Multi-Ray Jumping: Comprehensible Group Navigation for Collocated Users in Immersive Virtual Reality

Tim Weissker* Alexander Kulik† Bernd Froehlich‡
Virtual Reality and Visualization Research, Bauhaus-Universität Weimar

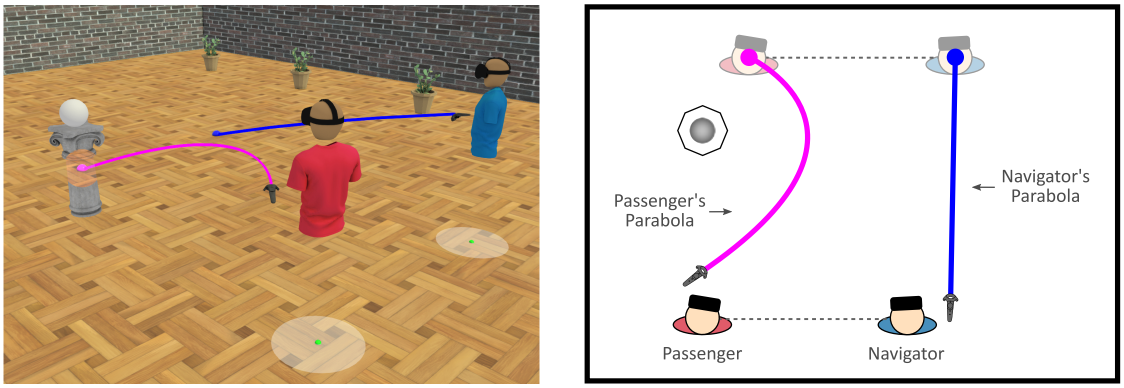

Figure 1: Target specification with *Multi-Ray Jumping* for two collocated users. While the navigator (blue shirt, right side) specifies a target using the blue parabola, the magenta curve adjusts accordingly to show the offset target location of the passenger (red shirt, left side). The initial direction of the magenta curve is defined by the passenger's controller.

Abstract

The collaborative exploration of virtual environments benefits from joint group navigation capabilities. In this paper, we focus on the design and evaluation of a short-range teleportation technique (jumping) for a group of collocated users wearing head-mounted displays. In a pilot study with expert users, we tested three naïve group jumping approaches and derived the requirements for comprehensible group jumping. We propose a novel Multi-Ray Jumping technique to meet these requirements and report results of two formal user studies, one exploring the effects of passive jumping on simulator sickness symptoms ($N = 20$) and a second one investigating the advantages of our novel technique compared to naïve group jumping ($N = 22$). The results indicate that Multi-Ray Jumping decreases spatial confusion for passengers, increases planning accuracy for navigators, and reduces cognitive load for both.

Keywords: Virtual reality, head-mounted displays, multi-user interaction, group navigation, collocation, teleportation, jumping.

Index Terms: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Virtual reality, I.3.6 [Computer Graphics]: Methodology and Techniques—Interaction techniques, H.5.3 [Information Interfaces and Presentation]: Group and Organization Interfaces—Collaborative computing

1 Introduction

Collaborative virtual reality systems enable collocated and distributed groups of people to meet and interact with each other in virtual environments. However, common virtual navigation techniques are not ideal for the joint exploration of large-scale 3D environments since they are designed for individual users and not for groups of users traveling together. In a guided virtual city tour, for example, the attendees might not always want to navigate individually to follow the guide along a path. Instead, the guide could navigate the whole group as a single entity, which leads to an asymmetric role distribution between the controlling *navigator* and multiple *passengers*. Passive locomotion, however, was shown to negatively affect spatial understanding in real-world settings [1, 28] and to increase the risk of simulator sickness in virtual environments [33, 35].

As a first step towards the joint navigation of user groups in immersive virtual reality, we investigated short-range teleportation techniques (jumping) for two collocated users. In a pilot study, six proficient virtual reality users compared three naïve approaches to group jumping. Based on their observations, we derived concrete requirements for *comprehensible* group jumping, that guided our development of *Multi-Ray Jumping*. This novel technique offers extended pre-travel information to help the passenger understand when a jump is planned and where the destination will be. It also facilitates jump planning for the navigator with information about the passenger's destination (Figure 1). *Multi-Ray Jumping* was evaluated in two formal user studies. First, we compared simulator sickness in the passenger role to active jumping as navigator ($N = 20$). Second, we quantified the benefits of the additional pre-travel information for navigator and passenger in comparison to a baseline from the pilot study ($N = 22$).

Workspace awareness as defined by Gutwin and Greenberg [13] emphasizes that an "up-to-the-moment understanding of another person's interaction with the shared workspace" is essential for group interaction. In particular, in situations when users are collocated in the real world and represented in the same spatial configuration in the virtual world, this is largely achieved by representing users by avatars. However, for group navigation, more than such an implicit awareness is needed as passengers are directly affected and not just the surrounding workspace. This is particularly important for navigation by jumping since large distances can be covered by a single jump. Passengers need to be able to anticipate the jump and comprehend where they are headed – potentially being able to mentally take the perspective of the target location. The navigator, on the other hand, must be aware of the passengers' locations and

*e-mail: tim.weissker@uni-weimar.de
†e-mail: kulik@uni-weimar.de
‡e-mail: bernd.froehlich@uni-weimar.de

understand how they will be affected by the planned jump. The research presented in this paper explores which mediations are necessary in this regard and how they can be realized for two collocated users. More specifically, we make the following contributions:

- the concrete requirements for *comprehensible* group jumping elicited by a pilot study of three initial approaches for joint jumping of two collocated users
- the design and implementation of *Multi-Ray Jumping*, a novel technique for comprehensible group jumping
- indications that simulator sickness symptoms are not affected by active or passive user roles during jumping
- evidence that *Multi-Ray Jumping* increases planning accuracy for the navigator and improves spatial awareness for the passenger while imposing lower cognitive load in both roles

A core lesson learned from our studies was that group navigation requires a robust understanding of all navigation-related activities and their consequences among all participants. Based on our results, we believe that *Multi-Ray Jumping* provides this understanding and is a promising first step towards comprehensible, effective, and comfortable group navigation in collaborative virtual reality.

2 Related Work

Navigation is an interplay of the motor component *travel* and the cognitive component *wayfinding* [6]. While physical movement is the most natural form of travel in virtual reality, it is usually limited by the size of the available tracking area. Walking in place [36], redirected walking for one [26] and two users [2], scaling [18], and resetting techniques [39] can help to overcome this limitation, but they become impractical and exhausting in large environments. As a result, virtual navigation techniques like *steering* and *target-based travel* move the user in the virtual environment without requiring physical locomotion.

In contrast to free exploration techniques without navigational constraints, prior research presented steering techniques guiding a user along interesting paths and features of the scene to explore. In the river analogy, for example, users automatically follow a predefined path while having active control over small deviations to investigate nearby features [12]. Additionally, automated guide avatars can draw user attention to previously specified points of interest [38]. In our research, navigation and guidance are provided by a human user rather than by the system.

In head-mounted displays (HMDs), especially, steering techniques have a high risk of inducing simulator sickness. One plausible reason for this is the sensory conflict between the visual and vestibular systems of a user [22, 27]. Fernandes and Feiner showed that field-of-view restrictions during movements can reduce these symptoms [11]. On the other hand, target-based travel by teleportation avoids the sensory conflict completely, and short-range teleportation (jumping) has become the de-facto standard for navigation in VR applications using HMDs. In the taxonomy of Bowman et al. [6], jumping can be described as a discrete selection of environmental/direct targets. Although jumping offers less spatial information for path integration, studies comparing jumping to steering confirmed effective navigation with significantly lower simulator sickness while spatial orientation and perceived presence did not seem to be affected [14, 37]. As a result, we focused on jumping techniques for effective and comfortable group navigation of HMD users. We used the classification scheme of Weissker et al. to describe our implementations in terms of target specification, pre-travel information, transition, and post-travel feedback [37].

Collaborative virtual reality systems allow multiple users to meet and interact with each other in real-time while exploring a shared 3D environment. This collaboration can be *collocated* for users in the same physical location and *distributed* between different locations via a network connection. Successful collaboration builds on mutual awareness and the effective negotiation of common goals. Gutwin and Greenberg introduced the concept of workspace awareness and suggested that it can emerge implicitly in real-world settings from consequential communication (perception of each other's activities), feedthrough (feedback of manipulated artifacts as perceived by others), and explicit communication. A fundamental basis for workspace awareness in HMD-based systems is the representation of users by avatars in the virtual world. Social meeting rooms like vTime[1] and AltspaceVR[2], for example, offer humanoid avatar representations for distributed users. For dispersed users in social virtual environments, Dodds and Ruddle presented methods to find and follow group members that are out of view [9, 10]. Our research focuses on navigation techniques supporting workspace awareness for groups of collocated rather than distributed users.

Collocated collaboration in virtual environments can be symmetric or asymmetric. In asymmetric settings, one or several users often use auxiliary information, e.g. presented on a desktop monitor, to guide an immersed user through the virtual environment [3, 24, 25]. In symmetric settings, with more than one head-mounted display operated in the same physical location, travel is often limited to physical walking. In the *EPICSAVE* project, for example, two tracked users could learn about and work on medical activities side by side [31]. The system of Roth et al. supported up to five users in a very large tracking space [29]. If multiple collocated users are provided with individual virtual navigation capabilities, the user configuration in the virtual world will diverge from the real-world situation. Lacoche et al. suggested mediators to convey the real-world positions of other users and visual barriers to avoid physical collisions [23]. In contrast, our approach is to retain consistent spatial user configurations between the real and the virtual interaction spaces and to provide techniques for joint navigation. Kulik et al. showed that this may lead to collisions with scene geometry during navigation and presented techniques to prevent those while steering through spatial constrictions [21].

Moving a group of collocated users implies an asymmetric role distribution of one *navigator* and multiple *passengers*. It was suggested to mitigate the resulting imbalance of awareness and control by a large stationary group navigation device, which facilitates the perception and negotiation of navigation control [4, 20, 21]. For jumping techniques, however, this approach appears overly laborious. Prior research also indicated a potentially increased risk of simulator sickness during passive locomotion through virtual environments [33, 35]. Moreover, passive locomotion can negatively affect the formation of spatial knowledge and scene understanding. Appleyard showed that participants who have actively explored a city by driving a car sketched more accurate maps than those who traveled by bus [1]. Chrastil and Warren's overview of studies on active and passive navigation in virtual environments lists examples that revealed similar disadvantages for passive navigation [7]. We aim to develop group navigation techniques that equally support spatial awareness and prevent simulator sickness for users in both roles.

3 Pilot Study: Expert Review of Group Jumping Techniques

In our design process of a group jumping technique for collocated users, we started with an initial expert review of three approaches motivated by related work (see Section 3.2) in a two-user virtual reality setup. Our system allowed both users to walk around in the shared tracking space but restricted the control of virtual navigation to a single user. Individual navigation for the passenger beyond walking was not possible, so the spatial user configuration in the virtual environment remained consistent with the real-world situation. Such a constrained collocated setup facilitates implicit awareness cues and explicit communication similarly to real-world settings.

[1]http://vtime.net/

[2]http://altvr.com/

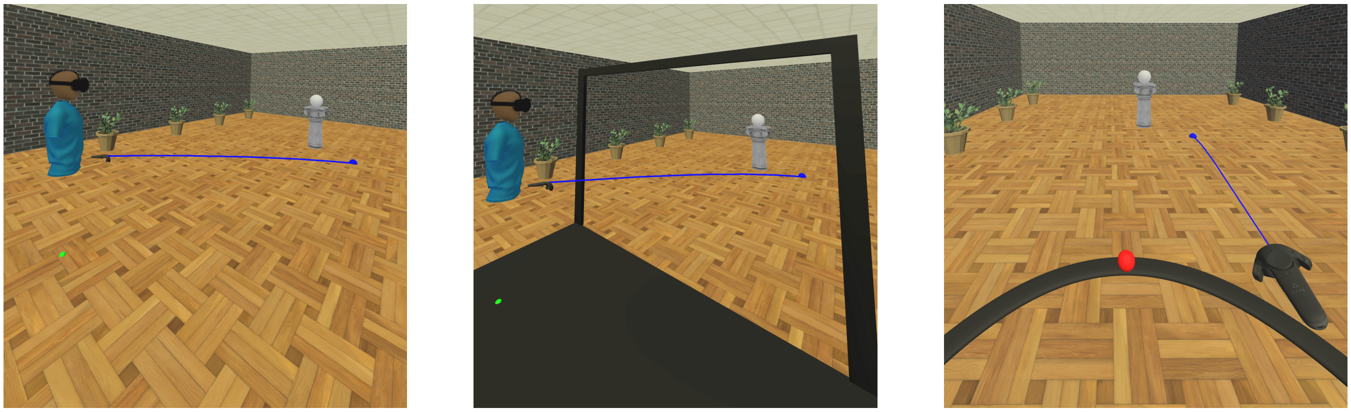

(a) In the *Coupled* condition, the same relative displacement as expressed by the navigator's input was also applied to the passenger's position to retain their spatial configuration.

(b) The *Vehicle* condition constrained *Coupled* jumping to locations that could be selected through a virtual window. Hence, larger changes of movement direction required virtual rotation.

(c) In the *Congruent* condition, the passenger was aligned with the navigator's virtual position for each jump. A compass widget (red sphere) indicated the navigator's viewing direction.

Figure 2: In our exploratory pilot study, we tested two-user implementations of three approaches for realizing group jumping. All screenshots are taken from the passenger's view and the navigator is visualized by a head and body avatar.

In this pilot study, we were interested to see (1) if existing jumping implementations can be directly used in such a setup and (2) if existing group navigation approaches from other systems can be adapted to target-based travel in head-mounted displays.

3.1 Experimental Setup

The VR setup consisted of two HTC Vive head-mounted displays offering both position and orientation tracking. Input was obtained using Vive handheld controllers. Two workstations were linked together and shared the same tracking space calibration. The tracking space was approximately 3m x 1.5m in size, allowing for a side-by-side user configuration in social space [15]. The virtual environment was rendered with a resolution of 1080x1200 pixels per eye and an update rate of 90 Hz. We measured an end-to-end latency of 12.5ms.

3.2 Conditions

Target specification for single-user jumping involves an egocentric selection of a target in the currently visible part of the scene (vista space), often using a parabolic pick ray (e.g. [5, 14, 37]). The ray and its intersection with the scene (pre-travel information) defines the navigator's position after the transition without applying changes to the viewing orientation. Post-travel feedback is usually not given in related work [37]. We evaluated three extensions of this basic single-user jumping for two collocated users with instant transitions called *Coupled*, *Vehicle*, and *Congruent* (see Figure 2). Our chosen avatar representation consisted of a wooden head with a head-mounted display, a shirt and a controller. We found this abstract representation suitable to support mutual awareness by providing more visual saliency than the representation of devices alone while not evoking uncanny feelings as known from imperfectly behaving avatars [30].

3.2.1 Coupled

The most straightforward extension of single-user jumping for two users is to apply the navigator's relative change of position to the passenger as well. Hence, the group's spatial user configuration remains identical to the real world during travel. We refer to this first condition of our pilot study as *Coupled* group jumping (Figure 2(a)) and use it to analyze the direct applicability of a single-user jumping technique to a two-user scenario.

3.2.2 Vehicle

Prior group navigation techniques for projection-based multi-user displays consider all users in front of the projection screen traveling together on a shared vehicle or viewing platform [4, 20, 21]. Most of the vehicle's movements are applied in the directions defined by the visible parts of the scene through the projection screen. Larger changes of movement direction, therefore, require virtual rotation methods. This offers the advantage that all users share a similar viewing direction, which is not necessarily true in collocated HMD setups as the screens rotate together with the user. In our *Vehicle* condition, we analyze the applicability of the vehicle metaphor to jumping in head-mounted displays. We adapt the navigational constraints from projection-based setups and indicate the possible movement directions by a virtual navigation window (Figure 2(b)). Left and right rotations at constant angular velocity were initiated and terminated by button presses on the controller.

3.2.3 Congruent

Passengers are always offset to the navigator, which can lead to virtual collisions with the scene geometry in confined environments [21]. This situation is uncomfortable and not considered in many studies on passive navigation summarized by Chrastil and Warren [7]. Here, passive navigation means watching motion recordings from the navigator's point of view. In our *Congruent* condition, we analyze how well the concept of seeing the navigator's view during travel is received in a two-user HMD scenario. Our implementation allowed the navigator to initiate navigation by pressing a button, which triggered a slow-in-slow-out animation from the passenger's actual position to the navigator with a duration of 0.5s. This allowed the passenger to see the controller and jumping ray from the navigator's perspective. A red sphere on a ring surrounding the passenger indicated the current viewing direction of the navigator for improved mutual awareness. When the navigator pressed the button again, the passenger was moved back to their tracked position.

3.3 Procedure

Participants arrived at our lab and signed an informed consent form. Afterwards, the three techniques motivated in Section 3.2 were tested in a counterbalanced within-subjects design. The role of the navigator was taken by the experimenter for comparable interaction sequences across all participants. The participants experienced the role of the passenger. Each technique was introduced with a short verbal explanation. Next, the experimenter performed nine jumps along a route through a virtual museum, thereby stopping and looking at various exhibits. The task of the passenger was to observe and understand the actions of the navigator as well as the exhibits being looked at. In the *Congruent* condition, participants were released to their actual positions at exhibits and moved back to the navigator before further travel. After testing each technique, users were asked for advantages, disadvantages, and general feedback in an open questionnaire. At the end, participants provided a preference ranking of the three techniques. The whole procedure took approximately 30 minutes to complete.

3.4 Participants

Six (two female and four male) student and research assistants between 21 and 28 years ($M = 24.0$, $\sigma = 2.65$) participated in this explorative study. All of them were proficient users of Virtual Reality and thus able to provide expert feedback.

3.5 Results and Discussion

The *Vehicle* condition was appreciated for being "fast and easy to understand" (P6) and for "jumps [that] cannot happen outside my field of view" (P4). On the other hand, participants complained about traveling through and standing in walls (P1, P5) because of their offset to the navigator. Their main point of criticism was the need for virtual rotation, which was considered "nauseating" (P4). One participant even had to "close the eyes for rotation to not get sick" (P5). Discrete rotations could be considered as an alternative to avoid continuous motion flow. However, it is subject to future research to find virtual rotation techniques that maximize spatial comprehensibility while minimizing simulator sickness. All participants named *Vehicle* as their least preferred technique.

Seeing the navigator's view in the *Congruent* condition was said to feel "almost like you're doing it yourself" (P4), and it was positively mentioned that "orientation was easy when in partner's view" (P3). One participant, on the other hand, noticed that navigation "feels lonely" (P4), which we attribute to the incorrect spatial representation of navigator and passenger in the virtual world. In addition, users mentioned that the transition moving the passenger to and away from the navigator's position "felt really tough" (P6), so a "slower transition animation" (P3, P6) was suggested. However, slower switching results in longer exposure times to visual motion flow, which can intensify simulator sickness symptoms. A jump-based transition, on the other hand, would need suitable feedback to minimize confusion about the immediate location changes. Half of the participants named *Congruent* as their most preferred technique.

The *Coupled* condition was mainly appreciated for its simplicity. Our participants deemed it "easy to understand" (P1, P6), "straightforward" (P2), and "more intuitive" (P3) than the other techniques. The problem of accidentally "standing in walls" (P5) resulting from the offset to the navigator was, nevertheless, also mentioned as a disadvantage here. Also, the inconsistency between indicated target of the navigator and actual landing position of the passenger made it "really difficult [...] to judge the target of the next jump" (P4). This was intensified by "obstacles" (P4) occluding the jumping ray and target. One participant also mentioned it's easily "possible to miss where the partner is pointing" (P3) since the passenger can still be busy looking at an exhibit while the navigator already plans the next jump in a different direction. In total, half of the participants named *Coupled* as their most preferred technique.

3.6 Requirements for Comprehensible Group Jumping

In summary, none of the tested implementations of group jumping was fully satisfying. A major complaint about the straightforward extension of single-user jumping in the *Coupled* condition was the frequently occurring confusion of passengers about their resulting position in the scene after a jump. Apparently, they often expected to arrive at the location indicated by the navigator since this was the only available target preview. This problem was intensified when parts of the navigator's parabola were occluded by the scene geometry. Moreover, our participants reported that they missed the planning phase of several jumps, which resulted in unexpected location changes and required spatial reorientation. They were also often placed into walls as it was difficult for the navigator to estimate the relative position of the participant and to incorporate this information into the planning process. The suggested constraints to the jumping direction (*Vehicle*) or the virtual passenger location (*Congruent*) solved some of these problems but introduced additional overheads and challenges to be solved separately. We thus review the observations in the *Coupled* condition in more detail and derive requirements for *comprehensible* group jumping.

Comprehensible group jumping techniques should foster the awareness of ongoing navigation activities and facilitate the predictability of their consequences for the navigator and all passengers. This implies the following interface requirements for passengers:

- a notification mechanism to raise attention when the navigator is planning a jump
- a clearly visible indication of the jump's target location for all participants (passengers and navigator) in order to make the jump predictable and avoid spatial confusion

For the navigator, the interface requirements for comprehensible group jumping can be summarized as follows:

- an indication of the current user configuration in the workspace to support the awareness of passengers and their agreement on the planned navigation
- a clearly visible indication of the jump's target location for all participants (navigator and passengers) to support a collision-free and precise placement of the group at the target location

4 The Multi-Ray Jumping Technique

Following our postulated requirements, comprehensible group jumping can be implemented in various ways. Most fundamentally, previews of the target locations of all involved participants must be provided. These previews can be simple location markers or semi-transparent copies of the current avatar representations at the planned location (ghost avatars). The latter implies an external view of the current group configuration, which offers additional situational awareness through body language. Passengers standing behind the navigator would become visible, and a jump could be interrupted if one or several avatars still seem to be busy at the current location. For the passengers, such an additional group representation might be less relevant. Instead, since we have noticed that the planning phase of a jump might go fully unnoticed, an always visible indicator about a planned jump may be more helpful.

Our *Multi-Ray Jumping* technique is an extension of *Coupled* group jumping with additional pre-travel information for enhanced comprehensibility. When the navigator plans a jump using this technique, a second parabola from the passenger's controller to the corresponding target position appears. To increase the awareness that the navigator is planning a jump, the secondary parabola starts from the passenger's controller in the respective pointing direction. This results in a curved path with a high chance that a considerable part is always visible. If necessary, the controller can also vibrate to attract attention when the navigator starts planning a jump. In certain situations, the secondary parabola may be occluded by scene geometry such that parts of it and the indicated target are not visible. We propose a see-through effect making scene geometry in front of the curve transparent to avoid these problematic situations. Target specification using *Multi-Ray Jumping* is shown in Figure 1.

Additional navigation rays connected to an input device per user afford equal access to navigation control. This can be useful for explorative scenarios without a clear role assignment of guide and attendee. In these cases, each user could simply take over the role of the navigator by pressing the button for target specification on their controller. If both users claim control over the group's navigation concurrently, the system can resolve the conflict in various ways. We suggest switching to individual target specification rays per user, which support the negotiation of a joint decision. Once this decision is made, one user decides to become the passenger by releasing their button again. The corresponding visualization returns to the secondary parabola introduced above, and the new navigator can jump the group to a specified target. If the passenger does not intervene, the navigator can assume agreement for the next jump.

5 Study 1: Simulator Sickness after Multi-Ray Jumping as Passenger

Jumping techniques were shown to reduce simulator sickness symptoms compared to steering in single-user scenarios [37], but potential differences between active and passive roles during group jumping are largely unexplored. Prior research on steering techniques indicated more severe sickness symptoms for passive over user-initiated steering [33, 35], which was also observed in real-world settings earlier [1, 28]. For our first formal study, we were wondering if similar effects of user control can be observed with *Multi-Ray Jumping*. Therefore, we compared simulator sickness in the passenger role of *Multi-Ray Jumping* to the baseline of active single-user jumping. For this study, we used the same experimental setup as our pilot study on group jumping techniques described in Section 3.1.

5.1 Conditions

In the *Active* condition, we tested a common implementation of jumping with a parabolic pick ray for target specification as an active navigator without any passengers. This single-user baseline is comparable to jumping implementations in related work, providing only the pick ray as pre-travel information, an instant transition, and no post-travel feedback. For the *Passive* condition, both the experimenter and the participant were present in the virtual environment, and the experimenter operated *Multi-Ray Jumping* in the navigator role. This ensured comparable interaction sequences to be observed by the participants in the passenger role. As in our pilot study, both users were represented with simple head and body avatars.

5.2 Procedure

Participants arrived at our lab and signed an informed consent form before testing *Active* and *Passive* exploration in a counterbalanced within-subjects design. Participants navigated or were navigated through 24 straight corridors of a virtual office building. The appearance of the corridors was similar to the environment shown in Figure 1. The distance to be covered through all corridors was 720m, but the lengths of the individual corridors varied. After each corridor, participants had to physically rotate by 90 degrees to face the next corridor, which ensured their attention during the experiment. After both conditions, participants were asked to fill in a Simulator Sickness Questionnaire (SSQ) [19], where 16 sickness symptoms are quantified on a 4-point Likert scale. Participants had a break of 5 minutes between the two conditions, and the whole procedure took approximately 30 minutes to complete. In accordance with previous findings in literature, we hypothesized that the *Passive* condition would lead to higher simulator sickness than the *Active* one.

5.3 Participants

20 participants (11 males, 9 females) aged between 22 and 46 ($M = 27.55$, $\sigma = 5.25$) with diverse backgrounds participated in this study. All participants received an expense allowance of 10 Euros for the successful completion of the experiment.

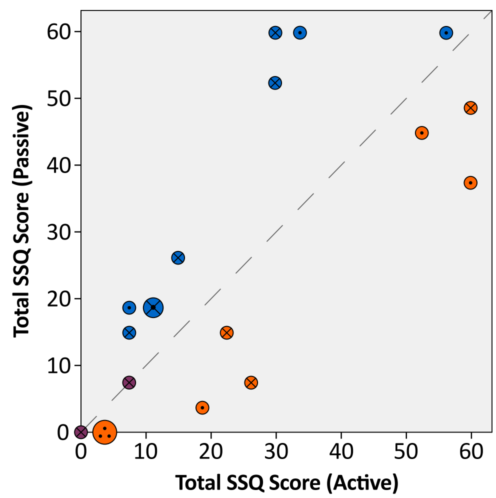

Figure 3: Per-participant scatterplot of the total SSQ scores after active and passive jumping. The dashed diagonal line represents no difference between both conditions. Circles in orange and blue mark participants with higher simulator sickness in the *Active* and *Passive* condition, respectively. A dot within a circle refers to the order *Active-Passive* while a cross represents the order *Passive-Active*. The larger circles subsume two and three identical cases.

5.4 Results and Discussion

The total simulator sickness scores resulting from the SSQ were approximately normally distributed in both conditions. A paired-samples t-test did not show a significant difference between the means of the *Active* ($M = 23.0$, $\sigma = 20.02$) and the *Passive* ($M = 24.68$, $\sigma = 22.22$) conditions, $t(19) = 0.531$, $p = 0.602$, $r = 0.121$. Similar non-significant results were obtained for the subscales nausea (N), oculomotor disturbance (O), and disorientation (D). Figure 3 shows a per-participant scatterplot of the total SSQ scores in the *Active* and *Passive* conditions. We did not observe systematic order effects between both conditions. The mean value of all $Passive - Active$ differences was 1.68 ($\sigma = 14.18$) with a 95% confidence interval of $[-4.95; 8.32]$.

Due to the scaling factor of the total SSQ score, an increase by one on any symptom results in an increase of the total score by at least 3.74. The four symptoms *General discomfort*, *Difficulty focusing*, *Difficulty concentrating*, and *Blurred vision* are taken into account twice, resulting in an increase of 7.48. As a consequence, even the upper bound of the confidence interval can already be achieved by a one-step increase on two of the 16 symptoms. Since the rating of symptoms is very subjective and also dependent on external factors, we therefore consider the differences represented by the confidence interval minimal. As a result, the data of this study provides an indication that the amount of simulator sickness perceived during *Multi-Ray Jumping* in the passenger role is close to the one during active single-user jumping. A significant negative effect of passive navigation, as in related work on steering, could not be observed.

6 Study 2: Advantages of Multi-Ray Jumping

We argued that the comprehensibility of group jumping techniques largely depends on clear previews of target locations for all involved participants. This additional pre-travel information facilitates the navigator's task of planning jumps and a passenger's anticipation of the next location in the scene. We believe that *Multi-Ray Jumping* offers significant benefits in that regard over a naïve implementation of *Coupled* jumping for two participants. In order to quantify these benefits, we conducted a second formal user study in which participants first experienced the passenger role (passenger task) before operating the techniques as the navigator (navigator task). To increase the reproducibility of our study, user activities in the corresponding other role were pre-defined. This means that the passenger was static in the navigator task while the navigator was animated with previously captured motion recordings in the passenger task.

6.1 Experimental Setup

The VR setup consisted of one HTC Vive Pro head-mounted display offering both positional and orientation tracking. Input was obtained using a Vive handheld controller. The tracking space was approximately 3m x 1.5m in size, and the virtual environment was rendered with a resolution of 1440x1600 pixels per eye and an update rate of 90 Hz. We measured an end-to-end latency of 12.5ms.

6.2 Conditions

Participants tested two jumping variants in a within-subjects design. The *Single-Ray* condition served as a baseline and represented a straightforward extension of single-user jumping for two participants. It was mostly identical to the *Coupled* implementation of our pilot study with the addition of a see-through effect when the navigator's ray and target were occluded by scene geometry from the perspective of the passenger. In the *Multi-Ray* condition, participants tested the implementation of *Multi-Ray Jumping* described in Section 4, also with an instant transition and no post-travel feedback. The order of techniques was counterbalanced across participants. However, both techniques were first presented in the passenger and afterwards in the navigator role.

In order to ensure similar distances to the second (simulated) user, participants were asked to stay within a circular area of diameter 0.75m in both conditions (see white circles on the floor in Figure 1). When a participant left this area, the scene lights of the virtual environment turned red to request the user to return. A small sphere on the floor always showed the user's projected head position to simplify this process. The distance of both circle centers was 2.4m, which guaranteed that both users within the circles always kept a distance in social space [15].

6.3 Experimental Tasks

We implemented two parametrizable tasks to quantify the passenger's spatial awareness after the jump as well as the navigator's planning accuracy and efficiency. In both tasks, four distinct spatial configurations of the two avatars were tested. They were either standing side by side or behind each other. These configurations were chosen as they frequently occur when starting in a side-by-side configuration and performing turns of 90° in the virtual world, e.g. while traveling through rectangular grid structures that are typical of many cities and office buildings. Figure 4 illustrates these spatial configurations for the passenger task.

6.3.1 Passenger Task

In the passenger task, we were interested to see if participants can anticipate the resulting spatial configuration after a jump and how long they need in order to reorient themselves. In a similar study on spatial awareness after passive navigation, Bowman et al. suggested measuring the time after travel to find a previously seen object in the scene and answering a simple two-option question on it [6]. We followed this approach but replaced the question on visible information with a rapid aimed movement towards a static object in the scene. At the beginning of each trial, the participant was placed 5m in front of a clearly visible pillar with a sphere as the target object on top. The navigator's avatar stood either to the left, to the right, in front of, or behind the participant (Figure 4). Participants had to press a button on the controller to request a recorded jump from the navigator, which moved them to one of five target positions around the pillar with a distance of 0.75m. After this jump, the task of the participant was to touch the sphere on the pillar with the controller

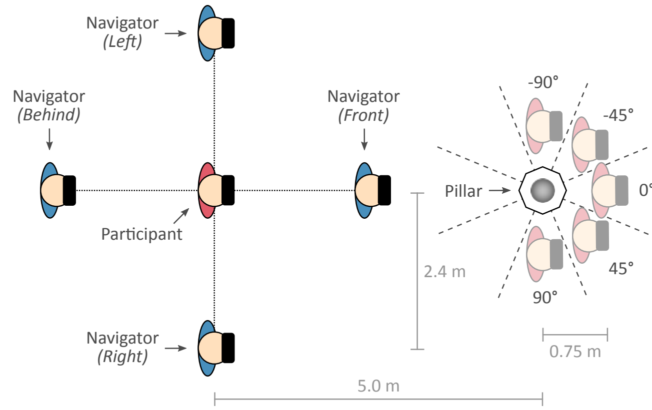

Figure 4: Participants in the passenger role were jumped by a pre-recorded navigator standing in one of four positions. The resulting passenger location after the jump was one of the five indicated positions around a pillar. We captured the reaction time to touch a sphere on top of this pillar as a measure of spatial awareness. The dashed lines emerging from the pillar illustrate the invisible sector borders that were used for our post-hoc task analysis in Section 6.7.1.

as fast as possible. The dependent variable, hence, was the time between the jump and touching the sphere. We deliberately excluded target positions in front of the pillar as we intended to test spatial understanding rather than pure reaction to a visible target location. The pre-recorded actions of the navigator followed a strict procedure for each jump. First, the target specification ray was activated for two seconds while pointing downwards. Next, the parabola was moved towards the target position over a duration of two seconds before initiating the jump. The placement deviation from the target location was lower than 0.05m in all recordings.

6.3.2 Navigator Task

The navigator task was motivated by guiding a museum tour in which the participant should move a passenger to specific locations relative to the exhibits. To remain consistent with the passenger task, we used a similar environment and spatial setup. In a trial of the navigator task, the simulated user was placed 5m in front of a pillar, and the participant in the role of the navigator appeared to the left, to the right, in front of, or behind them. One of five positions around the pillar was highlighted using a circular target. The task of the participant was to place the simulated user as close to the target's center as possible using a single jump (see Figure 5). The two dependent variables were the distance from the target's center (placement error) and the activation time of the target specification ray. In this part of the experiment, the simulated user was static.

6.4 Procedure

Participants arrived at our lab and signed an informed consent form. In the first part of the experiment, participants took the passenger role. They received an introduction sheet explaining the pre-recorded second user and the first jumping technique this navigator will use to move both users around. After putting on the head-mounted display, three jumps could be experienced without any specific task followed by seven training trials of the passenger task. During this phase, the experimenter was allowed to answer questions. Afterwards, participants completed 40 recorded trials in randomized order resulting from two repetitions of each combination of the four navigator positions (left, right, front, behind) and five target positions ($-90°$, $-45°$, $0°$, $45°$, $90°$) illustrated in Figure 4. Lastly, we measured cognitive load using the Raw Task Load Index (RTLX), a simplified version of the NASA Task Load Index without subscale weight-

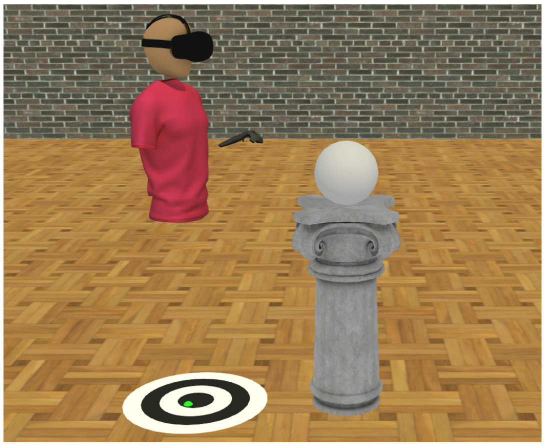

Figure 5: In the navigator task, the participant's task was to place the static passenger as close to a target's center as possible using only a single jump. Afterwards, a green sphere on the floor showed the passenger's projected head position and was used to compute the placement accuracy.

ing [16, 17]. The procedure was repeated for the second jumping technique before the passenger part concluded with a break of 5 minutes. In the second part of the experiment, participants switched to the navigator role. For each jumping technique, they could first complete three free jumps and seven training trials of the navigator task. Afterwards, they completed the same 40 trial configurations as in the passenger task in a new randomized order. Again, navigating with each jumping technique was followed by the RTLX questionnaire. The study ended with a concluding questionnaire on overall technique preferences and demographics. The whole procedure took approximately 60 minutes to complete.

6.5 Hypotheses

In comparison to the *Single-Ray* implementation, we expected the *Multi-Ray* condition to improve the predictability of target locations for the passenger as well as the planning accuracy and rapidity for the navigator. In the passenger task, we hypothesized faster reaction times and lower cognitive load for the *Multi-Ray* condition:

H_1: The average reaction times in the *Multi-Ray* condition will be shorter than in the *Single-Ray* condition.

H_2: The cognitive load scores in the *Multi-Ray* condition will be lower than in the *Single-Ray* condition for the passenger task.

For the navigator task, we hypothesized smaller placement errors, faster target specification times, and lower cognitive load for the *Multi-Ray* condition:

H_3: The average placement errors in the *Multi-Ray* condition will be lower than in the *Single-Ray* condition.

H_4: The average target specification times in the *Multi-Ray* condition will be lower than in the *Single-Ray* condition.

H_5: The cognitive load scores in the *Multi-Ray* condition will be lower than in the *Single-Ray* condition for the navigator task.

6.6 Participants

22 participants (11 males, 11 females) aged between 21 and 30 years ($M = 25.95$, $\sigma = 2.69$) participated in the user study. All of them were either students or employees of our university. On a Likert scale from 1 (rarely) to 7 (often), participants rated their everyday usage of head-mounted displays very low ($Mode = 1$, $Mdn = 1$). All participants received an expense allowance of 10 Euros. To further increase motivation, the user with the best performance won a gift voucher worth 30 Euros.

6.7 Statistical Results

For presenting the results of our user study, we abbreviate the means, medians, and standard deviations by M, Mdn and σ, respectively. When analyzing data for normality, visual inspections of the normal QQ-plots were used in combination with Shapiro-Wilk Tests [32]. When data was non-normally distributed, we tried to apply a log_{10}-transformation to satisfy the assumptions of parametric tests. If this did not succeed, we used a non-parametric equivalent. For each test, we computed the effect size r and applied the threshold values 0.1 (small), 0.3 (medium) and 0.5 (large) introduced by Cohen [8].

6.7.1 Passenger Task

The reaction times of all 40 recorded trials were averaged to a single time per participant and condition. The average time was $M_M = 0.90s$ ($\sigma_M = 0.42s$) in the *Multi-Ray* condition and $M_S = 1.44s$ ($\sigma_S = 0.48s$) in the *Single-Ray* condition. Using a paired-samples t-test, the log_{10}-transformed reaction times in the *Multi-Ray* condition were significantly lower than the ones in the *Single-Ray* condition, $t(21) = 8.08, p < 0.001, r = 0.757$ (large effect), which supports H_1. The RTLX questionnaire outputs an overall cognitive load score ranging between 0 and 100. A paired-samples t-test revealed that the cognitive load in the *Multi-Ray* condition ($M_M = 28.11$, $\sigma_M = 12.89$) was significantly smaller than in the *Single-Ray* condition ($M_S = 43.26$, $\sigma_S = 14.34$), $t(21) = 6.344, p < 0.001, r = 0.811$ (large effect), which supports H_2. *Multi-Ray* was preferred by 18 participants ($\hat{=}$ 81.8%) over *Single-Ray* for the passenger task.

In a post-hoc analysis, we investigated which task configurations in our study were particularly difficult to solve in the *Single-Ray* condition. For this purpose, we expected incongruent tasks like *[left, 90]*, where the navigator ray points to the left of the pillar but the participant lands right of it, to be more challenging than tasks without mismatches (e.g. *[left, -90]*). To formalize these task difficulties, we considered eight sectors around the pillar (illustrated as dashed lines in Figure 4) to define task difficulty by the sector distance between expected pillar direction when interpreting the navigator's ray and actual pillar direction after the jump. This resulted in a difficulty score between 0 (no mismatch) and 4 (maximum mismatch) for each task. Figure 6 shows the mean reaction times separated by task difficulty for both conditions. An overall Kruskal-Wallis test on the data showed significant differences in the reaction time distributions of the *Single-Ray* condition, $H(4) = 21.51$, $p < 0.001$. Post-hoc stepwise step-down analyses identified the two homogeneous subsets $[0, 1, 4]$ and $[4, 2, 3]$. For the *Multi-Ray* condition, no significant differences in distributions were observed, $H(4) = 2.664, p = 0.615$.

6.7.2 Navigator Task

In the navigator task, the placement errors and target specification times of all 40 recorded trials were averaged to single scores per participant and condition. A Wilcoxon signed-rank test showed that the median of placement errors in the *Multi-Ray* condition ($Mdn_M = 0.08m$, $\sigma_M = 0.03m$) was significantly lower than in the *Single-Ray* condition ($Mdn_S = 0.42m$, $\sigma_S = 0.42m$), $W = 0$, $z = -4.107$, $p < 0.001$, $r = 0.876$ (large effect), which backs H_3. However, the medians of the target specification times in the *Multi-Ray* ($Mdn_M = 2.97s$, $\sigma = 1.20s$) and *Single-Ray* condition ($Mdn_S = 3.29s$, $\sigma = 4.26s$) were not significantly different, $W = 80$, $z = -1.51$, $p = 0.131$. This contradicts H_4 although a medium effect size is visible ($r = 0.322$). Figure 7 shows a scatterplot of target specification time and placement error supplemented by boxplots for the individual variables (outliers excluded). Regarding cognitive load, a paired-samples t-test showed a significantly lower mean in the *Multi-Ray* condition ($M_M = 24.47$, $\sigma_M = 15.55$) compared to the *Single-Ray*

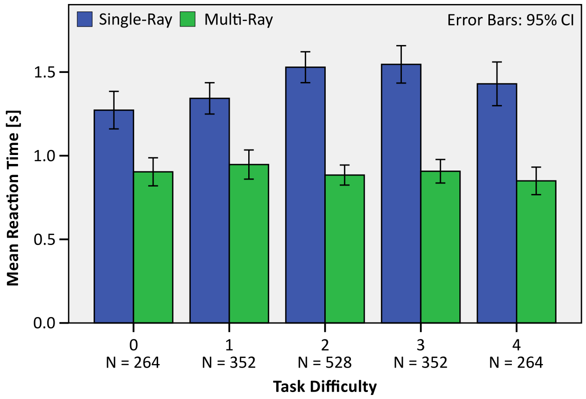

Figure 6: Mean reaction times with 95% confidence intervals over $N = 22 \cdot 40 \cdot 2 = 1760$ trials in the *Single-Ray* (blue) and *Multi-Ray* conditions (green) separated by expected task difficulty.

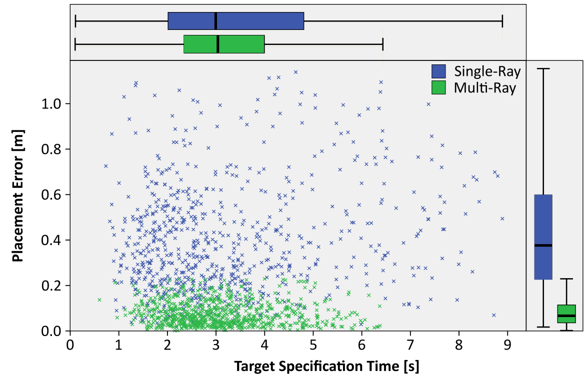

Figure 7: Time-Accuracy scatterplot of all trials of the navigator task without outliers and extreme values ($N = 1760 - 221 = 1539$). Additional boxplots show the distributions of values in the *Single-Ray* (blue) and *Multi-Ray* (green) conditions for both variables.

condition ($M_S = 53.07$, $\sigma_S = 13.29$), $t(21) = 9.668$, $p < 0.001, r = 0.904$ (large effect), thereby supporting H_5. *Multi-Ray* was preferred by 21 participants (≙ 95.5%) over *Single-Ray* for the navigator task.

6.8 Discussion

In the passenger task, *Multi-Ray* clearly outperformed *Single-Ray* in terms of significantly shorter reaction times and lower cognitive load with large effect sizes, which confirmed H_1 and H_2. This implies that the additional pre-travel information given by the secondary parabola could be properly interpreted and beneficially used to spatially comprehend upcoming jumps, thereby also reducing cognitive load. Our post-hoc task analysis revealed that there are indeed task difficulty differences when using the *Single-Ray* technique. However, trials with the largest mismatches were not as difficult as expected. It seems that the plain 180°-mismatches in category 4 (*[left, 90]*, *[right, -90]*, and *[behind, 0]*) were more obvious to recognize and hence easier to integrate than expected. Using the *Multi-Ray* technique made the significant differences in task difficulty vanish, which indicates helpful mediations in both simple and more complex task configurations. In the navigator task, participants showed significantly lower placement errors and cognitive load in the *Multi-Ray* condition with large effect sizes, thereby confirming H_3 and H_5. Contrary to our expectations, however, the time spent for target specification was not significantly different in both conditions, which led to rejecting H_4. Nevertheless, Figure 7 shows that the data range in the *Single-Ray* condition is more than 2 seconds greater than in the *Multi-Ray* condition, which could be an explanation for the observed medium effect size. Hence, it seems that only a subset of participants spent more time for planning jumps in the *Single-Ray* condition, but this did not lead to an overall significant difference between the two techniques. In both cases, the correlation coefficients of target specification time and placement errors were small ($r = 0.178$ for *Single-Ray* and $r = -0.061$ for *Multi-Ray*), indicating that spending more time during target specification did not systematically help to improve placement accuracy.

7 Conclusion and Future Work

Collaborative virtual environments require comprehensible navigation techniques for both collocated and remote user groups. In this paper, we derived the requirements for comprehensible group jumping for collocated users wearing head-mounted displays. Comprehensible techniques need to foster awareness of ongoing navigation actions and make their consequences predictable for the navigator and passengers of a group. Our *Multi-Ray Jumping* technique implements these requirements using additional pre-travel information and consequently showed significant advantages over a straightforward extension of single-user jumping for two users. We therefore conclude that *Multi-Ray Jumping* is more comprehensible as it decreases spatial confusion for the passenger while increasing passenger awareness and thus planning accuracies for the navigator. In addition, *Multi-Ray Jumping* reduces cognitive load in both user roles, which makes it highly beneficial for the joint exploration of virtual environments. For the passengers, future work should investigate the effects of using *Multi-Ray Jumping* in more complex environments on higher levels of spatial awareness like landmark, route and survey knowledge [34].

Our research was primarily motivated by guided tours in virtual reality, where the role distribution of guide (navigator) and attendees (passengers) is inherent and does not change throughout the experience. In other scenarios, dynamic role assignments and cooperative planning can be more relevant. If voice communication enables negotiation, passengers can simply ask the navigator to choose a different navigation target or to stop executing a jump. We also discussed how *Multi-Ray Jumping* affords the fluent negotiation of control, making it well suited for scenarios with more balanced user contributions. A formal evaluation of interaction techniques for collaborative jump planning and dynamic exchange of roles is subject to future work.

Future work also includes extending *Multi-Ray Jumping* to more than two users and investigating its applicability to distributed scenarios. Regarding the former, adding an individual parabola for multiple passengers can easily lead to visual clutter. A solution for the passengers could be to hide the parabolas of the other passengers, which reduces the amount of curves to two as in our presented study. The navigator, however, should see at least the target positions of all involved passengers in order to plan meaningful jumps for the whole group, for example, by ghost avatars. In distributed scenarios, *Multi-Ray Jumping* needs to be complemented by effective coupling and decoupling mechanisms for enabling dynamic changes between phases of individual and group navigation.

Acknowledgments

Our research has received funding from the European Unions Horizon 2020 Framework Programme for Research and Innovation under the Specific Grant Agreement No. 785907 (*Human Brain Project SGA2*) and from the German Ministry of Education and Research (BMBF) under grant 03PSIPT5A (*Provenance Analytics*). We would like to thank *Pauline Bimberg* for her support in conducting our user studies, the participants of our studies, and the members of the Virtual Reality and Visualization Research Group at Bauhaus-Universität Weimar (`http://www.uni-weimar.de/vr`).

References

[1] D. Appleyard. Styles and methods of structuring a city. *Environment and behavior*, 2(1):100–117, 1970.

[2] M. Azmandian, T. Grechkin, and E. S. Rosenberg. An evaluation of strategies for two-user redirected walking in shared physical spaces. In *2017 IEEE Virtual Reality (VR)*, pp. 91–98, 2017.

[3] F. Bacim, E. D. Ragan, C. Stinson, S. Scerbo, and D. A. Bowman. Collaborative navigation in virtual search and rescue. In *2012 IEEE Symposium on 3D User Interfaces (3DUI)*, pp. 187–188, 2012.

[4] S. Beck, A. Kunert, A. Kulik, and B. Froehlich. Immersive Group-to-Group Telepresence. *IEEE Transactions on Visualization and Computer Graphics*, 19(4):616–625, 2013.

[5] J. Bhandari, P. MacNeilage, and E. Folmer. Teleportation without Spatial Disorientation Using Optical Flow Cues. In *Proceedings of Graphics Interface*, pp. 162–167, 2018.

[6] D. A. Bowman, D. Koller, and L. F. Hodges. Travel in immersive virtual environments: an evaluation of viewpoint motion control techniques. In *Proceedings of IEEE 1997 Annual International Symposium on Virtual Reality*, pp. 45–52, 1997.

[7] E. R. Chrastil and W. H. Warren. Active and passive spatial learning in human navigation: Acquisition of survey knowledge. *Journal of experimental psychology: learning, memory, and cognition*, 39(5):1520, 2013.

[8] J. Cohen. A power primer. *Psychological Bulletin*, 112(1):155, 1992.

[9] T. J. Dodds and R. A. Ruddle. Mobile Group Dynamics in Large-Scale Collaborative Virtual Environments. In *2008 IEEE Virtual Reality Conference (VR)*, pp. 59–66, 2008.

[10] T. J. Dodds and R. A. Ruddle. Using Teleporting, Awareness and Multiple Views to Improve Teamwork in Collaborative Virtual Environments. In *Proceedings of the 14th Eurographics Conference on Virtual Environments*, pp. 81–88, 2008.

[11] A. S. Fernandes and S. K. Feiner. Combating VR sickness through subtle dynamic field-of-view modification. In *2016 IEEE Symposium on 3D User Interfaces (3DUI)*, pp. 201–210, 2016.

[12] T. A. Galyean. Guided Navigation of Virtual Environments. In *Proceedings of the 1995 Symposium on Interactive 3D Graphics (I3D)*, pp. 103–105, 1995.

[13] C. Gutwin and S. Greenberg. A Descriptive Framework of Workspace Awareness for Real-Time Groupware. *Computer Supported Cooperative Work*, 11(3):411–446, 2002.

[14] M. P. J. Habgood, D. Moore, D. Wilson, and S. Alapont. Rapid, Continuous Movement Between Nodes as an Accessible Virtual Reality Locomotion Technique. In *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 371–378, 2018.

[15] E. T. Hall. *The hidden dimension*. Doubleday, 1966.

[16] S. G. Hart. NASA-Task Load Index (NASA-TLX); 20 Years Later. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 50(9):904–908, 2006.

[17] S. G. Hart and L. E. Staveland. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. In P. A. Hancock and N. Meshkati, eds., *Human Mental Workload*, vol. 52 of *Advances in Psychology*, pp. 139 – 183. 1988.

[18] V. Interrante, B. Ries, and L. Anderson. Seven league boots: A new metaphor for augmented locomotion through moderately large scale immersive virtual environments. In *2007 IEEE Symposium on 3D User Interfaces (3DUI)*, 2007.

[19] R. S. Kennedy, N. E. Lane, K. S. Berbaum, and M. G. Lilienthal. Simulator sickness questionnaire: An enhanced method for quantifying simulator sickness. *The international journal of aviation psychology*, 3(3):203–220, 1993.

[20] A. Kulik, A. Kunert, S. Beck, C.-F. Matthes, A. Schollmeyer, A. Kreskowski, B. Fröhlich, S. Cobb, and M. DCruz. Virtual Valcamonica: collaborative exploration of prehistoric petroglyphs and their surrounding environment in multi-user virtual reality. *Presence: Teleoperators and Virtual Environments*, 26(03):297–321, 2018.

[21] A. Kulik, A. Kunert, S. Beck, R. Reichel, R. Blach, A. Zink, and B. Froehlich. C1x6: A Stereoscopic Six-user Display for Co-located Collaboration in Shared Virtual Environments. In *Proceedings of the 2011 SIGGRAPH Asia Conference (SA)*, pp. 188:1–188:12, 2011.

[22] J. R. Lackner. Motion sickness: more than nausea and vomiting. *Experimental Brain Research*, 232(8):2493–2510, 2014.

[23] J. Lacoche, N. Pallamin, T. Boggini, and J. Royan. Collaborators Awareness for User Cohabitation in Co-located Collaborative Virtual Environments. In *Proceedings of the 23rd ACM Symposium on Virtual Reality Software and Technology (VRST)*, pp. 15:1–15:9, 2017.

[24] T. T. H. Nguyen, T. Duval, and C. Fleury. Guiding Techniques for Collaborative Exploration in Multi-Scale Shared Virtual Environments. In *International Conference on Computer Graphics Theory and Applications (GRAPP)*, 2013.

[25] S. Notelaers, T. D. Weyer, P. Goorts, S. Maesen, L. Vanacken, K. Coninx, and P. Bekaert. HeatMeUp: A 3DUI serious game to explore collaborative wayfinding. In *IEEE Symposium on 3D User Interfaces (3DUI)*, pp. 177–178, 2012.

[26] S. Razzaque, Z. Kohn, and M. C. Whitton. Redirected Walking. In *Eurographics 2001 - Short Presentations*. Eurographics Association, 2001.

[27] J. T. Reason and J. J. Brand. *Motion sickness*. Academic press, 1975.

[28] A. Rolnick and R. E. Lubow. Why is the driver rarely motion sick? The role of controllability in motion sickness. *Ergonomics*, 34(7):867–879, 1991.

[29] D. Roth, C. Klelnbeck, T. Feigl, C. Mutschler, and M. E. Latoschik. Beyond Replication: Augmenting Social Behaviors in Multi-User Virtual Realities. In *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 215–222, 2018.

[30] A. P. Saygin and W. Stadler. The role of appearance and motion in action prediction. *Psychological research*, 76(4):388–394, 2012.

[31] J. Schild, D. Lerner, S. Misztal, and T. Luiz. EPICSAVE - Enhancing vocational training for paramedics with multi-user virtual reality. In *IEEE 6th International Conference on Serious Games and Applications for Health (SeGAH)*, pp. 1–8, 2018.

[32] S. S. Shapiro and M. B. Wilk. An analysis of variance test for normality (complete samples). *Biometrika*, 52(3-4):591–611, 1965.

[33] S. Sharples, S. Cobb, A. Moody, and J. R. Wilson. Virtual reality induced symptoms and effects (VRISE): Comparison of head mounted display (HMD), desktop and projection display systems. *Displays*, 29(2):58–69, 2008.

[34] A. W. Siegel and S. H. White. The development of spatial representations of large-scale environments. *Advances in child development and behavior*, 10:9–55, 1975.

[35] K. M. Stanney and P. Hash. Locus of User-Initiated Control in Virtual Environments: Influences on Cybersickness. *Presence: Teleoperators and Virtual Environments*, 7(5):447–459, 1998.

[36] J. N. Templeman, P. S. Denbrook, and L. E. Sibert. Virtual Locomotion: Walking in Place Through Virtual Environments. *Presence: Teleoperators and Virtual Environments*, 8(6):598–617, Dec. 1999.

[37] T. Weissker, A. Kunert, B. Frohlich, and A. Kulik. Spatial Updating and Simulator Sickness During Steering and Jumping in Immersive Virtual Environments. In *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 97–104, 2018.

[38] E. A. Wernert and A. J. Hanson. A Framework for Assisted Exploration with Collaboration. In *Proceedings of the Conference on Visualization (VIS) 1999*, pp. 241–248, 1999.

[39] B. Williams, G. Narasimham, B. Rump, T. P. McNamara, T. H. Carr, J. Rieser, and B. Bodenheimer. Exploring Large Virtual Environments with an HMD when Physical Space is Limited. In *Proceedings of the 4th Symposium on Applied Perception in Graphics and Visualization (APGV)*, pp. 41–48, 2007.