Ungrounded Vari-Dimensional Tactile Fingertip Feedback for Virtual Object Interaction

Eric M. Young

University of Pennsylvania
Philadelphia, Pennsylvania, USA
Max Planck Institute for Intelligent Systems
Stuttgart, Germany
yoeiric@seas.upenn.edu

Katherine J. Kuchenbecker

Max Planck Institute for Intelligent Systems
Stuttgart, Germany
kjk@is.mpg.de

Figure 1: Examples of the visual scene (top) and corresponding tactile cues (bottom) during virtual shape discrimination (left) and mass discrimination (right) tasks. In these examples, the wearable device is rendering 6-DOF tactile feedback to the user's index finger; the platform orients itself to be tangent to the surface, contacts the appropriate position on the fingertip with varying normal force, and simultaneously applies tangential and rotational shear caused by the interaction taking place.

ABSTRACT

Compared to grounded force feedback, providing tactile feedback via a wearable device can free the user and broaden the potential applications of simulated physical interactions. However, neither the limitations nor the full potential of tactile-only feedback have been precisely examined. Here we investigate how the dimensionality of cutaneous fingertip feedback affects user movements and virtual object recognition. We combine a recently invented 6-DOF fingertip device with motion tracking, a head-mounted display, and novel contact-rendering algorithms to enable a user to tactilely explore immersive virtual environments. We evaluate rudimentary 1-DOF, moderate 3-DOF, and complex 6-DOF tactile feedback during shape discrimination and mass discrimination, also comparing to interactions with real objects. Results from 20 naive study participants show that higher-dimensional tactile feedback may indeed allow completion of a wider range of virtual tasks, but that feedback

dimensionality surprisingly does not greatly affect the exploratory techniques employed by the user.

CCS CONCEPTS

• **Human-centered computing** → **Haptic devices**; *User studies*.

KEYWORDS

tactile feedback, fingertip, wearable, six-degree-of-freedom, haptic exploration

ACM Reference Format:

Eric M. Young and Katherine J. Kuchenbecker. 2021. Ungrounded Vari-Dimensional Tactile Fingertip Feedback for Virtual Object Interaction. In *CHI Conference on Human Factors in Computing Systems (CHI '21)*, May 8–13, 2021, Yokohama, Japan. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3411764.3445369>

1 INTRODUCTION

Daily physical interactions, such as flipping a switch or inserting a key into a lock, often involve complex tactile cues at the fingertips (Fig. 2). These soft, distal parts of the human hand are densely enervated with mechanoreceptors that enable precise perception of diverse object properties as well as competent control of contact [16]. Yet the haptic feedback provided by mainstream handheld devices, such as smartphones or virtual reality controllers, is typically limited to vibration [6]; these common haptic rendering approaches can catch the user's attention and communicate the

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 ACM 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.

CHI '21, May 8–13, 2021, Yokohama, Japan

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8096-6/21/05...\$15.00

<https://doi.org/10.1145/3411764.3445369>

intensity of the event taking place, but they cannot apply the same contact conditions to the skin as most of the original interactions being simulated. Even within academic research, the importance of reducing the weight and encumbrance of wearable haptic interfaces means that cutting-edge devices are well-suited to render only particular haptic interactions, rather than a more general set that ideally approaches all haptic interactions. To the best of our knowledge, no existing wearable haptic device can authentically render all three of the common interactions portrayed in Fig. 2.

We are curious to discover whether six-degree-of-freedom (6-DOF) approximation of tactile contact offers any benefits over the lower-dimensional tactile rendering approaches that have been previously explored, as well as how strongly various dimensions of tactile feedback affect human motions during virtual object interaction. Although compelling haptic rendering can be achieved by applying feedback at other parts of the user's hand or body, the goal of this paper is to examine the role of tactile dimensionality at the fingertip, as the fingertip plays an essential role in manipulation tasks and is commonly the focus of haptic device research. As shown in Fig. 1, we are purposefully focusing on rich tactile cues and omitting the grounded forces and torques that physically oppose finger movements during real contact in order to achieve a large workspace while limiting system cost and user encumbrance; this popular strategy is sometimes called sensory subtraction [35]. Importantly, even our complex 6-DOF consideration of finger-surface interactions completely disregards the local shape and material properties of the object being contacted, which would require infinite dimensions for proper portrayal. Instead, we aim to represent any physical contact with the surface's local tangent plane as a first-order approximation, in much the same way that computer graphics models often represent surfaces as polygon meshes.

After analyzing related work (Sec. 2), we present our approach to creating a novel integrated system that can render 6-DOF fingertip tactile cues to a user exploring an immersive virtual environment (Sec. 3). This setup is capable of recreating a wide range of dynamic tactile interactions that have previously been achieved by individual devices published in the wearable haptics literature. We utilize this single system to simulate systems with fewer degrees of freedom in order to examine how the number of degrees of freedom of a device affects the user (Sec. 4).

The results of our study (Sec. 5) show that the perceptual performance of users is boosted by more complex rendering when the additional DOFs are relevant to the task at hand, in some cases approaching performance associated with real physical interactions. However, even this sophisticated system with 6-DOF tactile cues falls short of imitating the perceived realism associated with physical object interactions. Furthermore, users did not seem to adjust their exploratory movements based on the available tactile feedback DOFs, raising questions about whether naive individuals can take advantage of richer contact-based cues with minimal training and interaction time. We discuss our results (Sec. 6), reflect on the limitations of our approach (Sec. 7), and summarize our findings (Sec. 8) to enable other researchers to learn from this investigation. Our contributions include:

- Strategies for creating a system that allows a participant to explore virtual environments while feeling tactile fingertip feedback of variable dimensions.

Figure 2: Examples of three real-world interactions (upper panels) in which different dimensions of tactile feedback are experienced at distinct locations on the index fingertip (lower panels). In the first example of pushing a button, only normal force is felt, near the tip of the finger. Holding a heavy cup brings in translational shear in addition to normal force, with contact happening on the flat area of the fingerpad. The third example consists of pinching a card between the thumb and the side of the index finger while the card is being pulled away (black arrow); this scenario causes normal force, translational shear, and rotational shear all collocated on the radial side of the index finger.

- A framework for rendering 1-DOF, 3-DOF, and 6-DOF fingertip tactile cues during virtual interactions with stationary surfaces and mobile objects.
- Quantitative evidence to demonstrate that users can complete representative virtual tasks (shape discrimination and mass discrimination) with cutaneous feedback alone.
- An analysis of the effect of the dimensionality of tactile fingertip feedback on perceptual performance and exploratory motions in the two studied tasks.

2 RELATED WORK

2.1 Object Manipulation

Detailed studies of grasping and manipulation demonstrate that common tasks elicit complex haptic sensations on the fingertips. Kamakura et al. [17] classified grips that humans use to grasp common objects, finding that different objects consistently elicited contact with different parts of the index finger, including the volar, radial, and distal aspects. Atzori and Müller [1] presented the Ninapro database, which included over 20 grasping movements that again elicited normal force and shear force at multiple locations on the index finger. Feix et al. [8] compiled existing human grasp taxonomies into a single GRASP taxonomy, which divided grasps into power, intermediate, and precision categories. Power grasps tended to require the volar aspect of the index finger, intermediate grasps tended to require the radial aspect of the index finger, and precision grasps often required the distal-most tip of the index finger. Figure 2 illustrates a small sample of common interactions discussed in these references, also depicting the corresponding forces that act at different locations on the index finger.

2.2 Fingertip Haptic Devices

Wearable one-degree-of-freedom (1-DOF) haptic devices have been presented for rendering contact location [36], tangential skin stretch [40] and normal indentation into the fingerpad [44]. Researchers have combined multiple 1-DOF devices into one system to render more complex interactions, such as using one voice coil on each finger to enhance object manipulation [24], using up to 12 asymmetric actuators to render 6D directional cues to the hand [7], or using a single motor to move an electrode array around the finger to present both skin stretch and pressure distribution with high spatial resolution [48]. This idea has also been extended to the design of various haptic hand exoskeletons and gloves that apply 1D haptic feedback to each fingertip, e.g., [13, 26, 41].

The field of wearable haptics has also fostered the development of devices with more degrees of freedom. Two-DOF fingertip devices have provided 1D kinesthetic force feedback with 1D vibrotactile feedback [5], 1D tangential shear with 1D normal displacement [27, 43, 46], and 2D tangential shear [12, 21, 42, 45]. Two 2-DOF tangential shear devices have been worn on the index finger and thumb to successfully communicate 5-DOF directional cues [14]. Three-DOF fingertip devices have demonstrated the ability to move a contact element along a spherical surface around the fingertip and translate in the direction normal to the sphere [10, 38], tilt a contact surface in place with 2 DOF and then extend into the volar side of the fingertip [2, 32], or provide 3-DOF feedback that couples orientation and position in a more complex manner [25, 34]. At least one design has offered 3-DOF translation of a contact element with a fixed orientation to render more precise shear cues [37], and another has presented 3-DOF translation via asymmetric vibrations from three orthogonal actuators [19].

We refer the reader to a review by Pacchierotti et al. [31] for more information on the breadth of wearable haptic systems for the fingertip and hand.

2.3 Implementation of Tactile Fingertip Feedback

Several studies have demonstrated that providing cutaneous and kinesthetic force feedback to the fingertip can improve curvature discrimination [3, 9, 34] or the perception of edges and texture [10], compared to providing kinesthetic feedback alone. Cutaneous feedback alone has been shown to reduce grasping forces during manipulation tasks [25, 30], or allow participants to find and follow a virtual surface [38], but the addition of kinesthetic feedback often further improves performance [30, 38].

Previous studies have shown that augmenting kinesthetic feedback with cutaneous skin stretch or slip can enhance the perception of friction [40] or reduce the force applied during a manipulation task [45]. Even in the absence of kinesthetic force feedback, skin stretch can allow participants to distinguish weight [27, 37] or directional cues [12, 21, 42].

A few wearable devices have rendered kinesthetic feedback to the finger itself, rather than the entire limb, by grounding the device on the base of the finger or on the hand. Benko et al. showed that the combination of kinesthetic and cutaneous feedback in this manner allowed the realistic presentation of virtual haptic interactions such as localizing targets or tracing lines [2]. Chinello et al. found

that cutaneous feedback alone yielded better performance than no haptic feedback during robotic palpation and VR exploration tasks, and that the addition of kinesthetic feedback further improved performance in these tasks [3].

Even from this brief summary of selected previous studies, it is clear that the addition of cutaneous feedback can improve performance in many tasks. However, these past works considered cutaneous feedback as a binary variable that is either present or not present, rather than considering and comparing the different types of cutaneous feedback that can be provided. In this work, we look more closely at how different dimensions of cutaneous feedback affect user performance to inform the design of future devices and the creation of future fingertip tactile feedback systems.

2.4 The Fingertip Puppeteer

A potential solution for a higher-dimensional wearable haptic device was recently demonstrated in the Fingertip Puppeteer, or Fuppeteer for short [50]. This unusual tactile display employs a parallel continuum manipulator to move a flat surface around the user's fingertip with six DOF (see Fig. 1 and this paper's supplementary video). As such, this device has the potential to render the entire 6-DOF set of possible interactions between a stationary index finger and a flat surface: one DOF for how hard the finger presses into the surface (making contact and normal force), two DOFs for rotating the finger such that a different location on the fingertip contacts the surface, two DOFs for sliding the finger along the surface (translational shear), and one DOF for twisting the finger on the surface (rotational shear).

The Fuppeteer can advantageously also simulate any other fingertip device that contacts the finger with a flat surface, since the ranges of motion for all such devices fall within its range of motion. We thus decided to use this unique device to explore how the dimensionality of cutaneous fingertip feedback affects user movements and virtual object recognition. Importantly, this device has never before been integrated into a haptic rendering system; it was previously only characterized from an engineering standpoint and used in a perceptual validation study where it remained stationary on a table and delivered tactile cues to the user's fingertip (passive touch) [50].

Despite its novelty and appeal, a few important limitations of the Fuppeteer were previously identified [50]. As mentioned in Sec. 1, the Fuppeteer is able to render the tangent plane of a surface; it is not designed to display edges or highly curved surfaces. We previously noted that end-effector pose error tends to accumulate over time, and we presented a quick calibration routine through which accumulated error could be eliminated. Due to the compliance of the parallel continuum manipulator design, some end-effector motions result in sustained oscillations. However, we previously demonstrated that trajectory constraints minimized these vibrations, resulting in no noticeable oscillatory forces once the platform contacts the user's fingertip; we implement the same trajectory constraints in this paper. Although the Fuppeteer is position controlled, the device's compliance also means that forces acting on the end-effector cause some amount of deformation, resulting in a complex relationship between configuration and stiffness. However, it was shown that feed-forward pose commands could be used to

Figure 3: A labeled diagram of the 6-DOF tactile fingertip feedback device [50] that we adapted for use in this project.

render recognizable contact and shear cues at all locations on the fingertip. We apply the same assumption in this paper. Although more work is needed to understand the complexities of 6-DOF tactile cues presented by the Fuppeteer, we can nonetheless use this device to render meaningful tactile feedback and answer our research questions.

3 SYSTEM INTEGRATION

The working principle of the selected 6-DOF device was initially proposed and demonstrated through a manually actuated prototype [49] and later refined and validated in a motorized device [50] that we adapted for use in this project. As illustrated in Fig. 3, six flexible, superelastic rods pass through a hand-held base platform and connect to a rigid end-effector that makes contact with the user's fingertip. We independently control the length of each rod, which enables us to set the position and orientation of the end-effector in a 6D workspace. As described below, a few key improvements were necessary to adapt the original design for use in a virtual environment.

3.1 Counterweight

For the study discussed in this paper, we ask users to hold the device while exploring a virtual environment for approximately 60 minutes. However, the mass of the entire device with its six motors is 0.45 kg. When developing the study, we observed user fatigue before all of the trials were complete. We reduce user fatigue via a 1.5 kg counterweight attached to the device's hand-held base platform via a three-pulley system that has a 4:1 reduction ratio. The pulley system reduces the acceleration of the counterweight itself, limiting the effect of the counterweight's inertia on the force applied to the device, and preventing slack in the counterweight line unless the device is accelerated upward at over about 39 m/s^2 (four times gravitational acceleration).

The device's entire workspace and a schematic of the pulley system are shown in Fig. 4. The counterweight force always includes

Figure 4: The workspace for the human-subject study (not drawn to scale). The counterweight pulley system is located above the middle of the back wall of the box, and it pulls upward on the haptic device to reduce the weight that must be supported by the participant.

a vertical component that reduces the effective weight of the device, but it also includes a horizontal component, since the device is not always located directly beneath the point where the cable exits the last pulley. The highest ratio of horizontal component to vertical component in the counterweight force occurs when the device is positioned in an upper, front corner of the workspace, where a static device feels a vertical force of 3.6 N and a horizontal force of 0.8 N. Although the non-vertical component of the counterweight force may affect the user's motion, we empirically found the benefit of reducing the effective weight to greatly exceed the drawback of slightly pulling the device in the horizontal direction.

3.2 Device Positioning

The calibration routine we devised [50] to eliminate accumulated error requires a designated calibration pose. To ensure that error is consistently minimized throughout our human-subject study, we calibrate the device before each trial. During the majority of our study, the user is wearing a headset and is thus unable to see the physical world. We provide virtual cues to aid the subject in reaching the calibration pose, but small errors in our tracking or in the subject's trajectory can make this task difficult. To facilitate reaching the calibration pose, we designed a holder that physically funnels the device to the correct position and orientation, as shown at the bottom of Fig. 3.

3.3 Finger Positioning

Another important improvement to the mechanical design of the selected device was the inclusion of a two-degree-of-freedom mechanism to position the fingertip. In [50], one strap was approximately located at the distal interphalangeal (DIP) joint of the index finger, and a second strap was located on the proximal phalanx. Although these two straps allowed the device to be quickly donned and doffed,

Figure 5: Using two adjustable straps to secure the fingertip gives us independent control of the fingertip’s left/right position and the strap’s tightness. Fingers of different sizes and curvatures can be positioned such that the fingertip is in a consistent location.

the fingertips of different subjects were held in different positions, depending on the curvature and size of the individual finger. We previously found that the size of the user’s fingertip had a significant effect on their perception of tactile cues [50], which suggests that fingertip position would also have a significant effect on perception with this device. Our new design replaces the distal strap with two narrow straps, each of which can be shortened or lengthened via two independent screw mechanisms. This setup beneficially allows independent adjustment of the overall tightness and the lateral position of the fingertip, as shown in Fig. 5.

3.4 Sensing and Communication

We use an NDI 3D Guidance trakSTAR electromagnetic tracking system to measure the position and orientation of the hand-held portion of the device, and we then apply a constant transform to estimate the fingertip pose. A MATLAB program, which maintains the virtual environment parameters, uses the fingertip pose to compute the desired tactile interaction and the six associated rod lengths that are required to render this cue to the user.

We created a visual representation of our virtual environment in Unity, which is displayed to the user via an Oculus Quest VR headset. We pair the Quest with a Bluetooth ESP32 module, allowing fingertip pose to be transmitted to the Quest such that a virtual hand is visually presented at the same position and orientation as the user’s actual hand. This wireless communication channel also allows MATLAB to receive user inputs from the Oculus Touch controller.

As in [50], the Fuppeteer is tethered for power and communication purposes. The tethered setup yields a range of motion for the user’s fingertip of over 1 m in any direction, which is suitable for our purposes of examining tactile dimensionality in a lab setting. Future work related to portable power, computation, and sensing would be required to adapt the Fuppeteer to a consumer product with no tethering requirements.

4 HUMAN-SUBJECT STUDY

We designed a human-subject study in which a user wears the haptic device to perform two perceptual tasks: a hidden surface task related to shape discrimination and a slider task related to mass discrimination. The user completes several trials of each task while the device provides tactile feedback to the right index finger using one of four rendering methods. Following a within-subjects

Figure 6: Example tactile cues that can be presented by our 1-, 3-, and 6-DOF rendering methods. Our 1-DOF method can provide contact and normal force at one fingertip location, the 3-DOF method can render contact and normal force at any fingertip location, and our 6-DOF method can provide contact, normal force, and shear force at any fingertip location.

design, each user completed one block of trials for each of the four rendering methods.

4.1 Haptic Rendering Methods

Three rendering methods simulate tactile feedback that would be provided by an n -DOF fingertip device, where $n \in \{1, 3, 6\}$, and one rendering method allows the user to interact with real objects. Our specific definitions for the three n -DOF haptic rendering methods are depicted in Fig. 6 and shown in the supplementary video. Our ultimate goal is to compare the subjects’ performance and motion when each of these rendering methods is implemented, to better motivate what types of fingertip devices are best suited to represent diverse haptic interactions.

Our 1-DOF controller simulates a device that moves a flat surface with fixed orientation to make contact with a single location on the fingertip and then apply variable normal force, similar to the VRTouch by Go Touch VR [44]. In this study, the 1-DOF contact location is defined by rotating 30 degrees from the volar aspect, or bottom, of the fingertip toward the distal tip. In general, we expect this 1-DOF rendering method to provide a natural representation of contact and penetration depth, especially when the pose of the user’s hand would cause virtual contact to occur near the selected tactile feedback point.

Our 3-DOF controller simulates a device that rotates a flat surface around the finger center and then translates the surface into the finger to apply variable normal force. This 3-DOF rendering method can similarly be used to display contact and penetration depth, but at different orientations and locations on the fingertip. Given the sensitivity of the fingerpad, the contact location is relatively easy to perceive. We expect this method to allow better understanding of a surface’s local slope compared to the 1-DOF method.

Our 6-DOF controller builds upon the 3-DOF controller by including three-dimensional shear cues. Once contact is made with the fingertip, the flat surface can then translate up to 4 mm in any direction tangent to the fingertip surface, and it can rotate up to 20 degrees in either direction about the vector normal to the fingertip surface. The friction between the skin and the platform causes

translational and rotation traction that are difficult to see by eye but are relatively easy to feel.

For the trials with real-world haptics, we remove the end-effector of the device such that no contact is made. We then have the participant interact with a physical model of the virtual environment to complete the trial. Whereas the 1-, 3- and 6-DOF rendering methods provide only tactile feedback, the real rendering method also provides kinesthetic feedback (grounded forces and torques) because the finger contacts real objects.

4.2 Hidden Surface Task

In the hidden surface task, participants explore and try to recognize smoothly curving shapes defined by the equation

$$z = ax^2 + by^2, \quad (1)$$

where the positive z -direction is up and a and b are assigned different values for each of six predetermined surfaces. In each trial, the subject explores a randomly selected surface while receiving tactile feedback from one of the four rendering methods. Images of all six possible surfaces, as shown in Fig. 7, are visible on the back wall of the virtual environment. Once the participant believes they know which surface is present, they use an Oculus controller in their left hand to select the corresponding image.

4.2.1 Hidden Surface Motivation. These shapes were inspired by Kappers et al. [18], who asked participants to identify a wide range of three-dimensional quadric surfaces. Kirkpatrick and Douglas [20] used five of these surfaces to evaluate point force devices, and Kuchenbecker et al. [22] used a similar set of five surfaces to evaluate how a passive tactile display device affected three-dimensional shape recognition.

Our six surfaces and their corresponding a and b values are shown in Fig. 7. Using the definition of curvedness in [18], our five non-flat surfaces have a curvedness of 4 m^{-1} , which is within the range of curvedness for their presented shapes. Each of the quadric surfaces spans $x \in [-200 \text{ mm}, 200 \text{ mm}]$ and $y \in [-150 \text{ mm}, 150 \text{ mm}]$, making them similar in size to the 200 mm diameter surfaces presented in [18]. Furthermore, it was shown in [20] that size only accounted for modest changes in the time needed to recognize the surface.

4.2.2 Hidden Surface Haptics. For our 1-DOF, 3-DOF, and 6-DOF haptic rendering methods, we first compute the distance between the virtual fingertip center and the nearest point on the virtual surface, d_V . We define d_H as the haptic distance from the real fingertip center to the device's platform, and we set the pose of the end-effector such that d_H equals d_V . The value of d_H is restricted to a range of 4 mm to 12 mm, which allowed contact to be made and broken for all fingertip sizes observed in this study. For the 3-DOF and 6-DOF rendering methods, we also assign the orientation of the end-effector such that the flat contact surface is tangent to the virtual surface at the nearest point.

Our 6-DOF rendering method includes three-dimensional shear cues. When the finger makes contact with and moves tangent to the surface, the end-effector moves up to 4 mm in the opposite direction so to render translational shear. When the finger rotates about the surface's normal vector, the end-effector rotates up to 20 degrees about the same vector in the opposite direction, such that rotational

Figure 7: The six quadric surfaces presented in the hidden surface task and the corresponding coefficients that are substituted into Eq. (1). The shorthand names for each surface are given in parenthesis and represent the curvatures in the front/back and left/right directions. Identical shapes were 3D-printed for the real-world haptic rendering method. All virtual and real surfaces are positioned such that the point at $(x, y, z) = (0 \text{ mm}, 0 \text{ mm}, 0 \text{ mm})$ is located at the same position.

shear is presented. We reduce end-effector translational shear by a factor of five relative to finger motion, in order to reduce the tactile sensation noise related to finger drift, which can be relatively large when a user holds his or her finger in free space. We also have the end-effector pose creep toward the current finger pose through a discrete 1 Hz low-pass filter so that we can continue to present shear during longer motions.

For our real-world rendering method, we remove the end-effector and place a 3D-printed surface in the same position as the virtual surface, allowing the participant to receive natural haptic feedback as the environment is explored. In these tasks, the participant continues to wear the headset so that the physical surface is not visible. An additional face shield was used to prevent viewing of the surface or the device through gaps around the headset.

4.2.3 Hidden Surface Hypotheses. We expect participants to implement different exploration techniques depending on the haptic rendering method. The 1-DOF controller provides no slope information; in this condition we thus expect participants to determine the shape of the surface by making and breaking contact multiple times, comparing the height of the surface at each location. With the 3-DOF and 6-DOF controllers, however, we expect the participant to immediately sense the slope of the surface during each contact, leading participants to move more horizontally along the surface. We also expect participants to complete the task more quickly and accurately with the 3-DOF and 6-DOF controllers than with the 1-DOF controller, as prior literature suggests that haptic curvature perception is dominated by local surface orientation [33, 47].

4.3 Slider Task

In the slider task, participants must determine whether the mass of a disk is less than or greater than the mass of a reference disk, which is provided for the subject to interact with at their own discretion. The light, reference, and heavy sliders have masses of 13 g, 18 g, and 23 g respectively, which pilot studies suggested are difficult, but possible, to discern. Each disk has a diameter of 50 mm

and a height of 10 mm. Subjects are not allowed to pick up the disks, but rather they must estimate the mass by pressing their index finger into the disk and sliding it on the table anywhere in the workspace. To give a response, the participant must slide the disk to a region marked ‘Light’, or a region marked ‘Heavy’, both of which are approximately 300 mm from the starting disk location.

4.3.1 Slider Motivation. Many previous studies have explored how effectively wearable fingertip devices can render the mass of an object [4, 11, 27, 28, 37, 39]. In most of these studies, participants pinched virtual objects between their thumb and index finger while haptic feedback was provided to both fingers. We present a similar scenario in which participants slide a disk on a flat table, such that we must provide haptic feedback to only the index finger. Although this scenario does not present the weight of a held object, we still render the inertial forces associated with accelerating the object. Exploratory motions with acceleration peaks were required in [11] and were observed without prompting in [39] for conditions in which only skin deformation was rendered.

4.3.2 Slider Haptics. For our 1-DOF, 3-DOF, and 6-DOF haptic rendering methods, the virtual disk slides on the table only when the virtual fingertip presses into the top of the disk by a required penetration depth, p . This condition follows from the assumption that a real disk moves only when static friction between the finger and the disk is greater than static friction between the disk and the table, as determined by the mass, finger force, and coefficients of friction. For the three real disks, the required penetration depths are within 1 mm of each other. For the virtual disks, we scale p to account for measurement error and finger drift associated with trying to keep a finger stationary in free space. Moving our virtual disks requires a penetration depth, p , of 2 mm for the light disk, 4 mm for the reference disk, and 6 mm for the heavy disk. Once the required penetration depth is met, the virtual slider moves with the finger.

Similar to the hidden surface task, we define d_V as the distance between the top surface of the slider and the virtual fingertip center, and we define d_H as the distance between the haptic contact surface and the real fingertip center. Fig. 8 shows the relationship between d_H and d_V for the slider task. This piecewise mapping ensures that no haptic feedback is given before the virtual finger touches the disk and that every participant receives the same 2 mm range of normal displacement between first contact and maximum required virtual penetration depth of 6 mm. For every rendering method, the participant has to press farther into heavier disks and feel larger normal forces in order to move the disk.

Our 6-DOF controller also provides feedback related to the inertia of the disk. We pass the disk’s acceleration through a discrete 5 Hz low-pass filter, to reduce noise associated with our finger pose measurements, and we multiply by the disk’s mass to estimate applied force. For the trials with real disks, participants receive some amount of kinesthetic force feedback in addition to tactile skin stretch feedback. Suchoski et al. reported that the Weber Fraction, or the just noticeable difference in proportion to the reference value, of virtual mass was approximately three times greater for skin deformation feedback than for kinesthetic force feedback [39].

Figure 8: In the slider task, the relationship between virtual distance d_V and haptic distance d_H is given by a piecewise function with four segments. a) When the virtual finger is far from the disk, no haptic feedback is provided. b) When the virtual finger presses into the disk, but a mass-dependent penetration threshold is not met, haptic normal displacement proportional to penetration is provided and the disk remains static. c) When the penetration threshold is exceeded, haptic normal displacement remains proportional to penetration and the disk moves with the finger. d) Once the virtual fingertip center penetrates the disk, maximum haptic normal displacement is rendered.

To account for this ratio, we scaled the difference between the virtual masses and the reference mass by a factor of three. After implementing this scaling in a pilot study, we found that a shear displacement to force ratio of 40 mm/N resulted in the task being difficult, but possible. This ratio of skin stretch to force is much higher than was reported in [39] and [37], but as was discussed in Sec. 2.4, the Fuppeteer’s end-effector is not perfectly rigid. Our effective stiffness accounts for fingertip stiffness and end-effector stiffness in series, resulting in a lower effective stiffness (0.025 N/mm) than for the fingertip shear stiffness alone, which Nakazawa et al. report as approximately 0.2 N/mm when contact force is 1 N [29].

For our natural haptic feedback controller, we provide physical disks for the participant to interact with. Similar to the hidden surface task, we have the subject wear the device, but we remove the end-effector. It is difficult to align a virtual slider with a real slider well enough that it can be effectively manipulated, so we do not have the participant wear the VR headset. We 3D-printed disks with different amounts of infill, such that they are externally identical but have the specified masses of 13 g, 18 g, and 23 g for the light, reference, and heavy disks respectively. We also attached a layer of 0.3 mm Teflon to the base of each disk to reduce friction, since our virtual representation of this environment did not include friction.

4.3.3 Slider Hypotheses. For both the 1-DOF and 3-DOF rendering methods, we expect participants to be able to distinguish slider mass only from the penetration threshold, as these conditions provide no skin stretch feedback. Users can compare the different penetration distances by making slow and careful motions, pressing into the top surface of the slider while trying to move it laterally. We do not expect the slope information provided by the 3-DOF

controller to substantially benefit the user in this task. In contrast, the 6-DOF controller provides tactile shear that is proportional to the disk's mass, which we anticipate will help with the task. Accordingly, we expect users to move the sliders horizontally with large accelerations to emphasize this cue.

4.4 Procedure

Each subject first gave informed consent and filled out an initial survey. Twenty people participated in the study, which was approved by the Institutional Review Board of the University of Pennsylvania under protocol number 843647. Participation took between 60 and 90 minutes; individuals not employed by the Max Planck Society were paid 8 euros per hour for their time. Our participant pool included thirteen females and seven males, with ages between 25 and 56 (mean of 32, standard deviation of 8.6). Their experience with haptic devices included 'None' (3 users), 'Limited' (8), and 'Moderate' (9), and their experience with VR included 'None' (2), 'Limited' (9), 'Moderate' (8), and 'Extensive' (1).

The device was worn on the subject's right index finger across all trials. A visual representation of the subject's hand was depicted in the virtual environment and followed the user's real motion. Both tasks took place in a virtual box with a width of 400 mm, depth of 300 mm, and height of 300 mm. We did not physically restrict the user's motion, but if the user moved beyond the boundaries, the penetrated wall turned red.

The study was broken into four blocks, each associated with one of the four rendering methods. Block order was assigned by a balanced Latin square for every four participants. We are interested in observing how participants naturally use these rendering methods to try to complete the tasks, rather than training participants to employ new exploratory techniques in order to succeed in each specific task. As such, participants were told that each rendering method was unique, but they were not given specific details as to how the methods differed.

Within each block, the participant first donned the Fingertip Puppeteer on their right index finger, and the experiment administrator adjusted the straps such that the fingertip was securely held in the proper location. The subject then put on the VR headset and began the hidden surface task, consisting of two training trials, in which the surface was shown to the participant, and four test trials, in which the surface was invisible. After the hidden surface task, subjects completed the slider task, which again consisted of two training trials and four test trials. Surface types and slider masses were randomized, although the two training types were not allowed to be the same. Sample trials are shown in Fig. 9. Before each trial, the device was automatically calibrated using a trakSTAR sensor embedded in the end-effector, as is described in [50].

We recorded the duration, accuracy, and motion trajectory associated with each trial. Following completion of the slider task, the participant removed the headset and the fingertip device. They then completed a survey specific to the most recent rendering method, with five-point scales for the realism of each task and the overall comfort of the rendering method. Two free-response questions allowed participants to note any aspects of the rendering method that worked well or poorly. After all four blocks had been completed,

subjects were able to provide additional comments in a closing free-response question.

5 RESULTS

We focused our analysis of the study results on performance metrics, motion metrics, and qualitative feedback.

5.1 Performance Metrics

Success rates and completion times for each rendering method are shown in Table 1. We computed the accuracy for each of the twenty participants within each task and rendering method as an interval variable, and we performed one-tailed one-sample t-tests to compare the responses to chance, with Bonferroni correction. All four rendering methods significantly outperformed chance (16.7%) in the hidden surface task, whereas only the 6-DOF and real rendering methods significantly outperformed chance (50%) in the slider task, with $p < 0.05$.

We also compiled the participants' responses into confusion matrices for each rendering method during the hidden surface task (Fig. 10) and the slider task (Fig. 11). Perfect identification would yield a diagonal confusion matrix; the diagonal is almost perfect for real rendering on the surface identification task but not for any of the other confusion matrices. We note that 1-DOF tactile rendering resulted in 14 surfaces incorrectly labeled as the flat surface (20.9% of non-flat surfaces), whereas the 6-DOF method yielded the next highest number of flat surface false-positives at three (4.8% of non-flat surfaces).

For further analysis, we conducted a two-way ANOVA to examine the effect of task and rendering method on success rate. We found a statistically significant interaction between task and rendering method ($F = 4.49, p < 0.01$). We then conducted a one-way ANOVA within each task and found a significant effect of rendering method on success rate for both the hidden surface task ($F = 14.55, p < 0.001$) and the slider task ($F = 2.78, p = 0.047$). In the hidden surface task, Tukey's post-hoc analysis yielded significant differences between the real rendering method and each of the other three methods ($p < 0.001$), but no statistically significant differences were found between the 1- and 3-DOF methods ($p = 0.48$), the 1- and 6-DOF methods ($p = 0.89$) or the 3- and 6-DOF methods ($p = 0.89$).

In the slider task, Tukey's post-hoc analysis yielded no statistically significant differences between rendering methods. The 1- and 3-DOF ($p = 0.96$), and the 6-DOF and real ($p = 0.99$) were the pairings least likely to be significantly different, followed by the 1- and 6-DOF ($p = 0.39$), 1-DOF and real ($p = 0.22$), 3- and 6-DOF ($p = 0.17$), and 3-DOF and real ($p = 0.09$).

The mean ratings from qualitative surveys administered after the completion of each rendering method are provided in Table 2. Using the Kruskal-Wallis test, we found that hidden surface task realism, slider task realism, and overall comfort were each significantly affected by rendering method ($p < 0.01$). Post-hoc analysis showed that the real rendering method yielded significantly higher hidden surface realism ratings than each of the other three rendering methods, and significantly higher slider realism and overall comfort ratings than the 1-DOF rendering method ($p < 0.05$). Other comparisons were not significant.

Figure 9: Side-by-side images of the VR and real-world setups during the human-subject study. The hidden surface task is depicted in (a) and (b), and the slider task is shown in (c) and (d). Subfigures (a) and (c) both show training scenarios in VR and physical setups for the 1-DOF, 3-DOF, or 6-DOF tactile rendering methods, whereas (b) and (d) show non-training trials and physical setups for real-world rendering. Note that the real-world haptic rendering of the slider task in (d) has no VR environment.

		Reported Surface						
Rendered Surface		00	+0	-0	++	+-	--	
	00	0.77	0.15	0	0.08	0	0	1-DOF
	+0	0.41	0.29	0	0.18	0.06	0.06	3-DOF
	-0	0.31	0	0.38	0.08	0	0.23	6-DOF
	++	0.15	0.15	0.08	0.62	0	0	Real
	+-	0	0.18	0	0	0.45	0.36	
	--	0.08	0	0.08	0	0.15	0.69	

Figure 10: Confusion matrices of the responses given in the hidden surface task for each rendering method. Each confusion matrix represents 80 trials, but because trial types were randomized, the number of trials within each row are not equal.

5.2 Exploratory Motions

In general, distinct users implemented a wide variety of exploratory techniques during the human-subject study. In the hidden surface task, some participants implemented sweeping horizontal motions (Fig. 12(a)), whereas other participants instead favored vertical poking motions (Fig. 12(b)). Similarly, in the slider task some subjects vigorously shook the slider side-to-side (Fig. 12(c)), whereas others made more precise, slow motions (Fig. 12(d)).

We measured the total horizontal and vertical distances travelled by the user's finger, which we normalized with respect to trial duration, resulting in average horizontal and vertical speeds for each

trial. We also computed the horizontal and vertical accelerations by passing the cubic interpolation of the position data through a 5 Hz low-pass filter and taking the second derivative. The horizontal and vertical speeds and accelerations across all trials are shown in Fig. 13. Finally, to highlight the variation between participants, we show the vertical speeds and horizontal accelerations for each participant during the hidden surface and slider tasks, respectively, in Fig. 14.

We performed a one-way repeated measures ANOVA with Tukey post-hoc analysis on the motion measurements within each task.

Rendering Method	Surface		Slider	
	Accuracy	Time (s)	Accuracy	Time (s)
1-DOF	52.5% $\pm$ 5.4 [†]	51.8 $\pm$ 0.5	51.3% $\pm$ 6.1	26.0 $\pm$ 0.2
3-DOF	63.8% $\pm$ 7.6 [†]	48.3 $\pm$ 0.4	47.5% $\pm$ 5.4	28.8 $\pm$ 0.2
6-DOF	60.0% $\pm$ 5.3 [†]	48.9 $\pm$ 0.4	63.8% $\pm$ 5.0 [†]	30.8 $\pm$ 0.3
Real	98.8% $\pm$ 1.3 [†]	25.0 $\pm$ 0.2	66.3% $\pm$ 5.6 [†]	22.2 $\pm$ 0.2

Table 1: Overall success rates and completion times (mean $\pm$ S.E.) for each combination of rendering method and task. A total of 80 trials was completed for each combination. A [†] indicates performance significantly better than chance ($p < 0.05$), as computed via a one-tailed one-sample t-test with Bonferroni correction.

		Reported Slider			
		L	H	L	H
Rendered Slider	L	0.57	0.43	0.53	0.47
	H	0.53	0.47	0.38	0.63
1-DOF					
3-DOF					
6-DOF					
Real					

Figure 11: Confusion matrices of the responses given in the slider task for each rendering method. Each confusion matrix represents 80 trials, but because trial types were randomized, the number of trials within each row are not equal.

Rendering Method	Surface Realism	Slider Realism	Overall Comfort
1-DOF	2.95 $\pm$ 0.06	2.95 $\pm$ 0.06	3.00 $\pm$ 0.04
3-DOF	3.95 $\pm$ 0.04	3.30 $\pm$ 0.05	3.65 $\pm$ 0.03
6-DOF	3.70 $\pm$ 0.05	3.35 $\pm$ 0.05	3.50 $\pm$ 0.03
Real	4.85 $\pm$ 0.02	4.20 $\pm$ 0.05	3.95 $\pm$ 0.03

Table 2: Mean subjective ratings ($\pm$ S.E.) for each rendering method. Realism for each task was rated on a numeric scale from very unrealistic (1) to very realistic (5), and overall comfort was rated from very uncomfortable (1) to very comfortable (5).

In the hidden surface task, we found a significant effect of rendering method on horizontal speeds ($F = 11.21, p < 0.001$); the real rendering method prompted larger horizontal speeds than each of the other three methods ($p < 0.01$). In the slider task, we found a significant effect of rendering method on vertical speeds ($F = 11.35, p < 0.001$) and that the real rendering method resulted in smaller vertical speeds than each of the other methods ($p < 0.01$). No other individual comparisons were statistically significant.

We also looked at the relationship between exploratory motions and performance, but we found no significant differences between the speeds or accelerations implemented during successful trials versus unsuccessful trials. Across all trials for both tasks, we measured overall speeds (mean $\pm$ standard error) of 0.109 ± 0.003 m/s during successful trials and 0.110 ± 0.003 m/s during unsuccessful trials. Looking more closely at slider task trials with 6-DOF tactile

feedback, during which horizontal acceleration directly affects the amount of tactile shear rendered to the user, we measured horizontal accelerations of 0.798 ± 0.073 m/s² during successful trials and 0.858 ± 0.094 m/s² during unsuccessful trials.

5.3 Qualitative Feedback

Answers to the free-response questions largely upheld the quantitative results. Seventeen comments specified that the hidden surface task was well-rendered by the 3-DOF and 6-DOF tactile methods, with five comments directly mentioning the ability to feel the ‘slope’ or ‘curvature’ of the surfaces. Two positive comments for the 6-DOF method noted the ‘continuous’ feel of the surfaces (P2, P5), whereas three negative comments for the 1-DOF method noted the absence of a ‘continuous’ feeling (P2, P5, P6). For the 6-DOF rendering, five participants stated that the slider task felt somewhat easier, with comments such as “*I had the feeling I was now able to better determine... what disc was heavier, but only slightly*” (P12). Overall, participants consistently noted that they struggled during the slider task, with every rendering method (including real haptic feedback) receiving from four to eight comments related to slider task difficulty.

Free-response answers also provide additional insight into components of the study that were not directly measured. Two subjects disliked the virtual finger being able to penetrate the virtual surface: “*I didn’t like that the virtual hand was able to go through the virtual hidden surface...*” (P1) and “*the finger crossed the surface visually*” (P5). A third user noted a shortcoming of the system was that there was “*no modeling of infinite resistance like for real surfaces*” (P6).

Despite our efforts to minimize weight supported by the user, some participants still commented on the device’s mass. Two participants described discomfort caused by weight: “*The device is a bit heavy for me*” (P4) and “*the weight was a bit too high*” (P5). Three participants noted that the weight of the device may have affected performance (P2, P6, P19), and one attributed blame to the counterweight itself: “*I think that the floating counterweight made it harder for me to distinguish how hard I was pushing on the slider*” (P17).

While administering the study, we observed multiple subjects who spent more time exploring the edges of the hidden surfaces when interacting with the real physical surfaces (real haptic rendering) than when the surfaces were virtually rendered by the 1-, 3-, or 6-DOF methods. A few participants lamented that edges were not rendered for virtual surfaces: “*The edges of the surfaces were not very clear.*” (P18), and “*I couldn’t feel the edges of the surfaces for the hidden surface task*” (P13).

Although these critical comments may be more interesting to the reader and better motivate design decisions for future tactile fingertip feedback systems, we also received many written and verbal comments expressing an overall pleasant experience with the tested system, some of which we highlight here: “*Good resolution of force.*” (P6, 1-DOF), “*I really enjoyed the hidden surface task*” (P4, 6-DOF), and “*I really like the detail where the device swivels around the tip of the finger!*” (P17, closing).

Figure 12: Example trajectories observed during the two tasks. On the left, we show examples of (a) horizontal exploration and (b) poking of surface three with the 3-DOF method. On the right, we show examples of (c) more vigorous shaking and (d) precise, gentle pushing of the light slider with the 6-DOF method. Each depicted trajectory comes from a distinct user.

Figure 13: Horizontal and vertical speeds and accelerations during the two tasks, organized by rendering method.

Figure 14: Vertical speeds for each participant during the hidden surface task (left) and horizontal accelerations for each participant during the slider task (right) across all trials.

6 DISCUSSION

In the hidden surface task, we recorded higher mean success rates, completion times, and realism ratings for the 3- and 6-DOF rendering methods than for the 1-DOF method, which fits with prior

literature that suggests haptic curvature perception is dominated by local surface orientation [33, 47]. Non-flat surfaces were also incorrectly labeled flat over four times as frequently with the 1-DOF method than the other methods. However, all four of the rendering

methods yielded surface identification rates better than chance, and none of the pairwise differences between success rates of the three tactile-only rendering methods were statistically significant. These findings suggest that there is minimal benefit from providing higher-dimensional tactile feedback in the absence of kinesthetic force feedback for this particular task. The large extent of the surfaces may have enabled reasonably accurate recognition simply through recognition of the spatial regions that cause fingertip contact, the cue that is provided even by the simplest 1-DOF rendering scheme. The real rendering method significantly outperformed the three tactile-only rendering methods, suggesting that kinesthetic force feedback, which was presented in each of the past works cited as motivation [18, 20, 22, 33, 47], is superior to tactile feedback alone for this shape discrimination task.

In the slider task, success rates for the 1-DOF and 3-DOF rendering methods were approximately equal to chance (50%), suggesting that participants did not notice any difference between the required penetration depths. The 6-DOF and real rendering methods resulted in success rates greater than chance, showing that the relationship between shear skin stretch and inertia was more easily perceived by users. The similarity of the success rates for the 6-DOF and real rendering methods in the slider task also suggests that our shear representation was reasonably realistic, despite the incomplete modelling of end-effector stiffness as a function of configuration.

Interestingly, exploration techniques were not heavily influenced by rendering method but instead seem to depend more strongly on individual preference. For example, P13 consistently applied relatively large horizontal accelerations to the slider, even when no tactile representation of inertia was provided, whereas P11 consistently applied small horizontal accelerations to the slider, even when skin stretch proportional to acceleration was rendered. Although participants were only able to perform mass discrimination at rates above chance when inertia was presented through tactile shear, the presence of this shear did not elicit larger accelerations or any other change in exploratory motion. Naive users seemingly expected the tactile fingertip feedback to accommodate their exploratory tendencies, and as such, subjects did not adjust their motions to match the provided tactile sensations.

Success rates were also independent of exploratory motion, as users implemented nearly identical overall fingertip speeds during successful and unsuccessful trials. If system latency hindered the haptic rendering, we would expect performance to degrade when users move more quickly. Instead, performance is consistent regardless of user speed, suggesting participants are able to intuitively understand and compensate for any delays between virtual interaction and tactile feedback. During the 6-DOF rendering of the slider task, when tactile shear is directly related to user acceleration, we still do not find a significant effect of acceleration on performance. It is possible that during these interactions in which we expect large accelerations to yield better results, latency has a more negative impact and thus offsets the expected benefits of high-acceleration exploratory techniques.

Even when real tactile and kinesthetic feedback was provided, we measured significant differences in exploratory motion only with respect to horizontal speeds during the hidden surface task and vertical speeds during the slider task. The difference in vertical speeds during the slider task is expected, since the finger is unable

to penetrate the real disks, which remained at a constant height. A similar rationale may explain the increase in horizontal speed during the hidden surface task, as participants may focus on horizontal exploration when kinesthetic feedback limits vertical exploration. However, the increase in horizontal speed with real feedback may also be attributed to greater comfort or confidence while exploring the surface, or less concern regarding the bandwidth of the system. Multiple users commented that the weight of the device influenced their perception, suggesting that higher velocities and accelerations might have been measured if the device was lighter.

We also note that some participants poked downward on the slider rather than moving it, suggesting either confusion between the physical properties of mass and stiffness, or an expectation for virtual mass to be haptically rendered as virtual stiffness. Interestingly, two participants poked the real sliders at least once. The perceptual confusion between mass and stiffness provides additional insight to the findings of Schorr et al. [37], who previously reported that a change in virtual mass was sometimes perceived as a change in stiffness or friction.

Answers to the free-response questions provide deeper understanding of the strengths and shortcomings of the tactile rendering methods and our overall system. Free-responses that commented on discomfort caused by the virtual finger passing through the virtual surface imply a desire for, at a minimum, visual kinesthetic feedback, which is also often called pseudo-haptics [15, 23]. One comment noting the 'lack of infinite resistance' reaffirmed the request for kinesthetic force feedback. Other featured comments suggested potential improvements related to reducing the mass of the device or enabling the system to render edges as well as faces.

7 LIMITATIONS AND FUTURE WORK

This study was limited to four rendering methods and two tasks. We are interested in expanding our set of rendering methods, such as providing shear skin stretch at a single location on the fingertip, and continuing to explore how different dimensions of fingertip tactile rendering affect performance in additional tasks.

Importantly, the study instructions provided no information as to how the rendering methods differed, or how each rendering method could be used to complete the tasks. This approach allowed us to observe naive performance, but it also proved challenging for some participants who never discovered exploration techniques that yielded positive results. If the tactile rendering methods were explained, or if suggested exploratory techniques were provided, users would most likely be better able to discern which techniques work best for a particular task and rendering method, rather than consistently trying the same technique. Alternatively, more exposure time and feedback about answer correctness would most likely have enabled users to discover better techniques on their own but would have increased study duration and fatigue.

Additionally, we want to explore how the relative scaling of different tactile cues affects exploratory procedures. For example, our required penetration depths did not enable participants to discern lighter and heavier disks at rates greater than chance, but we expect that penetration depth would allow efficient discrimination at some scaling value. Furthermore, if the disks could be easily identified

from penetration depth, users who are naturally inclined to vigorously shake the disk may instead implement slower, more precise vertical motions. It would be interesting to explore what type and scaling of feedback are required to convince users to adjust their natural exploratory techniques, and whether this adjustment can consequently improve performance.

8 CONCLUSION

We integrated a recently invented 6-DOF fingertip tactile device into an immersive VR system with haptic feedback and programmed it to render feedback as would be provided by devices with varying degrees of freedom. We presented rendering algorithms for 1-, 3-, and 6-DOF tactile cues during a shape discrimination and a mass discrimination task. We found that only the 6-DOF tactile feedback allowed both tasks to be completed better than chance, but that the dimensionality of tactile feedback does not strongly affect the exploratory techniques implemented by the user. We hope our work will encourage future tactile device designs that can present a wide variety of virtual interactions and inspire additional research into how tactile dimensionality and individual preference influence the utility of wearable haptic devices.

ACKNOWLEDGMENTS

We thank Bernard Javot for guidance in manufacturing various components and Seyhan Sitti for assistance in 3D-printing surfaces and components. This research is supported in part by the US National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1321851.

REFERENCES

- [1] Manfredo Atzori and Henning Müller. 2015. The Ninapro Database: a Resource for sEMG Naturally Controlled Robotic Hand Prosthetics. In *Proceedings of the 2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (Milan, Italy) (EMBC '15)*. IEEE, 7151–7154. <https://doi.org/10.1109/EMBC.2015.7320041>
- [2] Hrvoje Benko, Christian Holz, Mike Sinclair, and Eyal Ofek. 2016. NormalTouch and TextureTouch: High-Fidelity 3D Haptic Shape Rendering on Handheld Virtual Reality Controllers. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (Tokyo, Japan) (UIST '16)*. Association for Computing Machinery, New York, NY, USA, 717–728. <https://doi.org/10.1145/2984511.2984526>
- [3] Francesco Chinello, Monica Malvezzi, Domenico Prattichizzo, and Claudio Pacchierotti. 2020. A Modular Wearable Finger Interface for Cutaneous and Kinesthetic Interaction: Control and Evaluation. *IEEE Transactions on Industrial Electronics* 67, 1 (Feb. 2020), 706–716. <https://doi.org/10.1109/TIE.2019.2899551>
- [4] 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 (Québec City, QC, Canada) (UIST '17)*. Association for Computing Machinery, New York, NY, USA, 119–130. <https://doi.org/10.1145/3126594.3126599>
- [5] 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 (Montreal QC, Canada) (CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3173574.3174228>
- [6] Seungmoon Choi and Katherine J. Kuchenbecker. 2013. Vibrotactile Display: Perception, Technology, and Applications. *Proc. IEEE* 101, 9 (Sept. 2013), 2093–2104. <https://doi.org/10.1109/JPROC.2012.2221071>
- [7] Heather Culbertson, Julie M. Walker, Michael Raitor, and Allison M. Okamura. 2017. WAVES: A Wearable Asymmetric Vibration Excitation System for Presenting Three-Dimensional Translation and Rotation Cues. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (Denver, Colorado, USA) (CHI '17)*. Association for Computing Machinery, New York, NY, USA, 4972–4982. <https://doi.org/10.1145/3025453.3025741>
- [8] Thomas Feix, Javier Romero, Heinz-Bodo Schmiedmayer, Aaron M. Dollar, and Danica Kragic. 2016. The GRASP Taxonomy of Human Grasp Types. *IEEE Transactions on Human-Machine Systems* 46, 1 (Feb. 2016), 66–77. <https://doi.org/10.1109/THMS.2015.2470657>
- [9] Antonio Frisoli, Massimiliano Solazzi, Fabio Salsedo, and Massimo Bergamasco. 2008. A Fingertip Haptic Display for Improving Curvature Discrimination. *Presence: Teleoperators Virtual Environments* 17, 6 (Dec. 2008), 550–561. <https://doi.org/10.1162/pres.17.6.550>
- [10] Massimiliano Gabardi, Massimiliano Solazzi, Daniele Leonardis, and Antonio Frisoli. 2016. A New Wearable Fingertip Haptic Interface for the Rendering of Virtual Shapes and Surface Features. In *Proceedings of the 2016 IEEE Haptics Symposium (Philadelphia, USA) (HAPTICS '16)*. IEEE, 140–146. <https://doi.org/10.1109/HAPTICS.2016.7463168>
- [11] Adrien Girard, Maud Marchal, Florian Gosselin, Anthony Chabrier, François Louveau, and Anatole Lécuyer. 2016. HapTip: Displaying Haptic Shear Forces at the Fingertips for Multi-Finger Interaction in Virtual Environments. *Frontiers in Robotics and AI* 3, 6 (April 2016), 1–15. <https://doi.org/10.3389/fct.2016.00006>
- [12] Brian T. Gleeson, Scott K. Horschel, and William R. Provancher. 2010. Design of a Fingertip-Mounted Tactile Display with Tangential Skin Displacement Feedback. *IEEE Transactions on Haptics* 3, 4 (Nov. 2010), 297–307. <https://doi.org/10.1109/TOH.2010.8>
- [13] Xiaochi Gu, Yifei Zhang, Weize Sun, Yuanzhe Bian, Dao Zhou, and Per Ola Kristensson. 2016. Dexmo: An Inexpensive and Lightweight Mechanical Exoskeleton for Motion Capture and Force Feedback in VR. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (San Jose, California, USA) (CHI '16)*. Association for Computing Machinery, New York, NY, USA, 1991–1995. <https://doi.org/10.1145/2858036.2858487>
- [14] Ashley L. Guinan, Nicholas C. Hornbaker, Markus N. Montandon, Andrew J. Doxon, and William R. Provancher. 2013. Back-to-back Skin Stretch Feedback for Communicating Five Degree-of-Freedom Direction Cues. In *Proceedings of the 2013 World Haptics Conference (Daejeon, South Korea) (WHC '13)*. IEEE, 13–18. <https://doi.org/10.1109/WHC.2013.6548377>
- [15] Inyoung Jang and Dongjun Lee. 2014. On utilizing pseudo-haptics for cutaneous fingertip haptic device. In *Proceedings of the 2014 IEEE Haptics Symposium (HAPTICS)*. IEEE, 635–639. <https://doi.org/10.1109/HAPTICS.2014.6775529>
- [16] Roland S. Johansson and J. Randall Flanagan. 2009. Coding and Use of Tactile Signals from the Fingertips in Object Manipulation Tasks. *Nature Reviews Neuroscience* 10, 5 (April 2009), 345–359. <https://doi.org/10.1038/nrn2621>
- [17] Noriko Kamakura, Michiko Matsuo, Harumi Ishii, Fumiko Mitsuboshi, and Yoriko Miura. 1980. Patterns of Static Prehension in Normal Hands. *American Journal of Occupational Therapy* 34, 7 (July 1980), 437–445. <https://doi.org/10.5014/ajot.34.7.437>
- [18] Astrid M. L. Kappers, Jan J. Koenderink, and Inge Lichtenegger. 1994. Haptic Identification of Curved Surfaces. *Perception and Psychophysics* 56, 1 (Jan. 1994), 53–61. <https://doi.org/10.3758/BF03211690>
- [19] Hwan Kim, HyeonBeom Yi, Hyein Lee, and Woohun Lee. 2018. HapCube: A Wearable Tactile Device to Provide Tangential and Normal Pseudo-Force Feedback on a Fingertip. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (Montreal QC, Canada) (CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3173574.3174075>
- [20] Arthur E. Kirkpatrick and Sarah A. Douglas. 2002. Application-based Evaluation of Haptic Interfaces. In *Proceedings of the 10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (Orlando, USA) (HAPTICS '02)*. IEEE, 32–39. <https://doi.org/10.1109/HAPTIC.2002.998938>
- [21] Rebecca L. Koslover, Brian T. Gleeson, Joshua T. de Bever, and William R. Provancher. 2012. Mobile Navigation Using Haptic, Audio, and Visual Direction Cues with a Handheld Test Platform. *IEEE Transactions on Haptics* 5, 1 (Oct. 2012), 33–38. <https://doi.org/10.1109/TOH.2011.58>
- [22] Katherine J. Kuchenbecker, David Ferguson, Michael Kutzler, Matthew Moses, and Allison M. Okamura. 2008. The Touch Thimble: Providing Fingertip Contact Feedback During Point-Force Haptic Interaction. In *Proceedings of the 2008 Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (Reno, USA) (HAPTICS '08)*. IEEE, 239–246. <https://doi.org/10.1109/HAPTICS.2008.4479950>
- [23] Anatole Lécuyer. 2009. Simulating Haptic Feedback Using Vision: A Survey of Research and Applications of Pseudo-Haptic Feedback. *Presence: Teleoperators and Virtual Environments* 18, 1 (2009), 39–53. <https://doi.org/10.1162/pres.18.1.39>
- [24] 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 (Glasgow, Scotland UK) (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3290605.3300301>
- [25] Daniele Leonardis, Massimiliano Solazzi, Ilaria Bortone, and Antonio Frisoli. 2015. A Wearable Fingertip Haptic Device with 3 DoF Asymmetric 3-RSR Kinematics. In *Proceedings of the 2015 IEEE World Haptics Conference (Evanston, USA) (WHC '15)*. IEEE, 388–393. <https://doi.org/10.1109/WHC.2015.7177743>
- [26] Zhou MA and Pinhas Ben-Tzvi. 2015. RML Glove - An Exoskeleton Glove Mechanism with Haptics Feedback. *IEEE/ASME Transactions on Mechatronics* 20, 2 (Feb. 2015), 641–652. <https://doi.org/10.1109/TMECH.2014.2305842>

- [27] Kouta Minamizawa, Souichiro Fukamachi, Hiroyuki Kajimoto, Naoki Kawakami, and Susumu Tachi. 2007. Gravity Grabber: Wearable Haptic Display to Present Virtual Mass Sensation. In *Proceedings of the ACM SIGGRAPH 2007 Emerging Technologies* (San Diego, California) (SIGGRAPH '07). Association for Computing Machinery, New York, NY, USA, 4 pages. <https://doi.org/10.1145/1278280.1278289>
- [28] Kouta Minamizawa, Domenico Prattichizzo, and Susumu Tachi. 2010. Simplified Design of Haptic Display by Extending One-point Kinesthetic Feedback to Multipoint Tactile Feedback. In *Proceedings of the 2010 IEEE Haptics Symposium* (Waltham, USA) (HAPTICS '10). IEEE, 257–260. <https://doi.org/10.1109/HAPTIC.2010.5444646>
- [29] Nobuaki Nakazawa, Ryojun Ikeura, and Hikaru Inooka. 2000. Characteristics of Human Fingertips in the Shearing Direction. *Biological Cybernetics* 82, 3 (Feb. 2000), 207–214. <https://doi.org/10.1007/s004220050020>
- [30] Claudio Pacchierotti, Francesco Chinello, Monica Malvezzi, Leonardo Meli, and Domenico Prattichizzo. 2012. Two Finger Grasping Simulation with Cutaneous and Kinesthetic Force Feedback. In *Proceedings of the International Conference on Haptics: Perception, Devices, Mobility, and Communication* (Berlin, Heidelberg). Springer, 373–382. https://doi.org/10.1007/978-3-642-31401-8_34
- [31] Claudio Pacchierotti, Stephen Sinclair, Massimiliano Solazzi, Antonio Frisoli, Vincent Hayward, and Domenico Prattichizzo. 2017. Wearable Haptic Systems for the Fingertip and the Hand: Taxonomy, Review, and Perspectives. *IEEE Transactions on Haptics* 10, 4 (May 2017), 580–600. <https://doi.org/10.1109/TOH.2017.2689006>
- [32] Alvaro G. Perez, Daniel Lobo, Francesco Chinello, Gabriel Cirio, Monica Malvezzi, José San Martín, Domenico Prattichizzo, and Miguel A. Otaduy. 2015. Soft Finger Tactile Rendering for Wearable Haptics. In *Proceedings of the 2015 IEEE World Haptics Conference* (Evanston, USA) (WHC '15). IEEE, 327–332. <https://doi.org/10.1109/WHC.2015.7177733>
- [33] Sylvia C. Pont, Astrid M. L. Kappers, and Jan J. Koenderink. 1999. Similar Mechanisms Underlie Curvature Comparison by Static and Dynamic Touch. *Perception and Psychophysics* 61, 5 (July 1999), 874–894. <https://doi.org/10.3758/bf03206903>
- [34] Domenico Prattichizzo, Francesco Chinello, Claudio Pacchierotti, and Monica Malvezzi. 2013. Towards Wearability in Fingertip Haptics: A 3-DoF Wearable Device for Cutaneous Force Feedback. *IEEE Transactions on Haptics* 6, 4 (Oct. 2013), 506–516. <https://doi.org/10.1109/TOH.2013.53>
- [35] Domenico Prattichizzo, Claudio Pacchierotti, and Giulio Rosati. 2012. Cutaneous Force Feedback as a Sensory Subtraction Technique in Haptics. *IEEE Transactions on Haptics* 5, 4 (April 2012), 289–300. <https://doi.org/10.1109/TOH.2012.15>
- [36] William R. Provancher, Mark R. Cutkosky, Katherine J. Kuchenbecker, and Günter Niemeyer. 2005. Contact Location Display for Haptic Perception of Curvature and Object Motion. *International Journal of Robotics Research* 24, 9 (Sept. 2005), 691–702. <https://doi.org/10.1177/0278364905057121>
- [37] Samuel B. Schorr and Allison M. Okamura. 2017. Fingertip Tactile Devices for Virtual Object Manipulation and Exploration. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 3115–3119. <https://doi.org/10.1145/3025453.3025744>
- [38] Massimiliano Solazzi, Antonio Frisoli, and Massimo Bergamasco. 2010. Design of a Novel Finger Haptic Interface for Contact and Orientation Display. In *Proceedings of the 2010 IEEE Haptics Symposium* (Waltham, USA) (HAPTICS '10). IEEE, 129–132. <https://doi.org/10.1109/HAPTIC.2010.5444667>
- [39] Jacob M. Suchoski, Aaron Barron, Connie Wu, Zhan Fan Quek, Sean Keller, and Allison M. Okamura. 2016. Comparison of Kinesthetic and Skin Deformation Feedback for Mass Rendering. In *Proceedings of the 2016 IEEE International Conference on Robotics and Automation* (Stockholm, Sweden) (ICRA '16). IEEE, 4030–4035. <https://doi.org/10.1109/ICRA.2016.7487593>
- [40] Nicholas D. Sylvester and William R. Provancher. 2007. Effects of Longitudinal Skin Stretch on the Perception of Friction. In *Proceedings of the Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems* (Tsukuba, Japan) (WHC '07). 373–378. <https://doi.org/10.1109/WHC.2007.45>
- [41] Cyber Glove Systems. 2020. *CyberGrasp*. Retrieved July 21, 2020 from <http://www.cyberglovesystems.com/cybergasp>
- [42] Nikos G. Tsagarakis, T. Horne, and Darwin G. Caldwell. 2005. SLIP AESTHEASIS: A Portable 2D Slip/Skin Stretch Display for the Fingertip. In *Proceedings of the First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems* (Pisa, Italy) (WHC '05). IEEE, 214–219. <https://doi.org/10.1109/WHC.2005.117>
- [43] Dzmity Tsetserukou, Shotaro Hosokawa, and Kazuhiko Terashima. 2014. Link-Touch: A Wearable Haptic Device with Five-Bar Linkage Mechanism for Presentation of Two-DOF Force Feedback at the Fingertip. In *Proceedings of the 2014 IEEE Haptics Symposium* (Houston, USA) (HAPTICS '14). IEEE, 307–312. <https://doi.org/10.1109/HAPTICS.2014.6775473>
- [44] Go Touch VR. 2020. *VRTouch*. Retrieved July 29, 2020 from <https://www.gotouchvr.com/copy-of-technology-devices-1>
- [45] Robert J. Webster, Todd E. Murphy, Lawton N. Verner, and Allison M. Okamura. 2005. A Novel Two-Dimensional Tactile Slip Display: Design, Kinematics and Perceptual Experiments. *ACM Transactions on Applied Perception* 2, 2 (April 2005), 150–165. <https://doi.org/10.1145/1060581.1060588>
- [46] 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* (Montreal QC, Canada) (CHI '18). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3173574.3173660>
- [47] Maarten W. A. Wijntjes, Akihiro Sato, Vincent Hayward, and Astrid M. L. Kappers. 2009. Local Surface Orientation Dominates Haptic Curvature Discrimination. *IEEE Transactions on Haptics* 2, 2 (Jan. 2009), 94–102. <https://doi.org/10.1109/TOH.2009.1>
- [48] Vibol Yem and Hiroyuki Kajimoto. 2017. Wearable Tactile Device Using Mechanical and Electrical Stimulation for Fingertip Interaction with Virtual World. In *Proceedings of 2017 IEEE Virtual Reality* (Los Angeles, USA) (VR '17). IEEE, 99–104. <https://doi.org/10.1109/VR.2017.7892236>
- [49] Eric M. Young and Katherine J. Kuchenbecker. 2017. Design of a Parallel Continuum Manipulator for 6-DOF Fingertip Haptic Display. In *Proceedings of the 2017 IEEE World Haptics Conference* (Fürstentfeldbruck (Munich), Germany) (WHC '17). IEEE, 599–604. <https://doi.org/10.1109/WHC.2017.7989969>
- [50] Eric M. Young and Katherine J. Kuchenbecker. 2019. Implementation of a 6-DOF Parallel Continuum Manipulator for Delivering Fingertip Tactile Cues. *IEEE Transactions on Haptics* 12, 3 (June 2019), 295–306. <https://doi.org/10.1109/TOH.2019.2920928>