

X-Rings: A Hand-mounted 360° Shape Display for Grasping in Virtual Reality

Eric J. Gonzalez
Stanford University
Stanford, USA
ejgonz@stanford.edu

Eyal Ofek
Microsoft Research
Redmond, USA
eyalofek@microsoft.com

Mar Gonzalez Franco
Microsoft Research
Redmond, USA
margon@microsoft.com

Mike Sinclair
Microsoft Research
Redmond, USA
sinclair@microsoft.com



Figure 1: X-Rings (A) is a 360° shape display designed to support grasping. The device is hand-mounted (B.1), allowing the fingers to grasp and release rendered objects freely (B.2). Four expanding layers enable rendering of many virtual objects, such as a tankard (C) and goblet (D).

ABSTRACT

X-Rings is a novel hand-mounted 360° shape display for Virtual Reality that renders objects in 3D and responds to user-applied touch and grasping force. Designed as a modular stack of motor-driven expandable rings (5.7-7.7 cm diameter), X-Rings renders radially-symmetric surfaces graspable by the user's whole hand. The device is strapped to the palm, allowing the fingers to freely make and break contact with the device. Capacitance sensors and motor current sensing provide estimates of finger touch states and gripping force. We present the results of a user study evaluating participants' ability to associate device-rendered shapes with visually-rendered objects as well as a demo application that allows users to freely interact with a variety of objects in a virtual environment.

CCS CONCEPTS

• **Human-centered computing** → **Haptic devices**; *Virtual reality*.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

UIST '21, October 10–14, 2021, Virtual Event, USA

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-8635-7/21/10...\$15.00
<https://doi.org/10.1145/3472749.3474782>

KEYWORDS

Grasping, Handheld Haptics, Shape Display, Virtual Reality, 360 Shape Rendering

ACM Reference Format:

Eric J. Gonzalez, Eyal Ofek, Mar Gonzalez Franco, and Mike Sinclair. 2021. X-Rings: A Hand-mounted 360° Shape Display for Grasping in Virtual Reality. In *The 34th Annual ACM Symposium on User Interface Software and Technology (UIST '21)*, October 10–14, 2021, Virtual Event, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3472749.3474782>

1 INTRODUCTION

Many virtual environments enable users to physically touch and grab objects of different shapes and materials. However, whether the user picks up a sword, a large jug, or a ball, in most systems today the user is physically holding the same fixed-shape controller.

Traditionally, there have been two main approaches to solve the issue of mismatch between virtual and physical shapes in virtual reality (VR). One solution is to use an *encountered-type* haptics paradigm [16], where different physical props are placed or robotically positioned to align with virtual content. This technique is sometimes combined with haptic retargeting to optimize the number or type of physical props needed [4, 19].

The second approach aims to achieve more general haptic rendering by redesigning handheld controllers to change shape and apply forces directly to (or resist forces from) the hand. Hardware prototypes for simulating objects of different shapes and sizes are numerous [5, 10, 11, 43]. Instrumented fixed geometry controllers

have also been able to simulate physical properties of held objects using wideband vibration [23, 31], but their use is limited by the extent of the haptic illusions they produce [6, 20]. Devices with moving parts often allow more versatile haptic rendering. For example, CLAW and CapstanCrunch allow users to grab objects of different sizes, either with motors or brake mechanisms; they demonstrate successful implementations of palm referenced, finger-actuated grasping [12, 39].

Exoskeletons such as Wolverine [11] can provide feedback at the fingertips of the user's whole hand. Such systems typically use externally- or body-mounted actuators with linkages or wires providing resistance to the movement of the fingers and hand [15]. We call such approaches "outside-in", as the actuators and device mechanisms are externally-mounted outside the grasp volume, often on the back of the hand. As these systems occupy significant space surrounding the user's hand, they are at higher risk of colliding with the user's physical environment and can cause occlusions with external motion tracking systems. More importantly, these systems do not leverage the fact that in many VR applications the user is already holding a controller within their grasp.

In this work, we propose X-Rings as a change in this paradigm toward an "inside-out" approach. Borrowing concepts from Shape Displays [32], which are arrays of actuated pins, it integrates shape output into the controller handle itself. To the best of our knowledge, X-Rings is the first hand-scale shape display capable of rendering 360° surfaces which are graspable by the user's whole hand (Figure 1). Its unique cylindrical arrangement enables rendering the object's shape over most of the palm and fingers, in contrast to most prior work which focuses on fingertips alone. Limiting the haptic rendering to four fingers, where the thumb is in opposition, allows the system to be designed using with only four actuators, making the controller lighter, simpler, and more robust.

X-Rings mounts to the palm via an adjustable Velcro strap around the knuckles, freeing the user's fingers to grasp rendered objects naturally and release their grip without the controller falling. Each of the four device layers is an expandable ring, and can grow/shrink between diameters of 5.7 cm and 7.7 cm. A maximum rendering time of about 100 ms allows the device to swiftly change its shape as the user releases an object and reaches to grab another. Analyzing the user's hand motion enables the system to identify potential objects of interest early, and allows the device to pre-render shapes and surfaces

The system is equipped with two types of input sensing. Capacitive sensing detects whether the user is grasping the device. This allows the system to limit the timing of shape-change to when the user's hand is open. By measuring the motor current actuating each layer, X-Rings is able to generate estimates of the force applied by each finger. This signal can be used both as a safety measure, preventing the motors from sustaining too much current, and as a sensor for deformation behavior of virtual objects, such as breaking, compliance, or crushing.

In the remainder of this work, we first review the state of the art in handheld haptic rendering and shape displays. We then describe the design and implementation of X-Rings, its rendering capabilities, and its implications for user interaction in virtual environments. Finally, we report the results of an evaluation measuring users' ability to both differentiate between shapes rendered by X-Rings

and correctly identify them from a set of graspable virtual objects. In order to better understand the impact of potential rendering limitations on the effectiveness of X-Rings and similar devices, we also study the effect of reducing the device's dynamic range on the above measures.

1.1 Contributions

The main contributions of our work are as follows:

- (1) A novel concept for a 360° shape-changing controller that renders radially-symmetric surfaces and enables whole-hand encountered-type grasping in VR.
- (2) A unique electromechanical design for a low-cost expanding circle mechanism, useful for radial shape output and touch/pressure sensing.
- (3) An evaluation of users' abilities to associate graspable surfaces rendered by X-Rings with virtual objects, including as a function of the device's dynamic shape rendering range.

2 RELATED WORK

In this section, we review the state-of-the-art in wearable and hand-held haptics for VR, as well as shape display technology.

2.1 Handheld and Wearable Haptics

In recent years, a growing number of works have explored using a handheld controller format for haptic rendering systems in VR.

For example, CLAW [12] controls the force on the index finger, while the thumb is free to change the semantics of the operation and manipulate the controls. Other controllers added rendering of normal force and texture [5, 43], shear [43] and simulation of some grounded forces such as gravity [10], inertia [28, 37, 47], drag [48], or propulsion [22]. These devices enable a dynamic haptic representation of the held object primarily at one or few fingertips, and generally render lower-level local cues rather than global shape over the full hand.

Inflatable bladders [21, 33, 45] are another interesting method of generating shapes in the user's hand. By changing the air pressure in a bladder, it can render different resistive forces and change its volume [36]. PuPop [42] used multiple bladders attached to the palm to enable switching between a few fixed shapes. Haptx covers the palm with many individually inflatable bubbles that can simulate soft stimuli on the skin, such as rain or a crawling insect, but not complete geometry [2]. In contrast with inflatables, some haptic systems use pressure change to induce particle [17, 38] or layer jamming [34], creating malleable and variable stiffness surfaces. Although inflatable and jamming systems are safe and low-cost, their need for an additional pneumatic or hydraulic pump, their slow rendering speeds, their reliance on a whole-hand glove, and their limited degrees of freedom are significant limitations to providing real-time dynamic haptic sensations.

Haptic exoskeletons provide kinesthetic feedback of rigid grasping using mechanical structures worn on the fingers [1, 8, 9, 11, 25, 29?]. Most exoskeletons focus on rendering at the fingertip, with some loading the fingers but not the palm. Their primary drawback is their cumbersome form factor, which makes them difficult to don and doff and increases the potential for collisions with the user's environment. By occluding the space surrounding the user's

hands, exoskeletons can also impede hand tracking systems which benefit haptic VR applications. Such systems often also depend on the user’s hand dimensions, which limits their use for consumers.

2.2 Shape Displays

Shape displays are devices designed to express a wide range of dynamic shapes and surfaces, typically through a grounded planar array of linearly-actuated pins which move up and down to generate a 2.5D surface [32, 40]. While impressive, most shape displays are complex, heavy, and limited in the area they can cover. Recently, researchers have miniaturized similar displays for handheld shape rendering in the palm [46], index fingertip [5], and along the edge of a mobile phone [26]. While these handheld displays are an important step toward versatile haptic rendering, the small planar shape displays they use still require many actuators. Additionally, although rather large compared to handheld controllers, they render haptic sensation to only a small part of the palm or a fingertip. Furthermore, these devices maintain constant contact with the user’s hand, whereas in the real world we regularly make and break contact with the various objects we grasp and surfaces we touch – such feedback is often referred to as *encountered-type* [13].

Toward addressing these issues, with X-Rings we design and implement a shape display that extrudes surfaces in 360° about a central axis, forming a cylindrical rendering region which can be grasped by the whole hand, as opposed to traditional planar displays. X-Rings also supports *encountered-type* grasping in VR by mounting to the user’s hand, enabling their fingers to make and break contact with rendered objects. While rendered surfaces are radially symmetric, the reduced degrees of freedom dramatically reduces the number of actuators required.

3 DESIGN & IMPLEMENTATION OF X-RINGS

In this section, we detail the development of X-Rings, including our design goals and rationale, selection of shape-change topology, hardware, assembly, and device functionality. Figure 2 illustrates the final overall design, and Figure 3 presents a block diagram of the system architecture.

3.1 Design Goals

The ideal controller should be able to rapidly render diverse shapes, be comfortably and safely manipulated, and support user input through physical interaction. Based on these considerations, we derived the following design goals and used them to guide the development of our controller. Ideally, the device should:

- (1) Render meaningfully diverse and expressive graspable shapes.
- (2) Withstand human-scale grasping forces, while maintaining surface backdrivability to maximize interactivity and ensure safety.
- (3) Sense user interaction with the rendered surfaces, in the form of touch and/or gripping force.
- (4) Keep the fingers free from rendering hardware, enabling the user to freely encounter rendered surfaces.
- (5) Maintain a compact, "graspable" form factor, meaning all actuation and shape-change mechanisms should fit within a small enough radius that can be enclosed by the hand.




Figure 2: (A) Overall diagram of the X-Rings controller. (B) Minimum and maximum diameter of individual expanding layers. (C) Depiction of the controller demonstrating full range of motion between layers: first and third layers fully expanded, second and fourth layers fully retracted.

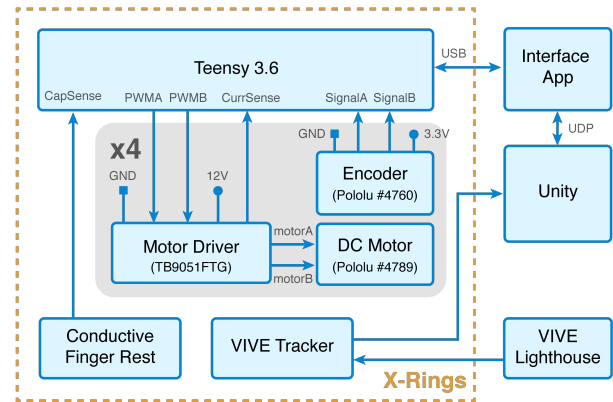


Figure 3: Block diagram of system detailing mechatronic components and relevant connections. For clarity, components are only shown for one of the four identical expanding layers (highlighted in gray).

- (6) Limit mechanical complexity and number of actuators, reducing required maintenance, cost, and weight.

3.2 Shape-Change Mechanism & Topology

The most critical design decision in the development of X-Rings was the selection of a shape-change mechanism and topology, or the arrangement of the shape-changing elements. The majority of existing shape displays use a planar array of individually actuated pins, resulting in 2.5D rendered surfaces. While highly expressive, these surfaces cannot render fully graspable objects as only half of a 3D object can be rendered at a time. This led us to explore mechanisms for radial pin extrusion, with the goal of rendering surfaces about a full 360° which can be grasped and enclosed by the hand.




Figure 4: Exploded view of a single expanding layer. As the motor rotates, a pinion gear drives a bevel gear, which rotates a cam. Spiral channels in the cam drive 5 arms in/out of the layer base.

The challenge with radial extrusion however is the limited available space to include the numerous actuation mechanisms often required by shape displays. We noted, however, that many commonly grasped objects such as tool handles, cups, and bottles have rough radial symmetry. This inspired us to pursue a mechanism for circular expansion, where one actuator can control the diameter of an entire layer, and layers can be stacked to create a 3D form. Circular expansion mechanisms such as Hoberman rings [24] have previously been used in shape-changing ambient displays [30] and data physicalizations [14], but to the best of our knowledge this topology has not been explored for shape rendering in VR.

Figure 4 shows an exploded view of a single circular expansion layer designed for X-Rings. It consists of 5 arms that are driven radially in/out of guide channels in the layer base by a custom spiral cam and bevel gear. The bevel gear is rigidly coupled to the cam, allowing it to be driven by a single gearmotor fit with a pinion gear. For low friction movement, a 2 mm ball-bearing is mounted to each arm and slides within a corresponding spiral channel in the cam. As the cam rotates, this bearing pushes the entire arm along its guide channel in the layer base. When the cam is turned, the arms extrude radially in unison changing the overall diameter of the layer. Each

layer has a minimum diameter of 5.7 cm to and a maximum of 7.7 cm when fully extruded. These dimensions were set by the size of the gearmotor used (see Section 3.4) – a smaller motor could reduce the minimum diameter, though at the cost of a reduced maximum diameter. The end of each arm is fit with a curved contact wedge (2.5 cm height) such that the layer forms a continuous ring when fully retracted and a segmented one when expanded. The concavity of each wedge helps ensure gripping forces are centrally applied to each arm. Additionally, the edges of each wedge are chamfered to prevent any skin pinching during shape-change.

3.3 Controller Assembly

A benefit of using these circular expansion layers is that they stack easily in the z-dimension, naturally taking on the form factor of a grip or handle. Additionally, each layer requires only a single motor which can be housed completely within the render volume of the mechanism itself, meaning no additional components need to be mounted externally on the hand or body.

We developed X-Rings leveraging these benefits, stacking four 2.5 cm thick layers (see Figure 2), fitting an average hand. Each layer provides feedback for a single finger, with the thumb and palm in opposition. We also include a simple 3D-printed proxy of a basic VR controller UI (trackpad + 2 buttons) mounted above the layers to demonstrate one way that traditional controller elements could be integrated with X-Rings. The base of each layer is fixed to a controller handle, creating a single graspable, radially shape-changing surface. A Velcro strap secures the ergonomically curved handle along the user’s knuckles, allowing them to freely grasp and release rendered shapes.

3.4 Hardware

X-Rings is assembled entirely from off-the-shelf parts and 3D printed components. With the exception of fasteners and bearings, all passive mechanical components on the device are 3D printed, using PLA filament on a hobbyist 3D printer (Creality CR-10S Pro V2).

Each layer of X-Rings is powered by a 12V DC gearmotor (Pololu #4789, 15:1 gear ratio). Motor rotation is measured using a magnetic encoder (Pololu #4760) mounted to the rear motor shaft, and controlled using a TB9051FTG motor driver through a software PID loop. Motor current is also monitored by the driver, and output as an analog voltage proportional to the motor current (500 mV/A). A set of 3D printed bevel gears (40:12 ratio) transmit motor power at 90 degrees to the extending arms via a spiral cam coupled to the bevel gear. A 3 mm ball-bearing aligns the rotation of the larger bevel gear. Each arm consists of an exterior wedge, a sliding pin, and an 2 mm ball-bearing which contacts the spiral cam. The four expanding layers are then mounted to the 3D printed controller handle via M2 screws. A Vive tracker [3] is mounted on the bottom of the controller handle to enable 6DOF tracking.

A Teensy 3.6 microcontroller governs all sensing and actuation on X-Rings, and receives commands from a PC via USB serial. Position control for each layer is maintained using a 1000 Hz PID loop, while four analog inputs are used to monitor the current of each motor. The FastTouch library [18] is used to sense user touch of an electrically-conductive finger rest (printed using conductive PLA,

Variable	Value
General	
Dimensions	95 x 95 x 185 mm
Weight	355 g (including Vive Tracker)
Power Draw	30 mA (idle), 3 A (max force) @ 12 V
Rendering	
Min Diameter	5.7 cm
Max Diameter	7.7 cm
Max Speed	20 cm/s (diametral)
Max Render Time	100 ms
Position Control	1000 Hz PID Loop
Max Holding Force	15 N per layer
Nominal Friction	~2 N per layer
Input	
Touch Sensing	Binary (grasped/not grasped)
Pressure Sensing	via motor current (500 mV/A), per layer
Motion Tracking	6DOF (Vive Tracker)

Table 1: Technical specifications of the X-Rings controller.

Proto-pasta CDP11705) in the top (index) layer; this is used to sense whether the user is currently grasping the device.

3.5 Device Functionality

Table 1 summarizes the technical specifications of the X-Rings controller. The primary functionality of the device is real-time rendering of graspable surfaces. Each layer can expand from 5.7 cm diameter to 7.7 cm diameter in approximately 100 ms, and can support radial loads of up to 15 N before backdriving the mechanism, gears and motor. Due to nominal friction in the 15:1 motor gearing (approx. 2 N), the device is not currently able to support high fidelity closed-loop force control. However, current readings from each motor driver do provide an indication of the grasping force applied to each layer, as the motors draw more current as a result of opposing forces.

X-Rings presently uses motor current information in three ways. First, when calibrating the device position, each layer is driven inward until a current spike is observed indicating full retraction and motor stall. To prevent damaging motors by drawing too much current, we also place a safety threshold on the maximum current used by a motor in response to a user squeezing the device or the mechanism causing motor stall. If the user causes a layer to draw more than 0.6 A (by applying approximately 10 N of force) for over 1 second, power is cut to that layer for 3 seconds, after which it will attempt to reach its desired position again. When this threshold is reached, users are notified either visually in VR or haptically through brief vibration of the motor. Finally, we can leverage current measurement to render "squeeze-responsive" objects in VR.

Unlike most existing shape displays which are purely output devices [40, 41], X-Rings is able to trigger visuo-haptic events in response to applied forces. This can be used to render objects that break or collapse above certain force loads. For example, Figure 5 shows a graspable clay jug which breaks when the user squeezes the controller beyond a set threshold on X-Rings. The breaking of the virtual object is accompanied by a retreat of the controller




Figure 5: A user squeezes X-Rings as it renders a virtual clay jug (left) until it reaches the measured motor current reaches a threshold, triggering the object to break (middle) and the device to retract its rings (right). Yellow lines indicate the device's original shape.

rings and a sudden reduction in resistance felt by the user's fingers. X-Rings can communicate this action to the virtual interface in order to trigger a visual animation of the object shattering, for example. Such force triggers can also be used to render the non-linear responses of haptic UI elements (e.g. buttons, switches) or simulate different object properties.

In addition to current sensing, X-Rings also leverages capacitive touch sensing to obtain additional context from the user. The present prototype uses a single touch sensor for the entire controller (a conductive finger rest located on the index finger layer) to determine whether the user is grasping the rendered object or not. We use this signal to ensure that a shape-change (between different objects) happens after the user has released the previous object.

3.6 Integration with VR

X-Rings connects to a PC via a USB Serial connection (115200 baud) and is powered via a 12V/3A wall power adapter. An HTC Vive setup (head-mounted display and two base stations) is used as the VR platform and controller tracker. VR applications are programmed in Unity 2019, which transmits appropriate device position commands to an interface application via UDP. The interface application then parses and forwards commands to the device via USB Serial. X-Rings sends a device status (including touch state, motor currents, and layer positions) back to Unity at 100 dataframes/second.

3.7 Shape Sampling

To determine the appropriate target shape for any given graspable object in VR, we use a technique we refer to as shape sampling. For each graspable object, a cylindrical volume is defined centered over the desired grasp region, which we call the grasp volume. This volume is the same height as the X-Rings controller's four layers, but twice as wide. Four rays are then cast inward from points on the grasp volume boundary corresponding to the four device layers. The




Figure 6: Shape sampling procedure for X-Rings. By casting a ray (red) from the boundary of the grasp volume toward the object, we measure the required radius of each ring (yellow). This value is then sent to X-Rings for rendering.

desired extension of each layer is then determined as the distance from the grasp volume’s central axis to the corresponding ray’s collision point on the object surface (see Figure 6). In the current application, the grasp volume for each object is fixed beforehand; however, for objects with variable grasp location (such as a staff), the grasp volume could be shifted in real-time.

In instances when the object surface lies outside the render volume of X-Rings, the rendered surface can be scaled such that the maximum (and/or minimum) object radius corresponds to the maximum (minimum) device radius. The dominance of the visual sense over haptic perception lends some flexibility to this mapping [35]. However, as shown by our evaluation, larger differences between the rendered and visual geometries can make rendered objects more difficult to identify.

4 USER STUDY: OBJECT RECOGNITION

The primary functionality requirement for X-Rings is the ability to render shapes that are distinguishable and recognizable when felt by a user. We performed an object recognition study to assess the rendering capabilities of X-Rings, evaluating participants’ ability to associate a felt, haptically-rendered shape with its visually-rendered counterpart. To further explore the impacts of physical rendering limitations on perception, we study two operating conditions: full dynamic range (where each layer can render diameters between 5.7-7.7 cm) and half dynamic range (where each layer is limited to a max diameter of 6.7 cm). This latter condition mimics the effect of a smaller device with reduced displacement capabilities, and was informed by our own development experience in which we had previously found that the dynamic range of the device impacted our shape perception.

4.1 Participants

We recruited 10 right-handed participants, ages 18 to 50 (5 male, 5 female) to take part in this study. To compensate for limitations in participation recruitment due to COVID-19, we designed a within-subjects study with multiple repetitions inspired by psychophysics studies [27]. A power analysis in G*Power confirmed that for a medium effect-size of $f = 0.25$ and error probability of $\alpha = 0.05$, an experiment with 2 groups (full and half dynamic range) and 18

repeated measures (6 objects with 3 repetitions) yields a power of above 0.8 for a total sample size of 10.

4.2 Methods

4.2.1 Objects. We tested 6 different virtual objects (shown in Figure 7), selected to represent a variety of commonly grasped and manipulated object geometries: a large egg, a tankard, a goblet, a sword handle, a bottle, and a chisel. These objects span a range of curvatures, from large round profiles (e.g. the egg) to narrow profiles with subtle features (e.g. the chisel). The dimensions of each object profile as rendered on X-Rings are shown in Figure 7.

Objects wider than the max diameter of the device, (such as the goblet) or narrower than the minimum diameter (such as the sword) were mapped to fit within the dynamic range of the X-Rings. In the full-dynamic range condition, object profiles were mapped between to a minimum diameter of 5.7 cm and max of 7.7 cm. In the half dynamic range condition the max rendered diameter was 6.7 cm.

4.2.2 Experimental Setup & Procedure. Participants were seated at a table and wore an HTC Vive head-mounted display (HMD), which showed all of the candidate objects in a virtual display before them. Each virtual object was labeled with a number, and marked with four green spheres indicating the grasp location for each finger, as shown in Figure 7.

The X-Rings controller was located on the table and was not visually displayed to the participant in VR. That is, the only feedback about rendered shapes was received through grasping and feeling the device. Prior to donning the HMD, participants strapped their right hand to the controller and were instructed to leave their hand/controller resting on the table throughout the study.

For each trial, participants were asked to grasp the device and select the virtual object which best matched the shape they felt. Participants were able to freely grasp and explore the rendered surface until they made their selection. Participants were not able to see their hand, the device, nor any visual representation of either. All candidate objects remained visible throughout the experiment.

Following their selection, participants were instructed to release their grasp. The device was then zeroed and the next shape was rendered. Within each dynamic range condition, participants experienced each of the 6 objects 3 times presented randomly (18 trials per condition) for a total of 36 experimental trials.

4.3 Results

Results of the user study are shown in Figure 8. For each condition, a confusion matrix indicates the distribution of selected objects for each physically rendered object. Diagonal elements indicate correctly identified objects, while off-diagonal elements indicate errors. The aggregated error rates for each object are also shown for each condition.

A two-way repeated measures ANOVA revealed a significant interaction effect between Object and Condition ($F(5, 45) = 3.749, p = 0.006$) on the number of selection errors made. Significant main effects of both Object ($F(5, 45) = 3.712, p = 0.007$) and Condition ($F(1, 9) = 31.016, p = 0.0003$) were also found. Post-hoc analysis using Bonferroni-corrected pairwise t-tests further showed that using half of the device’s dynamic range led to a significant increase in selection errors for the rendered goblet (Object 3; $p = 0.007$) and



Figure 7: Six objects tested in the user study. Each virtual object is shown alongside its corresponding rendered shape on X-Rings as well as the physically rendered dimension of each layer, as computed according to the Shape Sampling protocol detailed in Section 3.7. Dimensions for both full and half dynamic range conditions are given. Green circles on each virtual object indicate the grasp point for each finger.

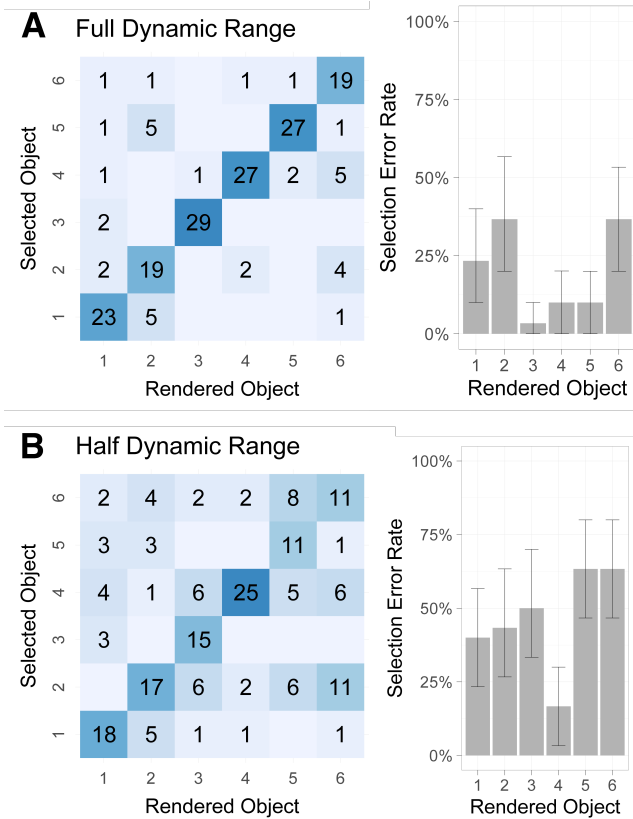


Figure 8: Results of the object recognition user study for the Full Dynamic Range condition (A) and the Half Dynamic Range condition (B). Confusion matrices show the number of times each object was selected for a given rendered object. The overall selection error rate for each physically rendered object is also shown, with 95% confidence intervals.

bottle (Object 5; $p = 0.003$) compared to the full dynamic range condition. Overall, reducing the device’s dynamic range from 2 cm to 1 cm more than doubled participants’ selection error rate from 20% to 46%.

By observation of the results, we also see that certain objects were more likely to lead to confusion regardless of the device’s dynamic range, such as the chisel (Object 6) and tankard (Object 2). These objects have subtle changes in their surface geometry, and thus were expected to be more difficult to identify than, say, the sword (Object 4) which has a consistent cylindrical geometry or the egg (Object 1) which has large, distinct changes in curvature.

The interaction effect between object and condition described above highlights the impact that a device’s dynamic range can have on object recognition. While the goblet (Object 3) and bottle (Object 5) were easily identifiable (error rate < 10%) with a full dynamic range of 2 cm, their features became less distinguishable with a dynamic range of only 1 cm. This highlights the challenge in distinguishing subtle surface differences using coarse geometry discretization within a limited dynamic range.

In both conditions, the error rate was well below Bayesian random chance selection. Given 6 objects the probability of correctly picking an object at random is 1/6, whereas the probability of committing an error is of 5/6 (83%). A total of 36 errors were made out of 180 trials (6 objects in 3 repetitions for 10 participants) in the full dynamic range condition, which shows an error rate of 20% – well below that expected for random selection. In the half dynamic range condition, the errors increased to 83, representing a 46% error rate, which nonetheless was still well below random selection.

5 DEMO APPLICATION: SELECTING AND GRASPING OBJECTS IN A VIRTUAL SCENE

To illustrate the potential of X-Rings as a shape-changing haptic controller for VR, we developed an application scenario in Unity that allows users to freely pick up and interact with a variety of




Figure 9: VR view of the virtual forge scene, with objects prepared for interaction with X-Rings. Grasp sensing volume and grasp points are shown here for illustrative purposes and not visible by the user.



Figure 10: Object prediction uses a temporally smoothed hand velocity vector (red arrow) to find the nearest graspable object to this ray (marked by a yellow dot).

objects common for a medieval game scenario while shape feedback is rendered by X-Rings (shown in Figure 9).

In the scene, the user has a selection of objects available to pick up from a table in front of them: a wooden tankard, a goblet, a sword, an axe, a bottle, and a chisel. (All but the axe are identical to objects investigated in our user study.) As the user reaches for any of these objects, the system predicts their target and sends the appropriate commands to X-Rings in order to render the appropriate shape.

5.1 Target Object Prediction

To anticipate which object the user is reaching for, we measure the vector direction of the user's hand motion as tracked by the controller. To reduce noise, we smooth this vector using an exponential filter with a constant of 0.95. The intended target object is predicted as the nearest graspable object nearest to the line defined by the current hand motion direction (see Figure 10). If the user's hand is within 50 cm of the predicted target object, the shape of the object is commanded to X-Rings for rendering. New object shapes are only commanded if the touch sensor on X-Rings indicates the user is not already grasping the device. Because the device can rapidly update its shape (< 0.1 seconds), this approach works well even as the user reaches between two objects that are very close together.

5.2 Interactivity

We leverage animations of a fully articulated virtual hand to further increase immersion in the application. When an object is registered as a graspable target and the user's hand enters the grasp volume, an animation is triggered rendering the grasping of the target object as the user closes their real hand (Figure 11). If the user releases their



Figure 11: As the user touches the controller (left), a quick animation of their virtual hand closing its fingers around the virtual object is triggered (right). A similar animation of the virtual hand releasing the object is displayed when the user releases their grasp.

grasp on X-Rings, the change is sensed via capacitive sensing and the virtual hand opens, releasing the virtual object. It is also possible to leverage pressure sensing (via changes in measured motor current) to enable effects such as squishing or breaking held objects if the current is higher than a set threshold. Once this threshold is reached on any layer, a breaking animation can be triggered and the device fully retracts, as previously illustrated in Figure 5.

6 DISCUSSION, LIMITATIONS, & FUTURE WORK

The results of our study indicate that X-Rings is effective at rendering distinguishable and recognizable graspable objects, particularly when operating at its full dynamic range of 2 cm. Moreover, its rapid rendering speed (approx. 100 ms for 2 cm expansion), graspable form factor, and capacity for physical input via touch and pressure sensing make X-Rings promising as a tool for physical interaction in VR. Our study also highlights the negative impact of reduced device dynamic range on grasped shape perception. However, the observed errors are likely also a result of the physical discretization inherent to shape displays, or perhaps differences between the user's expectations and rendered haptic feedback.

We also note that the task of identifying virtual objects by felt shape alone involves greater scrutiny on the device than might exist in normal interaction scenarios. For example, while the rendered chisel shape yielded higher selection error rates, this does not necessarily imply that the rendered shape is unsuitable for the virtual chisel. Rather, it suggests that this shape may be suitable for other objects as well. In this way, we expect that the results of our evaluation are quite conservative.

While X-Rings has demonstrated potential as a hand-worn shape-rendering controller for grasping virtual objects, several challenges and opportunities for improvement still exist. Crucially, the radially-symmetric nature of the expansion mechanism prevents it from rendering asymmetric object profiles. However, this comes at the benefit of using only 4 actuators compared to the large number of actuators used by traditional shape displays. This also results in reduced cost, power requirement, and weight while increasing the reliability of the system. The use of 4 approximately finger-width rings leads to naturally graspable shapes, where each finger is supported by a single ring. Future work will explore using smaller motors/rings so that each finger may be support two ring mechanisms, enabling the device to vary the contact surface normal and better approximate more general object surfaces. Additionally, future work can also explore specifically designed cam profiles that expand asymmetrically (such as an ellipse) or asynchronously,

allowing some asymmetry in rendered object without increasing the number of actuators required.

Another interesting option is to add two additional rings – one above the hand and one below – to provide additional haptic context to the sides of the hand and fingers. This would enable the controller to render features that lie outside the explicit grasp volume but provide important cues about the grasped object, such as the guard of a sword hilt felt by the top of the hand, or the base of a goblet felt by the bottom of the hand. These subtle additional cues would likely increase the perceived realism of certain rendered objects.

Additionally, the current X-Rings system is not wireless. While all motors and actuation mechanisms are housed within the controller, the device is tethered to an external control board which connects to a PC for USB communication and a wall adapter for power. By using a smaller microcontroller and motor controllers as well as wireless communication and battery power, we believe that future versions of X-Rings could incorporate all components into an untethered controller.

The "inside-out" paradigm for shape rendering enables a compact form-factor where all system elements are housed within a controller handle that can easily be picked up and manipulated and better fit a wide range of hand sizes. However, this paradigm inherently limits the minimum size of a rendered object, due to the need to contain the all components within the rendered shape. In the current implementation of X-Rings, the smallest rendered objects have a diameter of 5.7 cm. The dominance of the visual sense, however, enables us to believably render grasped objects slightly smaller or larger than the dynamic range of the device, as long as the shape is similar to the object. Our user study was partially designed towards understanding how much smaller we could make our design; however, the results suggest that significant differences between the scale of the rendered and observed shapes can hinder recognition. Other expansion mechanisms, such as telescoping components, can potentially be explored to achieve greater dynamic range, though likely at the cost of significant complexity. Furthermore, it may also be possible to position actuators in the controller handle and transmit mechanical power through tendons or linkages in order to reduce the minimum renderable diameter.

While much less dependent on the user's hand size than traditional "outside-in" haptic exoskeletons, there remains room to increase the flexibility of the X-Rings design. Presently, the device assumes that each layer provides feedback for a single finger. While the 2.5 cm layer thickness was suitable for all participants in our study to naturally rest one finger on each layer, individuals with smaller hands may contact some layers with multiple fingers. The ability to adjust layer size and/or spacing would further assist in ensuring broad accessibility to users.

In the current controller design, the device is strapped to the user's hand allowing them to freely grasp and release rendered objects. This design however uses a narrow handle that rests against the palm in order to secure the strap, preventing shape feedback from being rendered to the palm. While shape feedback along the fingers is perhaps most important for grasp rendering, X-Rings could also be adapted for a handle-less design that would allow the rendering of complete graspable surfaces.

Future work should further study the importance of human perceptual thresholds (e.g., shape discrimination) for graspable shape

generation. In our study we have focused on the device's physical dynamic range, but we believe that combining a device such as X-Rings with haptic retargeting techniques [4, 44] could further amplify its perceived rendering capabilities. Such an approach leverages visual dominance [35], or the tendency of vision to dominate perception in sensory conflicts. Previous work has demonstrated that one's grasp can be successfully "resized" in VR within certain thresholds [7]. In the case of X-Rings, it would be beneficial to understand the extent to which vision dominates proprioception during grasped-shape perception, such that the allowable "mismatch" between haptically rendered and visually presented objects can be better evaluated. We hypothesize that additional visual cues, such as animated virtual fingers, would strengthen this illusion by increasing the user's sense of embodiment, perhaps further increasing the range of allowable mismatch between the rendered and visual shape. We see X-Rings as a promising tool to aid in the investigation of these important questions.

Further research could also shed light on additional use cases of X-Rings, such as for bimanual interaction, either with multiple controllers or using one hand to interact with surfaces rendered on the other. Additionally, while the current implementation allows some valuable physical input to the device through motor current and touch sensing, higher quality gearmotors and/or sensors could enable richer haptic rendering of surface properties such as compliance through closed-loop force control.

Additionally, while we have focused our initial evaluation on the rendering capabilities of X-Rings alone, we believe a comparison against other haptic controllers that enable grasping (such as Wolverine [11] or CLAW [12]) would provide important context and highlight the pros and cons of each rendering approach.

Finally, while X-Rings has been tested and designed for VR, it may be a useful tool for augmented reality (AR) applications as well. Its compact, inside-out design could be particularly helpful in avoiding environmental occlusions while interacting with virtual objects situated in the real world.

7 CONCLUSION

We introduced X-Rings, a novel hand-mounted 360° shape display for grasping in VR. To the best of our knowledge, X-Rings is the first hand-scale shape display to support whole-hand encountered-type grasping of virtual objects and surfaces. Through X-Rings, we explore the use of an "inside-out" paradigm for graspable shape rendering, leveraging the VR controller form factor to house actuation components within the controller handle, as opposed to traditional exoskeletons which require cumbersome external equipment. Additionally, by strapping to the user's knuckles, X-Rings affords the ability to naturally grasp and release virtual objects freely. X-Rings is also able to sense direct input through touch and pressure sensing, both of which are important for compelling dynamic physical interactions with virtual objects. Our evaluation shows that X-Rings is effective in rendering identifiable shapes and highlights the impact of renderable dynamic range on grasped shape recognition. We hope that our design and implementation of X-Rings helps influence and inspire future research and development of VR and AR controllers that provide natural and expressive shape rendering.

REFERENCES

- [1] 2020. Dexmo. <https://origin.dextarobotics.com/>.
- [2] 2020. Haptx. <https://haptx.com/>.
- [3] 2020. Vive. <https://www.vive.com/>.
- [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 2016 chi conference on human factors in computing systems*. ACM, 1968–1979.
- [5] Hrvoje Benko, Christian Holz, Mike Sinclair, and Eyal Ofek. 2016. Normaltouch and textretouch: High-fidelity 3d haptic shape rendering on handheld virtual reality controllers. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. 717–728.
- [6] Christopher C Berger, Mar Gonzalez-Franco, Eyal Ofek, and Ken Hinckley. 2018. The uncanny valley of haptics. *Science Robotics* 3, 17 (2018).
- [7] Joanna Bergström, Aske Mottelson, and Jarrod Knibbe. 2019. Resized grasping in vr: Estimating thresholds for object discrimination. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*. 1175–1183.
- [8] Jonathan Blake and Hakan B Gurocak. 2009. Haptic glove with MR brakes for virtual reality. *IEEE/ASME Transactions On Mechatronics* 14, 5 (2009), 606–615.
- [9] 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 (2002), 256–263.
- [10] Inrak Choi, Heather Culbertson, Mark R Miller, Alex Olwal, and Sean Follmer. 2017. Grability: 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.
- [11] 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 *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 986–993.
- [12] Inrak Choi, Eyal Ofek, Hrvoje Benko, Mike Sinclair, and Christian Holz. 2018. Claw: A multifunctional handheld haptic controller for grasping, touching, and triggering in virtual reality. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [13] Heather Culbertson, Samuel B Schorr, and Allison M Okamura. 2018. Haptics: The present and future of artificial touch sensation. *Annual Review of Control, Robotics, and Autonomous Systems* 1 (2018), 385–409.
- [14] Maxime Daniel, Guillaume Rivière, and Nadine Couture. 2019. Cairnform: A shape-changing ring chart notifying renewable energy availability in peripheral locations. In *Proceedings of the Thirteenth International Conference on Tangible, Embedded, and Embodied Interaction*. 275–286.
- [15] Cathy Fang, Yang Zhang, Matthew Dworkman, and Chris Harrison. 2020. Wireality: Enabling Complex Tangible Geometries in Virtual Reality with Worn Multi-String Haptics. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–10. <https://doi.org/10.1145/3313831.3376470>
- [16] George W Fitzmaurice, Hiroshi Ishii, and William AS Buxton. 1995. Bricks: laying the foundations for graspable user interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 442–449.
- [17] 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.
- [18] Adrian Freed. 2018. FastTouch. <https://github.com/adrianfreed/FastTouch>.
- [19] Eric J Gonzalez, Parastoo Abtahi, and Sean Follmer. 2020. REACH+ Extending the Reachability of Encountered-type Haptics Devices through Dynamic Redirection in VR. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. 236–248.
- [20] Mar Gonzalez-Franco and Jaron Lanier. 2017. Model of illusions and virtual reality. *Frontiers in psychology* 8 (2017), 1125.
- [21] C. Harrison and S. E Hudson. 2009. Providing Dynamically Changeable Physical Buttons on a Visual Display. In *CHI '09*. 299–308.
- [22] Seongkook Heo, Christina Chung, Geehyuk Lee, and Daniel Wigdor. 2018. Thor's hammer: An ungrounded force feedback device utilizing propeller-induced propulsive force. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–11.
- [23] Seongkook Heo, Jaeyeon Lee, and Daniel Wigdor. 2019. PseudoBend: Producing haptic illusions of stretching, bending, and twisting using grain vibrations. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*. 803–813.
- [24] Charles Hoberman. 1991. Radial expansion/retraction truss structures. US Patent 5,024,031.
- [25] HyunKi In, Kyu-Jin Cho, KyuRi Kim, and BumSuk Lee. 2011. Jointless structure and under-actuation mechanism for compact hand exoskeleton. In *2011 IEEE International Conference on Rehabilitation Robotics*. IEEE, 1–6.
- [26] Sungjune Jang, Lawrence H Kim, Kesler Tanner, Hiroshi Ishii, and Sean Follmer. 2016. Haptic edge display for mobile tactile interaction. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. 3706–3716.
- [27] Lynette A Jones and Hong Z Tan. 2012. Application of psychophysical techniques to haptic research. *IEEE transactions on haptics* 6, 3 (2012), 268–284.
- [28] Robert Kovacs, Eyal Ofek, Mar Gonzalez-Franco, Alexa F Siu, Sebastian Marwecki, Christian Holz, and Mike Sinclair. 2020. Haptic PIVOT: On-Demand Handhelds in VR. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology (UIST '20)*. Association for Computing Machinery, New York, NY, USA.
- [29] Ken'ichi Koyanagi, Yuki Fujii, and Junji Furusho. 2005. Development of VR-STEF system with force display glove system. In *Proceedings of the 2005 international conference on Augmented tele-existence*. 91–97.
- [30] Dávid Lakatos. 2012. AMPHORM: form giving through gestural interaction to shape changing objects. Ph.D. Dissertation. Massachusetts Institute of Technology.
- [31] Jaeyeon Lee, Mike Sinclair, Mar Gonzalez-Franco, Eyal Ofek, and Christian Holz. 2019. TORC: A virtual reality controller for in-hand high-dexterity finger interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [32] Daniel Leithinger, Sean Follmer, Alex Olwal, and Hiroshi Ishii. 2015. Shape displays: Spatial interaction with dynamic physical form. *IEEE computer graphics and applications* 35, 5 (2015), 5–11.
- [33] Jifei Ou, Mélina Skouras, Nikolaos Vlavianos, Felix Heibeck, Chin-Yi Cheng, Jannik Peters, and Hiroshi Ishii. 2016. aeroMorph-heat-sealing inflatable shape-change materials for interaction design. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. 121–132.
- [34] Jifei Ou, Lining Yao, Daniel Tauber, Jürgen Steimle, Ryuma Niiyama, and Hiroshi Ishii. 2014. jamSheets: thin interfaces with tunable stiffness enabled by layer jamming. In *Proceedings of the 8th International Conference on Tangible, Embedded and Embodied Interaction*. 65–72.
- [35] Michael I Posner, Mary J Nissen, and Raymond M Klein. 1976. Visual dominance: an information-processing account of its origins and significance. *Psychological review* 83, 2 (1976), 157.
- [36] Harpreet Sareen, Udayan Umapathi, Patrick Shin, Yasuaki Kakehi, Jifei Ou, Hiroshi Ishii, and Pattie Maes. 2017. Printflatables: printing human-scale, functional and dynamic inflatable objects. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 3669–3680.
- [37] Jotaro Shigeyama, Takeru Hashimoto, Shigeo Yoshida, Taiju Aoki, Takuji Narumi, Tomohiro Tanikawa, and Michitaka Hirose. 2018. Transcalibur: weight moving VR controller for dynamic rendering of 2D shape using haptic shape illusion. In *ACM SIGGRAPH 2018 Emerging Technologies*. 1–2.
- [38] Timothy M Simon, Ross T Smith, and Bruce H Thomas. 2014. Wearable jamming mitten for virtual environment haptics. In *Proceedings of the 2014 ACM International Symposium on Wearable Computers*. 67–70.
- [39] Mike Sinclair, Eyal Ofek, Mar Gonzalez-Franco, and Christian Holz. 2019. CapstanCrunch: A Haptic VR Controller with User-Supplied Force Feedback. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology* (New Orleans, LA, USA) (UIST '19). Association for Computing Machinery, New York, NY, USA, 815–829. <https://doi.org/10.1145/3332165.3347891>
- [40] Alexa F Siu, Eric J Gonzalez, Shenli Yuan, Jason B Ginsberg, and Sean Follmer. 2018. Shapeshift: 2D spatial manipulation and self-actuation of tabletop shape displays for tangible and haptic interaction. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 291.
- [41] Yuqian Sun, Shigeo Yoshida, Takuji Narumi, and Michitaka Hirose. 2019. PoCoPo: A handheld vr device for rendering size, shape, and stiffness of virtual objects in tool-based interactions. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [42] Shan-Yuan Teng, Tzu-Sheng Kuo, Chi Wang, Chi-huan Chiang, Da-Yuan Huang, Liwei Chan, and Bing-Yu Chen. 2018. PuPoP: Pop-up Prop on Palm for Virtual Reality. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology* (Berlin, Germany) (UIST '18). Association for Computing Machinery, New York, NY, USA, 5–17. <https://doi.org/10.1145/3242587.3242628>
- [43] Eric Whitmire, Hrvoje Benko, Christian Holz, Eyal Ofek, and Mike Sinclair. 2018. Haptic revolver: Touch, shear, texture, and shape rendering on a reconfigurable virtual reality controller. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [44] Jackie Yang, Hiroshi Horii, Alexander Thayer, and Rafael Ballagas. 2018. VR grabbers: ungrounded haptic retargeting for precision grabbing tools. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. 889–899.
- [45] Lining Yao, Ryuma Niiyama, Jifei Ou, Sean Follmer, Clark Della Silva, and Hiroshi Ishii. 2013. PneuUI: pneumatically actuated soft composite materials for shape changing interfaces. In *Proceedings of the 26th annual ACM symposium on User interface software and Technology*. 13–22.
- [46] Shigeo Yoshida, Yuqian Sun, and Hideaki Kuzuoka. 2020. PoCoPo: Handheld Pin-based Shape Display for Haptic Rendering in Virtual Reality. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [47] Andre Zenner and Antonio Krüger. 2017. Shifty: A weight-shifting dynamic passive haptic proxy to enhance object perception in virtual reality. *IEEE transactions on visualization and computer graphics* 23, 4 (2017), 1285–1294.

[48] André Zenner and Antonio Krüger. 2019. Drag: on: A virtual reality controller providing haptic feedback based on drag and weight shift. In *Proceedings of the*

2019 CHI Conference on Human Factors in Computing Systems. 1–12.