

Propping up Virtual Reality with Haptic Proxies

Niels Christian Nilsson

Aalborg University

André Zenner

Saarland University & German Research Center for Artificial Intelligence (DFKI)

Adalberto L. Simeone

KU Leuven

Abstract—Physical props serving as proxies for virtual objects (*haptic proxies*) offer a cheap, convenient, and compelling way of delivering a sense of touch in virtual reality (VR). To successfully use haptic proxies for VR, they have to be both similar to and co-located with their virtual counterparts. In this paper, we introduce a taxonomy organizing techniques using haptic proxies for VR into eight categories based on when the techniques are deployed (offline or real-time), what reality is being manipulated (physical or virtual reality), and the purpose of the techniques (to affect object perception or the mapping between real and virtual objects). Finally, we discuss key advantages and limitations of the different categories of techniques.

■ **CONSUMER-GRADE VIRTUAL REALITY** (VR) has made it easier than ever for users to immerse themselves in compelling audiovisual virtual environments (VEs). However, it remains challenging to provide a realistic sense of touch in VR. Haptic interfaces, including those used for VR, are often divided into two categories based on whether they provide *active haptic feedback* through computer-controlled actuators exerting forces on the user, or *passive haptic feedback* originating from users' physical interaction with tangible objects.

Haptic proxies are physical props used as proxies for virtual objects during interaction with VEs. These props resemble their virtual coun-

terparts in terms of relevant haptic properties (e.g., shape, weight, or texture) and by definition provide passive haptic feedback. Moreover, haptic proxies can be augmented with computer-controlled actuators; thus combining active and passive feedback (i.e., *mixed haptic feedback* [26]). Nevertheless, it is not trivial to use haptic proxies as a source of touch for VR, and a growing body of work has explored different approaches to doing so. In this paper we discuss two criteria for deploying haptic proxies, present a taxonomy of previous work, discuss strengths and weaknesses of the different categories of techniques, and suggest some directions for future research.

SUCCESS CRITERIA

The benefits of using haptic proxies can be attributed to users' interaction with physical objects. Physical objects eliminate the need for simulating properties such as texture, hardness, weight, shape, and size. However, the utility of haptic proxies decreases in proportion to the complexity of the VE. As VEs grow more complex, a larger number of haptic proxies with different properties is needed. As we have argued elsewhere [18], these constraints can be expressed in terms of two high-level criteria for successfully deploying haptic proxies for VR:

1. *Criterion of Similarity:* All haptic proxies touched by the user should feel sufficiently similar to their virtual counterparts with respect both to material properties (e.g., texture, hardness, and temperature) and geometric properties (e.g., shape, size, and weight).
2. *Criterion of Co-location:* When the user touches a virtual object, it should be co-located with a haptic proxy in a way that allows for seamless interaction (e.g., the transformation of the virtual object should correspond to the position and orientation of the haptic proxy).

Importantly, the two criteria are relevant to interactions involving both *direct touch* with one's hands and *indirect touch* mediated by an object. Also, the criteria are orthogonal, as it is possible to satisfy one without satisfying the other. For example, only a subset of the virtual objects may be co-located with perfect physical replicas, or all virtual objects may be represented by physical props that are not sufficiently similar to their virtual counterparts (see **Figure 1**). Finally, it is likely that the two criteria do not have to be perfectly satisfied to yield a compelling experience (see sidebar).

SIDEBAR: HOW MUCH SIMILARITY AND CO-LOCATION IS ENOUGH?

The benefits of haptic proxies have been demonstrated using props that greatly resemble their virtual representations (e.g., [10]). However, this does not imply that virtual substitutes need to perfectly match haptic proxies with respect to all haptic properties. In fact, recent evidence indicates that users are not able to identify smaller differences between the two. For example, when

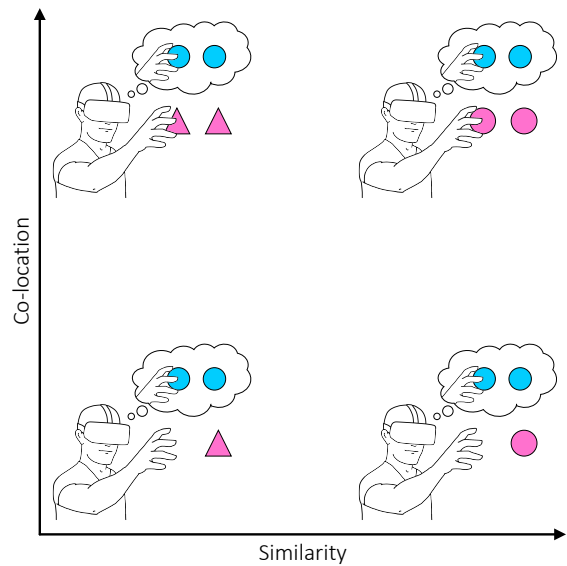


Figure 1. Visualization of the orthogonal nature of the criteria of similarity and co-location. Haptic proxies are shown with purple and virtual objects with blue.

grasping haptic proxies with the thumb and index finger, then users will not reliably detect mismatches of 5.8%, 43.8%, and 66.7% for the width, local orientation and local curvature, respectively [8]. It has yet to be established if these findings generalize to other haptic attributes, but the study indicates that some degree of mismatch is acceptable, even if the VR application aims for a high level of realism. Notably, work by Simeone et al. [20] suggests that the criterion of similarity sometimes can be relaxed even further. They present a user study indicating that, even though great mismatches in terms of shape, temperature, or weight will decrease believability, haptic proxies and virtual substitutes that share similar affordances and few discrepancies at points likely to be touched may still be sufficiently similar to maintain users' suspension of disbelief. A potential implication is that some VR applications (e.g., training systems) may demand high similarity to ensure skill transfer, whereas the criterion may be relaxed somewhat in relation to others (e.g., entertainment applications).

With respect to co-location, it will almost certainly be detrimental to users' experiences, if they attempt to grasp virtual objects and find nothing but thin air. But is some degree of dislocation tolerable? Recent work indicates that misalignments

of more than 1 cm may negatively influence both a user’s experience and performance [11]. However, even though it is possible that some misalignment can go unnoticed, the exact degree of tolerance has yet to be established.

TAXONOMY OF PREVIOUS WORK

Inspired by previous classifications of VR research [22], we propose a taxonomy that organizes techniques using haptic proxies into eight broad categories based on three dichotomous dimensions related to implementation strategies and the aim of the techniques (Figure 2):

- First, we distinguish between approaches based on *when* they are deployed: Are they deployed *offline* before the user is exposed to the VE or in *real-time* during exposure?
- Second, we distinguish between techniques based on *what reality* is being manipulated: Are *physical* or *virtual* objects and environments being manipulated?
- Third, we distinguish between techniques based on *what criterion* they are designed to address: Do they address the criterion of *similarity* or the criterion of *co-location*?

Even though the categories are distinct, they are not mutually exclusive and often complementary. Techniques belonging to different categories can be combined (e.g., real-time techniques often extending their offline counterparts), and even though each category describes techniques that primarily address one criterion, a few techniques are able to address both. In the following we introduce the eight categories, and provide examples of prior work belonging to each category.

OFFLINE VIRTUAL STRATEGIES

Offline techniques are deployed before the VR application is run. During development, virtual objects may be modelled to approximate the haptic properties of one or more physical objects (*Substitutional Reality*), and other techniques aim to ease the process of virtually representing entire physical environments (*assisted VE generation*).

Substitutional Reality

The concept of *Substitutional Reality* challenges the notion that virtual objects need to be replicas of haptic proxies. When introducing the

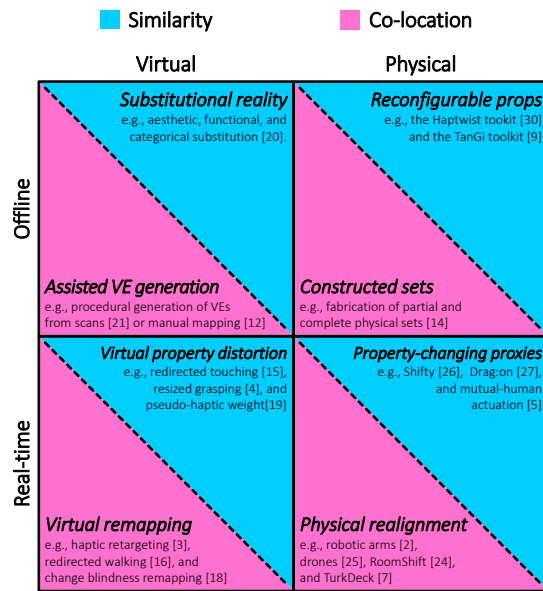


Figure 2. Taxonomy of techniques for deploying haptic proxies for VR. The horizontal axis subdivides the techniques based on what reality is being manipulated (virtual or physical). The vertical axis subdivides the techniques based on when the manipulation occurs (offline or real-time). The division of each cell separates techniques based on whether they address the criteria of similarity (blue) or co-location (purple).

concept, Simeone et al. [20] defined a model of substitution types, starting from the base *replica* level where the virtual substitute is a one-to-one representation of the haptic proxy. Each successive level introduces increasing mismatches between the two. That is, substitution of *aesthetic* features, such as keeping the overall shape and size identical but altering the appearance; substitution through *addition/removal* of physical details; substitution of *functional affordances* (e.g., substituting a book with a box); and *categorical* substitution where the haptic proxy has little or no resemblance to the virtual substitute. This model provides developers with larger creative freedom when designing VR applications.

Assisted VE Generation

Previous work has also sought to map VEs onto specific physical environments; thus ensuring co-location insofar as moving physical objects are tracked during runtime. To ease this process, Garcia et al. [12] developed a system enabling users to draw the volumes of space

where physical objects are located, which then can be substituted with virtual objects. Moreover, Sra et al. [21] present a system that enables users to create VEs from indoor physical environments of varying sizes and shapes. The system uses a Google Tango to construct a 3D map of the physical environment and procedurally generates a corresponding substitutional environment, including a moveable haptic proxy (a chair).

OFFLINE PHYSICAL STRATEGIES

It is usually easier to create virtual content that matches physical reality, than fabricating physical objects and environments that approximate a desired VE. Nevertheless, some works ensure similarity when reusing the same haptic proxy across scenarios (*reconfigurable props*). Moreover, to ensure co-location entire VEs may be recreated physically (*constructed sets*).

Reconfigurable Props

The complexity of VEs dictates how many different haptic proxies are needed. Reconfigurable haptic props help alleviate this problem to some extent. For example, Zhu et al. [30] proposed the *HapTwist* toolkit that, based on a given virtual object, generates a blueprint for creating a suitable physical prop using Rubik's Twists. Similarly, the *TanGi* toolkit by Feick et al. [9] combines primitive shapes with manipulable parts, enabling novices to build physical proxies that can support a variety of different interactions. However, the range of virtual objects that can be represented by such techniques is limited to objects of similar scale.

Constructed sets

The only offline physical strategy for ensuring co-location is the construction of complete physical sets. A notable example from research is Insko's [14] early work, documenting that passive haptic feedback may elicit stronger presence responses. Participants were exposed to a stressful VE where they were standing on the ledge of a 6 m deep pit. While the pit itself was not recreated physically, the rest of the VE, including the ledge, was constructed using wooden boards and styrofoam walls.

REAL-TIME VIRTUAL STRATEGIES

To improve the scalability of offline techniques, haptic perception can be manipulated during runtime. The experience of physically interacting with objects is inherently multisensory, but the sensory information is not always weighted evenly (e.g., vision tends to dominate spatial perception). Because head-mounted displays (HMDs) deprive users of information about the physical environment and their bodies, the audiovisual representation of VEs and the virtual bodies can be manipulated to improve similarity (*virtual property distortion*) or co-location (*virtual remapping*).

Virtual Property Distortion

The concept of *pseudo-haptic* feedback was introduced by Lécuyer et al. [17] when describing how haptic perception can be manipulated using visual feedback. Since then, pseudo-haptics have been extensively used to elicit impressions of haptic properties such as friction, stiffness, mass, and texture, and much of this work has focused explicitly on haptic perception in VR. For example, Samad et al. [19] showed that the perception of a virtual cube's weight can be altered by manipulating the control-display ratio while users are interacting with a haptic proxy (**Figure 3a**).

Virtual real-time strategies can also distort shape perception in VR. Kohli [15] proposed *redirected touching*, which leverages visual dominance over proprioception to address the criterion of similarity. The technique maps differently shaped virtual objects onto a single haptic proxy by introducing discrepancies between users' real and virtual hand motion. This ensures that the user's hand comes into contact with the haptic proxy whenever the virtual object is touched.

Furthermore, Bergström et al. [4] showed that users will tolerate mismatches between the size of real and virtual objects, if *resized grasping* is applied (**Figure 3b**). Specifically, they showed that smaller physical cuboids (3 cm wide) could be used to represent virtual cuboids with widths between 2.7 and 4.4 cm, and larger physical cuboids (9 cm wide) could represent virtual widths between 7.0 and 9.2 cm.

Virtual Remapping

Virtual remapping refers to the process of aligning virtual objects and haptic proxies in real time by warping the VE or the users' movements. To date, three broad approaches have been explored: redirected walking, haptic retargeting, and change blindness remapping. The three vary considerably in terms of their implementation, but they all enable two or more virtual objects to be mapped onto a smaller number of haptic proxies.

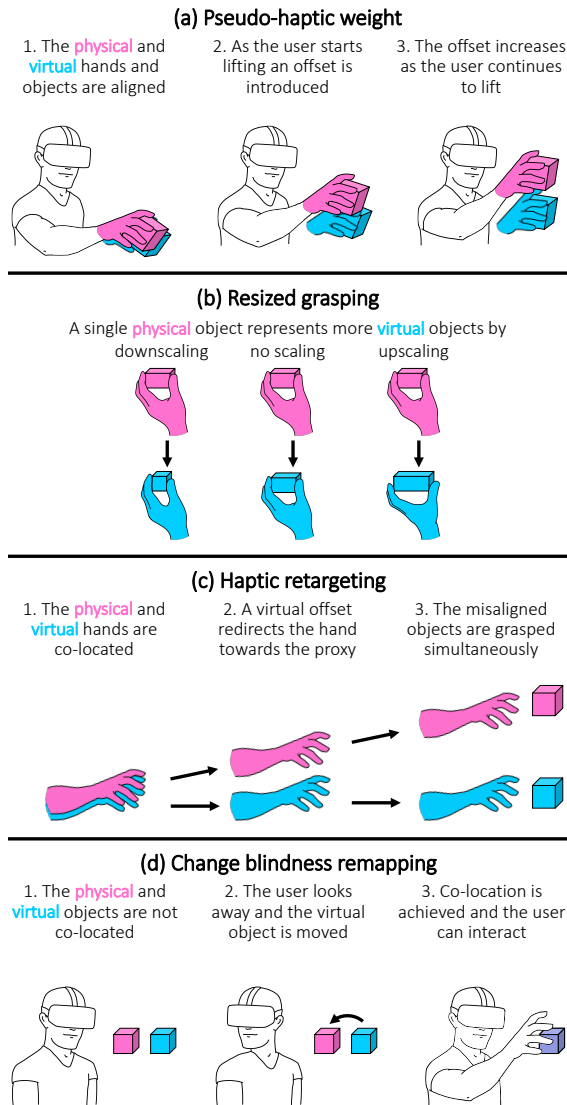


Figure 3. Two examples of virtual property distortion: (a) pseudo-haptic weight [19] and (b) resized grasping [4]; and two examples of virtual remapping: (c) haptic retargeting based on body warping [3] and (d) change blindness remapping [18]

The first approach combines haptic proxies with *redirected walking*. For example, the path of walking users can be changed by subtly manipulating the mapping between their real and virtual movement. Kohli [16] was the first to show that this approach can be used to repeatedly steer users back to a single haptic proxy while maintaining the sensation that they are interacting with a larger number of virtual objects. Another approach to redirected walking involves manipulation of the virtual architecture to produce overlapping virtual spaces. Suma et al. [23] showed that this approach can also be used for virtual remapping. They used architectural manipulations to ensure that users walked across the same patch of physical gravel when this surface was present in the VE. This form of manipulation is necessarily constrained to interior VEs, and remapping based on redirected walking is generally unable to handle scenarios where multiple virtual objects are presented within reach.

The second approach to virtual remapping, *haptic retargeting* [3], addresses this limitation. Like redirected touching, it leverages visual dominance to remap multiple virtual objects to a single haptic proxy. That is, the remapping is performed by dynamically aligning the haptic proxy and virtual objects through warping of the virtual environment, the user's virtual body, or both. A user study indicates that all three options elicit a stronger sense of presence compared to wand-based interaction, and the combination of body and world warping resulted in the highest self-reported presence and satisfaction [3]. Body warping is shown in **Figure 3c**. More recently, Cheng et al. [6] 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. Gonzalez and Follmer [13] explored bimanual haptic retargeting and found that users are more likely to notice the manipulation when their hands are redirected in different directions rather than in the the same direction.

Finally, Lohse et al. [18] proposed that visual change blindness may be leveraged for the purpose of virtual remapping. Change blindness is a phenomenon that occurs when an individual fails to detect changes in their environment, and people

are susceptible to visual change blindness if the changing feature is not visible when the change occurs. *Change blindness remapping* realigns virtual objects with appropriate haptic proxies when the virtual object is outside the user's field-of-view or when the user's view is occluded (**Figure 3d**). More recent work has explored similar ideas and used eye-tracking either to trigger realignment when the users' focus is elsewhere or to trigger discrete hand warping during blinks [28].

REAL-TIME PHYSICAL STRATEGIES

When users are wearing HMDs, it is also possible to introduce unseen changes to the physical environment, which can improve the scalability of passive haptic feedback provided when touching physical proxies, either by establishing similarity (*property-changing proxies*) or by ensuring co-location (*physical realignment*).

Property-Changing Proxies

Property-changing proxies encompasses a broad category of interfaces that help ensure sufficient similarity by physically manipulating haptic proxies. These interfaces may rely on either machine- or human-actuated manipulation to change the physical properties of haptic proxies.

Zenner and Krüger [26] proposed an approach to achieving sufficient similarity based on *dynamic passive haptic feedback* (DPHF), which combines the strengths of passive and active haptic feedback while minimizing their drawbacks. The approach leverages proxy objects augmented with computer-controlled actuators that can adjust the proxy's inherent passive haptic properties. An example for such a dynamic proxy is the weight-shifting VR controller *Shifty* that can change its inertial response by shifting an internal weight using a stepper motor [26]. This allows the proxy to represent various virtual tools or objects, by approximating the haptic stimuli expected by the user when handling objects of various sizes or shapes (**Figure 4a**). Another example is the property-changing proxy *Drag:on* [27]. Using basic actuation by servo motors, it can increase or decrease its surface area through two hand fans attached to the device. By manipulating the proxy's inertial response and air resistance, *Drag:on* approximates the haptic feeling of virtual objects differing in scale, material, or fill

state (**Figure 4b**); and it can be used to provide feedback during interactions involving resistance (e.g., turning a virtual dial).

A radically different approach is human actuation where the physical manipulation is performed by a person. For example, Cheng et al [5] proposed *mutual-human actuation* that delivers force feedback to two VR users while one is fishing and the other is flying a kite. This is accomplished by connecting the haptic proxies held by each user so that the forces exerted by one user are felt by the other.

Physical Realignment

Co-location can also be ensured through real-time physical strategies that realign real and virtual objects. This alignment can be performed by either a machine or a human.

The idea of delivering haptic feedback by presenting physical objects in a just-in-time manner is not new. However, the use of this form of encountered-type haptic feedback to ensure co-




Figure 4. Two examples of property-changing proxies delivering dynamic passive haptic feedback: (a) *Shifty* [26] and (b) *Drag:on* [27].

location is relatively recent. For example, the *Snake Charmer* [2] uses a robotic arm to ensure that haptic proxies are transported to, or held at, the exact location where the user makes contact with the virtual object. Also, this robot arm can automatically exchange between proxies with different material properties; thus, addressing the criterion of similarity. Unlike conventional robotic actuation, drones are not restricted to a limited interaction space. For example, Yamaguchi et al. [25] presented a system using flying drones to mediate different haptic sensations. Users rely on a grasped prop to indirectly touch a virtual creature, and a proxy surface attached to a drone hovering at the location of the creature, provides haptic feedback during interactions. Because they are ungrounded, drones are easily displaced when touched, and are best suited for displaying lightweight or soft virtual objects. Notably, RoomShift [24] is able to present heavy objects, including furniture and walls, by use of nine mobile robots augmented with mechanical scissor lifts. With the appropriate apparatus, the relocation of props can be implemented in fast and reliable ways. However, a general downside of this approach is the required specialized hardware that may constrain the interaction volume and limit the size of the proxies that can be presented.

Researchers have also explored ways to limit machine actuation during physical realignment. Cheng et al. [7] presented a system called *TurkDeck* which relies on human actuation to deliver a multi-sensory VR experience. The experience is made possible by non-VR users, who rearrange and reconfigure different haptic proxies to enable different interactions, including interaction with doors, walls, ledges, and switches. Meanwhile, the system guides this process through laser projections and auditory instructions.

DISCUSSION

The categories of techniques detailed in this paper offer distinct advantages and disadvantages related to their generalizability, complexity, and effects on users' experiences and behaviors.

Offline virtual strategies: Because Substitutional Reality and assisted VE generation permit discrepancies between real and virtual objects, a variety of virtual worlds can be presented in

the same physical environment. However, there are only so many VEs that can be presented in a single physical space, or virtual objects that can be mapped onto a single prop. Thus, the generalizability of offline virtual strategies is greatly constrained by the physical environment.

The design space for offline virtual strategies remains relatively unexplored. It is still difficult to dynamically generate VEs from physical environments; it is not straightforward to differentiate between objects that can be used for interaction and the background VE; and there is a need for authoring tools enabling the creation of virtual content that can be meaningfully deployed across varying physical environments. Finally, even though an increasing volume of work focuses on the extent to which users tolerate mismatches between real and virtual objects and how varying levels of discrepancy affect behavior and performance, these effects are not fully understood, and it remains uncertain how they vary across applications demanding different levels of realism.

Offline physical strategies: Even though constructed sets help ensure co-location, this approach is impractical for most applications; thus generalizability is very limited. On the other hand, reconfigurable props offer a promising approach to ensuring sufficient similarity as long as the prop is the user's main point of physical contact with the environment. The reconfiguration of props is usually limited to the global shape of the object and in a few cases also the weight. Nevertheless, reconfigurable props can be useful in relation to a broad range of scenarios, as long as the scenarios rely on repeated interaction with a single or a limited number of props.

Real-time virtual strategies: Virtual property distortion and virtual remapping take advantage of visual dominance, change blindness, and inattention blindness. Therefore these techniques are only useful insofar as they can be deployed without the user noticing it, which in turn implies that the flexibility of real-time virtual manipulation is constrained by users' capacity for detecting (or tolerating) the manipulations. Moreover, because these techniques rely on visual manipulation, they only work when interactions occur in the user's visual field, and they are unlikely to work when the visuals are impoverished

(e.g., in cases of poor or absent illumination) or when interactions purely rely on haptic sensations (e.g., when touching an object that is out of sight).

Under those circumstances, auditory pseudo-haptics may be of use, and it also seems possible that auditory feedback may be used to improve virtual property distortion; however, research on this topic remains scarce. Moreover, all three approaches to virtual remapping (redirected walking, haptic retargeting, and change blindness remapping) require information about what object the user will interact with next. Thus, it is necessary to rely on scripted scenarios where the system more or less explicitly dictates what actions the user should perform, or else the system will need to predict the user's future actions. Real-time virtual strategies are promising because they can extend the possibilities offered by offline techniques. However, the improved generalizability comes at the expense of computational complexity.

Real-time physical strategies: Property-changing proxies and physical realignment can greatly improve generalizability. Property-changing proxies can help ensure sufficient similarity, and physical realignment guarantees co-location. However, as with virtual strategies, the improved generalizability comes at the expense of computational complexity. Machine-actuated manipulations can be deployed at scale and a single piece of hardware can be programmed to function in different contexts. However, the mechanical actuators introduce additional development costs, and moving parts are subject to damage. Contrarily, human-actuated manipulations do not require complex hardware, apart from suitable reconfigurable props, and human-scale forces can be produced with ease. Nevertheless, to provide compelling experiences, developers have to carefully design the interactions to seamlessly integrate the manual reconfiguration steps, and human-actuated realignment also demands seamless orchestration of human actuators. Thus, human actuation is mainly useful in relation to scripted scenarios.

Hybrid Strategies: As noted, the different strategies are not mutually exclusive. In fact, many real-time techniques augment or extend their offline counterparts. However, very little work has combined real-time virtual and physical

strategies. Recent work indicates that through such combinations, virtual techniques can compensate for the limitations of the physical counterparts, and such hybrid strategies also make it possible to simultaneously address the criteria of similarity and co-location. Abtahi et al. [1] combined encounter-type haptics and haptic retargeting to enable interaction during a virtual shopping scenario. Drones with attached props allowed users to interact with fabrics, hangers, and shoeboxes; and dynamic haptic retargeting [6] was used to compensate for limited drone control accuracy. Even more recently, Zenner et al. [29] showed that the weight-shifting proxy *Shifty* [26] can be combined with haptic retargeting [6] to improve haptic rendering capabilities. Specifically, they showed that the combination can be used to render significantly greater perceived shifts and allow for significantly greater spatial offsets between virtual and real objects.

CONCLUSION

In this paper we argued that two criteria should be met when haptic proxies are used to deliver a sense of touch in VR; namely, the haptic proxies have to be both similar to and co-located with their virtual counterparts. Furthermore, we presented a taxonomy categorizing techniques for using haptic proxies for VR based on three dichotomous dimensions related to when the techniques are deployed (offline or real-time), what reality is being manipulated (physical or virtual reality), and the criterion addressed by the techniques (similarity or co-location). The criteria of similarity and co-location can be addressed both using offline and real-time techniques, and using physical and virtual strategies. However, few techniques are able to simultaneously ensure sufficient similarity and co-location. This suggests the need for more future work exploring hybrid strategies, combining techniques belonging to the eight categories. Additionally, quantification of the extent to which mismatches in terms of similarity and co-location remain unnoticeable, or at least tolerable, to users could drive the development of refined techniques. A better understanding of how such mismatches affect users' experiences, behaviour, performance, and skills transfer, will also help determine when it is beneficial to prop up VR with haptic proxies.

Acknowledgment

This work was supported in part by Internal Funds KU Leuven (C14/20/078) and the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – project number 450247716.

■ REFERENCES

1. P. Abtahi, B. Landry, J. J. Yang, M. Pavone, S. Follmer, and J. A. Landay. Beyond the force: Using quadcopters to appropriate objects and the environment for haptics in virtual reality. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, pages 359:1–359:13, New York, NY, USA, 2019. ACM.
2. B. Araujo, R. Jota, V. Perumal, J. X. Yao, K. Singh, and D. Wigdor. Snake Charmer: Physically enabling virtual objects. In *Proc. TEI*, pages 218–226, New York, NY, USA, 2016. ACM.
3. M. Azmandian, M. Hancock, H. Benko, E. Ofek, and A. D. Wilson. Haptic retargeting: Dynamic repurposing of passive haptics for enhanced virtual reality experiences. In *Proceedings of the 2016 chi conference on human factors in computing systems*, pages 1968–1979. ACM, 2016.
4. J. Bergström, A. Mottelson, and J. Knibbe. Resized grasping in vr: Estimating thresholds for object discrimination. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, pages 1175–1183. ACM, 2019.
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*, pages 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*, pages 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*, pages 417–426. ACM, 2015.
8. X. de Tinguy, C. Pacchierotti, M. Emily, M. Chevalier, A. Guignardat, M. Guillaudeux, C. Six, A. Lécuyer, and M. Marchal. How different tangible and virtual objects can be while still feeling the same? In *2019 IEEE World Haptics Conference (WHC)*, pages 580–585. IEEE, 2019.
9. M. Feick, S. Bateman, A. Tang, A. Miede, and N. Marquardt. Tangi: Tangible proxies for embodied object exploration and manipulation in virtual reality. In *2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 195–206. IEEE, Nov 2020.
10. A. Franzluebbbers and K. Johnsen. Performance benefits of high-fidelity passive haptic feedback in virtual reality training. In *Proceedings of the Symposium on Spatial User Interaction*, pages 16–24. ACM, 2018.
11. S. Fremerey, M. S. Suleman, A. H. A. Paracha, and A. Raake. Development and evaluation of a test setup to investigate distance differences in immersive virtual environments. In *2020 Twelfth International Conference on Quality of Multimedia Experience (QoMEX)*, pages 1–4. IEEE, 2020.
12. J. F. Garcia, A. L. Simeone, M. Higgins, W. Powell, and V. Powell. Inside looking out or outside looking in? an evaluation of visualisation modalities to support the creation of a substitutional virtual environment. In *Proceedings of the 2018 International Conference on Advanced Visual Interfaces*, pages 1–8, 2018.
13. E. J. Gonzalez and S. Follmer. Investigating the detection of bimanual haptic retargeting in virtual reality. In *25th ACM Symposium on Virtual Reality Software and Technology*, VRST '19, New York, NY, USA, 2019. Association for Computing Machinery.
14. 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.
15. L. Kohli. *Redirected touching*. PhD thesis, The University of North Carolina at Chapel Hill, 2013.
16. 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*, pages 253–254. ACM, 2005.
17. A. Lecuyer, S. Coquillart, A. Kheddar, P. Richard, and P. Coiffet. Pseudo-haptic feedback: Can isometric input devices simulate force feedback? In *Virtual Reality, 2000. Proceedings. IEEE*, pages 83–90. IEEE, 2000.
18. A. L. Lohse, C. K. Kjær, E. Hamulic, I. G. Lima, T. H. Jensen, L. E. Bruni, and N. C. Nilsson. Leveraging change blindness for haptic remapping in virtual environments. In *2019 IEEE 5th Workshop on Everyday Virtual Reality (WEVR)*, pages 1–5. IEEE, 2019.
19. M. Samad, E. Gatti, A. Hermes, H. Benko, and C. Parise. Pseudo-haptic weight: Changing the per-

- ceived weight of virtual objects by manipulating control-display ratio. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, page 320. ACM, 2019.
20. 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*, pages 3307–3316. ACM, 2015.
 21. M. Sra, S. Garrido-Jurado, C. Schmandt, and P. Maes. Procedurally generated virtual reality from 3d reconstructed physical space. In *Proceedings of the 22nd ACM Conference on Virtual Reality Software and Technology, VRST '16*, page 191–200, New York, NY, USA, 2016. Association for Computing Machinery.
 22. E. A. Suma, G. Bruder, F. Steinicke, D. M. Krum, and M. Bolas. A taxonomy for deploying redirection techniques in immersive virtual environments. In *2012 IEEE Virtual Reality Workshops (VRW)*, pages 43–46. IEEE, 2012.
 23. 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*, pages 33–35, 2011.
 24. R. Suzuki, H. Hedayati, C. Zheng, J. L. Bohn, D. Szafir, E. Y.-L. Do, M. D. Gross, and D. Leithinger. Roomshift: Room-scale dynamic haptics for vr with furniture-moving swarm robots. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–11, 2020.
 25. 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*, pages 43–46. ACM, 2016.
 26. A. Zenner and A. Krüger. Shifty: A weight-shifting dynamic passive haptic proxy to enhance object perception in virtual reality. *IEEE transactions on visualization and computer graphics*, 23(4):1285–1294, 2017.
 27. A. Zenner and A. Krüger. Drag:on – a virtual reality controller providing haptic feedback based on drag and weight shift. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI '19*, pages 211:1–211:12, New York, NY, USA, 2019. ACM.
 28. A. Zenner and A. Regitz, Kora Persephone and Krüger. Blink-suppressed hand redirection. In *2021 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE.
 29. A. Zenner, K. Ullmann, and A. Kruger. Combining dynamic passive haptics and haptic retargeting for enhanced haptic feedback in virtual reality. *IEEE Transactions on Visualization and Computer Graphics*, 2021.
 30. K. Zhu, T. Chen, F. Han, and Y.-S. Wu. Haptwist: creating interactive haptic proxies in virtual reality using low-cost twistable artefacts. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–13, 2019.

Niels Christian Nilsson is an associate professor at Aalborg University Copenhagen, Denmark. His research is broadly focused on locomotion, perception, and cognition in immersive virtual environments. Contact him at ncn@create.aau.dk.

André Zenner is a PhD candidate at Saarland University and the German Research Center for Artificial Intelligence (DFKI), Germany. His research focus lies on techniques overcoming the drawbacks of proxy-based haptics for VR using dynamic passive haptics and virtual manipulations (e.g., through hand redirection). Contact him at andre.zenner@dfki.de.

Adalberto L. Simeone is an assistant professor at the KU Leuven, Belgium. His research strives to find creative answers to the fundamental interaction challenges of XR. Contact him at adalberto.simeone@kuleuven.be.