Leveraging Change Blindness for Haptic Remapping in Virtual Environments

Andreas L. Lohse Christoffer K. Kjær Ervin Hamulic Ingrid G. A. Lima Tilde H. Jensen *
Luis E. Bruni Niels C. Nilsson†

Aalborg University Copenhagen

ABSTRACT

Passive haptic feedback provides an inexpensive and convenient approach to virtual touch. However, this approach requires that all virtual objects represented by physical props. In this paper we present *change blindness haptic remapping*—a novel approach that leverages change blindness to map two or more virtual objects onto a single physical prop. We describe a preliminary evaluation comparing the proposed approach to a control condition where all virtual objects where mapped to physical props. The study revealed no notable differences in terms of the participants' experience and less than one fourth of the participants noticed the manipulation. However, the participants did perform interaction errors when exposed to haptic remapping. Based on the findings, we discuss improvements to the proposed approach and potential directions for future work.

Index Terms: I.3.7 [Computer Graphics]: Three-Dimenshional Graphics and Realism—Virtual Reality;

1 INTRODUCTION

Current virtual reality (VR) systems offer compelling audio-visual experiences, but it remains difficult to provide users with a realistic sense of touch. The sense of touch plays an integral role during people's everyday interaction with their surroundings [25], and many VR applications demand physical interaction with the virtual environment (VE). Furthermore, it has been demonstrated that, during expsure to VR, haptic feedback positively influences factors including task performance, perceived task performance, realism, and presence [2, 22, 26, 37]. Passive haptics (i.e., the use of physical props serving as proxies for virtual objects) constitutes an inexpensive and convenient approach to virtual touch that offers a number of advantages. Unlike grounded haptic devices, physical props enable the user to interact with virtual objects while moving around the tracked space, and unlike most wearable haptic interfaces, physical props enable kinesthetic perception of object properties (e.g., stimulation of the musculature in response to weight). Finally, the position and orientation of physical props can be tracked using commodity VR systems.

To successfully deploy physical props in VR at least two criteria have to be met: (1) *The criterion of similarity*: The physical props have to be sufficiently similar to the virtual objects they serve as proxies for with respect to features such as shape, weight, and texture. (2) *The criterion of co-location*: All virtual objects the user chooses to interact with should be represented by a co-located physical proxy.

The criterion of similarity poses a problem because it is impractical to change the physical prop when the virtual object changes. *Redirected touching* [13, 15, 16] alleviates this problem by exploiting visual dominance. That is, discrepancies between the user's real and virtual hand motions are introduced in order to generate mappings between a single physical prop and virtual objects of various shapes. Recently, the technique has been extended to complex arbitrary shapes and multi-finger interaction [40].

The challenges associated with the criterion of co-location have rarely been addressed explicitly. Previous work often focused on relatively constrained scenarios involving interaction with a single virtual object represented by a single physical prop (e.g., a tool [8] or a piece of furniture [19]), or circumvented the issue by exposing users to fairly simple interior VEs that were almost fully represented using physical props (e.g., styrofoam walls and wooden boards [11]). Notably, the concept *substitutional reality* [27–29] turns the challenge imposed by the criterion of co-location on its head. Rather than introducing additional props to match the layout of the VE, the VE is adapted to the physical environment to ensure that all physical objects are represented by virtual proxies.

Regardless of whether the physical world is molded to match the VE or vice versa, the utility of physical props decreases in proportion to the number of virtual objects the user can interact with. Thus, the ability to develop complex VEs is constrained by the number of physical props at the user's disposal [1]. For this reason, it has been proposed that the same prop can be used as a proxy for multiple virtual objects [14]. In this paper we present a novel approach to mapping multiple virtual objects onto the same physical prop (i.e., *haptic remapping*) by leveraging *change blindness*—a phenomenon that occurs when individuals fail to detect changes in their environment [20]. Specifically, we propose that visual change blindness can be used to subtlety align virtual objects with physical props when the virtual objects are not visible to the user. Finally, we present a preliminary evaluation suggesting that the proposed approach can be used to map multiple virtual objects onto the same physical prop, albeit in very restricted contexts.

2 RELATED WORK

It is possible to distinguish between at least three general approaches to repurposing the physical props in VR: *dynamic props*, *redirected walking*, and *haptic retargeting*.

2.1 Dynamic Props

Many VR applications rely on head-mounted displays (HMDs) for presenting visual stimuli to the user. Because HMDs deprive users of visual information about their surroundings, it is possible to subtlety rearrange items in the physical environment and thereby repurpose physical props to serve as proxies for multiple virtual objects. The dynamic rearrangement of physical props has previously been performed using mechanical systems operated by a computer. Specifically, Franzluebbers and Johnson [8] describe that mechanical systems, such as robotic arms [10, 21] and drones [12, 39], previously have been used to physically align props with their virtual counterparts. More recently, it has been proposed that other humans can be used to dynamically rearrange physical props. For example, the *TurkDeck* [7] enabled a single user to explore a VE while a group of people ensured that appropriate passive haptic feedback was provided during a range of different interactions, including interactions

*e-mails: {alohse,ckjar ehamul,isampa,tijens}15@student.aau.dk
†e-mails: {leb,ncn}@create.aau.dk

with walls, doors, ledges, steps, and switches. Similarly, the *Haptic Turk* [4] leverages human actuation to generate physical movement during virtual flight, and mutual human actuation has allowed two immersed users, one fishing and one flying a kite, to provide each other with force feedback [5]. Finally, rather than involving other humans, the *iTurk* [3] presents users with a scenario that forces them to reconfigure and animate physical props themselves. The props can be repurposed across different virtual rooms that physically overlap because redirected walking is employed.

2.2 Redirected Walking

Redirected walking, originally proposed by Razzaque et al. [24], refers to a collection of techniques that enable walking users to freely navigate VEs while remaining within a comparatively smaller physical space. Approaches to redirected walking can be divided into two broad categories [23]: First, techniques that manipulate the mapping between the user's real and virtual movements (e.g., through application of translation [38], rotation [24], curvature [30], or bending gains [17]). Second, techniques that produce self-overlapping virtual spaces by manipulating the properties of the VE (e.g., the location of rooms, corridors, and doors [33, 35]). Importantly, redirected walking enables different parts of a VE to occupy the same physical space which creates opportunities for repurposing physical props. Specifically, Steinicke et al. [31] proposed that gains can be used to subtlety steer users toward a physical prop (i.e., a table) when a similar object is encountered in the VE. In a similar manner, Langbehn et al. [18] used bending gains to produce overlapping architecture that allowed users to interact with two virtual items represented by one physical prop, while inside two different virtual rooms connected by a corridor. Finally, Suma et al, [34] showed that *change blindness redirection* [32] can be used to change users' physical path to ensure that they walk across the same patch of physical gravel whenever this surface is presented in the VE.

2.3 Haptic retargeting

The third approach, *haptic retargeting* [1], leverages visual dominance to distort users' perception in a manner similar to redirected touching (see Section 1). However, rather than manipulating perception of an object's shape, haptic retargeting dynamically aligns physical and virtual objects in one of three ways: (1) *Body warping*: The mapping between the user's real and virtual hands is manipulated to ensure that the former reaches the physical prop when the latter reaches the virtual object. (2) *World warping*: The virtual world is continuously rotated around the user to ensure that the virtual and physical objects are co-located. (3) A *hybrid technique*: A combination of body and world warping. Han et al. [9] extended this work by exploring habituation to the technique and different configurations of offset magnitudes, offset directions, and object locations. Finally, Cheng et al. [6] have combined haptic retargeting with on-the-fly target remapping, thereby enabling physical interaction with a single *sparse haptic proxy* providing passive haptic feedback when the user touches several different virtual objects.

3 Change Blindness Haptic Remapping

The technique we propose, dubbed *change blindness haptic remapping* (CBHR), was heavily inspired by change blindness redirection [32]. As allude to earlier, change blindness is a phenomenon that occurs when an individual fails to detect changes in their environment [20]. Specifically, people are susceptible to visual change blindness if their visual field is occluded when the change to the environment is introduced (refer to Suma et al. [33] for a more detailed discussion of change blindness). Change blindness has proven to be very effective for redirection of walkers in interior VEs. Across two user studies, Suma et al. [33] found that only one in 77 participants noticed that the location of doors and corridors were changed behind their backs. Therefore, it seemed plausible that change blindness could be exploited for subtle realignment of real and virtual objects.

In general terms, CBHR is accomplished by aligning a virtual object with an appropriate physical prop when the virtual object is outside the user's field-of-view or when the user's view of the scene is occluded. Thus, like haptic retargeting, the proposed approach is likely to be useful when a user is faced with two or more virtual objects located in relatively close proximity (e.g. a collection of tools on a table). Specifically, the technique resembles haptic retargeting based on world warping (Section 2.3), as both approaches rely on manipulation of the VE. However, unlike world warping, the alignment of real and virtual objects is performed discretely.

The ability to deploy CBHR is constrained by number of static and dynamic factors: The static factors include the users' ability to detect the manipulation, the number of available physical props, and the number of virtual objects that can plausibly be mapped to the physical props (i.e., the objects that fulfill the criterion of similarity described in Section 1). The dynamic factors include access to information about what object the user will interact with next, the user's current viewing direction, and the current positions and orientations of the relevant virtual and physical objects. While the technique can be implemented in a variety of different ways, it is likely to involve at least the following three steps:

1. *Identify target:* The next object the user will interact with is identified. This target may be dictated by the scenario and communicated to the user through explicit instructions or indirect cues. Alternatively, the system may predict the target based on users' behavior and the current state of the scenario.
2. *Create/await opportunity for remapping:* To enable haptic remapping the system may nudge or force the user to look away from the target or introduce a virtual object occluding the relevant part of the scene. Again, these cues may take the form of explicit instructions or indirect cues. Opportunities for haptic remapping may also be produced by delaying information about what object the user should interact with.
3. *Perform remapping:* Once the target is out of view or occluded, the virtual object is translated and rotated to align with the physical prop. If multiple props are available, the target is likely to be aligned with the prop requiring the smallest translation and rotation.

4 User Study

As a preliminary evaluation of CBHR, we performed a between-subjects study comparing CBHR to a control condition where no remapping was performed (i.e., the number of virtual and physical objects were the same). This control condition was chosen because an ideal haptic remapping technique should yield a user experience that is indistinguishable from exposure to a VE where all virtual objects are mapped to physical proxies.

4.1 Participants and Procedure

A total of 26 participants were recruited from the student body at Aalborg University Copenhagen and randomly assigned to one of the two conditions ($n = 13$). The sample comprised 22 males and 4 females, and the participants were aged between 19 and 27 years ($M = 22.8$, $SD = 1.9$). When asked about their previous experience with VR, 6 reported that they had never tried it, 18 had tried it a few times, and 2 had more extensive experience. All participants gave written informed consent prior to participation and were informed that they could opt out at any point during the study. No participants chose to do so. After an introduction to the equipment and virtual interaction, the participants were exposed to the virtual scenario which lasted approximately 15 minutes. Once the participants had completed the scenario they were asked to fill a questionnaire pertaining to their experience.

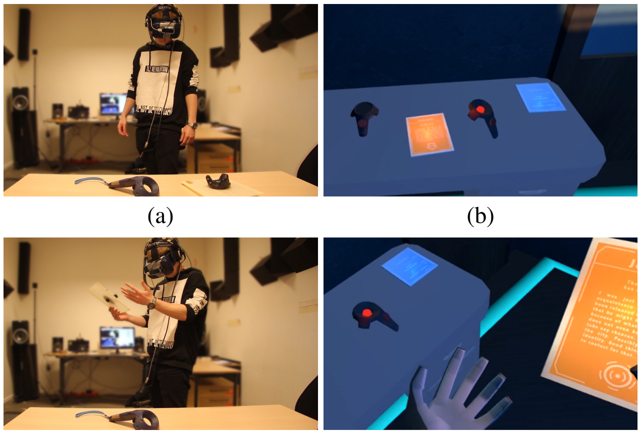

(a) (b) (c) (d)

Figure 1: (a) A user facing the table with two physical props, and (b) the corresponding view of the VE where four virtual objects are presented. (c) The same user interacting with the acrylic plate, and (d) the corresponding view of the virtual hand and journal.

4.2 Equipment

The VE used for the study was developed in Unity 3D and presented using a HTC Vive Pro and a pair of semi-open circumaural headphones. A Leap Motion mounted on the front of the HMD was used to track the users' hand movements which were represented in the VE using a simple grey hand model. The physical props comprised two Vive Controllers and two tablet-sized acrylic plates. Both acrylic plates were equipped with a Vive Tracker. While all four props were used during the control condition, CBHR only involved one acrylic plate and one Vive controller. The study was run on a PC with a i7-6700k processor and a Nvidia Geforce 1070 graphics card. The tracking area was 3m $\times$3m and included two physical tables which also were represented in the VE. Figure 1 shows a user interacting with the system and the corresponding views of the VE.

4.3 Virtual Scenario

The participants in the CBHR group and the control group were exposed to the same virtual scenario; namely, a narrative-driven puzzle game. Specifically, the participants experienced a futuristic detective story where they assumed the role of a detective tasked with solving the case of a murder which happened in a nearby alley. There are three suspects and the detective has to solve the case based on four pieces of evidence as well as information provided by his AI assistant. The story unfolds over three scenes: (1) While in his office, the detective is informed of the crime and likely subjects by his AI assistant. (2) At the crime scene, the detective searches for clues and finds an audio recorder and the murder weapon. (3) The detective returns to his office where he inspects the two pieces of evidence retrieved from the crime scene and two tablets containing a journal and a diary. Throughout the story the user is guided by the AI assistant and the inner monologue of the detective. Additionally, the office contains a large information screen used to display information pertinent to the investigation. The story is concluded by the user having to select which of the three subjects that committed the murder using the information screen.

4.4 Haptic Remapping

The two conditions were identical except from the third scene where the haptic remapping based on change blindness was performed. The remapping involved two pairs of objects: the two tablets displaying the journal and calendar, and the audio recorder and murder weapon. The tablets matched the acrylic plates in shape and size, and both the audio recorder and murder weapon were shaped so that they resembled Vive controllers. The two pairs of items were deliberately made visually distinct. That, is the journal and calendar were bright orange and blue, respectively, and the futuristic murder weapon had a distinct red sphere above the handle whereas the audio recorder was marked with a "W" (Figure 1b).

The scenario required the participants to interact with each of the four virtual objects once. Thus, haptic remapping was performed twice in the CBHR condition which only involved two physical props (one Vive controller and one acrylic plate). That is, initially the calendar is mapped to the acrylic plate and the murder weapon is mapped to the Vive controller. Both instances of haptic remapping are made possible by the AI assistant who explicitly instructions the user to direct attention toward the information screen. This forces the user to turn away from the table where the virtual objects are located and creates an opportunities for haptic remapping. Figure 2 illustrates how the calendar and journal were mapped to the acrylic plate during the scenario.

4.5 Measures

After exposure to the VE, the six-item version of the Slater-Usoh-Steed (SUS) presence questionnaire [36] was administered along with six additional items designed for the current study. The items of SUS questionnaire are answered using rating scales ranging from '1' to '7' where a high rating would be indicative of presence. Four of the six custom items were answered on similar rating scales and asked the participants to evaluate (1) how easy they found the interaction with the virtual objects, (2) whether the presence of the physical props made it easier to complete the task, (3) to what degree they felt that the voice lines had affected their actions, and (4) to what degree they were able to understand the narrative. Finally, items 5 and 6 were binary ('yes' or 'no') and asked whether the participants had noticed changes or movement in the VE. If the participants' answers were affirmative, they were asked to elaborate. In addition to the self-reports, we established the *number of interaction errors* occurring during exposure to CBHM. That is, based on video recordings of each session, we counted the number of times the participants attempted to interact with the calendar, journal, murder weapon, or audio recorder when these objects were not mapped to a physical prop.

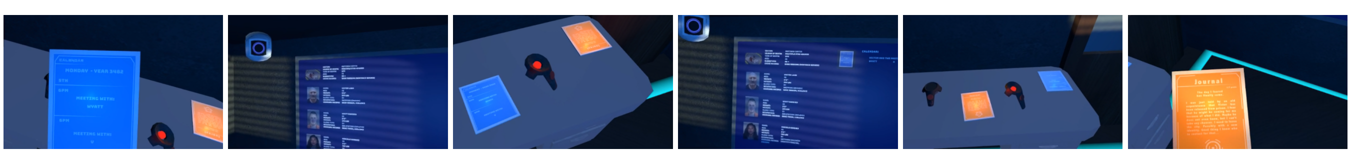

(a) (b) (c) (d) (e) (f)

Figure 2: Illustration of how the blue calendar and orange journal were mapped to a single acrylic plate, as seen from the user's perspective: (a) The user is interacting with the calendar. (b) If the calendar is not returned to the table, the AI assistant instructs the user to do so. (c) The calendar is returned to the table. (d) The AI assistant creates an opportunity for remapping by instructing the user to direct attention to the screen on the wall. (e) The calendar and the journal are swapped. (f) The user interacts with the journal.

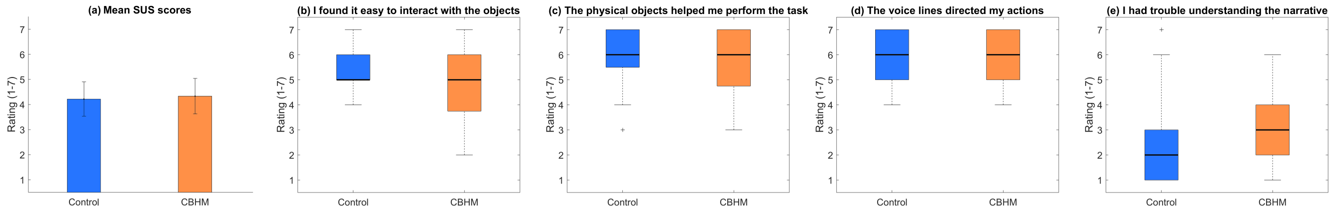

Figure 3: (a) Bar chart showing the grand mean of the mean SUS scores ± 1 SD, and (b-e) boxplots visualizing the results of the four custom items in terms of medians, interquartile ranges, minimum and maximum ratings, and outliers.

5 Results

Figure 3(a) shows the results pertaining to the mean SUS score derived from the six items of the SUS questionnaire. There were no outliers in the data, as assessed by inspection of a boxplot; the scores were normally distributed for both conditions, as assessed by Shapiro-Wilk's test; and there was homogeneity of variances, as assessed by Levene's test for equality of variances. An independent-samples t-test found no statistically significant difference in mean SUS scores between CBHR and the control condition.

The data obtained from the four custom rating scale items was analyzed using Mann-Whitney U tests to determine if there were statistically significant differences in scores between CBHR and the control condition. The tests did not indicate that the median scores differed significantly for any of the four questionnaire items. As apparent from Figure 3(b-e), the two conditions yielded the same median scores for three of the four questionnaire items (b-d), and the median scores pertaining to the fourth item (e) differ slightly, but the corresponding distributions of scores are relatively high.

When the participants who were exposed to CBHR were asked if they noticed that something changed or moved in the last scene, 3/13 (23.1%) correctly noted that the journal and calendar had changed places. This suggests that less than a fourth of the participants noticed that the haptic remapping took place.

Turning to the number of interaction errors during exposure to CBHR: 1 participant made no errors, 1 made a single error, 4 made two errors, and 6 made three errors ($M = 2.25$, $SD = 0.9$). In other words, 12/13 (92.3%) of the participants attempted to interact with the calendar, journal, murder weapon, or audio recorder when these objects were not mapped to a physical prop.

6 Discussion

The descriptive statistics pertaining to the rating scale items (Figure 3) indicate that the participants had relatively similar experiences when exposed CBHM and the control condition. That is, they report having had similar experiences with respect to how present they felt in the VE, how easy they found it to interact with objects, how helpful they found the physical props, how well they were guided by the voice lines, and how well they understood the narrative. This may be viewed as a positive indication considering that an ideal haptic remapping technique should yield user experiences that are indistinguishable from the ones accompanying exposure to a VE where all interactive objects are mapped to physical proxies, as they were in the control condition. Moreover, the results indicate that less than one fourth of the participants noticed the realignment of virtual objects during exposure to CBHR. While this is a promising indication, the fraction is considerably higher than the one in 77 walking users who noticed being manipulated during change blindness redirection [33]. These superior results suggest that it may be possible to decrease the noticeability of CBHR. Thus, it is necessary for future work to explore the factors that may reduce the noticability of CBHM. These factors include the arrangement of the virtual objects that can plausibly be mapped to the same physical prop (e.g., the distances between and number of objects), the visual appearance of the virtual objects, the amount of time between interaction with virtual objects mapped to the same prop, and the cognitive load imposed during this interval.

Notably, all but one participant attempted to interact with virtual objects that were not mapped to a physical prop. This finding may seem to conflict with the relatively low number of participants who noticed the manipulation. However, if the participants did not notice the manipulation, then the interaction errors may have been attributed to tracking problems or similar technical issues. Nevertheless, the number of interaction errors does point to important limitations and highlight potential directions for future work.

The majority of the interaction errors appears to be the result of spontaneous exploration or uncertainty about what object the voice lines encouraged them to interact with. Thus, even though the participants generally reported that the voice lines did direct their actions, they frequently attempted to intact with another object than the one suggested by the voice lines. For example, when instructed to interact with the blue calendar, some participants would attempt to pick up the orange journal instead. They seemingly did so because the instructions did not include a reference to the color of the two objects, making it necessary to read the text on the tablets in order to differentiate between the two. This highlights how crucial unambiguous instructions are if the next target is dictated by the system during *target identification*. Moreover, it seems likely that the explorative nature of the detective scenario indirectly encouraged spontaneous exploration. This suggests that CBHM may be less useful in relation to scenarios designed to encourage exploration.

Taken together these limitations highlight that CBHM probably only is usable in very restricted contexts, and for more generalized use it is likely that CBHM will need to be combined with other approaches, such as haptic retargeting [1] and remapping based on redirected walking [31]. The limitations also suggest the need for expanding the three steps involved in CBHM (see Section 2.3). Specifically, it seems likely that many implementations of CBHM will benefit from including fail-safe mechanisms that intervene if the user is about to interact with virtual objects that have not been remapped to a physical prop (e.g., if no opportunity for haptic remapping has occurred or if the distance between the virtual object and physical prop is too great to enable subtle remapping).

7 Conclusion

In this paper we proposed change blindness haptic remapping—a novel approach leveraging change blindness to repurpose physical props in VEs. The approach involves three steps: target identification, creating/awaiting opportunities for remapping, and discrete remapping (instantaneous realignment of unseen virtual objects with physical props). Moreover, we reported the findings of a user study indicating that the proposed approach can be deployed subtlety. However, additional work is needed to decrease the risk of detection. The study also revealed the importance of clear instructions to the user during target identification, the need for more work on how to create opportunities for remapping, and the need for introducing fail-safe mechanisms that intervene if the user is about to interact with virtual objects that have not been mapped to a physical prop.

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] S. J. Biggs and M. A. Srinivasan. *Haptic interfaces*, pp. 93–116. London: Lawrence Earlbaum, Inc, 2002.

[3] L.-P. Cheng, L. Chang, S. Marwecki, and P. Baudisch. iturk: Turning passive haptics into active haptics by making users reconfigure props in virtual reality. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, p. 89. ACM, 2018.

[4] L.-P. Cheng, P. Lühne, P. Lopes, C. Sterz, and P. Baudisch. Haptic turk: a motion platform based on people. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 3463–3472. ACM, 2014.

[5] L.-P. Cheng, S. Marwecki, and P. Baudisch. Mutual human actuation. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*, pp. 797–805. ACM, 2017.

[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] L.-P. Cheng, T. Roumen, H. Rantzsch, S. Köhler, P. Schmidt, R. Kovacs, J. Jasper, J. Kemper, and P. Baudisch. Turkdeck: Physical virtual reality based on people. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*, pp. 417–426. ACM, 2015.

[8] A. Franzluebbers and K. Johnsen. Performance benefits of high-fidelity passive haptic feedback in virtual reality training. In *Proceedings of the Symposium on Spatial User Interaction*, pp. 16–24. ACM, 2018.

[9] D. T. Han, M. Suhail, and E. D. Ragan. Evaluating remapped physical reach for hand interactions with passive haptics in virtual reality. *IEEE Transactions on Visualization & Computer Graphics*, (1):1–1, 2018.

[10] H. Hoshino, R. Hirata, T. Maeda, and S. Tachi. A construction method of virtual haptic space (ii). *Proceedings of ICAT/VRST*, 95:211–220, 1995.

[11] 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.

[12] P. Knierim, T. Kosch, V. Schwind, M. Funk, F. Kiss, S. Schneegass, and N. Henze. Tactile drones-providing immersive tactile feedback in virtual reality through quadcopters. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, pp. 433–436. ACM, 2017.

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

[14] L. Kohli, E. Burns, D. Miller, and H. Fuchs. Combining passive haptics with redirected walking. In *Proceedings of the 2005 international conference on Augmented tele-existence*, pp. 253–254. ACM, 2005.

[15] 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.

[16] 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.

[17] E. Langbehn, P. Lubos, G. Bruder, and F. Steinicke. Bending the curve: Sensitivity to bending of curved paths and application in room-scale vr. *IEEE transactions on visualization and computer graphics*, 23(4):1389–1398, 2017.

[18] E. Langbehn, P. Lubos, and F. Steinicke. Redirected spaces: Going beyond borders. *IEEE Virtual Reality (VR)*, 2018.

[19] M. Lee, K. Kim, S. Daher, A. Raij, R. Schubert, J. Bailenson, and G. Welch. The wobbly table: Increased social presence via subtle incidental movement of a real-virtual table. In *Virtual Reality (VR), 2016 IEEE*, pp. 11–17. IEEE, 2016.

[20] M. W. Matlin. *Cognition*. John Wiley & Sons, Inc., 7 ed., 2009.

[21] 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.

[22] M. Meehan, B. Insko, M. Whitton, and F. P. Brooks Jr. Physiological measures of presence in stressful virtual environments. *ACM Transactions on Graphics (TOG)*, 21(3):645–652, 2002.

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

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

[25] G. Robles-De-La-Torre. The importance of the sense of touch in virtual and real environments. *Ieee Multimedia*, 13(3):24–30, 2006.

[26] E.-L. Sallnäs, K. Rassmus-Gröhn, and C. Sjöström. Supporting presence in collaborative environments by haptic force feedback. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 7(4):461–476, 2000.

[27] A. L. Simeone. Substitutional reality: Towards a research agenda. In *Everyday Virtual Reality (WEVR), 2015 IEEE 1st Workshop on*, pp. 19–22. IEEE, 2015.

[28] A. L. Simeone and E. Velloso. Substitutional reality: bringing virtual reality home. *XRDS: Crossroads, The ACM Magazine for Students*, 22(1):24–29, 2015.

[29] A. L. Simeone, E. Velloso, and H. Gellersen. Substitutional reality: Using the physical environment to design virtual reality experiences. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pp. 3307–3316. ACM, 2015.

[30] F. Steinicke, G. Bruder, J. Jerald, H. Frenz, and M. Lappe. Estimation of detection thresholds for redirected walking techniques. *IEEE transactions on visualization and computer graphics*, 16(1):17–27, 2010.

[31] F. Steinicke, G. Bruder, L. Kohli, J. Jerald, and K. Hinrichs. Taxonomy and implementation of redirection techniques for ubiquitous passive haptic feedback. In *Cyberworlds, 2008 International Conference on*, pp. 217–223. IEEE, 2008.

[32] E. A. Suma, S. Clark, S. L. Finkelstein, and Z. Wartell. Exploiting change blindness to expand walkable space in a virtual environment. In *Virtual Reality Conference (VR), 2010 IEEE*, pp. 305–306. IEEE, 2010.

[33] 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.

[34] E. A. Suma, D. M. Krum, and M. Bolas. Redirection on mixed reality walking surfaces. In *IEEE VR workshop on perceptual illusions in virtual environments*, pp. 33–35, 2011.

[35] 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.

[36] M. Usoh, E. Catena, S. Arman, and M. Slater. Using presence questionnaires in reality. *Presence: Teleoperators & Virtual Environments*, 9(5):497–503, 2000.

[37] R. Viciana-Abad, A. R. Lecuona, and M. Poyade. The influence of passive haptic feedback and difference interaction metaphors on presence and task performance. *Presence: Teleoperators and Virtual Environments*, 19(3):197–212, 2010.

[38] B. Williams, G. Narasimham, T. P. McNamara, T. H. Carr, J. J. Rieser, and B. Bodenheimer. Updating orientation in large virtual environments using scaled translational gain. In *Proceedings of the 3rd symposium on Applied perception in graphics and visualization*, pp. 21–28. ACM, 2006.

[39] K. Yamaguchi, G. Kato, Y. Kuroda, K. Kiyokawa, and H. Takemura. A non-grounded and encountered-type haptic display using a drone. In *Proceedings of the 2016 Symposium on Spatial User Interaction*, pp. 43–46. ACM, 2016.

[40] Y. Zhao and S. Follmer. A functional optimization based approach for continuous 3d retargeted touch of arbitrary, complex boundaries in haptic virtual reality. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, p. 544. ACM, 2018.