This is the authors' version of the article. See https://doi.org/10.1109/VR.2019.8798143 for the final version of record in IEEE Xplore.

Estimating Detection Thresholds for Desktop-Scale Hand Redirection in Virtual Reality

André Zenner* Antonio Krüger†

German Research Center for Artificial Intelligence (DFKI)
Saarland Informatics Campus, Saarbrücken, Germany

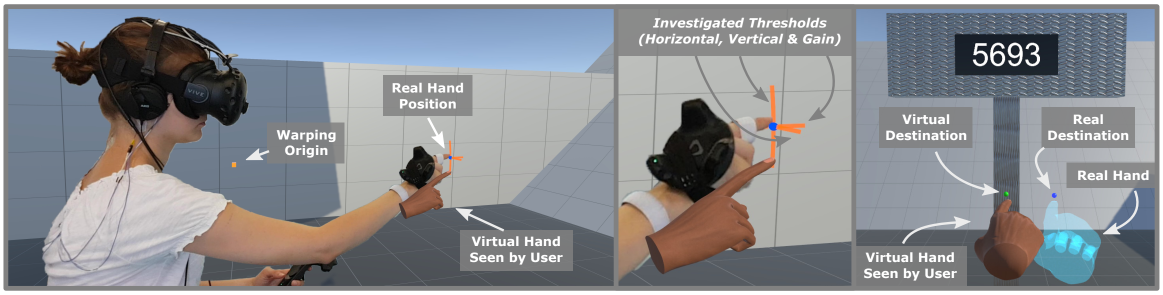

Figure 1: Left & Center: Hand redirection in VR. A warping algorithm displaces the rendered virtual hand from the real hand. According to our results, the virtual fingertip can unnoticeably be displaced to any point along the orange lines. Right: The *Vis-DT* condition in our experiment. The user reaches for the green target sphere with the displaced virtual hand while reading out the displayed number. The real hand location is added for illustration.

Abstract

Virtual reality (VR) interaction techniques like haptic retargeting offset the user's rendered virtual hand from the real hand location to redirect the user's physical hand movement. This paper explores the order of magnitude of hand redirection that can be applied without the user noticing it. By deriving lower-bound estimates of detection thresholds, we quantify the range of unnoticeable redirection for the three basic redirection dimensions, horizontal, vertical and gain-based hand warping. In a two-alternative forced choice (2AFC) experiment, we individually explore these three hand warping dimensions each in three different scenarios: a very conservative scenario without any distraction and two conservative but more realistic scenarios that distract users from the redirection. Additionally, we combine the results of all scenarios to derive robust recommendations for each redirection technique. Our results indicate that within a certain range, desktop-scale VR hand redirection can go unnoticed by the user, but that this range is narrow. The findings show that the virtual hand can be unnoticeably displaced horizontally or vertically by up to 4.5° in either direction, respectively. This allows for a range of ca. 9°, in which users cannot reliably detect applied redirection. For our gain-based hand redirection technique, we found that gain factors between $g = 0.88$ and $g = 1.07$ can go unnoticed, which corresponds to a user grasping up to 13.75% further or up to 6.18% less far than in virtual space. Our findings are of value for the development of VR applications that aim to redirect users in an undetectable manner, such as for haptic retargeting.

Keywords: Virtual reality, hand redirection, detection thresholds.

Index Terms: H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems—Artificial, augmented, and virtual realities; H.5.2 [Information Interfaces and Presentation]: User Interfaces—Evaluation/methodology, Interaction styles

*e-mail: andre.zenner@dfki.de
†e-mail: krueger@dfki.de

1 Introduction

Virtual reality (VR) allows users to experience virtual worlds in immersive ways. Besides visual and auditory feedback, systems today also focus on haptics to achieve truly immersive VR. In this context, we also perceive the posture, speed and position of our own limbs in proprioception. Here, our brain merges information gained from the visual, auditory and haptic channels and what we see can influence what we perceive with other senses, for example haptically. This strong influence of vision is called visual dominance [9].

Novel VR interaction concepts like *redirected touching* [17] and *haptic retargeting* [4,7] are based on hand redirection. These techniques exploit visual dominance to change how we interact in virtual spaces and are commonly used in the context of haptic feedback for immersive virtual environments (IVEs). In systems that use passive haptic feedback [12–14], physical objects, typically low-fidelity props (also called proxies), represent virtual objects and provide haptic feedback when touched by the user. While computationally cheap, simple and typically low-cost, the approach of passive haptics is inherently inflexible and does not scale well. In a naive implementation, for example, each modification of a virtual object would require an adaptation of the corresponding physical prop. This motivates research on more advanced techniques that compensate for the drawbacks of passive haptics like the lack of proxy generality, scalability and reusability. Besides mixed haptic approaches that combine active and passive haptics [3,23,32,33], hand redirection promises to address some of these challenges by leveraging the visual dominance effect.

Hand redirection concepts refrain from a 1-to-1 mapping from real to virtual space. Instead, the virtual hand seen by the user in the IVE is displaced from the real hand position. Perceiving the displaced hand visually, the user automatically compensates for the displacement and adapts the path of the real hand accordingly. Combined with passive haptic feedback, this technique can greatly enhance the generality and reusability of physical proxies. The real hand can be redirected to touch haptic props at different locations than the virtual hand seen in the IVE. However, important research questions regarding detectability of hand redirection are still to be investigated [17]. To employ hand redirection in interactive VR applications, it is essential to know about detection thresholds. This allows design of appropriate virtual and physical environments.

While informal investigations [17] and recent results [7] brought first insights about how much hand redirection is *tolerated* by users, which could mean users notice the redirection, we determine the order of magnitude at which hand redirection *can go unnoticed* when reaching for a virtual target under redirection. We conduct psychophysical experiments to derive lower-bound estimates for desktop-scale hand redirection thresholds in VR. Our investigation is motivated by the haptic retargeting use case [4, 7]. It focuses on the mid-air movement of a user's index finger approaching a retargeted location. In a VR application this could be, for example, a retargeted haptic prop. For this, we investigate horizontal, vertical and gain-based redirection. Moreover, we investigate three different conservative interaction scenarios differing in the user's distraction from the redirection. We also combine the results of all three scenarios to derive general recommendations. Our results are of value to researchers and developers of applications that make use of hand redirection, as they provide information crucial for the development of unnoticeable retargeting.

2 Related Work

This section reviews work on hand redirection and detection threshold estimation, and puts our investigation into context.

2.1 Hand Redirection and Haptic Retargeting

Hand redirection is based on non-isomorphic mappings from user input to displayed output. Such concepts have been studied in the context of 3D user interfaces [27] and many 3D interaction techniques decoupling real and virtual hand movement are known (e.g. [2, 26]). Modifications of the control-display ratio are also used for pseudo-haptics [8].

This paper investigates hand redirection as it is used for techniques like *redirected touching* [15–20] or *haptic retargeting* [4, 7]. Redirected touching can change how virtual objects and haptic props are perceived when touched [17]. Similarly, haptic retargeting allows for the reuse of a single haptic prop to provide feedback for different, spatially separated virtual objects [4, 11]. To achieve this, both techniques refrain from a 1-to-1 mapping of real and virtual space and instead warp the virtual space. When warping a complete scene, object shapes can be distorted [15, 17]. Applying the warp only to the movement of a tracked object (e.g. the hand) results in the virtual object moving on paths offset from the physical trajectory [16, 17]. Due to visual dominance [9], the displacement of the virtual hand leads to a compensation by the user which adapts and redirects the movement of the *real* hand. As a result, the user's real hand can end up at locations different from the virtual hand. Combined with passive haptics, this can be used to control how and where users touch haptic props [5, 17]. The concept is related to redirected walking [28], where visual modifications trick users into walking along physical paths different from the virtual paths. Kohli investigated how redirected touching can be used in combination with passive haptics [17], how warped spaces can enhance the generality of a passive proxy [15, 16, 18] and how they can be used to train movements [17, 20]. Spillmann et al. investigated how space warping can be used in an arthroscopy surgical simulator [30]. Azmandian et al. demonstrated how a single proxy can provide haptic feedback for multiple spatially separated virtual objects using hand redirection for haptic retargeting [4]. Cheng et al. investigated the combination of hand redirection and a generalized low-fidelity prop [7].

By reviewing existing work on redirected touching and haptic retargeting, we identified 3 main conceptual approaches: (1) distortions of the complete virtual space or parts of it [17, 24, 30, 34]; (2) rotating/translating the IVE (world warping) [4, 19] and (3) virtual-real hand offsets (body warping) [4, 7, 11].

We investigate body warping based on algorithms which define how the virtual hand is offset from the tracked physical hand. Cheng et al. [7] assessed tolerance thresholds for hand redirection using such techniques through questions on a rating scale and report that users tolerate virtual hand deviations of up to 40°. However, as tolerance is independent of detectability, users might tolerate certain redirections while at the same time being fully aware of them. Thus, it remains to be explored how much redirection can be applied *without the user even noticing it*. As described by Kohli, there is a need for formal investigation of detection thresholds for hand redirection in VR [17]. In this paper we derive corresponding estimates.

2.2 Estimating Detection Thresholds

Only a few works investigate the detection of visual-proprioceptive discrepancies in VR; even fewer concentrate on intentional hand redirection. Burns et al. [6] derived detection thresholds for gradually growing angular offsets of the real and virtual hand ($\approx 19.1°$ or $\approx 19cm$). However, we believe these results can only serve in a limited way as general lower-bound thresholds for hand redirection because a partial method-of-limits with only an ascending test series was used, and the experiment was embedded in a game probably affecting users' attention. This is in contrast to our more formal investigation employing a different methodology with trial-wise constant stimuli and different user distraction. Our experiment investigates fixed angular offsets and allows for very conservative lower-bound estimates as users do not play a game but focus primarily on detecting hand offsets. Lee et al. [21] investigated tolerance thresholds for finger tracking errors. The authors derived just noticeable differences (JNDs) of visual-proprioceptive conflict to derive requirements for finger tracking systems. Compared to Burns et al.'s results [6], the authors report a much lower JND value of $\approx 5.2cm$ and found tactile feedback at the fingertips to increase JNDs. It is important to note, however, that both the real fingertip position and a randomly displaced virtual position were shown to the user at the same time. Moreover, Lee et al. did not use an immersive virtual hand representation, as is common in VR applications, but abstract spheres indicating the fingertip positions. To study intended hand redirection, this is an unrealistic condition, and thus our approach displays only one hand at a time, using a virtual human hand model. Closely related to our investigation is recent work by Abtahi and Follmer [1]. While methodologically similar to our investigation, the tested conditions are crucially different. Abtahi and Follmer investigated fingertip movements along the edges of a physical proxy. As a result, the derived thresholds (horizontal remapping $\approx 49.5°$, horizontal scaling $\approx 1.9x$, vertical scaling $\approx 3.2x$) relate to a combination of visual-proprioceptive conflict and the continuous haptic signal felt by the user while moving the hand over the proxy. In contrast, our study investigates mid-air movements as the hand approaches a virtual target without haptic feedback at the hand.

In summary, we are not aware of any previous study that (1) derives *conservative lower-bound estimates* for hand redirection in VR and (2) investigates the movement of a hand approaching a retargeted location in *mid-air without continuous haptic feedback*, while (3) *rendering only the displaced virtual human hand*, as is common practice in VR applications. This work aims to fill this gap.

Several methodologies to estimate detection thresholds exist, and it was shown for redirected walking that estimated thresholds significantly vary with the methodology used [10]. Our experimental design is an adaptation of the established estimation methodology used by Steinicke et al. [31] to estimate detection thresholds for redirected walking. The authors conducted a two-alternative forced choice (2AFC) experiment (or pseudo-2AFC task [10]) and subjects repeated a certain movement in the IVE while the virtual head movement was manipulated to different extents. After each movement, users chose between two possible answers, for example: "Was the virtual movement *smaller* or *greater* than the physical move-

This is the authors' version of the article. See https://doi.org/10.1109/VR.2019.8798143 for the final version of record in IEEE Xplore.

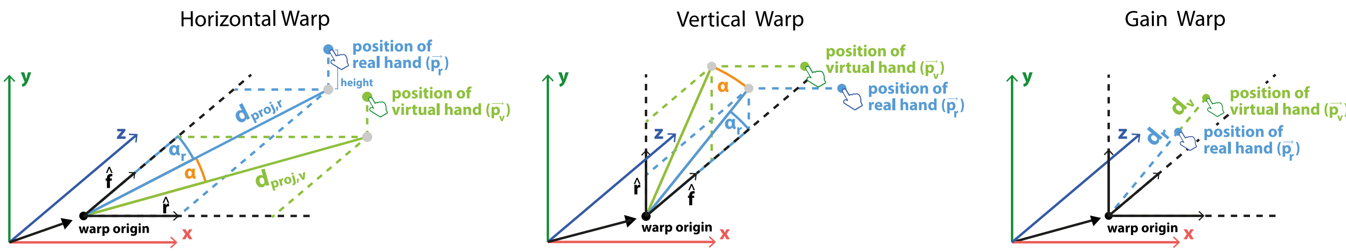

Figure 2: Illustration of the three investigated hand redirection types. Left: horizontal warping. Center: vertical warping. Right: gain warping

Algorithm 1 Pseudocode of the Rotational Warp Algorithm

Input: real hand position $\vec{p}_r$, warp origin $\vec{o}$, unit forward vector $\hat{f}$, unit redirection vector $\hat{r}$, redirection angle α
Output: warped virtual hand position $\vec{p}_v$

procedure ROTATIONALWARP($\vec{p}_r, \vec{o}, \hat{f}, \hat{r}, \alpha$)
  $\hat{h} = \hat{f} \times \hat{r}$ ▷ compute unit height vector
  $height = (\vec{p}_r - \vec{o}) \cdot \hat{h}$ ▷ save height
  $\vec{p_{proj}} = \vec{p}_r - height \cdot \hat{h}$ ▷ project on redirection plane
  $\vec{d_{proj,r}} = \vec{p_{proj}} - \vec{o}$ ▷ unwarped offset in plane
  $\alpha_r = atan2(\vec{d_{proj,r}} \cdot \hat{r}, \vec{d_{proj,r}} \cdot \hat{f})$ ▷ angle rel. to $\hat{f}$ & $\vec{o}$
  $\alpha_v = \alpha_r + \alpha$ ▷ adding angular offset
  $\vec{d_{proj,v}} = sin(\alpha_v) \cdot |\vec{d_{proj,r}}| \cdot \hat{r} + cos(\alpha_v) \cdot |\vec{d_{proj,r}}| \cdot \hat{f}$ ▷ warped offset in plane
  $\vec{p}_v = \vec{o} + \vec{d_{proj,v}} + height \cdot \hat{h}$ ▷ final warped position
  return $\vec{p}_v$
end procedure

ment?" [31]. Participant responses were then used to fit a psychometric function modeling the discrimination performance, allowing the derivation of detection thresholds.

3 The Investigated Hand Redirection Techniques

Our study focuses on body warping as it is easy to implement, versatile and characterized by a low complexity. Body warping approaches typically offset the virtual hand rendered to the user from the real hand position retrieved from the tracking system, with warping algorithms specifying the amount of 3D offset applied. Commonly, the offset is gradually increased as the user approaches a retargeted destination, implemented, for example, through interpolation methods [7]. To study detection thresholds in a controlled fashion, we split up 3D redirection in 3 intuitive dimensions to investigate them separately and assume a desktop VR setting, with the user being seated and interacting in the limited space in front. The following sections describe the corresponding warping algorithms that proved well suited for our study, as they allow for horizontal or vertical angular redirection and gain displacement.

3.1 Horizontal Hand Displacement

The first redirection type horizontally offsets the virtual hand by a warp angle α as the real hand moves away from a warp origin. To compute the warped position, the hand is projected on a horizontal plane, its angle relative to a forward direction and the warp origin is incremented by α and then projected back in 3D space.

For this, we define a general rotational warp algorithm allowing for displacements in arbitrary planes defined by a unit forward vector $\hat{f}$ and an orthogonal unit redirection vector $\hat{r}$ indicating the direction of positive displacement. Further inputs are the location of the warp origin $\vec{o}$ (e.g. virtual hand position when the warp starts) and the redirection angle α. Algorithm 1 shows the pseudocode of the rotational warp algorithm. For horizontal displacement as

Algorithm 2 Pseudocode of the Gain Warp Algorithm

Input: real hand position $\vec{p}_r$, warp origin $\vec{o}$, gain factor g
Output: warped virtual hand position $\vec{p}_v$

procedure GAINWARP($\vec{p}_r, \vec{o}, g$)
  $\vec{d}_r = \vec{p}_r - \vec{o}$ ▷ unwarped offset from origin
  $\vec{d}_v = g \cdot \vec{d}_r$ ▷ warped offset from origin
  $\vec{p}_v = \vec{o} + \vec{d}_v$ ▷ final warped position
  return $\vec{p}_v$
end procedure

investigated here, the algorithm is instantiated with $\hat{f} = +\vec{z}$ (z-axis) and $\hat{r} = +\vec{x}$ (x-axis). The effect of the horizontal warp algorithm is sketched on the left in Fig. 2 and in Fig. 3 (a).

3.2 Vertical Hand Displacement

Vertical redirection offsets the virtual hand up or down as the real hand moves away from the warp origin. For this, Algorithm 1 is instantiated with $\hat{f} = +\vec{z}$ (z-axis) and $\hat{r} = +\vec{y}$ (y-axis). The center of Fig. 2 and Fig. 3 (b) show the effect of the vertical displacement.

3.3 Gain-Based Hand Movement

The third algorithm scales the distance of the hand from warp origin. It computes the distance vector $\vec{d}_r$ to the unwarped position of the real hand and applies a gain factor g, effectively decreasing (if $0 < g < 1$) or increasing (if $g > 1$) the speed of the hand moving away from $\vec{o}$. Algorithm 2 sketches the pseudocode. The effect of the warp is illustrated on the right in Fig. 2, as well as in Fig. 3 (c).

4 Experiment

We conducted an experiment investigating the 3 individual redirection dimensions, each in 3 different scenarios. In the experiment, participants were immersed in a simple IVE with their hand tracked to interact therein. In 9 conditions, they repeatedly performed a simple interaction with different warps applied, and had to state the direction of the hand displacement. From the results, we derived how much redirection could go unnoticed. The experiment was approved by the ethical review board of our faculty.

4.1 Introduction

The experiment has a 2AFC (or pseudo-2AFC [10]) design in which participants are repeatedly exposed to different amounts of hand redirection while performing a pointing gesture with their hand in mid-air in front of themselves. In all conditions, the main task of the participants was to determine the direction of the hand offset and they were instructed to fully concentrate on that. Moreover, they were informed about the investigated redirection technique, and knew how it worked. Thus, all scenarios are classified as very conservative, representing a worst-case scenario for unnoticed redirection. To derive meaningful detection limits, we decided to not only investigate low-level perceptibility in a single very conservative

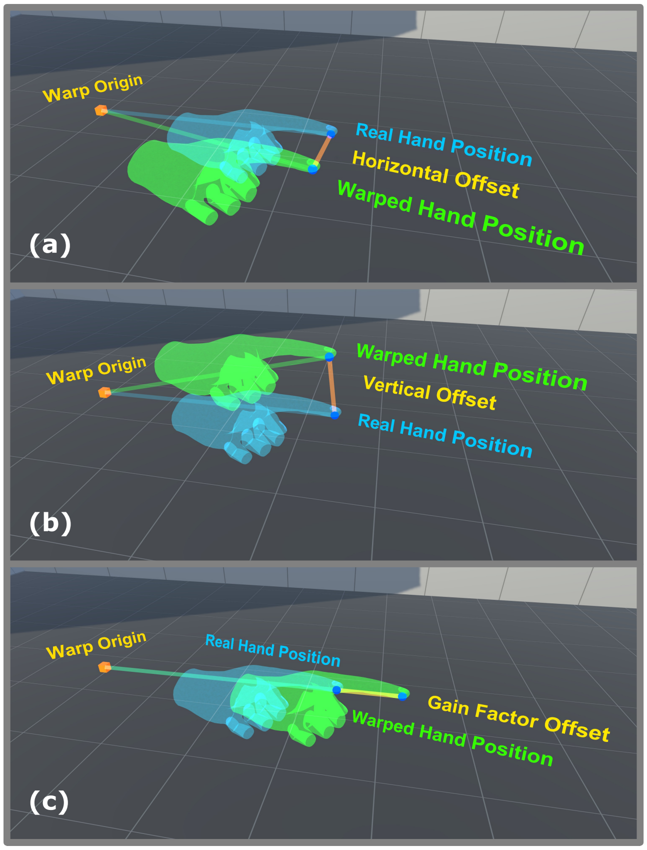

Figure 3: Visualization of the three investigated redirection techniques: (a) horizontal warping, (b) vertical warping and (c) gain-based warping. The location of the warp origin is shown as a yellow dot, the displaced hand is visualized in green and the real hand location is shown in blue. In the experiments, users only saw a single hand displayed at the warped hand location, rendered with a realistic texture.

scenario where participants only focused on detecting the offset; instead, we additionally considered noticeability of redirection in two further scenarios employing secondary tasks and visual distraction, as well as auditory and vibrotactile distractions, respectively. The experiment is described in detail in the following sections.

4.1.1 Scenario 1: No Distraction

The first and most conservative scenario did not distract the participants from the main task at all. They did not hear anything, they had no second task and no vibrotactile cues were present. Thus, the scenario is suited well to derive very conservative lower bounds for the detection thresholds, but at the same time is less realistic.

4.1.2 Scenario 2: Audio & Vibration Distraction

The second scenario better represents realistic application scenarios. As distractions might influence detection thresholds, here, we distract users in a substantial way using a combination of two additional modalities likely used in VR: auditory and vibrotactile feedback. Spatial sound and 4 head-mounted vibration cells let participants experience a virtual bee orbiting their head during redirection. We used vibrotactile actuation at the head as we expected this to yield strong distraction and believe it will be included in future generations of head-mounted displays (HMDs). Our vibration cell placement was based on the results of Myles and Kalb which showed that the forehead, occipital and temple regions are suited best for vibrotactile cues by being most sensitive to vibration [25].

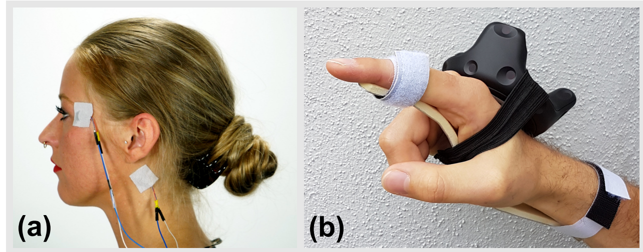

Figure 4: Hardware setup: (a) the vibration cells placed at the temple regions and the lower occipital region close to the neck; (b) the finger splint and the Vive Tracker used in our hand tracking solution.

4.1.3 Scenario 3: Visual & Dual-Task Distraction

Besides distraction by special effects, distraction can also ensue due to increased cognitive load. Thus, the third scenario included both a second task to solve *in parallel* to the interaction and a visual distraction forcing users to look away from the virtual hand at least once. For this task, we were inspired by a common application scenario for redirected touching, the simulation of cockpit procedures [17]. Here, users might be required to look away from the hand in order to read a number displayed in the cockpit. Thus, in our scenario, participants had to look at a number displayed on a virtual panel in front of them to read it out loud. Similar to Scenario 2, this represents a more realistic use case.

4.2 Participants

12 participants volunteered to take part in the experiment (6 f, 6 m, avg. 28 years, between 20 years and 61 years old). 4 participants wore glasses or contact lenses, but all participants confirmed that they have normal or corrected-to-normal vision, that they do not suffer from hearing impairments and that their sensation of vibrotactile feedback is not in any form negatively affected. 11 participants were right-handed and 1 participant was left-handed. The participants rated how regularly they play 3D video games on a scale from 1 (= never) to 7 (= regularly). Here, different frequencies were present, with answers between 1 and 7 ($M = 2.33$, $SD = 2.15$). Participants also stated how often they use VR technology on the same scale. We obtained answers between 1 and 5 ($M = 2.09$, $SD = 1.38$). A third question on the same rating scale asked how often participants perform precise handicrafts. Here, all different ratings were obtained ($M = 3.75$, $SD = 2.05$).

4.3 Apparatus

The experiment was conducted in our lab. We used an HTC Vive HMD with two hand-held controllers and a separate Vive Tracker. One controller was used by the participants to record the 2AFC answers; the other one was used for initial finger calibration. The IVE was developed with Unity3D 5.6 and represents a small, plain terrace-like scene as shown in Fig. 1. It was intentionally kept simple to prevent uncontrolled distraction. A desktop computer executed the VR application and logged the 2AFC responses. The experimenter logged all additional answers in a separate spreadsheet.

Headphones were used for auditory feedback and 4 small vibration cells (Adafruit Vibrating Mini Motor Disc #1201; 10mm x 2.7mm) controlled by a WeMos D1 mini microcontroller delivered tactile feedback. They were fixed with medical tape at the two temple regions and at the lower occipital region close to the neck [25] as depicted in Fig. 4 (a). The VR application controlled the vibration by communicating wirelessly with the microcontroller. In informal previous tests, we tested several vibration patterns and strengths to find suitable parameters for the experiment. We fixed the vibration strength to a comfortable but still readily noticeable level. When called by the VR application, the microcontroller repeatedly activated a random set of the 4 cells for a random duration between 150*ms* and 500*ms* until the command to stop was received.

We used the Vive Tracker for hand tracking, and during calibration attached the tracker on the back of the user's hand with a rubber band. We additionally used a finger splint that allowed participants to comfortably maintain a pointing hand position with the index finger pointing forward, as shown in Fig. 4 (b). Besides being comfortable and reducing fatigue, this splint ensured that the real and virtual hand stayed spatially registered at all times as the rigid structure of the splint did not allow participants to move the index finger relative to the tracker on the back of the hand. Thus, the offset from tracker location to fingertip remained fixed. The touch-sensitive trackpad of one of the Vive controllers was used in an initial calibration step to calibrate this offset and to align the virtual hand model with the real hand. We rendered either a female or a male and a right or a left virtual hand to account for the participant's gender and handedness. Both male and female hands had a fixed but realistic size and the male hand was slightly larger than the female hand model.

4.4 Procedure

Each participant was initially briefed with a prepared text explaining the concept of hand redirection and the experiment. The experimenter attached the finger splint, the Vive Tracker on the back of the participant's hand and the 4 vibration cells at the participant's head. Participants sat on a chair wearing headphones and the HMD.

Entering the IVE, participants were introduced to all 3 redirection types in a short training session. To ensure that the participant understood and noticed the applied redirection, we demonstrated the largest redirection used in the tests during training (values determined by informal pre-testing; $\alpha = +/-14°$ for horizontal and vertical redirection, $g = 0.75$ and $g = 1.25$ for gain redirection).

At the same time, the interaction to be performed was practiced: seeing a small green sphere appearing at the start location 30 *cm* beneath and 30 *cm* in front of the head, the participant was supposed to touch the sphere with the virtual index finger. When inside the sphere, redirection was applied with the warp origin set to this start location. We chose the origin to be at this comfortable position in front of the user, as it is just outside the zero-warp zone defined by Cheng et al. [7], and we see this as a representative location for desktop setups. Upon activation of the warp, the sphere relocated to the target position. The distance to the target was likewise chosen to be representative for typical desktop-scale interaction distances. For the horizontal and vertical redirection, this target position was another 40 *cm* away from the start position straight in front of the user. To account for gain factors $g < 1$, the target position, when the gain was applied, was only 30 *cm* from the start location to ensure reachability. Upon relocation of the sphere, participants were to move their hand in order to touch the target sphere with the virtual finger. This required participants to compensate for the redirection warp. To ensure consistency and comparable hand movement speeds, each participant was asked to perform this movement within around 3*s* to 4*s*. Finally, when reaching the target with the virtual hand, a question appeared on the controller held in the second hand. This 2AFC question asked participants to state in which direction the virtual hand was displaced during the movement. Participants had to decide between the answers *right* or *left* for *horizontal* displacement and between *up* or *down* for *vertical* displacement. In the *gain* condition, participants had to state whether the virtual hand moved *faster* or *slower* than the real hand. Answers were logged by pressing the corresponding controller button. While equivalent to the experiment procedure, no data was recorded during training. In a final training session, participants could practice reading out numbers displayed in front of them to become acquainted with the dual task.

Once the training phase was completed, the actual experiment started. To test the 3 introduced types of redirection (in the following abbreviated *Horiz* for horizontal displacement, *Vert* for vertical displacement and *Gain* for the gain warp) in the 3 introduced scenarios (abbreviated *None* for no distraction, *Audio-Vib* for audio/vibration distraction and *Vis-DT* for visual/dual-task distraction), each participant performed 9 runs in sequence. In each run, the interaction was performed repeatedly with different redirection parameters α or g, and after each interaction, a 2AFC question was answered.

In *Vis-DT* runs, participants had to read out a random 4-digit positive integer during the movement, appearing on a display in front when the hand progressed 20% along the way from the start to the target sphere, as shown on the right in Fig. 1. This forced them to look away from the hand at least once. We logged the correctness of the read numbers.

In contrast, the distracting virtual bee in *Audio-Vib* did not require participants to take any specific action. They just had to try to focus on the main task of determining the displacement direction.

After completing a run, participants denoted their agreement with 9 concluding statements on a Likert scale from 1 (= strongly disagree) to 7 (= strongly agree), displayed in the IVE. These statements assessed their subjective impressions of interacting under redirection in the experienced condition. After completion of all 9 runs, the participants filled out a SUS presence questionnaire [29] and a final post-study questionnaire. The duration of the experiment was ca. 90 *min*, including introduction, calibration, training, all experimental runs, final questionnaires and debriefing.

4.5 Design

The 2AFC experiment is a within-subjects study. We distinguish 2 independent variables: 1) hand redirection type and 2) user distraction type, with 3 different implementations each. Using a full-factorial design, we investigated 9 conditions.

We had 11 dependent variables: (1) the perceived offset of the virtual hand as a 2AFC answer for each sample, (2) – (10) the 9 subjective measures assessed as ratings on the Likert scale after each run and (11) the interaction time measured to reach the target. We studied the 3 hand redirection techniques, *Horiz*, *Vert* and *Gain*, in the 3 scenarios *None*, *Audio-Vib* and *Vis-DT*. The order of the conditions was counterbalanced across participants. We used a Latin square counterbalancing over the 3 distraction types and for each distraction type, we used an additional Latin square counterbalancing over the 3 redirection types. The resulting counterbalancing of size $n = 6$ was completed exactly twice with 12 participants. In each of the 3 *Gain* conditions, the interaction was performed 22 times resulting in 22 samples (2AFC answers). Using a step size of 0.05, we tested all 11 values between $g = 0.75$ and $g = 1.25$ (inclusive) twice in a randomized order. Similarly, each condition applying *Horiz* or *Vert* redirection collected 30 samples testing all redirection angles between $\alpha = -14°$ and $\alpha = 14°$ (inclusive) in steps of 2° twice in a randomized order. Thus, each participant contributed $6 \cdot 30 + 3 \cdot 22 = 246$ samples. With 12 participants, we obtained $12 \cdot 246 = 2952$ samples for the 2AFC analysis in total.

4.6 Results

First, we summarize our estimates for the hand redirection detection thresholds of all 9 individual conditions. We also derive overall thresholds for the 3 redirection techniques from the pooled samples of all 3 tested scenarios. Secondly, we summarize the results of the subjective responses.

4.6.1 Detection Thresholds for Hand Redirection

To analyze the performance of our participants in discriminating the hand offset direction, we used the obtained samples to fit a psychometric function modeling the discrimination performance over the applied redirection. Analogous to Steinicke et al.'s [31] derivation of redirected walking thresholds, we used a psychometric sigmoid function as its shape is a good approximation to model

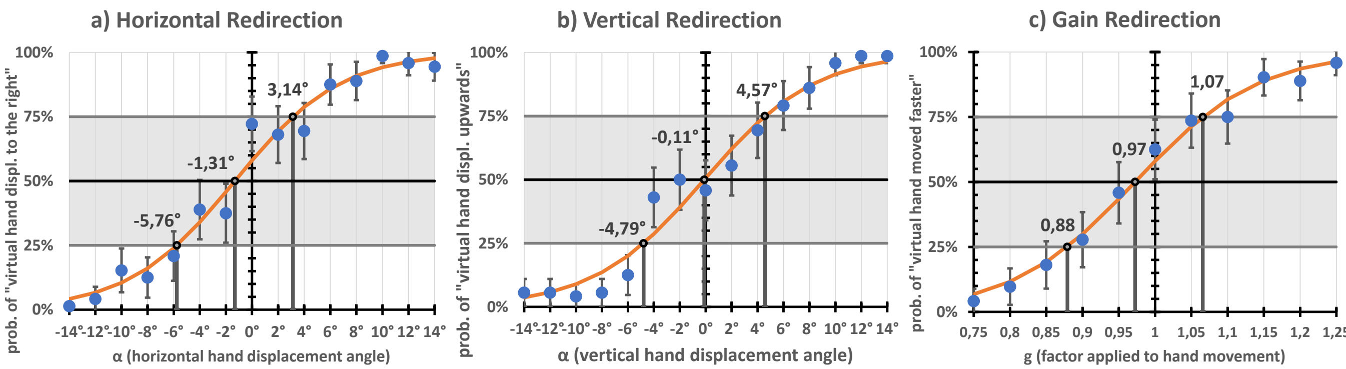

Figure 5: Pooled results of all 3 distraction scenarios for each redirection type. The charts additionally show the 95% confidence intervals, the fitted psychometric function f, the PSE and the derived detection thresholds for (a) horizontal redirection, (b) vertical redirection and (c) gain redirection.

Table 1: The Derived Detection Thresholds for VR Hand Redirection.

		None	*Audio-Vib*	*Vis-DT*	**Mixed**
Horiz	**Right**	+3.81°	+2.26°	+2.94°	+3.14°
	PSE	−0.28°	−1.67°	−2.11°	−1.31°
	Left	−4.38°	−5.60°	−7.17°	−5.76°
	Range	8.19°	7.86°	10.11°	8.90°
Vert	**Up**	+4.48°	+4.55°	+4.65°	+4.57°
	PSE	+0.04°	−0.09°	−0.31°	−0.11°
	Down	−4.40°	−4.74°	−5.28°	−4.79°
	Range	8.88°	9.29°	9.93°	9.36°
Gain	**Faster**	1.07	1.05	1.07	1.07
	PSE	0.98	0.96	0.98	0.97
	Slower	0.88	0.86	0.89	0.88
	Reach +	+13.38%	+15.71%	+11.92%	+13.75%
	Reach –	−6.50%	−5.11%	−6.63%	−6.18%

human response. Plotting the overall probability of our participants answering "The virtual hand was offset to the *right*/*up*/was *faster*" against the applied amount of virtual hand offset yields an s-shaped distribution of our samples. To derive detection thresholds we fitted the sigmoid function $f(x) = \frac{1}{1+e^{-\frac{x-a}{b}}}$ to our sample distribution by optimizing the parameters a and b. We computed the point of subjective equality (**PSE**) indicating when the virtual movement is perceived as equal to the physical movement. This is where f intersects the 50% probability. Additionally, we computed the discrimination thresholds, i.e. where f intersects the probability halfway between the random guessing level and the correct answer. For the upper redirection limit (**Right** for *Horiz*, **Up** for *Vert* and **Faster** for *Gain*), this is the redirection at which f intersects the 75% probability. Complementarily, for the lower redirection limit (**Left** for *Horiz*, **Down** for *Vert* and **Slower** for *Gain*), this is where the 25% probability is intersected. The range in between these two amounts is the range of redirection that can go unnoticed, as users could not with certainty tell about the redirection. With the results of the *Gain* condition, we also computed the range within which the user's *real* hand can be unnoticeably redirected to reach further (**Reach +**) or less far (**Reach –**) than the virtual hand. All these results are summarized in Table 1.

We additionally used the samples of all 3 scenarios to derive even more robust estimates of the detection thresholds. As the scenarios used different feedback channels for the distractions that can thus be regarded as orthogonal to each other, we mixed the samples of *None*, *Audio-Vib* and *Vis-DT*. We thereby derived more realistic but still conservative estimates while profiting from the tripled amount of samples per tested redirection. The resulting plots for this mixed scenario, the fitted function f and the derived thresholds, are depicted in Fig. 5 and summarized in Table 1 (column *Mixed*).

4.6.2 Distraction & Subjective Impressions

For each of the 9 Likert scale ratings assessed after each run, and applying a significance level of $\alpha = .05$, we conducted non-parametric Friedman tests to check for significant differences among conditions. To check for significant effects of the two factors (distraction type and redirection technique), we additionally conducted Friedman tests comparing the results of the 3 distraction scenarios and the 3 redirection types, respectively. Significant differences indicated by the Friedman tests were investigated with pairwise Bonferroni-corrected Wilcoxon signed-rank tests in post-hoc analysis.

To verify our user distraction, we assessed the responses to

- **Distraction** ($M = 3.21$, $SD = 0.75$)[1]: "*I felt distracted from the main task (determine hand offset direction).*"

and found significant differences among the 9 conditions ($\chi^2(8) = 59.455$, $p < .001$), among distraction types ($\chi^2(2) = 17.733$, $p < .001$) and among redirection types ($\chi^2(2) = 7.515$, $p = .023$). Post-hoc analysis (corrected significance level set at $p < .017$) confirmed the distraction of *Audio-Vib* ($M = 3.83$, $SD = 1.61$) ($Z = -2.941$, $p < .001$, $r = 0.60$) and *Vis-DT* ($M = 4.19$, $SD = 0.93$) ($Z = -3.063$, $p < .001$, $r = 0.62$) to be significantly higher than the distraction in our baseline scenario *None* ($M = 1.61$, $SD = 0.71$), as shown in Fig. 6. Concerning redirection types, *Vert* ($M = 3.47$, $SD = 1.00$) was found to be significantly more distracting than *Gain* ($M = 2.97$, $SD = 0.73$) ($Z = -2.588$, $p = .008$, $r = 0.53$).

We did not find any meaningful significant differences in the participants' agreement regarding the following 8 statements:

- **Disturbing Offset** ($M = 3.41$, $SD = 1.14$): "*The fact that the virtual hand representation was offset from the real hand position disturbed me.*"
- **10 Min.** ($M = 4.88$, $SD = 1.24$): "*I would not mind working under these conditions for a short amount of time (ca. 10 minutes).*"
- **2 Hours** ($M = 3.08$, $SD = 1.14$): "*I would not mind working under these conditions for a longer amount of time (ca. 2 hours).*"
- **Physical Exertion** ($M = 2.78$, $SD = 1.39$): "*The interaction was physically demanding.*"
- **Mental Exertion** ($M = 3.75$, $SD = 0.90$): "*The interaction was mentally demanding and I had to concentrate a lot.*"
- **Body Ownership** ($M = 4.78$, $SD = 1.04$): "*I had the feeling of interacting with my own hand in the virtual environment.*"
- **Hand Control** ($M = 4.88$, $SD = 1.20$): "*I had full control over the movements of the virtual hand at all times.*"

[1] The single M and SD values reported here for each measure are of the *overall* ratings, i.e. participant-wise averages over all 9 conditions.

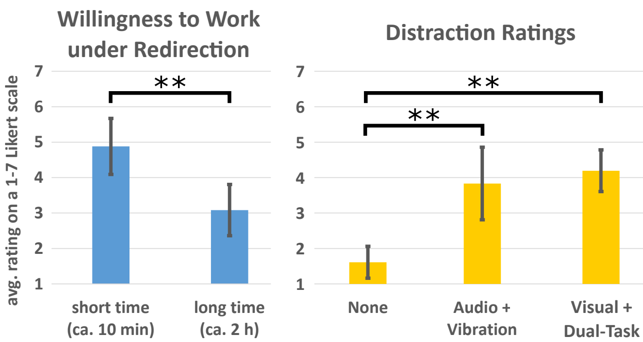

Figure 6: Left: Willingness to work for a short time (**10 Min.**) and a long time (**2 Hours**) with hand redirection. Right: Perceived **Distraction** in our scenarios. Error bars depict 95% confidence intervals. Brackets indicate pairwise significant differences considering the normal p-value for the left chart and the Bonferroni-adjusted p^*-value for the right chart ($\{p, p^*\} < .05$ (*); $\{p, p^*\} < .01$ (**)).

- **Sickness** ($M = 1.29$, $SD = 0.63$): "*When interacting, I felt uncomfortable (e.g. nausea, dizziness).*"

We further conducted a Wilcoxon signed-rank test comparing the overall average willingness to work for a short period of time (**10 Min.**) ($M = 4.88$, $SD = 1.24$) against the willingness to work for a longer period of time (**2 Hours**) ($M = 3.08$, $SD = 1.14$) in a warped space and found the difference to be significant ($Z = -2.983$, $p = .001$, $r = 0.61$), as shown in Fig. 6. Regarding interaction **Times**, we requested participants to complete the movement in $3s$ to $4s$. The obtained timing measurements ($M = 3172ms$, $SD = 599ms$) verified comparable hand movement speeds as participants on average managed to perform the hand movement in the desired time. The performance in the dual task, i.e. correctly reading out the displayed number during interaction, was in each of the 3 *Vis-DT* conditions $> 95.5\%$. A Friedman test did not discover a significant difference across redirection techniques.

Post-study SUS presence scores ($M = 1.33$, $SD = 0.89$) ranged from 0 to 3. In a concluding questionnaire, we also asked participants on a scale from 1 (= not at all) to 7 (= I became very sick) for any experienced discomfort during the study. Results ($M = 1.75$, $SD = 1.21$) did not indicate any problems and are in line with participants' debriefing comments. Overall, in verbal debriefing, most participants were enthusiastic about VR hand redirection.

5 Discussion

The results of the experiment provide insights into the order of magnitude at which hand redirection in VR can go unnoticed. The following sections discuss these results and provide a thorough comparison to results of related research.

5.1 Detection Thresholds for Hand Redirection

The primary goal of our investigation was to capture the range of hand redirection that can be applied in VR applications without the user being aware of it. The experiment showed that for both *Horiz* and *Vert*, the range of unnoticed redirection is very similar. As a general reference, we propose the estimates derived from the pooled results of all 3 distraction types, i.e. the thresholds denoted in column *Mixed* in Table 1. Considering the *Horiz* redirection, we believe the PSE bias (increasing to the left from $-0.28°$ to $-2.11°$ with increasing distraction) to be due to the predominance of right-handed participants. To derive recommendations, we thus take into account the derived angles relative to the PSE and conclude: *in desktop-scale hand redirection, the virtual hand can unnoticeably be displaced horizontally or vertically by* $\approx$ $4.5°$ *to the left/downwards or to the right/upwards respectively, covering a range of* $\approx 9°$.

Similar results were obtained for the *Gain* redirection, as here too, no noticeable increase of the redirection range was found with increasing user distraction. Users, however, seemed to detect an accelerated virtual hand better than a decelerated virtual hand. The possible downscaling range within limits ($g = 0.88 \rightarrow -12\%$) is almost double the possible upscaling range ($g = 1.07 \rightarrow +7\%$). From the thresholds, we can derive how the *real* hand movement is affected when redirection is applied. This knowledge is of immediate value for haptic retargeting. Compensating for the discrepancy between real and virtual hand location, users grasp farther when trying to reach a virtual target with deceleration ($g < 1$) applied. Knowing from our results that a factor of $g = 0.88$ is still within limits implies that the user reaches $\frac{1}{0.88} \approx 1.1375$ times the distance to the virtual target. In turn, when accelerating the virtual hand ($g > 1$), the real hand only needs to reach $\frac{1}{1.07} \approx 0.9382$ times the virtual distance. We summarize our recommendations for hand redirection detection thresholds using the investigated *Gain* warp technique as follows: *factors between* $g = 0.88$ *and* $g = 1.07$ *can unnoticeably be applied to the hand distance from the warp origin. This in turn means that the user's real hand can be redirected unnoticeably to grasp up to* 13.75% *further or up to* 6.18% *less far than the virtual hand.*

Based on previous results [6], we anticipated detection performance to decrease rapidly as the distraction increased. Our results (Table 1), however, did not deliver strong evidence for that, although the distraction worked, as proven by the reported **Distraction** ratings (Fig. 6). While we could generate significantly higher distraction than in the *None* scenario, all our scenarios were likely so conservative that the difficulty of detecting offsets did not increase substantially. We believe the significant differences to Burns et al.'s [6] results ($\approx 19.1°$ when users were primed) to stem from differences in methodology, scenario, degree of user distraction and type of offset investigated (gradually growing vs. fixed α).

However, our results support the JND values found by Lee et al. [21] ($\approx 5.2cm$) which simultaneously rendered both the real and the displaced fingertip position as a sphere to derive requirements for finger tracking systems. Assuming the distance of $40cm$ from origin to target tested in our experiment, our estimation of $4.5°$ yields similar thresholds of $\approx 3.1cm$. Comparing our estimates with the results by Abtahi and Follmer [1] (horizontal remapping $\approx 49.5°$, horizontal gain factor of ≈ 1.9, vertical gain factor of ≈ 3.2) yields additional, very interesting insights. We believe the significantly higher detection thresholds recorded in their experiment, compared to our results, to stem from the additional, continuous haptic feedback provided to the user's fingertip during the movement, a conclusion further supported by the results of Lee et al. [21]. Feeling and seeing an edge under the fingerpad in Abtahi and Follmer's experiment probably served as haptic guidance that increased the user's confidence in *not being redirected*. In this regard, we think that Abtahi and Follmer's study and ours complement each other. Considered together, they provide a more complete picture of the roles of visual and haptic feedback in VR hand redirection.

Compared to the *tolerance* ranges for hand redirection identified by previous research ($\approx 40°$) [7], our derived *imperceptibility* ranges are significantly smaller. We consequently conclude that human hand-eye coordination is so good that even small discrepancies can be detected when the assessment of the hand movement only relies on visual feedback (i.e. no haptic guidance is present [1]) and attention is paid.

5.2 Distraction & Subjective Impressions

The results for **Distraction** confirmed our scenario design as **Distraction** was significantly higher in *Audio-Vib* and *Vis-DT* compared to the baseline scenario *None*, even if *Audio-Vib* and *Vis-DT* were likely not distracting enough to show that detection thresholds can be increased by increasing user distraction. However, we still are confident that decreasing the user's attention on redirection (e.g. by an engaging second task or additional stimuli), or not telling the user about applied redirection, can allow for a larger redirection to go unnoticed. This is supported by the comparison of our findings to related research [1, 6, 22] and worthy of further investigation.

Regarding most remaining subjective scores, we found only minor differences between the tested conditions and did not find meaningful effects of the distractions or redirection techniques used. The obtained SUS presence scores verified that the IVE was sufficiently immersive. While most participants stated their willingness to work under redirection for a short amount of time (**10 Min.**), they were significantly less willing to use the experienced redirection for a longer amount of time (**2 Hours**). However, the lack of a control condition without redirection and the fact that many participants were rather inexperienced with VR makes it unclear how much this difference is due to a negative impact of redirection. Future work should compare the results to the general willingness to be in VR. **Body Ownership** ratings were high throughout all conditions, as were **Hand Control** ratings. This indicates that, despite the fixation of the finger, the illusion was appropriate and participants could reasonably identify with the virtual hand model.

Post-study results on discomfort were in line with the **Sickness** ratings assessed throughout the experiment and verified the absence of cyber-sickness problems. Overall, our subjective results emphasize that hand redirection in IVEs can work well in different distraction scenarios.

5.3 Limitations

We conducted the experiment with $N = 12$ participants, which, although at the lower end of the spectrum for psychophysical experiments, is a common participant set size in related research on redirected walking thresholds (cf. $N = 14$ reported in Steinicke et al. [31]). Having each participant contributing 246 samples to the 2AFC analysis in our experiment, and the fact that we did not experience large variations in detection performance between participants, as can be seen from the 95% confidence intervals in Fig. 5, supports our experimental design. For this reason, we are convinced that increasing N would not result in significantly different threshold estimations and would not change the conclusions of the paper.

Concerning the different distraction types, we did not test the derived thresholds for statistically significant differences. One way to do this would be to derive thresholds for each individual participant in order to obtain a distribution of upper and lower thresholds and PSEs. By comparing these threshold distributions between distraction scenarios, statistically significant differences could be detected. This, however, requires significantly more samples per participant in order to derive individual thresholds, which would have rendered our experiment unfeasible, especially regarding time and user fatigue.

Our work is an important step towards the derivation of hand redirection guidelines that apply even in worst-case application scenarios. While more complex 2D or 3D warps could be constructed by executing the investigated algorithms consecutively, a formal investigation of such combined redirection was out of the scope of this paper, but certainly is a promising next step in the context of this research direction.

6 Conclusion

We investigated how much hand redirection in VR can go unnoticed. For this, we individually investigated 3 basic hand redirection dimensions, namely horizontal, vertical and gain-based VR hand warping, by studying 3 corresponding desktop-scale body warping techniques. Motivated by the lack of general lower-bound detection thresholds for hand redirection in IVEs, we designed a psychometric 2AFC experiment to derive detection thresholds for the 3 basic redirection dimensions. We explored them in a very conservative IVE and in two conservative but more realistic scenarios with significantly higher user distraction. These scenarios employed auditory and tactile distraction and combined visual and dual-task distraction, respectively. We also combined the results of all 3 scenarios to derive general recommendations for conservative hand redirection thresholds. The results are of value for haptic retargeting and indicate that in the investigated techniques, the virtual hand can be displaced *horizontally* or *vertically* by up to $\approx 4.5°$ in either direction respectively, covering a range of $\approx 9°$, without users being able to reliably detect redirection. Regarding the *gain* technique, we found that factors between $g = 0.88$ and $g = 1.07$ can go unnoticed. Thus, the user's real hand can be redirected to grasp up to 13.75% further or up to 6.18% less far than the virtual hand.

Our investigation was able to show that there is a certain range of hand redirection in VR that can go unnoticed, even in worst-case scenarios, but that this range is narrow. Our recommendations quantify this range for the 3 basic dimensions of horizontal, vertical and gain-based hand displacement. Having said this, we also want to emphasize that the derived thresholds are likely not ultimate, exact limits as they might slightly vary as a function of the VR user and the IVE that is experienced – an expectation to be further explored in future work. Instead, they are meant to serve as useful and well-founded reference points that underline the limited range of *unnoticeable* hand redirection in worst-case scenarios. The derived recommendations can serve as a baseline for developers to adjust the redirection applied in their applications. Taking redirection angles and thresholds into account during the development of IVEs, for example, can inform developers as to which redirections might go unnoticed by the user, and which scene compositions result in redirections beyond thresholds. Moreover, by keeping track of the relative angles from the user's hand to props in the real environment, systems can use the thresholds to decide when best to start redirecting the hand. To ensure unnoticed redirection, such systems would trigger redirection early enough to stay within thresholds.

7 Future Work

Motivated by the comparison of our estimates with Burns et al.'s results [6], and by Grechkin et al.'s [10] findings which showed that threshold estimates vary with the methodology used, we plan future investigations using alternative methodologies to compare derived thresholds. Here, a formal exploration of how detection thresholds vary with different reaching distances and hand movement speeds could lead to the formulation of corresponding models. Follow-up research will also investigate the influence of adaptation and distraction in more depth. We can imagine that more engaging tasks, e.g. in the context of a simulation or game, can better control the user's attention than the tested distraction types. In this context, an investigation of adaptive distractions that, for example, adapt to the redirection type, strength and direction, could be fruitful. Apart from that, future work will investigate the combination of the 3 individual redirection dimensions. Complementing the 1D hand redirection ranges derived in this paper, investigating 2D and 3D hand displacement might eventually lead to the formulation of areal or volumetric threshold ranges. These could map a physical hand location to a 2D area or a 3D volume of possible, unnoticeable warped virtual hand locations. Moreover, we will also explore hand redirection in other setups, e.g. while the user is standing, lying down or during locomotion, to see if similar thresholds apply.

Acknowledgments

The authors thank Eva Gehring for her help in creating illustrations and Dominik Cermann for his support in making the accompanying video. Special thanks goes to all participants of the study. This research was funded in part by the German Federal Ministry of Education and Research (BMBF) under grant number 01IS17043 (project ViRUX).

References

[1] P. Abtahi and S. Follmer. Visuo-haptic illusions for improving the perceived performance of shape displays. In *Proc. CHI*, pp. 150:1–150:13. ACM, New York, NY, USA, 2018. doi: 10.1145/3173574.3173724

[2] M. Achibet, A. Girard, A. Talvas, M. Marchal, and A. Lécuyer. Elastic-Arm: Human-scale passive haptic feedback for augmenting interaction and perception in virtual environments. In *Proc. VR*, pp. 63–68. IEEE Computer Society, Mar 2015. doi: 10.1109/VR.2015.7223325

[3] B. Araujo, R. Jota, V. Perumal, J. X. Yao, K. Singh, and D. Wigdor. Snake Charmer: Physically enabling virtual objects. In *Proc. TEI*, pp. 218–226. ACM, New York, NY, USA, 2016. doi: 10.1145/2839462.2839484

[4] M. Azmandian, M. Hancock, H. Benko, E. Ofek, and A. D. Wilson. Haptic Retargeting: Dynamic repurposing of passive haptics for enhanced virtual reality experiences. In *Proc. CHI*, pp. 1968–1979. ACM, New York, NY, USA, 2016. doi: 10.1145/2858036.2858226

[5] Y. Ban, T. Kajinami, T. Narumi, T. Tanikawa, and M. Hirose. Modifying an identified angle of edged shapes using pseudo-haptic effects. In *Proc. EuroHaptics*, pp. 25–36. Springer-Verlag, Berlin, Heidelberg, 2012. doi: 10.1007/978-3-642-31401-8_3

[6] E. Burns, S. Razzaque, A. T. Panter, M. C. Whitton, M. R. McCallus, and J. Frederick P. Brooks. The hand is more easily fooled than the eye: Users are more sensitive to visual interpenetration than to visual-proprioceptive discrepancy. *Presence: Teleoperators and Virtual Environments*, 15(1):1–15, Feb. 2006. doi: 10.1162/pres.2006.15.1.1

[7] L.-P. Cheng, E. Ofek, C. Holz, H. Benko, and A. D. Wilson. Sparse Haptic Proxy: Touch feedback in virtual environments using a general passive prop. In *Proc. CHI*, pp. 3718–3728. ACM, New York, NY, USA, 2017. doi: 10.1145/3025453.3025753

[8] L. Dominjon, A. Lécuyer, J. M. Burkhardt, P. Richard, and S. Richir. Influence of control/display ratio on the perception of mass of manipulated objects in virtual environments. In *Proc. VR*, pp. 19–25. IEEE Computer Society, March 2005. doi: 10.1109/VR.2005.1492749

[9] J. J. Gibson. Adaptation, after-effect and contrast in the perception of curved lines. *Journal of Experimental Psychology*, 16(1):1–31, 1933. doi: 10.1037/h0074626

[10] T. Grechkin, J. Thomas, M. Azmandian, M. Bolas, and E. Suma. Revisiting detection thresholds for redirected walking: Combining translation and curvature gains. In *Proc. SAP*, pp. 113–120. ACM, New York, NY, USA, 2016. doi: 10.1145/2931002.2931018

[11] D. T. Han, M. Suhail, and E. D. Ragan. Evaluating remapped physical reach for hand interactions with passive haptics in virtual reality. *IEEE Transactions on Visualization and Computer Graphics*, 24(4):1467–1476, April 2018. doi: 10.1109/TVCG.2018.2794659

[12] K. Hinckley, R. Pausch, J. C. Goble, and N. F. Kassell. Passive real-world interface props for neurosurgical visualization. In *Proc. CHI*, pp. 452–458. ACM, New York, NY, USA, 1994. doi: 10.1145/191666.191821

[13] H. G. Hoffman. Physically touching virtual objects using tactile augmentation enhances the realism of virtual environments. In *Proc. VRAIS*, pp. 59–63. IEEE, 1998. doi: 10.1109/VRAIS.1998.658423

[14] B. E. Insko. *Passive Haptics Significantly Enhances Virtual Environments*. PhD thesis, University of North Carolina at Chapel Hill, USA, 2001.

[15] L. Kohli. Exploiting perceptual illusions to enhance passive haptics. In *IEEE VR Workshop on Perceptual Illusions in Virtual Environments*, pp. 22–24, 2009.

[16] L. Kohli. Redirected Touching: Warping space to remap passive haptics. In *Proc. 3DUI*, pp. 129–130. IEEE, 2010. doi: 10.1109/3DUI.2010.5444703

[17] L. Kohli. *Redirected Touching*. PhD thesis, University of North Carolina at Chapel Hill, USA, 2013.

[18] L. Kohli. Warping virtual space for low-cost haptic feedback. In *Proc. I3D*, pp. 195–195. ACM, New York, NY, USA, 2013. doi: 10.1145/2448196.2448243

[19] L. Kohli, E. Burns, D. Miller, and H. Fuchs. Combining passive haptics with redirected walking. In *Proc. ICAT*, pp. 253–254. ACM, New York, NY, USA, 2005. doi: 10.1145/1152399.1152451

[20] L. Kohli, M. C. Whitton, and F. P. Brooks. Redirected Touching: The effect of warping space on task performance. In *Proc. 3DUI*, pp. 105–112. IEEE, 2012. doi: 10.1109/3DUI.2012.6184193

[21] Y. Lee, I. Jang, and D. Lee. Enlarging just noticeable differences of visual-proprioceptive conflict in VR using haptic feedback. In *Proc. WHC*, pp. 19–24. IEEE, June 2015. doi: 10.1109/WHC.2015.7177685

[22] Y. Matsuoka, S. J. Allin, and R. L. Klatzky. The tolerance for visual feedback distortions in a virtual environment. *Physiology & Behavior*, 77(4-5):651–655, 2002.

[23] W. A. McNeely. Robotic Graphics: A new approach to force feedback for virtual reality. In *Proc. VRAIS*, pp. 336–341. IEEE Computer Society, Sep 1993. doi: 10.1109/VRAIS.1993.380761

[24] R. A. Montano Murillo, S. Subramanian, and D. Martinez Plasencia. Erg-O: Ergonomic optimization of immersive virtual environments. In *Proc. UIST*, pp. 759–771. ACM, New York, NY, USA, 2017. doi: 10.1145/3126594.3126605

[25] K. Myles and J. T. Kalb. Guidelines for Head Tactile Communication. Technical report, DTIC Document, 2010.

[26] I. Poupyrev, M. Billinghurst, S. Weghorst, and T. Ichikawa. The Go-Go interaction technique: Non-linear mapping for direct manipulation in VR. In *Proc. UIST*, pp. 79–80. ACM, New York, NY, USA, 1996. doi: 10.1145/237091.237102

[27] I. Poupyrev, S. Weghorst, and S. Fels. Non-isomorphic 3D rotational techniques. In *Proc. CHI*, pp. 540–547. ACM, New York, NY, USA, 2000. doi: 10.1145/332040.332497

[28] S. Razzaque. *Redirected Walking*. PhD thesis, University of North Carolina at Chapel Hill, USA, 2005.

[29] M. Slater, M. Usoh, and A. Steed. Depth of presence in virtual environments. *Presence: Teleoperators and Virtual Environments*, 3(2):130–144, 1994. doi: 10.1162/pres.1994.3.2.130

[30] J. Spillmann, S. Tuchschmid, and M. Harders. Adaptive space warping to enhance passive haptics in an arthroscopy surgical simulator. *IEEE Transactions on Visualization and Computer Graphics*, 19(4):626–633, 2013. doi: 10.1109/TVCG.2013.23

[31] F. Steinicke, G. Bruder, J. Jerald, H. Frenz, and M. Lappe. Estimation of detection thresholds for redirected walking techniques. *IEEE Transactions on Visualization and Computer Graphics*, 16(1):17–27, Jan 2010. doi: 10.1109/TVCG.2009.62

[32] S. Tachi, T. Maeda, R. Hirata, and H. Hoshino. A construction method of virtual haptic space. In *Proc. ICAT*, pp. 131–138, Jul 1994.

[33] A. Zenner and A. Krüger. Shifty: A weight-shifting dynamic passive haptic proxy to enhance object perception in virtual reality. *IEEE Transactions on Visualization and Computer Graphics*, 23(4):1285–1294, 2017. doi: 10.1109/TVCG.2017.2656978

[34] Y. Zhao and S. Follmer. A functional optimization based approach for continuous 3D retargeted touch of arbitrary, complex boundaries in haptic virtual reality. In *Proc. CHI*, pp. 544:1–544:12. ACM, New York, NY, USA, 2018. doi: 10.1145/3173574.3174118