

Detectability of ETHD Position and Speed Redirection for VR Haptics

Yuqi Zhou*
Purdue university

Voicu Popescu†
Purdue University

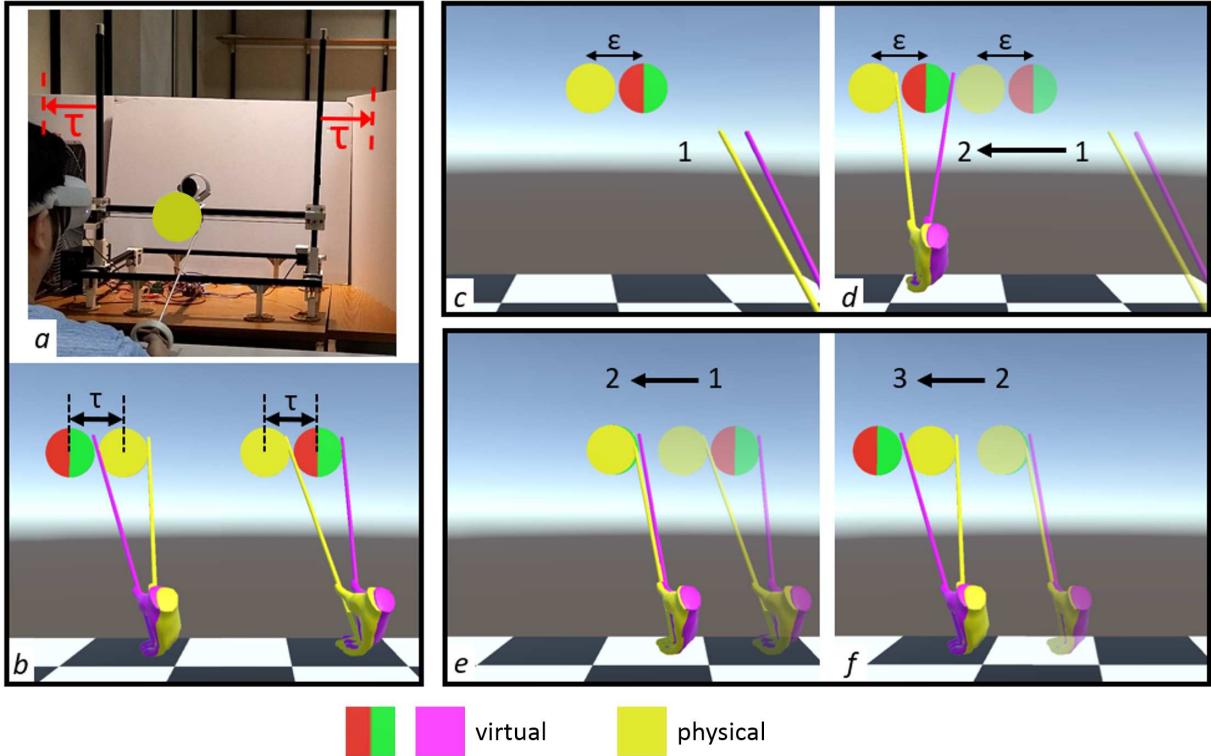


Figure 1: ETHD reachability extension (a, b), and user study experiments (c-f). The reach of our ETHD is virtually extended by τ in all directions (a). The extension is implemented by redirecting the virtual stick (magenta in b) to cover the gap between the physical (yellow) and virtual (red/green) disks (b). The physical disk and stick (yellow) are shown here for illustration purposes and are not shown to the user of the virtual reality application. In a first experiment (c, d), the virtual and physical disks move with the *same* speed, with a constant gap ε between them (c), and the user is asked to tap the virtual disk; the virtual stick is redirected to synchronize the virtual and physical contacts (d); the experiment measures the position redirection detection threshold τ , i.e., the largest undetectable ε . In a second experiment (e, f), the virtual and physical disks move with *different* speeds, and the user is asked to hold the stick in contact with the disk; here the virtual disk is faster than the physical disk and redirection maintains virtual contact; the experiment measures the speed redirection detection threshold, i.e., the largest undetectable speed difference between the disks.

ABSTRACT

An encountered-type haptic device (ETHD) moves a physical object to align it with the virtual object with which the user of a virtual reality application makes contact, providing haptic feedback. One limitation of ETHDs is their limited reachability due to mechanical constraints. One approach to extending the reachability of an ETHD involves redirection, achieved by altering the pose of the user's body, hand, or handheld prop in the virtual world. While previous studies have quantified the detection thresholds of redirection in the context of stationary objects, ETHD's present the opportunity and the need to study redirection in the context of dynamic objects. This paper presents a user study ($N = 25$) with two experiments

aimed at investigating whether dynamic object properties, such as direction and speed, significantly affect redirection detection thresholds. The first experiment finds that the mere presence of dynamic objects does not decrease the detection threshold. Consequently, previously measured detection thresholds remain applicable for ETHD reachability extension. However, the second experiment unveils a crucial relationship between the speed of dynamic objects and the original reachability of the ETHD. Although the virtual object can move 15 cm/s faster than the ETHD, this increase in speed is insufficient to compensate for the extended reachability enabled by the detection threshold.

Index Terms: Virtual Reality, Haptics, Redirection, ETHD

*e-mail: zhou1168@purdue.edu

†e-mail: popescu@purdue.edu

1 INTRODUCTION

In recent years, consumer-level virtual reality (VR) headsets have made remarkable advancements, granting users immersive experiences within intricately crafted virtual worlds. In order to allow the user not only to see, but also to touch the virtual environment (VE),

researchers have proposed several approaches for providing haptic feedback to the user. One approach relies on *active* haptic actuators worn by the user, such as gloves [2] or vests [1]. The advantage is that users take the haptic actuators with them as they move. The disadvantage is user encumbrance, as well as limited haptic feedback resolution, intensity, and duration. An approach that has the potential to alleviate these problems is to provide *passive* haptic feedback through physical objects aligned with the virtual objects that the user touches.

The fundamental challenge in passive haptic feedback is the requirement that a physical object be aligned with each virtual object that the user might ever touch. One approach for addressing this challenge is to modify the virtual world through redirection. Redirection changes the position and orientation of the user’s virtual hand [19], or of the virtual object with which the user is about to make contact [26], such that the physical and virtual contacts are aligned. The goal is to make these changes in a subtle way in order to elude user detection. However, staying below detection thresholds limits the size of the gap between the virtual and the physical object that redirection can bridge. Another approach is to use a mechanical system to align a physical object with the virtual object the user is about to touch. Such an encountered-type haptic device (ETHD) can provide haptic feedback for a continuum of virtual object positions and orientations [32]. The challenges of ETHDs are high cost and limited ranges of motions and speeds. The redirection and the ETHD approaches are orthogonal and can be combined to provide haptic feedback in a wider range of scenarios than each approach can cover individually. Using redirection, the ETHD does not have to align precisely the carried (physical) object to the virtual object. Conversely, the ETHD reduces the size of the virtual to physical gap that has to be bridged by redirection, so redirection can proceed unbeknownst to the user.

In this paper we conducted a study to measure how much the mechanical capabilities of an ETHD can be extended through redirection. Specifically, our study measured the detection thresholds for position and speed redirection in the context of a dynamic virtual object and a dynamic ETHD. The ETHD was implemented with a 3D Cartesian robot with which the user interacted using a handheld stick (Fig. 1a). Compared to other ETHD implementations, the Cartesian robot provides an advantageous trade-off between reach and payload on one hand, and cost on the other.

We ask the user to interact through a stick as opposed to directly, with their hand, for the following reasons: First, there are numerous VR applications where the user does not interact with the virtual environment directly with their hands, but rather with a handheld prop (e.g., bat, sword, hammer, axe, or club), and our research is directly applicable in such contexts. Second, the stick interaction modality is suitable for a “desktop haptics” scenario, where the user is seated in front of a virtual workbench and comfortably interacts with it through the stick, which extends the user’s physical reach, improving ergonomics, and all but eliminates safety concerns from reaching into the ETHD volume without seeing the ETHD. Third, haptic feedback for dynamic virtual objects has not been studied extensively, so our study design aims for maximizing generality. Indeed, we set out to measure position and speed redirection detection thresholds in the absence of potentially strong confounding factors arising from the discrepancy between the physical object surface properties and the surface properties of the virtual object that user expects. Avoiding such discrepancies when the user directly touches the carried object would require one study for every group or even every single type of surface, e.g., a tennis ball, a wooden board, a rough/smooth surface, a cold/hot surface, and so on. The stick blocks the user’s high resolution and multidimensional perception of surface properties and allows these early experiments to focus on the detectability of redirection.

In a first experiment we measured how large *the difference in*

position ε between the virtual object and the ETHD can be without the user noticing the redirection (Fig. 1c, d). Experiment 1 considered three ETHD speeds, two directions of motion, and two relative positions between the virtual object and the ETHD, to measure $3 \times 2 \times 2 = 12$ detection thresholds. Each threshold was measured by analyzing the average participant detection rate at five ε values. The results indicate that detection thresholds remain consistent across various directions and speeds of ETHD movement. This suggests that the detection thresholds established in prior studies [46] can be directly applied to tasks involving touching or tapping.

In a second experiment we measured how large *the difference in speed λ* between the virtual object and the ETHD can be without the user noticing the redirection (Fig. 1e, f). Experiment 2 considered two ETHD speeds, two directions of motion, and two relative positions between the virtual object and the ETHD, to measure $2 \times 2 \times 2 = 8$ detection thresholds. Each threshold was measured by analyzing the average participant detection rate at six λ values. When the virtual object is faster than the ETHD, the detection thresholds indicate that the maximum ETHD speed can be increased virtually through redirection from 22 cm/s to 37 cm/s, i.e., by 68%.

Given an ETHD with specific mechanical reach and speed capabilities, the results of the two experiments allow estimating the virtual object speed for which haptic feedback can be provided, as well as the duration of the haptic feedback. Conversely, given a desired virtual object speed and haptic feedback duration, the experiment results allow prescribing the required minimal ETHD mechanical capabilities. We also refer the reader to the video accompanying our paper which illustrates our experiments and our results.

2 RELATED WORK

Users wearing a VR headset can experience haptic feedback from a variety of sources, which can be categorized into indirect and direct forces applied to their body. An example of indirect haptic feedback is the use of mid-air tactile feedback delivered to the user’s bare hand through the application of ultrasound technology [27]. Haptic feedback can also be provided by imparting to the user direct forces through physical contact. Direct haptic feedback approaches are classified into two categories: *active* and *passive*.

2.1 Active vs passive haptic feedback

Active haptic devices require an external source of power to generate and apply the force to the user’s body. Examples include wearable devices like haptic gloves [2, 4], haptic controllers [41], and finger-mounted haptic devices [8]. Advantages of wearable devices in large virtual environments include continuous haptic feedback for numerous virtual objects encountered during exploration. However, drawbacks include the burden of wearing the device and limited range of haptic intensities. For instance, when touching a virtual wall, the device must generate appropriate forces to mimic realistic sensations without causing discomfort, such as avoiding a pinch sensation. Yet, elastic gloves may distribute opposing forces across the hand, potentially leading to discomfort as they tighten.

On the other hand, passive haptic devices do not require an external source of power, as they rely on the user’s actions to provide the force needed [25, 29]. This haptic feedback method closely replicates real-world sensations, enhancing realism. For instance, when touching a virtual wall, aligning a physical object with it generates authentic feedback, rather than a pinch-like sensation.

The fundamental challenge for passive haptic feedback is the requirement of having a physical object impersonate each virtual object with which the user might interact. Research has shown that there is some leeway in terms of shape, with users being convinced by the haptic feedback from a physical object with a similar but not identical shape to that of the virtual object [17, 11, 12].

Synchronizing the virtual and physical contacts presents challenges [13, 14]. For instance, if a user’s virtual hand reaches an object before their physical hand, the haptic system faces two undesirable options: halt the virtual hand until physical contact catches up, or allow the virtual hand to intersect the object. Likewise, if physical contact precedes virtual contact, the user may observe their virtual hand move without their physical hand moving, or their virtual hand stop short of the target object. Synchronizing the virtual and physical contacts is an important focus of VR haptics research, and it has been approached from two directions: redirection of the virtual entities to adjust the virtual contact (Sec. 2.2), and positioning of the physical object to adjust the physical contact (Sec. 2.3).

2.2 Redirection

The redirection method aims to bridge the gap between physical objects and their virtual counterparts by manipulating either the user’s body rendering in the virtual world [19, 10, 15], that of their handheld prop [37, 45], or the entire virtual world itself [26]. This manipulation can be done within a threshold that makes it difficult for users to perceive any significant warping or distortion, i.e., any anomaly, when virtual objects are positioned near physical ones. A well-designed physical workspace used in conjunction with redirection can find applications across various domains [30].

Redirection not only deceives users about the position of virtual objects but can also alter the object’s perceived shape [44, 46]. Redirection has been used, for example, to expand the resolution of a haptic display for it to render haptic feedback for a larger variety of virtual object shapes [6]. Another application of shape redirection is reducing the need for complex physical props in specific tasks, such as surgical simulators [36], where a generic physical prop can impersonate several virtual surgical instruments. When the user does not touch the virtual object directly with their hand, but rather through a handheld prop, the detection threshold for redirection can be even higher [46], expanding the possibilities for haptic feedback to a broader range of applications [34]. For example, redirection has been used to enhance the realism of tool-based interactions in a virtual workspace, providing haptic feedback evocative of impact and resistance when using tools like a hammer, a saw, or a screwdriver [37].

2.3 Encountered-type Haptic Devices (ETHD)

Besides redirection, another way of synchronizing the virtual and physical contacts is to rely on an active mechanical positioning system to place the physical object in alignment with the virtual object. Such encountered-type haptic devices (ETHDs) allow providing haptic feedback for many virtual object positions. There are several types of ETHDs, including ungrounded ETHDs, such as a drone [7], and grounded ETHDs, such as a robotic arm [16, 40] or such as a table-top robot [22]. Drones have drawn attention in recent years as a way of delivering haptics feedback in VR. Compared to grounded solutions such as a robot arm, a drone can easily cover a larger 3D volume without an extra mechanical support system. Drones have been used as a versatile *haptic puppet* to provide feedback not only around the user’s upper body, but also to the user’s feet, as needed for example to simulate kicking a ball [20]. The drone versatility as a haptic feedback device has also been demonstrated in a study where drones provided reactionary forces from a wide range of directions [9]. Whereas ungrounded ETHDs such as drones have the advantage of flexibility and extended reach, grounded ETHDs have the advantage of higher precision, lower latency, and larger payloads.

Recent surveys indicate that the most commonly used grounded ETHD is the robotic arm, due to its speed, precision, and versatility [32]. A robotic arm offers versatile manipulation with six degrees of freedom and can simulate various textures using attachments like textured tips [31]. These tips can also provide fric-

tion feedback by rolling as the user interacts with virtual surfaces. However, extending the arm’s reach significantly raises costs due to the need for stronger materials and more powerful motors. Alternatively, reachability can be increased by mounting the arm on tracks [16]. A separate robot has been employed to transport the robotic arm anywhere within a room-sized environment [33]. Prior work has also investigated relying on the user to take their ETHD with them as they move through the VE [24]. A Cartesian robot, on the other hand, has the advantage of a larger reach range and payload compared to robot arms of the same cost. Researchers have begun examining the potential of three-axes Cartesian robots for haptic feedback [38] in the context of stationary virtual objects. Our work examines the use of a three-axes Cartesian robot to provide haptic feedback for stationary and dynamic virtual objects, and we investigate the extension of the physical capabilities of the robot through redirection.

2.4 Combination of redirection and ETHD

Researchers have begun examining the benefits of combining ETHDs with redirection. For example, Lee et al., explored the combination of redirection and a robot arm to extend the reach range of an ETHD [28]. Another prior study employed a robot that could move on a planar surface, similar to a robotic floor sweeper, to provide feedback to a user touching stationary virtual objects on a table in front of them [22]. When the robot is unable to reach the desired position in time, redirection is applied to bridge the gap between the robot and the virtual object. Other than the robot arm and sweeper robot, redirection can also be applied to a drone [7].

Previous studies have demonstrated that redirection effectively extends the mapping range of an ETHD from its current position. Consequently, it is evident that redirection can readily enhance the reachability of ETHDs. Since the target in the tasks of previous studies was stationary, the reach range of an ETHD can be extended by the value determined by the detection threshold measured for hand redirection [43, 23, 18] or for stick redirection [46].

Haptic redirection for dynamic objects has not widely studied. Moving objects may impact the detection threshold due to the uncertainty caused by dynamic stimuli. If a dynamic object draws the user’s attention more towards changes in their hand or handheld stick, the detection threshold might decrease, consequently reducing the extendable range of the ETHD compared to tasks involving stationary objects.

Furthermore, while the detection threshold of stationary objects can extend the reach range of an ETHD, the speed disparity between virtual and physical objects becomes apparent when the virtual object transitions between points outside the ETHD. If users can readily detect this speed difference, the extendable reach range will be smaller than that facilitated by the detection threshold of stationary objects.

Therefore, we conducted experiments to measure the detection threshold for two tasks involving dynamic virtual objects: a tapping task and a following task. The tapping task replicates previous studies [43, 23, 18, 46], with the distinction that the object can be either stationary or dynamic. The following task assesses a person’s ability to discern speed differences between virtual and physical objects. We work under the assumption that the virtual object the user intends to touch is known. Our focus is on conducting extensive experiments to measure detection thresholds for multiple values of independent variables such as ETHD speed, direction of motion, and relative position of the virtual and physical objects. Our study aims to investigate how the dynamic properties of virtual objects can impact the reach range extended beyond the detection threshold typically measured for stationary objects.

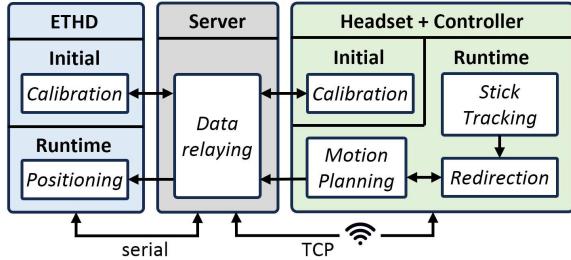


Figure 2: System overview: the ETHD communicates with the headset through a server (laptop).

3 USER STUDY

We conducted a user study, with the approval of our Institutional Review Board, to evaluate how much the reachability of an ETHD can be extended without the user noticing. The study had two experiments. The first experiment investigated detection thresholds when redirection was used to cover the positional difference between the ETHD and the virtual object with which the user interacted. The second experiment investigated detection thresholds when redirection was used to cover the speed difference between the ETHD and the virtual object. We first describe the methods common to both experiments (Sec.3.1), then experiment 1 (Sec.3.2), followed by experiment 2 (Sec.3.3), and we conclude the section with a summary of the study results (Sec.3.4).

3.1 Methods common to both experiments

Participants. We recruited $N = 25$ participants with an average age of 27.2, 17 men and 8 women, 1 never used VR before, 2 used VR once, 18 occasionally, and 4 frequently. All participants were right-hand dominant. The participants completed the two experiments in two separate days, half starting with the first experiment, and the other half with the second. Each experiment took about 30 minutes, for a total of 60 minutes.

Implementation and setup. The participant wore a Meta Quest 2 VR headset [3] and held the controller in their right hand (Fig. 1a). The stick the participant used to interact with the ETHD was attached to the controller. The participant was seated in front of an empty table 0.5m wide, which prevented them from reaching the ETHD with their hand. The VR application was built with Unity 3D [5], version 2022.3.4.

ETHD. The ETHD was implemented with a Cartesian robot that could reach anywhere in a $50\text{ cm} \times 50\text{ cm} \times 30\text{ cm}$ volume. The robot had a speed range of 10 cm/s to 22 cm/s along the x and y axes, and of 3 cm/s to 5 cm/s along the z axis. The ETHD was controlled by the headset through a server, i.e., a laptop, as shown in Fig. 2. The server was connected to the ETHD through a wired serial connection and to the headset through a wireless TCP connection. The ETHD coordinate system was registered to the headset coordinate system in a calibration step that relied on touching the disk with the stick at several ETHD positions. The carried object was a thin 3D printed PLA disk with 14 cm in diameter that served as the physical replica of a virtual disk of the same size. The identical disks minimized the influence of shape for the various impact directions. The virtual contact between the virtual stick and the virtual disk was synchronized to the physical contact between the physical stick and the physical disk through redirection. We chose the carried object to be a disk for two reasons. One is that the redirection detection thresholds measured for virtual and physical disks disk are likely to be similar to those for other pairs of virtual and physical objects. Second, prior work has demonstrated shape redirection in the context of haptic feedback for VR [46]. In other words, the physical and virtual objects do not need to be of

the same shape for believable haptic feedback, and the physical disk has proven to be a versatile proxy for many shapes.

The physical stick, measuring 55 cm in length, was securely attached to the right-hand controller using a 3D-printed part. Constructed from a hollow aluminum rod, the stick possesses a lightweight yet sturdy structure, offering both rigidity and elasticity suitable for the tapping and following tasks. During Experiment 1’s tapping task, participants were directed to tap the disk’s green side. Participants were directed to initiate the task by positioning their hand and stick 1 meter away from the ETHD. Holding the stick vertically, they were instructed to gradually and continuously advance their hand and stick towards the display, as illustrated in the supplementary video. In Experiment 2, participants were asked to place their stick on the virtual disk to initiate movement.

Redirection. We use the prior art FreeHand redirection algorithm [46], which we extend to the novel context of dynamic virtual objects. FreeHand is a redirection algorithm that guarantees the synchronization of the physical and virtual contacts. In other words, the user sees the virtual stick touching the virtual object precisely when the user feels the physical stick touching the physical disk. Freehand relies on a continuous prediction of the two virtual contact points, i.e., between the virtual stick and the virtual object, and of the two physical contact points, i.e., between the physical stick and the physical disk. Based on the prediction of the contact points, the virtual stick pose is adjusted to synchronize the virtual and physical contacts. The contact point prediction is based solely on the previous frame position of the four contact points and on the current position of the virtual and physical sticks and objects; the contact point prediction *does not* rely on predicting the motion of the user’s hand or the motion of the virtual object, which avoids delays between the virtual and physical contacts due to motion prediction errors.

Data collection and analysis. Our study measures redirection detection thresholds by varying the discrepancy between the physical and virtual worlds and asking the participant whether they noticed the difference. We employ a binary (yes/no) design for the user study as outlined in [42], aiming to discern any abnormalities in stick or hand movements perceived by users compared to their daily experiences or expectations. For each discrepancy value a detection rate is computed as the percentage of “yes” answers over all participants and all repetitions. The detection threshold is computed by first fitting a sigmoid psychometric function to the (discrepancy, detection rate) data points. The sigmoid is given in Equation 1, where x stands for the discrepancy magnitude. The fitting process determines parameters α , β , and δ : α determines the horizontal position of the sigmoid; β determines the inverse of the slope of the sigmoid; and δ determines whether the sigmoid increases or decreases (from left to right). Then the detection threshold is computed by intersecting the sigmoid with the horizontal line of 50% detection rate [42]. Unlike for a the two-interval forced-choice (2IFC) design where the threshold is read at the intersection with the 75% line, using the 50% line is appropriate for our “yes / no” detection threshold experimental design. Indeed, the “yes” rate decreases to 0% as the stimulus (i.e., the amount of redirection) decreases to 0, and it increases to 100% as the stimulus increases, so the midpoint, when a majority of participants can detect the redirection, is at 50%.

$$\Sigma(x; \alpha, \beta, \delta) = \frac{1}{1 + e^{\delta \frac{x-\alpha}{\beta}}} \quad (1)$$

In addition to determining the detection thresholds, we also conduct statistical comparisons of the participant detection answers for several independent variables such as ETHD speed, virtual object speed, and discrepancy direction. We use a significance level of $\alpha = 0.05$. Two groups of the categorical data are compared using Fisher’s exact test [21]. Three groups are compared with the

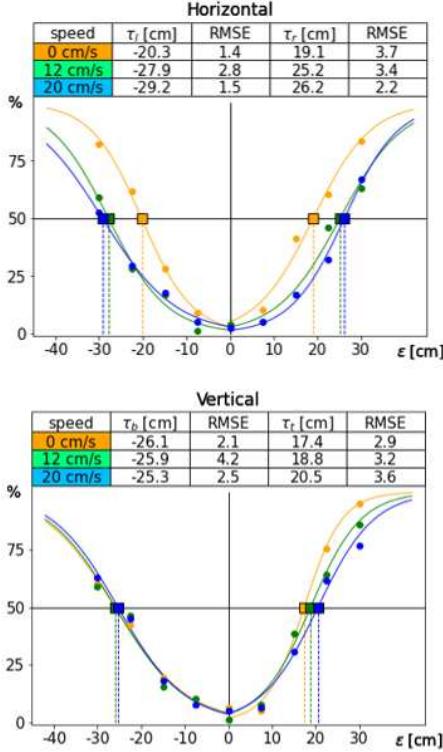


Figure 3: Sigmoid psychometric functions fit to find the 12 detection thresholds for Experiment 1. In each graph, the y axis gives the detection rate, and the x axis gives the distance ϵ between the virtual and physical objects. Negative ϵ values correspond to the physical disk trailing the virtual disk, and positive values correspond to the virtual disk trailing the physical disk. Each graph shows six sigmoids: three for negative ϵ values, one for each speed, and three for positive ϵ values. The detection thresholds τ are computed at the intersection between a sigmoid and the 50% line (square dots). The τ values are given together with the root mean square error (RMSE) of the fit.

Chi-square [35] test first; if the difference is significant, pairwise comparisons are performed with the Fisher’s exact test using a Bonferroni correction of $\times 3$ to account for the three pairs. We used the SciPy statistical package [39].

3.2 Experiment 1

The first experiment investigates how far apart the physical and virtual disks can be before participants notice the redirection.

Task. A participant performed multiple trials. A trial consists of a single tapping task. The participant was asked to tap the green side of the virtual disk using the handheld stick. The redirection algorithm changed the orientation and length of the stick to synchronize the virtual and physical contacts. After the tap, the participant was asked “Did the virtual stick move abnormally?”, and the participant selected either the “yes” or the “no” answer.

Independent variables. One independent variable is the ETHD speed, for which the experiment investigates three values: 0 cm/s, i.e., the ETHD does not move, 12 cm/s, and 20 cm/s. A second independent variable is the ETHD moving direction, which is either horizontal or vertical. The virtual object always moves with the same speed and direction as the ETHD; a third independent variable encodes whether the virtual object precedes or follows the ETHD. Consequently, experiment measured 12 detection thresholds (3 speeds, 2 directions, and 2 relative positions).

Data collection and analysis. The experiment collected the

“yes” / “no” answers of participants for five values of the gap ϵ between the physical and virtual disks: {0 cm, 7.5 cm, 15 cm, 22.5 cm, 30 cm}. No redirection should be noticeable for $\epsilon = 0$ cm when no redirection is applied, and the redirection is clearly noticeable for $\epsilon = 30$ cm. The data was analyzed with the sigmoid function in Eq. 1, for which the x variable was the gap ϵ .

Procedure. A participant performed 162 counterbalanced trials: 12 independent variable value combinations \times 4.5 ϵ values (the 0 cm value is used for both relative positions) \times 3 repetitions. Prior to each trial, the ETHD moved to the starting position, and participants were instructed to hold the stick close to their body, i.e., right fist close to right shoulder. At the beginning of each trial, the virtual disk appeared and started moving, in sync with the ETHD. To prevent memorization, a random offset between -5 cm and 5 cm was applied to the initial disk position. Participants were asked to tap the disk on its green side and the green side was always the trailing side of the disk. In other words, the stick moved in the same direction as the disk, catching up to it. For the trials when the speed was 0 cm/s, the entire disk was green, indicating that the participant could tap it on either side. If a participant attempted to tap the disk against its motion, the ETHD and the virtual disk stopped, and the invalid trial was discarded and repeated. In all trials, the ETHD moved for a duration of 2 seconds, providing participants sufficient time to see and tap the virtual disk. If this time limit was exceeded, the trial was restarted. A countdown timer was displayed above to inform participants of the remaining time. Before the actual experiment, a participant completed 18 practice trials, two for each of the 9 ϵ values, for various directions and speeds.

Results. Fig. 3 shows the sigmoids fit to find the 12 detection thresholds for Experiment 1. The RMSE indicates the average vertical distance between data points and the sigmoid. The RMSE values are in the [1.4%, 4.2%] range, which indicates a tight fit. The smallest absolute τ value is 17.4 cm, which indicates that redirection can increase the reachability of the ETHD by 17 cm horizontally and vertically with fewer than 50% of the participants noticing. This corresponds to increasing the 50 cm linear reach of the Cartesian motor to 84 cm, a 68% increase. Furthermore, the detection thresholds do not decrease with ETHD speed, and the horizontal detection thresholds increase with speed. The horizontal detection thresholds are symmetrical with respect to $\epsilon = 0$, i.e., τ_l and τ_r have similar values, but the vertical detection thresholds are not. We further investigate this difference below.

Influence of ETHD speed. Next we investigate whether for a given virtual to physical gap ϵ the redirection detection rate depends on the ETHD velocity. Tab. 1 compares the detection errors for all 9 ϵ values between the three ETHD velocities, for both the horizontal and the vertical moving direction. The only significant differences are for large gaps, and between the stationary ETHD (0 cm/s) and the non-stationary ETHD (12 cm/s and 20 cm/s).

Influence of the relative position of the virtual and physical objects. The detection threshold τ_b is consistently smaller than τ_l for all ETHD speeds, as shown in Fig. 3. We have investigated this asymmetry by comparing detection rates between $+\epsilon$ and $-\epsilon$ (Tab. 2). The differences are not significant for the horizontal direction, but they are significant for the vertical direction, except for the smallest absolute value of $\epsilon = \pm 7.5$ cm. We attribute this asymmetry to the asymmetrical influence of gravity, which is aligned with the vertical direction and not with the horizontal direction. When a participant hits down, the physical stick accelerates until it hits the physical disk. When the physical disk is lower than the virtual disk, redirection *decelerates* the virtual stick to delay the virtual contact, which is more unexpected and perceptually more salient than the virtual stick accelerating faster than expected.

		Virtual to physical position difference ε [cm]									
		-30	-22.5	-15	-7.5	0	7.5	15	22.5	30	
Horizontal	"yes" answer [%]	0 [cm/s]	82.1	61.5	28.2	9	2.6	10.3	41	60.3	83.3
		12 [cm/s]	59	28.2	16.7	1.3	3.8	5.1	16.7	46.2	62.8
		20 [cm/s]	52.6	29.5	17.9	5.1	2.6	5.1	16.7	32.1	66.7
	Chi-square test	p	<0.001	<0.001	0.152	0.093	0.863	0.342	<0.001	0.002	0.011
		statistic	16.393	23.234	3.769	4.743	0.295	2.147	16.551	12.484	8.997
	Pairwise comparison	0 vs 12 [cm/s]	0.008	<0.001	-	-	-	-	0.004	0.325	0.019
		12 vs 20 [cm/s]	1.558	3	-	-	-	-	3	0.301	2.213
		0 vs 20 [cm/s]	<0.001	<0.001	-	-	-	-	0.004	0.002	0.077
Vertical	"yes" answer [%]	0 [cm/s]	60.3	42.3	19.2	7.7	6.4	5.1	38.5	75.6	94.9
		12 [cm/s]	59	46.2	15.4	10.3	1.3	7.7	38.5	64.1	85.9
		20 [cm/s]	62.8	44.9	17.9	7.7	5.1	6.4	30.8	61.5	76.9
	Chi-square test	p	0.882	0.886	0.813	0.804	0.257	0.808	0.512	0.136	0.006
		statistic	0.251	0.242	0.414	0.437	2.716	0.427	1.337	3.987	10.372
	Pairwise comparison	0 vs 12 [cm/s]	-	-	-	-	-	-	-	-	0.302
		12 vs 20 [cm/s]	-	-	-	-	-	-	-	-	0.649
		0 vs 20 [cm/s]	-	-	-	-	-	-	-	-	0.006

Table 1: Comparison of detection rates between three ETHD speeds (0 cm/s, 12 cm/s, and 20 cm/s), for each ETHD direction (horizontal and vertical), and for each gap ε . The three speeds are first compared together (Chi-square), and then, if differences are significant, pairwise (Fisher's exact). Significant differences are shown in red.

3.3 Experiment 2

The second experiment investigates how large the difference in speed between the virtual and physical objects can be before participants notice the redirection.

Task. The participant was asked to place the stick on a stationary virtual disk. Then, once contact was made, the virtual disk started moving, and the participant was asked to keep the stick on the moving virtual disk. The physical disk carried by the ETHD is stationary at the beginning, and then when contact is made, the ETHD starts moving in the same direction as the virtual disk, but with a different speed. When the ETHD is slower than the virtual disk, the ETHD starts ahead of the virtual disk. When the ETHD is faster, the ETHD starts behind the virtual disk. The gap between the physical and virtual disks starts out large, then it shrinks to 0 as the faster disk catches up to the slower disk, and then it increases again. The motion of the disks lasted for 1.5 s. The gap is bridged through redirection to show the virtual stick in contact with the virtual disk at all times. After the 1.5 s, the participant was asked "Did the virtual disk have a different speed from the physical disk?", and the participant selected either the "yes" or the "no" answer.

The physical and virtual objects are not aligned at the beginning of Experiment 2 to focus on the detection of speed redirection while minimizing the confounding effect of position redirection. Specifically, we chose an initial position for the physical and virtual objects that is within the position redirection detection threshold established by Experiment 1. This ensures that as the objects move, their positions remain below this threshold. If the objects were placed at the same initial position, noticeable position differences would quickly arise, especially when the relative speed between the objects is large. Such position discrepancies could influence the detection threshold for speed redirection. Therefore, by

		Fisher's exact test p for $-\varepsilon$ vs ε [cm]			
		-30 vs 30	-22.5 vs 22.5	-15 vs 15	-7.5 vs 7.5
Horizontal	0 [cm/s]	1	1	0.129	1
	12 [cm/s]	0.743	<0.001	1	0.367
	20 [cm/s]	0.102	0.862	1	1
Vertical	0 [cm/s]	<0.001	<0.001	0.013	0.746
	12 [cm/s]	<0.001	0.036	0.002	0.78
	20 [cm/s]	0.08	0.054	0.092	1

Table 2: Detection rate comparison for the same virtual to physical gap size (i.e., absolute value of ε) but different virtual to physical relative position (i.e., different ε sign).

allowing the objects to start from different positions, Experiment 2 can more reliably measure speed redirection detection thresholds across a wider range of relative speeds.

Independent variables. One independent variable is the ETHD speed, for which we investigated two values, i.e., 12 cm/s and 20 cm/s. A second independent variable is the motion direction, which was either horizontal or vertical. A third independent variable was whether the virtual disk was faster or slower than the ETHD.

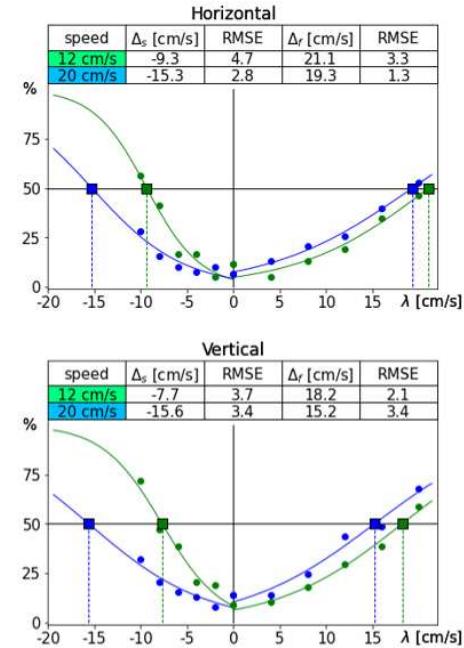


Figure 4: Sigmoid psychometric functions fit to find the 8 detection thresholds Δ for Experiment 2. In each graph, the y axis gives the detection rate, and the x axis gives the speed difference λ between the virtual and physical objects. Negative λ values correspond to the virtual disk being slower than the physical disk (Δ_s), and positive values to the virtual disk being faster (Δ_f). Each graph shows four sigmoid: two for negative λ values, one for each ETHD speed, and two for positive λ values. The Δ values are given together with the root mean square error (RMSE) of the fit.

		Virtual to physical speed difference λ [cm/s]										
		-10	-8	-6	-4	-2	0	4	8	12	16	20
Horizontal	”yes” answer [%]	12 [cm/s]	56.4%	41%	16.7%	16.7%	5.1%	11.5%	5.1%	12.8%	19.2%	34.6%
		20 [cm/s]	28.2%	15.4%	10.3%	7.7%	10.3%	6.40%	12.8%	20.5%	25.6%	39.7%
	Fisher’s exact test p		<0.001	<0.001	0.348	0.14	0.368	0.402	0.159	0.283	0.443	0.619
Vertical	”yes” answer [%]	12 [cm/s]	71.8%	47.4%	38.5%	20.5%	19.2%	9%	10.3%	17.9%	29.5%	38.5%
		20 [cm/s]	32.1%	20.5%	15.4%	12.8%	7.7%	14.1%	14.1%	24.4%	43.6%	48.7%
	Fisher’s exact test p		<0.001	<0.001	0.002	0.283	0.059	0.453	0.625	0.433	0.096	0.258

Table 3: Comparison of detection rates between two ETHD speeds (12 cm/s and 20 cm/s), for each ETHD direction (horizontal and vertical), and for each relative speed gap λ . The two speeds are compared using Fisher’s exact test. Significant differences are shown in red.

Data collection and analysis. Experiment 2 collected the ”yes” / ”no” answers of participants for six speed differences λ between the ETHD and the virtual disk. When the virtual disk was not slower than the ETHD, the λ values are 0 cm/s, 4 cm/s, 8 cm/s, 12 cm/s, 16 cm/s, and 20 cm/s. The maximum value of 20 cm/s was chosen to yield an initial gap between the virtual and physical disks of 15 cm, which is below the detection threshold for stationary virtual objects [46]. Indeed, over the 1.5 s of motion, a 20 cm/s speed difference corresponds to a drift between the two disks of 30 cm, which distributed evenly between the starting and ending position corresponds to a 15 cm initial gap. When the virtual disk was not faster than the ETHD, the λ values are -10 cm/s, -8 cm/s, -6 cm/s, -4 cm/s, -2 cm/s, and 0 cm/s. The largest speed gap of -10 cm/s was chosen such that the virtual disk moves with a speed of at least 2 cm/s when the ETHD speed is 12 cm/s.

We investigated ETHD speeds that were both slower and faster than the speed of the virtual object because the ETHD does not only have a maximum speed but also a minimum speed: the Cartesian robot cannot move slower than 10 cm/s due to physical limitations (asking the stepper motors to produce lower speeds causes mechanical vibration and noise). This way the experiment informs both on redirection detection thresholds for a virtual extension of the maximum ETHD speed, as well as for overcoming the minimum speed limitation of the ETHD.

The ”yes” answer rate data was analyzed with the sigmoid function in Eq. 1, for which the x variable is now the speed difference λ . Experiment measures 8 detection thresholds, for each ETHD speed (12 cm/s or 20 cm/s), motion direction (horizontal or vertical), and virtual/physical relative position (virtual disk following or being followed by physical disk).

Procedure. Each participant completed 14 practice trials for various λ values, ETHD speeds, and moving directions. Then each participant completed 132 trials in counter-balanced order, covering the 11 distinct λ values, two directions, two speeds, and three repetitions. Like for experiment 1, at the beginning of the trial the participant is asked to retract the stick. There is no time limit, as the trial is guided by the motion of the disk. If the participant moves the stick off the disk before the 1.5 s, the trial is repeated.

Results. Fig. 4 shows the 8 detection thresholds Δ measured by experiment 2: two motion directions (horizontal and vertical) \times two ETHD speeds (12 cm/s and 20 cm/s) \times two relative speeds between the virtual disk and the ETHD (Δ_s for when the virtual disk was slower than the ETHD, and Δ_f for when it was faster).

The virtual disk can be faster than the ETHD by as much as 15 cm/s before the discrepancy is noticed by more than 50% of the participants. This reveals that the ETHD with redirection can be useful at providing haptic feedback for dynamic objects that move considerably faster than the ETHD. In other words, redirection can bring a virtual increase of the speed of the ETHD, without the cost of actually increasing the speed of the Cartesian robot. The graphs show that the detection rate barely reaches 50% for the highest positive λ values, which could suggest trying even larger λ values to establish reliably even higher detection thresholds Δ_f . However, λ cannot be further increased as that would either reduce the contact time

below 1.5 s and participants will not have enough time to perceive the redirection, or it would increase the initial gap between disks beyond the detection threshold of 15 cm, revealing the redirection from the beginning. For example, a $\lambda = 30\text{cm/s}$ either implies a contact time of $(15\text{ cm} \times 2) / 30\text{ cm/s} = 1\text{ s}$, or, to maintain the 1.5 s contact time, an initial gap of $(30\text{ cm/s} \times 1.5\text{ s}) / 2 = 22.5\text{ cm}$.

When the virtual disk is slower than the ETHD, redirection can support virtual objects moving with low speed, circumventing the Cartesian robot minimum speed limitation. For example, for the ETHD speed of 12 cm/s, the virtual object can move with a speed of 12 cm/s - 9.3 cm/s = 2.7 cm/s. We decided to investigate negative λ values in the range [-10 cm/s, 0 cm/s] based on a small scale pilot study where participants consistently detected the redirection for $\lambda = -10\text{cm/s}$. The actual experiment reveals that most participants cannot detect redirection at -10 cm/s, so the thresholds Δ_s are computed by extrapolation.

Like for experiment 1, the detection thresholds are similar for the horizontal and vertical directions, with slightly lower thresholds for the vertical direction, especially for larger ETHD speeds.

Tab. 3 reveals that the ETHD velocity doesn’t influence the detection rate, i.e., the pairwise comparisons between the two speeds for a given λ do not reveal significant differences. The only exceptions are for the cases when the virtual disk is considerably slower than the ETHD, i.e., $\lambda \in \{-10\text{ cm/s}, -8\text{ cm/s}\}$. In these cases, when the ETHD speed is 12 cm/s, the virtual disk moves slowly, i.e., with absolute speeds of 2 cm/s and 4 cm/s; when the ETHD speed is 20 cm/s, the virtual disk absolute speeds are 10 cm/s and 12 cm/s; the results show that the detection rate is significantly higher when the virtual disk moves slowly, e.g., 56.4% vs 28.2% and 71.8% vs 32.1%. When the virtual object is faster than the ETHD, the detection rate is not influenced by the ETHD speed. In other words, users are sensitive to the absolute speed difference λ and not the relative speed difference (i.e., λ divided by the ETHD speed). While this indicates that our results might remain valid for higher ETHD speeds, it also means that these higher ETHD speeds might not increase the range of speed differences that can be hidden through direction.

3.4 Study results summary

Experiment 1 demonstrates that regardless of the speed and direction of the ETHD, the detection thresholds for dynamic objects consistently equal or exceed those for stationary objects. This suggests that the presence of dynamic objects does not increase the user’s focus on changes to their virtual hand or handheld stick. Therefore, the detection thresholds measured for stationary objects in previous studies, such as [46], can be directly applied to extend the reachability of an ETHD. This finding can be extrapolated to other stick lengths or types of ETHD, while acknowledging that the absolute thresholds may vary.

Experiment 2 shows that redirection can extend the speed range of the ETHD from $[e_{min}, e_{max}]$ to $[5, e_{max} + 15]$, in cm/s. For our ETHD $e_{min} = 10\text{ cm/s}$ and $e_{max} = 22\text{ cm/s}$, so $[10, 22]$ can be extended to $[5, 37]$, in cm/s.

Fig. 5 leverages the results of experiments 1 and 2 to show for how long haptic feedback can be provided for a virtual object mov-

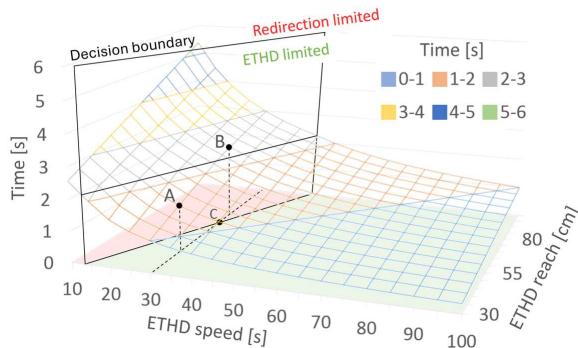


Figure 5: Connection between ETHD speed, reach, and longest possible haptic feedback duration, based on the detection thresholds $\tau = 17\text{cm}$ and $\Delta = 15\text{cm/s}$ established by experiments 1 and 2. For the red region of the $[\text{ETHD speed} \times \text{ETHD reach}]$ plane, the ability to provide haptic feedback is limited by the user’s ability to detect the redirection. For the green region, the ability to provide haptic feedback is limited by the ETHD mechanical capabilities.

ing at the $e_{\max} + 15\text{ cm/s}$ speed. The graph connects the reach and speed capabilities of the ETHD to the maximum virtual object speed for which the ETHD can provide haptic feedback, and for how long. For example, if a VR application wants to provide haptic feedback for at least 2 s for virtual objects moving with speeds up to 45 cm/s, the ETHD maximum speed has to be at least 30 cm/s; furthermore, the ETHD reach has to be at least 60 cm. In Fig. 5, point A corresponds to an ETHD maximum speed of 30 cm/s and a reach of 45 cm, which cannot meet the 2 s requirement. Point B corresponds to an ETHD with the same maximum speed of 30 cm/s, but with a 75 cm reach, which meets the 2 s requirement.

The plane of ETHD mechanical capabilities is partitioned by a line in a region where the haptic feedback is limited by the ETHD capabilities (green in Fig. 5) and a region where the haptic feedback is limited by the user’s ability to detect the redirection (red). For A, the haptic feedback is limited by the ETHD capabilities: the ETHD can move for 45 cm / 30 cm/s = 1.5 s, whereas the virtual object can move for $(45+34) \text{ cm} / (30+15) \text{ cm/s} = 1.76 \text{ s}$. For B, the haptic feedback is limited by the user’s ability to detect redirection: the ETHD can move for 75 cm / 30 cm/s = 2.5 s, whereas the virtual object can move for $(75+34) \text{ cm} / (30+15) \text{ cm/s} = 2.42 \text{ s}$.

The graph in Fig. 5 can also be used when building an ETHD. Given a maximum ETHD speed, the reach should be sufficient to avoid that the reach be the limiting factor when providing haptic feedback. For example, for an ETHD maximum speed of 30 cm/s, the reach should be at least 68 cm (i.e., point C on graph). The decision boundary is denoted as $\text{ETHD reach} = \frac{2\tau}{\Delta_f} \text{ETHD speed}$.

In our context, this fraction $\frac{2\tau}{\Delta_f}$ is determined to be $\frac{34}{15}$. While τ and Δ_f can vary due to factors such as stick length or prolonged usage with redirection, the established relationship remains consistent. Essentially, this equation elucidates how the reachability of an ETHD correlates with the disparity in speed between virtual and physical objects during following tasks.

In summary, given a desired duration of interaction, the application cannot fully benefit from the user’s inability to detect speed redirection unless the ETHD provides a minimum reachability. In other words, a fast moving virtual object will run out of ETHD reach before the desired duration of interaction is reached, even though the low ETHD speed is not a problem per se.

4 CONCLUSIONS, LIMITATIONS, AND FUTURE WORK

We have conducted a study that investigated how much the haptic feedback capabilities of an ETHD implemented by a Cartesian

robot can be extended through redirection. Specifically, we established the redirection detection thresholds for the position and the speed differences between a virtual and a physical object, considering multiple ETHD speeds, motion directions, and relative positions of the virtual and physical objects.

Our study has several limitations. One limitation stems from the fact that, surprisingly, participants were not able to detect the redirection consistently even when the virtual object was moving slower than half the speed or faster than twice the speed of the physical object. Additional experiments should be run to improve the robustness of the measurement. Our study focuses on a single carried object shape, with the shape having little impact on our tapping or following tasks, but future work should replicate the experiments for other shapes. Our study included three trials per parameter set, providing robust measurements of detection thresholds across 25 participants, resulting in 75 yes/no responses for each virtual-to-physical position (ε) or speed differences (λ). Future research could expand this by increasing trial numbers per participant for precise threshold measurement. To manage participant fatigue, a between-subjects design may be necessary. Once individual thresholds are established, inferential statistics can determine significant differences among ETHD speeds or movement directions.

Our ETHD has a fairly high minimum speed, so it does not support very slow virtual objects. One question that future work should answer is what is the highest minimum speed of an ETHD to support arbitrarily small virtual object speeds, i.e., approaching zero. Our experiments indicate that redirection detection acuity increases for slowly moving virtual objects, so it could be that a minimum ETHD speed of 3 cm/s is still too fast to provide extended haptic feedback for virtual objects moving ten times slower at 0.3 cm/s. Since ETHDs typically struggle to render low speeds, redirection is essential to extend their usability for such scenarios. Additional studies are needed to investigate how slowly a virtual object can move using an ETHD with a given minimum speed, without participants noticing the speed redirection.

Our experiments were conducted with a rigid stick of 55 cm in length, but other stick stiffnesses and lengths might result in different thresholds, which future work should examine. Our experiments used identical physical and virtual objects, and future work could also investigate redirection detection thresholds for physical and virtual objects of different shape, leveraging prior work on shape redirection [46]. Finally, future studies could also investigate allowing the user to touch the virtual object directly, for matching or redirected surface properties such as stiffness and texture.

Our ETHD is well-suited for outside-looking-in interactions, where the user engages with a desktop-sized virtual world from the outside. However, it is not suitable for inside-looking-out interactions, where the user walks through the virtual environment. For such interactions, other designs, such as drones, are preferable. Although we chose to use a grounded Cartesian robot due to its large reach, payload capacity, and low cost, our study could also inform the development of other ETHD designs, including drones, movable robotic arms, and other types of Cartesian robots.

We have demonstrated that virtual redirection can significantly and robustly extend the position and speed range of an ETHD, which is likely to benefit in the context of VR applications across various domains such as education, simulation and training, and entertainment. For example, our work contributes to the infrastructure needed to stand up virtual laboratory applications where the learner does not only see the objects but can also touch them, strengthening the embodied cognition that has been shown to promote learning.

ACKNOWLEDGMENTS

This material is based upon work supported by the United States National Science Foundation under Awards No. 2212200 and 2309564.

REFERENCES

[1] bhaptics vr suits. <https://www.bhaptics.com/>. Accessed: 2023-09-22. 2

[2] Haptx vr gloves. <https://haptx.com/>. Accessed: 2023-09-22. 2

[3] Quest 2 virtual reality headset. <https://www.meta.com/quest2>. Accessed: 2023-10-3. 4

[4] touchdiver. <https://www.weart.it/>. Accessed: 2023-9-8. 2

[5] Unity 3d engine. <https://unity.com/>. Accessed: 2023-09-29. 4

[6] P. Abtahi and S. Follmer. Visuo-haptic illusions for improving the perceived performance of shape displays. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pp. 1–13, 2018. 3

[7] P. Abtahi, B. Landry, J. Yang, M. Pavone, S. Follmer, and J. A. Landay. Beyond the force: Using quadcopters to appropriate objects and the environment for haptics in virtual reality. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pp. 1–13, 2019. 3

[8] O. J. Ariza Nunez, A. Zenner, F. Steinicke, F. Daiber, and A. Krüger. Holitouch: Conveying holistic touch illusions by combining pseudo-haptics with tactile and proprioceptive feedback during virtual interaction with 3dusis. *Frontiers in Virtual Reality*, 3:879845, 2022. 2

[9] M. I. Awan, A. Raza, and S. Jeon. Dronehaptics: Encountered-type haptic interface using dome-shaped drone for 3-dof force feedback. In *2023 20th International Conference on Ubiquitous Robots (UR)*, pp. 195–200. IEEE, 2023. 3

[10] 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 *Proceedings of the 2016 chi conference on human factors in computing systems*, pp. 1968–1979, 2016. 3

[11] Y. Ban, T. Kajinami, T. Narumi, T. Tanikawa, and M. Hirose. Modifying an identified angle of edged shapes using pseudo-haptic effects. In *Haptics: Perception, Devices, Mobility, and Communication: International Conference, EuroHaptics 2012, Tampere, Finland, June 13–15, 2012. Proceedings, Part I*, pp. 25–36. Springer, 2012. 2

[12] Y. Ban, T. Kajinami, T. Narumi, T. Tanikawa, and M. Hirose. Modifying an identified curved surface shape using pseudo-haptic effect. In *2012 IEEE Haptics Symposium (HAPTICS)*, pp. 211–216. IEEE, 2012. 2

[13] E. Burns, S. Razzaque, A. T. Panter, M. C. Whitton, M. R. McCallus, and F. P. Brooks Jr. The hand is more easily fooled than the eye: Users are more sensitive to visual interpenetration than to visual-proprioceptive discrepancy. *Presence: teleoperators & virtual environments*, 15(1):1–15, 2006. 3

[14] E. Burns, S. Razzaque, M. C. Whitton, and F. P. Brooks. Macbeth: The avatar which i see before me and its movement toward my hand. In *2007 IEEE Virtual Reality Conference*, pp. 295–296. IEEE, 2007. 3

[15] 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 *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pp. 3718–3728, 2017. 3

[16] S. Dai, J. Smiley, T. Dwyer, B. Ens, and L. Besancon. Robohapalytics: a robot assisted haptic controller for immersive analytics. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):451–461, 2022. 3

[17] X. de Tinguy, C. Pacchierotti, M. Emily, M. Chevalier, A. Guignardat, M. Guillaudeux, C. Six, A. Lécuyer, and M. Marchal. How different tangible and virtual objects can be while still feeling the same? In *2019 ieee world haptics conference (whc)*, pp. 580–585. IEEE, 2019. 2

[18] S. Esmaeili, B. Benda, and E. D. Ragan. Detection of scaled hand interactions in virtual reality: The effects of motion direction and task complexity. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 453–462. IEEE, 2020. 3

[19] M. Feick, K. P. Regitz, A. Tang, and A. Krüger. Designing visuo-haptic illusions with proxies in virtual reality: Exploration of grasp, movement trajectory and object mass. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pp. 1–15, 2022. 2, 3

[20] M. Feick, A. Tang, and A. Krüger. Hapticpuppet: A kinesthetic mid-air multidirectional force-feedback drone-based interface. In *Adjunct Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, pp. 1–3, 2022. 3

[21] R. A. Fisher. On the interpretation of χ^2 from contingency tables, and the calculation of p. *Journal of the royal statistical society*, 85(1):87–94, 1922. 4

[22] E. J. Gonzalez, P. Abtahi, and S. Follmer. Reach+ extending the reachability of encountered-type haptics devices through dynamic redirection in vr. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*, pp. 236–248, 2020. 3

[23] E. J. Gonzalez and S. Follmer. Investigating the detection of bimanual haptic retargeting in virtual reality. In *Proceedings of the 25th ACM Symposium on Virtual Reality Software and Technology*, pp. 1–5, 2019. 3

[24] A. Horie, M. Y. Saraiji, Z. Kashino, and M. Inami. Encounteredlimbs: A room-scale encountered-type haptic presentation using wearable robotic arms. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*, pp. 260–269. IEEE, 2021. 3

[25] B. E. Insko. *Passive haptics significantly enhances virtual environments*. The University of North Carolina at Chapel Hill, 2001. 2

[26] L. Kohli, M. C. Whitton, and F. P. Brooks. Redirected touching: Training and adaptation in warped virtual spaces. In *2013 IEEE Symposium on 3D User Interfaces (3DUI)*, pp. 79–86. IEEE, 2013. 2, 3

[27] P. Kourtesis, F. Argelaguet, S. Vizcay, M. Marchal, and C. Pacchierotti. Electrotactile feedback applications for hand and arm interactions: A systematic review, meta-analysis, and future directions. *IEEE Transactions on Haptics*, 2022. 2

[28] C.-G. Lee, G. L. Dunn, I. Oakley, and J. Ryu. Visual guidance for a spatial discrepancy problem of in encountered-type haptic display. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 50(4):1384–1394, 2017. 3

[29] R. W. Lindeman, J. L. Sibert, and J. K. Hahn. Towards usable vr: an empirical study of user interfaces for immersive virtual environments. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 64–71, 1999. 2

[30] B. J. Matthews, B. H. Thomas, G. S. Von Itzstein, and R. T. Smith. Towards applied remapped physical-virtual interfaces: Synchronization methods for resolving control state conflicts. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pp. 1–18, 2023. 3

[31] V. Mercado, M. Marchal, and A. Lécuyer. Entropia: towards infinite surface haptic displays in virtual reality using encountered-type rotating props. *IEEE transactions on visualization and computer graphics*, 27(3):2237–2243, 2019. 3

[32] V. R. Mercado, M. Marchal, and A. Lécuyer. “haptics on-demand”: A survey on encountered-type haptic displays. *IEEE Transactions on Haptics*, 14(3):449–464, 2021. 2, 3

[33] S. Mortezaipoor, K. Vasylevska, E. Vonach, and H. Kaufmann. Cobodeck: A large-scale haptic vr system using a collaborative mobile robot. In *2023 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*, pp. 297–307. IEEE, 2023. 3

[34] N. C. Nilsson, A. Zenner, and A. L. Simeone. Propping up virtual reality with haptic proxies. *IEEE Computer Graphics and Applications*, 41(5):104–112, 2021. 3

[35] K. Pearson. X. on the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 50(302):157–175, 1900. 5

[36] J. Spillmann, S. Tuchschnid, 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. 3

[37] P. L. Strandholt, O. A. Dogaru, N. C. Nilsson, R. Nordahl, and S. Serafin. Knock on wood: Combining redirected touching and physical props for tool-based interaction in virtual reality. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–13, 2020. 3

[38] B. van Deurzen, P. Goorts, T. De Weyer, D. Vanacken, and K. Luyten. Hapticpanel: An open system to render haptic interfaces in virtual reality for manufacturing industry. In *Proceedings of the 27th ACM Symposium on Virtual Reality Software and Technology*, pp. 1–3, 2021. [3](#)

[39] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2 [5](#)

[40] E. Vonach, C. Gatterer, and H. Kaufmann. Vrrobot: Robot actuated props in an infinite virtual environment. In *2017 IEEE Virtual Reality (VR)*, pp. 74–83. IEEE, 2017. [3](#)

[41] E. Whitmire, H. Benko, C. Holz, E. Ofek, and M. Sinclair. Haptic revolver: Touch, shear, texture, and shape rendering on a reconfigurable virtual reality controller. In *Proceedings of the 2018 CHI conference on human factors in computing systems*, pp. 1–12, 2018. [2](#)

[42] F. A. Wichmann, F. Jäkel, and J. Wixted. Methods in psychophysics. *Stevens handbook of experimental psychology and cognitive neuroscience*, 5:265–306, 2018. [4](#)

[43] A. Zenner and A. Krüger. Estimating detection thresholds for desktop-scale hand redirection in virtual reality. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 47–55. IEEE, 2019. [3](#)

[44] Y. Zhao and S. Follmer. 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*, pp. 1–12, 2018. [3](#)

[45] Y. Zhou and V. Popescu. Tapping with a handheld stick in vr: Redirection detection thresholds for passive haptic feedback. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 83–92. IEEE, 2022. [3](#)

[46] Y. Zhou and V. Popescu. Dynamic redirection for vr haptics with a handheld stick. *IEEE Transactions on Visualization and Computer Graphics*, 29(5):2753–2762, 2023. [2](#), [3](#), [4](#), [7](#), [8](#)