Examining Effects of Technique Awareness on the Detection of Remapped Hands in Virtual Reality

Brett Benda* Benjamin Rheault† Yanna Lin‡ Eric D. Ragan§

University of Florida, Department of Computer and Information Science and Engineering

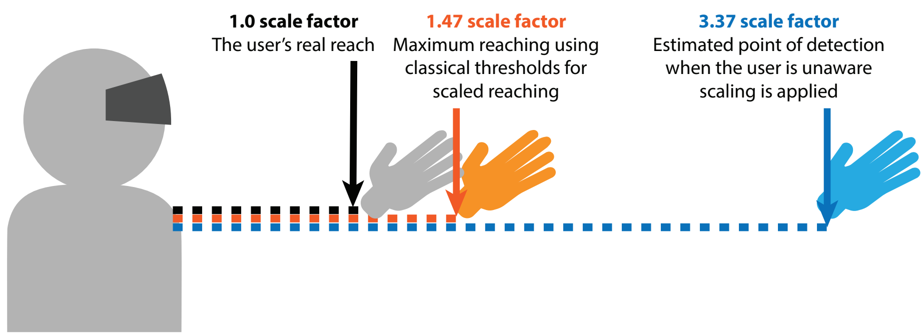

Figure 1: A to-scale image demonstrating the increase in undetectable technique usage when users do not know it is being used. The blue hand represents the average first-detected scale when users are not aware of the technique, while the orange hand represents the threshold derived from classical methods (in which users know about the technique).

Abstract

Input remapping techniques have been widely explored to allow users in virtual reality to exceed both their own physical abilities, the limitations of physical space, or to facilitate interactions with real-world objects. Often considered is how these techniques can be applied to achieve maximum utility, but still be undetectable to users to maintain a sense of immersion and presence. Existing psychophysical methods used to determine these detection thresholds have known limitations: they are highly conservative lower bounds for detection and do not account for complex usage of the technique. Our work describes and evaluates a method for estimating detection that reduces these limitations and yields meaningful upper bounds. We present the findings of our work where we apply this method to a well-explored hand motion scaling technique. In wholly unaware cases, we determined that users may detect their hand speed as abnormal at around 3.37 times the normal speed, compared to a scale factor of 1.47 that was estimated using traditional methods when users knew the motion scaling was occurring. A considerable number of participants in unaware cases (12 of 56) never detected their hand speed increasing at all, even at the maximum scale factor of 5.0. The study demonstrates just how conservative the thresholds generated by traditional psychophysical methods can be compared to detection during naive usage, and our method can be modified and applied easily to other techniques.

Keywords: virtual reality, interaction techniques, human perception

1 Introduction

A wide variety of interaction techniques exist to support common tasks such as selection, manipulation, and travel in virtual reality (VR). An established question for implementing 3D interaction techniques is whether the interaction should maintain realism compared to the real world and intentionally non-realistic techniques might be preferable [7]. For example, variations of raycasting [4, 35, 57] are commonly adopted methods of selection and manipulation despite lacking realism, and joystick-based techniques are frequently used for travel [10, 44]. While these techniques have clear benefits for practicality and convenience for many forms of applications, users often prefer techniques that maintain higher similarity to real-world interactions—such as those using real hand interaction [28] or physical walking [55]. Realistic and natural forms of interaction are easy to learn since their metaphors are used in the real world (i.e., if you know how to reach or walk in the real world, doing so in VR is trivial) and can increase the user's sense of presence [28, 32, 55].

However, the impact of practical constraints persists in adapting natural interaction techniques in VR. External factors like physical room size, occlusion by other objects, or object weight can affect whether a desired action will be possible. And just as actions in the real world are limited by their personal physical characteristics and motor abilities, certain forms of interaction in VR can be challenging while preserving high-fidelity interaction. In the context of reaching and grabbing an object, if a user's virtual hands perfectly matched their physical hands, these same limitations would carry over into the virtual world.

Thankfully, in VR we are permitted to break the rules of reality and modify the parameters of the virtual world to overcome some of these limitations while still maintaining a high level of realism for interaction metaphor. In this research, we focus on the study of techniques using *input remapping*, which modifies the translation of a user's physical movements into a different set of movements in the virtual world. For example, *hand remapping* is when these translations are applied to a user's hands so that the motion of a virtual hand in VR does not strictly match the movement of the user's physical hand. One such remapped hand technique is Go-Go [41], a generalized reaching technique that applies an increasing offset as a user extends their physical reach. Similar techniques also exist to support natural locomotion in VR by adjusting how a user's head is mapped (e.g., [43]).

However, seeing one's hand extend so far across a room and being able to tell your view is rotated as your walk are not natural by any means, and in cases where immersion and believability are important (such as gaming or training scenarios), it is important for

*e-mail: brett.benda@ufl.edu
†e-mail:brheault@ufl.edu
‡e-mail:linyanna@ufl.edu
§e-mail:eragan@ufl.edu

these augmented interactions to feel natural. Though humans can tolerate a certain amount of proprioceptive drift [54] before they notice a mismatch between the positions of their real and virtual hands or their real and virtual movements, at some point it becomes intolerable. Thus, limits are often imposed on the techniques which keep them within undetectable or unnoticeable ranges. These limits or *detection thresholds* can be estimated by psychophysical methods such as two-alternative forced-choice (2AFC) or Yes/No tasks [22, 30]. These methods have been widely applied to a variety of input remapping techniques, and often consider how other factors such as learning over time [31], the directionality of movement [17, 60], or hand representation [38] can effect detection.

However, for multiple reasons, these methods produce highly conservative thresholds that may be more limiting than what would be detected in a real application used by naive users [5, 17, 50, 60, 61]. Namely, users in these studies are made aware of the techniques they are evaluating while end users of a VR application may not know all the technical underpinning of their software. Additionally, the tasks used to determine these thresholds are focused solely on observing for the presence of a technique rather than taking place in more complex scenarios with a wider range of user goals.

Some work has recently attempted to use indirect means via physiological data (such as EEG, ECG, EDA, and RSP in conjunction with interaction data) to predict detection and identified a high variance in detection based on this data [19]. A notable limitation of this work was also the awareness of the technique occurring, even if the direct intervention found in classical psychophysical methods was not present. Other methods [48] assess techniques not to determine ranges of detectable intensities, but those that preserve immersion, though in those cases participants still experience a direct measurement intervention that is not present with real application usage.

Our work helps bridge this gap and seeks to address these biases of prior detection studies as they relate to the usage of input remapping techniques. In the presented research on detection, we study methodological differences using an experiment of a hand remapping technique that has been previously evaluated with a psychophysical method (scaled movement [17]). We address the research question: How much further can a hand be remapped without detection when a user does not know a hand remapping technique is being applied?

This allows us to provide strong upper-bound limits for detection in cases that may apply more broadly to the usage of natural interactions in real applications. We describe our method to determine when a user has detected that their hand has been remapped subtly and indirectly, and control for the aforementioned biases. We apply the method to a gain-based, fast-scaled hand movement technique due to its robustness under complex motion and utility in extending a user's reach [17, 61], though future work may apply the method to other techniques. We present the findings of our experiment, including detection rates for different amounts of prior information provided to users, comparisons to novel threshold analysis with our task (as well as previously understood thresholds), and implications for the body of detection-related work as it pertains to interaction techniques in VR. We find that remapping may be undetected even at values 180%–230% larger than the detection threshold when users are not aware it is happening.

2 Related Work

We first provide an overview of the wide range of uses for input remapping as a whole, including a discussion of head-based and hand-based techniques. We then cover traditional psychophysical detection studies for remapping techniques (2AFC and Yes/No), common limitations of these studies, and alternative detection methods for similar techniques.

Table 1: A selection of work using psychophysical methods, with techniques and thresholds.

Technique	Threshold Value
Horizontal Hand Scaling [17]	[0.809, 1.310]
Vertical Hand Scaling [17]	[0.869, 1.520]
Depth Hand Scaling [17]	[0.779, 1.380]
3D Hand Scaling [17]	[0.758, 1.430]
Horizontal Hand Offsetting [5]	[Left 10.27 cm, Right 9.40cm]
Vertical Hand Offsetting [5]	[Down 13.37 cm , Up 12.83cm]
Depth Hand Offsetting [5]	[Close 7.83cm, Far 13.25cm]
Horizontal Hand Warping [61]	[Left 4.5 deg, Right 4.5 deg]
Vertical Hand Warping [61]	[Down 4.5 deg, Up 4.5 deg]
Gain-based Hand Warping [61]	[0.88, 1.07]
Translational Walking Gain [50]	[14% slower, 28% faster]
Curvature Gain [50]	22m Circle
Strafing Walking Gain [60]	[5.57 deg Right, 4.68 deg Left]

2.1 Uses of Remapped Interaction Techniques in VR

While our work utilizes a hand remapping technique, it is work exploring both these techniques as well as those that remap the motions of other body parts (e.g., a user's head) as they use the same psychophysical methods.

2.1.1 Head-based Techniques

Input remapping has been applied to a user's virtual head (or view), both regarding its rotation and positioning. These are often used to support travel techniques or help with comfort and search by reducing the physical head movement needed to transform a user's virtual view.

These have largely been leveraged for purposes of travel, such as to enable *redirected walking* [29, 37, 43]. Three main forms of these techniques are *translational gain*, *rotational gain*, *curvature gain*, and *bending gain* [37]. Translation gain amplifies or reduces the translational movement of a user's head in VR [37]. For example, with a gain factor of 2.0 a user who walked 1 meter forward in the real world would have virtually moved 2 meters. Conversely, the same physical movement under a gain factor of 0.5 would have yielded a virtual movement of 0.5 meters. Modifications to the gain factor and direction can also be modified based on the predicted travel direction, as is used by the Seven League Boots metaphor [29]. Similarly, rotational gain scales the amount of virtual rotation a user experiences based on their real-world rotation [37]. Curvature gain applies a virtual rotation to the user's view as they translate position in the real world, with the goal being to have them correct the virtual rotation by physically rotating in the opposite direction [37]. The result is redirection onto a circular physical path, while the user navigates a linear virtual path. Bending gain works similarly to curvature gain, but applies rotation to an already curved path [37]. Another gain-based technique, *strafing gain* [60], exists which works similarly to curvature gains. However, rather than rotate the user's view as they perform translational real-world movements, a small horizontal displacement is applied. This causes a physical correction in the opposite direction, which redirects them onto a diagonal path which has the benefit of maintaining a consistent physical forward direction.

Different head-based techniques with the explicit goals of increasing comfort and enabling better search have also been developed. Bolte et al. applied amplified head rotations to the pitch axis of a user's head (i.e., raising or lowering your head to look up and down) which supports easier vertical search since the user does not need to rotate their head as much [6]. This works by adding additional virtual rotation when a user physically rotates so the visual effect is a larger turn. Ragan et al. [42] later explored rotations in the yaw axis (i.e., turning your head left and right), as well as learning effects of these techniques with consideration for display type and allowed amount of physical turning. Automated head rotations can also be applied to guide users back towards a more comfortable neutral position while maintaining a rotated virtual view [47].

2.1.2 Hand-based Techniques

In the real world, hands are a common means through which humans manipulate their physical environment (e.g., grabbing and moving objects), communicate with others (e.g., pointing and communicating gesturally, etc.), and interact with technology (e.g., typing on a keyboard or using mobile touchscreen devices). Given their widespread use in the real world, VR researchers have sought to develop techniques that maximize natural hand-based interactions in virtual environments. Many methods for hand-based interactions in VR rely on hand remapping. These techniques can be used to support techniques such as *augmented reaching* [41] or *haptic retargeting* [3], which can increase user ability, immersion, and overall satisfaction using VR applications [28].

Hand remapping has often been applied to support object selection and manipulation. Techniques such as Go-Go [41] virtually extend a user's reach by increasing the movement of a hand as the user physically extends their arm. This supports natural reaching and grabbing for selection and manipulation of objects from afar. Other techniques, such as Precise and Rapid Interactions through Scaled Movement (PRISM) support precision in object manipulation by changing the speed at which a selected object moves based on the user's hand speed [21]. These techniques (and others) can also be combined via chaining them sequentially [2].

Another benefit of remapping hands is that they can be used to reduce fatigue (e.g., [20, 36]). An effect of using natural input methods with arms and hands is the "gorilla arm" effect where muscle soreness increases over time [27]. Techniques like Erg-O use space partitioning methods to ideally modify a user's hand speed to different values depending on proximity to known areas of interest [36]. On the other hand, Ownershift simply rotates the user's virtual shoulder position upward so users can reach interfaces in front of their head without the need to fully move their arm forward [20].

Hand remapping can also be used to change movements for training and accessibility purposes. Rather than mapping all reachable physical locations to unique virtual locations, SnapMove maps all physical movements to a single preset virtual path [13]. While this mapping is perhaps not obviously interesting, it has been demonstrated by prior work that users will gradually align their positioning to match that of the avatar, a self-avatar follower effect (SAFE) [25]. If spatial discrepancies exist between a user's physical pose and that of their avatar, they will attempt to minimize it during movement. SnapMove leverages this phenomenon to gradually make users match their movements to that of their avatar. Lilija et al. has also used SAFE to train motor skills [33]. By providing a corrected virtual hand location, a user's physical movements better matched a desired physical path. Similarly, models such as the minimum jerk motion model can be applied to normally mapped hands as well as remapped hands to predict when users will reach a target location which can be used to optimize parameters of different techniques to better support reaching [23].

The inclusion of haptic feedback has strong positive impacts on a user's sense of presence in VR [28], and remapping has been applied to support haptic interactions with passive haptic props. Azmandian et al. identify three main classes of warping techniques used to align virtual objects with their corresponding physical prop: world warping (rotating or translating the virtual space to align the objects), body warping (manipulating the user's movements or positioning via remapping to align the objects), and hybrid warping (using both body and world warping) [3]. Space partitioning methods that modify the positioning of a virtual hand can be used to support haptics, with physical locations of objects as inputs [9]. Based on the mismatch between the virtual object and the prop, the virtual space is warped in order to move the virtual hand to the prop as the user reaches it. Han et al. explored translational and interpolative methods of aligning props [26]. In translational shifting, the virtual hand is linearly offset on the horizontal plane based on the difference between the virtual and real objects. During interpolated reach, the virtual and real hands begin at the same location, but an offset is gradually applied toward the virtual object.

2.1.3 Summary

Input remapping techniques are used for a wide range of purposes, of primary interest to this work are head-based and hand-based techniques. They can be used to overcome real-world limitations, either due to physical space or user ability, and can increase comfort and performance in virtual environments.

2.2 Detection of Remapped Interaction Techniques

While the techniques previously discussed highlight some common uses for input remapping techniques, when used to induce a large enough mismatch between a user's physical and virtual movements they can become distracting or confusing. An issue for techniques that modify a user's head rotations is an increase in simulator sickness [16] (e.g. curvature and rotation gains or amplified head rotations). Regarding hand techniques, there is evidence that even simple techniques like Go-Go [41] and PRISM [21] are less usable than a regular hand at certain higher levels of intensity [59]. For purposes of immersion and perceived realism, we also want to maintain a sense of virtual body ownership. Discontinuities and mismatches in a user's virtual and real bodies can decrease virtual body ownership [53], though a certain amount of mismatch can be tolerated [8].

To address these issues that can arise from remapped technique usage, psychometric detection studies are often performed to determine detection thresholds for each technique. By limiting the "intensity" of the technique to undetectable ranges, we can help a VR application user maintain a better sense of ownership and control.

2.2.1 Detection Methods for VR Interaction Techniques

VR researchers have adapted various methods from psychophysics in order to evaluate their own novel interaction techniques. While psychophysics has developed more robust methods to assess stimulus detection [58], the evaluation of VR interactions has largely used the classical methods. In these classical methods, two types of tasks can be used: *adjustment tasks* where users directly modify a stimulus until they cannot detect it, or *classification tasks* where users are presented a fixed set of stimuli and asked some binary question about it [18]. Adjustment methods are not well suited for evaluation with naive user but can be effective in calibrating a system by someone with technical knowledge (e.g., an experimenter) [18]. Therefore, different types of classification tasks have been used instead. Other factors, such as trial ordering within a task (e.g., increasing/decreasing, randomized, or stair-cased [14]) may also impact the estimated threshold values.

The main type of classification used for assessing the detection of different VR techniques is the *two-alternative forced-choice* (2AFC) methodology [18, 22]. In this design, users are presented with a technique, asked to perform some task while using it (e.g., reaching for an object), then are forced to classify it as one of two options presented to them (e.g., saying that their hand moved to the left or right). Intensities of the technique are varied, and responses to the 2AFC question are then fit to a psychometric sigmoid curve (e.g., the logistic function) which predicts when users will classify each intensity as one of the options. Threshold points can then be determined from this curve, with the 75% detection value and 25% value representing the upper and lower values of "undetectable" technique values.

Another method, Yes/No, instead asks users to state "Yes" if they thought the technique was present or "No" if not [18]. Yes/No methods are useful when concerned with the presence of a stimulus, but they can add some additional bias. [18] Participants may determine their own criteria for cases where they are uncertain (e.g., always defaulting to a "No" or "Yes" if there is any uncertainty). However, this can also apply to 2AFC. However, in cases such as [5] where multiple factors (6 directions) were simultaneously assessed using a Yes/No method simplifies what is being asked to participants. Other biases also exist. For instance, Abrahamyan et al. [1] identified that subjects in psychophysical studies may adopt strategies that are often reinforced over time if they receive feedback.

While these traditional psychophysical methods are beneficial for establishing conservative ranges for technique usage, there are some limitations to the methods that suggest that higher tolerances would be present in more complex and applied tasks. Primarily, prior work examining detection for hand remapping techniques typically introduces the technique to the user in a practice session and makes it directly known that their hand placement is going to be modified in the study. This is necessary for psychophysical methodologies as users must comprehend what it is they will be observing to make accurate judgments, but in real applications, this awareness may not be present. This can be understood as a bias of *awareness*. These methods also place a large emphasis on paying explicit and direct attention to a user's hand during simple and limited movements which induces a large bias of *vigilance*.

Other methods have used participant think-aloud (i.e., asking participants to verbalize their thinking during a study) in conjunction with guided interviews to estimate if a technique has been detected by users [52]. Questionnaires have also been used by having participants rate their agreement with certain statements (with some relating to the technique being evaluated) (e.g., [51] or [45]). However, in both cases, this is only useful for determining *if* a user detected a technique, not necessarily *when* they noticed, which is important for determining a threshold of detection. Nevertheless, these methods are more subtle than psychophysical methods which repeatedly and directly ask participants to examine a stimulus.

Alternative methods and methods besides detection have also been evaluated. Reinforcement learning with adaptive staircase has been demonstrated to produce useful detection thresholds [40]. Rather than investigate detection thresholds, Schmitz et al. [48] instead investigated *thresholds of immersion*, or ranges of remapping that maintain a user's sense of immersion regardless of if the technique is detected. In this methodology, participants are directly asked to report when something "feels strange or unnatural" while remapping is unknowingly applied. In this method, participants lack awareness of the assessed technique but are still vigilant toward detecting unexpected changes. In other cases, *tolerable remapping* has been explored [11].

2.2.2 Detection Thresholds

The previously discussed methods have been applied often to generate thresholds for a range of techniques (a subset are shown in Table 1). Work has also considered other factors that may influence detection with 2AFC studies. Ogawa et al. [38] examined the effects that an abstract hand representation has on detection compared to a realistic one and found that the abstract hand produced less detectable thresholds. Related, avatar representation can also cause users to change their movement behaviors in VR; more realistic avatars can cause users to avoid passing through virtual objects [39]. Different avatars also induce different levels of presence based on gender [49].

In an active approach to decreasing detection, change blindness has also been leveraged [62]. Moderately sized instantaneous shifts of a hand can be applied during a blink, yet remain undetected to users. This can also be applied in conjunction with a continuous technique while the eyes are open for maximum redirection. Effects of directionality regarding both physical movement [12] and technique direction [5, 17] have also been explored. Electrical stimulation of the arm can also be used to decrease detection rates, though this only had a significant effect in female participants [39].

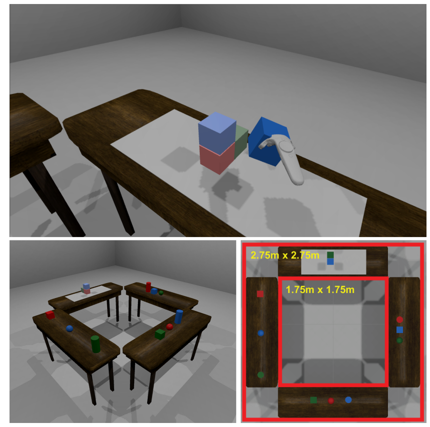

Figure 2: **Top:** the user's view in the experiment. Their virtual hand is represented as the controller they were holding. **Bottom Right:** a perspective view of the environment used in the experiment. **Bottom Left:** A top-down view of the environment with dimensions.

Distraction [61] and task complexity [17] have also been investigated for effects on detection, with distraction proving to increase the undetectable range of detection while task complexity did not. Work by Debarba et al. [15] has also used hand remapping to purposefully increase the difficulty of a task (which also has potential in rehabilitative contexts). In their experiment, participants were less likely to detect a remapped hand if the gains applied made the task easier. Additionally, over time user sensitivity to remapping but may not change, but their own behavior may change to account for the technique [31]. Users can also adapt to the use of techniques over time, becoming more accurate with their movements or more comfortable as the technique changes [26, 46]. Additionally, handedness and direction of remapping can impact detection [24].

3 Experiment

3.1 Research Goals

As noted in our discussion of related work, existing psychophysical methods are limited in their application to VR interactions for multiple reasons. Most significantly, as it relates to hand remapping, they impose a higher-than-normal amount of focus on a user's hand than what occurs in regular application or technique usage. They also instill prior knowledge of the technique in a user that may not be present when the technique is used outside an experimental context. Our work seeks to address these biases of prior detection studies and address the primary research question: How much further can a hand be remapped when a user does not know a hand remapping technique is being applied?

As the basis for the study, we designed an object search and manipulation task to provide the context for using the interaction technique to move objects around in VR during the experiment. The task involved pattern matching with a sequence of trials that displayed different configurations of colored objects, and participants had to move objects to match the patterns (see Figure 2). The experimental manipulations were applied over time to study detection. Because participants were only manipulating virtual objects, this study uses the virtual representation of a controller rather than a realistic hand. We do note that hand representation can affect detection [38].

3.2 Independent Variables

We designed a between-subjects experiment with one independent variable at three levels controlling for the experiment instructional biases (*BIAS*) present in traditional psychophysical methods. We note that our study of bias refers to **bias resulting from the research method** rather than personal or individual cognitive biases of human-subjects participants. The degree of methodological bias was manipulated by providing different sets of instructions to participants at the beginning of the experiment, with the key difference being whether the instructions explicitly informed participants they should be looking for something to change during the trials. The levels of bias examined are:

- *LOW*: Participants were not told their hands would be remapped. This is intended to be representative of application usage by users naive to the technique.
- *MED*: Participants were not told their hand is remapped, and they are told there may be intentional features of the study we were not telling them about but they should look for them. This is intended to instill more vigilance in participants in order to have them report their remapped hand movements more frequently.
- *HIGH*: Participants were made aware of the technique being applied, and were directly asked at the end of each trial if their hand felt "normal" or "not normal". This matches a traditional Yes/No design.

3.3 Measures

The experiment studies how the levels of *BIAS* affect:

- **Detection Rates**: how frequently participants report a remapped hand (regardless of *when*)
- **Earliest Detected Values**: the technique intensity at which participants first report a remapped hand

Specifically, we evaluated the effects of *BIAS* on the same scaled-movement technique described in [17], which is also similar to the gain-based technique described in [61]. This technique applies a variable scale factor to the user's hand in order to accelerate movements and consequentially increase reach (see Figure 3). In the technique, a user's virtual hand is placed according to Algorithm 1. We approximate the scale origin as a location 30 cm below the user's head location. This technique supports generic, target-agnostic reaching which is beneficial in cases where users want to select or manipulate an object out of their natural reaching distance and has known thresholds derived from an experiment using classical methods [17] we can compare to. We also conduct our own threshold analysis of a classical method (our *HIGH* bias condition) to verify these thresholds hold for our modified task and scale order. Specifically, the set of scale factors utilized was:

- **Scale Factor:** 1.0, 1.1, 1.2, 1.3, 1.4, 1.5, 1.6, 1.7, 1.8, 1.9, 2.0, 2.2, 2.4, 2.6, 2.8, 3.0, 3.33, 3.67, 4.0, 4.5, 5.0

We begin with small increments between scale factors close to 1.0, and increase this gap between higher scale factors.

3.4 Hypotheses

As it relates to the scaled movement technique, our experimental hypotheses are as follows:

- **H1 (Detection Rates):** *HIGH* $>$ *MED* $>$ *LOW*
- **H2 (Detection Values):** *LOW* $>$ *MED* $>$ *HIGH*

In **H1**, we expect to see more participants detect the technique when told to look for something unexpected (i.e., their remapped hand). Necessarily, all participants in the *HIGH* case will report their hand speed being abnormal. In **H2**, we expect the reported scale factors to be highest when no prior knowledge of the technique was given, followed by the scales reported when participants are vigilant. In our *HIGH* case, we expect to see higher sensitivity (lower scales being detected more often).

Algorithm 1 Scaled Hand Movement

Input: the scale origin O, the user's real hand position P_r, the scale factor s

Output: the user's scaled virtual hand position P_v

$\vec{d}_r \leftarrow O - P_r$ ▷ Get the offset between the origin and real hand.

$P_v \leftarrow O + s * \vec{d}_r$ ▷ Add the scaled offset to the origin.

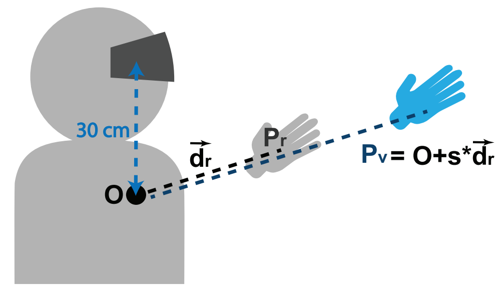

Figure 3: Illustration of the technique evaluated in this work. Scaling is applied from an origin point 30 cm below the user's head. The virtual hand (P_v) is scaled along the vector between the user's real hand and the origin (P_r) by a factor of s.

3.5 Design and Methodology

In this section, we describe the method used to detect when participants detected their increased hand speed. Because in both the *LOW* and *MED* conditions participants are not aware of the technique, it is necessary to design the task and experiment with elements that encourage participants to speak up if they notice something unexpected (in this case, their hand speed increasing). The four key elements of the design are **1)** using a deceptive task, **2)** gradually increasing the scaling of the remapped hand, **3)** having participants think-aloud during the experiment, and **4)** including intentional errors for participants to comment on.

A general overview of the method is as follows:

1. Participants were provided a "dummy" task that incorporates common interactions in VR (with a focus on visual search, object selection, and object manipulation). They complete this task without the awareness their hand movement is scaled.
2. Between trials, their hand speed is gradually increased without their knowledge.
3. While completing the task, they are instructed to think aloud about the task and other things they are observing. The trial and speed where they first comment on their hand is used to estimate the scale factor at which they detected the technique.
4. Intentional application/task errors are included early in the experiment to subtly encourage participants to report anything that is odd or abnormal (ideally, their hand speed).

3.5.1 User Task and Deception

First, in order to hide that we were scaling their hand speed, we utilized deception (approved by our organization's Institutional Review Board) so that participants were not informed of the focus on detection of interaction techniques. In our experiment, participants were told that we were examining how object placements and groupings affect search time and object manipulation. Specifically, the task was to examine a pattern consisting of objects with different shapes and colors, search for those objects on tables placed around them, grab them, then put them into the pattern. Once finished, a new pattern was loaded during a short fade to black (during which their hand speed was modified while not visible to them).

This task was chosen in order to implement use cases where hand remapping is useful (selection and manipulation of objects [32]) and other tasks common in VR (visual search and movement [32]).

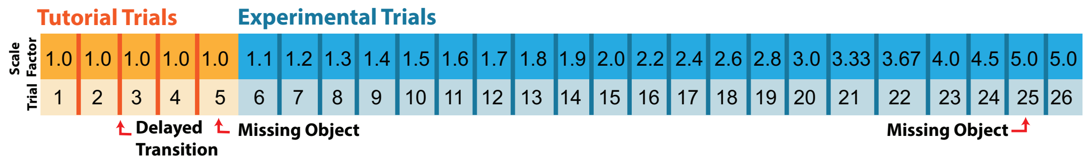

Figure 4: Ordering of scales in the experiments, and location of intentional application errors. These errors were to reinforce participant think aloud behavior and subtly encourage them to comment on things that were abnormal or strange that were not a stated part of the deceptive task (of interest to us, their hand speed changing). During the delayed transition, the fade effect between trials held indefinitely until the participant commented on it to the experimenter. During a trial with a missing object, an object belonging to the pattern was missing which prevented the participant from advancing until it was mentioned to the experimenter.

Other factors were also included in the design to provide a more complex task than those often used in 2AFC studies, including: 1. a precision constraint for how closely the objects had to match the pattern (5cm) and 2. 360 degree item placement around the user to require turning and small walking movements to reach objects when scaling was not applied

These ensure that the benefits/drawbacks of scaling the hand speed are observed subconsciously by participants. Over time, it becomes difficult to be precise as small physical movements become amplified by increasing amounts. However, simultaneously less walking and reaching are needed to grab the objects.

Participants were also told that the patterns would increase in difficulty over time, but all patterns observed while the movement scaling was applied were designed to be similar in difficulty (required searching for and moving four objects) and were modified versions of a few template patterns to maintain difficulty of patterns throughout the experiment. The patterns also accounted for the deceptive reasoning for the experiment: in some trials, objects were grouped by color, others by shape, and others randomly.

3.5.2 Gradually Increasing Scale Factor

Because participants are now aware the scaling is being applied, the order in which the scales are experienced by users can also affect detection. We carefully order them in an ascending order to prevent premature detection at a scale factor exceeding the participant's minimally detected scale. If a randomized ordering was used, then participants may be biased by seeing a large intensity early in the experiment causing them to become more vigilant for it later. Consider a participant who, by chance, saw their hand moving at five times the speed in the first trial: they become aware that this is something to look out for and pay more attention to their hand in the future. The ordering used in our study can be seen in Figure 4. We note that increases in the scale are not consistent; as the trials progress we add more scale to the hand between trials.

3.5.3 Think-Aloud

In our experiment, we are not directly intervening in every trial as is done in traditional 2AFC or Yes/No methods in order to reduce vigilance. Thus, we rely on participant utterances to determine if and when they detect the stimulus. We used a think-aloud [56] method during the experiment, where participants are asked to verbally explain their thought processes and observations as they complete a task or exercise. Eventually, given a hand moving at a noticeable speed, participants should comment on it.

Prior to beginning the trials, participants were asked to comment on things they were seeing and noticing, if the task felt easier or harder as they progressed, and if anything else came to mind. This is also where different levels of *BIAS* were instilled in users. In the *MED* case, participants received the same think-aloud prompt as participants in *LOW* but were also told that there may be aspects of the experiment they were not being told about that they should try to identify.

3.5.4 Inclusion of Intentional Application Errors

A risk of relying on a participant's think-aloud is that they may fixate on only commenting on elements of the task they are presented with (in this case, only about searching for and moving the objects to the pattern), rather than the stimulus. To counter this, purposeful errors and uncanny elements were added to the experiment to subtly communicate to the participant that they should also comment on anything unexpected (i.e., their hand moving quicker). Specifically, whenever one of these intentional glitches was reported, the experimenters responded, "Thank you for letting me know about that. Please let me know if anything else strange or abnormal happens.". This provides a subtle prompt to users to report anything which may also be perceived as an error like their hand moving faster than normal.

Specifically, we chose two error types that prevented participants from proceeding without mentioning it to the experimenter: a delayed transition during the fade-to-black between trials where the scene never reappeared (i.e., the application blacked out indefinitely), and missing virtual objects that were necessary to complete the trial. The delayed transition occurred between T2 and T3 after having seen the quick transition occur after the first trial. A missing object is present in T5, to reinforce the same behavior. Finally, another missing object is included in T25 to serve as a final check that participants have nothing abnormal to report. T26 repeats the final scale, in the event that the participant first reports the missing object before manipulating any objects.

3.6 Procedure

The procedure and deceptive nature of our study were approved by our Institutional Review Board (IRB). One hour was allotted for each participant to complete the study. When a participant arrived, they were asked to read through and sign a consent form. This form listed the false purpose of the study but described the task, risks, and benefits accurately.

Participants were then aided in donning the headset and handheld controller and told to begin completing the trials. To identify detection thresholds of the scaled hand movement participants were instructed to think aloud and verbally explain their thought process about things they were seeing and noticing, if the task felt easier or harder as they went on, and if anything else came to mind (*LOW BIAS*). In the second condition, participants were also told that there may be parts of the study they were not told about and to be on the lookout for anything unexpected or abnormal (*MED BIAS*). In keeping with standard think-aloud practices [56], if the participant stopped thinking aloud for an extended period of time during the experiment the investigator would remind them to continue thinking aloud, with no additional suggestion of what to talk about as to not bias the participant.. Follow-up questions were also asked if participant utterances were unclear or lacking in detail (e.g., if the participant uttered *"This is hard."*, they may be prompted to elaborate with more details with *"How is it hard?"*).

When a participant reached a stage with a built-in glitch the investigator would wait for them to verbally address the problem. Once it was addressed, the investigator would express confusion at the "glitch", apologize for the inconvenience and inform the participant that they would manually skip that trial, before telling the participant to continue addressing anything out of the ordinary in future trials. While the participant completed the trials, the investigator took notes of participant remarks that related to scaled hand movement, and the trial at which the remarks were made.

When the participant had completed the final trial they were asked to remove the headset and take a seat at a desktop computer to fill out a post-survey. This survey collected demographic data and responses to eight Likert questions asking to what extent participants noticed various abnormalities. These included irrelevant abnormalities that should not have been experienced such as hearing unexpected sounds in the virtual world, alongside relevant abnormalities such as strangeness in the participant's movements in the virtual world. This survey was followed up by verbal questions asked by the investigator to get detailed feedback on whether the participant perceived the task to get easier or harder as it went on, why they felt that way, and anything that felt strange. This also served as one last chance for participants who had not verbally addressed the hand scaling to comment on it.

Participants generally completed the entire procedure within 40–60 minutes. Immediately afterward, the investigator debriefed the participants on the true purpose of the study and told them of the hand scaling that had occurred. If the participant had not yet remarked on the hand scaling either implicitly or explicitly at any point during the study, they were asked if they recalled it in retrospect.

In the *HIGH* case, since deception and subtlety were not required, participants were provided with examples of a normal and fast hand before starting the experimental trials. They were instructed that after each trial they would be asked "Did your hand speed feel normal or not normal?", and that they would have to pick one of those choices. They were unaware that the scaled hand speed would be observed in a strictly increasing order. They also completed each trial without the intention glitches included for *LOW* and *MED*. Upon completion of the trial, the scene faded to black and they were then asked if their hand speed in the previous trial felt "normal" or "not normal". This was logged by the experimenter, the scene was restored, and the next trial started.

3.7 VR System and Environment

The VR system we utilized for the experiment consisted of a HTC Vive Pro Eye Office headset wired directly to the computer running the Unity project and tracked using three SteamVR base stations. The headset had a resolution of 1440×1600 pixels per eye at 615 PPI, a 110 degree horizontal field of view, and the system had a refresh rate of 60 FPS. The study application was run through the Unity game engine with the center of the virtual space aligned to be physically centered over an "X" marked in tape on the floor of the experiment room. Participants were instructed to start on "X". A six-foot radius was cleared around the center for participants to be able to move and reach toward any of the tables as needed.

The virtual environment was simple and utilized flat textures for everything except the tables, which used a realistic wooden texture. Several point lights that cast shadows were included to light the scene from above. See Figure 2 for images from the application.

Because users were manipulating only virtual objects, the representation of the virtual hand in VR was limited to the virtual representation of the controller.

3.8 Participants

We conducted this experiment with a total of 77 participants (self-reported as 50 male, 26 female, and 1 subject not willing to disclose gender identity) (*LOW*=28, *MED*=28, *HIGH*=21). All participants were conducted from undergraduate or graduate-level computer science and human-centered computing courses. Participant ages ranged from 18 to 37, with a median of 21 years. Participants self-reported handedness with 70 right-hand dominant, 6 left-hand dominant (2 in *LOW*, 1 in *MED*, 3 in *HIGH*, and 1 ambidextrous (*LOW)*. All completed the experiment with their right hand, though we note that handedness can impact detection [24]. Additionally, the gender, age, and educational background may also limit transfer to other populations.

4 Results

4.1 Yes/No Threshold Analysis

To establish a baseline comparison to classical psychophysical methods, we first conducted threshold analysis on our *HIGH BIAS* case to produce a meaningful point of comparison to our lower levels of *BIAS*. This analysis is necessary due to the different task and experiment designs used by Esmaeili et al. [17]. Our trials are longer and more complex and have increased exposure to the technique which may affect detection. Additionally, we cannot use a randomized order to evaluate *LOW* and *MED BIAS*, so this analysis also accounts for the constantly increasing scale ordering used.

We utilized the *quickpsy* [34] package in R to test several fits of curves. The logistic function had the minimal AIC of all candidate curves fit and was ultimately chosen to estimate this threshold. Of interest to us here is the scale factor that yields a 50% chance of a user detecting that their hand speed was not normal; exceeding this value indicates that users will be *more likely* to detect the hand speed is being scaled than not.

Fitting the curve to the data provided a threshold value of 1.47 (see Figure 5), with a 95% confidence interval of between 1.39 and 1.54. In comparison to Esmaeili's threshold, this value is slightly higher (1.47 vs. 1.43 [17]) and both confidence intervals overlap significantly (1.39–1.54 vs. 1.38–1.48 [17]). In the remainder of the paper, we compare directly against the threshold value from our work.

4.2 Detection Counts

The previous section covers the core analysis of the *HIGH BIAS* case. To analyze the *LOW* and *MED* bias cases, we first conducted an analysis of the counts for each type of detection under each condition (seen in Table 2). We are interested in seeing if the different levels of bias resulted in more participants detecting the technique (at any level). Participants in the *MED* case were told there were parts of the may be other features of the study they were not told about and to pay attention for anything unexpected or abnormal, so we would expect more participants to report that the hand scaling was present compared to *LOW* since the changing hand speed would be something unexpected.

To determine whether participants detected the remapping in our study method, we analyzed all participant utterances from the think-aloud recordings. After completing the experiment, we classified participants into one of three categories regarding detection based on their comments and responses during the experiment and debrief. An iterative coding method was used to determine these categories. We started with *a priori* codes for *detected* and *not detected*, and then early iterations expanded to four categories: 1) explicit detection of the technique (e.g., comments directly about the hand speed changing), 2) implicit detection (e.g., comments related to other "side effects" of the technique but not about the hand speed), 3) no detection, and 4) indeterminate. Ultimately, we decided to join the implicit cases with the explicit comments into a single group ("detected") for analysis, though future work may consider more carefully these user conceptualizations of remapping. The final three categories of detection types are:

- **Detected:** The participant either **explicitly verbalized seeing the technique** (e.g., "*My hand looks really far away...*" or **verbalized an indirect effect of the technique** (e.g., "*I'm having a really hard time placing these objects.*" or "*I kinda just have to stay in place... I don't have to move as much.*").

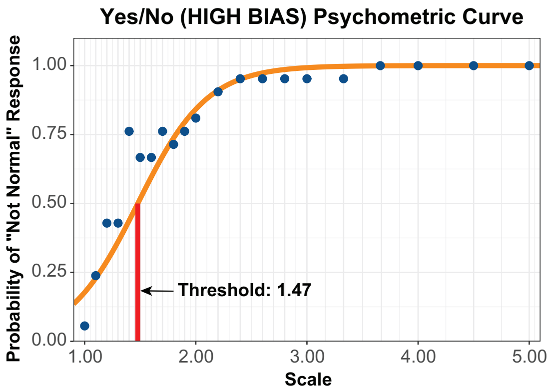

Figure 5: Detection threshold analysis of our Yes/No condition (*HIGH BIAS*). Presented are average detection probabilities at each scale, with the fit logistic curve and threshold marked as a dropdown line.

Table 2: Counts of Detection Types.

	Detected	Not Detected	Indeterminate	Total
LOW	20	5	3	28
MED	14	7	7	28
HIGH	21	0	0	21
Total	55	12	10	77

Table 3: Means, Medians, Standard Deviations, and 95% Confidence Intervals of Earliest Detected Scale Factors.

	Mean	Median	SD	95% CI
LOW	3.37	3.33	1.04	[2.91, 3.82]
MED	2.70	2.5	0.84	[2.27, 3.14]
HIGH	1.39	1.2	0.42	[1.22, 1.57]

- **Not Detected:** The participant did not verbalize anything about their hand during the experiment, and after being explicitly informed of the technique's presence during the debrief, the participant affirmed they did not notice it.
- **Indeterminate:** The participant did not verbalize anything related to their hand or interaction during the experiment, but after being informed of the technique's presence during the debrief, they indicated they did notice but did not feel the need to comment on it.

Pearson's Chi-Squared Test was used to determine if the distributions of detection classifications significantly varied between *LOW* and *MED*. The *HIGH* condition was excluded, as due to the method participants would necessarily detect the technique at some point.

As stated in our first hypothesis **H1**, we expect that when participants are told explicitly to report anything that is strange or abnormal more would detect the technique. The result of the test was $\chi^2 = 2.99$, df $= 2$, $p = 0.224$, which suggests that the type of bias induced in participants did not yield significantly different rates of detection; even when made vigilant for something unexpected (an increasing hand speed), participants did not comment on it more frequently. Because of this, we partially reject **H1**. Both *LOW* and *MED* did have lower levels of detection compared to *HIGH*, which logically follows from the experimental structure of the *HIGH* condition (all participants were aware of the technique and directly asked about it every trial so they necessarily must detect it at some point).

4.3 Earliest Detected Scale Factors

Next, we sought to determine whether the types of biases induced significantly affected the scale factor at which participants *first* noticed the hand scaling. Based on when the participant first uttered a statement suggesting detection (for *LOW* and *MED*) or the first time they responded their hand felt not normal (for *HIGH*), we consider the first scale reported as "not normal". This is intended to approximate the point where users *start* to notice the technique; they may not be confident, but it suggests that they have noticed it. Only data from participants who detected the technique were included in this analysis as they were the only ones who produced reliable, known detected values. The earliest detected scale factors for participants are shown in Figure 6, and averages among participants in each condition in Table 3.

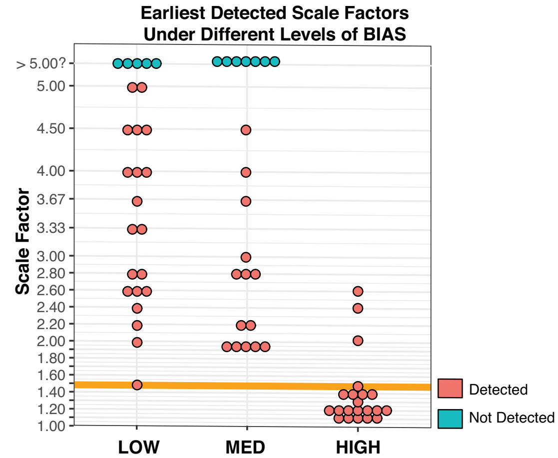

Figure 6: Earliest Detected Scale Factors by *BIAS* condition. Each dot represents the earliest detected scale for a participant, with "Not Detected" cases existing at some unknown point above our maximum scale of 5.0. The orange line at 1.47 represents the detection threshold determined by analysis of the *HIGH BIAS* case in Section 4.1.

Given the unequal group sizes, we utilized a non-parametric Kruskal-Wallis test to test for differences in these earliest detected scale factors between the levels of *BIAS*. Only data points from participants who detected the technique were included. The test result ($H = 34.35$, $p < 0.001$) suggests there is a difference significant difference between the groups. A pairwise Nemenyi test was used to determine pairwise differences; *HIGH* was significantly lower than both *LOW* and *MED* (both at $p < 0.001$), while *LOW* was not significantly different than *MED* ($p = 0.404$).

Then, individual one-sided Wilcoxon tests were used to compare these earliest detected scale factor against our estimated threshold of 1.47: *LOW*: $V = 210$, $p < 0.001$, *MED*: $V = 105$, $p < 0.01$, *HIGH*: $V = 61$, $p = 0.059$ We see that for the *HIGH* case, the average value is not significantly different from the threshold. For *LOW* and *MED*, these are significantly higher (this can be seen graphically in Figure 6).

Overall, **H2** is partially rejected. Though no differences between the earliest detected scale factors were observed between our *LOW* and *MED* cases, both were significantly lower than those found in the *HIGH* case. Compared to the threshold determined in Section 4.1, the converse is true. *LOW* and *MED* earliest detected scale factors are distributed well above the threshold, while *HIGH* is not.

5 Discussion

This study demonstrates how different levels of experimental bias (in the form of the prior awareness participants have regarding what is being evaluated) can impact detection rates and values. We considered a hand remapping technique where the user's hand moves at increasingly larger speeds, and found that when users were not aware that this was happening detection occurred at much higher values than when they were made aware.

5.1 Scaled Hand Movement Detection

Most importantly, this work provides much-needed upper bounds for detection that traditional psychophysical methods cannot estimate.

These methods require prior awareness of a technique and induce a higher level of attention toward the remapped interaction than in situ technique usage outside an experiment. While other work has examined the perception of remapping in looser contexts [48], our work further lessened these constraints (via the compared differences in procedure to provide varying levels of *BIAS)*, and did not make users aware of the technique.

We show that when unaware their hand is being remapped, users can unknowingly use the technique to a much larger extent than is estimated by 2AFC or Yes/No methods. We evaluated scaled hand movement and found that you can move a user's hand up to 3.37 times its normal speed on average before participants notice it on average when they were not made aware of its application. Compared to our threshold estimate from a classical Yes/No (*HIGH BIAS*) method (1.47), detection during cases lessened (*MED BIAS*) and no (*LOW BIAS*) amount of prior awareness was significantly higher with detection starting at approximately scales of 2.70 and 3.37 respectively. Compared to our estimated detection threshold (1.47), the *LOW* and *MED* average earliest detected values are 180% and 230% larger. Prior work using 2AFC and Yes/No methods often main claims that their thresholds are conservative (e.g., [5, 17, 50, 60, 61], but until now it has been difficult to understand exactly how restrictive they are.

Practically, this could translate into more utility from remapping techniques due to greater flexibility for deviation from realistic input-to-output mappings. For scaled movement, less physical movement is necessary when used at this upper bound which can increase the ease of reaching for users, while also letting them benefit from secondary effects like less physical movement being needed to grab something, lessened fatigue, and increased immersion through natural interactions. For other techniques, like those for walking, using upper bound values determined with similar experiments could also increase their utility and further reduce the physical space needed to simulate large virtual environments.

We also observed a wide range of reported detected scale factors, ranging from 1.5 to 5.0. All participants reported the technique occurring after the value of 1.47. Most interesting are the cases where participants never detected the technique. Approximately 1 in 5 participants in the *MED*/*LOW* conditions (12 of 56) completed the entire study without commenting on their hand moving faster and later confirmed they did not notice their hand moving faster after being made aware the technique was applied. This is quite surprising, and we suspect the engagement and focus of participants of the "dummy" object manipulation task sufficiently engaged or distracted these users from detecting their hand.

In summary, our key findings are: **1. Unaware users start detecting techniques at intensities much higher than detection thresholds.** Remapping may be undetected even at values 180%–230% larger than the detection threshold for the technique if users are brought up to that value over time. **2. Unaware users are widely varied in detection.** Even at our maximum value, 1-in-5 users did not notice the remapping being applied.

5.2 Limitations and Future Work

Our work represents a novel approach for estimating the detection of remapped interaction techniques, but we identify opportunities for further advancement of knowledge of detection through future research. For instance, there are trade-offs between using think-aloud utterances to generate detection points and the more direct data collection methods used by psychophysical methods. Think-alouds may introduce waste or delay in reporting, but allows us to evaluate technique usage without participants being aware. Repeated and direct asking of a 2AFC or Yes/No question posed to users every trial guarantees data collection, but these methods are also biased between users [18, 22, 30] and cannot evaluate techniques without participants being aware. Regarding our usage of think-aloud, 10 of 56 participants in the *LOW* and *MED BIAS* cases (approximately 18%) stated after the experiment that they did notice the technique but did not comment on it during the trials. However, a majority of participants did produce useful data points; as the technique intensity increased, 34 of 56 (approximately 61%) of participants reported noticing the technique during the experiment and 12 of 56 (approximately 21%) denied noticing the technique until after being told it was present. Combined, this is roughly 82% of participants in the *LOW* and *MED BIAS* cases.

As part of this design, we also examined possible acclimation effects of using a scaled hand that has been seen in other techniques (e.g., [26, 46]). As stated in Section 3.5.2, to find the smallest detected scale factor it was necessary to slowly increase the hand speed to prevent participants from prematurely becoming aware of hand modifications. If a participant saw their hand moving at the maximum scale factor in an early trial, then they would know to pay attention to their hand in later trials. Thus, our upper bounds may be detectable if applied immediately at the start of application usage. However, in our Yes/No condition, users still experienced an increasing hand speed order and we did not see drastically different thresholds than Esmaeili et. al [17] who used a random order. Regardless, work should examine how quickly the scale factor can increase to a threshold value without being noticeable.

Finally, this method should be refined further and generalized to accommodate the evaluation of other techniques. As discussed in Section 3.5.1, we chose to include a variety of commonly used tasks in VR (selection, manipulation, search, and navigation [32]). As is, this design may be applied to rotation gain applied to the head [37] since our task design requires 360-degree movement. Other techniques, such as those for travel, may require modifications such as moving the objects further away from the pattern to require walking.

6 Conclusion

The constraints of traditional detection methodologies (2AFC and Yes/No) limit the direct transfer of threshold values to more realistic cases. Under these methods, participants are explicitly made aware of the technique and their focus is on making a judgment regarding it. In applied scenarios, this is not the case; users are typically focused on a different primary task, and they may not actually know the technique is being used.

To address these limits, we designed and conducted an experiment to determine reasonable upper bounds for the detection of a scaled hand movement technique in virtual reality that may be used in higher-level applications and scenarios. In our experiment, participants were not told that their hand speed would be increasing over time (from a scale factor of 1.0 to 5.0) but were asked to think aloud about their observations while completing a pattern matching and search task. In the condition where users had the lowest amount of awareness, the average earliest detected scale factor was 3.37, which is a significantly larger value than both the Yes/No threshold we derived (1.47) and the average earliest detected scale value for the Yes/No condition (1.39).

Overall, this work starts to bridge the gap between highly controlled evaluations of remapping techniques and technique usage in realistic applications. The presented method simulates cases where users do not know remapped interaction techniques are used and demonstrates that detection in these cases varies greatly from detection in the limited scenarios previously investigated.

Acknowledgments

This work was supported in part by the DARPA Perceptually-enabled Task Guidance (PTG) Program under HR00112220005.

REFERENCES

[1] A. Abrahamyan, L. L. Silva, S. C. Dakin, M. Carandini, and J. L. Gardner. Adaptable history biases in human perceptual decisions. *Proceedings of the National Academy of Sciences*, 113(25):E3548–E3557, 2016.

[2] C. Auteri, M. Guerra, and S. Frees. Increasing precision for extended reach 3d manipulation. *International Journal of Virtual Reality*, 12(1):66–73, 2013.

[3] 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. ACM, New York, NY, USA, may 2016. doi: 10.1145/2858036.2858226

[4] M. Baloup, T. Pietrzak, and G. Casiez. Raycursor: A 3d pointing facilitation technique based on raycasting. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pp. 1–12, 2019.

[5] B. Benda, S. Esmaeili, and E. D. Ragan. Determining Detection Thresholds for Fixed Positional Offsets for Virtual Hand Remapping in Virtual Reality. In *2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 269–278. IEEE, nov 2020. doi: 10.1109/ISMAR50242.2020.00050

[6] B. Bolte, G. Bruder, F. Steinicke, K. Hinrichs, and M. Lappe. Augmentation techniques for efficient exploration in head-mounted display environments. In *Proceedings of the 17th ACM Symposium on Virtual Reality Software and Technology*, pp. 11–18, 2010.

[7] D. A. Bowman, R. P. McMahan, and E. D. Ragan. Questioning naturalism in 3d user interfaces. *Communications of the ACM*, 55(9):78–88, 2012.

[8] E. Burns, M. Whitton, S. Razzaque, M. McCallus, A. Panter, and F. Brooks. The hand is slower than the eye: a quantitative exploration of visual dominance over proprioception. *IEEE Proceedings. VR 2005. Virtual Reality, 2005.*, (June 2015):3–10, 2005. doi: 10.1109/VR.2005.1492747

[9] C. Carvalheiro, R. Nóbrega, H. da Silva, and R. Rodrigues. User Redirection and Direct Haptics in Virtual Environments. In *Proceedings of the 24th ACM international conference on Multimedia*, pp. 1146–1155. ACM, New York, NY, USA, oct 2016. doi: 10.1145/2964284.2964293

[10] S. S. Chance, F. Gaunet, A. C. Beall, and J. M. Loomis. Locomotion mode affects the updating of objects encountered during travel: The contribution of vestibular and proprioceptive inputs to path integration. *Presence: Teleoperators and Virtual Environments*, 7(2):168–178, 1998.

[11] 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.

[12] A. Clarence, J. Knibbe, M. Cordeil, and M. Wybrow. Investigating the effect of direction on the limits of haptic retargeting. In *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 612–621. IEEE, 2022.

[13] B. A. Cohn, A. Maselli, E. Ofek, and M. Gonzalez-Franco. SnapMove: Movement Projection Mapping in Virtual Reality. In *2020 IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR)*, pp. 74–81. IEEE, dec 2020. doi: 10.1109/AIVR50618.2020.00024

[14] T. N. Cornsweet. The staircase-method in psychophysics. *The American Journal of Psychology*, 75(3):485–491, 1962.

[15] H. G. Debarba, J.-N. Khoury, S. Perrin, B. Herbelin, and R. Boulic. Perception of Redirected Pointing Precision in Immersive Virtual Reality. In *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, number March, pp. 341–346. IEEE, mar 2018. doi: 10.1109/VR.2018.8448285

[16] M. S. Dennison and D. M. Krum. Unifying research to address motion sickness. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 1858–1859. IEEE, 2019.

[17] 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, 2020. doi: 10.1109/VR46266.2020.00066

[18] B. Farell and D. G. Pelli. Psychophysical methods, or how to measure a threshold and why. *Vision research: A practical guide to laboratory methods*, 5:129–136, 1999.

[19] M. Feick, K. P. Regitz, A. Tang, T. Jungbluth, M. Rekrut, and A. Krüger. Investigating noticeable hand redirection in virtual reality using physiological and interaction data. In *2023 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*, pp. 194–204. IEEE, 2023.

[20] T. Feuchtner and J. Müller. Ownershift: Facilitating overhead interaction in virtual reality with an ownership-preserving hand space shift. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*, pp. 31–43, 2018.

[21] S. Frees and G. Kessler. Precise and rapid interaction through scaled manipulation in immersive virtual environments. In *IEEE Proceedings. VR 2005. Virtual Reality, 2005.*, pp. 99–106. IEEE, 2005. doi: 10.1109/VR.2005.1492759

[22] G. A. Gescheider. *Psychophysics*. Psychology Press, 1997. doi: 10.4324/9780203774458

[23] E. J. Gonzalez, P. Abtahi, and S. Follmer. Evaluating the Minimum Jerk Motion Model for Redirected Reach in Virtual Reality. In *The Adjunct Publication of the 32nd Annual ACM Symposium on User Interface Software and Technology*, vol. 0, pp. 4–6. ACM, New York, NY, USA, oct 2019. doi: 10.1145/3332167.3357096

[24] 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.

[25] M. Gonzalez-Franco, B. Cohn, E. Ofek, D. Burin, and A. Maselli. The Self-Avatar Follower Effect in Virtual Reality. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 18–25. IEEE, mar 2020. doi: 10.1109/VR46266.2020.1580500165557

[26] 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, apr 2018. doi: 10.1109/TVCG.2018.2794659

[27] J. D. Hincapié-Ramos, X. Guo, P. Moghadasian, and P. Irani. Consumed endurance: a metric to quantify arm fatigue of mid-air interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1063–1072, 2014.

[28] Insko and B. Edward. Passive haptics significantly enhances virtual environments. 2001.

[29] V. Interrante, B. Ries, and L. Anderson. Seven league boots: A new metaphor for augmented locomotion through moderately large scale immersive virtual environments. In *2007 IEEE Symposium on 3D User interfaces*. IEEE, 2007.

[30] F. A. A. Kingdom and N. Prins. *Psychophysics: A Practical Introduction*. Academic, 2010.

[31] K. Kohm, J. Porter, and A. Robb. Sensitivity to hand offsets and related behavior in virtual environments over time. *ACM Transactions on Applied Perception*, 19(4):1–15, 2022.

[32] J. J. LaViola Jr, E. Kruijff, R. P. McMahan, D. Bowman, and I. P. Poupyrev. *3D user interfaces: theory and practice*. Addison-Wesley Professional, 2017.

[33] K. Lilija, S. Kyllingsbaek, and K. Hornbaek. Correction of avatar hand movements supports learning of a motor skill. *Proceedings - 2021 IEEE Conference on Virtual Reality and 3D User Interfaces, VR 2021*, pp. 455–462, 2021. doi: 10.1109/VR50410.2021.00069

[34] D. Linares and J. López-Moliner. quickpsy: An r package to fit psychometric functions for multiple groups. *The R Journal*, 8(1):122–131, 2016.

[35] Y. Lu, C. Yu, and Y. Shi. Investigating bubble mechanism for raycasting to improve 3d target acquisition in virtual reality. In *2020 IEEE Conference on virtual reality and 3D user interfaces (VR)*, pp. 35–43. IEEE, 2020.

[36] R. A. Montano Murillo, S. Subramanian, and D. M. Plasencia. Erg-O: Ergonomic optimization of immersive virtual environments. *UIST 2017 - Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*, pp. 759–771, 2017. doi: 10.1145/3126594.3126605

[37] N. C. Nilsson, T. Peck, G. Bruder, E. Hodgson, S. Serafin, M. Whitton, F. Steinicke, and E. S. Rosenberg. 15 years of research on redirected walking in immersive virtual environments. *IEEE computer graphics and applications*, 38(2):44–56, 2018.

[38] N. Ogawa, T. Narumi, and M. Hirose. Effect of Avatar Appearance on Detection Thresholds for Remapped Hand Movements. *IEEE Transactions on Visualization and Computer Graphics*, 27(7):3182–3197, jul 2021. doi: 10.1109/TVCG.2020.2964758

[39] N. Ogawa, T. Narumi, H. Kuzuoka, and M. Hirose. Do you feel like passing through walls?: Effect of self-avatar appearance on facilitating realistic behavior in virtual environments. In *Proceedings of the 2020 CHI conference on human factors in computing systems*, pp. 1–14, 2020.

[40] T. Porssut, Y. Hou, O. Blanke, B. Herbelin, and R. Boulic. Adapting virtual embodiment through reinforcement learning. *IEEE Transactions on Visualization and Computer Graphics*, 28(9):3193–3205, 2021.

[41] I. Poupyrev, M. Billinghurst, S. Weghorst, and T. Ichikawa. The Go-Go Interaction Technique: Non-linear Mapping for Direct Manipulation in VR. In *Proceedings of the 9th annual ACM symposium on User interface software and technology - UIST '96*, pp. 79–80. ACM Press, New York, New York, USA, 1996. doi: 10.1145/237091.237102

[42] E. D. Ragan, S. Scerbo, F. Bacim, and D. A. Bowman. Amplified head rotation in virtual reality and the effects on 3d search, training transfer, and spatial orientation. *IEEE transactions on visualization and computer graphics*, 23(8):1880–1895, 2016.

[43] S. Razzaque. *Redirected walking*. The University of North Carolina at Chapel Hill, 2005.

[44] B. E. Riecke and D. Feuereissen. To move or not to move: can active control and user-driven motion cueing enhance self-motion perception (vection) in virtual reality? In *Proceedings of the ACM Symposium on Applied Perception*, pp. 17–24. ACM, 2012.

[45] M. Rietzler, J. Gugenheimer, T. Hirzle, M. Deubzer, E. Langbehn, and E. Rukzio. Rethinking redirected walking: On the use of curvature gains beyond perceptual limitations and revisiting bending gains. In *2018 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 115–122, 2018. doi: 10.1109/ISMAR.2018.00041

[46] H. Sakono, K. Matsumoto, T. Narumi, and H. Kuzuoka. Redirected walking using continuous curvature manipulation. *IEEE Transactions on Visualization and Computer Graphics*, 27(11):4278–4288, 2021. doi: 10.1109/TVCG.2021.3106501

[47] S. P. Sargunam, K. R. Moghadam, M. Suhail, and E. D. Ragan. Guided head rotation and amplified head rotation: Evaluating semi-natural travel and viewing techniques in virtual reality. In *2017 IEEE Virtual Reality (VR)*, pp. 19–28. IEEE, 2017.

[48] P. Schmitz, J. Hildebrandt, A. C. Valdez, L. Kobbelt, and M. Ziefle. You spin my head right round: Threshold of limited immersion for rotation gains in redirected walking. *IEEE transactions on visualization and computer graphics*, 24(4):1623–1632, 2018.

[49] V. Schwind, P. Knierim, C. Tasci, P. Franczak, N. Haas, and N. Henze. " these are not my hands!" effect of gender on the perception of avatar hands in virtual reality. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pp. 1577–1582, 2017.

[50] 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, 2009.

[51] E. A. Suma, S. Clark, D. Krum, S. Finkelstein, M. Bolas, and Z. Warte. Leveraging change blindness for redirection in virtual environments. In *2011 IEEE Virtual Reality Conference*, pp. 159–166, 2011. doi: 10.1109/VR.2011.5759455

[52] E. A. Suma, Z. Lipps, S. Finkelstein, D. M. Krum, and M. Bolas. Impossible spaces: Maximizing natural walking in virtual environments with self-overlapping architecture. *IEEE Transactions on Visualization and Computer Graphics*, 18(4):555–564, 2012.

[53] G. Tieri, E. Tidoni, E. F. Pavone, and S. M. Aglioti. Mere observation of body discontinuity affects perceived ownership and vicarious agency over a virtual hand. *Experimental Brain Research*, 233(4):1247–1259, 2015. doi: 10.1007/s00221-015-4202-3

[54] M. Tsakiris and P. Haggard. The rubber hand illusion revisited: visuotactile integration and self-attribution. *Journal of experimental psychology: Human perception and performance*, 31(1):80, 2005.

[55] M. Usoh, K. Arthur, M. C. Whitton, R. Bastos, A. Steed, M. Slater, and F. P. Brooks. Walking > walking-in-place > flying, in virtual environments. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '99, p. 359–364. ACM Press/Addison-Wesley Publishing Co., USA, 1999. doi: 10.1145/311535.311589

[56] M. Van Someren, Y. F. Barnard, and J. Sandberg. The think aloud method: a practical approach to modelling cognitive. *London: AcademicPress*, 11:29–41, 1994.

[57] D. Vogel and R. Balakrishnan. Distant freehand pointing and clicking on very large, high resolution displays. In *Proceedings of the 18th annual ACM symposium on User interface software and technology*, pp. 33–42, 2005.

[58] A. B. Watson. Quest+: A general multidimensional bayesian adaptive psychometric method. *Journal of Vision*, 17(3):10–10, 2017.

[59] M. Weise, R. Zender, and U. Lucke. How Can I Grab That? *i-com*, 19(2):67–85, aug 2020. doi: 10.1515/icom-2020-0011

[60] C. You, B. Benda, E. S. Rosenberg, E. Ragan, B. Lok, and J. Thomas. Strafing gain: Redirecting users one diagonal step at a time. In *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 603–611, 2022. doi: 10.1109/ISMAR55827.2022.00077

[61] 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.

[62] A. Zenner, K. P. Regitz, and A. Kruger. Blink-Suppressed Hand Redirection. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*, pp. 75–84. IEEE, mar 2021. doi: 10.1109/VR50410.2021.00028