockings they performed. With 23 participants, completing 108 dockings our study captured 2,484 dockings in total. The interface we developed for testing selected all independent variables in a randomized order for all participants, this allowed us to counterbalance fatigue and training effects. See Figure 14 for an example of each target location and target rotation used.

For the docking task our main dependent variables were:

- **Time:** measured in seconds, it is the amount of time from the start of one docking to when the participant indicated they were at a sufficiently accurate alignment with the docking position.

- **Euclidean distance:** measured in centimeters (cm), it is the shortest 3D distance between the location a participant was satisfied with and the target location.

- **Rotational difference:** measured in degrees, it is the sum of x, y and z differences between the rotation a participant was satisfied with and the target rotation.

Through questionnaires we also captured several other variables. For each haptic approach-virtual object combination we collected match goodness ratings (how well a haptic approach matched the virtual object in the context of docking), and accuracy ease ratings (ease in which a participant could achieve their desired accuracy). Using the NASA Raw TLX questionnaire we assessed workload for each haptic approach through ratings for mental demand, physical demand, temporal demand, performance, effort, and frustration. Our heuristics questionnaire provided rankings on how each haptic approach met the four points of our Design Criteria (comfort, simplicity, shape diversity, and responsiveness). At the end of the study we captured data on participants' relative opinion of each haptic approach through ranking and eagerness measures.

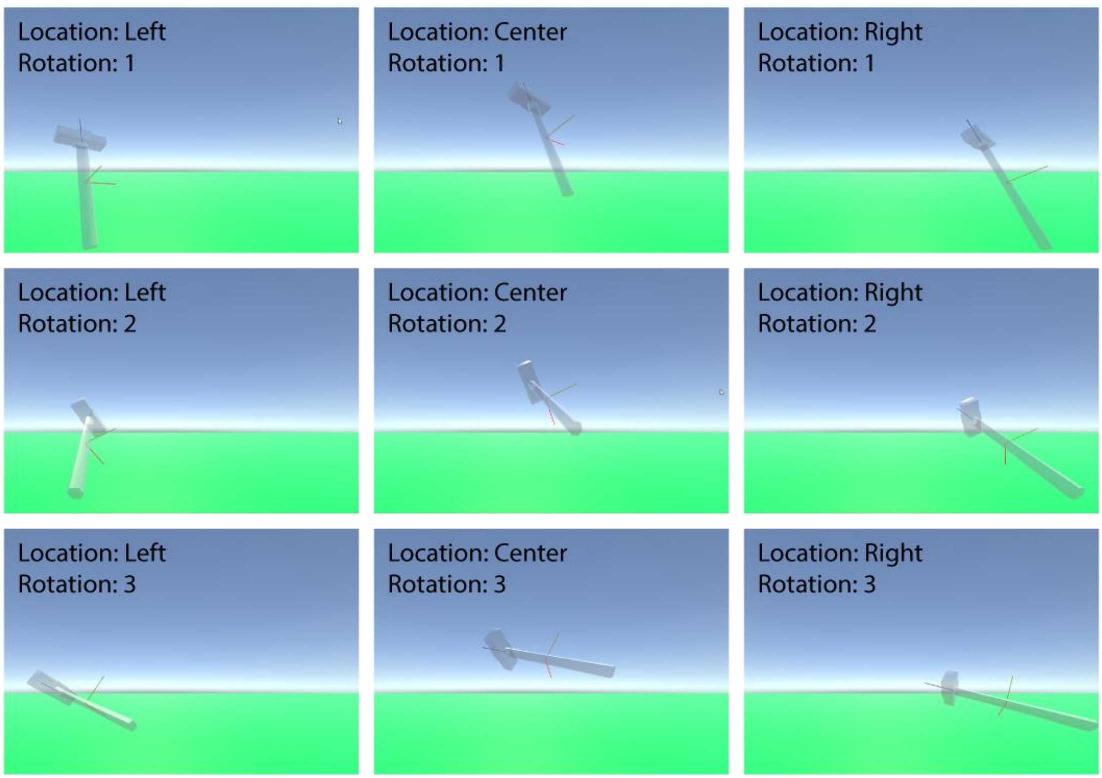

Figure 14. Each rotation-location combination for the 9 docking locations. They are shown with the same semitransparent copy of the virtual hammer used during the user study. We randomly generated the combinations for each participant who docked each virtual object at them, using each of the haptic approaches.

5.3 Results

We begin by reporting our performance-based data, time, Euclidean distance and rotational difference from the “Docking” section of the study. We examined each of these through a four-way repeated measures ANOVA. To assess if we needed to use corrections for a factor in our ANOVA we used Mauchly’s Test of Sphericity. If a factor was not significant in Mauchly’s Test, we assumed Sphereicity and made no correction, if we observed significance we made the Greenhouse-Geisser correction (if the Greenhouse-Geisser was 0.75 or below) or Huynh-Feldt correction (if the Greenhouse-Geisser was above 0.75). For factors of our ANOVAs showing significance, we ran post-hoc pairwise comparisons using Least Significant Differences.

Next, we cover questionnaire date covering match goodness, accuracy ease, the NASA RTLX questions, and the Heuristics questions. Our analysis of each of these started with using the Friedman Test to compare the data for each haptic approach. If our results showed significant differences, we followed with Wilcoxon Signed-Rank Test post-hoc using a Bonferroni correction. Last, we covered data from the “Ranking section” by examining the participants’ rankings further with a Chi-Squared Test.

5.3.1 Time

We collected data on the amount of time a participant took to complete each docking as one indicator of performance to compare each haptic approach. Overall the average time for Adaptic, the Hydra and the haptic props were 5.74 sec, 5.26 sec and 5.26 sec respectively. Figure 15 shows these results in greater detail. To examine time further we ran a four-way repeated measures ANOVA with haptic approach, virtual object, target location and target rotation as the independent variables, and time as the dependent variable. We found significant differences in the location factor and significant interactions between haptic approach and virtual object, haptic approach and location, and target location and target rotation (see Table 1). Due to the significant interactions involving location we do not explore its main effects further.

Source	df	error df	F	p-value	adjustment
haptic_approach	2.000	44.000	3.084	NS	Sphericity Assumed
object	3.000	66.000	2.354	NS	Sphericity Assumed
location	2.000	44.000	6.216	0.004	Sphericity Assumed
rotation	1.405	30.905	0.127	NS	Greenhouse-Geisser
haptic_approach * object	2.807	61.757	2.999	0.040	Greenhouse-Geisser
haptic_approach * location	4.000	88.000	2.541	0.045	Sphericity Assumed
object * location	3.540	77.874	1.659	NS	Greenhouse-Geisser
haptic_approach * rotation	4.000	88.000	2.368	NS	Sphericity Assumed
object * rotation	6.000	132.000	1.043	NS	Sphericity Assumed
location * rotation	2.885	63.478	4.635	0.006	Greenhouse-Geisser

Table 1. Four-way Repeated Measures ANOVA results for time taken to complete one docking. Significant p-values are highlighted in yellow.

We followed the ANOVA by investigating the significant interactions using post-hoc pairwise comparisons. For the significant interactions between the haptic approaches and virtual objects we uncovered several significant simple effects in the pairwise comparisons for haptic approaches within each virtual object (see Table 2 and Table 3). When using the tablet, we found Adaptic took significantly more time than both the Hydra and the haptic props. When participants used Adaptic, the book took significantly less time than the hammer, and the tablet. Looking at the haptic props approach, we see that the hammer took significantly more time than the book, tablet, and flashlight.

Object	haptic approach	mean (seconds)	std. error	Adaptic	Hydra	Props	F	df	error df
Hammer	Adaptic	5.776	0.442	.	NS	NS	1.342	2	21
	Hydra	5.324	0.363	.	.	NS			
	Props	5.710	0.43	.	.	.			
Book	Adaptic	5.337	0.387	.	NS	NS	0.725	2	21
	Hydra	5.227	0.421	.	.	NS			
	Props	5.041	0.39	.	.	.			
Tablet	Adaptic	6.209	0.632	.	**p = 0.047**	**p = 0.023**	2.927	2	21
	Hydra	5.169	0.343	.	.	NS			
	Props	5.009	0.358	.	.	.			
Flashlight	Adaptic	5.642	0.49	.	NS	NS	1.428	2	21
	Hydra	5.339	0.382	.	.	NS			
	Props	5.266	0.401	.	.	.			

Table 2. Pairwise comparisons post-hoc results for the interaction between virtual object and haptic approach comparing haptic approaches within each virtual object. Significant p-values are highlighted in yellow.

haptic approach	object	mean (seconds)	std. error	Hammer	Book	Tablet	Flashlight	F	df	error df
Adaptic	Hammer	5.776	0.442	.	**p = 0.007**	NS	NS	3.923	3	20
	Book	5.337	0.387	.	.	NS	NS			
	Tablet	6.209	0.632	.	.	.	NS			
	Flashlight	5.642	0.49	.	.	.	.			
Hydra	Hammer	5.324	0.363	.	NS	NS	NS	0.269	3	20
	Book	5.227	0.421	.	.	NS	NS			
	Tablet	5.169	0.343	.	.	.	NS			
	Flashlight	5.339	0.382	.	.	.	.			
Props	Hammer	5.710	0.43	.	**p = 0.006**	**p = 0.001**	**p = 0.017**	5.409	3	20
	Book	5.041	0.39	.	.	NS	NS			
	Tablet	5.009	0.358	.	.	.	NS			
	Flashlight	5.266	0.401	.	.	.	.			

Table 3. Time pairwise comparisons post-hoc results for the interaction between virtual object and haptic approach comparing virtual objects within each haptic approach. Significant p-values are highlighted in yellow.

Our pairwise comparisons post-hoc for the significant interaction between the haptic approaches and target locations also brought several significant simple effects to light (see Table 4 and Table 5). When docking at the center target location, Adaptic took significantly more time than the haptic props. The right target location mirrored this with Adaptic taking significantly more time than the haptic props. When using Adaptic, docking at the left target location took significantly less time than the center target location, and the right target location. For the Hydra controller, we found the left target location took significantly less time compared to the center target location.

location	haptic approach	mean (seconds)	std. error	Adaptic	Hydra	Props	F	df	error df
	Adaptic	5.442	0.433	.	NS	NS			
Left	Hydra	5.064	0.352	.	.	NS	1.492	2	21
	Props	5.236	0.379	.	.	.			
	Adaptic	5.906	0.475	.	NS	p = 0.024			
Center	Hydra	5.461	0.402	.	.		2.840	2	21
	Props	5.300	0.393	.	.	.			
	Adaptic	5.875	0.489	.	NS	p = 0.024			
Right	Hydra	5.269	0.346	.	.	NS	2.821	2	21
	Props	5.234	0.365	.	.	.			

Table 4. Time pairwise comparisons post-hoc results for the interaction between target location and haptic approach comparing haptic approaches with each target location. Significant p-values are highlighted in yellow.

haptic approach	location	mean (seconds)	std. error	Left	Center	Right	F	df	error df
	Left	5.442	0.433	.	p = 0.002	p = 0.027			
Adaptic	Center	5.906	0.475	NS	.	NS	3.923	2	21
	Right	5.875	0.489	.	NS	.			
	Left	5.064	0.352	.	p = 0.002	NS			
Hydra	Center	5.461	0.402	.	.	NS	0.269	2	21
	Right	5.269	0.346	.	.	.			
	Left	5.236	0.379	.	NS	NS			
Props	Center	5.300	0.393	.	.	NS	5.409	2	21
	Right	5.234	0.365	.	.	.			

Table 5. Time pairwise comparisons post-hoc results for the interaction between target location and haptic approach comparing target locations within each haptic approach. P-values for significant comparisons are highlighted in yellow.

5.3.2 Euclidean Distance

The second performance indicator we collected from participants was the Euclidean distance between a docking position's target location and the location of the haptic device when a participant completed that docking task. The average Euclidean distance for Adaptic, the Hydra controller and the haptic props were 1.14cm, 0.94cm and 1.01cm respectively. Figure 16 shows these and the averages for each virtual object. To gain a better understanding of Euclidean distance we ran a four-way repeated measures ANOVA with haptic approach, virtual object, target location and target rotation as the independent variables and Euclidean distance as the dependent variable. However, our results show no significant differences between any of the independent variables. Since no significant differences resulted from the ANOVA we did not pursue the pairwise comparisons post-hoc.

5.3.3 Rotational Difference

Our third performance indicator was rotational difference, the sum of x, y and z angle differences between the rotation a participant ended a docking task at and the target rotation. On average, using Adaptic, the Hydra controller and the haptic props resulted in rotational differences of 10.88°, 9.79° and 10.41° respectively. We show these averages and those for each of the individual virtual objects in Figure 17. To gather greater insight into rotational difference we ran a four-way repeated measures ANOVA with haptic approach, virtual object, target location and target rotation as the independent variables and rotational distance as the dependent variable. Our results show significance differences in virtual object, location, and the interaction between virtual object and location (see Table 6). Due to the significant interaction involving object and location we do not explore those main effects further.

Source	df	error df	F	p-value	adjustment
haptic_approach	1.643	36.153	0.539	0.554	Huynh-Feldt
object	1.074	23.620	10.553	0.003	Greenhouse-Geisser
location	2.000	44.000	12.162	< 0.001	Sphericity Assumed
rotation	1.201	26.426	2.162	0.150	Greenhouse-Geisser
haptic_approach * object	2.241	49.306	0.114	0.911	Greenhouse-Geisser
haptic_approach * location	2.265	49.820	1.114	0.342	Greenhouse-Geisser
object * location	6.000	132.000	2.695	0.017	Sphericity Assumed
haptic_approach * rotation	2.285	50.266	0.784	0.477	Greenhouse-Geisser
object * rotation	1.512	33.255	1.186	0.306	Greenhouse-Geisser
location * rotation	2.696	59.319	2.741	0.057	Greenhouse-Geisser

Table 6. Four-way Repeated Measures ANOVA results for rotational difference. Significant p-values are highlighted in yellow.

With the post-hoc pairwise comparisons, we looked further into the significant simple effects in the interactions between virtual objects and target locations (see Table 7 and Table 8). For dockings at the left target location we found the flashlight had a significantly higher rotational difference when compared to the hammer, the book, and the tablet. Using the middle target location resulted in a significantly higher rotational difference for the flashlight when compared to the hammer, book, and tablet. Dockings at the right docking location mirrored the other target locations with results that showed a significantly higher rotational difference for the flashlight when compared to the hammer, book, and tablet.

location	object	mean (degrees)	std. error	Hammer	Book	Tablet	Flashlight	F	df	error df
Left	Hammer	7.198	0.977	.	NS	NS	**p = 0.005**	3.923	3	20
	Book	6.800	0.711	.	.	NS	**p = 0.010**			
	Tablet	6.647	0.649	.	.	.	**p = 0.010**			
	Flashlight	13.757	2.91	.	.	.	.			
Center	Hammer	8.370	1.117	.	NS	NS	**p = 0.007**	0.269	3	20
	Book	8.469	1.017	.	.	NS	**p = 0.010**			
	Tablet	9.154	1.673	.	.	.	**p = 0.002**			
	Flashlight	16.778	3.554	.	.	.	.			
Right	Hammer	9.770	1.398	.	NS	NS	**p = 0.003**	5.409	3	20
	Book	8.933	1.19	.	.	NS	**p = 0.002**			
	Tablet	9.048	1.429	.	.	.	**p < 0.001**			
	Flashlight	19.421	3.776	.	.	.	.			

Table 7. Rotational difference pairwise comparisons post-hoc results for the interaction between target location and virtual object comparing target locations with each target rotation. Significant p-values are highlighted in yellow.

When participants used the hammer, the left target location resulted in significantly smaller rotational difference compared to the right target location. Using the book caused significantly less rotational error for the left target location compared to both the middle, and right target locations. Similar to the hammer, docking the tablet at the left target location caused significantly less rotational error compared to docking it at the right target location. Docking the flashlight in at the left target location produced significantly less rotational error compared to both the center target location, and the right target location. Finally, using the flashlight also produced significantly less rotational error for the center target location when compared to the right target location.

object	location	mean (degrees)	std. error	Left	Center	Right	F	df	error df
	Left	7.198	0.977	.	NS	0.010			
Hammer	Center	8.370	1.117	.	.	NS	3.988	2	21
	Right	9.770	1.398	.	.	.			
	Left	6.800	0.711	.	**p = 0.012**	**p = 0.004**			
Book	Center	8.469	1.017	.	.	NS	5.538	2	21
	Right	8.933	1.19	.	.	.			
	Left	6.647	0.649	.	NS	NS			
Tablet	Center	9.154	1.673	.	.	NS	2.401	2	21
	Right	9.048	1.429	.	.	.			
	Left	13.757	2.91	.	**p = 0.005**	**p < 0.001**			
Flashlight	Center	16.778	3.554	.	.	**p = 0.018**	8.303	2	21
	Right	19.421	3.776	.	.	.			

Table 8. Rotational difference pairwise comparisons post-hoc results for the interaction between target location and target location comparing target locations with each target rotation. Significant p-values are highlighted in yellow.

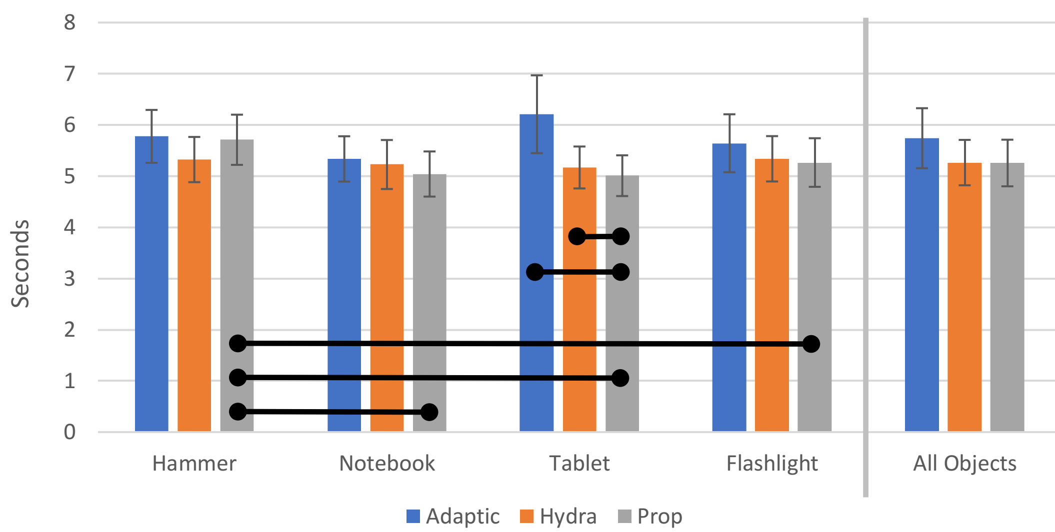

Figure 15. Average time taken to complete one docking with standard error bars and significant differences.

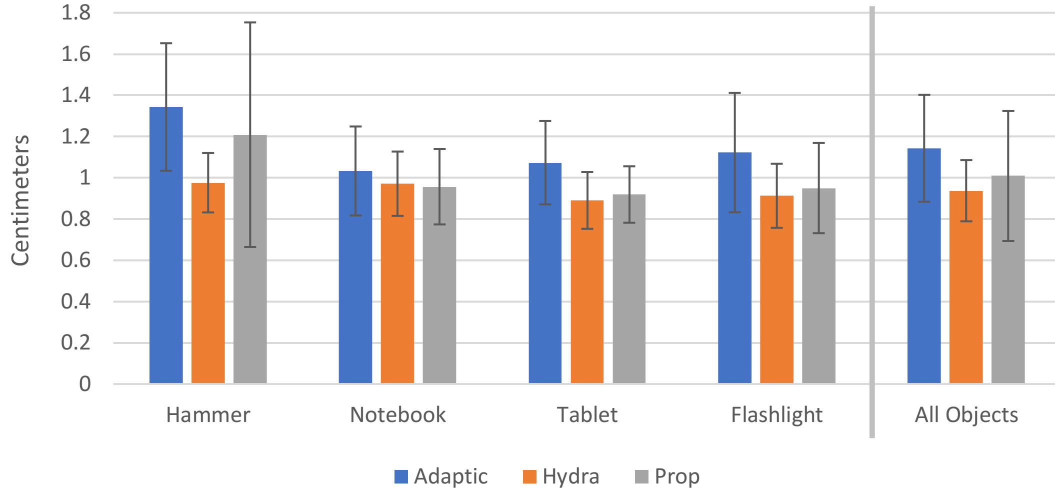

Figure 16. Average difference between the final location and the target location of a docking in Euclidean distance with standard error bars. We found no significant differences.

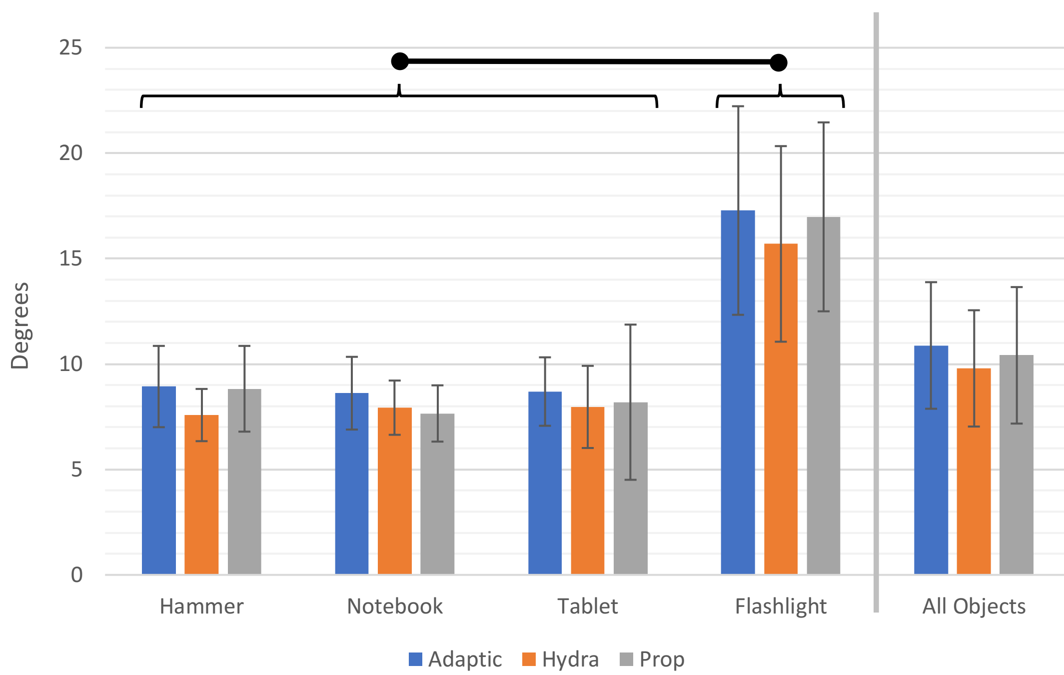

Figure 17. Average rotational difference for dockings with standard error bars and significant differences.

5.3.4 Match Goodness & Accuracy Ease

To compare how participants perceived each haptic approach for each of the virtual objects we collected feedback on match goodness and accuracy ease for each combination of the haptic approaches and virtual objects (seen in Figure 18 and Figure 19). To look closer at the ratings each haptic approach received with each virtual object we first used a Friedman Test (see Table 9 and Table 11), if its results were significant we then followed with a Wilcoxon-Signed Rank Test post-hoc using the Bonferroni adjustment[5] (see Table 10 and Table 12) to look for significant differences between the specific approaches.

5.3.4.1 Match Goodness

Following this process for the hammer we found the Freidman Test showed no significant differences between the haptic approaches. The Friedman Test for the book did show significant differences, and the following post-hoc revealed the haptic props received significantly higher ratings than both the Hydra, and Adaptic. The tablet was similar, with the Friedman test resulting in significant differences between the haptic approaches, and the haptic props showing higher ratings than the Hydra, and Adaptic in the following post-hoc. Significant differences also resulted from ratings for the flashlight after the Friedman Test, results from the post-hoc showed Adaptic received significantly lower ratings than both the Hydra controller, and the haptic props. See Table 9 and Table 10 for full results.

---

[5] **Adjusts the normal significance level of $p < 0.05$ by dividing it by the number of tests run. We run one test for each our 3 haptic approaches meaning our adjusted significance value is $p < 0.017$ (0.05/3 = 0.017)**

object	haptic approach	median	Chi-Square	df	p-value
Hammer	Adaptic	5.000	3.457	2.000	NS
	Hydra	6.000			
	Props	6.000			
Book	Adaptic	6.000	16.708	2.000	< 0.001
	Hydra	6.000			
	Props	7.000			
Tablet	Adaptic	5.000	14.026	2.000	0.001
	Hydra	6.000			
	Props	7.000			
Flashlight	Adaptic	6.000	18.968	2.000	< 0.001
	Hydra	6.000			
	Props	7.000			

Table 9. Friedman Test results for Match Goodness. Significant p-values are highlighted in yellow.

object	haptic approach (i)	median (i)	haptic approach (j)	median (j)	p-value	Z
Book	Adaptic	6.000	Hydra	6.000	NS	-1.642
	Hydra	6.000	Props	7.000	0.001	-3.468
	Props	7.000	Adaptic	6.000	0.002	-3.037
Tablet	Adaptic	5.000	Hydra	6.000	NS	-0.323
	Hydra	6.000	Props	7.000	0.002	-3.172
	Props	7.000	Adaptic	5.000	0.001	-3.365
Flashlight	Adaptic	6.000	Hydra	6.000	0.005	-2.793
	Hydra	6.000	Props	7.000	NS	-2.174
	Props	7.000	Adaptic	6.000	< 0.001	-3.557

Table 10. Wilcoxon-Signed Rank Test post-hoc for Match Goodness. Significant p-values after the Bonferroni adjustment are highlighted in yellow.

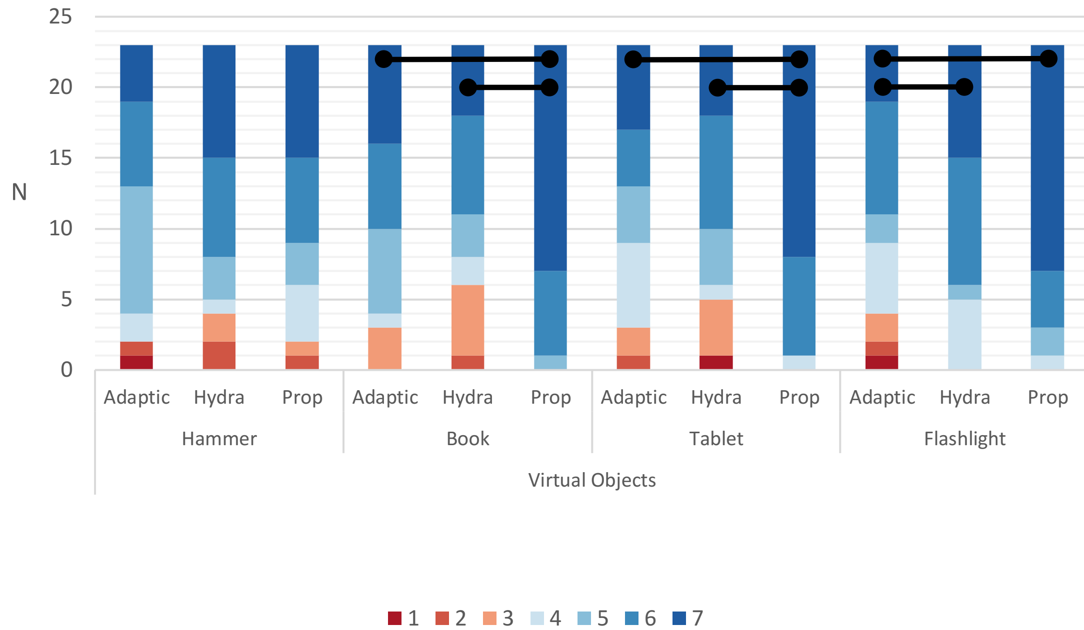

Figure 18. Shape Goodness ratings summed for each haptic approach within each virtual object with significant differences shown.

5.3.4.2 Accuracy Ease

After performing the Friedman test for accuracy ease ratings, both the hammer, and flashlight showed no significant difference in ratings when comparing the haptic approaches. The Friedman test for the book was significant, but the post-hoc did not reveal any significant differences after applying the Bonferroni adjustment. Ratings for the tablet also revealed significant differences, the following post-hoc showing that the haptic props received significantly higher ratings than Adaptic, and Prop. See Table 11 and Table 12 for full results.

object	haptic approach	median	Chi-Square	df	p-value
Hammer	Adaptic	5.000	2.493	2.000	NS
	Hydra	5.000			
	Props	5.000			
Book	Adaptic	6.000	6.028	2.000	0.049
	Hydra	6.000			
	Props	6.000			
Tablet	Adaptic	5.000	8.778	2.000	0.012
	Hydra	6.000			
	Props	7.000			
Flashlight	Adaptic	5.000	4.829	2.000	NS
	Hydra	6.000			
	Props	6.000			

Table 11. Friedman Test results for Accuracy Ease. Significant p-values are highlighted in yellow.

object	haptic approach (i)	median (i)	haptic approach (j)	median (j)	p-value	Z
Book	Adaptic	6.000	Hydra	6.000	NS	-0.275
	Hydra	6.000	Props	6.000	NS	-1.93
	Props	6.000	Adaptic	6.000	NS	-2.235
Tablet	Adaptic	5.000	Hydra	6.000	NS	-0.873
	Hydra	6.000	Props	7.000	NS	-2.285
	Props	7.000	Adaptic	5.000	0.002	-3.106

Table 12. Wilcoxon-Signed Rank Test post-hoc for Accuracy Ease. Significant p-values after the Bonferroni adjustment are highlighted in yellow.

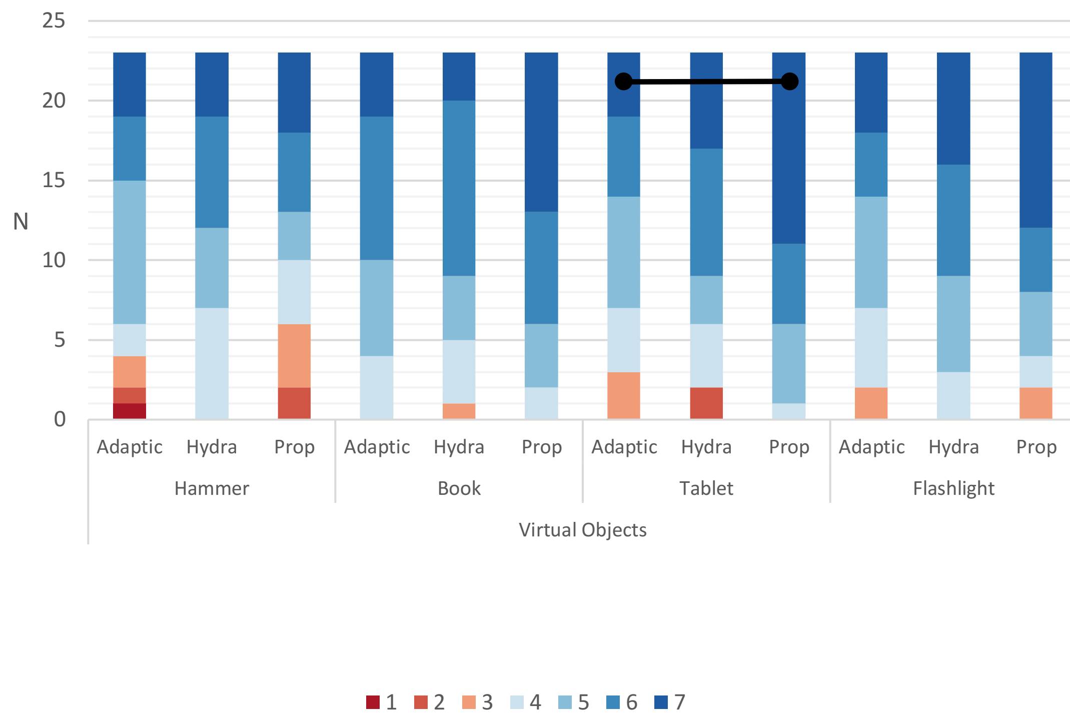

Figure 19. Accuracy ease ratings summed for each haptic approach within each virtual object with significant differences shown.

5.3.5 NASA RTLX

We compared the perceived workload of the haptic approaches using the NASA RTLX. As suggested by past research [26], we averaged the ratings for each of the 6 RTLX questions (Mental Demand, Physical Demand, Temporal Demand, Performance, Effort, and Frustration) individually, and summed those averages to arrive at a total measurement of workload for each haptic approach. The 6 questions rated on a 20-point scale meant 120 represented the highest workload level (6 questions * 20 point maximum rating). Adaptic scored 46.57, the Hydra controller scored 44.09, and the haptic props scored 44.57 (see Figure 20). To assess if the workload for each approach was significantly different we created an overall workload measurement for each participant by summing their responses to each question. Using the individual workload measurements, we ran a Friedman test that resulted in no significant differences between each approach ($\chi^2(2) = 1.295$, $p = 0.523$).

For more detail, we performed a Friedman Test comparing the ratings of each haptic device for each individual NASA RTLX question (see Figure 21 and Table 13). After finding significant differences between the haptic approaches for Physical Demand responses, the Wilcoxon Signed-Rank post-hoc (see Table 14) showed Physical Demand as significantly higher for Adaptic when compared to the Hydra controller. We also found no significant differences between the haptic approaches for Mental Demand, Temporal Demand, Performance, Effort, or Frustration.

measure	haptic approach	median	Chi-Square	df	p-value
Mental Demand	Adaptic	5.000	4.095	2.000	NS
	Hydra	4.000			
	Props	4.000			
Physical Demand	Adaptic	4.000	14.961	2.000	0.001
	Hydra	2.000			
	Props	3.000			
Temporal Demand	Adaptic	3.000	1.778	2.000	NS
	Hydra	3.000			
	Props	2.000			
Performance	Adaptic	12.000	5.564	2.000	NS
	Hydra	13.000			
	Props	11.000			
Effort	Adaptic	9.000	1.698	2.000	NS
	Hydra	8.000			
	Props	7.000			
Frustration	Adaptic	4.000	2.475	2.000	NS
	Hydra	3.000			
	Props	3.000			

Table 13. Friedman Test results for NASA RTLX questions. Significant p-values are highlighted in yellow.

measure	haptic approach (i)	median (i)	haptic approach (j)	median (j)	p-value	Z
Physical Demand	Adaptic	4.000	Hydra	2.000	0.003	-2.977
	Hydra	2.000	Props	3.000	NS	-2.015
	Props	3.000	Adaptic	4.000	NS	-0.829

Table 14. Wilcoxon-Signed Rank Test post-hoc for NASA RTLX questions. Significant p-values after the Bonferroni adjustment are highlighted in yellow.

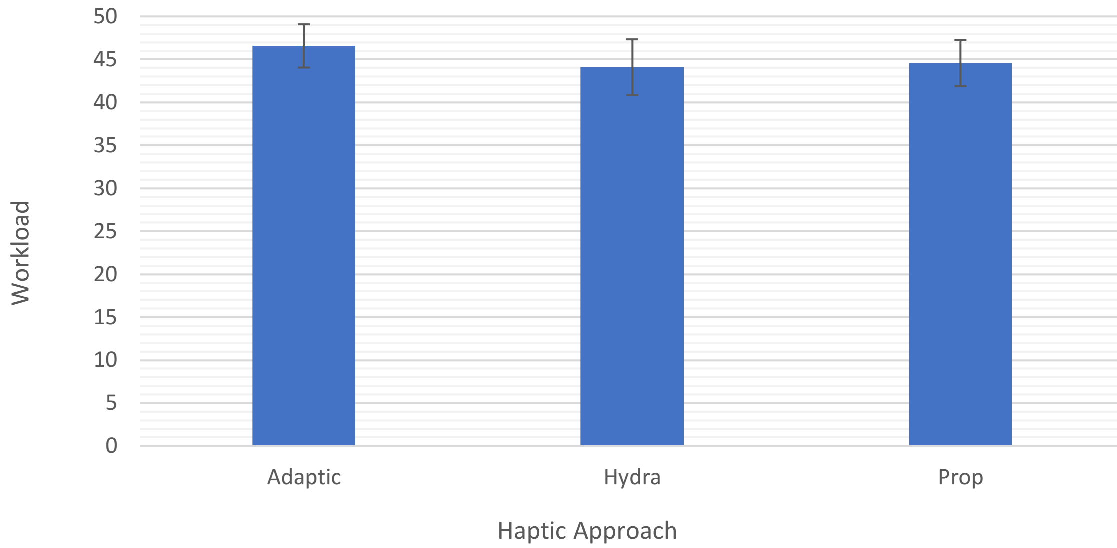

Figure 20. Total Workload for each haptic approach including standard error bars.

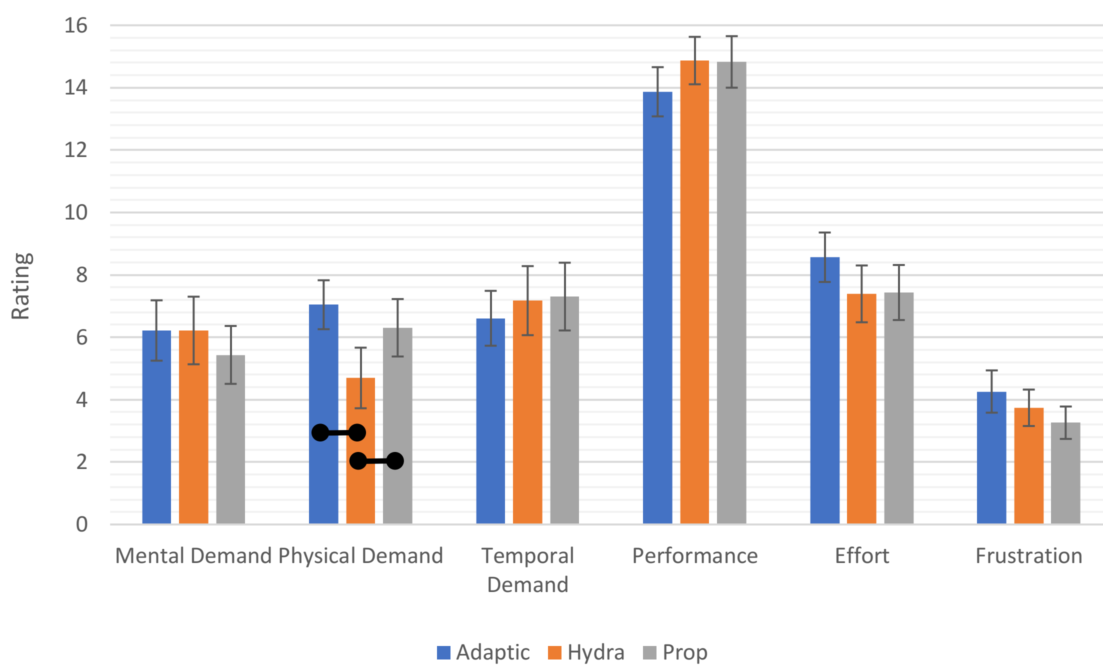

Figure 21. Averaged ratings for each haptic approach for each NASA RTLX question with standard error bars and significant differences.

5.3.6 Design Criteria Heuristics

Participants filled out a Heuristics Questionnaire to compare the haptic approaches along their adherence to our Design Criteria. The questionnaire consisted of four questions, one targeting each main point of the Design Criteria: Comfort, Simplicity, Shape Diversity and Responsiveness (see Figure 22). We performed a Friedman Test (see Table 15) on the responses to each question, followed by a Wilcoxon Signed-Rank post-hoc (see Table 16) if necessary. Using this process, we found significant differences between haptic approaches in Comfort ratings from the Friedman, but the post-hoc combined with the Bonferroni adjustment showed no significant differences between haptic approaches. We also found significance differences between the haptic approaches when analyzing the responses for Shape Diversity. The post-hoc revealed the haptic props performed significantly better than both Adaptic, and the Hydra controller. Results from both Simplicity, and Responsiveness showed no significant differences between haptic approaches.

measure	haptic approach	median	Chi-Square	df	p-value
Comfort	Adaptic	4.000	7.280	2.000	0.026
	Hydra	5.000			
	Props	5.000			
Simplicity	Adaptic	5.000	4.657	2.000	NS
	Hydra	6.000			
	Props	6.000			
Shape Diversity	Adaptic	5.000	14.000	2.000	NS
	Hydra	5.000			
	Props	7.000			
Responsiveness	Adaptic	5.000	2.676	2.000	NS
	Hydra	5.000			
	Props	6.000			

Table 15. Friedman Test results for Design Criteria Heuristics questions. Significant p-values are highlighted in yellow.

object	haptic approach (i)	mean (i)	haptic approach (j)	mean (j)	p-value	Z
Comfort	Adaptic	4.000	Hydra	5.000	NS	-2.35
	Hydra	5.000	Props	5.000	NS	-0.299
	Props	5.000	Adaptic	4.000	NS	-2.202
Shape Diversity	Adaptic	5.000	Hydra	5.000	NS	-0.832
	Hydra	5.000	Props	7.000	< 0.001	-3.569
	Props	7.000	Adaptic	5.000	0.006	-2.731

Table 16. Wilcoxon-Signed Rank Test post-hoc for Design Criteria Heuristics questions. Significant p-values after the Bonferroni adjustment are highlighted in yellow.

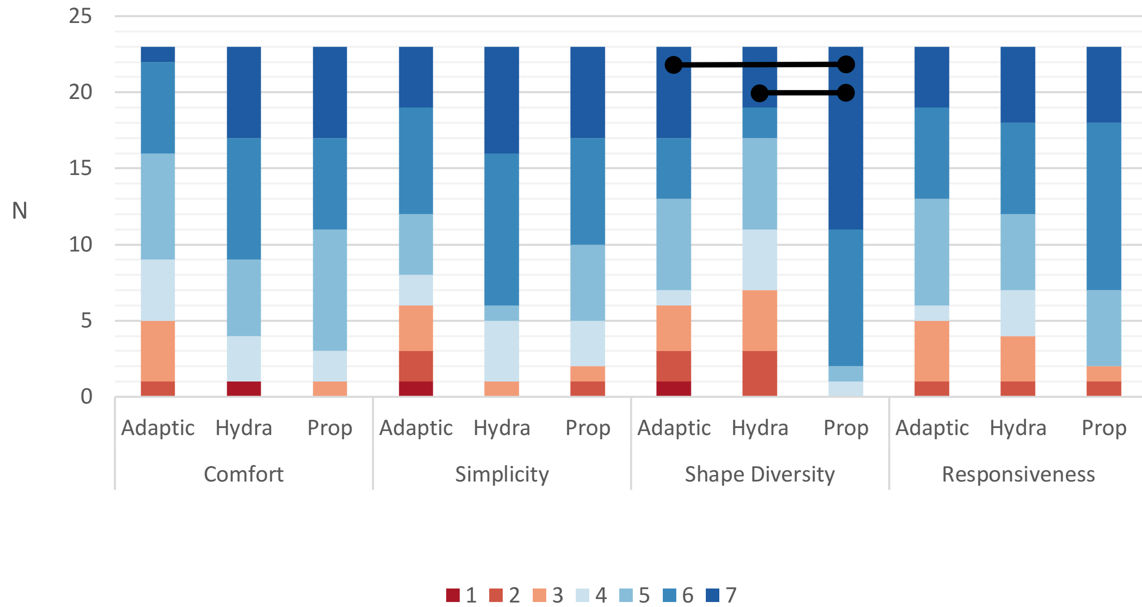

Figure 22. Summed ratings of each haptic approach for each Design Criteria Heuristics question with significant differences shown.

5.3.7 Ranking

At the end of the study we told participants they would do 15 additional dockings. We then asked them to rank the haptic approaches in the context of using them for these dockings (see Figure 23), informing them that they would use the approach ranked first. Three ranked Adaptic first, 12 ranked the Hydra controller first and 8 ranked the haptic props first. We analyzed significant differences between haptic approaches for the first, second and third rankings using a Chi-Squared test, but found no significant differences for any of them.

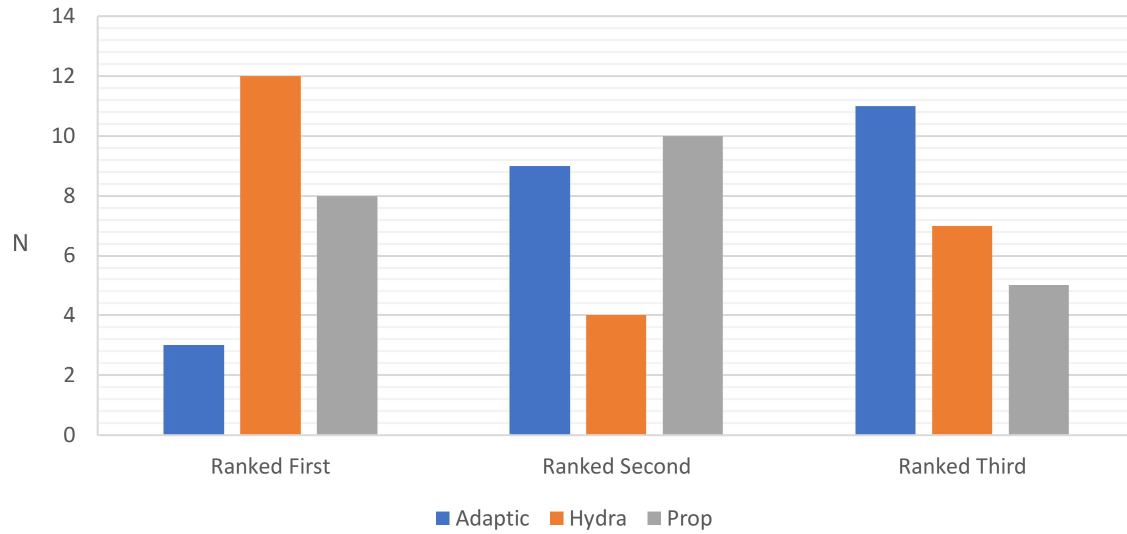

Figure 23. Summed participant rankings for each haptic approach.

5.4 Discussion

5.4.1 Question 1: Design Criteria

To satisfy the Design Criteria we outlined in Chapter 4 a haptic approach needs to perform well on the requirements of Comfort, Simplicity, Shape Diversity, and Responsiveness. We found that overall, Adaptic, and the other haptic approaches received positive ratings for all these requirements as measured through our Heuristics Questionnaire. Further examination of the requirements through other data we collected in the study also show positive results for all approaches, and bring attention to some areas of improvement for Adaptic.

5.4.1.1 Comfort

The Comfort requirement asks for a device to perform ergonomically in a variety of shapes and orientations. The heuristics question targeting comfort showed no significant differences in participant ratings. Expanding our view of comfort, we can draw on the

Physical Demand, Effort, and Frustration questions from the NASA RTLX as they measure related information. Of these, only Physical Demand shows significant differences, with the Hydra controller performing better than Adaptic. Increased physical demand for Adaptic conflicts with the work from Holman and Vertegaal [34], which we used to construct our Design Criteria. Specifically, it violates their "Form Equals Flow" requirement for OUIs. While this shows Adaptic has room for improvement, we do not believe that participants viewed it as uncomfortable overall. A closer look at participant ratings supports this with only 4 participants responding to the heuristics question with negative ratings (3 or lower) and only 7 participants responding to the Physical Demand question with negative ratings (11 or greater).

5.4.1.2 Simplicity

Overall, each haptic approach received positive ratings for Simplicity from the Heuristics Questionnaire and we found no significant differences between them. The Design Criteria's description of Simplicity includes maintaining a simple enough design so that VR interactions and performance are not limited. Looking at some of the performance indicators shows that docking the tablet with Adaptic as haptic feedback took more time to than the other two approaches and received worse accuracy ease ratings compared to the haptic props condition. While still receiving positive ratings, the relatively worse performance of Adaptic likely resulted from our inability to lock the shape into place due to the servos interfering with our magnetic tracking system. Because the shape matched with the tablet was completely flat, holding it without support from both ends resulted in it deforming. Contrary to this, the other virtual objects used more compact shapes that the natural stiffness of the servos maintained independently. By illustrating that the absence of shape-locking can lead to lower levels of performance, our results validate Holman and Vertegaal [34]'s requirement that "Function Equals Form" for OUIs and backed up feedback from our study with HaptoBend that requested a shape-locking function.

5.4.1.3 Shape Diversity

As one would expect, the haptic props' use of real-world objects as haptic feedback for corresponding virtual objects resulted in significantly better Shape Diversity ratings compared to Adaptic and the Hydra controller. However, given the Hydra controller's static shape we did find it surprising that it did not receive significantly lower Shape Diversity ratings when compared to Adaptic. Match goodness ratings for the book and tablet reflect the Heuristic Questionnaire results, with the haptic props performing significantly better than both Adaptic and the Hydra controller. To an extent, this conflicts with our findings from HaptoBend that appeared to show high performance for shape approximations due to visual dominance.

Match goodness also revealed another surprise, when docking the flashlight, the haptic props and the Hydra controller outperformed Adaptic. Looking at user comments, we see that participants described Adaptic's haptic shape for the flashlight as noticeably oversized, while the Hydra controller was coincidentally similar in size and shape to the flashlight making it a better match. The effect of size differences exhibits the limits of visual dominance and aligns with the findings of Simeone et al. [70] who demonstrated a larger, near identical physical prop used as PHF for a virtual object is significantly harder to use when compared a virtual object of equal size to the prop. The size mismatch also appears to conflict with the requirement of "Flow Follows Form" set out by Holman and Vertegaal [34] in their work outlining the preferred characteristics of OUIs.

5.4.1.4 Responsiveness

Our motivation to use of the same magnetic tracking system stemmed from maintaining an equal level of responsiveness to compare the haptic approaches across. Our heuristics question targeted at responsiveness confirms we achieved this with positive overall ratings and no significant differences seen between the approaches. Achieving the Responsiveness requirement outlined by our Design Criteria receives further reinforcement through the lack of differences in accuracy ease responses. Positive results for responsiveness also show Adaptic, and the other haptic approaches, satisfies Holman and Vertegaal [34]'s "Input Equals Output" requirement, especially important for achieving a sense of physical immersion.

5.4.2 Question 2: Performance

As a prototype filling the middle ground between the realism of identical haptic shapes and the simple convince of a static controller, we find Adaptic's performance very optimistic. Our results show that Adaptic is on an even playing field with both consumer level VR controllers (Hydra controller) and near identical objects used for PHF (haptic props). To compare Adaptic to the Hydra controller and the haptic props we collected data on each docking including time, Euclidean distance, and rotational difference. Each of the haptic approaches performed similarly. The haptic approaches showed no significant main effects for any of the performance variables and only appeared in significant simple effects when comparing docking times. By looking at the NASA RTLX performance question we also see that participants held positive perceptions of their performance with all the haptic approaches. Comparing performance through other questionnaires provides similar results with all the haptic approaches receiving positive ratings overall and few significant differences between them.

As expected with any prototype, there are areas of improvement for performance that we note as well. The interaction between haptic approaches and virtual objects for time reinforces the need for shape locking mentioned previously with Adaptic's less rigid tablet shape performing slower than the Hydra controller and haptic props. Slower performances when using Adaptic as haptic feedback for the tablet compared to the book adds support for shape locking as well.

We also find support for the haptic props' more realistic haptic feedback contributing to faster and more consistent performance when docking. While the haptic props showed no significant differences between target locations, the center and right target locations took more time to dock than the left location for both the Hydra controller and Adaptic. In addition, participants took a longer time to complete dockings with Adaptic compared to the haptic props in all but the left target location.

Our results also revealed the haptic props condition underperformed in other areas. Looking at the interaction between haptic approaches and virtual objects, we saw the hammer took significantly more time to dock than all the other virtual objects. Along with the support of participant comments we can confidently say this is a result of the weight of a real-world hammer, which was considerably heavier than any of the other haptic props. Simeone et al.'s [70] work on substitutional reality lead to similar results, showing poorer performance in a selection task and unfavorability among participants when using a heavier PHF prop. When comparing performance in VR with a flashlight and umbrella as PHF they found the heavier umbrella took more time to complete a selection task and users preferred the lighter flashlight more due to less physical exertion [70]. However, reduced performance from realistic weight for heavier virtual objects reveals a conflict with research suggesting that weight accuracy positively correlates with increases in perceived realism for VR [84]. Therefore, weight exhibits a notable trade-off between performance and realism for heavy objects, which we expect to increase for interactions over a longer time period by magnifying the "gorilla-arm" effect [6,30]. Leveraging visual dominance may offer solutions for this trade-off as past research [18,66] shows that visual feedback in a VE can distort a user's perception of weight.

5.4.3 Question 3: Preference

Participants did not appear to prefer Adaptic over either of the other haptic approaches they used. While participants gave the Hydra controller higher rankings and Adaptic lower rankings overall, our statistical analysis shows no significant differences between the three approaches. Our NASA RTLX scores echo an even preference between approaches, with statistically similar workload scores for each of them. Overall workload for our three haptic conditions also show consistency with similar research to ours by Besançon et al. [5] who reported total workload between 37 and 50 for each of their conditions. Therefore, while we cannot say there is a higher preference for deformable, devices, the similarity in ranking and other questionnaires suggests our participants saw each haptic approach with similar favorability. Combined with potential benefits of Adaptic not captured by our docking task, such as avoiding prop switching and added realism, a statistically equal preference to alternatives shows substantial support for continuing research in deformable, shape-changing haptics.

5.5 Limitations

We observed several limitations throughout the process of developing Adaptic and completing our docking task study. One area of limitations involved the design of our prototype. Adaptic's design became based on the size of servo motors powerful enough to actuate the device. This increased the size when compared to HaptoBend and according to several participants' comments the result was too large. Developing a way to use smaller rigid sections could alleviate this and possibly allow more shape options.

The magnetic tracking system we used also lead to limitations. As stated previously, we chose magnetic tracking to avoid occlusion issues, but this created a conflict with using the servo motors as it introduced too much magnetic noise. Without the use of the servos we were not able to test Adaptic as a shape-changing device. In addition, we couldn't use the servos for shape-locking which this study and the elicitation study both highlight as an important feature. We also encountered some inconsistencies in the tracking system when Adaptic's external wiring came too close to the Razor Hydra tracking base. Overall, we feel exploring alternatives to magnetic tracking that still avoid our occasion issues could benefit future iterations of Adaptic.

As mentioned in Chapter 4, our Design Criteria did not go through a process of evaluation due to an absence of past research in the area of applying deformable and/or shape-changing devices to VR haptics. Because we used the criteria to assess each haptic approach in the docking task study it is important to acknowledge its preliminary state may dilute some of our findings. We hope our work with HaptoBend and Adaptic will incentivize similar work in the future that refines our Design Criteria through evaluation.

Our docking task study assessed performance with tasks detached from real-world use by focusing on accuracy in short intervals. We feel this approach introduces a limitation by ineffectively measuring many of the benefits of deformation and shape-change, such as avoiding the need to switch props while maintaining realism through physical shape approximation. In addition, the docking task study only tested a small subset of objects that didn't cover a very large spectrum of shapes. We hope future research in this area explores performance in more realistic settings to better assess the benefits of our approach and compare them to others.

5.6 Summary

The stated goal of our second user study was to compare a deformable, shape-changing device with other approaches for VR haptics. To achieve this, we developed Adaptic, a deformable, shape-changing prototype based on our well-informed Design Criteria outlined in Chapter 4. Our study then used a docking task to compare Adaptic's performance with a completely static commercial VR controller, and near-identical haptic props for haptic feedback. Throughout the user study, we gathered data to answer three research questions that explore each haptic approach's adherence to our **Design Criteria**, their **Performance**, and the **Preference** of users. Participants gave each haptic approach positive ratings when assessing their adherence to our Design Criteria. However, improvements to Adaptic such as ensuring shape-locking is available and allowing smaller shapes could yield even better results. All the approaches performed at a high level allowing few differences between them in the speed and accuracy participants achieved when competing the docking tasks. Our indicators of performance also supported the need for shape locking, exhibited increased docking consistency with the more realistic haptic props, and showed a trade-off between realism and performance for weight. Participant rankings showed more of them preferred using the Hydra controller, followed by the haptic props second and Adaptic third. However, these differences appear marginal, as there were no significant differences between them, or the overall workload participants felt using them as measured by the NASA TLX. While Adaptic still has room for improvement, we draw a high level of optimism from it achieving an equal level of performance when compared with a commercial VR controller and near identical haptic props. These results show that the performance of deformable devices make them a viable option for VR haptics. Therefore, we recommend further research with prototypes like Adaptic to improve upon it and test its performance in a wider variety of VR interactions.

Chapter 6: Conclusion

6.1 Overview

Greater immersion [71], higher levels of presence [2,33,36], and improved performance in selection and manipulation tasks [5,74,77] highlight the value of haptic feedback in VR. Even with this level of importance, VR haptics still suffer from persistent issues that we define into three categories: *complexity*, *limiting interactions*, and *inadequate haptic feedback*. Targeting this gap, we developed an original approach to VR haptics by leveraging deformable, shape-changing devices, and the dominance of human vision over other senses to provide realistic haptic feedback with physical shape approximations [4,84]. Through researching this approach, we strived to answer one overarching question: Do deformable, shape-changing devices offer a positive alternative to conventional haptic feedback approaches for VR?

We started by creating our HaptoBend prototype to investigate preferences for using a deformable haptic device in VR. HaptoBend allows deformation into a variety of shapes with 4 foldable rigid sections connected in a row by hinges. Sensors track its shape and rotation in real time allowing users to interact with the device in VR. Using HaptoBend, we ran the first shape elicitation user study with 6 different virtual objects. Our results show participants enjoyed using HaptoBend and suggested a wide variety of applications to leverage its benefits. Participants also provided valuable design feedback on how to improve HaptoBend. Feedback from participants through agreement scores, goodness ratings, and ease ratings added support to past research showing visual dominance in VR allows physical shape approximations to serve as realistic haptic feedback, while also identifying limits in shape mismatch. Lastly, we identified influential factors for user preferences including virtual object's shape, size, weight and perceived function.

With a better understanding of user preferences for HaptoBend we created our Design Criteria to guide future development in deformable, shape-changing VR haptics. We informed it with data from our user study with HaptoBend, and Holman & Vertegaal [34]'s foundational work outlining requirements for OUI's, an area that encompasses deformable, shape-changing devices. Using our Design Criteria, we created Adaptic, a new prototype that improves on HaptoBend's design. Adaptic uses the same four-panel design but improves on it with double-hinged connections for more shape options. It also allows deformation and actuated shape-change using servo motors with 3D rotation tracking from an IMU, or just deformation with a magnetic system that allows 3D translation and rotation tracking.

To better understand Adaptic's performance, we compared it to alternatives in the first comparative docking task for assessing a deformable haptic device in VR. We compared Adaptic with a static Razor Hydra controller and near identical haptic props using docking speed, docking accuracy, workload, and compliance with our Design Criteria. Participant feedback showed all haptic conditions did a good job of meeting the Design Criteria, but Adaptic could improve through a smaller form factor and better incorporating shape-locking. Accuracy appeared even between each approach but we did observe small differences in speed. Overall participants rated each approach similarly, however, they rated the Hydra marginally higher and Adaptic marginally lower.

Using these cumulative findings to answer our research question, we can say that deformable devices offer a positive alternative to conventional haptic feedback approaches for VR. Users reflect positively on their experiences with deformable devices and viewed their physical shape approximations as realistic for a range of virtual objects. Deformable devices also appear to match the performance of commercial VR controllers and near identical props while showing the potential to mitigate the issues of unrealistic haptic shapes and prop switching they respectively embody. Unfortunately, we cannot apply our findings to shape changing devices because our research did not specifically address them. However, due to the similarities between deformable and shape-changing devices we would expect similar results. Therefore, the positive performance from our prototypes gives us a high level of optimism that deformable, shape-changing devices can fill current gaps in VR haptics, and hope our findings motivate further exploration in this area.

6.2 Future Work

We see many opportunities for future work in VR haptics with deformable, shape-changing devices. The two studies we completed looked at our prototypes in very confined contexts. To expand our knowledge on how they perform we would like to see deformable, shape-changing devices applied to areas that leverage their benefits in VR. As a start, solving the conflict we experienced between shape-change and 3D translation tracking would allow these future studies to move forward.

One application of interest for future research is using deformation as an input. Through bending a device like Adaptic into predefined shapes, users could quickly transition between different virtual objects, such as a shield triggered by a completely flat shape and a sword activated by rolling the device up into a wand-like shape. Applying research that supports capturing emotional states through deformation [75] is another area to explore that could expand interactions in VR. Past research promoting the use of VR to treat mental issues such as arachnophobia [11,22] and PTSD [17,67] shows that applying new methods for expressing emotions, like deformation, could be very valuable.

We also believe research into creating richer haptic feedback in VR with deformable, shape-changing devices could expand the value of our approach. We would like to see exploration into techniques such as selectively setting hinges to the deformable and shape-locking modes to create more realistic representations of an object's physical characteristics. For example, a user could open and close a virtual book by locking some hinges to mimic the hardcover and allowing the middle hinges to move freely to mimic the spine. Researching how shape-change could facilitate animated haptic experiences also shows potential given the connections between emotion and shape-change [43,47,59,75].

Finally, we see a lot of potential in combining redirected touching with deformable, shape-changing devices in VR. As mentioned in the Related Work, redirected touching warps a VE to allow the reuse of a single haptic device for several virtual objects [4,40]. Using Adaptic with redirected touching would increase this illusion by allowing the haptic device to change its shape and location.

Chapter 7: References

1. Merwan Achibet, Géry Casiez, and Maud Marchai. 2016. DesktopGlove: A multi-finger force feedback interface separating degrees of freedom between hands. In *3D User Interfaces (3DUI), 2016 IEEE Symposium on*, 3–12.

2. Laurent Aguerreche, Thierry Duval, and Anatole Lécuyer. 2010. Reconfigurable tangible devices for 3D virtual object manipulation by single or multiple users. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, 227–230.

3. Teemu T Ahmaniemi, Johan Kildal, and Merja Haveri. 2014. What is a device bend gesture really good for? In *Proceedings of the ACM Conference on Human Factors in Computing Systems, CHI '14* (CHI '14), 3503–3512. https://doi.org/10.1145/2556288.2557306

4. Mahdi Azmandian, Mark Hancock, Hrvoje Benko, Eyal Ofek, and Andrew D Wilson. 2016. Haptic retargeting: dynamic repurposing of passive haptics for enhanced virtual reality experiences. In *Proceedings of the ACM Conference on Human Factors in Computing Systems. CHI '16*, 1968–1979.

5. Lonni Besançon, Paul Issartel, Mehdi Ammi, and Tobias Isenberg. 2017. Mouse, tactile, and tangible input for 3D manipulation. In *Proceedings of the ACM Conference on Human Factors in Computing Systems, CHI '17*, 4727–4740.

6. Sebastian Boring, Marko Jurmu, and Andreas Butz. 2009. Scroll, tilt or move it: using mobile phones to continuously control pointers on large public displays. In *Proceedings of the 21st Annual Conference of the Australian Computer-Human Interaction Special Interest Group: Design: Open 24/7*, 161–168.

7. Mourad Bouzit, Grigore Burdea, George Popescu, and Rares Boian. 2002. The Rutgers Master II-new design force-feedback glove. *IEEE/ASME Transactions on mechatronics* 7, 2: 256–263.

8. Frederick P Brooks. 1999. What's real about virtual reality? *IEEE Computer graphics and applications* 19, 6: 16–27.

9. Sarah Buchanan, Bourke Floyd, Will Holderness, and Joseph J LaViola. 2013. Towards user-defined multi-touch gestures for 3D objects. In *Proceedings of the ACM International Conference on Interactive tabletops and surfaces*, 231–240.

10. Wolfgang Büschel, Patrick Reipschläger, and Raimund Dachselt. 2016. Foldable3D: Interacting with 3D Content Using Dual-Display Devices. In *Proceedings of the 2016 ACM on Interactive Surfaces and Spaces*, 367–372.

11. Albert S Carlin, Hunter G Hoffman, and Suzanne Weghorst. 1997. Virtual reality and tactile augmentation in the treatment of spider phobia: a case report. *Behaviour research and therapy* 35, 2: 153–158.

12. Tom Carter, Sue Ann Seah, Benjamin Long, Bruce Drinkwater, and Sriram Subramanian. 2013. UltraHaptics: Multi-point mid-air haptic feedback for touch surfaces. In *Proceedings of the ACM symposium on User interface software and technology, UIST '13*, 505–514.

13. Michael Chen, S Joy Mountford, and Abigail Sellen. 1988. A study in interactive 3-D rotation using 2-D control devices. *ACM SIGGRAPH Computer Graphics* 22, 4: 121–129.

14. Nicholas Chen, Francois Guimbretiere, Morgan Dixon, Cassandra Lewis, and Maneesh Agrawala. 2008. Navigation techniques for dual-display e-book readers. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1779–1788.

15. Inrak Choi, Heather Culbertson, Mark R Miller, Alex Olwal, and Sean Follmer. 2017. Grabity: A Wearable Haptic Interface for Simulating Weight and Grasping in Virtual Reality. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*, 119–130.

16. Inrak Choi, Elliot W Hawkes, David L Christensen, Christopher J Ploch, and Sean Follmer. 2016. Wolverine: A wearable haptic interface for grasping in virtual reality. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 986–993.

17. Joann Difede, Judith Cukor, Ivy Patt, Cezar Giosan, and Hunter Hoffman. 2006. The application of virtual reality to the treatment of PTSD following the WTC attack. *Annals of the New York Academy of Sciences* 1071, 1: 500–501.

18. Lionel Dominjon, Anatole Lécuyer, J-M Burkhardt, Paul Richard, and Simon Richir. 2005. Influence of control/display ratio on the perception of mass of manipulated objects in virtual environments. In *Proceedings of IEEE Virtual Reality*, 19–25.

19. Sean Follmer, Daniel Leithinger, Alex Olwal, Nadia Cheng, and Hiroshi Ishii. 2012. Jamming user interfaces: programmable particle stiffness and sensing for malleable and shape-changing devices. In *Proceedings of the 25th annual ACM symposium on User interface software and technology*, 519–528.

20. Eisuke Fujinawa, Shigeo Yoshida, Yuki Koyama, Takuji Narumi, Tomohiro Tanikawa, and Michitaka Hirose. 2017. Computational design of hand-held VR controllers using haptic shape illusion. In *Proceedings of the 23rd ACM Symposium on Virtual Reality Software and Technology*, 28.

21. David T Gallant, Andrew G Seniuk, and Roel Vertegaal. 2008. Towards more paper-like input: flexible input devices for foldable interaction styles. In *Proceedings of the 21st annual ACM symposium on User interface software and technology*, 283–286.

22. Azucena Garcia-Palacios, H Hoffman, Albert Carlin, T A u Furness, and Cristina Botella. 2002. Virtual reality in the treatment of spider phobia: a controlled study. *Behaviour research and therapy* 40, 9: 983–993.

23. Mathieu Le Goc, Lawrence H Kim, Ali Parsaei, Jean-Daniel Fekete, Pierre Dragicevic, and Sean Follmer. 2016. Zooids: Building blocks for swarm user interfaces. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, 97–109.

24. Antonio Gomes, Andrea Nesbitt, and Roel Vertegaal. 2013. MorePhone: A Study of Actuated Shape Deformations for Flexible Thin-Film Smartphone Notifications. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13*, 583–592. https://doi.org/10.1145/2470654.2470737

25. Antonio Gomes and Roel Vertegaal. 2015. PaperFold: evaluating shape changes for viewport transformations in foldable thin-film display devices. In *Proceedings of the International Conference on Tangible, Embedded, and Embodied Interaction, TEI '15*, 153–160.

26. Sandra G Hart. 2006. NASA-task load index (NASA-TLX); 20 years later. In *Proceedings of the human factors and ergonomics society annual meeting*, 904–908.

27. Elliot Hawkes, B An, Nadia M Benbernou, H Tanaka, S Kim, E D Demaine, D Rus, and Robert J Wood. 2010. Programmable matter by folding. *Proceedings of the National Academy of Sciences* 107, 28: 12441–12445.

28. Fabian Hemmert, Susann Hamann, Matthias Löwe, Josefine Zeipelt, and Gesche Joost. 2010. Shape-changing mobiles: tapering in two-dimensional deformational displays in mobile phones. In *CHI'10 Extended Abstracts on Human Factors in Computing Systems*, 3075–3080.

29. Atis Hermanis, Ricards Cacurs, and Modris Greitans. 2016. Acceleration and magnetic sensor network for shape sensing. *IEEE Sensors Journal* 16, 5: 1271–1280.

30. Juan David Hincapié-Ramos, Xiang Guo, Paymahn Moghadasian, and Pourang Irani. 2014. Consumed endurance: a metric to quantify arm fatigue of mid-air interactions. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, 1063–1072.

31. Ken Hinckley, Morgan Dixon, Raman Sarin, Francois Guimbretiere, and Ravin Balakrishnan. 2009. Codex: a dual screen tablet computer. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1933–1942.

32. Ken Hinckley, Randy Pausch, John C Goble, and Neal F Kassell. 1994. Passive real-world interface props for neurosurgical visualization. In *Proceedings of the ACM Conference on Human Factors in Computing Systems, CHI '94*, 452–458.

33. Hunter G Hoffman. 1998. Physically touching virtual objects using tactile augmentation enhances the realism of virtual environments. In *Proceedings of IEEE VRAIS*, 59–63.

34. David Holman and Roel Vertegaal. 2008. Organic user interfaces: designing computers in any way, shape, or form. *Communications of the ACM* 51, 6: 48–55.

35. Yingdan Huang and Michael Eisenberg. 2012. Easigami: virtual creation by physical folding. In *Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction*, 41–48.

36. Brent Edward Insko. 2001. Passive haptics significantly enhances virtual environments. University of North Carolina at Chapel Hill.

37. Bret Jackson, Tung Yuen Lau, David Schroeder, Kimani C Toussaint, and Daniel F Keefe. 2013. A lightweight tangible 3D interface for interactive visualization of thin fiber structures. *IEEE transactions on visualization and computer graphics* 19, 12: 2802–2809.

38. Richard D Joyce and Stephen Robinson. 2017. Passive haptics to enhance virtual reality simulations. In *AIAA Modeling and Simulation Technologies Conference*, 1313.

39. Yuichiro Katsumoto, Satoru Tokuhisa, and Masa Inakage. 2013. Ninja track: Design of electronic toy variable in shape and flexibility. In *Proceedings of the International Conference on Tangible, Embedded and Embodied Interaction, TEI '13* (TEI '13), 17–24. https://doi.org/10.1145/2460625.2460628

40. Luv Kohli, Mary C Whitton, and Frederick P Brooks. 2013. Redirected Touching: Training and adaptation in warped virtual spaces. In *3D User Interfaces (3DUI), 2013 IEEE Symposium on*, 79–86.

41. Eun Kwon, Gerard J Kim, and Sangyoon Lee. 2009. Effects of sizes and shapes of props in tangible augmented reality. In *Proceedings of IEEE ISMAR*, 201–202.

42. Byron Lahey, Audrey Girouard, Winslow Burleson, and Roel Vertegaal. 2011. PaperPhone: understanding the use of bend gestures in mobile devices with flexible electronic paper displays. In *Proceedings of the ACM Conference on Human Factors in Computing Systems, CHI '11*, 1303–1312.

43. Jung Min Lee, So Yon Jeong, and Da Young Ju. 2015. Emotional Interaction and Nofitication of Flexible Handheld Devices. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*, 2025–2030.

44. Minkyung Lee. 2010. Multimodal speech-gesture interaction with 3D objects in augmented reality environments. University of Canterbury. Department of Computer Science and Software Engineering.

45. Sang-Su Lee, Sohyun Kim, Bopil Jin, Eunji Choi, Boa Kim, Xu Jia, Daeeop Kim, and Kun-pyo Lee. 2010. How users manipulate deformable displays as input devices. In *Proceedings of the ACM Conference on Human Factors in Computing Systems, CHI '10*, 1647–1656.

46. Robert W Lindeman, John L Sibert, and James K Hahn. 1999. Hand-held windows: towards effective 2D interaction in immersive virtual environments. In *Proceedings of IEEE Virtual Reality*, 205–212.

47. David Lindlbauer, Jens Emil Grønbæk, Morten Birk, Kim Halskov, Marc Alexa, and Jörg Müller. 2016. Combining shape-changing interfaces and spatial augmented reality enables extended object appearance. In *Proceedings of the ACM Conference on Human Factors in Computing Systems, CHI '16*, 791–802.

48. Jessica Lo and Audrey Girouard. 2014. Fabricating bendy: Design and development of deformable prototypes. *IEEE Pervasive Computing* 13, 3: 40–46.

49. Pedro Lopes, Alexandra Ion, and Patrick Baudisch. 2015. Impacto: Simulating physical impact by combining tactile stimulation with electrical muscle stimulation. In *Proceedings*

of the 28th Annual ACM Symposium on User Interface Software & Technology, 11–19.

50. Pedro Lopes, Sijing You, Lung-Pan Cheng, Sebastian Marwecki, and Patrick Baudisch. 2017. Providing haptics to walls & heavy objects in virtual reality by means of electrical muscle stimulation. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 1471–1482.

51. Thomas H Massie and J Kenneth Salisbury. 1994. The phantom haptic interface: a device for probing virtual objects. In *Proceedings of the ASME*, 295–300.

52. John C McClelland, Robert J Teather, and Audrey Girouard. 2017. HaptoBend: shape-changing passive haptic feedback in virtual reality. In *Proceedings of the ACM Symposium Spatial User Interaction*, 9p.

53. Ken Nakagaki, Artem Dementyev, Sean Follmer, Joseph A Paradiso, and Hiroshi Ishii. 2016. ChainFORM: A Linear Integrated Modular Hardware System for Shape Changing Interfaces. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, 87–96.

54. Ken Nakagaki, Sean Follmer, and Hiroshi Ishii. 2015. LineFORM: actuated curve interfaces for display, interaction, and constraint. In *Proceedings of the ACM Symposium on User Interface Software & Technology, UIST '15*, 333–339.

55. Francisco R Ortega, Alain Galvan, Katherine Tarre, Armando Barreto, Naphtali Rishe, Jonathan Bernal, Ruben Balcazar, and Jason-Lee Thomas. 2017. Gesture elicitation for 3D travel via multi-touch and mid-air systems for procedurally generated pseudo-universe. In *IEEE Symposium on 3D User Interfaces (3DUI)*, 144–153.

56. Michaël Ortega, Jérôme Maisonnasse, and Laurence Nigay. 2017. EXHI-bit: a mechanical structure for prototyping EXpandable handheld interfaces. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services*, 4.

57. Ming Ouh-Young, David V Beard, and Frederick P Brooks. 1989. Force display performs better than visual display in a simple 6-D docking task. In *Robotics and Automation, 1989. Proceedings., 1989 IEEE International Conference on*, 1462–1466.

58. Young-Woo Park, Joohee Park, and Tek-Jin Nam. 2015. The trial of bendi in a coffeehouse: use of a shape-changing device for a tactile-visual phone conversation. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 2181–2190.

59. Esben W Pedersen, Sriram Subramanian, and Kasper Hornbæk. 2014. Is my phone alive?: a large-scale study of shape change in handheld devices using videos. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, 2579–2588.

60. Joel C Perry, Jacob Rosen, and Stephen Burns. 2007. Upper-limb powered exoskeleton design. *IEEE/ASME transactions on mechatronics* 12, 4: 408–417.

61. Thammathip Piumsomboon, Adrian Clark, Mark Billinghurst, and Andy Cockburn. 2013. User-defined gestures for augmented reality. In *Extended Abstracts on Human Factors in*

Computing Systems, CHI-EA '13, 955–960.

62. Raf Ramakers, Johannes Schöning, and Kris Luyten. 2014. Paddle: highly deformable mobile devices with physical controls. In *Proceedings of the ACM conference on Human factors in computing systems, CHI '14*, 2569–2578.

63. Majken Kirkegård Rasmussen, Timothy Merritt, Miguel Bruns Alonso, and Marianne Graves Petersen. 2016. Balancing user and system control in shape-changing interfaces: A designerly exploration. In *Proceedings of the TEI'16: Tenth International Conference on Tangible, Embedded, and Embodied Interaction*, 202–210.

64. Sharif Razzaque, Zachariah Kohn, and Mary C Whitton. 2001. Redirected walking. In *Proceedings of EUROGRAPHICS*, 105–106.

65. Christian Rendl, David Kim, Sean Fanello, Patrick Parzer, Christoph Rhemann, Jonathan Taylor, Martin Zirkl, Gregor Scheipl, Thomas Rothländer, Michael Haller, and others. 2014. FlexSense: a transparent self-sensing deformable surface. In *Proceedings of the ACM symposium on User interface software and technology, UIST '14*, 129–138.

66. Michael Rietzler, Florian Geiselhart, Jan Gugenheimer, and Enrico Rukzio. 2018. Breaking the Tracking: Enabling Weight Perception using Perceivable Tracking Offsets. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 128.

67. Barbara Olasov Rothbaum, Larry Hodges, Renato Alarcon, David Ready, Fran Shahar, Ken Graap, Jarrel Pair, Philip Hebert, Dave Gotz, and Brian Wills. 1999. Virtual reality exposure therapy for PTSD Vietnam veterans: A case study. *Journal of traumatic stress* 12, 2: 263–271.

68. Anne Roudaut, Abhijit Karnik, Markus Löchtefeld, and Sriram Subramanian. 2013. Morphees: toward high "shape resolution" in self-actuated flexible mobile devices. In *Philosophical Transactions of the Royal Society B: Biological SciencesProceedings of the ACM Conference on Human Factors in Computing Systems*, 593–602.

69. Makoto Sato. 2002. Development of string-based force display: SPIDAR. In *VSMM*.

70. Adalberto L Simeone, Eduardo Velloso, and Hans Gellersen. 2015. Substitutional reality: Using the physical environment to design virtual reality experiences. In *Proceedings of the ACM Conference on Human Factors in Computing Systems, CHI '15*, 3307–3316.

71. Mel Slater. 2009. Place illusion and plausibility can lead to realistic behaviour in immersive virtual environments. *Philosophical Transactions of the Royal Society B: Biological Sciences* 364, 1535: 3549–3557.

72. Jürgen Steimle, Andreas Jordt, and Pattie Maes. 2013. Flexpad: highly flexible bending interactions for projected handheld displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 237–246.

73. Frank Steinicke, Gerd Bruder, Jason Jerald, Harald Frenz, and Markus Lappe. 2008. Analyses of human sensitivity to redirected walking. In *Proceedings of the 2008 ACM symposium on Virtual reality software and technology*, 149–156.

74. Richard Stoakley, Matthew J Conway, and Randy Pausch. 1995. Virtual reality on a WIM: interactive worlds in miniature. In *Proceedings of the ACM Conference on Human Factors in Computing Systems, CHI '95*, 265–272.

75. Paul Strohmeier, Juan Pablo Carrascal, Bernard Cheng, Margaret Meban, and Roel Vertegaal. 2016. An Evaluation of Shape Changes for Conveying Emotions. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 3781–3792.

76. Sriram Subramanian, Sue Ann Seah, Hiroyuki Shinoda, Eve Hoggan, and Loic Corenthy. 2016. Mid-air haptics and displays: systems for un-instrumented mid-air interactions. In *Extended Abstracts on Human Factors in Computing Systems, CHI-EA '16*, 3446–3452.

77. Robert J Teather, Daniel Natapov, and Michael Jenkin. 2010. Evaluating haptic feedback in virtual environments using ISO 9241-9. *IEEE Virtual Reality*: 307–308.

78. M Turner, D Gomez, M Tremblay, and M Cutkosky. 1998. Preliminary tests of an arm-grounded haptic feedback device in telemanipulation. In *Proc. of the ASME Dynamic Systems and Control Division*, 145–149.

79. Radu-Daniel Vatavu and Jacob O. Wobbrock. 2015. Formalizing agreement analysis for elicitation studies. *Proceedings of the ACM Conference on Human Factors in Computing Systems, CHI '15*: 1325–1334. https://doi.org/10.1145/2702123.2702223

80. Julie Wagner, Stéphane Huot, and Wendy Mackay. 2012. BiTouch and BiPad: designing bimanual interaction for hand-held tablets. In *Proceedings of the ACM Conference on Human Factors in Computing Systems, CHI '12*, 2317–2326.

81. Jacob O Wobbrock, Htet Htet Aung, Brandon Rothrock, and Brad A Myers. 2005. Maximizing the guessability of symbolic input. In *Extended Abstracts on Human Factors in Computing Systems, CHI-EA '05*, 1869–1872.

82. Jacob O Wobbrock, Meredith Ringel Morris, and Andrew D Wilson. 2009. User-defined gestures for surface computing. In *Proceedings of the ACM Conference on Human Factors in Computing Systems, CHI '05* (CHI '09), 1083–1092. https://doi.org/10.1145/1518701.1518866

83. André Zenner. 2016. Investigating Weight Distribution in Virtual Reality Proxy Interaction. Saarland University & DFKI, Germany.

84. Andre Zenner and Antonio Kruger. 2017. Shifty: A Weight-Shifting Dynamic Passive Haptic Proxy. *IEEE Transactions on Visualization and Computer Graphics*: 1285–1294.

85. Yiwei Zhao, Lawrence H Kim, Ye Wang, Mathieu Le Goc, and Sean Follmer. 2017. Robotic Assembly of Haptic Proxy Objects for Tangible Interaction and Virtual Reality. In *Proceedings of the 2017 ACM International Conference on Interactive Surfaces and Spaces*, 82–91.

Chapter 8: Appendix

8.1 Consent Forms

8.1.1 Elicitation Study

CUREB clearance #106524

Carleton UNIVERSITY
Canada's Capital University

Consent Form: Sample

Title: HaptoBend: Bendable Shape-Changing Passive Haptic Feedback in Virtual Reality

Date of ethics clearance: To be determined by CUREB (as indicated on the clearance form)

Ethics Clearance for the Collection of Data Expires: To be determined by CUREB (as indicated on the clearance form)

I ____, choose to participate in a study on passive haptic virtual reality input devices. This study aims to explore the functionality of HaptoBend, a bendable input device prototype for virtual reality. **The researcher for this study is John McClelland at the School of Information Technology.** We are completing this experiment as part of our course work in Prof. Teather's course, ITEC 5200**.**

This study will explore the functionality of HaptoBend, a bendable input device prototype for virtual reality, and its ability to mimic a variety of objects.

The study involves using the HaptoBend with an Oculus Rift head mounted device. Participants will be asked to use HaptoBend to explore possible interactions with it in virtual reality. Your will also be asked to map the movement of Haptobend to manipulate several virtual objects and report on how well it performs. At the end of the study you will be asked to give overall feedback on your experience with HaptoBend. During the study the virtual environment will be recorded and you will be filmed if you specifically consent to it below.

All research data will be encrypted. Research data will only be accessible by my team and my supervisor. Data will be anonymous and will not be linked back to your name in any way.

Once the project is completed, all research data will be until conclusion of the course, after which time it will be deleted.

If you would like a copy of the finished research project, you are invited to contact the researcher to request an electronic copy which will be provided to you.

The ethics protocol for this project was reviewed by the Carleton University

Page **1** of **2**

This document has been printed on both sides of a single sheet of paper.
Please retain a copy of this document for your records.

Research Ethics Board-B, which provided clearance to carry out the research. Should you have any ethical concerns with the study, please contact Dr. Andy Adler (Chair, Carleton University Research Ethics Board-B (by phone: 613-520-2600 ext. 4085 or email: ethics@carleton.ca).

If you have any questions or concerns, please contact Professor Teather at 613-520-2600 ext. 4176, or at Rob.Teather@carleton.ca.

Researcher contact information:
John McClelland
School of Information Technology
Carleton University
Tel: 613-983-9234
Email: johnmcclelland@cmail.carleton.ca

Supervisor contact information:
Robert J. Teather
School of Information Technology
Carleton University
Tel: ext. 4176
Email: Rob.Teather@carleton.ca

Do you agree to be video-recorded: ___Yes ___No

_________________________ _______________
Signature of participant Date

________________________ _______________
Signature of researcher Date

This document has been printed on both sides of a single sheet of paper.
Please retain a copy of this document for your records.

8.1.2 Docking Task Study

CUREB clearance #108655

Carleton UNIVERSITY
Canada's Capital University

Appendix 1
Consent Form

Title: Adaptic: Adaptable Haptic Feedback in Virtual Reality Using Deformation and Shape-Change

Funding Sources: NSERC Discovery Grant 105849, NSERC CREATE CLUE 102107, Ontario Early Researcher Award 103495

Date of ethics clearance: To be determined by CUREB (as indicated on the clearance form)

Ethics Clearance for the Collection of Data Expires: To be determined by CUREB (as indicated on the clearance form)

I ____________________________________, choose to participate in a study on virtual reality haptic devices. This study aims to explore the functionality of a bendable input device prototype for virtual reality. **The researcher for this study is John McClelland at the School of Information Technology.** He will be working under the supervision of Dr. Audrey Girouard and Dr. Robert J. Teather.

This study will explore the functionality of a bendable input device prototype for virtual reality, and its ability to physically mimic a variety of virtual objects.

The study involves using several haptic devices while wearing a virtual reality head mounted display. Participants will be asked to align each haptic device in a specific location and orientation to compare the device's performance. Participants are also asked to fill out several questionnaires to gain a detailed understanding of their experience.

During the study the virtual environment will be recorded and audio and video of you will be recorded if you specifically consent to it below. Audio and video recording of yourself is not necessary to participate, but is a valuable tool for recording your feedback in greater detail. Choosing not to have audio or video of yourself recorded will not affect your compensation for participating in the study.

Interacting in virtual reality does hold a small risk of cybersickness, however you will be stationary in a seated position for the entire study making the likelihood of experiencing cybersickness very small. If you do experience any of the symptoms of cybersickness, which include headache, nausea, and dizziness, please inform the researcher immediately. If cybersickness is experienced we will pause the study and remove the head mounted display. After the symptoms subside we will only continue the study if you feel comfortable. Discontinuing the study for the reasons of cybersickness or any other reason will still result in the same level of compensation.

Page **1** of **2**

This document has been printed on both sides of a single sheet of paper.
Please retain a copy of this document for your records.

All research data will be anonymized. Research data will only be accessible by our team and my supervisors. Data will be anonymous and will not be linked back to your name in any way.

You have the right to end your participation at any time during the session, for any reason. However, because data is anonymized you must withdraw during your participation in the study so that we can identify the data connected collected from your participation. If you withdraw from the session, all information you have provided and data we collected during your participation will be destroyed. Withdrawing from the study or choosing not to participate in the study will not affect your compensation.

If you would like a copy of the finished research project, you are invited to contact the researcher to request an electronic copy which will be provided to you.

The ethics protocol for this project was reviewed by the Carleton University Research Ethics Board-B, which provided clearance to carry out the research. If you have any ethical concerns with the study, please contact Dr. Andy Adler, Chair, Carleton University Research Ethics Board-B (by phone at 613-520-2600 ext. 4085 or via email at ethics@carleton.ca).

If you have any questions or concerns, please contact Dr. Audrey Girouard at 613-520-2600 ext. 8816, or at audreygirouard@cunet.carleton.ca.

Researcher contact information:
John McClelland
School of Information Technology
Carleton University
Tel: 613-983-9234
Email: johnmcclelland@cmail.carleton.ca

Supervisor contact information:
Audrey Girouard
School of Information Technology
Carleton University
Tel: 613-520-2600 ext. 8817
Email: audreygirouard@cunet.carleton.ca

Do you agree to be video-recorded: ___Yes ___No

Do you agree to be audio-recorded: ___Yes ___No

________________________ ______________
Signature of participant Date

_______________________ ______________
Signature of researcher Date

This document has been printed on both sides of a single sheet of paper.
Please retain a copy of this document for your records.

8.2 Questionnaires

8.2.1 Elicitation Study

Think Aloud Recording Sheet

Shape: ____________________

Use: ____________________

____

____

____

Shape: ____________________

Use: ____________________

____

____

____

Shape: ____________________

Use: ____________________

____

____

____

Elicitation Recording Sheet

Virtual Object

2D Small ☐ 2D Same ☐ 2D Large ☐

3D Small ☐ 3D Same ☐ 3D Large ☐

The shape I picked is a good match for its intended purpose.

Strongly Disagree ☐ ☐ ☐ ☐ ☐ ☐ ☐ **Strongly Agree**

Comments:

____

____

the shape I chose was easy to perform.

Strongly Disagree ☐ ☐ ☐ ☐ ☐ ☐ ☐ **Strongly Agree**

Comments:

____

____

Virtual Object

2D Small ☐ 2D Same ☐ 2D Large ☐

☐ ☐ ☐

Final Thoughts Questionnaire

The shape I picked is a good match for its intended purpose.

Strongly Disagree ☐ ☐ ☐ ☐ ☐ ☐ ☐ **Strongly Agree**

Please explain the things you liked about HaptoBend:

____

____

____

____

____

Please explain the things you disliked about HaptoBend:

____

____

____

____

____

____

8.2.2 Docking Task Study

Object-Approach Combination Questionnaire

Virtual Object:____

Haptic Approach:

☐ **Adaptic**

☐ **Razer Hydra Controller**

☐ **Identical Prop**

The controller I used for this virtual object was a good match for this task.

Strongly Disagree ☐ ☐ ☐ ☐ ☐ ☐ ☐ **Strongly Agree**

Comments:

____

____

It was easy to achieve the level of accuracy I wanted.

Strongly Disagree ☐ ☐ ☐ ☐ ☐ ☐ ☐ **Strongly Agree**

Comments:

____

____

NASA Raw TLX

Haptic Approach:

☐ **Adaptic** ☐ **Razer Hydra Controller** ☐ **Identical Prop**

Mental Demand — How mentally demanding was the task?

Very Low — Very High

Physical Demand — How physically demanding was the task?

Very Low — Very High

Temporal Demand — How hurried or rushed was the pace of the task?

Very Low — Very High

Performance — How successful were you in accomplishing what you were asked to do?

Perfect — Failure

Effort — How hard did you have to work to accomplish your level of performance?

Very Low — Very High

Frustration — How insecure, discouraged, irritated, stressed, and annoyed wereyou?

Very Low — Very High

Heuristics Questionnaire

Please Rate these statements on the 7-point scale provided below them

I found the haptic approach comfortable to hold for all virtual objects.

Strongly Disagree — Strongly Agree

The design of this haptic approach felt inviting to use and would not limit my interactions in virtual reality.

Strongly Disagree — Strongly Agree

I felt this haptic approach provided good shape for each of the virtual objects in the study.

Strongly Disagree — Strongly Agree

The responsiveness of this haptic approach allowed me to feel fully immersed in the virtual environment.

Strongly Disagree — Strongly Agree

Preference Questionnaire

You will need to complete a series of docking tasks for 15 additional virtual objects. Please indicate which haptic approach you would like to use for these virtual objects by ranking it first. Then rank the remaining approaches to indicate your 2nd and 3rd choices.

___ Razer Hydra Controller   ___ Identical Prop   ___ Adaptic

Opt-In Questionnaire

If given the option, would you have completed a series of docking tasks for 15 additional virtual objects using your preferred haptic approach?

☐ **Yes**

☐ **No**