

Gaining the High Ground: Teleportation to Mid-Air Targets in Immersive Virtual Environments

Tim Weissker , Pauline Bimberg , Aalok Shashidhar Gokhale, Torsten Kuhlen , and Bernd Froehlich 



Fig. 1: Teleportation to a mid-air target using our *Simultaneous* or *Two-Step* technique. **Left:** Conventional target specification for teleportation using a selection parabola. A preview avatar shows the future position in the environment. **Middle:** *Simultaneous* allows to specify the new elevation with the controller’s touchpad while moving the parabola. *Two-Step* allows the user to select a reference position on the same elevation before tilting the controller to move the preview up or down. For both techniques, a portal window shows the preview avatar at the new elevation. **Right:** Once confirmed, the user is teleported to the specified position.

Abstract— Most prior teleportation techniques in virtual reality are bound to target positions in the vicinity of selectable scene objects. In this paper, we present three adaptations of the classic teleportation metaphor that enable the user to travel to mid-air targets as well. Inspired by related work on the combination of teleports with virtual rotations, our three techniques differ in the extent to which elevation changes are integrated into the conventional target selection process. Elevation can be specified either simultaneously, as a connected second step, or separately from horizontal movements. A user study with 30 participants indicated a trade-off between the simultaneous method leading to the highest accuracy and the two-step method inducing the lowest task load as well as receiving the highest usability ratings. The separate method was least suitable on its own but could serve as a complement to one of the other approaches. Based on these findings and previous research, we define initial design guidelines for mid-air navigation techniques.

Index Terms—Virtual Reality, 3D User Interfaces, 3D Navigation, Head-Mounted Display, Teleportation, Flying, Mid-Air Navigation.

1 INTRODUCTION

Teleportation has emerged as one of the most widely adopted forms of travel through immersive virtual environments as it minimizes the occurrence of sickness symptoms for many users [9, 10, 17, 45]. However, the locations that can be reached using conventional teleportation techniques are limited to the vicinity of scene objects that can be intersected with the selection ray. As a result, users can perform movements along the virtual floor or a series of adjacent rooftops, but they cannot move to targets in mid-air where no objects are nearby. This limitation may prevent the user from, for example, getting an overview of the scene from an elevated position, skipping over larger obstacles on a route, and maneuvering around taller objects of interest at different heights.

In this paper, we present three novel approaches to extend common

teleportation workflows with the ability to specify elevation changes, which enables users to navigate to both object-based and mid-air destinations. We compared our techniques in an empirical user study with 30 participants to investigate their usage in both a constrained route following task and a more exploratory search task. Based on our findings, we discuss the advantages and disadvantages of the three approaches and combine our findings with previous results to define initial design guidelines for mid-air navigation techniques.

Our work is motivated by related research on the specification of virtual rotation as an additional degree of freedom for teleportation. In particular, prior work has presented techniques that enable the user to specify rotation either simultaneously with a target position [5, 20, 44, 46], as part of a two-step process along with a target position [3, 25], or as an additional mechanism independent of movement to a target position [3, 38, 49]. Our three techniques for changing elevation in this paper are derived from these conceptual models. As a result, our central research question for this paper asks which concept is most effective, efficient, comfortable, and therefore suitable for specifying both object-based and mid-air teleports when they are required. To approach this question, our contributions can be summarized as follows:

- Tim Weissker and Torsten Kuhlen are with the Visual Computing Institute at RWTH Aachen University. E-mail: me@tim-weissker.de, kuhlen@vr.rwth-aachen.de
- Pauline Bimberg is with the Human-Computer Interaction Group at the University of Trier. E-mail: bimberg@uni-trier.de
- Aalok Shashidhar Gokhale and Bernd Froehlich are with the Virtual Reality and Visualization Research Group at Bauhaus-Universität Weimar. E-mail: aaloksg@gmail.com, bernd.froehlich@uni-weimar.de

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxx

- the derivation of three conceptual approaches for adding elevation change as an additional degree of freedom to common teleportation workflows
- the design of three novel teleportation techniques based on these concepts, which allow users to execute a variety of horizontal, vertical, and diagonal movements

- a discussion of the benefits and drawbacks of these three techniques as identified by a user study in a constrained route following as well as a more open search scenario ($N = 30$)
- a summary of initial design guidelines for mid-air navigation techniques combining prior as well as the presented results

Our results encourage the consideration of mid-air teleportation techniques for application scenarios in which ground-constrained navigation is not sufficient to meet the user's goals. They also provide relevant guidance for future research and development in this area.

2 RELATED WORK

Navigation is a fundamental task in virtual reality consisting of the motor component *travel* and the cognitive component *wayfinding* in order to explore the environment, search for particular target objects, and maneuver around objects of interest [4]. While several benefits have been demonstrated for travel by physical walking (e.g. [37, 41, 42, 47]), spatial and physiological constraints often require additional virtual travel techniques like steering or teleportation to traverse larger distances. Similar to physical walking, many implementations of virtual travel are restricted to ground-based movements in order to match the user's real-world expectations of the navigable space [31]. This is especially relevant in scenarios where a transfer of results between the real and virtual world is desired, for example for virtual training [7] or studies on human-environment interaction [26]. If the virtual environment consists of different floor levels, prior work suggested transitioning between them using virtual stairs [39], ramps [14], ladders [27, 39], or elevator cabins [43]. To get a better overview of the scene, the user's view may also be temporarily decoupled from the egocentric perspective of their self-avatar by moving it to a static elevated viewpoint for navigation before re-embodiment of the avatar at the new location [11, 21]. The avatar movements in this case are still constrained to the available scene objects. Enabling users to move themselves through mid-air freely, on the other hand, can be beneficial for offering novel perspectives onto the scene during exploration and search (e.g. [32]), moving to target locations more expeditiously (e.g. [31]), and enabling more precise maneuvering around potentially large objects for their close-up inspection (e.g. [25]). Our work in this paper builds upon prior approaches to realizing mid-air travel based on steering and teleportation (Section 2.1) and focuses on the design of three novel teleportation techniques for enabling mid-air travel as well. To do so, we draw inspirations from prior work on the integration of another degree of freedom, namely virtual rotation, into the target specification process of teleportation (Section 2.2).

2.1 Prior Approaches to Mid-Air Travel

Steering techniques are especially convenient to adapt for mid-air travel. Since the user has to continuously specify the intended movement direction using their gaze, body, or controller [1], the vertical component of the corresponding vector can be used to incorporate elevation changes into the movement. In the literature, this form of unconstrained 3D steering is often referred to as *flying* [1, 13, 31, 36]. To increase user comfort during flying, Medeiros et al. suggested to display a virtual floor proxy (the *magic carpet*) as well as a self-body representation to reduce cybersickness, fear of heights, and imbalance issues [31]. Chen et al. suggested to automatically fly the user along the straight-line path to a target selected in a deformed version of the environment that surrounded the user [8]. However, the visual motion flow introduced by steering techniques is often considered a plausible cause of sickness symptoms as it contradicts the vestibular cues perceived by the user [34]. Teleportation-based techniques prevent these contradicting cues and have consequently been shown to mitigate sickness symptoms for a large proportion of users compared to steering [9, 10, 17, 45]. Riecke et al. therefore extended pointing-directed flying with automatic teleports in the indicated direction when the user exceeded a certain velocity threshold [35]. Most implementations of teleportation, however, work without a continuous locomotion component and therefore require the initial selection of a ground-based target location to which the user will then be teleported. The vertical distance of

the user's viewing position to the ground is typically given by their tracked physical height to ensure consistency between the real and virtual ground. Matviienko et al. presented two central challenges in extending ground-based teleportation to mid-air travel [30]: the selection of a target location in mid-air (C1) and the sensible choice of a relative camera placement with respect to this location (C2) [30]. For target selection (C1), Drogemuller et al. suggested employing a straight ray emanating from the controller whose length could be adjusted to indicate the desired travel distance [15]. Lee et al. built upon this method and computed suitable travel distances automatically based on the proximity of surrounding objects [29]. Matviienko et al. themselves compared linear and parabolic variants of a depth cursor (cf. [22]) in combination with different transition modes (C1) and proposed to take the user's controller as a reference object that will exactly match the specified location after a teleport (C2). The study results indicated that the linear depth cursor in combination with instant transitioning led to the highest efficiency and accuracy in an obstacle-free scene. Nevertheless, the authors raised the concern that linear pointing with a freely movable depth cursor is in turn less suited for target specification on the ground. Therefore, research on effective hybrid techniques for specifying ground-based as well as mid-air targets in more populated scenes is mentioned as relevant future work [30]. In this paper, we derive, describe, and evaluate three novel ideas to seamlessly extend ground-based parabolic teleportation techniques with an option to incorporate elevation changes (C1). To improve user comfort during mid-air travel, we adapt the idea of a virtual floor proxy by Medeiros et al. to provide a stable visual representation of the real-world ground, which also serves as a reference object for placing the user's camera at the tracked physical height above it (C2).

2.2 Virtual Rotation as an Additional Degree of Freedom for Target-Based Travel

To extend conventional teleportation workflows with an option to incorporate elevation changes, our work draws inspirations from related research on the specification of virtual rotations. For the purpose of this paper, we distinguish prior work by the degree to which this additional degree of freedom was integrated into the target specification process of teleportation. A more comprehensive overview of virtual rotation approaches is given, for example, in the work of Zielasko et al. [49].

Simultaneous Specification The virtual rotation to be applied during a teleport is specified at the same time as a target position is selected with the selection ray. For example, the *Curved Teleport* technique by Funk et al. [20], the extended *Point & Teleport* technique by Bozgeyikli et al. [5], and the group navigation technique by Weissker et al. [44] all employ the otherwise unused roll angle of the controller to indirectly specify the rotation angle. Alternatively, additional inputs like joysticks or touchpads can be harnessed [20, 46, 49].

Two-Step Specification The position and rotation to be teleported to are specified in two consecutive steps, where the selection of a point of interest in the first step influences the user's rotation at the target position specified in the second step. A prominent example for this in the realm of projection-based systems is the *Navidget* technique by Knoedel et al. [25], which was adapted to the specific requirements of teleportation with head-mounted displays in the technique presented by Bimberg et al. [3].

Separate Specification The specification of a user's position and rotation is decoupled into two distinct techniques that are operated independently of each other. Analogously to the benefits of teleportation over steering for position changes, the study of Sargunam et al. suggested that user-triggered discrete rotations with fixed intervals are less sickness-inducing than continuous rotations [38]. This mechanism is also often referred to as *Rotation Snapping* [3, 49].

In our work for this paper, we applied each of these design paradigms to the specification of elevation and compared the resulting techniques empirically to gain a better understanding of their effectivity, efficiency, and usability. Our prototypes with simultaneous, two-step, and separate mechanisms will be explained in the following.




Fig. 2: Schematic illustration of the interaction workflows of our three proposed mid-air teleportation techniques called *Simultaneous* (see Section 3.2), *Two-Step* (see Section 3.3), and *Separate* (see Section 3.4). Solid user avatars represent the user’s current position in the virtual environment while colored outlines represent preview avatars that are displayed to the user during target specification. The star represents a point of interest the user aims to travel to.

3 TECHNIQUE DESIGN

Based on the different concepts of adding virtual rotations to the teleportation process, we developed three techniques that enable users to specify teleports with elevation changes using a single *HTC Vive* controller. For this purpose, we consider the user standing on a virtual movement platform (or magic carpet) that may teleport horizontally, vertically, or both at the same time. As suggested by Medeiros et al. [31], the platform is constantly visible to the user (see Figure 1 right) alongside a virtual representation of their own body. While all three of our techniques share the same mechanism to teleport the platform based on intersections of the selection ray with scene objects, they differ in the way that elevation changes can be specified in order to travel to targets in mid-air. In the following, we will therefore briefly summarize the basic teleportation features common to all techniques (Section 3.1) before detailing our three approaches to specify elevation changes simultaneously with (Section 3.2), as a consecutive second step after (Section 3.3), and separately from (Section 3.4) horizontal movements. These techniques are illustrated schematically in Figure 2.

3.1 Basic Teleportation Features

Pressing the controller’s trigger activates a typical parabolic pick ray for target selection. By moving the controller, the user can select target positions wherever the parabola intersects with scene objects, which also allows them to travel to smaller upward elevations where corresponding geometry (e.g., the roof of a house) can be reached. However, to avoid unpleasant teleports into walls, we only allow surfaces whose slope does not deviate more than 30 degrees from the horizontal. If the user is already located at a higher elevation (e.g., standing on a roof), the parabola intersects with an invisible proxy plane at the current elevation to allow for subsequent horizontal teleports into mid-air.

During target selection, a preview avatar visualizes how the user will be positioned if the teleport is executed (see Figure 1, left), which has shown to be beneficial in terms of predictability by related work [16, 44, 48]. The teleport can be executed by pressing the trigger fully until it clicks; releasing the trigger earlier cancels the teleport in case the user changes their mind.

3.2 Specifying Elevation Changes: Simultaneous

The *Simultaneous* technique is inspired by virtual rotation techniques that make users specify rotation via additional input channels in parallel to the target position. Applying this idea to mid-air teleportation, our *Simultaneous* technique enables the user to select a reference point on the current elevation level with their parabola while also having the option to move their preview avatar up or down relative to this point by operating the touchpad. To do so, the vertical distance of the user’s




Fig. 3: The transfer function to specify elevation in our *Two-Step* technique consists of three different zones around the initial pointing angle A_{press} at which the reference point on the current elevation level is selected, here 90° . Subsequent movements within $\pm DZ$ are neglected (red), movements within $0^\circ + VZ$ and $180^\circ - VZ$ are directly mapped to the preview’s elevation (orange), and pointing within the remaining zones controls the preview’s movement velocity (blue).

finger from the center of the touchpad is used to determine the speed at which the preview will be moved up or down ($v_{max} = 3m/s$). As shown in Figure 1 (center), an additional pillar geometry with equidistant black rings (distance: $1m$) illustrates the future platform position. If the preview happens to intersect with horizontal geometry on a lower elevation, further downward inputs are discarded to facilitate aligning with scene geometry and to prevent destinations inside of objects or below the ground.

Specifying larger elevation changes with this method can lead to uncomfortable situations in which the user has to strain their neck for prolonged periods of time in order to be able to see the preview. Moreover, the preview also gets increasingly difficult to see as the difference in elevation increases. We therefore suggest to display a portal window in front of the pillar on the user’s current eye level, which shows a third-person view onto the preview avatar at its currently selected elevation. (see Figure 1, center). Using this portal window, the user gets the same visual information they would get for a horizontal teleport, but they do not need to rotate their head or deal with poor avatar visibility for specifying large elevation changes. The portal window only appears when it does not occlude the actual preview avatar and maintains the same relative size independent of the pillar’s horizontal distance to the user.

3.3 Specifying Elevation Changes: Two-Step

The *Two-Step* technique is inspired by virtual rotation techniques that rely on the initial fixation of a point of interest with the controller before using further controller movements to define the user’s relative position and therefore also rotation towards that point. Applying this idea to mid-air teleportation, our *Two-Step* technique asks the user to select and lock a reference point on the current elevation level with their parabola in the subsequent second step, the elevation of the preview can then be adjusted by tilting the controller upward or downward.

The underlying transfer function for this process is illustrated in Figure 3 taking an initial pointing angle of $A_{press} = 90^\circ$ as an example at which the reference point is selected. Around this initial angle, the function defines a deadzone ($\pm DZ = \pm 3^\circ$, red area) within which no angle changes are applied to the preview. This ensures that minor movements caused by hand tremor or tracking errors do not prevent users from executing perfectly horizontal teleports. When leaving the deadzone, a linear mapping between pointing angle of the controller and elevation of the preview allows for direct and precise adjustments ($1^\circ \cong 0.1m$, orange area). With this linear mapping, however, larger elevation changes would require uncomfortable or even impossible arm movements. We therefore propose switching to a velocity-based trans-




Fig. 4: Interaction sequence to change elevation using our *Separate* technique. (a+b) A portal window is attached to the user’s field of view, which initially shows the same content as the main view. (c) The portal view can be moved up and down using the touchpad of the controller. The user can look around to inspect the surroundings within the portal. (d) After confirmation, the user is teleported to the specified elevation level.

fer function when the absolute controller angle approaches verticality within a certain range ($\pm VZ = \pm 30^\circ$, blue area). Within this velocity zone, the angular difference to the beginning of the zone defines the speed at which the preview avatar will move up or down ($1^\circ \triangleq 0.1m/s$). As a result, the user can move the preview to a coarse region using the velocity zone before performing more fine-grained adjustments using the direct mapping. Entering and exiting the velocity zone is communicated to the user via a controller rumble to increase awareness.

Analogously to the *Simultaneous* technique, a pillar highlights the future position of the platform, and a portal window shows the perspective on the preview avatar at the new elevation to offer the user the same view onto the preview as they would have for a horizontal teleport (see Figure 1). As before, downward inputs are discarded if the new platform position intersects with horizontal geometry. Releasing the previously held trigger executes the teleport.

3.4 Specifying Elevation Changes: Separate

The *Separate* technique is inspired by rotation techniques that operate independently of the parabolic selection ray and therefore allow the user to rotate in place. Applying this idea to mid-air teleportation, the most straightforward approach would be to apply the principles of *Rotation Snapping* to the vertical movement direction, meaning that button presses would teleport the platform vertically up or down in fixed increments. However, unlike for rotation techniques, we consider this approach too constraining for mid-air teleportation since the user cannot correct their position to a steady elevation between the increments by physical movements. Therefore, our *Separate* technique was designed to enable arbitrary up and down movements along the vertical axis. The corresponding interaction sequence is shown in Figure 4. To change elevation, the user presses the menu button above the touchpad, which attaches a quadratic portal window to their field of view, which initially shows the same content as the main view. The user can then move the view within the portal up or down using the touchpad analogously to the controls described for the *Simultaneous* technique ($v_{max} = 3m/s$). Since the portal is attached to the user’s view, they can freely look around at the specified elevation before traveling there by pressing the trigger. We decided for this portal-based approach over direct steering since it minimizes the motion flow in the user’s periphery as well as provides a stable rest frame outside of the portal – two promising strategies to minimize the occurrence of sickness symptoms (e.g., [6, 18]). As before, downward inputs are discarded if the new platform position intersects with horizontal geometry.

Unlike for the other two techniques, we explicitly chose to display a first-person instead of a third-person perspective within the portal here. While a third-person perspective is a consistent choice for *Simultaneous* and *Two-Step* as the user is already placing their avatar from an external viewpoint when the portal becomes visible, *Separate* does not fulfill this prerequisite. As a result, opening the portal would lead to a change of perspective “backwards” that we considered more difficult to grasp than the seamless transition offered by displaying a first-person perspective.

3.5 Discussion

Our three presented techniques pursue different strategies with respect to integrating elevation changes into conventional teleportation workflows. *Simultaneous* allows to manipulate all degrees of freedom at the same time, which could either turn out efficient or result in an inordinate task load and therefore lower usability. *Separate* follows the opposite approach by splitting off the elevation degree of freedom into another mechanism that is distinct from the conventional selection parabola. This clear division could turn out beneficial for task load and precision or, on the other hand, too constraining in terms of movement paths. *Two-Step* represents a compromise between the full integration and full separation of elevation specification by making it a connected second step based on the fixation of an initial reference position. This has the potential to combine either the benefits or the drawbacks of the other two approaches, and the movement-based transfer function for specifying elevation might be harder to learn and operate than the touchpad. Based on these considerations, we decided to compare the three techniques in a user study in order to better understand their effectivity, efficiency, and usability. For this purpose, the described implementations were designed to be as comparable as possible by using similar visualizations and similar velocity parameters where applicable. Our user study will be presented in the following.

4 EMPIRIC COMPARISON OF ELEVATION SPECIFICATION

Based on our discussion in Section 3.5, we compared our three techniques *Simultaneous*, *Two-Step*, and *Separate* in a user study focusing on their effectivity, efficiency, and usability.

4.1 Hardware Setup and Virtual Environment

We equipped a quiet corner of our laboratory with a workstation, an *HTC Vive Pro 2* system, and two base stations 2.0 mounted on the walls. Participants had a flat interaction space of approximately 3.8m x 2.5m to freely move within before virtual navigation was required. The virtual environment of this study was a low-poly cityscape covering an area of around 280m x 215m with several buildings and monuments that motivated navigation across multiple elevation levels. To prevent users from navigating too far into the sky, our three navigation techniques were clamped to a maximum elevation of 65m, which corresponded to the top antenna of the tallest building in the city. The scene was rendered using *Unity3D* at the native resolution of the head-mounted display (2448 x 2448 pixels per eye) and an update rate of 90Hz.

4.2 Experimental Tasks

We evaluated our techniques in two subsequent tasks with different requirement characteristics: a constrained and therefore more controlled route following task (Section 4.2.1) as well as a more explorative search task (Section 4.2.2).

4.2.1 Route Following Task

In the route following task, participants were shown a series of way-points through the environment and asked to navigate to each of them



Fig. 5: In the route following task, participants moved to and activated a series of virtual camera geometries at various elevations using our navigation techniques. Prior to starting the task, a World-in-Miniature was shown to participants illustrating the entire route.

in the given sequence. As shown in Figure 5, each waypoint was visualized in the form of a camera geometry that pointed towards an interesting perspective in the environment. A red cone and a white guiding line were shown to the user to indicate the next camera to navigate to. Intersecting the controller with the camera’s display activated the next camera, which was confirmed with a click sound. The total length of the route was 238m and included 17 cameras at different elevations from the ground to one of the floating hot-air balloons at an elevation of 50m in mid-air. 4 pairs of cameras had a minor elevation difference of less than 2m (4 upward, mainly at the beginning of the route), 8 pairs had an elevation difference between 2m and 10m (4 upward and 4 downward), and 4 pairs had a large elevation difference of more than 10m (2 upward and 2 downward). We therefore ensured that participants had to perform a large variety of different teleports in order to get a better feel for our navigation techniques. Since individual participants might suffer from fear of heights, a World-in-Miniature [40] was shown to participants illustrating the entire route before activating the first target (see Figure 5). This gave participants another option to opt out of the experiment if they did not feel comfortable with the task.

4.2.2 Search Task

We designed three different instances of a search task, in which participants were given a restricted subspace of the virtual city (size: 90m x 75m per task instance) to search for three cats that were hidden at a high elevation of more than 30m (e.g., on a hot air balloon or airplane), a medium elevation of around 10m (e.g., on a roof), and on the ground, respectively. To assist participants, we placed a virtual notice board at the starting location of each search that listed three one-sentence clues as to where the cats were located (see Figure 6). Participants had as much time as needed to study these clues in an otherwise empty environment before the experimenter started the task and the surrounding cityscape appeared around the notice board. While searching for the cats, participants could always return to the starting location to read the clues again if required. The task was terminated either when all three cats were found and touched with the controller or after a maximum time limit of five minutes had elapsed. In contrast to the route following task, this task was less constrained and therefore required participants to perform their own route planning. From this, we hoped to gain additional insights into the use of our techniques in a more exploratory usage context.

4.3 Procedure

Participants arrived at our lab, were informed about the purpose of the study, and signed a consent form. In the first part of the study, participants were asked to complete the route following task with the three navigation techniques in a counterbalanced order based on a Latin Square. In each of these runs, the experimenter started by



Fig. 6: In the search task, participants received clues as to where they could find three hidden cats in the virtual environment. They were asked to explore the surroundings to find and touch the cats with their controller.

demonstrating the current navigation technique and controls using the desktop control monitor. Participants then put the head-mounted display on and practiced its usage in a separate and considerably smaller tutorial version of the virtual city. They were also asked to complete a practice route of five virtual cameras for training the task procedure. Once all questions were clarified, participants were placed in the full version of the virtual city and completed the above-described route of 17 cameras, which was identical for all three navigation techniques. Afterwards, they put the head-mounted display off and completed a questionnaire on the desktop, which consisted of a single question to quantify overall discomfort including sickness (“On a scale from 0-10, 0 being how you felt coming in, 10 is that you want to stop, where are you now?” [2, 33]), the Raw TLX to quantify task load [23, 24], and a few custom questions to quantify ease of learning, ease of use, confusion, and overall impression. After this procedure was completed for all three navigation techniques, participants were asked to rank the techniques and to provide an optional written justification.

In the second part of the study, participants were asked to complete the three different instances of the search task with the navigation techniques in the same order as presented before. In each of these runs, participants had the chance to briefly recap the use of the current technique in the tutorial scene before being introduced to the clues of the current search task instance. To ensure that each of the task instances was completed equally often with each navigation technique across all participants, the instances always appeared in the same order such that the counterbalanced order of navigation techniques led to this equal distribution. Since each repetition featured different search locations in a different part of the city, it was not possible for participants to gain relevant knowledge for a future repetition. After the search was completed with a navigation technique, participants took the head-mounted display off and answered a similar questionnaire as in the route following task, where the question on ease of learning was removed. Additionally, we asked participants to think back to both the route following as well as the search task and complete the User Experience Questionnaire (UEQ) [28] regarding their navigation experience. After this procedure was completed for all three navigation techniques, participants were once again asked to rank the techniques and provide an optional written justification.

In the end, we asked participants to fill in a final questionnaire on demographics before thanking them for their participation. The entire procedure of the user study took between 60 and 90 minutes to complete. Participants did not receive any form of monetary compensation for taking part in the study.

4.4 Dependent Variables and Hypotheses

During both tasks, our system logged each teleport in terms of start and end position as well as specification time. The post-hoc questionnaires

after each technique in both tasks yielded a discomfort score between 0 and 10, an overall task load score between 0 and 100 derived from the answers of the Raw TLX, and individual question scores quantifying ease of learning (1 to 7), ease of use (1 to 7), confusion (1 to 7) as well as the overall suitability of the technique for the performed task as a grade (1 to 5). The ranking performed after all techniques were tested for a task provided an additional numeric indicator of suitability (1 to 3). Finally, the evaluation of the UEQ yielded six scores between -3 and 3 representing the perceived attractiveness, perspicuity, efficiency, dependability, stimulation, and novelty of the three techniques. Based on all these measures, we derived hypotheses before conducting our experiment as a prerequisite for inferential statistical tests. However, given our discussion about the unclear directions of potential effects in Section 3.5, we opted for an undirected formulation of the hypotheses. In the evaluation, we therefore also only considered the less powerful two-tailed significance values for the statistical tests that were based on a symmetric probability distribution. We separate these hypotheses into the ones related to logging data (Section 4.4.1), the standardized questionnaires (Section 4.4.2), and the custom questionnaires (Section 4.4.3).

4.4.1 Logging Data

We decided to take the number of teleports as well as the distance covered in the more comparable route following task as indicators of accuracy. We complemented this by analyzing the mean specification time for teleports involving elevation changes as a measure of technique efficiency in both tasks:

- The mean number of teleports (H_1) and distance covered (H_2) in the route following task will be different based on the used navigation technique.
- The mean specification time per teleport (H_3) will be different based on the used navigation technique and the task.

4.4.2 Standardized Questionnaires

We hypothesized influences of technique and task on task load as measured by the Raw TLX questionnaire. For the only-once conducted UEQ, we could only hypothesize an influence of technique:

- The mean scores for task load (H_4) will be different based on the used navigation technique and the task.
- The mean scores of the six UEQ subscales (H_5) will be different based on the used navigation technique.

We did not formulate a hypothesis for the discomfort score since all techniques were based on teleportation as the core travel metaphor, which was shown to be favorable in terms of sickness [34, 45]. Moreover, discomfort could also be affected by fear of heights, which we also expected to be similar across conditions due to the identical platform and teleport visualizations when in mid-air. We will therefore only analyze this variable descriptively, thereby providing validation that participants were in good shape throughout the study.

4.4.3 Custom Questionnaires

Regarding our custom questionnaires, we hypothesized influences of both technique and task for most of the resulting variables. The only exception was the ease of learning score, which was measured only once. We did not formulate a hypothesis for technique ranking and therefore will only perform descriptive analyses for this variable:

- The mean ease of learning score (H_6) will be different based on the used navigation technique.
- The mean scores for ease of use (H_7), confusion (H_8), and grading (H_9) will be different based on the used navigation technique and the task.




Fig. 7: Boxplots illustrating the usage data logged by our application. We compared our techniques by the number of teleports in the route following task, the distance covered in the route following task, and the average specification time for elevation changes in both tasks.

4.5 Participants

Our study was completed by 30 participants (17 male, 10 female, 2 diverse, 1 unknown) between 20 and 34 years of age ($M = 25.6$, $\sigma = 4.06$). We placed an emphasis on primarily recruiting users with VR experience in order to reduce a potential novelty bias and to obtain more in-depth feedback. Therefore, our sample consisted of only two first-time VR users while 5, 10, and 13 participants rated themselves as beginner, advanced, and expert users, respectively. Nobody decided to stop the experiment early, giving us $N = 30$ data points per variable for the analysis.

5 RESULTS

Based on our hypotheses, we analyzed our data using *IBM SPSS Statistics*. For hypotheses concerning measurements for each technique after each task, we performed a 3x2 factorial repeated-measures ANOVA to test for statistically significant main and interaction effects. For variables that were captured only once per technique, a one-way repeated-measures ANOVA was conducted instead. For all ANOVAs, we also computed the effect size η_p^2 with the threshold values of 0.01, 0.06, and 0.14 representing small, medium, and large effects, respectively [12, pp. 285–287]. Bonferroni-corrected post-hoc paired-samples t-tests were performed when the overall effect was statistically significant. If the interaction effect was also significant, we separated this post-hoc analysis by task and increased the Bonferroni correction factor accordingly. For all post-hoc tests, we computed the effect size d with the threshold values of 0.2, 0.5, and 0.8 for the above-mentioned effect magnitudes [12, pp. 24–26].

Conducting valid repeated-measures ANOVAs requires a normal sampling distribution as well as equal variances of the differences between conditions (sphericity). Given our sample size of $N = 30$, we can carefully tend to assume a normal sampling distribution based on the central limit theorem [19, pp. 170–172]. Regarding sphericity, we conducted Mauchly’s tests and reported Greenhouse-Geisser corrected results when sphericity was violated. Since there were only two tasks in the study, the criterion of sphericity did not apply to the main effects of the task.

5.1 Logging Data

Figure 7 shows the distributions of the data logged by our application, which will be supplemented by inferential analyses in the following.

Route Following Task Completion (H_1, H_2): We observed that the number of teleports in the route following task was significantly affected by technique ($F(2, 58) = 110.177$, $p < 0.001$, $\eta_p^2 = 0.792$). Post-hoc tests revealed significant differences between all pairs of techniques (all $p < 0.001$, $d > 0.85$). We also observed that the distance covered in the route following task was significantly affected by technique




Fig. 8: Left: Boxplots illustrating the distributions of task load scores that resulted from the Raw TLX questionnaire. Right: Estimated marginal means of the task load score demonstrating the interaction effect between technique and task.

($F(2, 58) = 190.230, p < 0.001, \eta_p^2 = 0.868$). Again, post-hoc tests revealed significant differences between all pairs of techniques (all $p < 0.005, d > 0.64$).

Specification Time for Elevation Changes (H_3): We observed a significant main effect of technique ($F(2, 58) = 10.856, p < 0.001, \eta_p^2 = 0.272$), a significant main effect of task ($F(1, 29) = 20.234, p < 0.001, \eta_p^2 = 0.411$), and a significant interaction effect between both ($F(2, 58) = 37.391, p < 0.001, \eta_p^2 = 0.563$). In both tasks, post-hoc tests revealed significant differences between *Simultaneous* and *Two-Step* ($p_{Route} = 0.007, d_{Route} = 0.662, p_{Search} = 0.001, d_{Search} = 0.813$) as well as *Simultaneous* and *Separate* ($p_{Route} = 0.023, d_{Route} = 0.757, p_{Search} = 0.023, d_{Search} = 0.572$). The difference between *Two-Step* and *Separate* was only significant in the search task ($p_{Route} = 1.0, d_{Route} = 0.217, p_{Search} < 0.001, d_{Search} = 1.211$).

5.2 Standardized Questionnaires

Task Load (H_4): A descriptive overview of the measured task load scores per technique and task is given in Figure 8. We observed a significant main effect of technique ($F(2, 58) = 5.236, p = 0.008, \eta_p^2 = 0.153$), a non-significant main effect of task ($F(1, 29) = 0.275, p = 0.604, \eta_p^2 = 0.009$), and a significant interaction effect between both ($F(2, 58) = 3.346, p = 0.042, \eta_p^2 = 0.103$). Post-hoc tests revealed only one significant difference between *Simultaneous* and *Two-Step* in the route task ($p_{Route} = 0.027, d_{Route} = 0.562, p_{Search} = 1.0, d_{Search} = 0.092$). The comparisons between *Simultaneous* and *Separate* ($p_{Route} = 1.0, d_{Route} = 0.190, p_{Search} = 0.643, d_{Search} = 0.303$) as well as *Two-Step* and *Separate* ($p_{Route} = 0.169, d_{Route} = 0.422, p_{Search} = 0.350, d_{Search} = 0.360$) were non-significant in both tasks.

UEQ Subscales (H_5): The results of the statistical tests as well as boxplots evaluating the six UEQ subscales are given in Figure 9. In summary, we observed overall significant differences on all six subscales with medium to large effect sizes. Post-hoc tests revealed significant differences between all pairs of techniques for attractiveness (all $d > 0.65$) and efficiency (all $d > 0.49$). For perspicuity and dependability, only the comparison of *Simultaneous* and *Two-Step* was significant (both $d > 0.5$) while this was the only non-significant comparison for stimulation ($d = 0.283$). For novelty, the comparison between *Simultaneous* and *Separate* was significant ($d = 0.529$).

Discomfort Score (descriptive only): The mean discomfort scores reported after the use of each technique in each task (with a possible range of 0 to 10) were all less than 1 with standard deviations of less than 1.5. Most individual scores (167) were distributed within the range from 0 to 2 with a few outliers (12) between 3 and 4. The remaining observation was an extreme value of 7 after using the *Simultaneous* technique in the route following task.

	Overall Test				Post-Hoc Tests		
	ϵ	F	p	η_p^2	d	d	d
Attractiveness	0.786	22.013	<0.001	0.432*	0.678*	1.072*	0.652*
Perspicuity	1.000	4.969	0.010	0.146*	0.627*	0.259	0.295
Efficiency	1.000	22.549	<0.001	0.437*	0.495*	1.078*	0.759*
Dependability	1.000	3.235	0.047	0.100*	0.532*	0.361	0.064
Stimulation	0.826	12.980	<0.001	0.309*	0.283	0.833*	0.594*
Novelty	1.000	5.509	0.006	0.160*	0.231	0.387	0.529*




Fig. 9: Top: Results of the statistical tests conducted on the six subscales of the User Experience Questionnaire (UEQ). The effect sizes η_p^2 and d are marked with an asterisk if the corresponding (corrected) test was significant at $p < 0.05$. Bottom: Boxplots illustrating the score distributions on the six subscales separated by technique.

5.3 Custom Questionnaires

A descriptive overview of responses given to our single-item questions on ease of learning, ease of use, confusion, and technique grading is given in Figure 10 and supplemented by inferential analyses in the following.

Ease of Learning (H_6): The assumption of sphericity was violated, so test results had to be corrected ($\epsilon = 0.825$). We observed that the ease of learning score was significantly affected by technique ($F(1.650, 47.849) = 6.046, p = 0.007, \eta_p^2 = 0.173$). Post-hoc tests revealed significant differences between *Simultaneous* and *Two-Step* ($p = 0.022, d = 0.528$) while the remaining comparisons were non-significant (both $p > 0.15, d < 0.38$).

Ease of Use (H_7): We observed a significant main effect of technique ($F(2, 58) = 4.884, p = 0.011, \eta_p^2 = 0.144$), a significant main effect of task ($F(1, 29) = 13.606, p = 0.001, \eta_p^2 = 0.319$), and a non-significant interaction effect between both ($F(2, 58) = 0.218, p = 0.805, \eta_p^2 = 0.007$). Post-hoc tests revealed overall significant differences between *Simultaneous* and *Two-Step* ($p = 0.030, d = 0.503$) as well as *Two-Step* and *Separate* ($p = 0.016, d = 0.552$). The comparison of *Simultaneous* and *Separate* was non-significant ($p = 1.0, d = 0.032$).

Confusion (H_8): We observed a non-significant main effect of technique ($F(2, 58) = 0.110, p = 0.904, \eta_p^2 = 0.003$), a significant main effect of task ($F(1, 29) = 32.537, p < 0.001, \eta_p^2 = 0.529$), and a non-significant interaction effect between both ($F(2, 58) = 0.775, p = 0.466, \eta_p^2 = 0.026$).

Grading (H_9): We observed a significant main effect of technique ($F(2, 58) = 19.588, p < 0.001, \eta_p^2 = 0.403$), a significant main effect of task ($F(1, 29) = 5.241, p = 0.03, \eta_p^2 = 0.153$), and a significant interaction effect between both ($F(2, 58) = 3.610, p = 0.033, \eta_p^2 = 0.111$). Post-hoc tests revealed significant differences between *Two-Step* and *Separate* in both tasks ($p_{Route} = 0.009, d_{Route} = 0.644, p_{Search} < 0.001, d_{Search} = 1.346$). The comparisons of *Simultaneous* and *Two-Step* ($p_{Route} = 0.709, d_{Route} = 0.294, p_{Search} = 0.003, d_{Search} = 0.710$) as well as *Simultaneous* and *Separate* ($p_{Route} = 0.288, d_{Route} = 0.377, p_{Search} < 0.001, d_{Search} = 0.905$) were only significant in the search task.

		Simultaneous		Two-Step		Separate	
		M	σ	M	σ	M	σ
Ease of Learning		5.13	1.55	5.97	1.07	5.60	1.22
1 very difficult – 7 very easy							
Ease of Use	R	5.80	1.19	6.27	1.02	5.83	1.15
1 very difficult – 7 very easy	S	6.33	0.92	6.77	0.50	6.23	0.86
Confusion	R	5.80	1.16	5.73	1.31	5.67	1.27
1 after every teleport – 7 never	S	6.37	0.72	6.57	0.86	6.50	0.73
Grading	R	1.73	0.70	1.47	0.69	2.07	0.79
1 very good – 5 very poor	S	1.53	0.51	1.14	0.35	2.10	0.71

Fig. 10: Means (M) and standard deviations (σ) of responses given to the single-item questions on ease of learning, ease of use, confusion, and technique grading. Where applicable, results are divided into the route following task (R) and the search task (S).

Technique Ranking (descriptive only): Figure 11 visualizes how often a certain rank was assigned to each of the techniques after both the route following and the search task. Overall, *Separate* was most often rated the worst while *Simultaneous* and *Two-Step* both appear frequently as best technique. For *Simultaneous*, however, the number of appearances as second and worst technique differs strongly between the route following (6 second, 11 worst) and the search task (13 second, 5 worst).

5.4 Discussion

Overall, all participants were able to complete the tasks of the user study with all three techniques. However, we observed significant differences with medium to large effect sizes on many of the dependent variables, which leads to a more careful consideration of the advantages and disadvantages of the individual approaches.

5.4.1 Time, Accuracy, and Task Load

It was not surprising that *Separate* required the highest number of teleports and the largest travel distance to complete the route following task, which was due to the strict decomposition of diagonal movements into their vertical and horizontal components. More interestingly, the comparison of *Simultaneous* and *Two-Step* was also significant on these two variables, indicating that *Simultaneous* led to more efficient diagonal movements than *Two-Step*. Based on the free texts provided by participants, the main challenge with *Two-Step* was the mental projection of the next waypoint on the current elevation plane in order to be able to lock an appropriate reference point with respect to which the elevation change should be applied. As *Simultaneous* allowed adjusting this reference point in parallel to setting the desired elevation with the touchpad, participants could use the preview to plan teleports more accurately and thereby reduce the need to perform follow-up teleports for correction. However, achieving this increased accuracy in the route following task also came with increases in specification time and task load compared to *Two-Step*, which indicates a trade-off when choosing between the two techniques. In the search task, both techniques required smaller specification times, but *Two-Step* was still significantly faster than *Simultaneous*. However, the significant difference in task load vanished in this task. We therefore conclude that operating *Simultaneous* was less demanding in the search task, which is likely due to the fact that participants did not have to focus on the exact acquisition of waypoints anymore. Conversely, however, the task load induced by *Two-Step* did not benefit from the more exploratory task context.

Looking at *Separate* in the route following task, the specification times for elevation changes were significantly smaller than for *Simultaneous*. While this appears promising at first sight, it is important to reiterate that *Separate* typically requires follow-up horizontal teleports if the destination is not directly above or below the user's current location, which adds to the total time required to get to a given destination. In the search task, *Separate* did not achieve overall smaller specification times like the other two techniques, which resulted in the fact that even the individual vertical teleports of *Separate* were significantly slower to specify than the diagonal teleports of *Simultaneous* and *Two-Step*. Moreover, the separation of vertical and horizontal movements was also




Fig. 11: Ranking submitted by participants after each task was completed with all three techniques. The size of the bars and the enclosed numbers indicate the frequencies with which the techniques were placed in each rank.

not accompanied by significant reductions in task load. We therefore conclude that *Separate* appears to be the least promising approach for realizing mid-air teleportation in terms of elevation specification time and induced task load.

5.4.2 Usability and Wellbeing

All three techniques yielded similarly low discomfort scores, which was expected due to the teleportation-based movements and comparable helper visualizations when in mid-air.

On the User Experience Questionnaire (UEQ), *Two-Step* was rated significantly higher than *Simultaneous* in terms of attractiveness, perspicuity, efficiency, and dependability. *Separate*, on the other hand, received significantly poorer scores than both *Two-Step* and *Simultaneous* in terms of attractiveness, efficiency, and stimulation. Nevertheless, the recorded scores of all techniques on all sub-scales ranged far up and almost always included the maximum score of +3.0 (only exception: *Separate* on dependability). Overall, we therefore conclude that *Separate* is the most controversial technique regarding its usability while the comparison of the other two techniques points towards a slight usability preference of *Two-Step* over *Simultaneous* with the exception of stimulation and novelty. We believe that the poorer usability scores of *Separate* are likely a consequence of the more laborious navigation sequence discussed before.

In our custom questionnaires, we observed no significant differences between techniques regarding confusion. Since the mean scores on this variable were rather high, we presume that participants were able to understand the provided preview visualizations and portal views rather well. While statistical non-significance does not allow to infer that there is no difference, the small main effect size of technique does provide an initial indication that the consistent choice of a first-person portal view for *Separate* as opposed to a third-person portal view for the other techniques did not lead to additional confusion. Looking at the other variables, *Two-Step* received significantly higher scores than *Simultaneous* regarding the ease of learning, the ease of use, and the awarded grade in the search task. *Simultaneous* and *Separate*, on the other hand, did not yield significant differences regarding the ease of learning, ease of use, and route task grading. However, the grade of *Simultaneous* was significantly better than the one of *Separate* in the search task. Based on these results, we can conclude advantages of *Two-Step* over the other two approaches as well as a slight advantage of *Simultaneous* over *Separate* for more exploratory task scenarios.

5.4.3 Preference

Based on the results discussed previously, it is reasonable that *Two-Step* and *Simultaneous* were often ranked in first or second place while *Separate* appeared most often in last place. Interestingly, however, *Simultaneous* also appeared quite often in last place for the route following task despite the demonstrated advantages in terms of accuracy.

We therefore believe that the higher task load imposed on participants was a relevant factor for the ranking. In the search task where the task load imposed by *Simultaneous* was reduced, it also appeared less frequently in last place. The resulting gap was almost exclusively filled by votes for *Separate*. This once again indicates that the separation of horizontal and vertical movements by *Separate* was not the ideal model for navigating to mid-air targets. However, several participants mentioned that the elevation specification of *Two-Step* and *Simultaneous* was more difficult to operate when the horizontal distance to the destination was small since the preview geometries were too close to the viewing position. It was therefore suggested to offer *Separate* as a complement to either *Simultaneous* or *Two-Step* for these situations. Interestingly, this suggestion is in line with the results of Bimberg et al. [3] as well as several commercial applications offering virtual rotations, where *Rotation Snapping* on the spot serves as a complement to more complex techniques that can integrate rotation changes into the conventional target specification process for teleportation.

6 INITIAL DESIGN GUIDELINES FOR MID-AIR NAVIGATION

Mid-air navigation techniques become relevant if users need to gain novel perspectives onto the scene from elevated viewpoints, move to target locations more expeditiously, or maneuver around tall objects to inspect them. In virtual reality applications in which these use cases are not relevant, developers can safely resort to conventional ground-based navigation techniques as known from prior research and applications. In the following, we offer some initial recommendations for the design of mid-air navigation techniques based on the results of our study as well as previous work in the field.

Generally, standing in mid-air can feel uncomfortable for the user and should therefore be avoided without additional mediation. It is advisable to display a stable visual representation of the real-world ground under the user's feet and provide a self-body avatar as a visual reference when looking downward [31]. To move through the three-dimensional space, steering-based techniques are the most straightforward choice since the vertical component of the user-specified direction vector can be directly mapped to corresponding movements. However, this method is only preferable for users who are not susceptible to sickness symptoms during continuous movements. For a large proportion of users, we therefore recommend teleportation-based techniques for mid-air navigation based on the reductions in sickness symptoms demonstrated by prior work [9, 10, 17, 45]. However, the free specification of target positions in mid-air is more challenging since there are no selectable scene objects in the vicinity [30]. If the virtual environment is only sparsely populated, linear pointing techniques with a depth selection method appear to be a good choice for mid-air target specification [15, 30]. If the virtual environment is more populated, however, the user must be able to specify both object-based as well as mid-air targets. It is therefore advisable to extend the commonly seen parabolic target selection method for object-based targets with options to incorporate elevation changes if they are required. Based on the results presented in this paper, we recommend allowing the user to manipulate all resulting degrees of freedom simultaneously when the scenario requires high specification accuracy in a low number of teleports, with the caveat that this might also result in longer specification times and higher task load. In less demanding scenarios, where arriving in the close vicinity of a target location is sufficient or where correction teleports can be tolerated, we suggest embedding elevation changes as a connected second step after selecting a reference position on the current elevation, which decreases task load and improves user comfort. Finally, we advise against the separation of position and elevation specification into two distinct techniques as the main travel method, but we recommend considering elevation changes on the spot as a complement to one of the other two approaches.

7 CONCLUSION AND FUTURE WORK

While conventional ground-based navigation is sufficient for many use cases, the ability to teleport to targets beyond the vicinity of scene objects can offer novel perspectives on and a more efficient traversal of the environment to meet more complex navigational demands. We

believe that mid-air teleportation techniques should work complementarily to the already established mechanisms for teleportation to add expressivity while building upon the familiarity of previously learned interaction workflows. In this paper, we presented and evaluated three novel approaches in this regard. Our results demonstrated that all three techniques were learnable and usable after a short training phase, yet they differed significantly on several of the recorded measures. In short, *Simultaneous* was the most accurate technique when it was required, but this also led to longer specification times and task load. *Two-Step* was less accurate but showed several advantages in terms of specification time, task load, usability, and preference. Finally, *Separate* appears to be most suited as a complement to one of the other two techniques. Based on these as well as previous results, we formulated initial design guidelines for mid-air navigation. We hope that these guidelines serve as a solid basis for researchers and developers to be further validated and refined in future work.

Overall, our overview of related work in this paper has shown that research on mid-air navigation is still at the beginning. Future work should therefore study more variations of the presented techniques, especially for different controller types, more specific usage scenarios, and a larger user base including absolute novices in virtual reality. Furthermore, a more in-depth analysis of the fear of (virtual) heights during mid-air navigation is necessary to build upon the initial mitigation ideas proposed by Medeiros et al. (floor proxy and self-avatar) that we adapted in our designs [31]. Another interesting design challenge for future research is the comprehensible integration of virtual rotation as well as elevation change into a single mid-air navigation technique, which would be particularly beneficial for seated or otherwise spatially constrained users. Finally, the interplay of mid-air teleportation and mid-air steering techniques for collaborators in multi-user virtual environments with different navigation preferences is an interesting aspect to be addressed in future work. All in all, we believe that future research on mid-air navigation will complement the large body of existing work on ground-based navigation and therefore lead to a richer set of design recommendations for effective, efficient, and comfortable navigation in virtual reality.

ACKNOWLEDGMENTS

We would like to thank *Daniel Rupp* for his diligent assistance in conducting a large part of the user studies for this paper. This work has mainly received funding from the Ministry of Economic Affairs, Industry, Climate Action and Energy of the State of North Rhine-Westphalia under grant *005-2108-0055* (project *VITAMINE_5G*). This work was also partially funded by the German Ministry of Education and Research (BMBF) under grant *16SV8716* (project *Goethe-Live-3D*) and the Thuringian Ministry of Economic Affairs, Science and Digital Society (TMWDDG) under grant *5575/10-5* (project *MetaReal*).

REFERENCES

- [1] M. Al Zayer, P. MacNeilage, and E. Folmer. Virtual Locomotion: A Survey. *IEEE Transactions on Visualization and Computer Graphics*, 26(6):2315–2334, 2020. doi: 10.1109/TVCG.2018.2887379 2
- [2] P. Bimberg, T. Weissker, and A. Kulik. On the Usage of the Simulator Sickness Questionnaire for Virtual Reality Research. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops*, pp. 464–467, 2020. doi: 10.1109/VRW50115.2020.00098 5
- [3] P. Bimberg, T. Weissker, A. Kulik, and B. Froehlich. Virtual Rotations for Maneuvering in Immersive Virtual Environments. In *Proceedings of the 27th ACM Symposium on Virtual Reality Software and Technology*, 2021. doi: 10.1145/3489849.3489893 1, 2, 9
- [4] D. A. Bowman, E. Kruijff, J. J. LaViola, and I. Poupyrev. An Introduction to 3-D User Interface Design. *Presence*, 10(1):96–108, 2001. doi: 10.1162/105474601750182342 2
- [5] E. Bozgeyikli, A. Raij, S. Katkooi, and R. Dubey. Point & Teleport Locomotion Technique for Virtual Reality. In *Proceedings of the 2016 Annual Symposium on Computer-Human Interaction in Play*, pp. 205—216, 2016. doi: 10.1145/2967934.2968105 1, 2
- [6] Z. Cao, J. Jerald, and R. Kopper. Visually-Induced Motion Sickness Reduction via Static and Dynamic Rest Frames. In *2018 IEEE Conference*

- on *Virtual Reality and 3D User Interfaces*, pp. 105–112, 2018. doi: 10.1109/VR.2018.8446210 4
- [7] D. W. Carruth. Virtual Reality for Education and Workforce Training. In *2017 15th International Conference on Emerging eLearning Technologies and Applications*, 2017. doi: 10.1109/ICETA.2017.8102472 2
 - [8] S. Chen, F. Miranda, M. Ferreira, M. Lage, H. Doraiswamy, C. Brenner, C. Defanti, M. Koutsoubis, L. Wilson, K. Perlin, and C. T. Silva. UrbanRama: Navigating Cities in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics (Early Access)*, 2021. doi: 10.1109/TVCG.2021.3099012 2
 - [9] C. G. Christou and P. Aristidou. Steering Versus Teleport Locomotion for Head Mounted Displays. In *International Conference on Augmented Reality, Virtual Reality and Computer Graphics*, pp. 431–446. Springer, 2017. doi: 10.1007/978-3-319-60928-7_37 1, 2, 9
 - [10] J. Clifton and S. Palmisano. Effects of Steering Locomotion and Teleporting on Cybersickness and Presence in HMD-based Virtual Reality. *Virtual Reality*, 24(3):453–468, 2020. doi: 10.1007/s10055-019-00407-8 1, 2, 9
 - [11] S. Cmentowski, A. Krehhov, and J. Krüger. Outstanding: A Multi-Perspective Travel Approach for Virtual Reality Games. In *Proceedings of the Annual Symposium on Computer-Human Interaction in Play*, pp. 287—299, 2019. doi: 10.1145/3311350.3347183 2
 - [12] J. Cohen. *Statistical Power Analysis for the Behavioral Sciences*. Routledge, second ed., 2013. doi: 10.4324/9780203771587 6
 - [13] M. Di Luca, H. Seifi, S. Egan, and M. Gonzalez-Franco. Locomotion Vault: The Extra Mile in Analyzing VR Locomotion Techniques. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 2021. doi: 10.1145/3411764.3445319 2
 - [14] J. L. Dorado and P. A. Figueroa. Ramps are Better than Stairs to Reduce Cybersickness in Applications Based on a HMD and a Gamepad. In *2014 IEEE Symposium on 3D User Interfaces*, pp. 47–50, 2014. doi: 10.1109/3DUI.2014.6798841 2
 - [15] A. Drogemuller, A. Cunningham, J. Walsh, M. Cordeil, W. Ross, and B. Thomas. Evaluating Navigation Techniques for 3D Graph Visualizations in Virtual Reality. In *2018 International Symposium on Big Data Visual and Immersive Analytics*, 2018. doi: 10.1109/BDVA.2018.8533895 2, 9
 - [16] C. Elvezio, M. Sukan, S. Feiner, and B. Tversky. Travel in Large-Scale Head-Worn VR: Pre-oriented Teleportation with WIMs and Previews. In *2017 IEEE Virtual Reality*, pp. 475–476, 2017. doi: 10.1109/VR.2017.7892386 3
 - [17] Y. Farmani and R. J. Teather. Evaluating discrete viewpoint control to reduce cybersickness in virtual reality. *Virtual Reality*, 24(4):645–664, 2020. doi: 10.1007/s10055-020-00425-x 1, 2, 9
 - [18] 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*, pp. 201–210, 2016. doi: 10.1109/3DUI.2016.7460053 4
 - [19] A. Field. *Discovering Statistics Using IBM SPSS Statistics*. Sage Publications, 4th ed., 2013. 6
 - [20] M. Funk, F. Müller, M. Fendrich, M. Shene, M. Kolvenbach, N. Döbertin, S. Günther, and M. Mühlhäuser. Assessing the Accuracy of Point & Teleport Locomotion with Orientation Indication for Virtual Reality Using Curved Trajectories. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 2019. doi: 10.1145/3290605.3300377 1, 2
 - [21] N. N. Griffin and E. Folmer. Out-of-Body Locomotion: Vectionless Navigation with a Continuous Avatar Representation. In *Proceedings of the 25th ACM Symposium on Virtual Reality Software and Technology*, 2019. doi: 10.1145/3359996.3364243 2
 - [22] T. Grossman and R. Balakrishnan. The Design and Evaluation of Selection Techniques for 3D Volumetric Displays. In *Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology*, pp. 3—12, 2006. doi: 10.1145/1166253.1166257 2
 - [23] S. G. Hart. Nasa-Task Load Index (NASA-TLX); 20 Years Later. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 50, pp. 904–908, 2006. 5
 - [24] S. G. Hart and L. E. Staveland. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. In *Advances in Psychology*, vol. 52, pp. 139–183. Elsevier, 1988. doi: 10.1016/S0166-4115(08)62386-9 5
 - [25] S. Knödel, M. Hachet, and P. Guittou. Navidget for Immersive Virtual Environments. In *Proceedings of the 2008 ACM Symposium on Virtual Reality Software and Technology*, pp. 47—50, 2008. doi: 10.1145/1450579.1450589 1, 2
 - [26] S. Kuliga, T. Thrash, R. Dalton, and C. Hölscher. Virtual reality as an empirical research tool — Exploring user experience in a real building and a corresponding virtual model. *Computers, Environment and Urban Systems*, 54:363–375, 2015. doi: 10.1016/j.compenvurbysys.2015.09.006 2
 - [27] C. Lai, R. P. McMahan, and J. Hall. March-and-Reach: A realistic ladder climbing technique. In *2015 IEEE Symposium on 3D User Interfaces*, pp. 15–18, 2015. doi: 10.1109/3DUI.2015.7131719 2
 - [28] B. Laugwitz, T. Held, and M. Schrepp. Construction and Evaluation of a User Experience Questionnaire. In *HCI and Usability for Education and Work, 4th Symposium of the Workgroup Human-Computer Interaction and Usability Engineering of the Austrian Computer Society*, pp. 63–76. Springer, 2008. 5
 - [29] J.-I. Lee, P. Asente, B. Kim, Y. Kim, and W. Stuerzlinger. Evaluating Automatic Parameter Control Methods for Locomotion in Multiscale Virtual Environments. In *26th ACM Symposium on Virtual Reality Software and Technology*, 2020. doi: 10.1145/3385956.3418961 2
 - [30] A. Matvienko, F. Müller, M. Schmitz, M. Fendrich, and M. Mühlhäuser. SkyPort: Investigating 3D Teleportation Methods in Virtual Environments. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, 2022. doi: 10.1145/3491102.3501983 2, 9
 - [31] D. Medeiros, M. Sousa, A. Raposo, and J. Jorge. Magic Carpet: Interaction Fidelity for Flying in VR. *IEEE Transactions on Visualization and Computer Graphics*, 26(9):2793–2804, 2020. doi: 10.1109/TVCG.2019.2905200 2, 3, 9
 - [32] R. Pausch, J. Snoddy, R. Taylor, S. Watson, and E. Haseltine. Disney’s Aladdin: First Steps toward Storytelling in Virtual Reality. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*, pp. 193—203, 1996. doi: 10.1145/237170.237257 2
 - [33] L. Rebenitsch and C. Owen. Individual Variation in Susceptibility to Cybersickness. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology*, pp. 309—317, 2014. doi: 10.1145/2642918.2647394 5
 - [34] L. Rebenitsch and C. Owen. Review on cybersickness in applications and visual displays. *Virtual Reality*, 20(2):101–125, 2016. 2, 6
 - [35] B. E. Riecke, D. Clement, A. Adhikari, D. Quesnel, D. Zielasko, and M. von der Heyde. HyperJumping in Virtual Vancouver: Combating Motion Sickness by Merging Teleporting and Continuous VR Locomotion in an Embodied Hands-Free VR Flying Paradigm. In *ACM SIGGRAPH 2022 Immersive Pavilion*, 2022. doi: 10.1145/3532834.3536211 2
 - [36] W. Robinett and R. Holloway. Implementation of Flying, Scaling and Grabbing in Virtual Worlds. In *Proceedings of the 1992 Symposium on Interactive 3D Graphics*, pp. 189—192, 1992. doi: 10.1145/147156.147201 2
 - [37] R. A. Ruddle and S. Lessels. The Benefits of Using a Walking Interface to Navigate Virtual Environments. 16(1), 2009. doi: 10.1145/1502800.1502805 2
 - [38] S. P. Sargunam and E. D. Ragan. Evaluating Joystick Control for View Rotation in Virtual Reality with Continuous Turning, Discrete Turning, and Field-of-View Reduction. In *Proceedings of the 3rd International Workshop on Interactive and Spatial Computing*, pp. 74—79, 2018. doi: 10.1145/3191801.3191815 1, 2
 - [39] M. Slater, M. Usoh, and A. Steed. Steps and Ladders in Virtual Reality. In *Virtual Reality Software And Technology*, pp. 45–54. World Scientific, 1994. 2
 - [40] R. Stoakley, M. J. Conway, and R. Pausch. Virtual Reality on a WIM: Interactive Worlds in Miniature. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 265—272, 1995. doi: 10.1145/223904.223938 5
 - [41] E. A. Suma, S. Babu, and L. F. Hodges. Comparison of Travel Techniques in a Complex, Multi-Level 3D Environment. In *2007 IEEE Symposium on 3D User Interfaces*, 2007. doi: 10.1109/3DUI.2007.340788 2
 - [42] M. Usoh, K. Arthur, M. C. Whitton, R. Bastos, A. Steed, M. Slater, and F. P. Brooks. Walking >Walking-in-Place >Flying, in Virtual Environments. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, pp. 359—364, 1999. doi: 10.1145/311535.311589 2
 - [43] K. Vasylyeva, H. Kaufmann, and V. Khrystyna. Influence of Vertical Navigation Metaphors on Presence. In *Challenging Presence-Proceedings of 15th International Conference on Presence*, pp. 205–212, 2014. doi: 10.500.12708/55149 2
 - [44] T. Weissker and B. Froehlich. Group Navigation for Guided Tours in Distributed Virtual Environments. *IEEE Transactions on Visualization*

and *Computer Graphics*, 27(5):2524–2534, 2021. doi: [10.1109/TVCG.2021.3067756](https://doi.org/10.1109/TVCG.2021.3067756) 1, 2, 3

- [45] T. Weissker, A. Kunert, B. Froehlich, 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*, pp. 97–104, 2018. doi: [10.1109/VR.2018.8446620](https://doi.org/10.1109/VR.2018.8446620) 1, 2, 6, 9
- [46] D. Wolf, M. Rietzler, L. Bottner, and E. Rukzio. Augmenting Teleportation in Virtual Reality With Discrete Rotation Angles. *CoRR*, June 2021, 2021. doi: [10.48550/arXiv.2106.04257](https://doi.org/10.48550/arXiv.2106.04257) 1, 2
- [47] C. Zambaka, B. Lok, S. Babu, A. Ulinski, and L. Hodges. Comparison of Path Visualizations and Cognitive Measures Relative to Travel Technique in a Virtual Environment. *IEEE Transactions on Visualization and Computer Graphics*, 11(6):694–705, 2005. doi: [10.1109/TVCG.2005.922](https://doi.org/10.1109/TVCG.2005.922)
- [48] Y. Zhang, N. Ladeveze, H. Nguyen, C. Fleury, and P. Bourdot. Virtual Navigation Considering User Workspace: Automatic and Manual Positioning before Teleportation. In *26th ACM Symposium on Virtual Reality Software and Technology*, 2020. doi: [10.1145/3385956.3418949](https://doi.org/10.1145/3385956.3418949) 3
- [49] D. Zielasko, J. Heib, and B. Weyers. Systematic Design Space Exploration of Discrete Virtual Rotations in VR. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces*, pp. 693–702, 2022. doi: [10.1109/VR51125.2022.00090](https://doi.org/10.1109/VR51125.2022.00090) 1, 2