

Retargeted Self-Haptics for Increased Immersion in VR without Instrumentation

Cathy Mengying Fang
Carnegie Mellon University
mengyinf@alumni.cmu.edu

Chris Harrison
Carnegie Mellon University
chris.harrison@cs.cmu.edu




Figure 1: In this example scene, a user first places their hand down onto a palm scanner, which flips the glass cover of a number keypad. In conventional, uninstrumented hand tracking VR systems, users would not experience any haptics (A). In this work, we explore self-haptics (B), where through retargeting the user's virtual hands, we can guide the physical hands together to create physical surfaces and objects for interaction – in this case, the user's right hand physically feels a keypad surface.

ABSTRACT

Today's consumer virtual reality (VR) systems offer immersive graphics and audio, but haptic feedback is rudimentary – delivered through controllers with vibration feedback or is non-existent (i.e., the hands operating freely in the air). In this paper, we explore an alternative, highly mobile and controller-free approach to haptics, where VR applications utilize the user's own body to provide physical feedback. To achieve this, we warp (retarget) the locations of a user's hands such that one hand serves as a physical surface or prop for the other hand. For example, a hand holding a virtual nail can serve as a physical backstop for a hand that is virtually hammering, providing a sense of impact in an air-borne and uninstrumented experience. To illustrate this rich design space, we implemented twelve interactive demos across three haptic categories. We conclude with a user study from which we draw design recommendations.

CCS CONCEPTS

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

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).
UIST '21, October 10–14, 2021, Virtual Event, USA
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8635-7/21/10.
<https://doi.org/10.1145/3472749.3474810>

KEYWORDS

Virtual Reality; Interaction Techniques; Haptics; Retargeting

ACM Reference Format:

Cathy Mengying Fang and Chris Harrison. 2021. Retargeted Self-Haptics for Increased Immersion in VR without Instrumentation. 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, 13 pages. <https://doi.org/10.1145/3472749.3474810>

1 INTRODUCTION

The developer and user base of virtual reality (VR) has grown tremendously in recent years due to successful consumer-oriented devices like the Oculus Quest 2¹ and HTC VIVE Cosmos². These systems offer immersive graphics and audio, but physical touch feedback continues to be limited. The very highest-end consumer systems feature dual handheld controllers with integrated vibrotactile actuators. Of course, a buzzing sensation applied to one's palm falls short of any realistic interaction with a physical object or surface. Research systems utilizing e.g., special room infrastructure and body exoskeletons are expensive, heavy, and generally limit mobility and consumer viability. More lightweight and mobile-friendly VR experiences prefer to avoid encumbering the user's hands and operate entirely in the air, offering no means for haptic feedback (e.g., *Waltz of the Wizard* on the Oculus Quest). In short, although we can build highly-detailed, near-photo-realistic digital worlds, we cannot yet reach out and feel them in ways practical for consumer

¹<https://oculus.com/quest-2/>

²<https://vive.com/us/product/vive-cosmos/overview/>

adoption. For this reason, innovative haptic feedback approaches in VR is an active area of research across several communities.

In this paper, we describe a new take on haptic delivery in VR – to use one’s own body for physical feedback (Figure 1). This “Self-Haptics” approach is highly practical, as we do not require any additional hardware to achieve a haptic effect. Instead, through graphical (i.e., software) manipulation of a user’s virtual hand position and pose, we can guide the actual hands together in a realistic way during bimanual tasks, what is known as retargeting or redirection in the literature [3, 35]. At this intersection point, with careful design, one hand can provide a physical shape or surface that corresponds to an interactive element in VR, allowing the other hand to physically feel and interact with a virtual element. We note that perfect realism is not our goal – our haptic feedback is inherently coarse (i.e., lacking appropriate texture, compliance, and contour). However, we found the visuo-audio-haptic fusion makes the interaction surprisingly fun and immersive, as deemed by our study participants. We also note that our effect is not illusory, as users can readily sense their own body is being co-opted for haptic effect. In some use cases, we can even provide appropriate haptic feedback to both hands – a haptic category that found to be most successful in our user study.

After reviewing key related work, we describe a design space we used to categorize our explorations. In total, we built twelve interactive demos that exemplify three haptic categories. To investigate if these paradigms were successful, we selected two demos per haptic category to show to participants and gather feedback. At a high level, our study shows imminent feasibility; the haptic feedback works and serves to boost realism and immersion compared to experiences with no haptics. That said, the interactive design tends to be highly bespoke (unlike vibration feedback, which can be easily applied to e.g., any collision), and thus is not going to be applicable in all VR contexts and interactions. Nonetheless, we believe our approach is an intriguing new way to deliver haptics to users in a practical manner, given that it requires no new or extra hardware. Indeed, any VR headset with uninstrumented hand tracking (such as the Oculus Quest 2) could support games and other experiences incorporating self-haptic interactions. We hope this paper serves to stimulate new ideas, and we have no doubt that creative VR developers could build unique and compelling experiences around this concept.

2 RELATED WORK

There is a significant body of research on creating haptic feedback in VR, from using props [61] to wearable mechanisms [21]. Closer to our work are systems that use the body as an input method, which innately include touch feedback to the body [65]. Most related to our work are implementations of redirected or retargeted haptics, which we review in greater detail. Lastly, we cover existing work on the concept of “self-haptics”.

2.1 Hand-Centric Haptics in VR

Current popular consumer systems use handheld controllers, which exclude natural, free-hand interaction with virtual objects. These controllers provide mostly vibrational haptic feedback to users. A large body of research has presented creative solutions to provide

haptics for the arm and hand, including passive [29, 51] and active props [2, 6, 31] that match the shape and location of virtual objects. Handheld systems like TORC [39], CapstanCrunch [52], and Haptic Links [56] add new haptic attachments to controllers. Other wearable systems constrain the fingers using strings [21, 60], exoskeletons [17–19], and electrostatic breaks [30]. While these prior systems increase immersion by simulating grasping and resistive forces, they also add hardware, cost, and weight (to a potentially mobile setup). Some systems limit the freedom of movement or are tethered to a larger apparatus, which is also less desirable.

Another body of research has looked into contactless haptics (i.e., without instrumenting the body), which allows free-hand interactions. One approach is to use pressurized air [58] or air vortexes [53] to deliver tactile sensation to the skin. Another approach uses ultrasonic phased arrays, which vibrate the air at high frequencies and generate localized regions of high pressure [32, 33, 49]. The haptic effect is akin to a tapping or vibration on the skin; kinesthetic forces that one might feel when grasping an object are not currently possible [46]. The \$3750 UltraHaptics STRATOS Inspire³ offers a 0.15m³ interactive volume, requires wall power, and weights 3.1 kg – all properties that preclude large-area hand interactions as well as free-movement VR uses.

In contrast to these prior methods, our self-haptics approach requires no instrumentation of the hand (or indeed any new hardware), offers a comparatively large interactive area, and is inherently mobile since it runs as software on existing VR headsets. Of course there are also limitations of our technique, as we will discuss later, making it well-suited for particular interactions.

2.2 On-Body & On-Skin Interaction

On-body interactions have unique design implications to accommodate operating on the skin [7, 27, 28]. In addition, touch interactions on the body innately involve “dual” tactile feedback that reinforces the interaction feedback loop. Inspired by how people naturally use and interact with different body parts, many researchers have designed speculative and functional interaction paradigms [16, 57]. Some systems are input only, such as [20, 25, 50], while other systems have used projectors to render graphical interfaces onto the skin’s surface [26, 42, 59]. Interfaces can also be rendered in VR that match body locations, so as to provide a physical surface for input, as shown in ActiTouch [65] and PinchType [22]. Natural poses and gestures also inspire metaphors and enable rich and intuitive interactions such as keyboard-free typing [22, 38] and general-purpose control [44, 47, 55]. Similarly, we also leverage the availability of our own body as a means for haptic feedback in VR.

2.3 Redirected & Retargeted Touch in VR

Visual dominance is the effect that when visual and tactile information conflict, the visual stimuli is dominant. Studies have shown that people often are not aware of the conflict and prefer the visual shape over the tactile shape [48]. Techniques like redirected touching or haptic retargeting leverage visual dominance and temporal coincidence [54] of visual-haptic feedback to change how virtual objects are perceived. Both techniques decouple the real and virtual hand position to provide haptic feedback with limited space and

³<https://ultraleap.com/product/stratos-inspire/>

	Two Hands Receive Appropriate Feedback	One Hand Receives Appropriate Feedback
Two-Handed Task	Examples Bow & Arrow Smartphone Ammo Magazine Baton	Examples Scanner & Keypad Striking Match Hammering Nail Cutting Vegetables
One-Handed Task	Not Applicable	Examples Launch Button Whisk Fingerprint Scanner Knocking Door

Figure 2: To help guide our efforts, we used a simple design space with two axes: one- vs. two-handed tasks, and one vs. two hands receiving appropriate haptic feedback. Note that one cell is not applicable (one-handed tasks wherein two hands receive haptic feedback). The dozen interactive demos we built are balanced across the three possible categories.

physical props. For instance, Ban et al. warped the hand/finger position to modify the perceived shape of a curved surface [4]. Kohli et al. explored the mapping of many virtual objects to one physical prop, and they demonstrated that users could adapt to the warped virtual world and train as well as in an unwarped one [35, 36]. Similar to redirected touching, haptic retargeting is a technique for repurposing passive props. For this, researchers have explored unimanual [1] and bimanual [23, 41] retargeting using a combination of world warping [15, 43] and body warping [3, 8, 63] methods. Later evaluations indicate a range of thresholds within which the warping of hand positions is imperceptible [1, 64, 66]. We build our approach on top of these body-warping/retargeting techniques, but instead of using physical props, we warp the hands of a user such that they come into contact with one another in a purposeful way in order to achieve a designed haptic effect.

2.4 Human-Generated Haptics

Prior work has considered the use of manual labor instead of machines to deliver haptic feedback. Notably, Cheng et al. explored the idea of having humans provide haptics for users in VR [13] as well as also groups of users in VR providing haptic feedback to one another [12, 14]. Our approach is related in that we also rely on human-generated actions and the human body to provide haptic feedback, though our work does not require other humans to achieve this goal, allowing for single-user experiences.

2.5 Self-Haptics

Most related to our work and the least explored is “self-haptics”, which is haptic feedback of touching one’s own body [10]. Bovet et al. highlights the lack of exploration into the benefits of passive haptic feedback provided by self-contact [11]. Only a handful of works have presented concepts that leverage self-contact and the availability of the body to provide passive haptic feedback. Kohli introduced the concept of “the haptic hand”, where the non-dominant hand hosts a virtual touch interface in VR, such that the dominant hand receives haptic feedback upon touch [37]. Bimanual tasks are not explored, nor are retargeting techniques employed, as the interfaces are overlaid on the non-dominant hand. In “Air Haptics”, Ban et al. pioneered the idea of retargeting/warping the thumb and index finger to create pinching and pulling sensations for virtual objects [5]. For example, when pinching a virtual ball, the fingers are rendered as in contact with the ball’s surface, but in reality “self join” (i.e., fingers that physically touch one another) to create a haptic effect. In our work, we extend the application space beyond the warping of pinched fingers to 1) retargeting the entire hand, 2) for unimanual and bimanual interactions, and 3) for more complex shapes and interactions.

3 DESIGN SPACE & EXAMPLE USES

To guide our brainstorming and development efforts, we formulated a simple design space organized along two key dimensions (Figure 2). First is whether tasks require one or two hands (i.e., uni- or bi-manual). For example, striking a match on a matchbook requires two hands, whereas pulling a lever requires one. Our second axis is whether one hand or both hands receive appropriate feedback. We define “appropriate” feedback as when the magnitude and direction of forces approximate the real-world action, a concept we will elucidate through examples discussed subsequently. In the case of one-handed tasks, two hands are not involved to receive appropriate haptics, and so the bottom left of our design space is not applicable.

We also use this section to describe twelve interactive demonstration applications we created to help us formalize our ideas and




Figure 3: In this bow and arrow target practice example, the user would either experience (a) no haptics or (b) self-haptics, where the fingers latch and exert pulling forces, simulating a bow under tension.




Figure 4: The user’s index finger taps on a smartphone screen either feels (a) nothing and pass through the 3D model of the phone or (b) self-haptics, where the finger is stopped by the palm, creating the sensation of a surface, while the phone-holding hand feels the appropriate downward impulse from finger taps.

illustrate our core concepts. These also serve to convey the generalizability of our approach across one- and two-handed interactions, as well as a variety of objects, including weapons (for e.g., games) and tools (e.g., training simulations), small handheld items (e.g., matchbook) to larger fixed objects (e.g., doorknob), and common interactive widgets (e.g., buttons, touchscreens).

3.1 Two-Handed Task / Two Hands Receive Appropriate Feedback (2T2F)

In this category of bimanual interactions, we aim for both hands to receive coordinated feedback that feels appropriate to the virtual task at hand. For instance, when pulling on the string of a bow, the hand holding the bow is pulled towards the string, and the hand holding the string is pulled towards the bow, with the opposing forces matching in magnitude. Thus, a virtual version of this interaction should also match these force vectors in order to




Figure 5: In this game, the user pushes the ammunition magazine to reload and feels either (a) no haptics or (b) self-haptics, where the two hands impact, generating a sensation that the magazine has been fully loaded.




Figure 6: The user swings a baton and receives either (a) no haptic feedback or (b) self-haptics where the fist falls into the other hand, simulating the impact of the baton swing.

feel appropriate. We built such a bow and arrow demo, where the user’s left hand holds a bow, which we re-target such that we place the left thumb (extended) to where the virtual bow string appears. When the right hand reaches out to hook the bow string and pull to fire an arrow, the user instead latches onto their own thumb (Figure 3). This offers a physical sensation similar to that of a bow under tension.

We also created a virtual smartphone with a self-haptic touchscreen. Specifically, we shift the virtual phone-holding hand in the palm normal direction by the thickness of our virtual phone model. When a finger on the other hand taps on the virtual touchscreen, rather than feeling nothing but air (Figure 4a) or having the finger pass through the 3D model, the finger is stopped by the palm, creating the sensation of a surface (Figure 4b). Although the palm receives a less convincing sensation, the overall downwards impulse matches the real-world force that would be transmitted through a phone to the hand holding it during touch input, and thus some realism is achieved.

In a game-oriented example we built, the user reloads a firearm by sliding an ammunition magazine into the grip. Due to the length




Figure 7: The user strikes a match against a matchbox and feels (a) nothing or (b) a resistive force caused by dragging the thumb across the opposing hand’s palm.




Figure 8: While chopping a carrot, the user (a) feels nothing and sees the knife pass through the cutting board, or (b) the knife physically stops at the board due to self-haptics.

of the magazine, a non-retargeted interaction would never have the user's hands come into contact (Figure 5a). However, by warping the two hands closer together in real-world space, we can cause them to intersect, and in turn, generate the sensation of the magazine being fully loaded (Figure 5b). Finally, we also built a demo of a baton that can be rapped in the other hand. By similarly warping the position of the hands in real-world space, we can cause the impact felt by both arms feels appropriate in terms of force and direction.

3.2 Two-Handed Task / One Hand Receives Appropriate Feedback (2T1F)

In this category of bimanual interactions, only one hand receives appropriate physical feedback (while the other receives incongruent feedback in service of the main haptic effect). For example, in the palmprint scanner and keypad demo we built, a user first places their hand onto the scanner, which causes a door keypad to flip open, after which they can enter the unlock code with a finger.




Figure 9: While holding a nail in the left hand, the user hits the nail with a hammer in the right hand and receives (a) no haptics or (b) self-haptics, where the right hand is halted by the left hand.




Figure 10: In this demo, a user can grab a whisk to mix in a bowl, either with (a) no haptics or (b) with self-haptics by presenting their left hand as a physical handle prop.

In a traditional VR experience, both hands would feel nothing interacting with these virtual elements (Figure 1a). However, by dynamically shifting the position of the virtual scanning hand to the left, the right hand ends up intersecting the back of the physical left hand at the location of the keypad. Now, when entering the unlock code, the user feels an actual physical surface (Figure 1b).

In another scenario we made, a user can strike a match and experience a resistive force while dragging the match across the coarse surface of a matchbox. To achieve this effect, we shift the matchbox-holding hand in the palm normal direction such that the thumb of the match-holding hand rubs against the other palm when striking the match (Figure 7). As the palm of the matchbox-holding hand does not expect any rubbing sensation, it is very much the case that only one hand receives appropriate feedback.

We also simulate the bimanual interaction of hitting a nail with a hammer. In our self-haptic implementation, we slide the virtual hammering hand downwards, following the orientation of the arm. The physical right hand then needs to reach higher in order for the hammer head to hit the nail. Now, as the right hand performs a hammering action, it impacts the left hand that is holding the




Figure 11: Here a user can place their finger on a fingerprint scanner, either with (a) no haptics or (b) with self-haptics, by presenting their left hand as a physical scanner prop.




Figure 12: In this VR scene, a red launch button is presented. In one case, (a) the user taps on the button with no haptic feedback. With retargeting, (b1) a translucent fist is also rendered in the scene, which (b2) turns green when the user matches the requested position and pose. (b3) Now, with their dominant hand, the user taps on the button and receives haptic feedback due to the other fist.

virtual nail (Figure 9). Although both hands receive haptic feedback, only the hammering hand receives appropriate feedback (i.e., while hammering in the real world, it is desirable not to receive an impact of equal magnitude on your hand holding the nail, as the nail is meant to be held lightly such that the nail slides through the fingers when hit). For this reason, we consider this a 2T1F interaction. Very similarly, in a vegetable chopping scenario, two virtual hands are warped such that the two physical hands touch when the virtual knife touches the cutting board, which would otherwise cut through the air in a non-retargeted, no-haptic scenario (Figure 8).

3.3 One-Handed Task / One Hand Receives Appropriate Feedback (1T1F)

Not all interactions require both hands, and so we also explored self-haptics for one-handed tasks. While we did consider how a single hand could provide haptics for itself, we found the design space to be much richer when the unused hand was appropriated as an “accessory prop”.




Figure 13: Here a user can knock on a door, either with (a) no haptics or (b) with self-haptics, by presenting their left hand as a physical door prop.

For example, when a red launch button is presented, typically in VR, the user would just reach out and “paw” at it with no haptic feedback (Figure 12, a). With our approach, a translucent fist is also rendered in the scene (Figure 12, b1), which turns green when the user matches the requested position and pose (Figure 12, b2). Now, with their dominant hand, the user can reach out to perform the one-handed interaction and receive some haptic feedback from the other fist, which is in similar shape and size to the button (Figure 12, b3).

Similarly, in a mixing whisk example, a translucent thumbs-up hand appears in the scene, and when the user matches it with their non-dominant hand, a whisk appears in the scene. When the user reaches out and grabs the virtual whisk, they end up grabbing the thumb of their non-dominant hand, which approximates the whisk handle (Figure 10). In two other examples we built, the unused palm serves as a backstop for interactions with the dominant hand (Figures 11 and 13). By explicitly lending one’s hand as a prop for an interaction, we found that most users are able to suspend disbelief. In doing so, haptic sensations on the prop hand tend to be ignored by the user, and are not perceived as particularly incongruent; in fact, they are expected. Instead, we found that users generally attend more to the primary haptic effect.

4 IMPLEMENTATION

For hardware, we used an Oculus Rift DK2⁴ with a Leap Motion⁵ mounted to the front, connected to an Intel Core i7 laptop. Our software pipeline and interactive scenarios were developed in the Unity 3D engine. For hand tracking, we used a Leap Motion (infrared stereo camera) and its Orion hand tracking API. We found its accuracy to be more than sufficient to prototype our example interactions, which are generally more gross-motor oriented. We note that our technique is not limited to any one tracking technology, and will broadly benefit from continued advances in hand tracking.

⁴<https://github.com/facebookarchive/riftdk2>

⁵<https://ultraLeap.com/product/leap-motion-controller/>

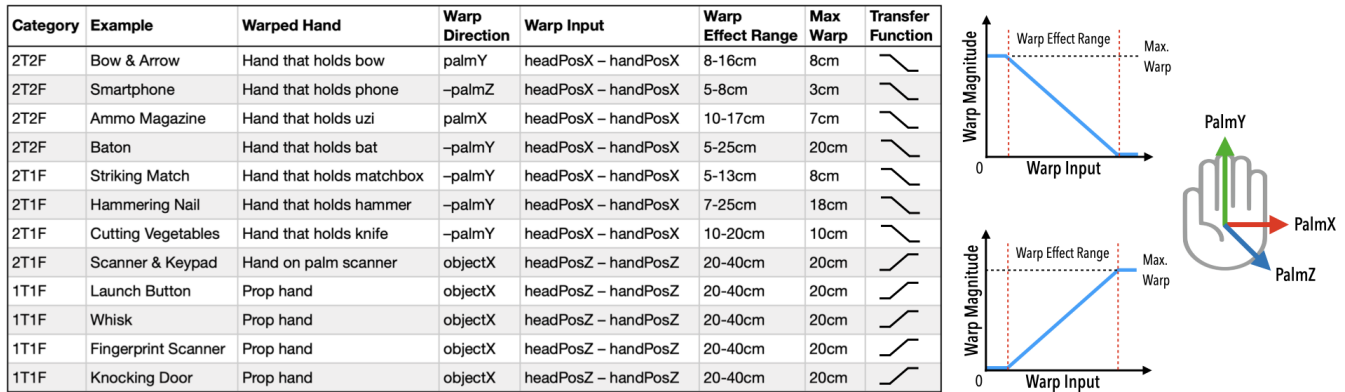


Figure 14: Retargeting parameters for each example scenario. Note the relationship between the warp input and the warp magnitude is one-to-one linear within the effect range.

4.1 Retargeting Parameters

For retargeting, we modified the “body-warping” technique described in prior work [3]. Instead of maintaining the real hand position when the virtual hand is retargeted to locations within an offset from the real hand position, we retarget the user’s real hand position by dynamically offsetting the virtual hand position. For example, if we want to shift the real hand to the right, we offset the virtual hand in the opposite direction so that the user would naturally move their hand to compensate for the offset and end in the desired position which enables self-haptics.

We use two ways to guide the non-dominant hand. In 2T2F and 2T1F, the non-dominant hand simply reaches out towards the target location as seen in the virtual world (e.g., the user reaches out to put their hand onto the palmprint scanner, which denotes the “correct position”). In 1T1F, the “correct position” is marked by a grey translucent hand, which turns green when correct. Following the approach in [3], we defined an interaction radius, inside of which linear warping is applied (Figure 14, right). The amount of warp is dependent on the distance to the target. The warp relationship is symmetric, and so it is applied and unapplied as the user approaches and departs the target. When the hand is not in the region, no redirection is applied. The retargeting parameters for each of our demos can be found in Figure 14. To further facilitate replication and encourage experimentation with VR self-haptics, we have open sourced a series of “hello world”-style example applications at <https://github.com/FIGLAB/RetargetedSelfHaptics>. These can be easily adapted to create a multitude of demos, often by simple substitution of media assets.

5 EVALUATION

To gather feedback on the self-haptic categories we identified, as well as evaluate aspects such as immersion and realism, we selected two exemplary scenarios from each of our three haptic categories: 1) For two-handed task / two hands receive appropriate feedback (2T2F), we selected *smartphone* and *ammo magazine*; 2) for two-handed task / one hand receives appropriate feedback (2T1F), we selected *striking match* and *scanner & keypad*; and finally, 3) for one-handed task / one hand receives appropriate feedback (1T1F),

we selected *launch button* and *knocking door*. For each scenario, there are two conditions: no haptic feedback and with self-haptic feedback. The only difference for the no-haptic version is that no redirection was applied to participants’ hands. Additionally, the no-haptic version of the 1T1F tasks did not require the non-interacting prop hand, and the participants were not shown a translucent guide for the prop hand. The order of the conditions, as well as the order of the scenarios within each category, was randomized across participants, and the order of the categories was randomized within participants. We recruited 11 participants (1 female, mean age 25) for the study, which lasted roughly 60 minutes. Participants were paid \$20 for their time. All participants self-reported as being right-handed and all but one participant had some prior VR experience.

5.1 Questionnaires

We devised questionnaires to investigate three key questions: 1) Is retargeted self-haptics able to provide tactile feedback? 2) What is its effect on ownership and agency of the virtual hands? 3) How do users rate the different self-haptic categories in terms of immersion, fun, and realism?

We use the corresponding questions (Q1-2 and Q6-9) from the Embodiment Questionnaire [24] and the given formula to aggregate the questions into the ownership and agency subscales. Q3-5 from the original questionnaire refer to looking in a virtual mirror and were not applicable to our study, and thus dropped. Questionnaire items were rated on a 7-point Likert scale ranging from -3 (strongly disagree) to +3 scale (strong agree); please refer to [24] for complete details of this measure.

Unlike the questions on agency and ownership, we did not find a questionnaire in the literature compatible with our investigation on tactile sensation. Thus, we designed a new scale drawing inspiration from [1, 3, 37, 66] and in consultation with Dr. Roberta Klatzky. Following best practices, this questionnaire went through several rounds of iteration, utilizing 12 pilot participants separate from our later study participants. The result of this design process was the following four questions: T1) I felt the virtual object; T2) The virtual object felt like it was there; T3) The virtual object felt real; and T4) I did not feel anything when my hand touched the object.

Note that T4 is a negative question, and thus the value is flipped in statistical analysis (a technique also used in [24]). We use the same Likert scale (-3 to +3) as [24]. We then aggregate these four questions into a single tactility subscale by taking their mean.

The above ten questions were given immediately after participants completed each scenario, whether it was with no haptics or self-haptics. At the end of a haptic category (i.e., after both no and self-haptic, and both exemplary scenarios in the category), we administered an additional seven-question Likert-scale questionnaire to measure immersion, fun, realism: 1) These two “air feedback” examples felt realistic; 2) These two “self feedback” examples felt realistic; 3) These two “air feedback” examples made me feel more immersed in the scene; 4) These two “self feedback” examples made me feel more immersed in the scene; 5) These two “air feedback” examples were fun; 6) These two “self feedback” examples were fun; and 7) I preferred the “air feedback” over “self feedback” examples. To mitigate order effects, we counterbalanced the ordering of questions 1-6, and the condition text in question 7.

We note that our decision to describe no feedback as “air feedback” was informed by iterative pilot testing of our questionnaire (see above). Specifically, we found the term “no feedback” to be confusing to participants, as they still experienced proprioception and air resistance, and also felt their own hands when performing grasps. Calling this category “no feedback” in the actual presence of feedback had a negative connotation and introduced bias. We found the colloquial “air feedback” term to effectively mitigate this issue.

5.2 Procedure

Participants were asked to fill out a pre-study survey including demographic information and three exemplary Likert-scale questions [40] to acquaint themselves with the questionnaire format. They were then introduced to the concept of “using one’s own body to provide tactile feedback in VR”. During the experiment, participants remained seated in a chair in the center of a $1 \times 1 \text{ m}^2$ open area and in front of the Oculus tracking camera. Participants tried each scenario in two conditions (no feedback and self-haptic feedback). Upon entering the scene, they followed verbal instructions to perform the interaction. The participants could repeat the interaction as many times as they wish, and they were encouraged to think-aloud during the interaction. After each scenario (no haptics and self-haptics), participants completed the ten-question questionnaire. At the end of each haptic category (both scenarios with no haptics and self-haptics), they completed the seven-question questionnaire. After the participants had completed all of the scenarios in both conditions, they were asked to rank the self-haptic versions of all of the scenarios (1-most preferred, 6-least preferred). The study concluded with a semi-structured exit interview.

6 RESULTS & DESIGN RECOMMENDATIONS

We performed statistical analysis on the quantitative questionnaire data, and we also coded and thematically clustered participants’ think aloud and exit interview statements to extract trends. This data directly informed a series of design recommendations, which we now elucidate. We start with more general findings, before transitioning to more fine-grained considerations.

6.1 Tactile Sensation

The most immediate question to answer was whether our self-haptic manipulation successfully created the feeling of touching virtual objects. To analyze the success of our tactile sensations, we used our aforementioned tactility subscale and performed a two-way ANOVA on data that was Aligned Rank Transformed [62], as the Likert scores of the questions did not have a normal distribution (Shapiro-Wilk normality test $p < .001$). The haptic feedback condition significantly affected tactility ($F_{1,110} = 161.406$, $p < .001$), with self-haptics rating significantly higher in tactility ($M = 3.5$, $SD = .7$). There was no significant difference between scenarios ($F_{1,110} = .510$, $p = .768$) and no interaction effect ($F_{1,110} = .502$, $p = .740$) (Figure 15). Participants feedback strongly supported this result, e.g., “it was nice to get a real feedback that I am pressing something!” (P6). As there was no significant difference across our self-haptic design space, we recommend all categories for use, which are useful and applicable in different interactive contexts.

6.2 Ownership & Agency

Feeling embodied in a virtual avatar helps users to engage in the virtual experience. Ownership and agency are two important aspects that affect the elicitation of the embodiment illusion in a virtual environment [24]. In comparison to conventional VR interactions, our technique manipulates virtual hand position, which presumably would affect users’ feeling of ownership and agency. To analyze the ownership and agency subscales, we performed 2 two-way ANOVAs on Aligned Rank Transformed data [62] as the Likert scores of the questions did not have a normal distribution (Shapiro-Wilk normality test $p < .001$). The analysis revealed no significant difference for the main effect haptic condition ($F_{1,110} = 3.636$, $p = .059$ and $F_{1,110} = 3.783$, $p = .054$; respectively), no significant main effect for scenario ($F_{1,110} = 1.000$, $p = .421$ and $F_{1,110} = 1.485$, $p = .200$; respectively), nor an interaction effect ($F_{1,110} = 1.110$, $p = .359$ and $F_{1,110} = 1.257$, $p = .288$; respectively) (Figure 15).

This result suggests that our retargeted self-haptic technique does not have a significant influence on the ownership or the agency of the virtual hands, which is important for task completion and usability in VR. However, anecdotally, we found that mismatches




Figure 15: Results of tactile sensation, ownership, and agency. Error bars represent standard error (SE). 2T2F: two-handed task / two hands receive appropriate feedback; 2T1F: two-handed task / one hand receives appropriate feedback; 1T1F: one-handed task / one hand receives appropriate feedback. The haptic feedback condition significantly affected tactility ($p < .001$).

introduced through the warping of the virtual hand becomes salient when the two hands physically touch (i.e., but are not shown to be touching in the VR scene). This caused some users to dissociate the virtual hand as their own. Although this was not a statistically significant effect, nor remarked on by most participants, we do see that our 2T1F self-haptic category received a lower rating for ownership compared to other categories (Figure 15). Additionally, as with other retargeting systems, the threshold of the acceptable displacement requires careful tuning. However, at a high level, our findings suggest that the less displacement needed, the more successful the illusion. Beyond some threshold, the visuo-haptic mismatch becomes too great for a believable experience.

6.3 Immersion & Realism

Immersion and realism are two aspects that influence the feeling of presence in the virtual environment. For the immersion subscale, a significant difference was found for the main effect haptic condition ($F_{1,110}=7.696$, $p=.010$), with self-haptics rating significantly higher ($M=4.8$, $SD=1.7$). No significant difference was found for the main effect self-haptic category ($F_{1,110}=.109$, $p=.474$), nor for the interaction effect ($F_{1,110}=2.625$, $p=.143$). Similarly, for realism, a significant difference was found for the main effect haptic condition ($F_{1,110}=7.758$, $p=.008$), with self-haptics again rating higher ($M=5.2$, $SD=1.7$). As before, there was no significant difference found for the main effect self-haptic category ($F_{1,110}=7.169$, $p=.897$), nor for the interaction effect ($F_{1,110}=2.024$, $p=.082$). These results are shown in Figure 16.

From this, we can conclude that our self-haptic technique can increase immersion and realism, and among the three categories, the 2T2F category was rated the highest. From user feedback, this result was due to 2T2F's providing appropriate haptic feedback to both hands, while making minimal modification to the action itself, which we explain in more detail in the following sections. For example, P2 commented that loading the ammo magazine (a 2T2F example) where the left hand provided self-haptic feedback "felt good because [they were] expecting a push on the right hand". At a high level, our studies and user feedback show that self-haptics can increase immersion and realism, and if this is an interactive dimension one wishes to improve in your VR application's design, this technique is worth exploration.

6.4 Haptic Condition User Preference

In traditional redirected touch and haptic retargeting, the warping of the hand is imperceptible within a threshold, as the primary source of sensory information comes from visual perception. When a second layer of information is provided (in this case, tactile feedback from another body part), the redirection illusion can break, which can feel disorienting at first. However, participants reported that as they continued to repeat the interaction, they became acquainted with the mismatched information and appreciated the feedback from the "haptic hand". One way to mitigate the unexpectedness or minimize the mismatch is to incorporate the shift into a natural gesture. Thus, we expected our 2T2F category to be the most preferred as both hands receive appropriate feedback and the bimanual interaction was the least modified. To test our hypothesis, we conducted a non-parametric Friedman test (normality




Figure 16: Results of immersion, realism, fun, and preference. Error bars represent standard error (SE). 2T2F: two-handed task / two hands receive appropriate feedback; 2T1F: two-handed task / one hand receives appropriate feedback; 1T1F: one-handed task / one hand receives appropriate feedback. For Immersion and Realism, a significant difference was found for the haptic condition ($p<.05$).

test $W=.829$, $p<.001$) on the preference ratings for each category and rendered a Chi-square value of 1.19, which was not statistically significant ($p=.552$). See also Figure 16.

Even though the above analysis showed that there was no significant difference among the self-haptic categories for user's preference, we looked at the participants' ranking of the self-haptic demos in concert with anecdotal comments to draw design recommendations. From this data, we see a trend that the 2T2F category is ranked the highest and most successful overall. Participants shared that experiencing self-haptics was the "most natural" and unsurprising for examples in this category because there was little to no change to how the interaction would have occurred in the real world. In comparison, our 1T1F modality was less preferred, as the prop hand is not part of the conventional (i.e., real world) one-handed task and thus it is naturally less intuitive for it to be involved. However, after a few tries, participants reported the haptic feedback to the prop hand became less conspicuous, and some reported to have "stopped paying attention to my [prop] hand's position" (P3). For the same reasons, our 2T1F category was also rated lower, as only one hand receives appropriate haptic feedback despite it being a two-handed task.

6.5 Playfulness

A vast majority of haptics research aims to simulate fine-grained, precise tactile sensations (e.g., the contour of a bottle or the compliance of a pillow). Thus, the evaluation of such work is often focused on the perceptual similarity of the simulated sensation to that of the real world, as well as the performance metrics of the novel hardware. While our approach also seeks to increase realism by providing tactile feedback, it is inherently limited by solely relying on the user's own body, and not a new hardware intervention (e.g., gloves, exoskeleton).

As a consequence, our haptic sensations are inherently less precise and more coarse – almost a caricature of the real world (e.g., our demo involving a big red launch button) – giving our interactions a more gross-motor-oriented and playful flavor, and requiring some suspension of disbelief. We do not view this as a negative, as many VR experiences are created for play and entertainment. To

measure this, we included questionnaire questions about the “fun-ness” of experiences in both no feedback and self-haptic feedback conditions. We found a significant difference for the main effect haptic condition ($F_{1,110}=.113$, $p<.001$). Note that the self-haptics categories had no significant effect ($F_{1,110}=18.205$, $p=.882$) on fun, nor was there an interaction effect ($F_{1,110}=.816$, $p=.445$). Put simply, all three self-haptic categories were significantly more fun ($M=5.9$, $SD=1.4$) than their no-haptic counterparts (Figure 16), which was universally echoed in participants’ think-aloud feedback. Thus, as a design recommendation, we suggest pairing our approach with more playful VR application domains, such as games and immersive artistic experiences. In more formal domains, such as telepresence or telesurgery, where precision and fidelity is required, other haptic approaches may be better suited, or could be combined with self-haptics to create additional effects.

6.6 Contact Duration & Surface Area

Anecdotally, we observed that the duration of contact influenced realism. When users slowed down their interactions - for example pressing a button tentatively or dwelling on a button - they more frequently remarked on mismatches between what they were seeing and what they were feeling. However short-duration events, like our knocking on door and loading ammo magazine scenarios, were highly rated. We suspect that when interactions are short in duration, the dominant haptic effect is from the impact, and not the underlying geometry of the surfaces. However, when performing actions slowly, or when remaining in contact with a self-haptic surface, realism degrades due to geometric and textural mismatches. Thus, we recommend that retargeted self-haptics are best utilized for short and fast (impact-oriented) interactions.

Similar to our findings on contact duration, we also found users remarking on degraded realism when interacting with scenarios involving larger contact areas. For example, our keypad demo had users press buttons on the back of their hand using a fingertip (i.e., small contact area), which revealed little about the underlying geometry (i.e., the dominant tactile sensation was that it was a flat surface, much like the top surface of a flat button). In contrast, our whisk demo had users wrap one hand around the thumb of the other hand - a contact area involving all four non-thumb digits. The irregular geometry of the thumb, its limited length and softer compliance, all contributed to the degraded illusion of holding a metal-handled whisk. For this reason, we suggest designers tend towards smaller contact-area interactions.

Note that we define “duration” and “area” coarsely. The difference in surface area between a finger tip and a palm is more than an order of magnitude, and similarly, contact duration could be a brief instant (e.g., loading ammo) or a sustained grasp (e.g., grabbing the thumb as the whisk handle). Our findings focus on the polar ends of this spectrum, but not the continuous fine grained space in between, which merits careful future work.

6.7 Target Size & Tracking Error

Uninstrumented hand tracking with computer vision is a very active area of research. Self-haptic designs will have to contend with some level of tracking imprecision for the foreseeable future, which leads to another design consideration. We found that small tracking errors

could dramatically break realism. Most notably, in our whisk demo (where a user’s thumb was used as a handle prop), hand tracking errors could cause a user to bump into the “handle” prematurely or even miss the thumb entirely. This was because the thumb not only provided a relatively small and unforgiving target, but one hand often occluded the other when reaching out to grasp the virtual whisk, which can cause momentary tracking loss. In contrast, our keypad demo (which used the back of the hand to simulate a keypad surface), was much more forgiving. Hand tracking errors generally went unrealized by participants, because even with drift, the back of the hand provided a sufficiently large surface area to “catch” finger touch events. Moreover, its relatively flat surface meant that the sensation of pressing a flat button was present over a large area. Put simply, the self-haptic surface was harder to miss, and even with tracking errors, the haptic sensation remained consistent across the collision surface - important properties that designers should try to incorporate when possible. Similarly, we suggest that designers minimize building interactions utilizing small body features and also surfaces with irregular geometries.

6.8 Mapping Virtual Objects to Body Parts

Through our iterative development process, we identified five factors that are key to identifying the body part best suited to convey a haptic effect. We only considered the arms and hands in this work, as they can be easily brought into the interactive volume, but it may be other body parts could contribute haptics in the future (e.g., a virtual keyboard that uses the thighs as a haptic surface while seated). First, we found that matching size was key. If one is aiming to simulate the grasping of e.g., a soda can, grabbing ones thumb immediately breaks the realism due to a dramatic size mismatch. Instead, a hand formed into fist or using the forearm would provide a volume better approximating the size of a beverage can.

Second, and equally important in our estimation, is body contour. Returning to our example of a soda can, while a fist decently approximates the size, the surface irregularity created by bent fingers produces a decidedly non-smooth, cylinder-like object. The forearm, however, is naturally smooth and better emulates the surface geometry of an aluminum can. Conversely, the fist is an excellent proxy for a large irregular rock. This bring up a third point - whenever possible, we suggest selecting virtual objects without well-defined sizes. While the size of a soda can is well known by people, allowing them to readily detect comical mismatches, rocks come in any shape and size, offering designers much greater freedom and improving end-user realism.

Fourth is what we call appropriate anchoring. For instance, using one’s thumb to simulate the haptics of a virtual joystick works well, because the degrees of freedom match (i.e., two rotational axes with an origin at the base). Furthermore, users tend to manipulate the stick, but the whole joystick apparatus tends to stay fixed in reality. As such, in a VR setting, the user can self-manipulate their own thumb, without having to pull their other hand around (i.e., the thumb is kept loose, while the host hand largely stays still in the air). This is very different from our whisk example scenario, in which the thumb was also grasped, but then the whole hand had to be translated around with it to simulate whisking ingredients in a bowl. Participants noted how unnatural this felt, due to not

only the added mass of the “unused” hand but also the inadvertent haptics on that hand detracted from the overall experience. Thus, we recommend selecting virtual objects and body parts with similar degrees of freedom and mass, and to avoid translating parts of the body that are not part of the virtual object.

Finally, to create the feeling of a large, heavy or fixed object, we suggest pairing a large body element with a small one. For example, in our fingerprint scanner demo, an index finger tip is touched to the palm of the other hand. The larger mass of the other hand, combined with the relatively low impact force of a finger tap, provides a reasonable illusion the scanner is fixed to a (virtual) wall. However, in our door knocking demo, one hand knocks (a more vigorous action than a finger tap) against the other hand. The essentially equal masses of the two hands lessens the feeling of knocking on a big heavy door, as the other hand gets repulsed upon each knock. As noted in our section on Playfulness, the goal of our approach is not to create high fidelity haptics (compared with no haptics, feeling any knocking sensation improves immersion), but nonetheless we include this discussion to guide future designers towards interactions that will be more successful.

6.9 Haptic Valence

While most self-haptic scenarios led to positive emotional valence and arousal, we observed that certain scenarios induced a more intense emotional response from our participants. This was typified by our carrot chopping scenario, where the hand holding the carrot acts as a physical barrier simulating the cutting board surface. This sensation is appropriate for the cutting hand, but it is rather unexpected – and more importantly unwanted – to feel a downwards force on one’s hand while the other hand is holding a knife. Unlike most of our other interactions, where the other hand felt ambiguous, confusing, or irrelevant tactile sensations, in this instance, the carrot-holding hand experienced what might be described as negative “haptic valence”. This negative effect was less pronounced for our similarly configured hammering nail scenario, but we suspect it extends, at least to some degree, in any scenario where there is potential to invoke self-harm. Thus, we recommend avoiding self-haptics in such tasks, unless a negative valence is desired.

6.10 Somatosensory Attenuation

Another psychophysical effect likely at play in our technique is somatosensory attenuation [9], in which the brain attenuates haptic sensations in self-touch events. Prior work has studied whether the suppression is also true with a user’s real finger touching an illusory hand (i.e., rubber hand illusion), and also a illusory finger touching a user’s real hand (a sort of “rubber finger” illusion). In both cases, attenuation occurs, assuming the user experiences ownership of both appendages [34]. Our arrangement is slightly different: a real hand touching another real hand, but one hand is represented visually by a virtual non-hand object, and thus ownership is less well defined. It is unclear if somatosensory attenuation occurs across our self-haptic categories, which poses interesting questions for future work.

7 LIMITATIONS & FUTURE WORK

A recurring issue we faced was poor hand tracking when the hands overlap (i.e., occlusion) as part of our interactive sequences. This is an open problem in computer-vision-based hand tracking, and many research groups are making tremendous strides. However, at present, the current best practice is to “avoid the overlap of two hands due to current computer vision limitations. A good way around this is to design interactions that can be performed with just one hand” [45]. Of course, our technique requires the two hands to touch each other, so we had to implement some basic mitigation strategies to overcome this current limitation. First, we had to avoid interactions that led to heavy occlusion – for instance, we had the idea of an ajar door (simulated by the left hand) being pulled open (by the right hand). Second, when we detect tracking loss, rather than the virtual hand disappearing, we freeze it in place, as more often than not, it was holding the same pose. Although a basic approach, this worked well in practice.

Since our technique utilizes a visuo-haptic mismatch, the surprise and disorientation when the two hands first touch is expected. Our technique requires a brief period of adaptation, but we found during our studies that all participants were able to quickly acquaint themselves with the shifted position of the hand and the self-haptic feedback. On the note of unexpectedness, the category of self-haptics where the non-dominant hand is used as a prop is inherently less intuitive and slower than the unmodified unimanual tasks.

We also note that our technique is not suitable for all interactions. For example, the technique would not be suitable for training tasks that require small haptic details and high precision. Additionally, our technique retargets different parts of the hands and arms, which have a limited surface area. Thus, this technique is more suitable for handheld objects instead of large and fixed obstacles (such as wall surfaces). The human body also has a limited pallet of textures and compliance (boney to fatty), which inherently feels different from commonplace object materials (e.g., metal, plastic). However, other body features such as hair and nails might be used in new and creative ways.

In our interaction design space, we have only explored scenarios that retarget the upper body, specifically the arms and hands. Yet to be explored are other body parts, such as the shoulder and legs. The thigh, in particular, offers promising opportunities for its large surface area and availability for touch interactions [27]. Moreover, in this paper, we only explored the interaction that involves tapping or touching the skin. In “Human–Computer Interaction on Skin”, Bergström and Hornbæk laid out four types of skin input [7]. We imagine future work could explore self-haptics that leverage gestures such as pinching, scratching, twisting, and swiping the skin.

Finally, future work could look into retargeting the fingers to enable single-handed self-haptics where the hand provides haptics for itself (extending the preliminary work in [5]). For example, when the user flicks a lighter, their index finger could serve as the surface of the lighter. Retargeted thumb-typing is another example in this category. Similarly, single-handed self-haptics could add interesting sensations to unimanual grips/grasps, which is a frequent interaction modality in VR that currently lacks haptic feedback.

8 CONCLUSION

In this paper, we have described a method for providing haptic feedback using the users' own body by retargeting the position of the hands. We crafted twelve exemplary demos, split across three haptic categories, where either one or both hands receive appropriate feedback during unimanual and bimanual interactions. We validated our approach through a user study which showed that our technique can effectively provide touch feedback and increase immersion and realism while maintaining presence. While this technique is not applicable in all use scenarios, it can nonetheless add rich haptics where there were previously none, requiring no new hardware.

ACKNOWLEDGMENTS

The authors would like to thank Prof. Roberta Klatzky and Prof. Sven Mayer for their guidance on the user study protocol and evaluation of the study results.

REFERENCES

- Parastoo Abtahi and Sean Follmer. 2018. Visuo-Haptic Illusions for Improving the Perceived Performance of Shape Displays. 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/10.1145/3173574.3173724>
- Bruno Araujo, Ricardo Jota, Varun Perumal, Jia Xian Yao, Karan Singh, and Daniel Wigdor. 2016. Snake Charmer: Physically Enabling Virtual Objects. In *Proceedings of the TEI '16: Tenth International Conference on Tangible, Embedded, and Embodied Interaction* (Eindhoven, Netherlands) (TEI '16). Association for Computing Machinery, New York, NY, USA, 218–226. <https://doi.org/10.1145/2839462.2839484>
- 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* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 1968–1979. <https://doi.org/10.1145/2858036.2858226>
- Yuki Ban, Takuji Narumi, Tomohiro Tanikawa, and Michitaka Hirose. 2012. Modifying an Identified Position of Edged Shapes Using Pseudo-Haptic Effects. In *Proceedings of the 18th ACM Symposium on Virtual Reality Software and Technology* (Toronto, Ontario, Canada) (VRST '12). Association for Computing Machinery, New York, NY, USA, 93–96. <https://doi.org/10.1145/2407336.2407353>
- Yuki Ban, Takuji Narumi, Tomohiro Tanikawa, and Michitaka Hirose. 2015. Air Haptics: Displaying Feeling of Contact with AR Object Using Visuo-Haptic Interaction. In *ACM SIGGRAPH 2015 Emerging Technologies* (Los Angeles, California) (SIGGRAPH '15). Association for Computing Machinery, New York, NY, USA, Article 5, 1 pages. <https://doi.org/10.1145/2782782.2792498>
- 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>
- Joanna Bergström and Kasper Hornbæk. 2019. Human–Computer Interaction on the Skin. *ACM Comput. Surv.* 52, 4, Article 77 (Aug. 2019), 14 pages. <https://doi.org/10.1145/3332166>
- 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.
- Sarah-Jayne Blakemore, Daniel Wolpert, and Chris Frith. 2000. Why can't you tickle yourself? *Neuroreport* 11, 11 (2000), R11–R16.
- Sidney Bovet and Ronan Boulic. 2017. Ensuring Self-Haptic Consistency for Immersive Amplified Embodiment. (2017).
- Sidney Bovet, Henrique Galvan Debarba, Bruno Herbelin, Eray Molla, and Ronan Boulic. 2018. The Critical Role of Self-Contact for Embodiment in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics* 24, 4 (2018), 1428–1436. <https://doi.org/10.1109/TVCG.2018.2794658>
- Lung-Pan Cheng, Li Chang, Sebastian Marwecki, and Patrick Baudisch. 2018. Iturk: Turning passive haptics into active haptics by making users reconfigure props in virtual reality. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–10.
- Lung-Pan Cheng, Patrick Lühne, Pedro Lopes, Christoph Sterz, and Patrick Baudisch. 2014. Haptic turk: a motion platform based on people. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 3463–3472.
- Lung-Pan Cheng, Sebastian Marwecki, and Patrick Baudisch. 2017. Mutual human actuation. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*. 797–805.
- Lung-Pan Cheng, Eyal Ofek, Christian Holz, Hrvoje Benko, and Andrew D. Wilson. 2017. Sparse Haptic Proxy: Touch Feedback in Virtual Environments Using a General Passive Prop. 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, 3718–3728. <https://doi.org/10.1145/3025453.3025753>
- Changeeok Cho and Huichul Yang. 2002. Body-Based Interfaces. In *Proceedings of the 4th IEEE International Conference on Multimodal Interfaces (ICMI '02)*. IEEE Computer Society, USA, 466. <https://doi.org/10.1109/ICMI.2002.1167040>
- Inrak Choi, Heather Culbertson, Mark R. Miller, Alex Olwal, and Sean Follmer. 2017. Gravity: 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>
- 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. <https://doi.org/10.1109/IROS.2016.7759169>
- 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>
- Niloofar Dezfuli, Mohammadreza Khalilbeigi, Jochen Huber, Florian Müller, and Max Mühlhäuser. 2012. PalmRC: Imaginary Palm-Based Remote Control for Eyes-Free Television Interaction. In *Proceedings of the 10th European Conference on Interactive TV and Video* (Berlin, Germany) (EuroITV '12). Association for Computing Machinery, New York, NY, USA, 27–34. <https://doi.org/10.1145/2325616.2325623>
- Cathy Fang, Yang Zhang, Matthew Dworman, 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>
- Jacqui Fashimpaur, Kenrick Kin, and Matt Longest. 2020. PinchType: Text Entry for Virtual and Augmented Reality Using Comfortable Thumb to Fingertip Pinches. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI EA '20). Association for Computing Machinery, New York, NY, USA, 1–7. <https://doi.org/10.1145/3334480.3382888>
- Eric J. Gonzalez and Sean Follmer. 2019. Investigating the Detection of Bimanual Haptic Retargeting in Virtual Reality. In *25th ACM Symposium on Virtual Reality Software and Technology* (Parramatta, NSW, Australia) (VRST '19). Association for Computing Machinery, New York, NY, USA, Article 12, 5 pages. <https://doi.org/10.1145/3359996.3364248>
- Mar Gonzalez-Franco and Tabitha C Peck. 2018. Avatar embodiment. towards a standardized questionnaire. *Frontiers in Robotics and AI* 5 (2018), 74. <https://doi.org/10.3389/frobt.2018.00074>
- Sean Gustafson, Christian Holz, and Patrick Baudisch. 2011. Imaginary Phone: Learning Imaginary Interfaces by Transferring Spatial Memory from a Familiar Device. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology* (Santa Barbara, California, USA) (UIST '11). Association for Computing Machinery, New York, NY, USA, 283–292. <https://doi.org/10.1145/2047196.2047233>
- Chris Harrison, Hrvoje Benko, and Andrew D. Wilson. 2011. OmniTouch: Wearable Multitouch Interaction Everywhere. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology* (Santa Barbara, California, USA) (UIST '11). Association for Computing Machinery, New York, NY, USA, 441–450. <https://doi.org/10.1145/2047196.2047255>
- Chris Harrison and Haakon Faste. 2014. Implications of Location and Touch for On-Body Projected Interfaces. In *Proceedings of the 2014 Conference on Designing Interactive Systems* (Vancouver, BC, Canada) (DIS '14). Association for Computing Machinery, New York, NY, USA, 543–552. <https://doi.org/10.1145/2598510.2598587>
- Chris Harrison, Shilpa Ramamurthy, and Scott E. Hudson. 2012. On-Body Interaction: Armed and Dangerous. In *Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction* (Kingston, Ontario, Canada) (TEI '12). Association for Computing Machinery, New York, NY, USA, 69–76. <https://doi.org/10.1145/2148131.2148148>
- Anuruddha Hettiarachchi and Daniel Wigdor. 2016. Annexing Reality: Enabling Opportunistic Use of Everyday Objects as Tangible Proxies in Augmented Reality.

- 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, 1957–1967. <https://doi.org/10.1145/2858036.2858134>
- [30] Ronan Hinchet, Velko Vechev, Herbert Shea, and Otmar Hilliges. 2018. DextrES: Wearable Haptic Feedback for Grasping in VR via a Thin Form-Factor Electrostatic Brake. 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, 901–912. <https://doi.org/10.1145/3242587.3242657>
- [31] Matthias Hoppe, Pascal Knierim, Thomas Kosch, Markus Funk, Lauren Futami, Stefan Schneegass, Niels Henze, Albrecht Schmidt, and Tonja Machulla. 2018. VRHapticDrones: Providing Haptics in Virtual Reality through Quadcopters. In *Proceedings of the 17th International Conference on Mobile and Ubiquitous Multimedia* (Cairo, Egypt) (MUM 2018). Association for Computing Machinery, New York, NY, USA, 7–18. <https://doi.org/10.1145/3282894.3282898>
- [32] Thomas Howard, Maud Marchal, Anatole Lécuyer, and Claudio Pacchierotti. 2019. PUMAH: Pan-tilt Ultrasound Mid-Air Haptics for larger interaction workspace in virtual reality. *IEEE transactions on haptics* 13, 1 (2019), 38–44.
- [33] I. Hwang, H. Son, and J. R. Kim. 2017. AirPiano: Enhancing music playing experience in virtual reality with mid-air haptic feedback. In *2017 IEEE World Haptics Conference (WHC)*. <https://doi.org/10.1109/WHC.2017.7989903>
- [34] Konstantina Kilteni and H Henrik Ehrsson. 2017. Body ownership determines the attenuation of self-generated tactile sensations. *Proceedings of the National Academy of Sciences* 114, 31 (2017), 8426–8431.
- [35] L. Kohli. 2010. Redirected touching: Warping space to remap passive haptics. In *2010 IEEE Symposium on 3D User Interfaces (3DUI)*. 129–130. <https://doi.org/10.1109/3DUI.2010.5444703>
- [36] Luv Kohli, Eric Burns, Dorian Miller, and Henry Fuchs. 2005. Combining Passive Haptics with Redirected Walking. In *Proceedings of the 2005 International Conference on Augmented Tele-Existence* (Christchurch, New Zealand) (ICAT '05). Association for Computing Machinery, New York, NY, USA, 253–254. <https://doi.org/10.1145/1152399.1152451>
- [37] Luv Kohli and Mary Whitton. 2005. The Haptic Hand: Providing User Interface Feedback with the Non-Dominant Hand in Virtual Environments. In *Proceedings of Graphics Interface 2005* (Victoria, British Columbia) (GI '05). Canadian Human-Computer Communications Society, Waterloo, CAN, 1–8.
- [38] Falko Kuester, Michelle Chen, Mark E. Phair, and Carsten Mehring. 2005. Towards Keyboard Independent Touch Typing in VR. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology* (Monterey, CA, USA) (VRST '05). Association for Computing Machinery, New York, NY, USA, 86–95. <https://doi.org/10.1145/1101616.1101635>
- [39] 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>
- [40] Rensis Likert. 1932. A technique for the measurement of attitudes. *Archives of psychology* (1932).
- [41] Brandon J Matthews, Bruce H Thomas, Stewart Von Itzstein, and Ross T Smith. 2019. Remapped Physical-Virtual Interfaces with Bimanual Haptic Retargeting. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 19–27. <https://doi.org/10.1109/VR.2019.8797974>
- [42] Pranav Mistry, Pattie Maes, and Liyan Chang. 2009. WUW - Wear Ur World: A Wearable Gestural Interface. In *CHI '09 Extended Abstracts on Human Factors in Computing Systems* (Boston, MA, USA) (CHI EA '09). Association for Computing Machinery, New York, NY, USA, 4111–4116. <https://doi.org/10.1145/1520340.1520626>
- [43] Roberto A. Montano Murillo, Sriram Subramanian, and Diego Martinez Plasencia. 2017. Erg-O: Ergonomic Optimization of Immersive Virtual Environments. 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, 759–771. <https://doi.org/10.1145/3126594.3126605>
- [44] Masa Ogata and Michita Imai. 2015. SkinWatch: Skin Gesture Interaction for Smart Watch. In *Proceedings of the 6th Augmented Human International Conference* (Singapore, Singapore) (AH '15). Association for Computing Machinery, New York, NY, USA, 21–24. <https://doi.org/10.1145/2735711.2735830>
- [45] Facebook Inc. Best Practices. [n.d.]. <https://developer.oculus.com/learn/hands-design-bp/>
- [46] Ismo Rakkolainen, Euan Freeman, Antti Sand, Roope Raisamo, and Stephen Brewster. 2020. A survey of mid-air ultrasound haptics and its applications. *IEEE Transactions on Haptics* 14, 1 (2020), 2–19.
- [47] J. Rekimoto. 2001. GestureWrist and GesturePad: Unobtrusive Wearable Interaction Devices. In *Proceedings Fifth International Symposium on Wearable Computers*. 21–27. <https://doi.org/10.1109/ISWC.2001.962092>
- [48] Irvin Rock and Jack Victor. 1964. Vision and Touch: An Experimentally Created Conflict between the Two Senses. *Science* 143, 3606 (1964), 594–596.
- [49] Antti Sand, Ismo Rakkolainen, Poika Isokoski, Jari Kangas, Roope Raisamo, and Karri Palovuori. 2015. Head-mounted display with mid-air tactile feedback. In *Proceedings of the 21st ACM Symposium on Virtual Reality Software and Technology*. 51–58.
- [50] Munehiko Sato, Ivan Poupyrev, and Chris Harrison. 2012. Touché: Enhancing Touch Interaction on Humans, Screens, Liquids, and Everyday Objects. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, Texas, USA) (CHI '12). Association for Computing Machinery, New York, NY, USA, 483–492. <https://doi.org/10.1145/2207676.2207743>
- [51] Adalberto L. Simeone, Eduardo Velloso, and Hans Gellersen. 2015. Substitutional Reality: Using the Physical Environment to Design Virtual Reality Experiences. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) (CHI '15). Association for Computing Machinery, New York, NY, USA, 3307–3316. <https://doi.org/10.1145/2702123.2702389>
- [52] 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>
- [53] Rajinder Sodhi, Ivan Poupyrev, Matthew Glisson, and Ali Israr. 2013. AIREAL: interactive tactile experiences in free air. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 1–10.
- [54] Barry E Stein and M Alex Meredith. 1993. *The merging of the senses*. The MIT Press.
- [55] Steven Strachan, Roderick Murray-Smith, and Sile O'Modhrain. 2007. BodySpace: Inferring Body Pose for Natural Control of a Music Player. In *CHI '07 Extended Abstracts on Human Factors in Computing Systems* (San Jose, CA, USA) (CHI EA '07). Association for Computing Machinery, New York, NY, USA, 2001–2006. <https://doi.org/10.1145/1240866.1240939>
- [56] Evan Stranick, Christian Holz, Eyal Ofek, Mike Sinclair, and Hrvoje Benko. 2018. Haptic Links: Bimanual Haptics for Virtual Reality Using Variable Stiffness Actuation. 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/10.1145/3173574.3174218>
- [57] Paul Strohmeyer, Juan Pablo Carrascal, and Kasper Hornbæk. 2016. What Can Doodles on the Arm Teach Us about On-Body Interaction?. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (San Jose, California, USA) (CHI EA '16). Association for Computing Machinery, New York, NY, USA, 2726–2735. <https://doi.org/10.1145/2851581.2892544>
- [58] Yuriko Suzuki and Minoru Kobayashi. 2005. Air jet driven force feedback in virtual reality. *IEEE computer graphics and applications* 25, 1 (2005), 44–47.
- [59] Emi Tamaki, Takashi Miyak, and Jun Rekimoto. 2010. BrainyHand: A Wearable Computing Device without HMD and It's Interaction Techniques. In *Proceedings of the International Conference on Advanced Visual Interfaces* (Roma, Italy) (AVI '10). Association for Computing Machinery, New York, NY, USA, 387–388. <https://doi.org/10.1145/1842993.1843070>
- [60] Hsin-Ruey Tsai and Jun Rekimoto. 2018. ElasticVR: Providing Multi-Level Active and Passive Force Feedback in Virtual Reality Using Elasticity. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI EA '18). Association for Computing Machinery, New York, NY, USA, 1–4. <https://doi.org/10.1145/3170427.3186540>
- [61] Michael White, James Gain, Ulysse Vimont, and Daniel Lochner. 2019. The Case for Haptic Props: Shape, Weight and Vibro-Tactile Feedback. In *Motion, Interaction and Games* (Newcastle upon Tyne, United Kingdom) (MIG '19). Association for Computing Machinery, New York, NY, USA, Article 7, 10 pages. <https://doi.org/10.1145/3359566.3360058>
- [62] Jacob O Wobbrock, Leah Findlater, Darren Gergle, and James J Higgins. 2011. The aligned rank transform for nonparametric factorial analyses using only anova procedures. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 143–146.
- [63] Jackie (Junrui) 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* (Berlin, Germany) (UIST '18). Association for Computing Machinery, New York, NY, USA, 889–899. <https://doi.org/10.1145/3242587.3242643>
- [64] André Zenner and Antonio Krüger. 2019. Estimating Detection Thresholds for Desktop-Scale Hand Redirection in Virtual Reality. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 47–55. <https://doi.org/10.1109/VR.2019.8798143>
- [65] Yang Zhang, Wolf Kienzle, Yanjun Ma, Shiu S. Ng, Hrvoje Benko, and Chris Harrison. 2019. ActiTouch: Robust Touch Detection for On-Skin AR/VR Interfaces. 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, 1151–1159. <https://doi.org/10.1145/3332165.3347869>
- [66] Yiwei Zhao and Sean Follmer. 2018. A Functional Optimization Based Approach for Continuous 3D Retargeted Touch of Arbitrary, Complex Boundaries in Haptic 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, 1–12. <https://doi.org/10.1145/3173574.3174118>