Evaluating Remapped Physical Reach for Hand Interactions with Passive Haptics in Virtual Reality

Dustin T. Han, Mohamed Suhail, and Eric D. Ragan

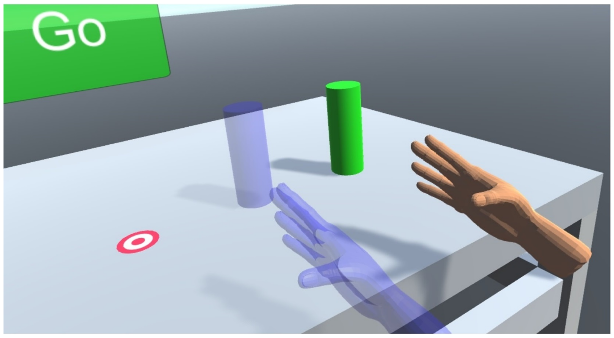

Fig. 1. Applications used in the research. *Right*: The immersive game designed in the case study uses remapped reaching techniques to interact with multiple virtual objects. *Left*: A representation of translational shift. The opaque hand and green cylinder show the virtual objects, and the partially transparent blue objects show where the hand and object are in the real world. This image demonstrates the simple environment used in Experiments 1 and 2, though the experiments did not include the blue representations of the real locations.

Abstract— Virtual reality often uses motion tracking to incorporate physical hand movements into interaction techniques for selection and manipulation of virtual objects. To increase realism and allow direct hand interaction, real-world physical objects can be aligned with virtual objects to provide tactile feedback and physical grasping. However, unless a physical space is custom configured to match a specific virtual reality experience, the ability to perfectly match the physical and virtual objects is limited. Our research addresses this challenge by studying methods that allow one physical object to be mapped to multiple virtual objects that can exist at different virtual locations in an egocentric reference frame. We study two such techniques: one that introduces a static translational offset between the virtual and physical hand before a reaching action, and one that dynamically interpolates the position of the virtual hand during a reaching motion. We conducted two experiments to assess how the two methods affect reaching effectiveness, comfort, and ability to adapt to the remapping techniques when reaching for objects with different types of mismatches between physical and virtual locations. We also present a case study to demonstrate how the hand remapping techniques could be used in an immersive game application to support realistic hand interaction while optimizing usability. Overall, the translational technique performed better than the interpolated reach technique and was more robust for situations with larger mismatches between virtual and physical objects.

Index Terms—Virtual reality, 3D interaction, passive haptics, hand interaction, remapped reach, 3D object selection

1 Introduction

Virtual reality (VR) systems often combine interactive 3D graphics with motion tracking input to support realistic or natural interaction with virtual environments [4]. In many applications, tracking the user's hand or using a tracked controller enables the use of physical hand movements for selection and manipulation of virtual objects. Researchers and developers have created many interaction techniques that rely on such hand controls, and common implementations include a virtual hand with button activation to indicate selection or ray casting that allows selection of distant objects [3]. For higher tactile fidelity, a user would ideally be able to physically interact with virtual objects through direct reaching and hand grasping. One way to achieve this is through *passive haptics*, which is the use of real-world physical objects that correspond to virtual objects [9]. Passive-haptic props can be used to allow direct hand interaction with real-world physical objects to provide tactile feedback and physical grasping. However, unless a physical space is custom configured to match a specific VR experience, the ability to perfectly match the physical and virtual objects is limited.

To address this challenge, our research studies methods that allow one physical object to be mapped to multiple virtual objects that can exist as different virtual locations in the same egocentric reference frame. Use of a single physical prop for multiple virtual objects can provide great flexibility for using a VR to simulate a variety of scenarios with different configurations of objects, but this approach is not without its problems for hand interaction. One obvious issue is that virtual objects can be represented in any form or size, which makes it difficult for a single prop to be physically accurate for all cases, but studies have found that a perfect physical representation of a virtual object is not necessary to increase the sense of presence and realism [9, 13].

Another major challenge is dealing with varying locations of virtual objects while having limited physical props. A discrepancy of positions between physical and virtual objects can cause users to miss or inaccurately touch a physical prop when reaching for the corresponding virtual representation. Although research has been done using robotics to enable a physical prop to move around depending on specified virtual locations [8], this method requires expensive equipment and is impractical for many VR users with common setups. Rather than dynamically modifying the physical world, our research focuses on remapping the virtual space. By adjusting the location of the virtual hand during reaching or touching actions, it can be possible to redirect a user's physical

• *Dustin T. Han is with Texas A&M University. E-mail: dthan@tamu.edu.*
• *Mohamed Suhail is with Texas A&M University. E-mail: mohamedsuhail@tamu.edu.*
• *Eric D. Ragan is with Texas A&M University. E-mail: eragan@tamu.edu.*

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxxx/TVCG.201x.xxxxxxx

reach towards an intended target [1, 11]. It has been shown that, to a certain degree, the visual sense can dominate the proprioceptive senses in scenarios with mismatching physical and virtual objects, which can allow for some virtual body warping to occur unnoticed [5].

In order to effectively take advantage of remapping techniques in VR applications, it is important to understand how remapping can be implemented in different ways and in cases with different object locations. In this research, we study two remapping techniques for reaching: *translational shift*, which introduces a static offset between the virtual and physical hand before a hand reach, and *interpolated reach*, which dynamically interpolates the position of the virtual hand during a reaching motion. Although virtual remapping can technically be applied to all cases of physical-virtual object mismatch, there will be cases where the offset is so large that the techniques are no longer practical for their initial intended use of enabling more realistic interaction. A separate issue to consider is the physical location of the haptic prop in relation to the user, as certain combinations of locations and virtual-physical offsets are likely more problematic than others. Thus, we aim to address the following research questions:

- How do translational shifting and interpolation techniques influence performance in reaching tasks involving hand interaction with physical objects, and how do the techniques compare in terms of task efficiency, comfort, and overall usability?
- To what extent do different distances from the user to the prop and distances between physical and virtual objects influence reaching?
- How can remapped reach techniques be configured for practical use in VR applications?

To address these questions, this paper presents two controlled experiments and a case study to evaluate remapped reach techniques. The first experiment evaluated how different configurations of the techniques influence reaching performance (speed and errors), and the second experiment studied how people adjust to remapping techniques with continued use of the same technique. Finally, to learn more about how remapped reach might be used to support realistic hand interactions for VR applications, we conducted a case study in designing a VR game that uses the remapped reaching techniques for object interaction.

2 Related Work

This research leverages previous research of passive haptics and interaction techniques that make use of remapping physical motion input to different virtual positions.

2.1 Haptics in Virtual Reality

Haptics have a long history of use in many industries including aerospace, medical, and entertainment [22]. For VR, haptics are used to provide additional feedback to users, as many systems focus on visual and auditory feedback. There are two general categories of haptics: *passive haptics* and *active haptics*. Passive haptics is a category of haptics where objects or devices provide tactile feedback naturally through a user's sense of touch of their physical properties. In contrast, *active haptics* generates forces through controlled actuators.

Haptics is an area where robotics plays a role, as movable physical props is possible. McNeely discussed the concept of robotic graphics using robots to enable force feedback for VR [14]. McNeely presented the idea of robotic shape displays (RSD), which utilize robots that are specifically used to provide tactile feedback by repositioning and reorienting themselves based on anticipated movements of the user. He et al. [8] also used robots to provide haptic feedback in research that used haptic proxies to support collaborative tasks in VR. In their work, a robot moves around on a table in the physical world to accommodate changes done by another user who is collaborating on the task.

Although robots and other active haptic devices have the benefit of dynamically adapting to match changes in a virtual environment, these methods are costly and not practical for most common VR setups such as home use of commercially available systems. Passive haptics have the advantage of being easily accessible, as any convenient household items (e.g., bottles, cups, or boxes) could easily be used as passive-haptic props. Insko [9] provided evidence that passive haptics can enhance a user's virtual experience in various ways. Passive haptics can provide users with more control and allow for better correlation between virtual and real world tasks, such as handwriting in VR as demonstrated by Poupyrev et al. [18]. Visual feedback can also be used with passive haptics to help users perceive different forces and tactile sensations that do not match the physical forces or shape of a prop [13].

2.2 Visual and Proprioceptive Mismatch

Evidence has shown that visual feedback can dominate proprioceptive senses in cases where there is a mismatch between the virtual and physical world [5]. Many researchers have taken advantage of this visual dominance to provide virtual experiences that would not be possible within the user's physical environment. This has commonly been done for travel techniques, such as with redirection methods that subtly adjust the view while the user is physically walking (e.g., [19])) or virtually moving (e.g., [20]). For example, Razzaque et al. [19] studied the use of slight rotational adjustments in the head to enable redirected walking and create the illusion that the user is moving through a larger space. Research has also explored methods that change the mapping between physical and virtual environments while the user is not looking [24, 25].

For our study of reaching, we focus on remapping physical hand motions to different virtual motions. Closely related to the interpolation technique in our research, Poupyrev et al. [16] proposed the Go-Go technique, which allows users to interact with distant objects through amplified hand movements. A comparison of Go-Go to a ray-casting method found Go-Go to be more accurate for object selection and comparable for object positioning [17]. In a study of remapped selection cursors, Paljic et al. [15] studied the addition of a virtual depth offset to a 3D cursor controlled by hand pointing. The results showed found minimal effects for 20 cm offsets, but larger offsets of 40 cm significantly slowed selection performance. Our work uses direct remappings for reaching rather than using a cursor, and we study offsets in multiple directions.

Other work studied the effects of mismatch for physical reaching tasks. Ebrahim et al. [7] found that although tactile feedback can provide more accurate depth judgment in physical reaching tasks, performance can be significantly affected when visual and proprioceptive information do not match. While the scope of their study focused on a single reaching task, we consider task performance and effects of different types of visual-proprioceptive mismatch.

Understanding how people can adapt to mismatches overtime is also important. In addition to finding the effects of the presence of an arm or full body avatar in task performance, Bodenheimer et al. [2] found that people adapted overtime to a lateral displacement of the visual environment for a throwing task. Similarly relevant, Kohli et al. [12] found that although there are indications of reduced performance in warped virtual spaces, after adaptation, users can perform tasks when their visual senses do not fully match their proprioceptive sense with performance similar to a one-to-one virtual environment. Our research also studies the effects of remapping techniques on task performance over time for reaching, and in contrast to previous studies, we study the effects over a large set of motions and offsets with two different techniques.

2.3 Redirecting Reach

This idea of virtual remapping can be used on a virtual hand or limb to redirect reaching or touching movements towards an intended physical object. Kohli et al. [10] explored the concept of *redirected touching*, where virtual space was warped to provide haptic feedback for multiple virtual objects with a single physical object. Kohli's research went on to study the effects of virtual warping on task performance in a target-touching task with a physical board in front of participants [11].

Exploring interaction with multiple objects, Suhail et al. [23] demonstrated the use of offsets between physical and virtual objects for interaction with a physical passive-haptic prop in a VR game. This work demonstrated the use of a single prop for different types of object interaction (e.g., opening a door, pulling a switch, moving an object) with multiple interaction points at different locations throughout the virtual environment, and they explored the use of travel adjustments to help align users with the physical objects for easier reaching. In other work, Cheng et al. [6] discussed applications of hand redirection through the use of *sparse haptic proxies*, which are geometric primitives that provide touch feedback for various objects in a virtual environment. The study explored the use of the haptic props in two vastly different environments to show how the approach can be used to provide tactile feedback for a variety of scenarios.

Also highly relevant to our research, Azmandian et al. [1] studied the use of haptic remapping where different techniques are used to remap a hand to reach and grab a physical object. Their research considers the use of *body warping*, *world warping*, and a hybrid of both. *Body warping* manipulated the virtual arm and hand offset as it reaches towards a virtual object, whereas *world warping* remapped a virtual object by rotating the virtual world around the user's head, with this adjustment applied as the user looks away from the object. The researchers found evidence that the remapped hand techniques generally provided benefits for perceived realism and sense of presence. While this work addressed usability and feedback of their remapping techniques in general, our research focuses more on interaction performance and provides a more detailed "stress test" of how different situations and deviations between the virtual and physical world can influence performance. That is, our research explicitly tests different configurations of offset sizes, offset directions, and object locations, and we also study habituation to remapped techniques over time. Relatively little is known about how different remapping techniques influence performance and usability for different types of reaching actions with different technique configurations.

3 TECHNIQUES

Our research studies two techniques for redirecting reach by remapping the position of the virtual hand based on the position of the physical hand. This section describes the VR system, physical setup, and techniques used for two experiments and a case study.

3.1 Equipment and Study Setup

The techniques were tested with a task that involved reaching for an object (a plastic bottle) on a table. All participants sat in a stationary (non-rotating) chair at a 72 x 30 inch table. The studies used an Oculus Rift CV1 head-mounted display (HMD). An Oculus Rift remote was used by the experimenter to trigger tasks and log any errors that occurred such as missing or bumping the bottle. The study software applications were developed in the Unity engine and ran on a computer running 64-bit Windows 10, a 3.4 GHz Quad Core processor, and a GeForce GTX 1070 graphics processing unit.

Six degree-of-freedom (DOF) motion tracking was handled using an OptiTrack capture system with 12 Flex 13 cameras. Based on manufacturer reports, each camera recorded tracking data at a frame rate of 120 frames per second and operated with 8.33 ms latency. Six DOF head-tracked viewing was enabled for the HMD, and hand tracking was enabled by attaching a rigid body to the participant's right wrist via a velcro strap. A rigid body attached to an elastic band was also worn on the tip of the middle and ring fingers, which was used to control the grasping action and animation of the virtual hand. With this configuration, the thumb was not tracked, and the other four fingers moved together to open or close during grasping motions.

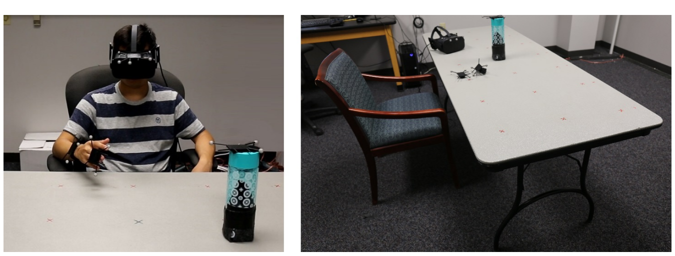

Fig. 2. The physical setup for the experiments includes a table, a stationary chair, a tracked bottle, and hand markers, and an HMD.

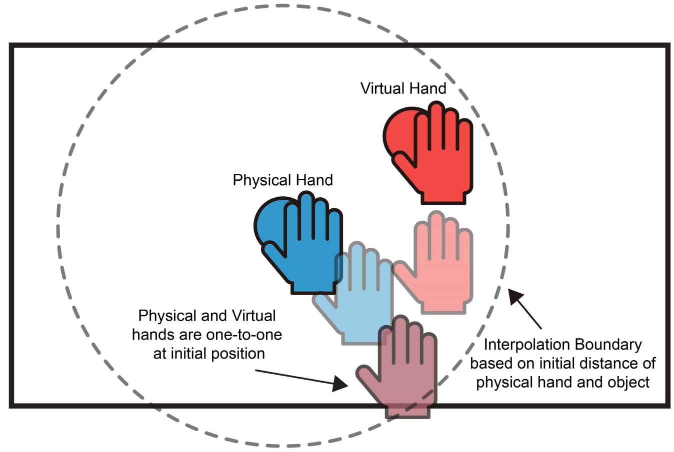

Fig. 3. A top-down diagram demonstrating interpolated reach shows the physical hand and object in blue and the virtual hand and object in red. The dotted circle has a radius equal to the initial distance of the physical hand and object, and it denotes the region where interpolation is in effect. Outside this boundary, a one-to-one mapping is used.

The studies also used a cylindrical plastic bottle as a passive haptic prop, which was tracked by a rigid body on top. Soft fabric was added to the bottom of the physical bottle to reduce the noise when the experiment placed it on the table. This noise reduction was especially important for this experiment in order to eliminate any sound cues as the experimenter was required to move the bottle to different positions on the table during the experiments. All participants used their right hands for interactions in all experiments, and they did not use any type of controller. Figure 2 shows the physical setup.

3.2 Translational Shift

Translational shift is a remapping technique that involves relocating the virtual hand based on the positional offset between the real world object and the virtual object. Figure 1, left, demonstrates the basic concept using semi-transparent blue objects to show where the positions of tracked objects would appear with a one-to-one mapping. The calculation for the virtual hand position, P_{vh}, using translation offset is given by:

$$P_{vh} = P_{ph} + (P_{po} - P_{vo})$$

Here, P_{ph} is the physical hand's current world position and P_{po} and P_{vo} are the world positions of the physical object and virtual object respectively. Translational offset can be applied in any axis depending on the offset. In the case of our studies, we focused on offsets only in the horizontal plane.

3.3 Interpolated Reach

Interpolated reach is a technique that aims to initially keep the virtual hand at a starting position that matches the physical real world hand so the virtual hand appears to be where it would be in relation to the user's body. As the physical hand moves, an offset is gradually applied to the virtual hand as it moves toward the virtual object until it reaches the maximum offset at the time the physical hand reaches the physical object. It is expected that the user will gradually adjust his hand position during reaching so the virtual hand will reach the virtual target, and this will guide the physical hand to the physical prop at the same time. The technique uses the calculated offset between the virtual and physical object to dynamically blend between the physical real world hand position to its offset position. Because of this dynamic changing of hand position, the virtual hand experiences manipulations in its speed and trajectory during reaching tasks to compensate for the offset. Figure 3 demonstrates the concept.

The virtual hand position, P_{vh}, for interpolation is given by the following equation:

$$P_{vh} = \begin{cases} P_{ph}, & D \geq B \\ P_{ph} + (P_{po} - P_{vo}) * (1 - \frac{D}{B}), & D < B \end{cases}$$

Here, P_{ph} is again the physical hand's current world position, and P_{po} and P_{vo} are the world positions of the physical object and virtual object, respectively. We have an additional term, $(1 - \frac{D}{B})$, which is the offset control value that determines how much of the offset is applied to the virtual hand. Here D is the distance between the physical object and physical hand. As the hand moves closer to the real world object, D gets closer to 0. The term B is the interpolation boundary, which is a constant value that determines the region where interpolation is in effect, as given by: $B = |P_{po} - P_{ph}| + C$.

In this implementation, the boundary offset is given by the initial distance between the physical hand and physical object plus a small buffer value C. The added buffer is needed to ensure the hand stays within the interpolation region regardless of initial hand movement. Our implementation used $C = 0.1$ meters. Without the buffer, the hand starts at the very edge of the boundary region, and there is potential for the user to reach in a direction that moves away from area of effect causing them to miss the object completely as the offset is not being applied. When $D \geq B$, the hand is outside the region of effect, so the virtual hand reverts to a one-to-one mapping with the physical hand.

4 Experiment 1: Performance and Usability

Although past research has shown that remapping reach techniques can be applied to reuse physical objects for providing tactile feedback of multiple virtual objects, there is still the question of how different techniques compare to each other, and how these techniques affect a user's ability to perform tasks when compared to one-to-one mapping of the virtual and physical hand. Experiment 1 compares two remapping techniques. It is important to understand the limitations of these techniques in order to use them effectively and comfortably within applications. The goal of this experiment was to test various conditions of the two remapping techniques to determine what combinations and ranges of offsets and table positions for the physical prop are best in terms of task performance and comfort.

4.1 Experimental Task and Environment

The study used the setup described in Section 3.1. The table was marked to indicate different physical locations for the physical prop during the study. The table locations were placed based on 10-inch intervals starting from the center of the table and extending both to the left and right as well as towards and away from the user's position. The chair and table were positioned the same way for all participants using markers on the floor for alignment.

To make sure participants were focused on the task, Experiment 1 used a simple virtual environment that consisted of an open environment with only a white rectangular table placed in front of them that matched the size and shape of the real world table. The environment also had a cylindrical virtual object used for the reaching task (see Figure 1, left).

Directly in front of the participant and beyond the virtual table, a message panel was shown to indicate when to begin each trial. The task for each trial was to use the virtual hand to reach and grab the virtual object when the object color changed from red to green and the message panel showed a "GO!" message. Directly after grabbing the object, the participant placed the object on a red target that appeared on the virtual table.

4.2 Experimental Design

When mapping virtual to physical objects for real applications, it may be desirable to use remapping for a number of different configurations that depend on the position of the user's physical hand, physical props in the real world, and virtual objects in VR. We hypothesized that various configurations would change the overall difficulty and usability of reaching. Experiment 1 focused on evaluating reaching performance (time and accuracy). We hypothesized that between the translational technique and the interpolation technique, the translational technique would perform better overall because it maintains one-to-one hand motion in contrast to the changing of speed and trajectory of the hand in interpolated mapping. Logically, we believed larger offsets and physical prop positions that were farther away from the user would cause decrease in task performance. In this experiment, we compared the following conditions:

- **Technique:** As previously described, the study tested two remapping techniques: *translational offset* and *interpolation*. The study also included tasks using a normal (non-shifted, non-interpolated) one-to-one mapping as a basis for comparison.

- **Offset Size:** Offset size is the distance measured between the physical and virtual object. We studied a total of 4 variations of offset size: 10, 14, 20, and 30 inches. The 14-inch offsets were exclusively used for diagonal offsets, as each marker on the table was placed 10 inches apart from any adjacent marker, making the diagonal distance approximately 14.14 inches.

- **Offset Direction:** Along with the 2D object distribution of table locations, it was possible to have different 2D directions for the offset vector between the real and virtual objects. We find it most helpful to classify two cases of offset direction, which we call *toward* and *away*. An offset is considered *toward* a user if the virtual object appears closer to the user than the actual physical object (in other words, the virtual object is between the physical object and the user's physical hand). An offset is considered *away* when the virtual object appears farther than the physical object from the user's point of view (in other words, the physical object is between the virtual object and the user's physical hand).

- **Table Position:** We use the term *table position* to refer to the location on the table where the physical object is. The experiment used 17 possible locations on the table. Figure 6 shows the approximate configuration of locations by square cells. Each location was separated from others by 10 inches to the left and right, and by 10 inches in the direction away from the user. During the experiment, the experimenter adjusted object placement at table positions by manually moving the object.

The experiment was conducted with a repeated-measures design. In other words, each participant was exposed to all possible task variations during the experiment. The experimental procedure was broken up into three parts, each using a different hand mapping technique: normal *one-to-one* reach, *translational shift*, and *interpolated reach*. All participants started with the one-to-one version for the first block of trials. The *one-to-one* block consisted of 17 tasks which covered each possible marked location on the table. For the *translational shift* and *interpolated reach* blocks, participants were asked to perform a set of different tasks based on unique variations of the independent variables. Because of the numerous possible combinations of the independent variables, a subset of 42 tasks for each technique was chosen for the experiment. Each of the 42 tasks were unique with a different combination of offset size, offset direction, and object position. The tasks were derived to cover all possible table locations with offset directions alternating between every other location. To ensure the virtual object always appeared as it was resting on top of the table, the virtual object was never offset beyond the virtual table borders.

Although each participant was given the one-to-one session to begin with, the other two sessions with the remapping techniques were counter balanced so that half of the participants started with *translational shift* and the other half started with *interpolated* technique. Task ordering within each technique block (including the one-to-one session) was randomized for each participant.

To measure the effects of the independent variables on reach performance, we considered time and errors as dependent variables. Reach time was measured from when the task started, signified by the "Go" sign and the virtual bottle changing from red to green, to the end of the task, when the participants successfully grasped the object and all five fingers made contact with the physical object. Times were recorded by the experimenter where both the start and end of the reach task were manually marked through clicking a button. Because people might normally reach at different speeds, and each table positions will have varying times due to different distances, we analyzed reach time for each condition by taking the difference between the measured reach time with the remapping technique and the participant's one-to-one reach time for the corresponding table position measured in the first block of trials.

To measure reach errors, we considered two types of errors: bumping and missing the physical object. A bump error was recorded if the hand made contact with the bottle, including knocking it over, without a successful grasp. A miss error was recorded if the hand reached past the physical object without a successful grasp. Errors were manually noted by the experimenter observing the participant. The error metric uses the total number of all errors (of either type) observed per trial.

4.3 Procedure

At the beginning of the experiment, participants were seated at a table and given an overview of the study and equipment used including the tracking system and Oculus Rift. They signed an informed consent form before proceeding with the experiment. Each participant filled out a brief background questionnaire that included questions regarding demographic information (age, gender, and occupation) and their overall experience with 3D games and VR.

The participants were then assisted in putting on the HMD and markers to track the hand. Before beginning the main trials, a practice session was provided to get the participant accustomed to the task and virtual environment before the main experiment. The practice session included five one-to-one reaching tasks, so participants were not introduced to the remapping techniques beforehand. After the practice session, participants took a short break before beginning the main trials. Participants were not told about the techniques or how reaching would be modified throughout the study, but the experiment was not designed to be subtle or unnoticeable with the application of the techniques, as the extent of remapping was often intentionally large.

For each trial, the virtual hand and object were hidden from view before reaching. When the task began, the virtual hand, object, and target fade into view and the word 'READY' was displayed on a gray instructions panel. At this time ,the virtual object was red. During this time, the participant was asked to place their hands on the arm rest, to not move, and simply acknowledge the location of the virtual objects. After approximately 2 seconds, the experimenter initiated the task, which changes the gray 'READY' sign to a green 'GO' sign and changes the virtual object from red to green. This signals the participant to reach the object and place it onto the red target. Once the object is placed on the target, the hand, target, and virtual object fade away until the next task is started. Participants were asked to perform the tasks as quickly as they could in a way that was comfortable and natural to them while avoiding error. The main reaching trials were divided into three sessions (as explained in the previous section) with mandatory five-minute breaks in between sessions.

After the three sessions, a semi-structured interview was conducted by the experimenter to collect additional feedback from the participant. The interview consisted of questions regarding how they felt about the different versions and tasks they experienced and what differences they noticed. Experiment 1 took approximately 1 hour.

4.4 Participants

Sixteen university students (8 male, 8 female) took part in Experiment 1. Participants' ages ranged from 20 to 29 with a median of 25 years. Participants were students in degree programs related to computing, engineering, and art. Nine participants reported spending at least one hour a week playing 3D video games, and 11 reported having prior experience with VR.

4.5 Results and Discussion

Experiment 1 tested two different reach techniques with different offset sizes. The experiment also considers many different table positions and offset directions, which we summarize based on *toward* and *away* offset directions and by *close* and *far* table positions.

4.6 Performance Results

Considering the effects of all combinations of possible factors together would be difficult to comprehend in any meaningful way. Therefore, we present multiple analyses focusing on different relevant subsets of factors. For our analysis and presentation of results, we use repeated-measures ANOVA tests and graphical plots to represent differences due to experimental variables. For brevity, we only report test statistics for significant effects. We present results using box-and-whisker plots where colored rectangles represents the interquartile range (IQR) and a horizontal black band marks the median value. Black "whisker" lines extend from the rectangle to the most extreme value that falls an additional half-IQR beyond the IQR, or no whisker is shown if no points fall in this range. Black dots denote values beyond this range.

For Experiment 1, the time results met the assumptions of normality and sphericity for parametric testing. However, error results did not meet the assumption of normality due to many trials having low errors. We were unable to correct the distributions with data transformations or to identify appropriate non-parametric alternatives for our sample sizes. For this reason, and considering the robust nature of ANOVA tests for normality violations (e.g., [21]), we chose to rely on parametric repeated-measures ANOVAs with the acknowledgement of increased risk of type 1 errors (e.g., [26]).

For statistical testing, we first consider performance for the two reach techniques, different offset sizes, and offset direction (the virtual object offset being *away* from user or *toward* the user). We ran three-way repeated-measures ANOVAs for reach time (the difference between the condition and average times for the position with baseline one-to-one reaching) and reach errors. The distribution of times are shown in Figure 4. Overall, the *interpolated reach* technique was significantly slower overall than the *translational shift* technique with a significant main effect yielding $F(1,15) = 4.83$ and $p = 0.048$, but this is better explained by interaction effects (to be explained in the following paragraph). The effect of offset direction was also significant $F(1,15) = 148.49$ and $p < 0.001$, showing that reach speeds were significantly faster when the virtual object was offset away from the body than when offset toward the body. Additionally, offset direction had a significant main effect with $F(3,45) = 27.83$ and $p < 0.001$, with posthoc Bonferonni-corrected pairwise comparisons showing offsets of size 30 were significantly slower than all other offset sizes. It is interesting to note the failure to detect significant performance differences with smaller offset sizes, which suggests that people may generally be able to tolerate and compensate for substantial mismatches between virtual and physical objects without major problems. Further experimentation with finer-grain offset sizes would be needed to more precisely assess any tolerance threshold.

Further evidence of the relationship among the experimental factors is given by a significant three-way interaction among technique, offset size, and offset direction, with $F(3,45) = 5.48$ and $p = 0.002$, and by a significant two-way interaction between offset size and direction, with $F(3,45) = 10.14$ and $p < 0.001$. From Figure 4, we can see that slower reach times were observed for *towards* reaches compared to *away* reaches, and the problems with the largest offsets were significantly worse with the *away* movements. Most interestingly, the *translational* technique resulted in relatively good speeds for *away*, and this combination seemed to be most tolerant to large (30 in) offset sizes. On the other hand, the worst combination for speed was the *interpolated reach* in the *toward* direction and with 30-in offsets.

We similarly tested for effects on reach errors due to technique, offset size, and offset direction with a three-way repeated-measures ANOVAs. The test found significant main effects for both technique and offset size, but these effects are better explained by a significant interaction between the two factors with $F(3,45) = 4.62$ and $p = 0.007$. The interaction shows that greater numbers of errors were made in

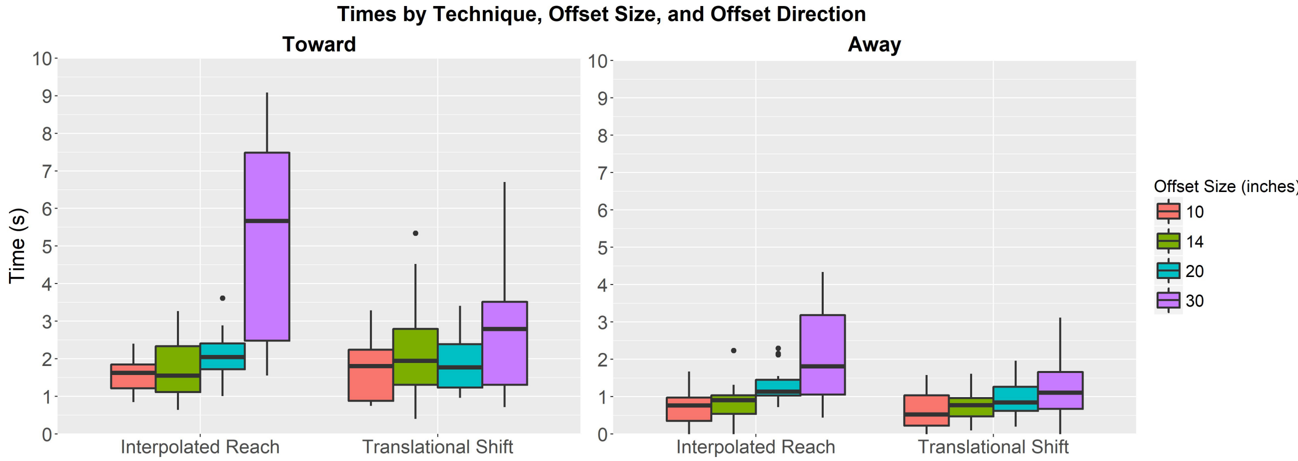

Fig. 4. Reach times from Experiment 1 separated by technique type, offset direction, and offset size. Speeds are similar for small to moderate offset sizes. The translational techniques are significantly more robust to large offset sizes, and reaching away from the body is significantly faster than reaching towards the body.

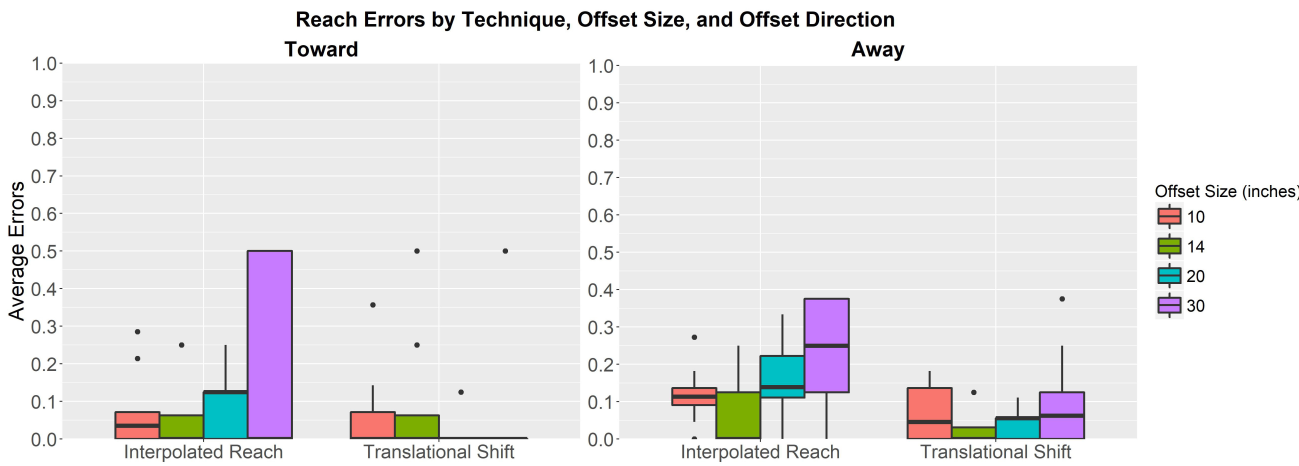

Fig. 5. Errors for the reach task in Experiment 1 were significantly worse with the interpolation technique when used with high offset sizes.

conditions using *interpolated reach* with higher offset sizes. In contrast, with the *translational* technique, errors were much more similar among offset sizes (see Figure 5).

Considering the time and error results together provides consistent evidence that the *translational shift* technique is more robust for better performance with larger mismatches between virtual and physical targets. Also, the *interpolated* technique did especially poorly for 30-inch offsets in terms of both time and errors. This suggests the existence of a limit to the offset size before seeing a dramatic penalty to performance. Though reaching away from rather than towards the body allowed significantly faster reaching, the ANOVA for errors failed to detect a significant effect of reach direction.

We also considered effects on performance based on table position. Unsurprisingly, reaching times were longer and more errors were observed for reaches when the prop was positioned far from the default position of the user's right side. A distribution of performance results by table position is summarized in Figure 6. This distribution was observed independent of other factors.

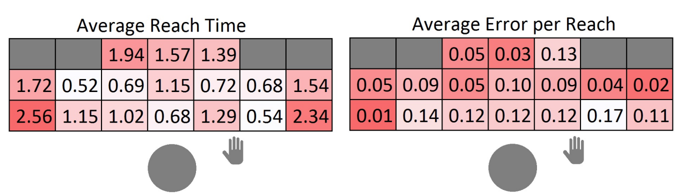

Fig. 6. Average reach times and errors from Experiment 1 based on table position, where the shape of the data table corresponds to the shape of the physical table, with the user's location represented by the gray circle. Gray cells represent areas of the table not used in the study.

4.6.1 Qualitative Results

In addition to the performance results, we gathered qualitative feedback from participants through a semi-structured interview at the end of the study. In terms of overall preference of techniques, 9 participants preferred the *translational shift* technique, 5 preferred *interpolated reach*, and 2 felt they were about equal in terms of difficulty of use.

In general, participants did not like tasks with *toward* offset, where the virtual object seemed visually closer than the physical object actually was. This was especially true for *interpolated reach* where the hand would move slower due to a small virtual distance and large physical distance; four participants stated this slow motion control made them feel disoriented or nauseated. Most did not mind the *interpolated reach* with fast interpolation (where the virtual distance exceeded the physical distance). One participant noted that in a subset of tasks with 10 inch offsets in the *interpolated reach* session, it felt like they were doing one-to-one motion. However, interpolation was more disorienting to a majority of participants, and there was a more obvious disconnect with the physical world as the hand was not moving the way they expected.

The main criticism for *translational shift* was for offsets that took the virtual hand out of the peripheral view of the user, which caused confusion and required the participant to initially look around for the hand. Several participants commented on being less connected with the hand in *translational shift* compared to *interpolated reach* because the hand was constantly being placed in different locations relative to the body. Some related it to controlling a "puppet" or "ghost" hand and did not feel like it was their own. For both techniques, many found they had to rely more on visuals to guide their hand to the object as opposed to their usual intuition and proprioceptive senses; the impression was that the techniques were usable but needed practice over time. The randomness of the tasks and offsets were also noted by some to cause difficulty in adjusting to the techniques.

Participants felt the haptic prop was beneficial overall, explaining that it made the experience more immersive and gave them a real-world physical cue of when they have the object in their grasp. Some participants noted that being seated at a phsyical table that was similar in shape and size to the virtual table made them think about the real world more and look for relationships in the virtual world. There was also a fear of hitting or knocking over the bottle during the experiment. One participant suggested that they would probably perform better if the table was larger and the object was smaller to reduce the chances of error. A comment was also made that stretching and reaching across the table was uncomfortable. Some minor visual discrepancies, such as the thumb not being tracked or the hand clipping through parts of the virtual object, brought some participants out of the experience.

5 Experiment 2: Evaluating Acclimation

Experiment 1 required participants to complete a set of consecutive reaching tasks where each task used a different random configuration of each technique. Each task had a different combination of offset direction, offset size, and physical object table position. The main purpose of Experiment 1 was to understand how these different independent variables can affect a reaching performance. However, due to the randomness of condition ordering, participants could never know the extent of the remapping before they started each reaching task. Consequently, it was likely difficult for participants to ever adapt to any of the technique configurations. We therefore conducted a second experiment to help understand how well users might be able to adapt to continued use of the remapping techniques. Experiment 2 aimed to study how task performance changed over time with the different remapping techniques when given tasks with the same offset effect.

5.1 Experimental Task and Environment

Experiment 2 used the same virtual environment as Experiment 1. However, for Experiment 2, the virtual table length and width were increased by 1.5 times the original to allow for virtual offset positions that would otherwise appear off the table. In this experiment, participants were again asked to reach and grab the virtual object when it changed from red to green and the sign in front of them signaled for them to "Go". However, only the reaching portion of the task was required for this experiment with no target placement after. Once the object was reached, the virtual object and hand would fade away, indicating the successful completion of that task.

5.2 Experimental Design

The experiment was designed to first have participants complete consecutive reaching tasks with the same remapping configuration for different object positions. Then, after multiple reaches with a remapped technique, the remapping was disabled, and participants immediately completed additional reaching tasks with normal one-to-one movement. We hypothesized that performance would improve over time with continued use of the same technique. We also hypothesized that after continued reaching with a remapped technique, it would take time to adjust back to normal reaching with one-to-one hand movements.

The experiment followed a repeated-measures design with four technique configurations: (1) translational shift with a 20-inch offset to the user's right, (2) translational shift with 20-inch offset in the forward direction (away from the user), (3) fast interpolated reach, and (4) slow interpolated reach. Note that the *interpolated reach* technique manipulates the speed and trajectory of the hand depending on the offset between the virtual and physical objects. In order to constrain the behavior of the hand for interpolated reach to study acclimation, the virtual object was offset from the physical object position based on the vector from the participants hand to the physical object. From this vector, the object was either offset 10 inches in the direction of the vector for *fast interpolated reach* or offset 10 inches in the opposite direction (toward the participant) for *slow interpolated reach*. Technique order was balanced using a Latin square, and task order within each technique configuration was randomized for each participant.

As in Experiment 1, the reaching task in Experiment 2 involved different bottle positions on the table. The number of possible locations was reduced from 17 in Experiment 1 to 10 in Experiment 2 because the focus was studying adaption to techniques over time rather than to evaluate reaching difficulty and performance of each technique. All participant again started with a session with a normal one-to-one mapping as a baseline for comparing reach times. The one-to-one session consisted of 20 tasks where a task for each of the 10 locations was performed twice. The average of the two one-to-one tasks was used as the base time for that table location.

5.3 Procedure

Participants sat at a table with a stationary chair and first completed a consent form. Participants filled out a background questionnaire just as in Experiment 1. Next, the experimenter explained the reaching task, and participants completed a practice session where they had to reach and grab the virtual object with a matched one-to-one mapping five times at varying locations. After the practice session, participants were given a short 5-minute break before starting the main trials.

The main reaching tasks for the study were broken into five sessions. In the first session, all participants performed 20 reaching tasks with one-to-one mapping. In each of the the following four sessions, participants performed 30 consecutive reach tasks. Of these 30 reaches, the first 20 reaches included the offset for one of the four possible technique configurations. Immediately after performing the 20 tasks with the remapping technique, the technique was switched to regular one-to-one mapping for 10 additional reach tasks that covered each of the possible ten table locations. Participants were not notified of the technique change during the sequence of tasks.

In between all sessions, participants were required to take a 5-minute break from the virtual environment. During these breaks, participants were asked to reach for the physical bottle and move it around the desk a few times to help them recalibrate their senses to the real world before moving on to the next session with another remapped technique After finishing all sessions with the different techniques, the experiment was complete. A participant's completion of Experiment 2 took approximately 45 minutes.

5.4 Participants

Twelve university students (5 male, 7 female) took part in Experiment 2. The participants' ages ranged from 24 to 29 with a median of 26 years. Participants were graduate students in degree programs related to computing, engineering, and art. Eight participants reported spending at least one hour a week playing 3D video games, and all participants reported having prior experience with VR.

5.5 Results and Discussion

Because the focus of Experiment 2 is the time needed to adjust to technique changes, the results are the times taken to reach the target object over the sequence of trials. As in Experiment 1, the reported times are the difference between trial times and the corresponding reach times for each table position from the initial task block with normal one-to-one reaching. Here, a time difference of 0 indicates the reach time taken in trials with the remapped technique was the same as one-to-one reaching, and a negative value indicates a faster reach with the remapped technique than in the initial one-to-one reaches.

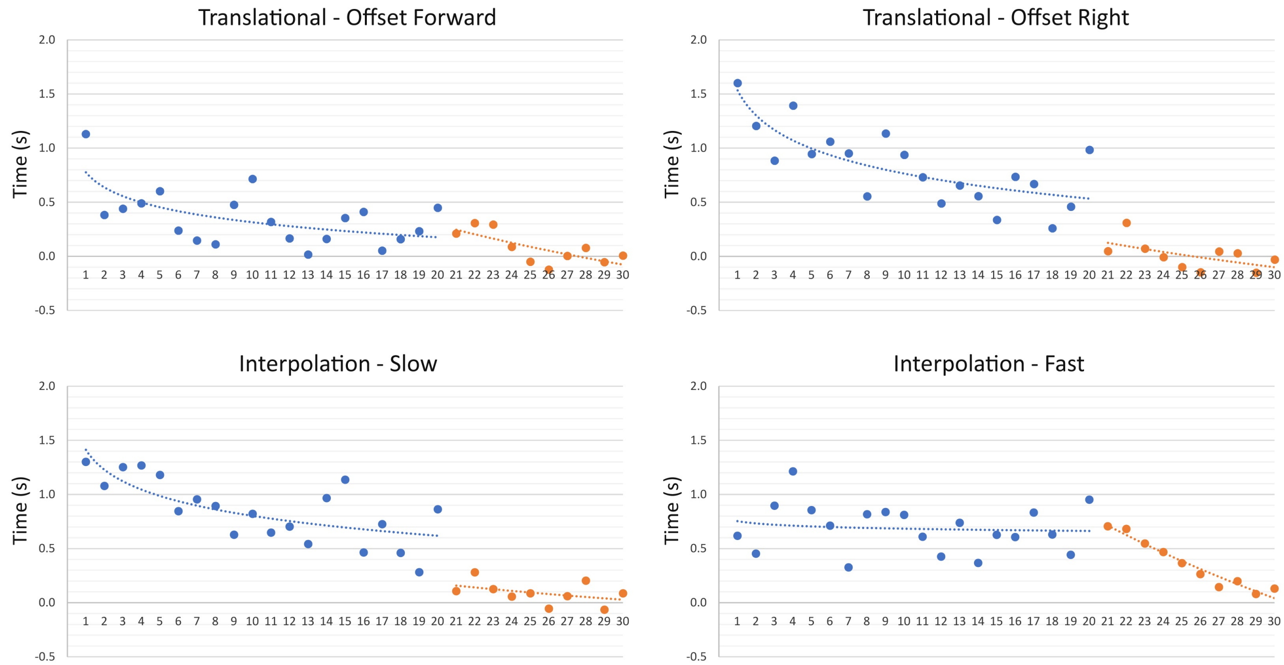

Fig. 7. Reach time results from Experiment 2. Plots show the difference between times for the remapping techniques and times for corresponding baseline one-to-one reaching tasks. Each plot shows a set of 30 reach tasks for a different technique in Experiment 2. The horizontal axis shows trial number. Blue points represent tasks done using the specified remapping configuration from tasks 1 through 20, and orange points represent tasks with one-to-one mapping from tasks 21 through 30. Trend lines are shown with logarithmic fit.

Figure 7 shows time differences as trials progressed, where further to the right of each plot indicates later trial numbers. Note that regardless of trial number, it is expected that times fluctuate due to some reaching tasks being more difficult than others. However, the overall results show that reach times did decrease with continued use of the remapped techniques over 20 reaches. These trends provide clear evidence that participants were adjusting to most techniques over time. Although, reach times did not often reach the baseline zero level within 20 reaches. It is possible that they would have with additional practice with the remapping techniques, but we cannot extrapolate based on these results.

An exception to the technique-acclimation effect was observed for the *fast interpolated reach*, which shows little change over the 20 reaches (see Figure 7, bottom right). From the results of Experiment 1, we found that users were generally comfortable with fast interpolation, and offsets away from the user saw better performance than towards the user, which may explain the lack of a noticeable improvement over time within 20 tasks. In contrast, with a remapping like for *right translational shift*, which takes the virtual hand out of the user's field of view, and *slow interpolated reach*, which may initially catch users by surprise, we see a noticeable improvement in speed from the first few tasks to the last.

When switching back to one-to-one movement for the last 10 reaches (trials 21–30), speed improves further and quickly returns to the zero baseline of normal reaching. Reacclimation to one-to-one motion occurred quickly—with average times returning to the baseline within approximately five reaches. However, the *fast interpolated reach* technique again was the exception for this results. Reacclimation to normal took longer after using *fast interpolated reach*, and although the mean time does drop close to the baseline, the average did not reach the baseline in the 10 tasks.

Overall, these results suggest that it does take time to adjust to using remapped techniques, and improvements can be achieved even after a short period of practice. Based on these results, it is expected that it would be difficult to change to different remapping approaches in the same application.

Furthermore, the different outcomes with the *fast interpolated reach* technique demonstrate that it can be easier or harder to adjust to different remapping techniques. Some techniques may require more practice than others. Additionally, the higher number of trials needed to reacclimate to normal reaching suggests it may take longer to "unlearn" remapped motions that are more difficult to learn. We suspect that if a technique is less natural, it will require more attention and conscious effort to effectively control the interaction, and it will therefore take more time to stop using the remapped movement strategy. Also note that participants were not notified of the change from remapped technique back to normal one-to-one motion after the first 20 tasks, so participants could not be sure how the motion would work. It would also be interesting to study normal one-to-one motion in the real world after practicing with the remapped technique (similar to the approach by Bodenheimer et al. [2]).

6 Case Study

After the two controlled experiments, we conducted a case study to better understand practical issues for applying the two remapping techniques in VR applications. The case study primarily involved (1) the design and development of an immersive game, and (2) an informal usability evaluation of the game. The game design also demonstrates how the remapped reaching techniques can be used with multiple virtual objects that can dynamically adjust based on gaze direction.

6.1 Game Design

For the test application, we created an immersive game with two parts: (1) a weapons-crafting scenario where users mix and match different objects to craft weapons, and (2) a monster combat scenario where the user must use the crafted weapons to defeat various creatures. Figures 1 and 8 show screenshots from the game. The game takes place in a forest environment populated with trees and rocks, and the player is seated at a virtual workbench that aligns with the physical table. The game was developed with assets from the Unity Asset Store.

The controls for the game required the user only to reach, grab, and move objects with their tracked right hand. Object selection among multiple targets was done using user gaze based on the view orientation (similar to [6]). Selected objects were highlighted with a green outline to indicate which object was currently selected for interaction. The

Fig. 8. In the game developed in the case study, players combine items to "craft" weapons that they can swing at monsters.

selected object was locked once the hand was no longer close to the body to prevent accidental offset switching during reaching. In order to select a new object, players needed to retract their hand close to their body to allow the gaze selection to be re-enabled.

Instructions for the game were presented at the start of the application which explain to the user how to craft weapons to be used to fight various monsters. For the weapon crafting section of the game, the player is presented with four interactive objects: a sword, a candle, a blue bottle, and a green bottle. In order to craft weapons, the player must grab objects and place them on three different visual markers on the workbench. In order to proceed with the game, the player must combine the four ingredients in all possible combinations of three, which leads to a total of four possible weapons to craft. Each successfully crafted weapon is presented hovering in the background.

Once all four weapons are created, the next phase of the game starts with a monster appearing in front of the player. The player can then select any of the crafted weapons. Once the player grabs a weapon, the player can then swing it at the creature to attack it (as shown in Figure 8). Each creature can only be damaged by a specific weapon, so the player must find which weapon is effective on a specific enemy. Monsters have a health bar indicating current health and amount of damage taken, and if the weapon is ineffective, the health bar will flash without depleting. Once all four enemies are defeated, the player has successfully completed the game.

6.2 Usability Evaluation

We conducted a small usability evaluation of the game using the *translational shift* and *interpolated reach* techniques. Four participants (two male, two female) played through the game twice where each session used a different remapping technique. Two participants started with *translational shift* while the other two participants started with *interpolated reach*. The study was conducted using the same setup as the two experiments with the same VR system, table, chair, and physical prop. All four participants had previously experienced the remapping techniques from the previous experiments. However, other than the written instructions presented in-game, participants were not given any instructions or explanations of how to play the game or what to expect.

The participants played through both the weapon crafting and monster battling portions of the game and took a short break between the two play sessions. Participants were encouraged to express their thoughts while playing the game and ask questions if they were stuck or confused about game mechanics. After completing both sessions, a semi-structured interview was conducted to collect qualitative feedback from the participant on how they felt about the game and their thoughts on using the two different techniques. Each playthrough took approximately 5 to 10 minutes to complete. The second session was always faster than the first since the game was identical for both sessions with the exception of the technique being used, so participants understood the goal of the game by then.

6.3 Discussion and Design Considerations

From implementing the hand remapping techniques in a game application, we came across some unique technical hurdles that did not come up in the experimental scenarios and should be addressed in applications using remapping techniques. One of the main challenges was how to handle object selection with multiple virtual objects. We opted to keep the entire experience hands-free with no controllers so the entire experience would only require physical hand and head interaction. Because object selection was based solely on head orientation, user who prefer to move their eyes to look at objects rather than adjust their head orientation might have difficulty or annoyance with achieving accurate selection. Eye tracking could be an improvement for this and limit the need for additional head and neck movement. Object occlusion is also an issue for gaze-based selection.

With regard for the remapping techniques, the main challenge for *translational shift* was to find an unobtrusive way to transition between offset changes when the user changed target selections. We were initially concerned that abrupt changes to the position of the virtual hand would be jarring or confusing for participants. Fading the hand in and out between updates worked reasonably well, but additional exploration of other options would be interesting for future work.

For *interpolated reach*, one of the biggest problems was dealing with large offsets when reaching distant objects such as hovering weapons in the game (such as the hammer in the left half of Figure 8). Because of the linear nature of the hand movements with the *interpolated reach* technique, hand movements became very sensitive for larger offsets. For example, when reaching for the weapons in the game that were placed at a far distance from the user, the virtual hand moved much faster than the real world hand, making it difficult to control. Because of this sensitive movement, minor differences of where the hand was placed on the bottle would translate to a significant difference of where the virtual hand was positioned in the game. This could be addressed in future iterations of the technique by using an easing function for more gradual interpolation near the hand's starting and ending locations.

Reducing the overall learning curve to play the game is something that should be addressed as well. Participants came in to the study assuming they could interact with the virtual objects similarly to the real world where they can move their hand from one object to another. Participants took some time to realize they needed to bring their hand back to select a new object, which is a limitation of common implementations of remapping techniques (e.g., [1]). More descriptive instructions, visual cues, or practice could help with this problem.

Another note is that because users were manipulating multiple virtual objects in the same virtual space, a physical object could end up far from the user or be placed off the physical table depending on the order of selecting and moving objects. This is a major concern for making remapped reach viable for multiple virtual objects using a single passive-haptic prop. Potential solutions include: taking advantage of change blindness to allow the environment and virtual objects move to more convenient locations when they are not in view (similar to [24]), teleporting the user to a more appropriate view that would reduce difficult reaches or object placement (as in [23]), or providing additional visual cues and warnings. Another possibility is further adjusting the control-display ratio to scale virtual movements such that physical objects could remain within a desirable physical range.

Finally, we note that both remapped reach techniques rely on visually substituting the virtual hand for the physical hand, so the studied techniques and findings may not apply to other types of VR systems (e.g., CAVEs, projection screens, or displays other than HMDs) where you can always see the physical hand and prop.

7 Conclusion

We evaluated two remapping techniques that allow hand-reaching for selection and manipulation of virtual objects using a passive-haptic props. Our findings demonstrate that remapping techniques can work for use with reach, but there are clear performance implications in certain cases. Overall, the *translational shift implementation* was generally superior to *interpolation* techniques for the cases tested in our research.

In Experiment 1, we compared the *translational shift* and *interpolated reach* remapping techniques and compared their performance under various conditions in terms of time and errors, and Experiment 2, we studied acclimation time to adjust using the modified reach techniques. Experiment 1 considered performance implications of different offset sizes, offset direction, and physical object locations. Overall, the *translational shift* technique performed better than interpolated reach and proved to be more robust for situations with larger mismatches between virtual and physical objects. No significant performance penalties were detected with between offsets of 10 and 14 inches for hand reaching, demonstrating that people can generally tolerate some mismatches between virtual and physical content.

Direction of offset was found to have a significant effect on reach time, where actions were slower when the virtual objects were shifted towards the user. Larger offsets (30 inches) were found to be significantly slower than all other offset sizes, and the use of *translational shift* with offsets away from the user were the most tolerable to larger offsets for reach performance. Our results also indicate that reaching errors increased with *interpolated reach* using larger offsets as opposed to *translational shift*, where errors occurred less frequently and more consistently across situations with different offset sizes. Unsurprisingly, the results also show that objects positioned farther from a user can produce more error and longer reach times when the virtual hand is remapped to accommodate mismatch.

Feedback from participants generally matched the performance results, showing a preference for *translational shift* from a majority of participants. However, *translational shift* had problems with participants needing to look for objects and feeling a lack of ownership of the virtual hand. Still, *translational shift* was considered more comfortable overall, and several participants commented on the disorienting nature of *interpolated reach*—especially when the virtual hand moved slower than expected.

The acclimation results from Experiment 2 show that, with practice, users can adjust to remapped motions to improve proficiency. These results suggest that use of modified reach techniques might become easier or feel more natural over time with continued use. On the other hand, users did not adjust to all techniques the same way, and not all participants acclimated to the remapped techniques to the point that reach times reached one-to-one speeds within the scope of the study. In addition, *fast interpolation* showed noticeable acclimation problems when compared to the other remapping techniques. We also tested for reacclimation to normal one-to-one mapping after using the remapped techniques, and results showed that after using *fast interpolated reach*, it took longer to reacclimate back to normal hand movements.

In addition, we presented a case study in the design and development of an immersive game using the remapped reach techniques, and we discussed several practical implications. For use of a single passive-haptic prop with multiple virtual objects and dynamic offset changes during a continued experience, there are challenges to address in terms of object selection and keeping the physical object within physical reach and within bounds of a physical surface area.

Overall, the results of the two experiments and case study demonstrate that remapped reaching with passive-haptic props is a promising approach for enabling realistic tactile feedback for direct hand interaction, but consideration for technique limitations is important for avoiding substantial penalties to performance, comfort, and usability.

REFERENCES

[1] 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, 2016.

[2] B. Bodenheimer, S. Creem-Regehr, J. Stefanucci, E. Shemetova, and W. B. Thompson. Prism aftereffects for throwing with a self-avatar in an immersive virtual environment. In *Virtual Reality (VR), 2017 IEEE*, pp. 141–147. IEEE, 2017.

[3] D. A. Bowman and L. F. Hodges. An evaluation of techniques for grabbing and manipulating remote objects in immersive virtual environments. In *Proceedings of the 1997 symposium on Interactive 3D graphics*, pp. 35–ff. ACM, 1997.

[4] 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.

[5] E. Burns, S. Razzaque, A. T. Panter, M. C. Whitton, M. R. McCallus, and F. P. Brooks. The hand is slower than the eye: A quantitative exploration of visual dominance over proprioception. In *Virtual Reality, 2005. Proceedings. VR 2005. IEEE*, pp. 3–10. IEEE, 2005.

[6] 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. ACM, 2017.

[7] E. Ebrahimi, B. Altenhoff, L. Hartman, J. A. Jones, S. V. Babu, C. C. Pagano, and T. A. Davis. Effects of visual and proprioceptive information in visuo-motor calibration during a closed-loop physical reach task in immersive virtual environments. In *Proceedings of the ACM Symposium on Applied Perception*, pp. 103–110. ACM, 2014.

[8] Z. He, F. Zhu, A. Gaudette, and K. Perlin. Robotic haptic proxies for collaborative virtual reality. *arXiv preprint arXiv:1701.08879*, 2017.

[9] B. E. Insko, M. Meehan, M. Whitton, and F. Brooks. *Passive haptics significantly enhances virtual environments*. PhD thesis, University of North Carolina at Chapel Hill, 2001.

[10] L. Kohli. Redirected touching: Warping space to remap passive haptics. In *3D User Interfaces (3DUI), 2010 IEEE Symposium on*, pp. 129–130. IEEE, 2010.

[11] L. Kohli, M. C. Whitton, and F. P. Brooks. Redirected touching: The effect of warping space on task performance. In *3D User Interfaces (3DUI), 2012 IEEE Symposium on*, pp. 105–112. IEEE, 2012.

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

[13] A. Lécuyer. Simulating haptic feedback using vision: A survey of research and applications of pseudo-haptic feedback. *Presence: Teleoperators and Virtual Environments*, 18(1):39–53, 2009.

[14] W. A. McNeely. Robotic graphics: A new approach to force feedback for virtual reality. In *Virtual Reality Annual International Symposium, 1993., 1993 IEEE*, pp. 336–341. IEEE, 1993.

[15] A. Paljic, S. Coquillart, J.-M. Burkhardt, and P. Richard. A study of distance of manipulation on the responsive workbench (tm). In *Immersive Projection Technology Workshop IPT*, 2002.

[16] 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*, pp. 79–80. ACM, 1996.

[17] I. Poupyrev, T. Ichikawa, S. Weghorst, and M. Billinghurst. Egocentric object manipulation in virtual environments: empirical evaluation of interaction techniques. In *Computer graphics forum*, vol. 17, pp. 41–52. Wiley Online Library, 1998.

[18] I. Poupyrev, N. Tomokazu, and S. Weghorst. Virtual notepad: handwriting in immersive vr. In *Virtual Reality Annual International Symposium, 1998. Proceedings., IEEE 1998*, pp. 126–132. IEEE, 1998.

[19] S. Razzaque, Z. Kohn, and M. C. Whitton. Redirected walking. In *Proceedings of EUROGRAPHICS*, vol. 9, pp. 105–106. Citeseer, 2001.

[20] 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 *Virtual Reality (VR), 2017 IEEE*, pp. 19–28. IEEE, 2017.

[21] E. Schmider, M. Ziegler, E. Danay, L. Beyer, and M. Bühner. Is it really robust? *Methodology*, 2010.

[22] R. J. Stone. Haptic feedback: A brief history from telepresence to virtual reality. In *Haptic Human-Computer Interaction*, pp. 1–16. Springer, 2001.

[23] M. Suhail, S. P. Sargunam, D. T. Han, and E. D. Ragan. Redirected reach in virtual reality: Enabling natural hand interaction at multiple virtual locations with passive haptics. In *3D User Interfaces (3DUI), 2017 IEEE Symposium on*, pp. 245–246. IEEE, 2017.

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

[25] 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.

[26] M. W. Vasey and J. F. Thayer. The continuing problem of false positives in repeated measures anova in psychophysiology: A multivariate solution. *Psychophysiology*, 24(4):479–486, 1987.