Virtually-Extended Proprioception: Providing Spatial Reference in VR through an Appended Virtual Limb

Yang Tian[1], Yuming Bai[2], Shengdong Zhao[3], Chi-Wing Fu[1,2], Tianpei Yang[2], Pheng Ann Heng[1,2]
The Chinese University of Hong Kong, Hong Kong, China[1]
Shenzhen Key Laboratory of Virtual Reality and Human Interaction Technology, SIAT, CAS, Shenzhen, China[2]
NUS-HCI Lab, National University of Singapore[3]
{ytian,cwfu,pheng}@cse.cuhk.edu.hk, zhaosd@comp.nus.edu.sg,{baiyumingc, ytp327}@gmail.com

ABSTRACT

Selecting targets directly in the virtual world is difficult due to the lack of haptic feedback and inaccurate estimation of ego-centric distances. Proprioception, the sense of self-movement and body position, can be utilized to improve virtual target selection by placing targets on or around one's body. However, its effective scope is limited closely around one's body. We explore the concept of *virtually-extended proprioception* by appending virtual body parts mimicking real body parts to users' avatars, to provide spatial reference to virtual targets. Our studies suggest that our approach facilitates more efficient target selection in VR as compared to no reference or using an everyday object as reference. Besides, by cultivating users' sense of ownership on the appended virtual body part, we can further enhance target selection performance. The effects of transparency and granularity of the virtual body part on target selection performance are also discussed.

CCS Concepts

•**Human-centered computing → Virtual reality; Empirical studies in HCI;**

Author Keywords

appended limb; spatial reference; target selection; virtual reality; proprioception

INTRODUCTION

Reaching out to 3D targets in the real world directly with bare hands is generally fast and natural, given the continuous and reliable visual and haptic feedback that we can receive. In contrast, reaching out to 3D targets in virtual worlds is not that straightforward. One possible reason behind this is the lack of haptic feedback to guide us to reach out to the virtual targets. Another is the visual perception in the virtual world, which is different from that in the real world. The egocentric distances perceived by an observer to virtual targets are known to be unreliable. As reported in [37, 28, 18, 2], such distances are usually shorter in the virtual environments. However, the concrete reason behind the problem is still unclear [5, 25].
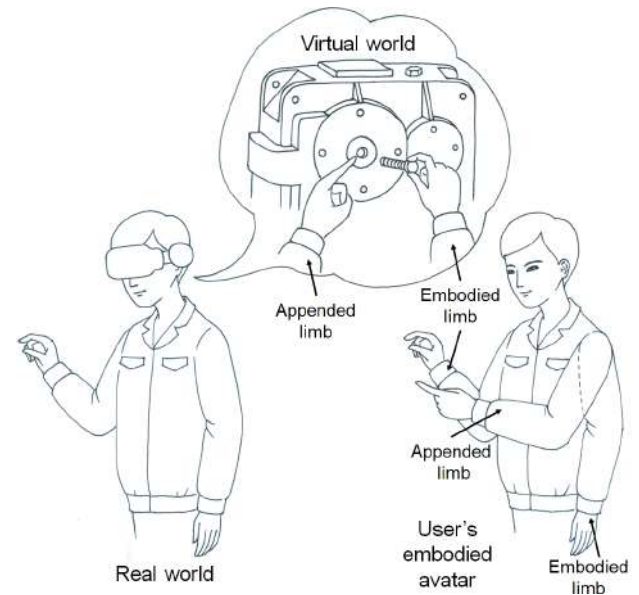
Figure 1. Virtually-extended proprioception concept. We append virtual body parts mimicking real ones to the user's embodied avatar to provide spatial reference to objects or locations in VR. In the example scenario (a virtual assembly) depicted in this figure, we make an appended non-dominant upper limb point on the position of a screw hole. The user lifts his dominant embodied upper limb to insert a screw to the screw hole with the spatial reference provided by the appended upper limb.

Proprioception refers to the sense of one's body posture and actions, even without the aid of visual information. By this sensation, we can use one hand to touch the other one or use a thumb to flick the other fingers with high accuracy. Many works [22, 19, 9, 13] have explored these capabilities to compensate for the problem of egocentric distance perception and the lack of haptic feedback in the virtual world. For example, Mine et al. [22] made use of proprioceptive information to assist VR users in accessing virtual widgets held on their non-dominant hands. They showed that compared with virtual widgets floating in mid-air, it is easier for users to access the hand-held widgets. Lindeman et al. [19] aligned a virtual planar UI with a physical paddle held on one's non-dominant hand, so that the proprioceptive information, together with the haptic feedback, enhances the planar UI interaction. However, under the above methods, users have to lift both arms, which may make the users tired if they need to use the hand-held UIs frequently or for a long time. Gugenheimer et al. [9] attached a back-of-device touchpad on the front side of the VR user's headset for touch input. Very recently, Jiang et al. [13]

placed pressure sensors on fingers for text entry in VR via thumb-to-fingers touches.

Existing works [22, 19, 9, 13] only explored proprioception to aim users to reach out to 2D/3D targets on or closely around the physical body (e.g., hand and head) of the user. In many scenarios, users need to reach out to targets over diverse locations in their peripersonal space [4] in the virtual world, e.g., selecting an installation point on a machine in a virtual assembly, marking the location of a lesion on a virtual patient, dragging a virtual object in the 3D egocentric space, etc.

Motivated by works [39, 35] on adding virtual appendages, e.g., a third arm, a tail, etc., to a user's embodied virtual avatar in VR[1], we propose the concept of *virtually-extended proprioception*, i.e., appending a virtual upper limb (a non-dominant hand with arm), which mimics the user's real one, to the user's embodied virtual avatar; and making the appended upper limb point on or hold the targets located in the user's peripersonal space to provide spatial reference to the targets. Through this setup, the user may acquire targets using their dominant hand more efficiently; see Figure 1 for illustration.

We elaborated the relation between the appended upper limb and proprioception, and conducted three experimental studies to explore the above concept and some related practical factors. Results of Studies 1 and 2 showed that applying the concept of virtually-extended proprioception facilitates the participants to access the target balls of random sizes at various locations (Study 1) and icons distributed on a virtual touchscreen (Study 2) in front of the participants more efficiently. Later, by asking the participants to naturally move their non-dominant embodied upper limb in front of them (Study 3) to cultivate ownership on the appended upper limb, we could further enhance its spatial reference effect. We also explored and discussed the effects of changing the transparency of the appended upper limb and its granularity in Study 3.

RELATED WORK

In this section, we discuss the following three areas of works that are related to our research.

Egocentric Distance Perception Problem in VR

Extensive research [37, 28, 18, 2] has underscored that users tend to underestimate the egocentric distances from themselves to the virtual objects in VR. Though the problem might relate to various factors (see [28]), e.g., the limited field of view (FOV) [40, 3, 16], the quality of graphics in VR [26, 28, 36], etc., the concrete reason is still unclear [25, 5].

Virtual reality has been applied to provide training in many domains, such as surgery, rehabilitation, and industrial assembly. It not only offers standardized, reproducible, and controllable training environments [23], but also helps motivate the trainees [27, 29] and enhance the training performance by providing augmented feedbacks [33]. Regardless of these benefits, the problem of egocentric distance perception persists, hindering how the users perceive in the VR environments and interact with the virtual objects.

---

[1] An embodied virtual avatar is a virtual body in VR that is visuo-proprioceptively matched with the user's physical body.

Several recent works [5, 6, 24] have attempted to alleviate the problem. Finnegan et al. [5] intentionally shifted the positions of visual information and the corresponding audio information to compensate for the deviation in egocentric distances artificially. Later, they proposed an egocentric distance prediction model [6] to better determine the offset between the visual and audio information. However, it applies only to the scenarios where the virtual objects actively emit audio signals. Peer et al. [24] suggested measuring the perception error in personal egocentric distances, based on which the distances to the virtual objects are adjusted. However, in VR training systems, some virtual objects are supposed to be directly manipulated by trainees. If we artificially adjust the virtual objects, their positions may not match the real counterparts, so the kinematic properties of the trainees (e.g., their arm actions) may no longer match those expected in the real world.

Proprioceptive Interaction

Proprioception has been utilized for supporting various input modalities in HCI. Harrison et al. [11] adopted human skin as an input interface by locating finger taps on arms/hands using the vibration signals caused by finger tapping. Wolf et al. [38] showed that, due to the proprioceptive sense on fingers, there is no significant difference in accuracy between selecting targets on the back- and front-side of a tablet. Serrano et al. [31] explored the design space of "hand-to-face" gesture inputs using optical cameras and markers for gesture sensing. Later, Yamashita et al. [41] proposed a more compact device for hand-to-face gestures. Lopes et al. [20] showed that proprioception could also be used as an output modality by actuating a part of the user's body via electrical muscle stimulation.

Due to the egocentric distance problem and the lack of haptic feedback in VR, proprioception received much attention from VR researchers. Many works, e.g., [22, 19, 9, 42, 13], have utilized proprioception to aid the user to reach out to targets that are held on their hands, attached to physical props, or worn on their body. However, the scope where proprioception is effective in facilitating target acquisition is only limited to regions that are closely around the physical body parts. For targets located beyond the above scope, acquiring them in the virtual world is still more difficult than that in the real world. Also, making use of proprioception on the non-dominant hand [22, 19] may require a user to lift their non-dominant arm and keep it still frequently or for a long time, which may bring fatigue to the user. Different from the existing works, using the appended virtual limb to provide spatial reference enables users to select targets more efficiently in their peripersonal space, as compared to no reference or using a familiar-size everyday object as a reference. This phenomenon is like extending the limited effective scope of proprioception from the space closely-around the user's body to their peripersonal space.

Ownership Illusion on Avatar Appendages in VR

The body transfer illusion, e.g., the rubber hand illusion [1], makes the participant feel owning a body part or an entire body that is not of their own. Yuan et al. [43] demonstrated that a relevant illusion in VR, the virtual arm ownership illusion, made the participants feel having an embodied virtual arm. Beyond the ownership of regular virtual avatars, some works

showed that humans could develop ownership on novel virtual avatars with appendages, e.g., a sixth finger [12], a third long arm [39], a tail [35], etc. Won et al. [39] added a third virtual arm longer than the arms of a user's embodied virtual avatar between its two arms. Steptoe et al. [39] added a tail that is also longer than the arms of a user's embodied virtual avatar at its coccyx. By rotating his/her physical arms or head [39], or moving the hip [35], the user can control the movement of the long virtual appendages to reach out to targets located out of the user's physical arm reach.

Compared with the above works [39, 35], we explore how to facilitate VR users to use a real hand to reach out to virtual targets more efficiently by providing spatial reference to targets through virtual appendages. Besides, we demonstrated that actively cultivating the ownership of an appended limb mimicking a user's real one may further enhance the spatial reference effect of the appended limb.

VIRTUALLY-EXTENDED PROPRIOCEPTION

In this section, we first elaborate on the virtually-extended proprioception hypothesis and then describe our design of the appended upper limb in VR to support the concept.

The Virtually-Extended Proprioception Hypothesis

Through a person's rich experience of sensing the proprioception on their own limb (this work considers upper limb), while simultaneously observing the limb's retinal appearance in the real world, we conjecture that *the pose of the upper limb* and its corresponding *retinal appearance* are correlated. Here, the pose of the upper limb refers to one's proprioception sensation, while the retinal appearance of the upper limb refers to the appearance of the limb observed in one's egocentric view. Based on this correlation, we hypothesize that we can induce *a sense of virtually-extended proprioception (VE-Proprioception)* by

> *appending a virtual upper limb (which is rendered like a real one) to a VR user's embodied virtual avatar, such that based on the observed retinal appearance of the virtual limb, the user would develop some level of ownership of the limb, which triggers the user's brain to utilize the above correlation to instinctively acquire plausible poses of the limb in their egocentric view. In this way, we can extend the effective scope of proprioception to some extent in VR through the appended virtual limb.*

If the hypothesis stands, it means that the user can better estimate the pose of the appended virtual limb given the limb's appearance in their egocentric view. In this sense, the user should have acquired a certain 3D spatial context around the appended limb in their view in VR space. Therefore, by specifically posing the appended limb around some virtual targets (e.g., the screw hole in Figure 1), the appended limb can serve as a spatial reference to allow users to estimate the 3D locations of the virtual targets better and help improve the efficiency in selecting and manipulating them. For instance, in VR training systems such as virtual assembly and virtual surgery, the trainees could access the virtual targets through the guidance of the appended virtual limb more efficiently.

The Design of the Appended (Virtual) Upper Limb

To empirically explore the effects of applying the VE-Proprioception concept, we should first consider how to design a virtual upper limb that mimics a real one. To this end, we have made the following considerations.

(i) How to pose the virtual upper limb in the user's view?

First, the pose of the virtual upper limb should be *anatomically plausible*, so that it can appear like a real upper limb. Further, to induce a sensation of ownership, we propose to *extend the virtual upper limb from the user's embodied virtual avatar* by appending it to one of the avatar's shoulders.

(ii) Which shoulder to append the virtual upper limb onto?

If we append the virtual upper limb to the shoulder on the side of the user's dominant hand, when the user physically lifts their dominant upper limb to reach out to virtual targets, the corresponding dominant embodied virtual upper limb could intersect with the appended upper limb in the user's view, thereby producing a conflicting sensation. Although we might fade out the appended upper limb and make it disappear, after or while the user is lifting their dominant upper limb to avoid the conflict, the strategy cannot provide continuous spatial reference in the whole course of the interaction. Therefore, we *append the virtual limb to the shoulder of the user avatar on the side of their non-dominant hand*.

In this manner, the non-dominant appended upper limb can serve as a spatial reference to virtual targets and guide the user to access the targets with the dominant embodied hand. The scenario is similar to everyday real-world situations where we often use our non-dominant hand to hold an object, such as a smartphone, then use the dominant hand to access and interact with the object. The design is also consistent with the Guiard's theory of asymmetric division of labor in human skilled bimanual action [10]. Such similarity between our design and real-world situations may help users bring the existing skills developed in these situations into VR and further develop ownership of the appended upper limb.

(iii) What is the appearance of the appended upper limb?

Ideally, we should provide personalized virtual upper limbs tailored to individual users to enhance the development of ownership. However, current techniques to reconstruct human body parts in VR with high accuracy are rather cumbersome and are not prevalent in general VR systems. Hence, we follow existing VR applications to adopt a universal upper limb model to create and render the appended upper limb in our studies.

Further, in familiar real-world situations where one uses the dominant hand to interact with an object held on the non-dominant hand, one would generally perceive consistent appearance between the dominant and non-dominant upper limbs. The two limbs are said to be *symmetrically congruent*. In VR, we also need to consider symmetrical congruence between the *non-dominant appended upper limb* and the *dominant embodied upper limb* being associated with the user's real dominant upper limb. Hence, in VR, we make the appearance of the dominant embodied upper limb symmetrically congruent with that of the appended upper limb.
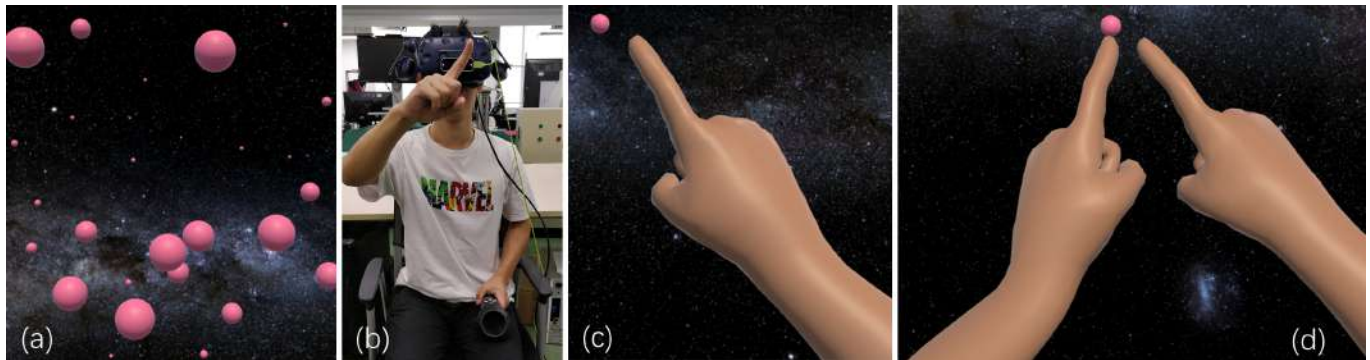
Figure 2. Study 1: reaching out to target balls in the user's egocentric view. (a) Target balls have random positions and sizes. (b) In the experiment, the participant uses the index finger of his dominant hand to reach out to a target ball in VR, and holds a joystick on the non-dominant hand to record the starting time of each trial; note that the participant never lifts their non-dominant limb. (c) Participant's egocentric view in VR, when reaching out to a target ball in the *no reference* condition. (d) Participant's egocentric view in VR, when reaching out to a target ball in the *appended limb* condition.

Overall, the appended upper limb should be anatomically plausible, naturally extended from the user's shoulder on the side of their non-dominant hand, and symmetrically congruent with the user's dominant embodied upper limb.

STUDY 1: REACHING TARGET BALLS

In the formal study, we posed the appended upper limb to point on 3D target balls in VR and explored if the participants can use the index fingertip of the dominant embodied hand to better access the balls in front of them under the appended limb's spatial reference, compared to no reference provided.

One possible condition to include in this study is to use an everyday object to verify if that alone can provide useful spatial reference to improve target selection performance. However, both the literature (Gerig et al. used an HTC VIVE joystick as the familiar object [7]) and our own pilot study with six participants (using a Starbucks coffee mug as the familiar object) failed to reveal any performance gain. We believe that using ordinary everyday objects as spatial references only allows the participants to perceive rough egocentric distances to targets that are probably not enough to enhance selection performance. Therefore, we focused our investigation in using the appended upper limb as a spatial reference in our study.

We started our exploration with the following questions:

Q1: Is the appended limb an effective spatial reference for virtual targets in terms of target selection performance? If the appended upper limb provides effective spatial reference, participants could pay less mental effort to estimate the positions of the target objects, and adopt a more efficient kinematic path to raise a finger of the dominant embodied hand to access the target objects. Hence, we use the average time taken to select targets as the measurement of target selection performance.

Q2: Can participants gain ownership of the appended limb? In the study, we should pay attention to whether the participants can develop ownership of the appended upper limb or not. To facilitate this, we should pose the appended limb to point on the target balls with anatomically plausible poses.

Apparatus
We used an HTC VIVE VR system to provide VR experience and attached a Leap Motion sensor on the front side of the VR headset. Then, we could track the poses of the headset, participants' hands, and forearms; see Figure 2 (b). The software was developed using Unity engine version 2018.3.8f1 in C#.

Participants
We recruited 14 participants: aged 21 to 31; four females and ten males; all are right-handed, and eight had VR experience.

Prepare the target balls and appended limb poses
If we simply randomize the target ball positions, it is challenging to determine anatomically plausible poses of the appended upper limb pointing on target balls located at these positions. Thus, before experimenting with the participants, an experimenter used the index finger of his non-dominant real hand to specify a set of random positions in front of him; meanwhile, we could also get the corresponding plausible poses of his non-dominant real hand with forearm tracked by the Leap Motion sensor. In detail, the experimenter sat on a chair, wore the VR headset, used his dominant real hand to grab a joystick. Then he lifted his non-dominant real hand and used the index fingertip to locate 35 random target positions. For each position, he pressed the joystick's trigger to confirm it. By doing so, we could obtain 35 locations for placing reachable target balls, and the corresponding poses of the non-dominant real hand with forearm. Among the 35 positions, five were randomly chosen for the practice session, and the rest were used in the formal experiment. Furthermore, we used inverse kinematics on the acquired hand and forearm poses to determine the 35 poses of the appended upper limb, such that the appended upper limb could point on each target ball with an anatomically plausible pose.

Task
The task was to use the index fingertip of the dominant embodied virtual hand to reach out to the virtual target balls in mid-air as fast as possible. The participant was seated and told to rest their non-dominant physical upper limb on their thigh on the non-dominant side throughout the experiment. Also, he/she needed to use their non-dominant physical hand to grab
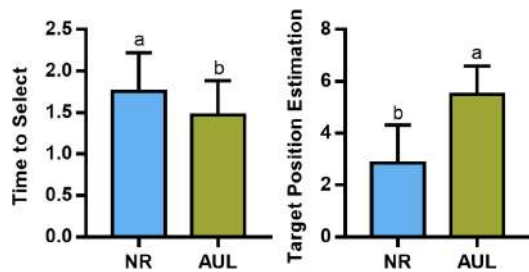
Figure 3. Study 1 results. AUL denotes Appended Upper Limb (See Figure 2 (d)), whereas NR denotes No Reference (See Figure 2 (c)). The error bars in the plots indicate the standard deviation (SD).

a joystick (Figure 2 (b)). They pressed its trigger to make a new target ball appear at a random 3D position and notify our software to record the starting time of each trial. In each trial, the participant first put the index fingertip (dominant embodied limb) within a virtual yellow ball that marked the initial position fixed in front of them. After that, he/she pressed the joystick trigger. At the same time, he/she should immediately use the index fingertip (dominant embodied hand) to select the current target ball. If the fingertip entered the sphere of the ball, the ball disappeared, and our software recorded the completion time of this trial automatically.

Experimental design

Study 1 investigated whether the appended limb provides effective spatial reference to the target balls. We compared the conditions of (i) using the appended upper limb as a reference vs. (ii) no reference. In each trial, a target ball would appear at a random position, with diameter randomly selected in range 3mm to 40mm. Each participant performed the task as described previously under the above two conditions. The order of the two conditions was counter-balanced across the participants. For each condition, the participants had to complete two blocks of trials, each of which was a random permutation of acquiring the target balls at the 30 pre-determined positions.

At the beginning of each condition, we gave each participant a two-minute practice session. Then he/she conducted the formal trials. After finishing the experiment under one condition, he/she received a two-minute rest. After the formal trials, we asked each participant to rate on the following statements in a 7-point Likert scale (from 1 (fully disagree) to 7 (fully agree)): (i) I can estimate the positions of the target balls accurately with *the appended upper limb* (or *no reference*); and (ii) I feel that the appended upper limb is part of my body. In the end, we interviewed each participant about the experiment experience and recorded the discussion for further analysis.

Thus, Study 1 included a total of 14 participants $\times$ 2 conditions (*appended upper limb* vs. *no reference*) $\times$ 2 blocks of trials $\times$ 30 target positions = 1680 trials. The dependent variables were (i) the time between the moment when the joystick trigger was pressed and when a target ball was successfully accessed; (ii) subjective ratings on target position estimation; and (iii) subjective ratings on the ownership of the appended limb.

Results

Figure 3 presents the means (M) and standard deviations (SD) under conditions (i) and (ii). We adopted the t-test to analyze

the *time to select the targets* and the Wilcoxon Signed Rank test to analyze the subjective ratings.

Time to select targets. The time to select the targets under condition (i) appended upper limb ($M=1.47s$, $SD=0.41$) was significantly shorter than that under condition (ii) no reference ($M=1.75s$, $SD=0.47$) with $t(13)=5.076$, $p=0.0002$. This result suggested that the appended limb enabled the participants to estimate the target ball positions more efficiently, as compared to the situation without any spatial reference.

Ratings on target position estimation. The ratings on target position estimation accuracy under condition (i) appended upper limb ($M=5.5$, $SD=1.1$) was significantly better than that under condition (ii) no reference ($M=2.9$, $SD=1.5$) with $Z=-3.096$, $p=0.002$. This result implied that the appended upper limb provides effective spatial reference for participants to feel that they could better estimate the positions of target balls.

Ratings on the ownership. The mean rating (5.2) on the ownership of the appended limb was larger than the middle value (4) in the 7-point Likert scale, indicating that the participants might be able to develop some level of ownership on the appended limb. However, the standard deviation was relatively high (0.94), showing that the ownership might not be steady.

Qualitative feedback. Most participants (13/14) commented that finger-pointing was a very familiar gesture in the real world; thus, a virtual finger pointing on a target ball was intuitive for them, and the fingertip served as a good reference to the target balls. Participant 5 (P5)'s comment was representative, "*In daily life, we often use an index finger to point to objects to guide others to see them; with the guide of the virtual finger, I feel that I can point to the ball instinctively without too much thinking on where the target balls are. When there is no reference, I have to think about the positions of the target balls a little more.*" Several participants (3/14) described the experience of using the dominant embodied index fingertip to flick the index fingertip of the appended upper limb as an experience of "gradually developing ownership on the appended upper limb." For example, P12 commented that "*The more I flick the fingertip of the appended upper limb, the more I feel that the finger is part of my body.*" Several participants (5/14) also mentioned an "ownership transfer phenomenon." They said that they knew well the retinal appearance of the dominant embodied limb after moving and using it and felt it is part of their body. Since the 3D appearance of the appended upper limb was symmetrically congruent with that of the dominant embodied one, they knew well the retinal appearance of the appended upper limb and felt it is part of their body.

Discussion

Q1: Is the appended limb an effective spatial reference for virtual targets in terms of target selection performance? A: Our results have demonstrated that the appended limb was effective in spatially-referencing the virtual targets for faster target selection ($\sim$15% faster) compared to the situation of no reference. The subjective ratings on target position estimation also supported a positive answer to Q1.
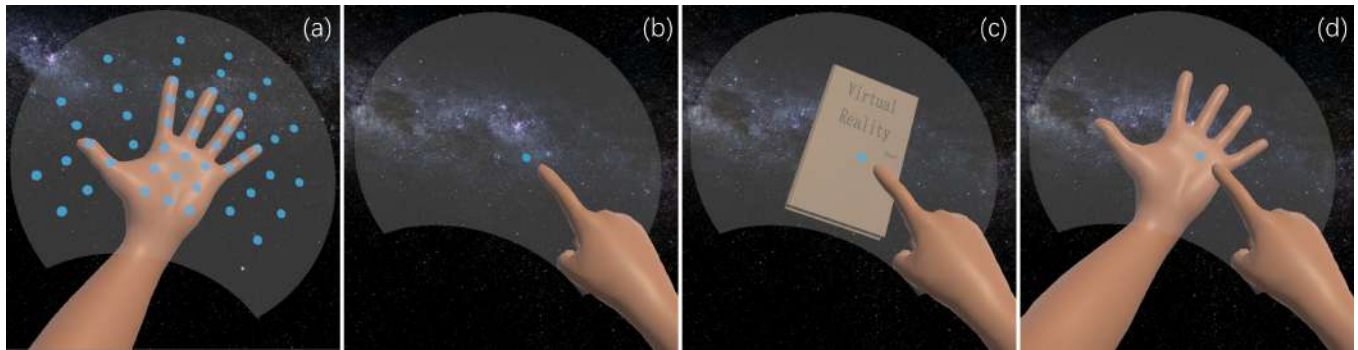
Figure 4. Study 2: selecting icons on planar UI. (a) We sample 50 icon positions roughly evenly on a 2D plane extended from the hand palm of the appended upper limb: twenty-eight icons are within the range of the hand, and the rest (twenty-two) are around the hand. (b)-(d) The participant selects a target icon (b) without spatial reference, (c) with a familiar size object (book) as reference, and (c) with appended limb as reference.

Q2: Can participants gain ownership of the appended limb?
A: The ratings by the participants on appended limb ownership suggested that they might be able to develop ownership on the virtual limb. Also, we observed interesting mechanisms that may facilitate the development of ownership: (i) through interaction between the embodied virtual limb and the appended upper limb, as mentioned in [1]; and (ii) through the usage of the dominant embodied virtual limb [43], participants' sense of ownership of the dominant embodied virtual limb could be symmetrically transferred to the appended upper limb.

Because the appended upper limb mimics a real one (see 3.2), we believe that participants can perceive plausible poses of the limb egocentrically. Besides, the asymmetric division of labor [10] between the non-dominant appended limb and the dominant embodied limb may enable participants to utilize the skills developed in real-world situations (see 3.2) to select targets better in VR. Furthermore, people are familiar with the appearance of upper limbs and their poses because people estimate hand position by vision and proprioception [32] over the years. Thus our VE-proprioception hypothesis, i.e., the effective scope of proprioception can be virtually extended via an appended limb, may be one potential reason why the VE-Proprioception concept works. Similar effects have been observed by extending proprioception into amputees' prostheses [21], which is a phenomenon called extended physiological proprioception (EPP) [34], although their settings focus on extending physical prostheses to disabled people, while VE-Proprioception appends virtual limbs to ordinary people. Although we discussed a few possible reasons behind the performance gain observed in using VE-Proprioception, we do not know how exactly it works in our brain. To fully understand why it works requires more in-depth investigation into our perception and the brain in the future.

STUDY 2: INTERACTING WITH 2D SURFACES
The appended limb affords not only pointing on a 3D position using the index fingertip but also holding a flat surface using the palm. Many interactions are performed on flat surfaces, e.g., writing, drawing, moving objects on tables and floors, etc. Interacting with widgets on 2D surfaces is also very common on touch devices. In VR, similar interactions occur when users need to interact with 2D virtual surfaces for various purposes.

Compared with physical surfaces, virtual surfaces give no haptic feedback. Then it is hard to interact with them. Study 1 has shown that the appended virtual limb provides effective spatial reference to points. In Study 2, we would expand the scope of the VE-Proprioception concept and explore whether our approach can also facilitate interactions with 2D virtual surfaces. We began our exploration with the following questions.

Q3: Can the appended upper limb reduce workload and error when interacting with 2D targets on flat virtual surfaces?
Without haptic feedback, users lost the most direct immediate feedback on whether or not they have successfully touched the surface. They need to rely on additional visual or auditory feedback to help them to make that decision, which can result in delay, hesitation, and inaccuracy in the interactions. We hypothesize that by placing a flat virtual surface on the open hand of the appended upper limb, the limb can provide effective spatial reference to 2D targets on that surface, increase users' confidence, and reduce errors in the interactions.

Q4: Is there any difference between interacting with targets inside and outside the open hand area? The appended virtual limb naturally offers some flat surfaces, e.g., open hand, but with limited area. Thus, we explore the possibility of increasing the effective hand area. If the appended hand can provide effective spatial reference to targets outside, but near the hand, we can arrange more or larger widgets within/around the hand.

Apparatus
We used the same hardware setup as the one in Study 1 (see Figure 2 (b)). The experiment software was also developed using Unity engine version 2018.3.8f1 in C#.

Participants
We recruited 21 participants: aged 20 to 29; 10 females and 11 males; all are right-handed, and 15 had VR experience.

Prepare the Appended Limb Pose and Icon Positions
Figure 4 shows the setup. Mimicking how one holds a tablet using a non-dominant hand in the real world, we rendered a half-transparent flat touchscreen (transparency = 10/255) on the open hand of the appended upper limb with the inside of the hand facing the user. This pose was verified to be anatomically plausible before the experimenter, following a procedure similar to the one in Study 1. We sampled 28 icon
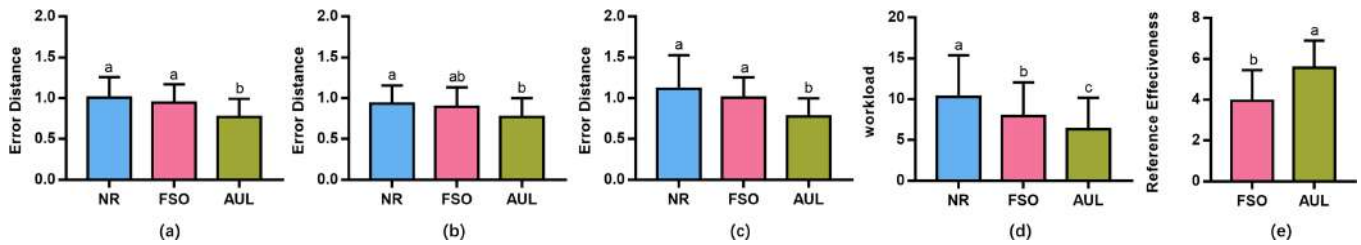
Figure 5. Study 2 results for *NR* (No Reference), *FSO* (Familiar Size Object), and *AUL* (Appended Upper Limb). (a)-(c) show the error distance for the three referencing methods, when considering (a) all 50 icon positions, (b) the 28 icons within the appended virtual hand, and (c) the 22 icons outside but around the virtual hand. (d) The participant workload. (e) The ratings on the effectiveness of *FSO* and *AUL* to enhance the participants' ability to estimate the virtual touchscreen pose. Data is expressed as mean ± standard deviation, so the error bars in the plots indicate the standard deviation.

positions within the open hand, covering positions on palm, on fingers, and between fingers, then sampled another 22 positions outside but near the hand region. We sampled ten additional icon positions for the practice session.

Task

The task procedure was the same as the one in Study 1, but the targets here were random icons on the flat touchscreen held on the open hand of the appended upper limb.

Experimental design

Study 2 explored the effects of spatial referencing on workload and error distance when selecting icons on a 2D virtual touchscreen. We considered three spatial referencing methods: (i) *AUL*: Appended Upper Limb holding the touchscreen; (ii) *FSO*: attaching a Familiar Size Object to the touchscreen; and (iii) *NR*: no reference. For *FSO*, we followed the width and height ($11cm \times 18cm$) of a real book to set up a virtual book. Its thickness ($1.8cm$) was set to the average palm thickness of the participants, after making appointments with them. The area and thickness of the virtual book roughly matched those of the appended virtual hand. Before the experiment, we showed the real book to the participants, so they could see and touch it to understand its width and height better. Also, we explained to them how we determine the virtual book size. We set the icon diameter as 1cm. Each participant performed the task with the three referencing methods. Their order was counterbalanced across the participants using a balanced Latin square design. When each method started, the participant practiced two blocks of trials using the ten additionally sampled icons. For each referencing method, the participant had to complete two blocks of trials, each of which was a random permutation of acquiring the target icons at the 50 pre-sampled positions. After the formal trials, we asked the participants to rate their workload based on NASA-TLX, and a statement using a 7-point Likert scale (from 1 (fully disagree) to 7 (fully agree)): this referencing method effectively enhances my ability to estimate the touchscreen pose (*AUL* or *FSO*).

Thus, Study 2 included a total of 21 participants × 3 referencing methods × 2 blocks of trials × 50 icon positions = 6300 trials. The dependent variables were (i) error distance between the participant's index fingertip and the icon center; (ii) NASA-TLX workload ratings; and (iii) subjective ratings on the effectiveness of the referencing methods.

Results

Figure 5 shows the means (M) and standard deviations (SD) of the dependant variables. We adopted the one-way ANOVA with post hoc Tukey to analyze the error distance data. Then we adopted the Friedman test with post hoc Wilcoxon Signed Rank tests applying the Bonferroni correction to analyze the NASA-TLX workload data. Further, we adopted the Wilcoxon Signed Rank test to analyze the participants' ratings on the effectiveness of spatial referencing with *AUL* and *FSO*.

Error distance for all icons. We found significant difference: $F_{2,20}=14.54$, $p=0.0002$. The mean error distance under *AUL* ($M=0.769cm$, $SD=0.223$) was significantly smaller than that under *FSO* ($M=0.944cm$, $SD=0.226$) with $p=0.0041$ and that under *NR* ($M=1.01cm$, $SD=0.248$) with $p=0.0013$. The mean error distance under *FSO* was just marginally smaller than that under *NR* with $p=0.0761$. This result implied that the appended limb provided more effective spatial referencing than both *FSO* and *NR* for selection on 2D virtual surfaces.

Error distance for icons within hand. There was also significant difference: $F_{2,20}=5.533$, $p=0.0087$. The error distance under *AUL* ($M=0.767cm$, $SD=0.233$) was marginally better than that under *FSO* ($M=0.895cm$, $SD=0.239$) with $p=0.0878$ and significantly better than that under *NR* ($M=0.929cm$, $SD=0.226$) with $p=0.0084$. There was no significant difference between *FSO* and *NR* with $p=0.7641$.

Error distance for icons outside hand. There was also significant difference: $F_{2,20}=10.91$, $p=0.0011$. *AUL* ($M=0.772cm$, $SD=0.227$) was significantly better than *FSO* ($M=1.007cm$, $SD=0.249$) with $p=0.0003$ and *NR* ($M=1.113cm$, $SD=0.414$) with $p=0.0048$. *FSO* and *NR* had no significant difference: $p=0.332$.

Workload. We found significant difference: $\chi^2(2)=18.84$, $p<0.001$. For pairwise comparisons, workload under *NR* ($M=10.3$, $SD=5.1$) was significantly larger than that under *FSO* ($M=7.9$, $SD=4.2$) with $p=0.003$ and *AUL* ($M=6.3$, $SD=3.9$) with $p<0.001$. Workload under *AUL* was significantly lower than that under *FSO* with $p=0.014$. This result implied that *AUL* provides more effective spatial referencing that helped the participants reduce their workload in the task.

Subjective ratings on spatial referencing effectiveness. *AUL* ($M=5.6$, $SD=1.2$) was significantly better than *FSO* ($M=4.0$, $SD=1.5$) with $Z=-2.817$, $p=0.005$. This result meant that participants found that *AUL* gave better spatial referencing than *FSO* did.
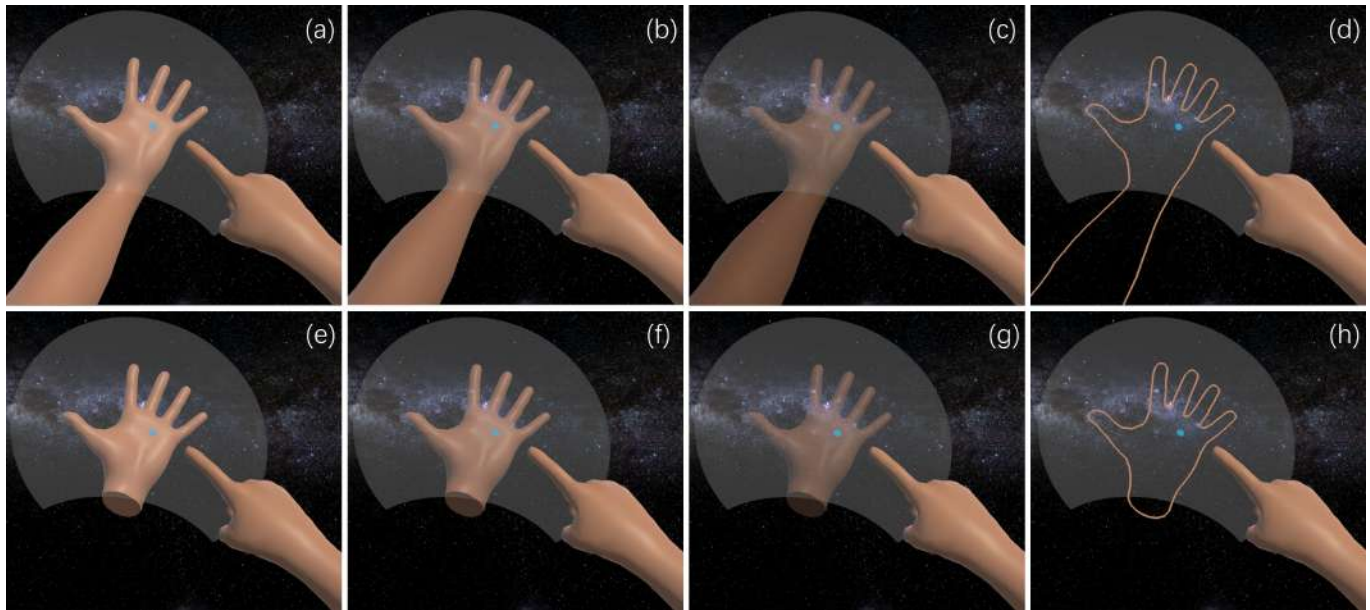
Figure 6. Study 3: exploring the appearance of the appended upper limb. Top row: appended upper limb with both hand and arm. Bottom row: we show only the hand part. (a) & (e): rendering in opaque; (b) & (f): 67% transparency; (c) & (g): 33% transparency; and (d) & (h): outline only.

Qualitative feedback. All the participants (21/21) commented that it was hard to select targets without spatial reference. They tended to move their fingers in and out of the virtual touchscreen a lot to ensure that they touched the targets. Ten participants mentioned that the curved hand profile helped them estimate the pose of the virtual screen. In contrast, although the virtual book was also 3D, they only focused on the book cover. Six participants commented that they felt "pseudo-haptic feedback" through seeing the interaction between the embodied index finger and the appended virtual hand. For example, P1 said, "Because the virtual touchscreen is closely attached to the skin of the appended virtual hand, as long as I poke the finger through the touchscreen a little, I have a feeling of my embodied finger is colliding with the appended virtual hand. Then I will stop my poking." Several participants (4/21) used the word "comfortable" to describe their experience under *AUL*. For example, P18 said, "I felt seeing an appended hand is very comfortable because I think it is soft, so I'm more willing to touch it." Eight others mentioned that they were more familiar with a hand than a book.

Discussion

Q3: Can the appended upper limb reduce workload and error when interacting with 2D targets on flat virtual surfaces? A: The above results showed that both workload and error distance could be significantly reduced with the appended limb providing spatial reference. Thus the VE-Proprioception concept facilitates not only 3D target selection but also selecting 2D targets on flat surfaces. In contrast, although the virtual book was effective in reducing workload, it failed to provide effective spatial reference to reduce the error distances significantly. According to participants' feedback, the virtual book offered less 3D cues than the appended hand did. Future studies should check if other everyday objects with more 3D cues can improve the performance, but due to people's familiarity

with their hands, we believe that it is difficult for ordinary everyday objects to beat the appended upper limb.

Q4: Is there any difference between interacting with targets inside and outside the open hand area? A: Results (inside hand area: M=0.767cm, SD=0.233; outside hand area: M=0.772cm, SD=0.227) showed that the spatial reference effect of the appended upper limb is well-preserved within an area of width ≈3.5cm outside the hand. The comparative advantages of *AUL* over *NR* were ∼31% (outside) and ∼17% (inside). Considering that the 22 icons outside the hand area were located almost at the edge of the participants' field of view (FOV); while the other icons were more in the central part of the FOV, the reason may be that the participants' ability to perceive spatial information without spatial reference within their FOV in VR is anisotropic, which was verified in the very recent work by Peillard et al. [25].

STUDY 3: PRACTICAL FACTORS

In Study 3, we explored additional practical factors when applying the VE-Proprioception concept in VR. First, let us start with the following questions.

Q5: Can the spatial reference effect be further enhanced by actively cultivating ownership on the appended limb? In Q2, we qualitatively evaluated the possibility of cultivating ownership of an appended limb and observed possible mechanisms that may further facilitate the development of ownership. Now, we quantitatively investigate whether actively cultivating ownership by moving and using the non-dominant embodied limb can further enhance the spatial reference effect.

Q6: How transparency affects the spatial reference effect of the appended virtual limb? Rendering the appended limb as an opaque object in VR may fully occlude other objects behind. If we render it with transparency, will it lose its spatial

referencing capabilities? If not, users can benefit from its spatial referencing capabilities while being able to see more of the VR world behind the appended limb. This motivates us to explore the transparency factor. Further, how if we draw only the contour of the limb?

Q7: How does removing the arm of the appended upper limb affect the spatial referencing effect? Another way to reduce the visual footprint of the appended limb is to simplify it by removing the arm and keeping only the hand part. However, we are not sure whether only an appended hand can provide sufficient spatial reference.

Apparatus
We used the same hardware setup and software as the ones used in the first two studies (see Sections 4 and 5).

Participants
We recruited 24 participants: aged 21 to 27; 10 females and 14 males; all are right-handed, and 16 had VR experience. None of them participated in the previous two studies.

Experimental design
Study 3 focused on the three factors mentioned in Q5 to Q7.

- For the first factor, *ownership cultivation*, we considered two conditions: (i) pose and move the non-dominant embodied limb before the experiment starts (*withCultivation*); and (ii) directly start the experiment (*noCultivation*).

- For the second factor, *transparency*, we considered four conditions: (i) opaque; (ii) 67% transparency; (iii) 33% transparency; and (iv) outline only (see the columns in Figure 6).

- For the third factor, *granularity*, we had two conditions: (i) *Hand+Arm* and (ii) *HandOnly* (see the rows in Figure 6).

The first factor, *ownership cultivation*, was a between-subject factor, i.e., twelve participants were under *withCultivation*, while the other twelve were under *noCultivation*. The remaining two factors were within-subject factors. We used balanced Latin square designs for counter-balancing these two factors. Each participant performed 16 sessions of selecting target balls and icons to cover all combinations of the two within-subject factors under one level of factor *ownership cultivation*. In each of the 16 sessions, each participant had to complete two blocks of random permutations of selecting 15 target balls or icons. When a session started, the participant could practice five trials. The target ball or icon positions used in the practice session were different from those in the formal trials. After each session, the participant took rest for 30 seconds. Thus, Study 3 included a total of 24 participants × 2 tasks (select balls and icons) × 4 transparency levels × 2 limb granularity levels × 2 blocks of trials × 15 target positions = 11,520 trials.

The dependent variables were (i) the time between the moment when the joystick trigger was pressed and when a target ball was successfully selected for the selecting target ball task; and (ii) error distance between the participant's index fingertip and the icon center for the selecting target icon task.

Results
We adopted the t-test to analyze the dependent variables for the factors of *ownership cultivation* and *granularity*. We adopted the one-way ANOVA with post hoc Bonferroni to analyze the above dependent variables for the *transparency* factor.

Target ball task (measurement: time to select targets). We found a significant difference with $t(22)=2.458$, $p=0.0223$ between *withCultivation* ($M=1.3s$, $SD=0.2$) and *noCultivation* ($M=1.56s$, $SD=0.31$). We found no significant difference among the four levels of transparency (opaque: $M=1.4s$, $SD=0.3$; 67% transparency: $M=1.36s$, $SD=0.28$; 33% transparency: $M=1.38s$, $SD=0.34$; and outline only: $M=1.46s$, $SD=0.29$) with $F_{3,23}=1.255$, $p=0.2968$. For granularity levels, there was also no significant difference between *Hand+Arm* ($M=1.45s$, $SD=0.32$) and *HandOnly* ($M=1.35s$, $SD=0.26$) with $t(23)=0.4651$, $p=0.6462$.

Target icon task (measurement: error distance). There was no significant difference between *withCultivation* ($M=0.81cm$, $SD=0.17$) and *NoCultivation* ($M=0.92cm$, $SD=0.21$) with $t(22)=1.424$, $p=0.1686$. There was significant difference among the four levels of transparency (opaque: $M=0.8cm$, $SD=0.27$; 67% transparency: $M=0.85cm$, $SD=0.28$; 33% transparency: $M=0.86cm$, $SD=0.29$; and outline only: $M=0.93cm$, $SD=0.31$) with $F_{3,23}=4.731$, $p=0.0133$. In the post hoc comparisons, the error distance under the opaque condition was marginally significantly smaller than that under the outline-only condition with $p=0.076$. There was no significant difference between *Hand+Arm* and *HandOnly* with $t(23)=1.159$, $p=0.2584$.

Discussion
Based on the experiment results, let us respond to Q5-Q7:

Q5: Can the spatial reference effect be further enhanced by actively cultivating ownership on the appended limb? A: Our results seemed to indicate that the effect of active ownership cultivation depended on tasks. For the target ball task, we found a significant enhancement of the time to select the targets (∼16.7% faster). For the target icon task, there was no significant difference between *withCultivation* and *noCultivation*. The reason may be that the hand of the appended upper limb closely behind the virtual touchscreen provided pseudo-haptic feedback, as reported by some participants in Study 2. The pseudo-haptic feedback alone could already avoid the participants from excessively poking their fingers through the virtual touchscreen.

Q6: How transparency affects the spatial reference effect of the appended virtual limb? A: Our results showed that, for the target ball task, there was no significant difference among the four transparency levels, including even the outline-only level. While for the target icon task, we only found a marginally significant difference between the opaque level and the outline-only level. The reason may be that, to select a target ball, the participant only needed to perceive its 3D position, but to select an icon, the participant needed to perceive first the 3D pose of the virtual touchscreen, then the 2D icon position on the touchscreen. A simple outline might provide effective

spatial reference to a 3D position, but might not be sufficient for revealing the 3D pose of a virtual surface.

Q7: How does removing the arm of the appended upper limb affect the spatial referencing effect? A: Our results showed that for both tasks, removing the arm of the appended limb did not affect its spatial reference effect. This result might suggest that, given the participants knew that the appended hand (without the arm) mimics their non-dominant hand, they would *mentally complete* the appended hand with an invisible non-dominant arm in mind. Hence, even an appended hand seemed to be sufficient for providing effective spatial reference. However, some participants (8/24) reported that they felt horrible to see a "broken hand" floating in mid-air. Hence, we should be cautious when applying this idea, e.g., we leave the choice of removing the arm or not to users.

In summary, we found that the benefit of applying the VE-Proprioception concept could be further enhanced by actively cultivating ownership on the appended upper limb; though effects of texture realism [15], skeleton [14] or "invisible body" [17] on ownership have been explored, their effects on the VE-Proprioception phenomenon have not yet been studied. Our results showed that making the appended upper limb half-transparent did not affect its effect of spatial reference to target balls, while only rendering it as an outline could marginally affect its effect of spatial reference to target icons; removing the arm of the appended upper limb did not affect its effect of spatial reference to both target balls and icons; however, we have to consider the emotional factors of doing so.

OVERALL DISCUSSION AND FUTURE WORK

The goal of this work is to explore the effects of applying our VE-Proprioception concept on virtual target selection. Through our user studies, we discuss some common insights and several potential areas worth exploring in the future.

Using Appended Non-dominant Limb to Guide the Interaction of the Dominant Embodied Limb in VR. The essence of the VE-Proprioception concept is to use a non-dominant appended virtual limb to guide or facilitate the interaction of a dominant embodied limb in VR. This VE-Proprioception concept is not limited to the appended upper limb we explored in this work. The VE-Proprioception phenomenon may also be induced on an appended lower limb, which facilitates the interaction of the dominant embodied lower limb. In our user studies, we only recruited right-handed participants to explore the VE-Proprioception concept. In the future, whether the VE-Proprioception concept works with left-handed participants needs further exploration.

Ownership May Be Cultivated Naturally. From Study 3, we observed that active ownership cultivation could further enhance the appended upper limb's spatial reference effect to the target balls. This result provides us an insight: if a VR user always uses a universal upper limb model for embodied upper limbs in all VR applications, he/she will naturally develop strong ownership on the embodied upper limbs [43]. Then a strong level of ownership of the appended upper limb with the same retinal appearance of the non-dominant embodied upper

limb may be naturally cultivated. Then the spatial reference effect of this appended upper limb may be naturally enhanced.

"Ownership Transfer" among Different Levels of Transparency. Study 3 told us that making the non-dominant appended hand (with arm) half-transparent did not affect its spatial reference effect. Surprisingly, rendering the appended upper limb as an outline also did not affect that to target balls. We wonder whether it is because of the within-subject design for the transparency factor. Although we counter-balanced the four levels, for three out of four participant groups, they would first experience at least one non-outline level. Then they might develop some degree of ownership on an opaque or half-transparent non-dominant appended hand (with arm). Then this ownership might be transferred to the outline of the appended hand (with arm), which might enable the participants to estimate the target ball positions still well. Although we need to use a between-subject design to explore the effects of the transparency factor further, we can get an insight for VR designers: the transparency of the appended upper limb can be gradually manipulated from opaque to half-transparent, and even to fully transparent with only an outline.

Is VE-Proprioception really ownership-based? The subjective ratings and qualitative feedback on ownership in Study 1 suggested that participants started to gradually develop ownership over the appended limb during the experiment, and the ownership might be further cultivated or transferred (Study 3). To further validate that the mechanism of the VE-Proprioception phenomenon is ownership/embodiment-based, future evidence provided by additional methods [30, 8] is needed.

CONCLUSION

In this work, we explored the concept of VE-Proprioception, i.e., appending a virtual upper limb mimicking a user's non-dominant real one to their embodied avatar to provide spatial reference to virtual targets. From the first two user studies, we found that applying the VE-Proprioception concept provided effective spatial reference to 3D target balls or 2D target icons. From the third user study, we found that actively cultivating ownership of the appended upper limb could further enhance its spatial reference effect to target balls. Also, we explored the effects of changing the transparency and granularity of the appended upper limb. At last, we discussed several insights and future work about the VE-Proprioception concept.

REFERENCES

[1] Matthew Botvinick and Jonathan Cohen. 1998. Rubber hands "feel" touch that eyes see. *Nature* 391, 6669 (1998), 756.

[2] Sarah H Creem-Regehr, Jeanine K Stefanucci, William B Thompson, Nathan Nash, and Michael McCardell. 2015. Egocentric distance perception in the Oculus Rift (DK2). In *Proceedings of the ACM SIGGRAPH Symposium on Applied Perception*. ACM, 47–50.

[3] Sarah H Creem-Regehr, Peter Willemsen, Amy A Gooch, and William B Thompson. 2005. The influence of restricted viewing conditions on egocentric distance perception: Implications for real and virtual indoor environments. *Perception* 34, 2 (2005), 191–204.

[4] James Cutting and Peter Vishton. 1995. *Perceiving layout and knowing distances: The interaction, relative potency, and contextual use of different information about depth*. Vol. 5. 69–177.

[5] Daniel J Finnegan, Eamonn O'Neill, and Michael J Proulx. 2016. Compensating for distance compression in audiovisual virtual environments using incongruence. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 200–212.

[6] Daniel J Finnegan, Eamonn O'Neill, and Michael J Proulx. 2017. An approach to reducing distance compression in audiovisual virtual environments. In *2017 IEEE 3rd VR Workshop on Sonic Interactions for Virtual Environments (SIVE)*. IEEE, 1–6.

[7] Nicolas Gerig, Johnathan Mayo, Kilian Baur, Frieder Wittmann, Robert Riener, and Peter Wolf. 2018. Missing depth cues in virtual reality limit performance and quality of three dimensional reaching movements. *PLoS one* 13, 1 (2018), e0189275.

[8] Mar Gonzalez-Franco and Tabitha C. Peck. 2018. Avatar Embodiment. Towards a Standardized Questionnaire. *Frontiers in Robotics and AI* 5, Article 74 (2018). DOI: `http://dx.doi.org/10.3389/frobt.2018.00074`

[9] Jan Gugenheimer, David Dobbelstein, Christian Winkler, Gabriel Haas, and Enrico Rukzio. 2016. Facetouch: Enabling touch interaction in display fixed uis for mobile virtual reality. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. ACM, 49–60.

[10] Yves Guiard. 1987. Asymmetric division of labor in human skilled bimanual action: The kinematic chain as a model. *Journal of motor behavior* 19, 4 (1987), 486–517.

[11] Chris Harrison, Desney Tan, and Dan Morris. 2010. Skinput: appropriating the body as an input surface. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 453–462.

[12] Ludovic Hoyet, Ferran Argelaguet, Corentin Nicole, and Anatole Lécuyer. 2016. "Wow! I have six Fingers!": would you accept structural changes of your hand in VR? *Frontiers in Robotics and AI* 3 (2016), 27.

[13] Haiyan Jiang, Dongdong Weng, Zhenliang Zhang, and Feng Chen. 2019. HiFinger: One-Handed Text Entry Technique for Virtual Environments Based on Touches between Fingers. *Sensors* 19 (07 2019), 3063. DOI: `http://dx.doi.org/10.3390/s19143063`

[14] Shunichi Kasahara, Keina Konno, Richi Owaki, Tsubasa Nishi, Akiko Takeshita, Takayuki Ito, Shoko Kasuga, and Junichi Ushiba. 2017. Malleable Embodiment: Changing Sense of Embodiment by Spatial-Temporal Deformation of Virtual Human Body. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 6438–6448. DOI: `http://dx.doi.org/10.1145/3025453.3025962`

[15] Konstantina Kilteni, Antonella Maselli, Konrad P. Kording, and Mel Slater. 2015. Over my fake body: body ownership illusions for studying the multisensory basis of own-body perception. *Frontiers in Human Neuroscience* 9, Article 141 (2015). DOI: `http://dx.doi.org/10.3389/fnhum.2015.00141`

[16] Joshua M Knapp and Jack M Loomis. 2004. Limited field of view of head-mounted displays is not the cause of distance underestimation in virtual environments. *Presence: Teleoperators & Virtual Environments* 13, 5 (2004), 572–577.

[17] Ryota Kondo, Maki Sugimoto, Kouta Minamizawa, Takayuki Hoshi, Masahiko Inami, and Michiteru Kitazaki. 2018. Illusory body ownership of an invisible body interpolated between virtual hands and feet via visual-motor synchronicity. *Scientific reports* 8, 1 (2018), 7541.

[18] Chiuhsiang Joe Lin, Bereket Haile Woldegiorgis, Dino Caesaron, and Lai-Yu Cheng. 2015. Distance estimation with mixed real and virtual targets in stereoscopic displays. *Displays* 36 (2015), 41–48.

[19] Robert W Lindeman, John L Sibert, and James K Hahn. 1999. Towards usable VR: an empirical study of user interfaces for immersive virtual environments. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. ACM, 64–71.

[20] Pedro Lopes, Alexandra Ion, Willi Mueller, Daniel Hoffmann, Patrik Jonell, and Patrick Baudisch. 2015. Proprioceptive Interaction. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 939–948. DOI: `http://dx.doi.org/10.1145/2702123.2702461`

[21] A. Mablekos-Alexiou, G. A. Bertos, and E. Papadopoulos. 2015. A biomechatronic Extended Physiological Proprioception (EPP) controller for upper-limb prostheses. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 6173–6178. DOI: `http://dx.doi.org/10.1109/IROS.2015.7354257`

[22] Mark R Mine, Frederick P Brooks Jr, and Carlo H Sequin. 1997. Moving objects in space: exploiting proprioception in virtual-environment interaction. In *SIGGRAPH*, Vol. 97. 19–26.

[23] M Morel, B Bideau, J Lardy, and R Kulpa. 2015. Advantages and limitations of virtual reality for balance assessment and rehabilitation. *Neurophysiologie Clinique/Clinical Neurophysiology* 45, 4-5 (2015), 315–326.

[24] Alex Peer and Kevin Ponto. 2019. Mitigating Incorrect Perception of Distance in Virtual Reality through Personalized Rendering Manipulation. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 244–250.

[25] Etienne Peillard, Thomas Thebaud, Jean-Marie Normand, Ferran Argelaguet, Guillaume Moreau, and Anatole Lécuyer. 2019. Virtual Objects Look Farther on the Sides: The Anisotropy of Distance Perception in Virtual Reality. In *IEEE VR Conference 2019*.

[26] Lane Phillips, Brian Ries, Victoria Interrante, Michael Kaeding, and Lee Anderson. 2009. Distance perception in NPR immersive virtual environments, revisited. In *Proceedings of the 6th Symposium on Applied Perception in Graphics and Visualization*. ACM, 11–14.

[27] David J Reinkensmeyer and Sarah J Housman. 2007. " If I can't do it once, why do it a hundred times?": Connecting volition to movement success in a virtual environment motivates people to exercise the arm after stroke. In *2007 Virtual Rehabilitation*. IEEE, 44–48.

[28] Rebekka S Renner, Boris M Velichkovsky, and Jens R Helmert. 2013. The perception of egocentric distances in virtual environments-a review. *ACM Computing Surveys (CSUR)* 46, 2 (2013), 23.

[29] Robert Riener, Mathias Wellner, Tobias Nef, Joachim Von Zitzewitz, Alexander Duschau-Wicke, Gery Colombo, and L Lunenburger. 2006. A view on VR-enhanced rehabilitation robotics. In *2006 International Workshop on Virtual Rehabilitation*. IEEE, 149–154.

[30] Daniel Roth, Jean-Luc Lugrin, Marc Erich Latoschik, and Stephan Huber. 2017. Alpha IVBO - Construction of a Scale to Measure the Illusion of Virtual Body Ownership. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '17)*. ACM, New York, NY, USA, 2875–2883. DOI: http://dx.doi.org/10.1145/3027063.3053272

[31] Marcos Serrano, Barrett Ens, and Pourang Irani. 2014. Exploring the Use of Hand-To-Face Input for Interacting with Head-Worn Displays. *Conference on Human Factors in Computing Systems - Proceedings* (05 2014). DOI:http://dx.doi.org/10.1145/2556288.2556984

[32] Brandon M. Sexton, Yang Liu, and Hannah J. Block. 2019. Increase in weighting of vision vs. proprioception associated with force field adaptation. *Scientific reports* 9, 10167 (2019). DOI: http://dx.doi.org/10.1038/s41598-019-46625-7

[33] Roland Sigrist, Georg Rauter, Robert Riener, and Peter Wolf. 2013. Augmented visual, auditory, haptic, and multimodal feedback in motor learning: a review. *Psychonomic bulletin & review* 20, 1 (2013), 21–53.

[34] DC Simpson and others. 1974. The choice of control system for the multimovement prosthesis: extended physiological proprioception (epp). *The control of upper-extremity prostheses and orthoses* (1974), 146–150.

[35] William Steptoe, Anthony Steed, and Mel Slater. 2013. Human tails: ownership and control of extended humanoid avatars. *IEEE transactions on visualization and computer graphics* 19, 4 (2013), 583–590.

[36] William B Thompson, Peter Willemsen, Amy A Gooch, Sarah H Creem-Regehr, Jack M Loomis, and Andrew C Beall. 2004. Does the quality of the computer graphics matter when judging distances in visually immersive environments? *Presence: Teleoperators & Virtual Environments* 13, 5 (2004), 560–571.

[37] David Waller and Adam R Richardson. 2008. Correcting distance estimates by interacting with immersive virtual environments: Effects of task and available sensory information. *Journal of Experimental Psychology: Applied* 14, 1 (2008), 61.

[38] Katrin Wolf, Christian Müller-Tomfelde, Kelvin Cheng, and Ina Wechsung. 2012. Does proprioception guide back-of-device pointing as well as vision?. In *CHI'12 Extended Abstracts on Human Factors in Computing Systems*. ACM, 1739–1744.

[39] Andrea Stevenson Won, Jeremy Bailenson, Jimmy Lee, and Jaron Lanier. 2015. Homuncular Flexibility in Virtual Reality. *Journal of Computer-Mediated Communication* 20, 3 (01 2015), 241–259. DOI: http://dx.doi.org/10.1111/jcc4.12107

[40] Bing Wu, Teng Leng Ooi, and Zijiang J He. 2004. Perceiving distance accurately by a directional process of integrating ground information. *Nature* 428, 6978 (2004), 73.

[41] Koki Yamashita, Takashi Kikuchi, Katsutoshi Masai, Maki Sugimoto, Bruce Thomas, and Yuta Sugiura. 2017. CheekInput: turning your cheek into an input surface by embedded optical sensors on a head-mounted display. 1–8. DOI:http://dx.doi.org/10.1145/3139131.3139146

[42] Xiaoqi Yan, Chi-Wing Fu, Pallavi Mohan, and Wooi Boon Goh. 2016. CardboardSense: Interacting with DIY Cardboard VR Headset by Tapping. In *ACM Conference on Designing Interactive Systems*. ACM, 229–233.

[43] Ye Yuan and Anthony Steed. 2010. Is the rubber hand illusion induced by immersive virtual reality?. In *2010 IEEE Virtual Reality Conference (VR)*. IEEE, 95–102.