THEME ARTICLE: Virtual and Augmented Reality

15 Years of Research on Redirected Walking in Immersive Virtual Environments

Niels Christian Nilsson
Aalborg University

Tabitha Peck
Davidson College

Gerd Bruder
University of Central Florida

Eric Hodgson
Miami University, Ohio

Stefania Serafin
Aalborg University

Mary Whitton
University of North Carolina at Chapel Hill

Evan Suma Rosenberg
University of Southern California

Frank Steinicke
University of Hamburg

Virtual reality users wearing head-mounted displays can experience the illusion of walking in any direction for infinite distance while, in reality, they are walking a curvilinear path in physical space. This is accomplished by introducing unnoticeable rotations to the virtual environment—a technique called *redirected walking*. This paper gives an overview of the research that has been performed since redirected walking was first practically demonstrated 15 years ago.

Locomotion, the act of moving from one location to another, is considered one of the most fundamental and universal activities performed during interaction within Virtual Reality (VR) To most people locomotion is a trivial and frequent everyday activity. However, allowing users to freely *walk* through virtual environments (VEs) presents a considerable challenge. Ideally the walker's movement within a VE should be constrained only by the virtual topography and architecture, not by the size of the available physical space.

Redirection techniques manipulate the physical transformations of the user's movement within the VE so that motion is no longer mapped 1:1, or they manipulate the physical characteristics, e.g., the architecture, of the VE. These manipulations make it possible to guide the physical path of walking users so they can travel through VEs larger than the available tracking space. The literature on virtual walking describes other locomotion techniques intended to facilitate unconstrained walking in large VEs, e.g., omnidirectional treadmills and walking-in-place techniques.

Published by the IEEE Computer Society
0272-1716/18/$33.00 ©2018 IEEE

However, redirection techniques have at least one considerable advantage: the user is physically walking and therefore experiences correct proprioceptive, kinesthetic, and vestibular stimulation.

The shared aim of all redirection techniques is to allow users to walk freely within VEs. On the most general level there are techniques that strive to maintain an illusion of walking freely, i.e., *redirected walking* (RDW), and techniques that provide users with a (non-walking) means of transportation to a place in the VE from which they can then navigate on foot to the extent that the physical space allows it. These *relocation techniques* are designed to enable walking and therefore fall under the broad category of redirection techniques.[1] However, because the primary mode of locomotion is not walking, they do not properly qualify as walking techniques. Contrarily, RDW focuses exclusively on enabling users to walk through VEs. The ideal RDW technique has at least four criteria: it is *imperceptible*, the user is unaware that redirection is taking place; it is *safe*, the walker is prevented from leaving the tracking space and colliding with physical obstacles and other users; it is *generalizable*, it is applicable within any VE and with any number of users; and it is *devoid of unwanted side effects*, it does not introduce cybersickness or interfere with primary and secondary tasks. The ability of a RDW method to meet these criteria is constrained by both static and dynamic factors. Static factors include the size and shape of the available tracking space and of the VE, the users' ability to detect the manipulation, the number of users, and access to information about the users' expected path. The dynamic factors are the users' current and past positions within the real and the virtual environment.

Even though no single RDW technique can claim to be imperceptible, safe, generalizable, and devoid of unwanted side effects, much work has addressed these challenges in the 15 years since RDW was first described. This article is a survey of the literature on RDW from Razzaque et al.'s original paper[2] to recent innovations.

APPROACHES TO REDIRECTED WALKING

While numerous different forms of RDW have been proposed, they are all based on one of two forms of manipulation: (1) RDW techniques that *manipulate the mapping* between the user's real and virtual translation and rotation and thereby steer the user away from the edges of the tracking space and physical obstacles, and (2) RDW techniques that *manipulate the architectural properties of the VE*, e.g., manipulate the location of rooms, hallways, and doors, to produce self-overlapping virtual spaces that make it possible to compress comparatively large VEs into smaller tracking spaces. It is possible to differentiate between redirection techniques that can be applied without the user's knowledge (subtle techniques) and techniques that are detectable by the user (overt techniques).[1] The former category fulfills the criteria of being imperceptible and is thus preferable for most applications. However, sometimes practical limitations or the need to adhere to the criteria of safety make the use of overt techniques necessary. In what follows, we present an overview of techniques that rely on subtle or overt manipulation of the mapping between the user's real and virtual transformation and subtle or overt manipulation of the virtual architecture.

Subtle Manipulation of Gains

The first implementation of RDW relied on manipulation of the mapping between the user's real and virtual rotation.[2] Particularly, this technique steered users down a virtual hallway along a pre-defined path through the application of imperceptible *rotation gains* to the user's natural head rotations, while in fact users were walking back and forth within the tracked-space. While a rotation gain in principle can be applied to each component of the rotation (pitch, yaw, roll), manipulation of yaw is commonly used for the purpose of redirection. Recently other types of gains have been proposed to control the walker's physical path (see Figure 1).

Translation gains can be used to scale the user's forward steps and thereby enable mapping of a larger VE to a smaller physical space. *Curvature gains* are a continuous rotation applied while the user is walking forwards and may allow users to walk infinitely along a virtual path while walking in circles in reality.

Bending gains in a similar manner bend the walkers' paths to the left or right while they are walking along a curved virtual path. Note that curvature and bending gains require *a priori* knowledge of the user's future path or a good prediction of that path generated by extrapolating from the recent path. In order for the applied gains to remain imperceptible they need to be below the users' perceptual threshold. Estimation of detection thresholds for different types of gains continues to be an active area of research (for a summary of this work see the sidebar).

Finally, most subtle manipulations of the mapping between the user's real and virtual transformation are applied continuously because the abrupt nature of instantaneous translation or rotation makes imperceptibility challenging. Nevertheless, Bruder et al.[3] demonstrated the possibility of masking instant translations using inter-stimulus images and visual optic flow effects in the user's peripheral vision, suggesting that repeatedly performed small discrete translations may enable continuous subtle scaling of a walking user's motion.[1] Moreover, it has been demonstrated that subtle translation of the user's viewpoint can be applied during eye blinks or saccades (i.e., the rapid eye movement accompanying changes in gaze direction). [4]

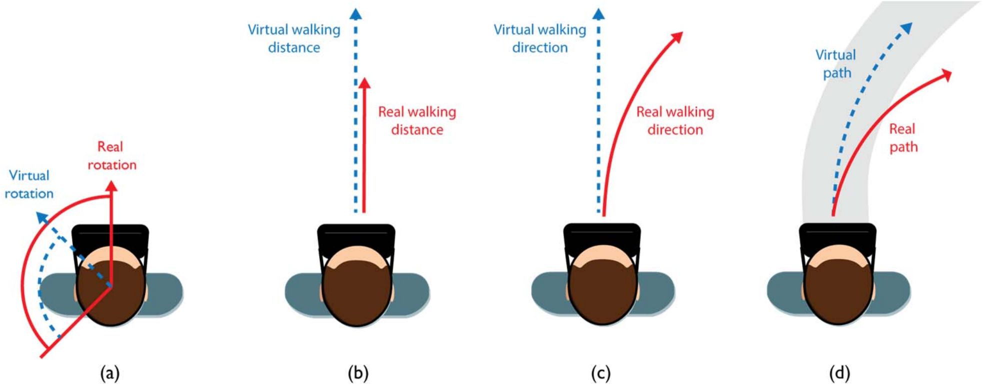

Figure 1. Illustration of four types of gains used to manipulate the mapping between the user's real and virtual movement: (a) rotation gains (user stationary), (b) translation gains (user moving forward), (c) curvature gains (user moving forward), and (d) bending gains (user moving on a curve).

Overt Redirection and Interventions

Even though RDW ideally should remain imperceptible, overt techniques have been proposed as a means condensing even larger VEs into the tracking space. The technique *seven league boots*[5] was intentionally designed as overt because it was based on the user enabling and disabling perceptible translation gains by pressing a button on a hand-held wand. While this form of redirection is rare, it is not uncommon to rely on overt techniques as a last resort when the user reaches the bounds of the tracking space. In these worst-case conditions, for the safety of the user and the equipment, the system must intervene and trigger execution of an overt technique to guide the user back toward the center of the tracking space without stopping the virtual experience. Williams et al.[6] introduced three overt techniques aimed at doing so: (1) The *freeze–backup technique* where the virtual experience is frozen, the experimenter guides the user to the center of the tracked space, and the virtual experience is resumed; (2) the *freeze–turn technique* where the display system is frozen, the user physically turns toward the center of the tracked space, and the virtual experience is resumed; (3) the *2:1 turn* where the user is instructed to stop and physically turn while the VE rotates at twice her speed. The user physically turns 180° and is rotated 360° in the virtual world. As a result the user no longer faces the bounds of the tracking space and is able to continue along the same virtual path.

Peck, Fuchs, and Whitton[7] describe distractors—an object in the VE for the user to visually focus on while the VE rotates during the user's head turns. A red ball (serving as the distractor) was used to guide the user's eyes during rotation, as well as distract the user from the rotation of the VE. Results suggest that distractors are preferred by participants and produce higher levels of subjective presence when compared to other overt reorientation techniques. The authors compared three types of distractors: the sight of a hummingbird, the sound of the bird, and the sound and sight of the bird. The results did not reveal whether the addition of sound decreased the likelihood of noticing the applied gain, but they did suggest that improving visual realism and adding sound positively influenced the participants' sense of presence in the VE.

It has been proposed that instead of overtly intervening it may be possible to capture the user's attention using events that are plausible within the narrative and use those events as opportunities to imperceptively apply translation, rotation, and curvature gains.[8] For example, Neth et al.[9] used agents walking in front of users to slow them down, and agents intersecting the users' path to cause changes in their heading.

(a) Change blindness redirection

(b) Impossible spaces

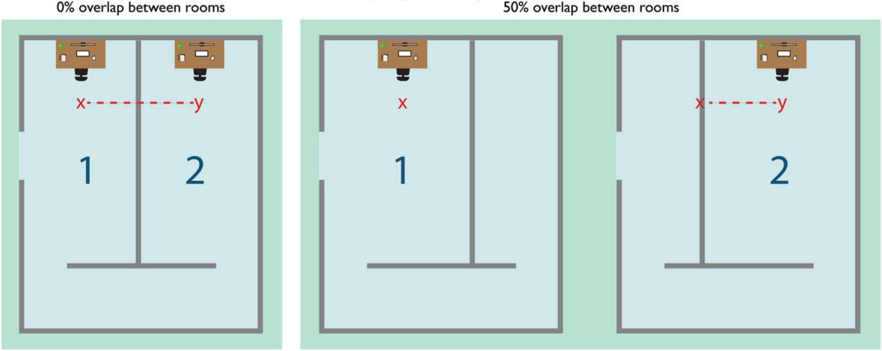

Figure 2: Examples of two forms of redirection through manipulation of the virtual architecture: (a) change blindness redirection used to present two virtual rooms in the same tracking space, and (b) impossible spaces used to present two overlapping rooms.

Subtle Manipulation of Virtual Architecture

Suma et al. proposed a radically different approach to redirection: they subtly manipulate the architecture of the VE, taking advantage of *change blindness*, i.e., an individual's inability to detect changes in the environment. Suma et al. were able to change the location of doorways and corridors behind users' backs and thereby manipulate their walking paths (Figure 2a). Impressively, only one of 77 participants in two user studies noticed the manipulation. This form of architectural manipulation enabled users to navigate a large dynamic virtual office building of approximately 219 sq. meters without leaving a tracking space of 4.3m × 4.3m. [10]

Suma et al. describe *Impossible Spaces*, another technique relying on architectural manipulation.[11] Impossible Spaces makes it possible to compress (larger) virtual interior environments into smaller physical spaces by means of self-overlapping architecture (Figure 2b). Two user studies revealed that relatively small virtual rooms can overlap by up to 56% without the user's knowledge, and larger virtual rooms mapped to a physical space of 9.14m × 9.14m may overlap by up to 31%. If the aim is not to replicate the spatial layout of a real space, then the technique called *Flexible Spaces* can provide unrestricted walking within a dynamically generated interior VE.[12] Work by Vasylevska and Kaufmann[13] demonstrated that it is possible to avoid detection and increase the amount of overlap between virtual rooms by connecting these rooms using longer corridors with more corners, and smoothly curved corridors may be more beneficial for spatial compression. Even though this work is encouraging, an inherent drawback of approaches relying on architectural manipulation is that they are limited to interior environments.[10]

REDIRECTION CONTROLLERS

Free exploration of virtual spaces with any one of the techniques described above is nearly impossible. Developers use redirection controllers to manage the application of redirection techniques based on user positions in both the real and the virtual spaces. Software libraries such as the Redirected Walking Toolkit[26] offer implementations of basic redirection techniques; redirection controllers manage redirection and repositioning at a higher level.

Scripted controllers steer users through a space along a physical and virtual path that is largely pre-determined by the developer.[2] An advantage of scripted controllers is that developers predefine a user's path and the redirection technique to be used. This ensures that the virtual experience fits within a desired physical space. Such controllers are not easily reusable across different VEs, can require extensive customization for each simulation, and can break down if users leave the predefined path. Controllers that rely on change blindness (e.g., moving doors that are out of sight),[10] or impossible overlapping spaces connected by planned corridors[11] can be considered scripted controllers.

Generalized controllers are reactive in nature, and use subtle continuous techniques to guide users toward a certain physical place or pattern regardless of a user's intended virtual travel. No knowledge of the VE or pre-scripting of user movements is required, making these controllers more flexible, but also less optimized. Razzaque[15] originally described three generalized algorithms: *steer-to-center*, which always guides users to the center point of the tracking space; *steer-to-multiple-waypoints*, which guides users to a set of points; and *steer-to-orbit*, which guides users toward and along a circular path around the center. Other variations of these methods have been proposed, including nested orbits or figure-eight patterns.[27] For a review and comparison of generalized algorithms, see the work of Hodgson and colleageus.[28]

Predictive controllers analyze the physical and/or virtual environments to determine where a user can go, cannot go, or is likely to go, then uses that information to guide redirection. For example, a user walking along a virtual hallway with no intersections will likely walk straight ahead. Predictive controllers may also consider the characteristics of typical human walking to make informed steering choices.[29] Zmuda et al.'s[30] FORCE algorithm uses subtle continuous techniques combined with changes in redirection strategy at key decision points. The VE is pre-processed into a series of navigable paths and decision points, and the controller calculates optimal physical locations to steer the user toward. Similar to scripted controllers, predictive controllers require pre-processing and cataloging available paths for each VE, or real-time estimation of likely paths. More recent approaches, such as that proposed by Zank and Kunz,[31] simplify this process by automatically calculating navigable paths and decision points for a given VE so that developers need not specify them manually.

Resetting controllers are redirection techniques that are most often used to reorient users away from a nearby tracking space boundary. Examples are the overt 2:1 turn and freeze-turn techniques, and visual distractors. Resetting controllers are effective in any size tracking space, but cause increasingly frequent interruptions as the size of the tracking space diminishes.

Note that these controller categories are not mutually-exclusive. Generalized controllers can use resetting to catch failure cases, e.g., *Redirected Free Exploration with Distractors*.[7] Predictive controllers can use generalized algorithms and change strategies as needed. For example, the work on bending gains described in the sidebar combines scripted and generalized approaches enabling users free choice among branching VE paths, while walking along connected, scripted arc segments to form a predefined pattern in the physical world. Different applications may lend themselves toward different combinations of controllers and techniques. Similarly, Nescher et al.[32] have presented an algorithm for dynamically choosing and weighting different redirection controllers in order to optimize space and minimize costs. Regardless of controller type, the RDW literature has focused almost exclusively on guiding a single immersed user, but recent work has sought to address the challenge of simultaneously redirecting more than one user (see sidebar "Redirection of Multiple Users").

UNWANTED SIDE EFFECTS OF REDIRECTED WALKING

RDW techniques intentionally create subtle or overt discrepancies between the spatial senses, distort the perception-action mappings during movements in a VE, or dynamically rearrange the geometry and affordances of a virtual world. This has the potential to (a) introduce simulator sickness, (b) interfere with spatial learning and memory, and (c) lead to higher cognitive load than walking in the real world. For RDW to be widely adopted in VR, we must document these effects, understand their causes, and develop mitigation strategies.

Simulator Sickness: When first introducing RDW to the scientific community Razzaque et al.[2] described a pilot user study suggesting that redirection could be performed without notable simulator sickness. Since then several studies exploring detection thresholds have also involved assessment of simulator sickness based on questionnaires administered before and after the user's VE experience (see work cited in the sidebar on estimation of detection thresholds). The results of these self-reports generally indicate an increase in simulator sickness, albeit to varying degrees. Because participants were exposed to both noticeable and imperceptible gains for extended periods of time, it is impossible to conclude, based on these data alone, whether subtle redirection elicits simulator sickness. Similarly, work on architectural manipulation tends to find increases in simulator sickness when comparing scores before and after studies. Despite the frequent assessment of simulator sickness, it is currently not possible to conclude to what the degree of sickness varies across populations, across different gains and across methods relying on manipulation of gains or virtual architecture. The work exploring masking redirection using inter-stimulus images and optic flow effects did not reveal a difference in simulator sickness reported before and after exposure to the VE;[3] there is some indication that redirection using both sound and visuals may elicit stronger simulator sickness than purely acoustic redirection; and the addition of passive haptic cues might reduce sickness during exposure to curvature gains (see work cited in the sidebar on estimation of detection thresholds).

Spatial Performance: Relatively little empirical data exists on the effects of redirection techniques on spatial cognition and performance. Hodgson et al.[36] examined the effect of different subtle continuous reorientation techniques on spatial learning and memory. A crucial finding of this work is the lack of any negative impact of RDW on spatial memory. In no case did they find performance decrements for landmark learning while being redirected. Indeed, there were small numerical trends in all dependent measures toward better performance for RDW over normal navigation where no redirection was applied. Similar effects have been shown by Peck et al.[7] finding that participants experiencing subtle reorientation with occasional overt reorientation (resetting) perform significantly better than participants using a non-walking locomotion interfaces.

These findings stand in contrast to a related study by Williams et al.[6] Their work examined the impact of overt reorientation or repositioning techniques on spatial memory compared to no reorientation. Their overt interventions did increase angular errors and response latencies in a spatial pointing task. It appears that the strong sensory-motor conflicts that are created during overt interventions may not be resolved at a low, perceptual level and can interfere with spatial learning or introduce distortions into a user's memory. Work by Suma et al.[10] indicates that even when redirection is performed by dynamically changing the layout of the environment, then participants' mental maps resembled the static layouts the authors intended them to perceive. Nevertheless, recent work has shown that variations in the complexity of paths joining overlapping rooms may influence distance perception within these rooms.[13]

Cognitive Load: Human working memory draws from finite cognitive resources, including verbal and spatial resources. Marsh et al.[37] evaluated different types of semi-natural locomotion user interfaces in VR and found that locomotion with less natural interfaces increased spatial working memory demands. Bruder et al.[38] analyzed the effects of a subtle continuous reorientation technique based on curvature gains on cognitive load. They employed a dual tasking method to determine the mutual influence between a locomotion task using RDW and a concurrent task that draws from either verbal or spatial cognitive resources. The results showed that the radius of the circular path on which participants were redirected with curvature gains had a significant effect on verbal and spatial working memory. The smaller the circle was, the more this technique drew from the participants' finite cognitive resources. However, significant effects on cognitive load were only shown when participants were clearly able to detect the manipulation. In cases where users have to perform a complex cognitive task in the virtual world, the authors recommend using curvature gain manipulations only when the physical space available is more than 10m × 10m or combining multiple subtle redirection techniques if the physical space is smaller.

CURRENT CHALLENGES AND FUTURE WORK

The ideal RDW technique should be imperceptible, safe, generalizable, and devoid of unwanted side effects. Perhaps not surprisingly, these criteria align closely with the broader themes and research questions addressed by the scientific community since redirected walking was first introduced. This section outlines central lessons learned, current challenges, and potential directions for future work.

Subtlety: A recurring theme of research on RDW has been identification of detection thresholds for rotation gains. This work has provided ample evidence that the mapping between users' real and virtual spatial movement can be manipulated while avoiding detection. However, our understanding of the perceptual limits is incomplete, suggesting the need for future work. Much of the work on detection thresholds was performed using displays that do not compare to current generation HMDs, and threshold estimates vary significantly depending on estimation methods. This suggests the need for developing reliable estimation methods and determining whether previously identified thresholds apply to current HMDs. Additionally, each user has different perceptual thresholds and tolerances. However, currently we do not know the distribution of sensitivity across the user population (or sub-populations). Such knowledge could enable higher thresholds for perceptually tolerant users, and lower thresholds for sensitive users. Similarly, we do not know enough about the interaction between detection thresholds and task engagement, i.e., the attentional resources devoted to the task. Understanding this relationship would enable dynamic adjustment of gains based on the current task or measures of users' attentional state.

As this review points out, some redirection techniques decrease noticeability by deliberately *masking* the manipulation by using visual distractors during resetting, or inter-stimulus images and optic flow effects, or by performing redirection during naturally occurring masking caused by saccades and eye blinks or during narrative events. Identifying effective ways of masking gains is an important direction for future research. Examples of unexplored approaches include masking gains through manipulation of the display field of view or by affecting optic flow with virtual snow or other particle effects.

Redirection through manipulation of the VE is a relatively new, yet promising, approach to RDW. So far, work has focused exclusively on architectural models of indoor environments. Nevertheless, it seems possible that the layout of outdoor VEs could be manipulated to similar effects, e.g., relocation of virtual structures and walkways outside the user's field-of-view. Additionally, utilizing these illusions often requires the manual design and modeling of a custom-built VE. This points to the need for research into new methods for automatically manipulating the layout of static environment models created by artists or from 3D scans of a physical environment.

Safety: Redirection techniques relying on gain manipulations do not guarantee safety unless the system intervenes once subtle redirection is no longer an option. This form of resetting is inherently disruptive. Visual distractors partially alleviate this problem, and recent work has proposed techniques that integrate into the narrative of the VE.[8] Such techniques could provide opportunities for subtle redirection before overt interventions become necessary. However, the effectiveness of narrative interventions is not well-documented, nor do we understand trade-offs between overt and narrative interventions.

Recent advances in inside-out tracking and environment mapping technologies are driving a progression towards immersive experiences at “building-scale.” Unbounded by the limitations of a pre-defined motion tracking volume, virtual reality seems poised to graduate from controlled lab spaces and move out into the world. While this presents a compelling vision, the unpredictability of operating in an ad hoc environment introduces new challenges for user safety. RDW may be a promising approach for locomotion in cluttered physical spaces and avoiding collisions with real world obstacles such as furniture.

Generalizability: The RDW controllers proposed to date are suboptimal in one or more ways: (1) scripted controllers are effective but the user must follow a designed path. (2) Generalized controllers offer greater user freedom but resetting is often necessary. (3) Predictive controllers require off-line pre-processing to generate possible paths and knowledge of typical human behavior to inform redirection when the user encounters branches between paths. Thus, automatic calculation of navigable paths and decision points from a given VE and prediction of user’s future paths remain major challenges. With a few recent exceptions, predictive controllers have relied on information about the user’s current walking behavior. However, future work might benefit from exploring real-time capture of biometrics, e.g., eye-tracking and brain-computer interfaces, for path and target prediction.

The challenges outlined above are even greater if the redirection controller needs to handle spontaneous alternation behavior and manipulate the paths of multiple users, or if the user is relying on consumer grade equipment with small active tracking spaces. The latter case may require novel combinations of RDW, relocation techniques, and other walking interfaces, such as walking-in-place techniques.

Unwanted side effects: Even if applied subtly, RDW requires the user to actively (consciously or subconsciously) compensate for the introduced manipulations. Further evaluation of the cognitive effects concerning cognitive load and spatial reasoning of RDW should be evaluated. Moreover, a common assumption is that subtle redirection can be deployed without great risk of simulator sickness. As pointed out by this review, the research community needs to bolster this assumption. We lack a complete understanding of what makes users sick, and we do not know what the population sensitivities are to those factors. Answering these questions are crucially important and a sizable challenge. More generally, research on the effects of RDW has commonly relied on fairly simple paradigms for studying spatial performance and cognitive resource demands. Thus, future research should explore some of the more practically important effects of RDW applied to actual use cases.

Despite the unanswered questions and challenges outlined above, we view the existing body of work on RDW as a testament to the potential of this approach. So far, however, most knowledge in the field is held by a small number of researchers. Considering the strong interest of VR developer communities all over the world, it will be important to bridge the gap between research prototypes and VR developer projects. It is important for the community to start focusing on more applied research while maintaining the tradition of innovation, exploration, and solid basic research that has characterized the past 15 years of work on RDW.

SIDEBAR: DETECTION THRESHOLDS FOR REDIRECTED WALKING

Researchers have adopted psychophysical methods for identifying detection thresholds for visual translation,[14] rotation,[15] curvature,[16] and bending gains.[17] Recent work has explored the effects of multi-sensory stimuli using similar methods.[18–19]

Translation gains: Work by Steinicke et al. suggests that imperceptible translation gains range from downscaling by 14% to upscaling by 26%.[16] If these thresholds apply to current commercially available head-mounted displays, then it should be possible to map a VE of about 5.8m × 5.8m onto the 4.6m × 4.6m tracking space offered by the first generation HTC Vive. Thus, if applied as the only form of redirection in a small tracking space, the benefits of translation gains are limited. Moreover, there is some indication that users unconsciously compensate for applied translation gains by altering their walking speed.[20]

Rotation gains: Razzaque's[15] early work demonstrated that small rotational changes are less likely to be detected, and that the rotation is less likely to be noticed if the user's head is already turning. Since then, several studies have sought to identify detection thresholds for rotation gains during head turns and full body turns. Work by Jerald et al.[21] indicates that users are less likely to notice gains applied in the same direction as head rotation rather than against head rotation. Specifically, users can be physically turned approximately 11.2% more and 5.2% less than the virtual rotation. Bruder et al.[22] evaluated full-body turns and found users can be physically turned approximately 30.9% more and 16.2% less than the virtual rotation.

Moreover, this study indicated that users are less likely to detect the manipulation for smaller virtual rotation angles. Similarly, Steinicke et al.[16] found that users can physically turn approximately 49% more and 20% less than virtual rotation.

Curvature gains: Experimental results vary with respect to the size of the space needed to imperceptibly redirect walkers using curvature gains. It has been suggested that users can be redirected in VEs along the circumference of circles with radiuses of 22m,[16] 11.6m,[23] or as small 6.4m.[23] It has been proposed that methodological differences may account for the variation in results, suggesting the need for further work to better define the limits of human perception.[23] Moreover, research suggests that users are less likely to notice curvature gains when walking at slower speeds,[9] and it has been proposed that curvature and translation gains can be applied simultaneously.[23]

Bending gains: In regard to bending gains, Langbehn et al.[17] found that a virtual curvature can be bent up to 4.4 times its radius in the real world. Moreover, the authors present an application of bending gains allowing users to cover a virtual space of approximately 25m × 25m within a physical space of about 4m × 4m.[24] However, in order to condense such a large VE into this small tracking space, the walker's movement is constrained to specific predefined paths.

Other Sensory Approaches to Redirection: A few studies have explored the hypothesis that the addition of haptic and auditory stimuli matching the visuals might allow unnoticeable manipulation while using higher gains. Matsumoto et al.[25] describe a subtle continuous redirection approach based on visio-haptic interaction. Users walked along a straight virtual path while touching a collocated real and virtual wall where the virtual wall was mapped onto a physical, convex surface. Through application of curvature gains users walked along a potentially infinite virtual wall while in fact they were walking in circles. Study results indicated that passive haptic feedback may increase the likelihood that users feel as if they are walking straight during the application of curvature gains.

Recent work explored whether users' ability to detect visual rotation gains is affected by the addition of spatialized sound when the location of the sound is also manipulated by the same gains.[19] The benefit of adding spatial sound has not been conclusively shown, possibly due to methodological differences as well as variations in the visual and auditory stimuli. For example, it seems possible that spatial sound is more likely to have an effect during scenarios involving richer soundscapes or scenarios where audition plays a more central role for spatial cognition (e.g., foggy or dimly lit environments).[19]

SIDEBAR: REDIRECTION OF MULTIPLE USERS

Multi-user redirection controllers guide two or more users away from each other and any obstacles in the same tracking space. A few studies have used simulations to explore this issue,[33–34] but no user studies have yet been conducted.

In its simplest form, a multi-user redirection controller can implement Razzaque's[15] steer-to-center algorithm, with each user having a different 'center' position, thus directing users away from each other and toward their own sub-areas of the tracking space. More sophisticated controllers are needed, however, to maximize the tracking space for each participant.

Multi-user redirection introduces a number of new issues. Primary among them is collision prediction and avoidance for dynamic and sometimes unpredictable objects, including the users themselves. Users can collide in different ways (head-on, rear-ending a slow user, crossing or converging paths, etc.), or users may simply pass closely without colliding. If a collision is predicted, RDW can resolve it by curving the users away from each other while, ideally, not steering them toward a wall.[33] Collisions can also be prevented by using translation gains to modulate a user's progress toward a potential collision point or help another reach a turning point more quickly. Impending collisions can also be prevented with resetting methods.

Accurately predicting collisions can be difficult due to users' ability to freely choose their path through the VE. It can also become computationally expensive with additional users: redirecting one user to prevent a collision could guide them into the path of another. Thus, all users' predicted paths and their various interactions must be modeled to find the optimal space in which to steer each user.

It is also important to consider social aspects of a multi-user experience. Because the positions in the virtual and physical worlds diverge as a result of RDW, there are likely to be scenarios in which two users are next to each other in the virtual world but far apart in the real world. Attempting a handshake or conversation will fail, as the users are not actually co-located. Conversely, users can be far apart virtually but next to each other physically. Much of this can be mitigated by using voice-chat over noise-canceling headphones (see [35]).

REFERENCES

1. E.A. Suma et al., "A taxonomy for deploying redirection techniques in immersive virtual environments," *Proceedings of IEEE Virtual Reality Short Papers and Posters* (VRW 12), 2012, pp. 43–46.
2. S. Razzaque, Z. Kohn, and M.C. Whitton, "Redirected walking," *Proceedings of Eurographics*, 2001, pp. 105–106.
3. G. Bruder et al., "Tuning self-motion perception in virtual reality with visual illusions," *IEEE Transactions on Visualization and Computer Graphics*, vol. 18, no. 7, 2012, pp. 1068–1078.
4. B. Bolte and M. Lappe, "Subliminal reorientation and repositioning in immersive virtual environments using saccadic suppression," *IEEE Transactions on Visualization and Computer Graphics*, vol. 21, no. 4, 2015, pp. 545–552.
5. V. Interrante, B. Ries, and L. Anderson, "Seven league boots: A new metaphor for augmented locomotion through moderately large scale immersive virtual environments," *Proceedings of the IEEE Symposium on 3D User Interfaces* (3DUI 07), 2007, pp. 167–170.
6. B. Williams et al., "Exploring large virtual environments with an HMD when physical space is limited," *Proceedings of the ACM Symposium on Applied Perception in Graphics and Visualization* (APGV 07), 2007, pp. 41–48.
7. T.C. Peck, H. Fuchs, and M.C. Whitton, "An evaluation of navigational ability comparing redirected free exploration with distractors to walking-in-place and joystick locomotion interfaces," *Proceedings of the 2011 IEEE Conference on Virtual Reality* (VR 11), 2011, pp. 55–62.
8. T. Grechkin et al., "Towards context-sensitive reorientation for real walking in virtual reality," *Proceedings of the 2015 IEEE Conference on Virtual Reality* (VR 15), 2015, pp. 185–186.
9. C.T. Neth et al., "Velocity-dependent dynamic curvature gain for redirected walking," *IEEE Transactions on Visualization and Computer Graphics*, vol. 18, no. 7, 2012, pp. 1041–1052.
10. E.A. Suma et al., "Leveraging change blindness for redirection in virtual environments," *Proceedings of the 2011 IEEE Conference on Virtual Reality* (VR 11), 2011, pp. 159–166.
11. E.A. Suma et al., "Impossible spaces: Maximizing natural walking in virtual environments with self-overlapping architecture," *IEEE Transactions on Visualization and Computer Graphics*, vol. 18, no. 4, 2012, pp. 555–564.
12. K. Vasylevska et al., "Flexible spaces: Dynamic layout generation for infinite walking in virtual environments," *Proceedings of the IEEE Symposium on 3D User Interfaces* (3DUI 13), 2013, pp. 39–42.

13. K. Vasylevska and H. Kaufmann, “Compressing VR: Fitting large virtual environments within limited physical space,” *IEEE Computer Graphics and Applications*, vol. 37, no. 5, 2017, pp. 85–91.
14. B. Williams et al., “Updating orientation in large virtual environments using scaled translational gain,” *Proceedings of the ACM Symposium on Applied Perception in Graphics and Visualization* (APGV 06), 2006, pp. 21–28.
15. S. Razzaque, “Redirected Walking,” thesis PhD Thesis AAI3190299, The University of North Carolina at Chapel Hill, 2005.
16. F. Steinicke et al., “Estimation of detection thresholds for redirected walking techniques,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 16, no. 1, 2010, pp. 17–27.
17. E. Langbehn et al., “Bending the curve: Sensitivity to bending of curved paths and application in room-scale VR,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 23, no. 4, 2017, pp. 1389–1398.
18. F. Meyer, M. Nogalski, and W. Fohl, “Detection thresholds in audio-visual redirected walking,” *Proceedings of the Sound and Music Computing Conference* (SMC 16), 2016, pp. 1–7.
19. N.C. Nilsson et al., “Estimation of detection thresholds for audiovisual rotation gains,” *Proceedings of the 2016 IEEE Conference on Virtual Reality* (VR 16), 2016, pp. 241–242.
20. R. Zhang, B. Li, and S.A. Kuhl, “Human sensitivity to dynamic translational gains in head-mounted displays,” *Proceedings of the ACM Symposium on Spatial User Interaction* (SUI 14), 2014, pp. 62–65.
21. J. Jerald et al., “Sensitivity to scene motion for phases of head yaws,” *Proceedings of the ACM Symposium on Applied Perception in Graphics and Visualization* (APGV 08), 2008, pp. 155–162.
22. G. Bruder et al., “Reorientation during body turns,” *Proceedings of the Joint Virtual Reality Conference of EGVE - ICAT - EuroVR*, 2009, pp. 145–152.
23. T. Grechkin et al., “Revisiting detection thresholds for redirected walking: combining translation and curvature gains,” *Proceedings of the ACM Symposium on Applied Perception* (SAP 16), 2016, pp. 113–120.
24. E. Langbehn et al., “Application of redirected walking in room-scale VR,” *Proceedings of the 2017 IEEE Conference on Virtual Reality* (VR 17), 2017, pp. 449–450.
25. K. Matsumoto et al., “Curvature manipulation techniques in redirection using haptic cues,” *Proceedings of the IEEE Symposium on 3D User Interfaces* (3DUI 16), 2016, pp. 105–108.
26. M. Azmandian et al., “The redirected walking toolkit: A unified development and deployment platform for exploring large virtual environments,” *2016 IEEE 2nd Workshop on Everyday VR* (WEVR 16), 2016, pp. 1–6.
27. T. Field and P. Vamplew, “Generalised algorithms for redirected walking in virtual environments,” *Proceedings of the International Conference on Artificial Intelligence in Science and Technology* (AISAT 04), 2004, pp. 58–63.
28. E. Hodgson, E. Bachmann, and T. Thrash, “Performance of redirected walking algorithms in a constrained virtual world,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 20, no. 4, 2014, pp. 579–587.
29. M. Zank and A. Kunz, “Where are you going? Using human locomotion models for target estimation,” *The Visual Computer*, vol. 32, no. 10, 2016, pp. 1323–1335.
30. M.A. Zmuda et al., “Optimizing constrained-environment redirected walking instructions using search techniques,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 19, no. 11, 2013, pp. 1872–1884.
31. M. Zank and A. Kunz, “Optimized graph extraction and locomotion prediction for redirected walking,” *Proceedings of the IEEE Symposium on 3D User Interfaces* (3DUI 17), 2017, pp. 120–129.
32. T. Nescher, Y.-Y. Huang, and A. Kunz, “Planning redirection techniques for optimal free walking experience using model predictive control,” *Proceedings of the IEEE Symposium on 3D User Interfaces* (3DUI 14), 2014, pp. 111–118.
33. M. Azmandian, T. Grechkin, and E.S. Rosenberg, “An evaluation of strategies for two-user redirected walking in shared physical spaces,” *Proceedings of the 2017 IEEE Conference on Virtual Reality* (VR 17), 2017, pp. 91–98.

34. E.R. Bachmann et al., “Collision prediction and prevention in a simultaneous two-user immersive virtual environment,” *Proceedings of the 2013 IEEE Conference on Virtual Reality* (VR 13), 2013, pp. 89–90.
35. M. Zank, T. Nescher, and A. Kunz, “Robust prediction of auditory step feedback for forward walking,” *Proceedings of the ACM Symposium on Virtual Reality Software and Technology* (VRST 13), 2013, pp. 119–122.
36. E. Hodgson, E. Bachmann, and D. Waller, “Redirected walking to explore virtual environments: Assessing the potential for spatial interference,” *ACM Transactions on Applied Perception*, vol. 8, no. 4, 2011, pp. 1–22.
37. W.E. Marsh et al., “Cognitive demands of semi-natural virtual locomotion,” *Presence: Teleoperators and Virtual Environments*, vol. 22, no. 3, 2013, pp. 216–234.
38. G. Bruder, P. Lubos, and F. Steinicke, “Cognitive resource demands of redirected walking,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 21, no. 4, 2015, pp. 539–544.

ABOUT THE AUTHORS

Niels Christian Nilsson is an assistant professor at Aalborg University Copenhagen. He received a Ph.D. and M.Sc. in Medialogy from the same institution in 2016 and 2011, respectively. He is a member of IEEE and his research is broadly focused on locomotion, perception, and cognition in immersive virtual environments. Contact him at ncn@create.aau.dk.

Tabitha Peck is an assistant professor in the Mathematics and Computer Science Department at Davidson College. Her research interests include virtual reality, locomotion interfaces, human perception, and virtual embodiment. Peck has a Ph.D. in computer science from the University of North Carolina at Chapel Hill and is a member of IEEE. Contact her at tapeck@davidson.edu.

Gerd Bruder is a research assistant professor at the Institute for Simulation and Training at the University of Central Florida. He received his Habilitation in Computer Science from the University of Hamburg in 2017, and his Ph.D. in Computer Science from the University of Muenster in 2011. His research is focused on understanding perception and cognition in Virtual and Augmented Reality, in particular stereoscopic displays, human surrogates, and 3D user interfaces. He is a member of ACM, IEEE, GI, and VSS. Contact him at bruder@ucf.edu.

Eric Hodgson is the director of the Smale Interactive Visualization Center at Miami University in Oxford, Ohio. He received his MA and Ph.D. in Cognitive Psychology from Miami in 2005 and 2008. His research is focused on spatial perception in virtual environments, navigation, redirected walking, and applied uses of virtual reality. Contact him at eric.hodgson@miamiOH.edu.

Stefania Serafin is a professor in sound for multimodal environments at Aalborg University Copenhagen. She received a Ph.D. in computer-based music theory and acoustics from Stanford University in 2004, and a Master in Acoustics, computer science and signal processing applied to music from IRCAM (Paris) in 1997. Her main research interests include sound models for interactive systems and multimodal interfaces, and sonic interaction design. Contact her at sts@create.aau.dk.

Evan Suma Rosenberg is a research assistant professor at the University of Southern California Institute for Creative Technologies. His research lies at the intersection of virtual reality and HCI, encompassing immersive technologies, 3D user interfaces, and spatial interaction techniques. He received his Ph.D. in 2010 from the Department of Computer Science at the University of North Carolina at Charlotte. He is a member of ACM and IEEE. Contact him at suma@ict.usc.edu.

Mary Whitton is a Research Professor of computer science at the University of North Carolina at Chapel Hill. For 15 years she has co-led the Effective Virtual Environments research group that investigates what makes virtual environment systems effective--what is

required of the system to create and maintain the VR illusion. The research group developed and evaluated techniques to make VEs more effective for applications such as simulation, training, and rehabilitation. She is a member of ACM, ACM SIGGRAPH, and a senior member of IEEE. Contact her at whitton@cs.unc.edu.

Frank Steinicke is a professor for Human-Computer Interaction at the Department of Informatics at the University of Hamburg. His research is driven by understanding the human perceptual, cognitive and motor abilities and limitations in order to reform the interaction as well as the experience in computer-mediated realities. He is a member of the GI, ACM, IEEE and member of the Steering committee of the ACM SUI Symposium and the GI SIG VR/AR, and currently editor of the IEEE Computer Graphics & Applications Department on Spatial Interfaces. Contact him at steinicke@informatik.uni-hamburg.de.