Designing Viewpoint Transition Techniques in Multiscale Virtual Environments

Jong-In Lee*
Simon Fraser University

Paul Asente†
Adobe Research, retired

Wolfgang Stuerzlinger‡
Simon Fraser University

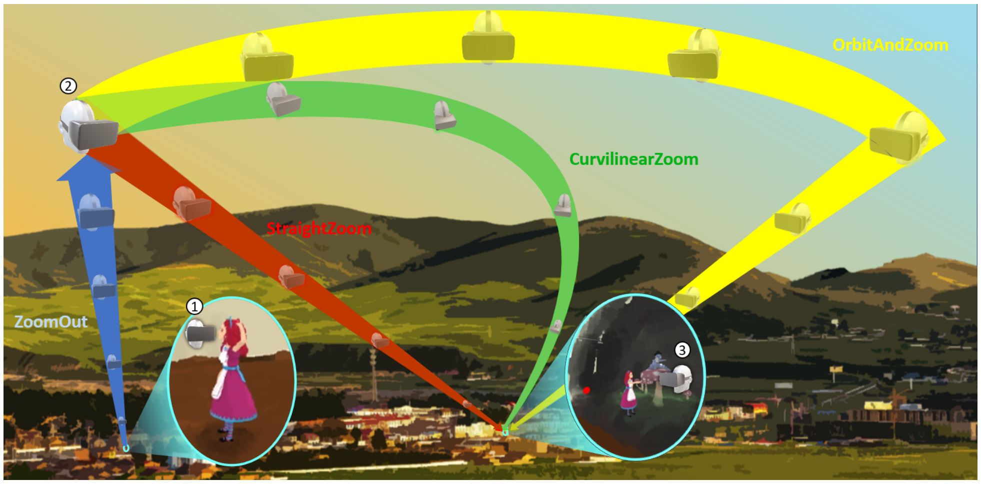

Figure 1: Three design variations of mixed viewpoint transition techniques (StraightZoom, CurvilinearZoom, and OrbitAndZoom) which first zoom-out from 1) initial viewpoint to 2) a larger Level-Of-Scale (LoS) and then zoom in to 3) another LoS. All variations zoom out in the same way but zoom in along different paths.

ABSTRACT

Viewpoint transitions have been shown to improve users' spatial orientation and help them build a cognitive map when they are navigating an unfamiliar virtual environment. Previous work has investigated transitions in single-scale virtual environments, focusing on trajectories and continuity. We extend this work with an in-depth investigation of transition techniques in multiscale virtual environments (MVEs). We identify challenges in navigating MVEs with nested structures and assess how different transition techniques affect spatial understanding and usability. Through two user studies, we investigated transition trajectories, interactive control of transition movement, and speed modulation in a nested MVE. We show that some types of viewpoint transitions enhance users' spatial awareness and confidence in their spatial orientation and reduce the need to revisit a target point of interest multiple times.

Index Terms: Human-centered computing—Human computer interaction (HCI)—Interaction techniques

1 INTRODUCTION

To enable a user to travel within a large virtual environment (VE) within a limited physical tracking space in Virtual Reality (VR), different locomotion and navigation interfaces have been developed in VR research and games [17]. For this, many steering and target-based locomotion techniques have been designed. Steering techniques let users continuously adjust the direction and speed, while physically staying stationary [38]. The most prominent example is the flying technique, in which a user controls the flying direction and speed using controllers or various body parts [18, 31, 48, 51, 60]. Target-based techniques, on the other hand, let users specify the target position and, sometimes, orientation before executing a teleportation to instantly move them there [38]. With target-based techniques, users can either specify the target in their surroundings through simple pointing or through other means if the target location is not visible from their current position.

Many current VR applications use point-and-teleport to let users specify a visible target location in their surroundings and quickly move there. This technique reduces motion sickness, as it changes the user's viewpoint instantly rather than using continuous movement like steering-based techniques [11, 24], but this instantaneous change in the user's view can cause disorientation [39, 44]. To mitigate this, researchers have integrated viewpoint transition methods into target-based navigation techniques, which automatically move a user from an initial to a target view through an animation. Studies demonstrated that viewpoint transitions facilitate users' path integration and help maintain spatial orientation [8, 39, 47]. Other research developed transition methods for target-based techniques that enable travel to targets *beyond* their surroundings, such as World-In-Miniature (WIM) [35] and bookmark interfaces [59], which let users specify travel targets in the VE by interacting with a 3D map or selecting a bookmark icon, respectively. Still, there is currently no in-depth study that evaluated such transition methods in large and/or multiscale VEs. Multiscale virtual environments (MVEs), especially with nested structures, have recently become more common in VR applications, including geographical simulations [25], medical visualizations [35, 41], cosmological applications [19, 64], and VR painting software [26, 32, 53]. For navigating in nested MVEs, viewpoint transitions are crucial. Targets may be invisible due to structures or walls blocking the line of sight, or they may simply be too small (or sometimes even too large) to be visible at

---

*e-mail: jonginl@sfu.ca
†e-mail: research@asente.com
‡e-mail: w.s@sfu.ca

the current LoS or within the current field of view.

We performed two user studies on novel viewpoint transition methods in MVEs with nested structures, with 16 and 18 participants, respectively. In these studies, we compare variations of three key design components of viewpoint transitions: the transition trajectory, the interactive control of the timing of the transition, and the transition speed. For the transition trajectory, we explore different trajectory options in study 1 and 2 and evaluate how they affect users' spatial understanding and their overall experience. For interactive control of the timing of the transition and the transition speed, we compare conditions with or without a timing control and automatic speed modulation in study 1. Finally, based on the findings from the two user studies, we identify the design options that lead to the best navigation performance and usability.

2 RELATED WORK

In this section we review previous work on target-based navigation interfaces and viewpoint transition techniques in VEs.

2.1 Target-Based Navigation in Virtual Environments

Since teleportation was first studied in a virtual environment by Bowman et al. [10], it has been widely investigated in 3D and VR user interfaces. Teleportation instantly moves a user to a new location. Point-and-teleport [11] is the most prominent teleportation technique, used frequently in VR games and applications. In previous work, researchers found that instantly changing a user's viewpoint can lead to disorientation [9–11, 39, 44]. To mitigate this, researchers and game designers developed viewpoint transitions for point-and-teleport, which moves the user's viewpoint continuously to the target instead of instantly teleporting there [8, 47, 57, 61, 62]. While point-and-teleport lets users freely select a visible target in their surroundings, it is limited to travel within a local area since the target must be visible and close enough to be selected accurately. To enable teleportation to out-of-sight or distant targets, several researchers have proposed alternative interfaces. The AutoTeleport technique lets users point in the direction of a distant target, but the distance of the jumps is automatically controlled to prevent under- or overshooting [39]. WIMs allow users to specify a distant target location by interacting with a miniature widget [20, 21, 35, 37]. Bookmark interface allows users to select a predefined bookmark to instantly travel to the corresponding location [7, 55, 59].

Rahimi et al. [54] studied the effect of the trajectory and the continuity of viewpoint transition on usability and spatial orientation, the transition methods they studied were limited to visible targets in a small area at a single scale. Beyond that, few studies have investigated the effect of viewpoint transitions for target-based navigation to out-of-sight targets. For MVEs with a nested structure, Kopper et al. [35] and Bacim et al. [5] developed transition methods for target-based navigation. While they incorporated transition methods into target-based navigation interfaces for MVEs, it only interpolated position and scale to transition from the current to the target Level-of-Scale (LoS). Moreover, there was no discussion about the importance of viewpoint transitions while navigating between different LoS, nor an evaluation of the transitions per se.

2.2 Viewpoint Transition Techniques

Several viewpoint transition techniques for different use cases and scenarios have been presented. Depending on the type of target viewpoint, these can be divided into two categories: transitioning to user-defined and to pre-defined target viewpoints.

Viewpoint transition techniques for *user-specified* viewpoints are often integrated into a navigation interface. Buchholz et al.'s technique [12] smoothly transitions between user-controlled locomotion and physics-based automatic movement. Hachet et al. [28] developed a transition technique that automatically moves and rotates the camera towards the target view specified by the user. Target-based locomotion interfaces by Medeiros et. al. [47] and Cmentowski et al. [16] also provided viewpoint transition upon target specification. Others developed hybrid techniques that transition the view between different navigation modes. Griffin et al. [27] developed a technique that shows a viewpoint transition upon switching from first to third-person's perspectives to control the user's avatar. Krekhov et al. [36] used a similar approach for natural walking-based locomotion.

Transition techniques for *pre-defined* target viewpoints are usually presented through a bookmark interface that lets a user browse saved viewpoints and select one to visit. Veas et al. [63] presented an interface that presents a list of views with different perspectives and a transition technique that first moves the camera up to an overview and then smoothly interpolates to the selected view. Similarly, Sukan et al. [58] developed a technique that quickly interpolates the current view to the target view of a selected bookmark in an indoor AR application. Both techniques simply interpolate directly to a viewpoint, similar to the user-specified target view techniques.

Yet, pre-defined viewpoints are sometimes far away and may thus not even be visible. Simple interpolation-based transitions, as presented in previous work, do not support the user in maintaining their spatial orientation during the transition, as they will typically rotate the user while moving, which makes it hard to maintain spatial orientation and can induce motion sickness. The "JackIn" technique [33] addresses this issue by first approaching the target viewpoint, then orbiting around it to match the desired orientation, and finally moving closer to the target from that direction, but this technique was not evaluated in a study. Similarly, and even though Moghadam and Ragan [52] compared three viewpoint transitions in terms of spatial awareness and simulator sickness, they only considered transitions within a small area at a single LoS. Here, we present new viewpoint transition techniques between saved viewpoints across different LoS and investigate what properties of these transitions affect the spatial awareness and user experience in MVEs.

Also, most techniques do not give the user a chance to stop or rewind a transition, e.g., to let users check their spatial understanding. Burtnyk et al. [13, 14] let the viewer scrub back and forth through the transition animation with a mouse. Similar to their approach, we designed a touch-based input method to interactively control the transition, but extend their work towards MVEs where the transition not only affects the position and orientation but also the scale.

3 VIEWPOINT TRANSITION TECHNIQUES FOR MULTISCALE ENVIRONMENTS

We developed variations for three core design components of viewpoint transition techniques: 1) transition trajectory, 2) transition timing control, and 3) transition speed modulation. The following subsections describe each component and the multiscale bookmark interface we designed for user studies.

3.1 Different Types of Transitions

When designing an MVE, the designer also implicitly creates a hierarchical structure in which smaller places, such as a house, are nested within larger places, such as a village. To support ease of understanding and effective navigation, the designer can create a hierarchically-structured map (HiSMap) [5], where the designer specifies all important places and which place nests within which other place. This HiSMap then allows the system to determine what type of transition is to be executed when the user selects the place to travel to, either by selecting the node directly in the HiSMap or, like in our study, through a bookmark icon in a list (Fig. 3(a)). The possible options for viewpoint transitions then include zoom-in, zoom-out, and mixed transitions.

Zoom-in transitions move the user from the current LoS to a smaller one, often across multiple LoS, (see the yellow line in Fig. 2). Previous work [40] developed multiple variations, with a variant, called ORBITANDZOOM that first orbits around the target until

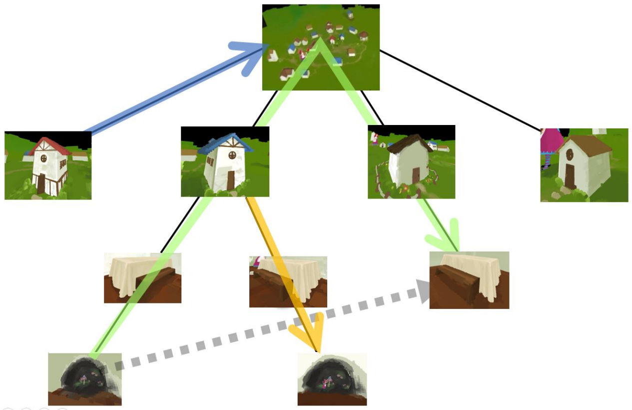

Figure 2: An example LoS hierarchy for a MVE with 4 LoS: The root node represents the largest LoS (i.e., the village). The two bottom leaf nodes represent the smallest LoS (i.e., the caterpillar lairs). Edges describe the associated spatial nesting hierarchy, e.g., that the two left tables, are nested in the blue house. The colored and dashed arrows show different types of transitions.

the current view direction matches the target orientation, and then zooms in by gradually translating and scaling the view to the target performing best. We use OrbitAndZoom for User Study 1 and motivate this choice further in Sect. 4.1.

In contrast, zoom-out transitions go from the current LoS to a larger LoS (see the blue line in Fig. 2). Here, users do not need to choose a target, the system instead brings the user automatically to a system-defined vantage point. We experimented with different variations, but an informal pilot with colleagues showed that simultaneously rotating and zooming-out in transitions led to severe disorientation and confusion for the user. As an example, consider passengers in a moving vehicle, but facing backwards, who will lose track of where they are heading as the vehicle turns and have difficulty orienting themselves. Due to the lack of vestibular and proprioceptive feedback, this can be even worse in VR. We then tried a zoom-out technique that first rotated the view direction to align with the view direction of a predetermined overview bookmark at the largest LoS and then started position and scale interpolation, but this approach only let users see the destination at the end, without any meaningful information around it. We thus developed a new zoom-out technique with no initial rotation phase, immediately zooming out backwards to an automatically computed vantage point. That vantage point is determined by moving the user backwards to a position that is at the higher LoS and behind and above the user's head (see the blue line and user's head at (1) in Fig. 1). This approach ensures that the user always faces the initial location while zooming out backward to the overview at the larger LoS.

Mixed transitions first zoom out to the lowest common ancestor of the current and the target LoS and then zoom into the target (see the green line in Fig. 2). We initially tested a simple interpolation of position, orientation, and scale, but found this to be confusing because it often caused users to pass through the surfaces of multiple objects without useful visual clues for the spatial context (see the grey dashed line in Fig. 2). We then designed three mixed transition techniques that first zoom out to the common parent LoS to let users see both the initial and target LoS from an overview, and then zoom in to the target LoS. All three techniques use the zoom-out technique described above, but use different zoom-in approaches. Fig. 1 shows these as yellow, green, and red lines. See Sect. 5.1 for details.

Beyond the three types of viewpoint transition mentioned above, there are other options for transitions. Linear interpolation while disregarding the spatial context is the simplest. Yet, this approach is problematic in an MVE cluttered with a large number of objects and complex structures, as a linear transition would only provide optical flow, but not enough information to help users gain spatial understanding and awareness. Another option is to utilize route planning algorithms to determine a transition trajectory that prevents collision with objects [46]. Yet, in MVEs with a nested structure, where continuous surfaces completely enclose a part of the scene, such an approach is insufficient because getting into or out of a location that is encapsulated by one or more surfaces is then impossible.

3.2 Interactive Control of Transition Timing

We also implemented a method that lets the user interactively control the progress of the transition using a touchpad on a controller. As previous work had identified active control to be beneficial [45] but potentially can cause cognitive overhead [15], we designed a simple transition control method that lets a user actively control the progress and pace of the transition with minimum overhead by simply sliding their thumb on the touchpad. The user can advance the transition by moving towards the right side and rewind the transition by moving towards the left. The distance from the center to the touchpoint directly determines the "play" speed of the transition, similar to a video scrubbing interface [43]. Our user study compared this active control to passive observation to understand how it affects spatial updating and usability.

3.3 Adaptive Modulation of Transition Speed

Extending the constant transition speed explored in previous work [13, 14, 29, 33, 54, 58, 70], we created a new method that modulates the speed based on the distance to objects. Previous work [40] had observed that users felt uncomfortable when moving slowly through a wall, even in the presence of a cut-away visualization of the target location, as it blocks most of their view except for a small hole open in front of the target, which gave them an impression of being trapped inside the wall, negatively impacting the user experience. Inspired by proximity-based speed control [4, 67] and distance field-based navigation [39], but using the distance measure in an *inverse* way, we developed a new transition speed modulation method. It accelerates transitions when near or inside an object, e.g., a wall or a roof, to make the viewpoint quickly pass through such obstacles during transition (Fig. S.1 in the supplemental material). We started with a simple linear transition:

$$S_{n+1} = S_i + P_n * (S_t - S_i) \tag{1}$$

Equation (1) describes the scene scale in the next step S_{n+1}, as determined by the progress variable P_n. The value of P_n is 0 at the beginning of the transition and 1 at the end, which means that the current scale S_n is equal to the initial scale S_i at the beginning and to the target scale S_t at the end of the transition.

$$P_{n+1} = P_n + \frac{S_n}{S_t} g(n) \tag{2}$$

The progress for the next step P_{n+1} is the result of an accumulation of the second term, based on the current scale to target scale ratio and the gain value computed by equation (3). The use of the ratio slows down progress when the current scale gets relatively bigger and vice versa. This makes the objects in the scene look as if they approach or move away at a constant speed when the scene expands or shrinks around the bookmark position.

$$g(n) = 1 - (f(\mathbf{a}_n) - k \cdot \mathbf{d} \cdot \nabla f(\mathbf{a}_n))/S_n \tag{3}$$

The transition gain $g(n)$ is based on the proximity to the surface of the closest object. $f(\mathbf{a}_n)$ is the value of the distance field at the current position $\mathbf{a}_n$, $\nabla f(\mathbf{a}_n)$ the gradient, and k is a weight that controls how much the field gradient affects the transition speed. We performed a pilot study with three colleagues to determine the optimal

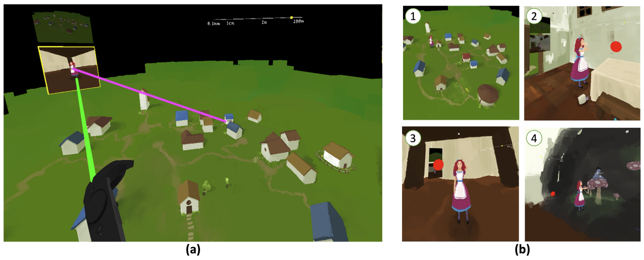

Figure 3: (a) Our bookmark interface allow user selects a thumbnail icon using raycasting. A magenta line connecting the selected bookmark to the location of the corresponding target viewpoint indicates the transition destination. (b) Alice at different LoS: 1) Overview with Alice at the 1st LoS, 2) Alice at the 2nd LoS, 3)) Alice at the 3rd LoS, and 4) Alice at the 4th LoS.

k value for zoom-in transition, where we compared k values of 0.25, 0.5, 1, and 2 and found $k = 1$ works best, as it starts accelerating the transition before colliding with a surface, and slows it down before moving out of an object. The acceleration was too small to notice when k was smaller than 1 and when k as larger than 1, the speed was too fast to potentially stop before moving into the next smaller LoS. Then, the user's view direction $\mathbf{d}$ influences the final value of $g(n)$, increasing when the view direction aligns with the gradient vector and decreasing when it is the opposite. By negating the sum of the distance field value and the inner product of view direction and the gradient, this method "pulls" the user into a wall more quickly while the transition approaches it. After entering the wall, it also prevents them from being "launched" out the wall's other side. This still minimizes the time when inside a wall, which obscures the view, yet maximizes the time for viewing the scene without obstruction. As mentioned above, we hypothesize that our speed modulation feature reduces visual discomfort and frustration. In Sect. 4, we test this hypothesis by comparing our adaptive speed transition technique with a constant speed version, with $g(n)$ a constant.

3.4 Multiscale Bookmark interface

A bookmark interface allows users to browse through saved viewpoints and revisit them by selecting the corresponding icon with raycasting (Fig. 3(a)). Similar to TrailMap [72], our interface displays bookmark icons as distinct snapshots of the places where the bookmarks were created. If different bookmarks show similar-looking objects or spatial configurations, snapshots alone might be ambiguous. To address this issue, and inspired by previous work [50], we display a straight magenta line connecting the selected icon to the saved viewpoint (Fig. 3). When the user first selects the icon, its frame is highlighted in yellow and the connecting line appears, with a cut-away visualization [22] where a hole is cut out between the current and target viewpoint not just to reveal the hidden point of interest (POI) but also to reduce the discomfort while passing through walls, see the zoomed-in view of (3)-(6) in Fig. S.1 in the supplemental material. Then, when the icon is selected again, the frame changes to red and the transition begins.

4 USER STUDY 1: ZOOM-IN TRANSITION

We conducted a user study to test two main hypotheses:

- **H1:** Giving users interactive control of viewpoint transitions in MVEs helps them to understand the space better
- **H2:** Accelerating transition speed while passing through walls improves the user experience of viewpoint transition in MVEs

We base **H1** on previous work that showed that active navigation improves navigation performance compared with passive viewing of automatically generated tours or movements controlled by others in VEs [6, 15]. **H2** is based on user feedback in prior work [40], where participants felt discomfort while passing through walls during a zoom-in transition and wanted such moments to be minimized.

4.1 Motivation for Trajectory Design

In multiscale 2D map applications, such as Google Maps, when the user selects a saved place or presses the "my location" button, the view smoothly pans and zooms from an overview to a view where the corresponding location appears in the center of the screen at a reasonable scale. These continuous transitions in translation and scale help users understand both the relative position between two places and the relative LoS in the LoS hierarchy of the scene. For instance, continuously panning and zooming from an overview of the Pacific ocean into a view of a local park in London illustrates not only where the area is located relative to the ocean but also how much the overview has to zoom in to get to the local level, and similar for 3D MVEs. Besides panning and zooming, the view direction also needs to continuously rotate towards the orientation corresponds to the saved bookmark to prevent disorientation due to an instant rotation change. While simultaneous interpolation of position, rotation, and scale is feasible, previous work identified that simultaneous translation and rotation deteriorates spatial awareness [52, 54].

Hence, we decided to separate transitions into two phases, with an orbiting (rotation) and a zoom-in phase (translation + scaling), extending previous viewpoint control approaches [42] and a point-of-interest based technique [33] that use only rotation and translations. For this we follow previous work which developed ORBITANDZOOM [40], a technique that starts with orbiting at a constant distance around the bookmark, until the current view direction aligns with the target orientation, and then gradually translates and scales the viewpoint to the target location and LoS (Fig. S.1 in the supplemental material). Such an ORBITANDZOOM transition helps the user to stay oriented [40] better than other alternatives. Another way to separate transitions into rotation and zoom-in would be to first zoom in and then to rotate at the target location to match the target orientation. Yet, previous work found that the rotation at the small LoS causes a large amount of discomfort and disorientation [40].

4.2 Experimental Design

We used a two factor within-subject design. The first variable was the *Control* method for transition: PASSIVE and INTERACTIVE. The second was the *Speed* modulation method: CONSTANT and ADAPTIVE. For PASSIVE, users only had to select a bookmark to initiate automatic transition with the controller. With INTERACTIVE, users could control the direction of the animation of the transition by pressing the right part of the touchpad on the controller to advance it and the left part to rewind it. CONSTANT adapts the transition speed proportionally to users' scale, speeding up at a larger LoS and slowing down at a smaller LoS. As shown in previous work [34, 71], basing the speed on the current scale lets users pass big objects faster at a larger LoS and move slower at a smaller LoS, letting them see the details that enable spatial understanding at the corresponding scale. ADAPTIVE uses the same basic speed calculation as CONSTANT but adds acceleration near obstacles as described in Sect. 3.3.

4.3 Participants

We recruited 16 participants from the local university, seven male and five female, and the average age was 26.7 (SD=5.21). All had experience with 3D games and were compensated with 15 CAD.

4.4 Apparatus and Environment

Our experiment was conducted with an HTC Vive HMD, with 1080 × 1200 pixels for each eye, and Vive controllers for input. The experiment was performed in a space of 3 m × 3 m, spacious enough to let participants turn freely and perform controller-based navigation in any direction. To evaluate transition control and speed modulation methods, we designed a MVE based on the narrative and character design of *Alice in Wonderland* by Lewis Carroll [68]. Using the Canvox system[1], we sketched a small town in which several instances of the main character Alice, scaled to different sizes, are present at different locations. Each size of Alice represents a LoS that users need to navigate to. The biggest one at the 1st LoS was 10 m tall, and the ratios of the instances at the 1st, 2nd, 3rd, and 4th LoS were 512:64:8:1.

The bookmark interface was inspired by Zhao et al. [72] displaying a list of icons that shows thumbnails of places where the bookmarks were saved (Fig. 3(a)). For interactive navigation, we adopted a bi-manual technique from previous work [56], which separates panning and zooming functions onto each controller. We initially implemented zooming similarly Satriadi et al. [56], where the scaling center coincides with the controller. However, in an informal test with lab colleagues, most people struggled with specifying the correct scaling center to zoom into a house. They frequently specified the scaling center in front of or behind the house. To address this issue, we evaluated different scaling centers, such as a position in front of the user's head or the center between the two controllers in a subsequent pilot with four participants. We found that the position in front of the user was most reliable and let users get into the house easier and quicker, which let users complete NAVIGATION four times faster than the initial design.

4.5 Task

There were two main tasks: BOOKMARK and NAVIGATION. The first task, BOOKMARK, was to start with a scene overview at the largest LoS of the environment and visit a bookmark showing a red target located near Alice at a smaller LoS. When the user selected the bookmark, a transition according to the current condition was initiated. During this period, participants were explicitly asked to pay attention to the shape, color, and layout of objects around the target, priming them to better understand the spatial context for the later transitions. They could teleport back to the overview by selecting the first bookmark icon and visit the target bookmark as many times as they needed to arrive at this understanding. We decided to give users the freedom of iterating as many times as needed (instead of limiting the number of transitions) to minimize the effect of the difference between participants' spatial ability on the results. Once they became familiar with the given structure and felt confident about locating the red target, they could move on to the second task, NAVIGATION, by pressing the menu button on the controller. Upon the menu button press, participants teleport back to the initial overview position.

When NAVIGATION began, the thumbnail of the target bookmark was displayed as a reference and participants were asked to move to this red target using bi-manual navigation. If they reached it within the time limit (60 seconds), we counted a successful trial, the system displayed a success message with a chime. If they could not reach it within time, the system played an error sound and displayed a failure message. These two tasks were repeated for different bookmarks with targets at different LoS.

4.6 Procedure

First, participants were asked to complete a brief background questionnaire. Then, they were asked to put the HMD on and to hold both controllers while standing in the middle of the experimental space. They were encouraged to rotate their body freely, but asked to stay roughly in the middle. Participants encountered the four conditions, (two *Control* conditions × two *Speed* modulation conditions) in counterbalanced order to reduce learning effects. In each condition participants were first instructed on using the current transition interface and conducted three practice tasks. As we noticed that different people have different preferences for overall transition speed, the experimenter manually adjusted a speed gain based on the participant's feedback during practice. We found this necessary to ensure every participant performed BOOKMARK comfortably at their own pace. Individual speed gains varied at most by 20%, which affected task completion time to a limited degree. Still, after the study, the data for navigation performance, such as the number of transitions, completion time, and navigation error were all normally distributed, which lets us compare the different conditions with inferential analysis. For practice tasks, the order of LoS at which the target appeared was randomized. Then, the three sets of main experimental tasks were presented, again in randomized order for LoS. After each condition, participants were asked to fill out a NASA task load index questionnaire (NASA-TLX) [30]. Then they experienced the other three experimental conditions in the order determined by counter-balancing. We used four similar-looking environments with different layouts with bookmarks at different locations (i.e., different locations of houses, objects, and Alice). All techniques had identical settings for the view parameters. The order of the four environments presented for each condition was randomly chosen. Once they were done with a condition, participants completed post-task questionnaires. For each questionnaire, we used a 7-point Likert scale with items ranging from "Strongly Disagree" to "Strongly Agree". Also, we performed semi-structured interviews asking about the advantages and disadvantages of each condition. The used questionnaires and interview questions are included in the supplemental material. The study lasted about 70 min on average.

4.7 Results

We analyzed the data for task completion time, error rate, and task load. We conducted inferential analysis using repeated-measures ANOVA with $\alpha = 0.05$ in R. If the data was not normally distributed we used ART [69] before ANOVA; see Table 1.

Number of transitions: For the BOOKMARK, we measured the number of transitions each participant triggered by selecting bookmarks. We counted any transitions that progressed from the beginning to the end as one incident of transition, which was trivial with PASSIVE. With INTERACTIVE, we counted one transition when the user played the transition forward all the way from the beginning until they reached the end. Once they got within a clear view of the target bookmark, we did not observe any participants rewinding all the way back to the beginning, which could have been counted as a totally new transition. We identified a significant effect of *Control* on the total number of transitions, with no effect of *Speed* nor an interaction between *Control* × *Speed*.

Task completion time: We analyzed completion times for each task type individually. ANOVA did not identify a significant effect of *Control* nor *Speed* on completion time for either task.

Pointing error: To analyze spatial orientation error, we first computed the average angle between the vector from the starting position to the target and the direction of the controller in the user's dominant hand for all the tasks and compared them between conditions. Two-way ANOVA did not identify a significant effect of *Control* nor *Speed* on pointing error.

User preferences: We measured user preferences for each condition using a 7-point Likert scale. We identified a significant effect of *Control* but no effect of *Speed* on preference. There was no interaction between *Control* × *Speed*.

Task Load Index: The two-way ANOVA did not identify a significant effect of *Control* nor *Speed* on the overall TLX score and most of the TLX sub categories. However, it identified a significant effect of *Control* on performance and effort scores. There was an interaction between *Control* × *Speed* on the effort score. Post-hoc tests with Tukey-HSD revealed only a significant difference between

---

[1] Canvox [32] affords fast octree traversal and hierarchical-distance field computation, which is crucial for a smooth navigation experience without introducing notable latency and/or errors in teleportation distance calculations.

		Descriptive Analysis				Inferential Analysis (ANOVA)		
		Control				***Control***	***Speed***	***Control X Speed***
		Interactive		***Passive***				
Measurements		**Adaptive**	**Constant**	**Adaptive**	**Constant**			
Number of transitions		$M = 2.17, SD = 1.10$	$M = 2.00, SD = 1.10$	$M = 2.67, SD = 1.28$	$M = 2.43, SD = 1.29$	$F(1,15) = 6.36, p < 0.05, \eta^2 = 0.30$	$F(1,15) = 4.31, p = 0.06, \eta^2 = 0.22$	$F(1,15) = 0.02, p = 0.89\ \eta^2 = 0.00$
Task completion time		$M = 22.42, SD = 11.28$	$M = 22.70, SD = 9.34$	$M = 25.81, SD = 11.75$	$M = 22.80, SD = 11.77$	$F(1,15) = 2.67, p = 0.12\ \eta^2 = 0.15$	$F(1,15) = 1.42, p = 0.25, \eta^2 = 0.05$	$F(1,15) = 0.08, p = 0.38, \eta^2 = 0.05$
Navigation error		$M = 0.04, SD = 0.10$	$M = 0.03, SD = 0.09$	$M = 0.03, SD = 0.10$	$M = 0.03, SD = 0.10$	$F(1,15) = 0.18, p = 0.68, \eta^2 = 0.01$	$F(1,15) = 0.00, p = 0.99, \eta^2 = 0.001$	$F(1,15) = 0.33, p = 0.52, \eta^2 = 0.02$
NASA-TLX	Overall	$M =14.76 , SD = 13.53$	$M = 20.90, SD = 14.61$	$M = 23.88, SD = 17.30$	$M = 19.77, SD = 18.66$	$F(1,15) = 1.59, p = 0.23, \eta^2 = 0.10$	$F(1,15) = 0.16, p = 0.69, \eta^2 = 0.01$	$F(1,15) = 4.95, p < 0.05, \eta^2 = 0.25$
	Mental	$M = 18.25, SD = 17.13$	$M = 29.25, SD = 18.62$	$M = 24.69, SD = 18.73$	$M = 23.25, SD = 27.78$	$F(1,15) = 0.00, p = 0.97, \eta^2 = 0.001$	$F(1,15) = 1.69, p = 0.21, \eta^2 = 0.10$	$F(1,15) = 2.34, p = 0.15, \eta^2 = 0.14$
	Physical	$M = 15.31, SD = 13.90$	$M = 21.36, SD = 15.46$	$M = 14.25, SD = 12.05$	$M = 12.31, SD = 10.73$	$F(1,15) = 1.84, p = 0.20, \eta^2 = 0.11$	$F(1,15) = 1.97, p = 0.18, \eta^2 = 0.12$	$F(1,15) = 9.55, p < 0.01, \eta^2 = 0.40$
	Temporal	$M =15.31 , SD =18.26$	$M =19.19 , SD = 18.89$	$M = 20-25, SD = 24.34$	$M = 19.63, SD = 20.72$	$F(1,15) = 0.29, p = 0.60, \eta^2 = 0.02$	$F(1,15) = 0.23, p = 0.64, \eta^2 = 0.01$	$F(1,15) = 0.78, p = 0.39, \eta^2 = 0.05$
	Frustration	$M = 8.63, SD = 11.95$	$M = 15.31, SD = 16.96$	$M = 19.50, SD = 24.68$	$M = 14.06, SD = 16.20$	$F(1,15) = 2.42, p = 0.14, \eta^2 = 0.14$	$F(1,15) = 0.04, p = 0.85, \eta^2 = 0.002$	$F(1,15) = 3.34, p = 0.08, \eta^2 = 0.18$
	Effort	$M = 15.94, SD = 14.56$	$M = 22.56, SD = 19.10$	$M = 30.81\ SD = 24.08$	$M = 20.00, SD = 22.08$	$F(1,15) = 4.96, p < 0.05, \eta^2 = 0.25$	$F(1,15) = 0.35, p = 0.56, \eta^2 = 0.02$	$F(1,15) = 5.90, p < 0.05, \eta^2 = 0.28$
	Performance	$M = 4.50, SD = 8.98$	$M = 9.69, SD = 12.89$	$M =14.25 , SD = 14.23$	$M =14.31 , SD = 21.92$	$F(1,15) = 6.66\ p < 0.05, \eta^2 = 0.31$	$F(1,15) = 2.03, p = 0.18, \eta^2 = 0.12$	$F(1,15) = 0.85, p = 0.37, \eta^2 = 0.05$
Subjective ratings	Preference	$M = 5.62, SD = 0.96$	$M = 5.69, SD = 1.66$	$M = 4.43, SD = 1.44$	$M = 4.19, SD = 1.52$	$F(1,15) = 15.84, p < 0.001, \eta^2 = 0.19$	$F(1,15) = 0.11, p = 0.74, \eta^2 = 0.001$	$F(1,15) = 0.20, p = 0.66, \eta^2 = 0.003$
	Ease of use	$M =5.50 , SD = 0.97$	$M = 5.25, SD = 1.53$	$M = 5.19, SD = 1.60$	$M = 6.00, SD = 0.97$	$F(1,15) = 1.77, p = 0.20, \eta^2 = 0.007$	$F(1,15) = 0.27, p = 0.28, \eta^2 = 0.01$	$F(1,15) = 3.62, p = 0.08, \eta^2 = 0.04$
	Spatial Awareness	$M = 5.81, SD = 1.05$	$M = 6.00, SD = 1.37$	$M = 5.63, SD = 1.36$	$M = 4.69, SD = 1.74$	$F(1,15) = 4.29, p = 0.06, \eta^2 = 0.02$	$F(1,15) = 1.52, p = 0.24, \eta^2 = 0.02$	$F(1,15) = 3.24, p = 0.09, \eta^2 = 0.04$
	Efficiency	$M = 5.86, SD = 0.96$	$M = 4.88, SD = 1.93$	$M = 5.13, SD = 1.86$	$M = 5.63, SD = 1.15$	$F(1,15) = 0.00, p = 0.99, \eta^2 = 0.001$	$F(1,15) = 1.07, p = 0.32, \eta^2 = 0.007$	$F(1,15) = 5.19, p < 0.05, \eta^2 = 0.06$
	Spatial Orientation	$M = 5.81, SD = 1.04$	$M = 6.00, SD = 1.36$	$M = 5.62, SD = 1.36$	$M = 4.69, SD = 1.74$	$F(1,15) = 4.29, p = 0.06, \eta^2 = 0.07$	$F(1,15) = 1.52, p = 0.24, \eta^2 = 0.02$	$F(1,15) = 3.24, p = 0.0p, \eta^2 = 0.04$
	Sickness reduction	$M = 5.75, SD = 1.29$	$M = 5.37, SD = 1.67$	$M = 5.44, SD = 1.36$	$M = 5.00, SD = 1.21$	$F(1,15) = 1.00, p = 0.33, \eta^2 = 0.02$	$F(1,15) = 3.14, p = 0.10, \eta^2 = 0.02$	$F(1,15) = 0.01, p = 0.93, \eta^2 = 0.001$

Table 1: Descriptive and inferential analysis results of study 1, with significant results highlighted.

Interactive+Adaptive vs. Passive+Adaptive ($p < 0.05$). There were no significant differences between all other pairs. There was also no difference between *Speed* modulation methods.

Post-task Interview Responses: Most participants preferred the Interactive *Control* method over Passive, and reported that it helped them to understand the space better. P12 and P13 stated that they could "*go back and forth*" to locate the target and learn the surrounding area. P16 and P17 also mentioned that they could go back if they "*missed*" any spatial information along the way. P7 and P9 said they liked Interactive because they could "*pause and look around*", to learn "*reference points*" (i.e. landmarks) and the spatial layout. Also, participants stated that Interactive alleviated motion sickness and reduced task load. P8 and P15 said that interactive control of the pace led to lower effort for "*learning the space*". Having interactive control seems to have also reduced the incidence of (mild) motion sickness-related symptoms, such as eye strain (P5), motion sickness (P15), and dizziness (P4, P6). Two participants did not like Interactive because the control was difficult to learn (P3) and prevented them from observing the space (P19).

For *Speed* modulation, participant opinions were divided: some preferred Adaptive and others Constant. The participants who preferred Adaptive explained that it was more efficient as it minimizes the moment where their vision was obscured "*by the wall*" (P10, P12, P13) and avoids the "*visual time-out*" where they lost track of landmarks for a moment (P9, P18, P19). One participant liked Adaptive only with Interactive, since the acceleration was perceived as jarring when it was not manually controlled (P17). Interestingly, a few participants mentioned that Adaptive feels smoother (P20) and more reactive, which made the experience more "*engaging*" and "*real*" (P6, P16, P18). On the other hand, participants who preferred Constant reported that the abrupt speed-up with Adaptive interfered with remembering object locations (P3, P14, P15) and made them feel dizzy (P4, P5).

4.8 Discussion

Interactive was strongly preferred over Passive and significantly reduced the total number of transitions without using more time to finish Bookmark. For the Bookmark at each LoS, Interactive significantly reduced the number of transitions at the 3rd and 4th LoS but not at the 2nd. The average completion time for Interactive was larger than for Passive at the 2nd LoS, although the difference was not significant. These results may indicate that the users did not seem to get much of an advantage when using Interactive for the Bookmark at the 2nd LoS, where there are only a few levels to zoom-in. Because they only had to pass through one layer (a wall of a house) to get into the 2nd LoS, it was easy to understand the spatial relationship between the 1st and 2nd LoS without rewinding the transition. Thus, we find that active control of transitions can lengthen the time for completing simple tasks. Yet, when users had to pass through two or more layers, they needed to remember the spatial relationships between multiple nested levels, a challenging task that rarely occurs in the real world. We believe this explains why participants rewound transitions more frequently with Interactive, rather than executing the whole transition more often, which led to higher transition counts with Passive. Yet, further study is warranted. Overall, while there was no significant difference in task completion time, participants judged Interactive to be more efficient and to facilitate spatial understanding more, as it helped them to observe the space better, letting them choose the transition pace and interactively control the transition direction. We hypothesize that in a more complex VE with additional LoS we might observe a stronger difference between *Control* methods.

Even though the difference was not significant, Adaptive required a higher number of transitions due to the sudden speed-up while passing through the layer that surrounds each LoS, which might have confused users and caused them to execute the transition again with Passive. Yet, the NASA-TLX results showed that Interactive increases users' confidence (e.g., performance score) while requiring less effort compared to Passive. While there was no significant difference between Adaptive and Constant in users' confidence, a majority of the participants made positive comments about the combination of Adaptive and Interactive as it is more efficient, engaging, and feels smoother. We believe that the strength of this combination lies in the ability not just to allow users to actively skip parts of a transition where occlusions occur but also the ability to rewind to a previous viewpoint, in case the user missed important features along the path.

Overall, our results partially support **H1** that giving users interactive control of viewpoint transitions in MVEs helps them to understand the space better when the MVE has a complex hierarchical structure. **H2** is also partially supported in that modulating the transition speed based on the object proximity improved the user experience in the MVE to the extent that object proximity-based speed modulation increased users' perceived performance and engagement if transition control is also provided. Overall, we believe that Interactive with Adaptive provides the best user experience among the four combinations of *Control* × *Speed* methods.

5 USER STUDY 2: MIXED TRANSITIONS

Mixed transition techniques can help users maintain orientation and build a cognitive map while transitioning to a distant POI. Previous work [40] had identified that OrbitAndZoom helped users maintain spatial orientation by showing how the view direction rotates at the global level rather than at the local level (ZoomAndRotate) or rotating simultaneously while zooming in (CurvilinearZoom). Thus, we investigate in our second study how similar transition techniques affect navigation performance and usability when they are integrated into mixed transition techniques that involve not only zoom-in but also zoom-out transitions.

5.1 Design

We developed three mixed transition techniques for our study (Fig. 1). MIXEDSTRAIGHTZOOM first zooms out to the lowest common ancestor LoS, then zooms in straight to the selected bookmark. Based on the findings of previous work, where egocentric virtual rotation caused severe disorientation and sickness at a small LoS [40], we decided to eliminate the last rotation component of ZOOMANDROTATE when designing MIXEDSTRAIGHTZOOM to help users to maintain their orientation better. Instead of virtually rotating the view direction at the target location, we let participants turn by themselves towards the target, supported by a corresponding direction indicator. MIXEDORBITANDZOOM was also built on top of the one of the zoom-in techniques, which first zooms out to the overview, orbits around the selected bookmark, and zooms in straight just as ORBITANDZOOM. MIXEDCURVILINEARZOOM is a hybrid technique, which first zooms out and then simultaneously zooms in and orbits around the bookmark like CURVILINEARZOOM. In addition to the three mixed transition techniques, we added a speed modulation feature that slows the transition to a stop as the zoom-out and orbit conclude and then speeds up for the next phase of the transition. We implemented this to address an issue brought up in comments to study 1, that abrupt changes in transition (e.g., from orbiting to zoom-in) were surprising. The speed modulation mentioned above addresses such issues by applying ease-in and ease-out effects to the transition speed between the phases to help users perceive the change in transition in a more unobtrusive way.

5.2 Hypotheses

Based on the results of user study 1, we hypothesize that:

- **H3:** MIXEDORBITANDZOOM is slower for the BOOKMARK than the other two techniques.
- **H4:** MIXEDORBITANDZOOM helps maintain spatial orientation better than MIXEDCURVILINEARZOOM and MIXEDSTRAIGHTZOOM.

H3 is motivated by previous work that compared three zoom-in transition techniques in an MVE [40]. That study showed that for completing a task similar to BOOKMARK, ORBITANDZOOM takes more time than a technique that simultaneously orbits and zooms in. MIXEDORBITANDZOOM is almost the same as ORBITANDZOOM, except for the initial zoom-out step. **H4** is also based on ORBITANDZOOM supporting spatial orientation better than other methods for zoom-in transitions in previous work [40]. We thus postulate that MIXEDORBITANDZOOM supports mixed transition scenarios, too.

5.3 Participants

Similar to previous work that assesses spatial awareness performance of transition techniques using pointing tasks, e.g., [52], we targeted 18 participants (n=18). We recruited new participants from the local university, seven male and five female, with average age 27.4 (SD=4.64). Six were heavy gamers but did not show any distinct behavior or difference in performance in the experimental tasks. All participants were compensated with 15 CAD.

5.4 Apparatus and Environment

We used the same hardware, bookmark interface, and physical setup as in Study 1. We created three versions of a MVEs with different spatial layouts and randomized their order for each participant.

5.5 Task

As in Study 1, there were two main tasks: BOOKMARK and NAVIGATION. However, unlike Study 1, the transitions did not begin at a scene overview. BOOKMARK and NAVIGATION were adapted to the scenario where users travel to a distant POI at a small LoS in a different part of the LoS hierarchy. For BOOKMARK, users had to visit the first and second bookmarks, going back and forth three times. While transitioning to a bookmark they selected, we asked participants to specifically focus on the spatial information that would be needed to perform the following NAVIGATION, including the location of the giant Alice, the locations and features of the two houses where the two bookmarks are located, the hierarchical LoS structure in the houses, and the locations of the two small Alices at the bookmarks. Also, we asked participants to match their view direction with the target orientation at the bookmark by turning their head or body when the transition finished. After BOOKMARK, they were asked to perform the NAVIGATION, in which they first pointed directly to a location in front of the giant Alice and then flew straight to it. The pointing task is similar to previous studies, where it has been shown to be effective in assessing spatial orientation performance [6, 49, 52, 65, 66]. Yet, we adapted the analysis to the 3D pointing scenario, and thus compare not only the pointing accuracy in yaw but also in pitch. Once participants reached the target at the giant Alice, they were asked to fly straight to the small Alice at the second bookmark.

5.6 Procedure

We followed a similar procedure to the study in previous work [40], i.e., two blocks of BOOKMARK and NAVIGATION for each condition. For each blcok, similar-looking environments with different object configurations were presented. The same pre-task questionnaires as in Study 1 for participants' background, and similar post-task questionnaires and interview questions were used to ask about their overall ratings for the three conditions and the reason why they gave the ratings. We also used questionnaires after each condition to record task load. The order of conditions was again counterbalanced.

5.7 Results

To analyze the collected data for angular error, response time, and task completion time, we conducted inferential analysis using repeated-measures ANOVA with $\alpha = 0.05$ in R. For NASA-TLX and Subjective Ratings, we conducted analysis non-parametric analysis using ART [69] before the ANOVA; see Table 2 for details.

Angular Error: To analyze spatial orientation, we computed the angular error between the target vector from the current position to the giant Alice and the first directional vector of the controller when users start pulling the trigger to initiate flying. This is an effective way to assess users' spatial orientation accuracy because the giant alice is not visible from the location at the caterpillar lair where NAVIGATION starts so they only have to depend on their sense of direction. The ANOVA identified a significant effect of *Technique* for the first block. For the second block, the ANOVA did not identify a significant effect of *Technique*.

Response Time: To assess response time for each technique in NAVIGATION, we measured the time between when each NAVIGATION started and when the user first pulled the trigger while pointing at the giant Alice. Previous work showed that response time in a spatial pointing task reliably measure users' cognitive load for spatial orientation [1]. The two-way ANOVA identified a significant effect of *Technique* for the first block, but not for the second block.

Task Completion Time: We analyzed the task completion time for each type of task individually. The ANOVA did not identify a significant effect of *Technique* on completion time for either one.

NASA-TLX: One-way ANOVA identified a significant effect of *Technique* on the overall NASA-TLX score as well as on the performance, effort, and frustration TLX categories, but not on mental, physical, nor temporal demand.

Subjective Ratings: We measured subjective ratings on user's preference, ease of use, spatial awareness, efficiency, spatial orientation, and sickness reduction for each condition using a 7-point

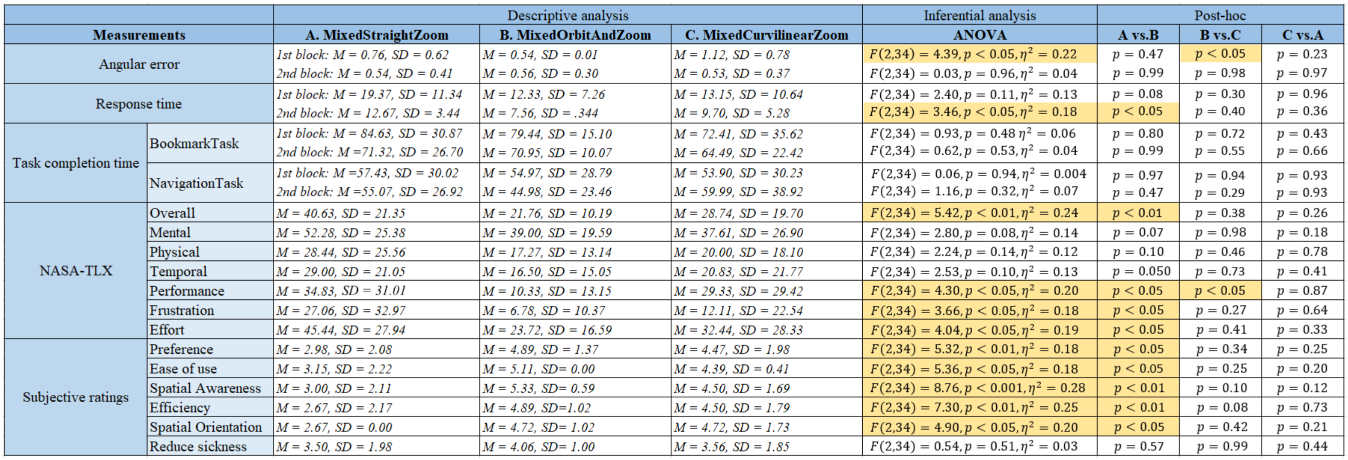

		Descriptive analysis			Inferential analysis	Post-hoc		
Measurements		**A. MixedStraightZoom**	**B. MixedOrbitAndZoom**	**C. MixedCurvilinearZoom**	**ANOVA**	**A vs.B**	**B vs.C**	**C vs.A**
Angular error		1st block: $M = 0.76, SD = 0.62$ 2nd block: $M = 0.54, SD = 0.41$	$M = 0.54, SD = 0.01$ $M = 0.56, SD = 0.30$	$M = 1.12, SD = 0.78$ $M = 0.53, SD = 0.37$	$F(2,34) = 4.39, p < 0.05, \eta^2 = 0.22$ $F(2,34) = 0.03, p = 0.96, \eta^2 = 0.04$	$p = 0.47$ $p = 0.99$	$p < 0.05$ $p = 0.98$	$p = 0.23$ $p = 0.97$
Response time		1st block: $M = 19.37, SD = 11.34$ 2nd block: $M = 12.67, SD = 3.44$	$M = 12.33, SD = 7.26$ $M = 7.56, SD = .344$	$M = 13.15, SD = 10.64$ $M = 9.70, SD = 5.28$	$F(2,34) = 2.40, p = 0.11, \eta^2 = 0.13$ $F(2,34) = 3.46, p < 0.05, \eta^2 = 0.18$	$p = 0.08$ $p < 0.05$	$p = 0.30$ $p = 0.40$	$p = 0.96$ $p = 0.36$
Task completion time	BookmarkTask	1st block: $M = 84.63, SD = 30.87$ 2nd block: $M =71.32, SD = 26.70$	$M = 79.44, SD = 15.10$ $M = 70.95, SD = 10.07$	$M = 72.41, SD = 35.62$ $M = 64.49, SD = 22.42$	$F(2,34) = 0.93, p = 0.48\ \eta^2 = 0.06$ $F(2,34) = 0.62, p = 0.53, \eta^2 = 0.04$	$p = 0.80$ $p = 0.99$	$p = 0.72$ $p = 0.55$	$p = 0.43$ $p = 0.66$
	NavigationTask	1st block: $M =57.43, SD = 30.02$ 2nd block: $M =55.07, SD = 26.92$	$M = 54.97, SD = 28.79$ $M = 44.98, SD = 23.46$	$M = 53.90, SD = 30.23$ $M = 59.99, SD = 38.92$	$F(2,34) = 0.06, p = 0.94, \eta^2 = 0.004$ $F(2,34) = 1.16, p = 0.32, \eta^2 = 0.07$	$p = 0.97$ $p = 0.47$	$p = 0.94$ $p = 0.29$	$p = 0.93$ $p = 0.93$
NASA-TLX	Overall	$M = 40.63, SD = 21.35$	$M = 21.76, SD = 10.19$	$M = 28.74, SD = 19.70$	$F(2,34) = 5.42, p < 0.01, \eta^2 = 0.24$	$p < 0.01$	$p = 0.38$	$p = 0.26$
	Mental	$M = 52.28, SD = 25.38$	$M = 39.00, SD = 19.59$	$M = 37.61, SD = 26.90$	$F(2,34) = 2.80, p = 0.08, \eta^2 = 0.14$	$p = 0.07$	$p = 0.98$	$p = 0.18$
	Physical	$M = 28.44, SD = 25.56$	$M = 17.27, SD = 13.14$	$M = 20.00, SD = 18.10$	$F(2,34) = 2.24, p = 0.14, \eta^2 = 0.12$	$p = 0.10$	$p = 0.46$	$p = 0.78$
	Temporal	$M = 29.00, SD = 21.05$	$M = 16.50, SD = 15.05$	$M = 20.83, SD = 21.77$	$F(2,34) = 2.53, p = 0.10, \eta^2 = 0.13$	$p = 0.050$	$p = 0.73$	$p = 0.41$
	Performance	$M = 34.83, SD = 31.01$	$M = 10.33, SD = 13.15$	$M = 29.33, SD = 29.42$	$F(2,34) = 4.30, p < 0.05, \eta^2 = 0.20$	$p < 0.05$	$p < 0.05$	$p = 0.87$
	Frustration	$M = 27.06, SD = 32.97$	$M = 6.78, SD = 10.37$	$M = 12.11, SD = 22.54$	$F(2,34) = 3.66, p < 0.05, \eta^2 = 0.18$	$p < 0.05$	$p = 0.27$	$p = 0.64$
	Effort	$M = 45.44, SD = 27.94$	$M = 23.72, SD = 16.59$	$M = 32.44, SD = 28.33$	$F(2,34) = 4.04, p < 0.05, \eta^2 = 0.19$	$p < 0.05$	$p = 0.41$	$p = 0.33$
Subjective ratings	Preference	$M = 2.98, SD = 2.08$	$M = 4.89, SD = 1.37$	$M = 4.47, SD = 1.98$	$F(2,34) = 5.32, p < 0.01, \eta^2 = 0.18$	$p < 0.05$	$p = 0.34$	$p = 0.25$
	Ease of use	$M = 3.15, SD = 2.22$	$M = 5.11, SD= 0.00$	$M = 4.39, SD = 0.41$	$F(2,34) = 5.36, p < 0.05, \eta^2 = 0.18$	$p < 0.05$	$p = 0.25$	$p = 0.20$
	Spatial Awareness	$M = 3.00, SD = 2.11$	$M = 5.33, SD= 0.59$	$M = 4.50, SD = 1.69$	$F(2,34) = 8.76, p < 0.001, \eta^2 = 0.28$	$p < 0.01$	$p = 0.10$	$p = 0.12$
	Efficiency	$M = 2.67, SD = 2.17$	$M = 4.89, SD=1.02$	$M = 4.50, SD = 1.79$	$F(2,34) = 7.30, p < 0.01, \eta^2 = 0.25$	$p < 0.01$	$p = 0.08$	$p = 0.73$
	Spatial Orientation	$M = 2.67, SD = 0.00$	$M = 4.72, SD= 1.02$	$M = 4.72, SD = 1.73$	$F(2,34) = 4.90, p < 0.05, \eta^2 = 0.20$	$p < 0.05$	$p = 0.42$	$p = 0.21$
	Reduce sickness	$M = 3.50, SD = 1.98$	$M = 4.06, SD= 1.00$	$M = 3.56, SD = 1.85$	$F(2,34) = 0.54, p = 0.51, \eta^2 = 0.03$	$p = 0.57$	$p = 0.99$	$p = 0.44$

Table 2: Descriptive and inferential analysis results of study 2, with significant results highlighted.

Likert scale. See section 1 in the supplemental material for the questionnaires. The one-way ANOVA identified a significant effect of *Technique* on preference, ease of use, spatial awareness, efficiency, and spatial orientation.

Post-task Interview Responses: Participants commented positively about MixedOrbitAndZoom supporting their spatial understanding and orientation and that it provides an "overview of the layout and topography" (P6, P8), and "*helps understand the location of each item and object*" (P1, P3). P1 specifically mentioned that it "*gave enough time to process all the thing(s) that were happening around me.*" Participants also appreciated that MixedOrbitAndZoom was "*easy to follow*" (P3, P11), "*much smoother*" (P6), and "*less effort*" (P3) to perform the spatial tasks. However, several participants who preferred MixedStraightZoom thought that the orbiting phase was confusing since it "*forces you to change perspective regarding Alice's position when [the] angle changes*" (P0), or "*rotation makes it harder to remember where you came from*" (P19). P0 and P19 commented similarly on MixedCurvilinearZoom.

MixedCurvilinearZoom got positive feedback for efficiency and smoothness, such as "*fast and efficient*" (P5, P15, P16, P17), "*feels smooth and flexible*" (P9, P12, P13, P14), and "*easy to use*" (P2, P7, P10, P15). P2 mentioned that the technique made it easier "*to remember orientation*" as it moves "*for you.*" Negative comments for this technique mentioned causing dizziness (P2, P4, P8), being confusing (P8, P14), being difficult to track the movement (P0, P19), and not being able to see the overview long enough (P13).

Some preferred MixedStraightZoom as it was intuitive (P0), and easy to understand the simple movement (P8, P14, P19). However, the majority did not like it due to the difficulty of maintaining a sense of direction with it. The need to physically turn at the end of the transition disoriented users (P1, P6, P7, P10, P16, P17), increased mental demand (P2, P3, P4, P12, P13, P15), and hindered spatial context understanding (P5, P14). People who disliked MixedStraightZoom stated that it did not allow them see the overview long enough (P1, P10), forced them to reorient after arriving at the target (P16, P17), and that they lost track of the landmark's locations outside the house (P1, P2, P6, P17)

5.8 Discussion

MixedOrbitAndZoom exhibited the best results in terms of angular error and response time for Navigation, performance, frustration, effort, and overall scores in the NASA-TLX, and most categories except for sickness reduction in subjective ratings. MixedOrbitAndZoom exhibited significantly better accuracy (i.e., angular error) for the initial navigation direction in the first block, which aligns with the subjective spatial awareness scores. Even though we did not see a significant difference in accuracy in the second block, the technique still significantly reduced the response time for the initial pointing in Navigation. These results indicate that participants' ability to maintain orientation with the three transitions was improved by repeating the task which led to similar pointing accuracy, but they spent significantly less time using MixedOrbitAndZoom to orient themselves once they adapted to the task itself. Many participants also expressed their confidence and ease in maintaining a sense of direction and an understanding of space in the post-task questionnaires and interviews. These results support **H4**. While MixedOrbitAndZoom supports spatial orientation and helps users understand the space better, it did not lead to longer times for finishing Bookmark, even though the transition took longer. This result rejects **H3**. We believe that MixedOrbitAndZoom's better support for spatial orientation made participants pause less to reorient themselves when the transition was done.

MixedCurvilinearZoom received similar scores with MixedOrbitAndZoom in many subjective measures and TLX categories. However, this technique significantly increased the pointing error over MixedOrbitAndZoom. These results resonate with the findings in a previous study [52] where simultaneous translation and rotation leads to disorientation and increase the pointing error. The MixedCurvilinearZoom technique combines translation and rotation components into one transition motion which caused significant disorientation and dizziness. While MixedOrbitAndZoom also combines translation and rotation components for the orbiting motion, the results of study 2 showed that this combination helps users understand their spatial relationship to the environment rather than causing disorientation for the mixed transition as well.

Interestingly, we found a distinct group of participants who prefer MixedStraightZoom over the other two techniques. This group seems to have used different strategies and we hypothesize that they were mainly depending on route knowledge (while the others were depending more on survey knowledge for wayfinding). In their post-task questionnaires and a follow-up interview, people in this group explained that they focus more on the direction where they are going instead of understanding their relative position in the space. However, the lack of rotational transition of MixedStraightZoom from the vantage point created a higher effort to assess their self-orientation after physical turning, which increased response time. We previously observed a similar problem with ZoomAndRotate, which motivated us to eliminate the rotational transition at the end when developing MixedStraightZoom. However, eliminating the rotational transition and giving users vestibular and proprioceptive feedback by asking them to rotate physically still did not address this challenge adequately.

6 OVERALL DISCUSSION FOR MULTISCALE TRANSITIONS

6.1 Trajectory

Separating the rotational component of zoom-in and mixed transition shows a number of benefits in terms of spatial awareness, orientation, and usability (task load). Moreover, when designing transitions that include rotations, the order of rotation and zoom-in also affected navigation performance. Orbiting at a large LoS before zooming in to a small LoS not only supported users in building a cognitive map of the environment but also helped them keep track of where they were heading with respect to their surroundings. Rotating after zooming in to a small LoS increased mental load and was prone to participants losing track of the landmarks regardless of whether the rotation was done physically by the user or virtually by the system. Thus we do not consider this a good design choice.

We also found that any rotational component during zoom-out transitions confused users and caused disorientation. For scenarios where zooming out is necessary, such as in mixed transitions between different LoS, we presented an effective design for zooming out to a larger LoS. However, our zoom-out technique does not apply in transitions to a target bookmark that is at the immediate parent LoS. In this case, we suggest designing a mixed transition that first zooms out to the lowest common ancestor of the current and the target LoS in the hierarchy and then zooms into the target (see Fig. 2).

Even though we pre-defined the LoS hierachy for user study 2, ORBITANDZOOM and MIXEDORBITANDZOOM can support the user's spatial awareness even without a pre-defined hierarchy. Depending on the visibility of the target LoS from the current LoS, the system can automatically determine whether a zoom-in or mixed-transition should be initiated. For example, if the target is at a smaller LoS and visible from the current viewpoint, the system can simply initiate ORBITANDZOOM. If it is not visible or is visible but at a larger LoS, the system can initiate a MIXEDORBITANDZOOM where it first zooms back out until both initial and target viewpoints are visible within the user's FOV, then execute ORBITANDZOOM.

6.2 Interactive Control

In general, users preferred having interactive control of transitions since it increased confidence by letting them determine the timing of the transition and even rewind the transition to understand the hierarchical relationships between LoS rather than executing the transition repeatedly. However, giving them the control of the transition may delay the user flow in a few scenarios, such as VR authoring, where users switch frequently among their bookmarks, e.g., to make consistent changes in two distant places, such as changing the design of Alice's hair band in two different places. Thus, we suggest designing the input method for bookmark navigation so that it does not overlap with other input modes for authoring tasks, so that users can stop, fast forward, or rewind the transition whenever they want to interrupt their main authoring tasks as little as possible.

6.3 Speed Modulation

Proximity-based speed modulation did not significantly improve navigation performance nor usability. However, ADAPTIVE showed potential when combined with INTERACTIVE where users can minimize the duration of their sight being blocked by a wall while passing through it. When there was no control over the timing, sudden changes in transition speed may be jarring to users, but our study 1 results also showed that when users have control over the timing of such changes, it can help improve the user experience without surprising them. Further study is warranted to investigate whether other speed modulation [2, 23] and scale adaptation methods [4] can enhance not just usability but also navigation performance when being used with interactive control.

Our speed adaptation method was designed for MVEs with a nested structure, i.e., nested objects, where each LoS is (potentially) completely enclosed by walls and ceilings or continuous surfaces without an opening, and where the user has to pass through a surface to get inside other objects. For simpler scene configurations other designs could work, too. For example, if the MVE consists only of a flat surface, such as a 2D map in VR, or a single sphere, such as Google Earth VR, a transition method could accelerate before stopping at a target bookmark. Also, if the MVE is an mostly open environment, such as a solar system, where all objects float in space, and one does not need to pass through walls, the transition speed could increase while passing by other planets or asteroid belts. In both approaches, these sudden speed-ups might be perceived as unpredictable, potentially disrupting users from maintaining spatial orientation. Yet, we had to disregard these simpler approaches, because we wanted to ensure that our method works in general MVEs that contain even geometrically fully nested structures. Future work is needed to improve our method so that it detects simpler cases and then potentially suppresses unnecessary acceleration. In addition, it would also be worthwhile to test different speed adjustment methods presented in previous work [3, 23] in such environments to assess their effectiveness.

6.4 Cut-away Visualization

The vast differences in scale and distances in MVEs make cut-away visualization challenging. We initially tried a fixed size cut-away and one that scales proportionally to the user, but both options yield results that are either too small or too large (encompassing even whole houses in our MVE). Thus, the size of a cut-away would also need to (at least) depend on the distance to the next object along the movement direction, and we will explore this in the future.

7 CONCLUSION AND FUTURE WORK

We studied viewpoint transitions for bookmark navigation in a complex virtual environment with multiple levels of scale, extending previous work that focused only on an environment with a single level of scale or on transitions method that simply interpolate position and scale. In study 1, we developed a zoom-in transition technique and evaluated a new control method that lets users interactively control the zoom-in speed and direction. The results showed that having active control improve users' spatial awareness and performance in spatial learning tasks in MVEs. We also evaluated a method that automatically modulates speed based on object proximity and found advantages when combined with interactive control. In study 2, we developed three mixed transition techniques that involve both zoom-out and zoom-in to transition to a distant target. The results showed that a technique that zooms out straight, orbits to reorient the user, and then zooms in showed the best performance in supporting spatial orientation, usability, and preference scores. For both our studies, each statistically significant result showed a large effect size, traditionally defined as $\eta^2 > 0.14$, which indicates that our results should be robust. Still, we did not see many power levels above 80%, so it may be necessary to replicate our work with larger participant numbers to verify the results. Some open questions remain. First, we want to follow up on the different navigation strategies we observed in study 2. Second, research on the effectiveness of different speed modulation methods within MVEs with different geometrical configurations, such as open spaces, would also be meaningful.

ACKNOWLEDGMENTS

The authors wish to thank Byungmoon Kim for research advice, Yeojin Kim for development support, and Jini Kwon for designing the environment.

REFERENCES

[1] A. Adhikari, D. Zielasko, I. Aguilar, A. Bretin, E. Kruijff, M. von der Heyde, and B. E. Riecke. Integrating continuous and teleporting vr locomotion into a seamless 'hyperjump'paradigm. *IEEE Transactions on Visualization and Computer Graphics*, 2022.

[2] F. Argelaguet. Adaptive navigation for virtual environments. In *2014 IEEE symposium on 3D user interfaces (3DUI)*, pp. 123–126. IEEE, 2014.

[3] F. Argelaguet. Adaptive navigation for virtual environments. In *2014 IEEE Symposium on 3D User Interfaces (3DUI)*, pp. 123–126, Mar. 2014. doi: 10.1109/3DUI.2014.7027325

[4] F. Argelaguet and M. Maignant. Giant: stereoscopic-compliant multi-scale navigation in ves. In *Proceedings of the 22nd acm conference on virtual reality software and technology*, pp. 269–277, 2016.

[5] F. Bacim, D. Bowman, and M. Pinho. Wayfinding techniques for multiscale virtual environments. In *2009 IEEE Symposium on 3D User Interfaces*, pp. 67–74. IEEE, 2009.

[6] J. Z. Bakdash, S. A. Linkenauger, and D. Proffitt. Comparing decision-making and control for learning a virtual environment: Backseat drivers learn where they are going. *Proceedings of the human factors and ergonomics society annual meeting*, 52(27):2117–2121, 2008.

[7] M. D. Barrera Machuca and W. Stuerzlinger. 3d camera pose history visualization. In *Proceedings of the 2016 Symposium on Spatial User Interaction*, pp. 183–183, 2016.

[8] J. Bhandari, P. R. MacNeilage, and E. Folmer. Teleportation without spatial disorientation using optical flow cues. In *Graphics interface*, pp. 162–167, 2018.

[9] B. Bolte, F. Steinicke, and G. Bruder. The jumper metaphor: an effective navigation technique for immersive display setups. In *Proceedings of Virtual Reality International Conference*, pp. 1–7, 2011.

[10] 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. IEEE, 1997.

[11] E. Bozgeyikli, A. Raij, S. Katkoori, 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.

[12] H. Buchholz, J. Bohnet, and J. Dollner. Smart and physically-based navigation in 3d geovirtual environments. In *Ninth International Conference on Information Visualisation (IV'05)*, pp. 629–635, 2005. ISSN: 2375-0138. doi: 10.1109/IV.2005.117

[13] N. Burtnyk, A. Khan, G. Fitzmaurice, R. Balakrishnan, and G. Kurtenbach. StyleCam: interactive stylized 3d navigation using integrated spatial & temporal controls. In *Proceedings of the 15th annual ACM symposium on User interface software and technology*, UIST '02, pp. 101–110. Association for Computing Machinery, 2002. doi: 10.1145/571985.572000

[14] N. Burtnyk, A. Khan, G. Fitzmaurice, and G. Kurtenbach. ShowMotion: camera motion based 3d design review. In *Proceedings of the 2006 symposium on Interactive 3D graphics and games*, I3D '06, pp. 167–174. Association for Computing Machinery, 2006. doi: 10.1145/1111411.1111442

[15] C. G. Christou and H. H. Bülthoff. View dependence in scene recognition after active learning. *Memory & Cognition*, 27(6):996–1007, 1999.

[16] S. Cmentowski, A. Krekhov, 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.

[17] 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*, pp. 1–10, 2021.

[18] 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 (BDVA)*, pp. 1–10. IEEE, 2018.

[19] Dynamoid. Cosm worlds. [steamVR], 2016.

[20] 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 (VR)*, pp. 475–476. IEEE, 2017.

[21] D. Englmeier, W. Sajko, and A. Butz. Spherical world in miniature: Exploring the tiny planets metaphor for discrete locomotion in virtual reality. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*, pp. 345–352. IEEE, 2021.

[22] S. K. Feiner and D. D. Seligmann. Cutaways and ghosting: satisfying visibility constraints in dynamic 3d illustrations. *The Visual Computer*, 8(5):292–302, 1992.

[23] S. Freitag, B. Weyers, and T. W. Kuhlen. Automatic speed adjustment for travel through immersive virtual environments based on viewpoint quality. In *2016 IEEE Symposium on 3D User Interfaces (3DUI)*, pp. 67–70. IEEE, 2016.

[24] J. Frommel, S. Sonntag, and M. Weber. Effects of controller-based locomotion on player experience in a virtual reality exploration game. In *Proceedings of the 12th international conference on the foundations of digital games*, pp. 1–6, 2017.

[25] Google. Google earth vr. [steamVR], 2016.

[26] Google. Google tilt brush. [steamVR], 2016.

[27] N. N. Griffin and E. Folmer. Out-of-body locomotion: Vectionless navigation with a continuous avatar representation. In *25th ACM Symposium on Virtual Reality Software and Technology*, pp. 1–8. ACM, 2019. doi: 10.1145/3359996.3364243

[28] M. Hachet, F. Decle, S. Knodel, and P. Guitton. Navidget for easy 3d camera positioning from 2d inputs. In *2008 IEEE Symposium on 3D User Interfaces*, pp. 83–89, 2008. doi: 10.1109/3DUI.2008.4476596

[29] M. Hachet, F. Decle, S. Knödel, and P. Guitton. Navidget for 3d interaction: Camera positioning and further uses. *International Journal of Human-Computer Studies*, 67(3):225–236, 2009. doi: 10.1016/j.ijhcs.2008.09.013

[30] S. G. Hart. Nasa task load index (tlx). volume 1.0; paper and pencil package, 1986.

[31] A. Hashemian, M. Lotfaliei, A. Adhikari, E. Kruijff, and B. Riecke. Headjoystick: Improving flying in vr using a novel leaning-based interface. *IEEE Transactions on Visualization and Computer Graphics*, 2020.

[32] Y. Kim, B. Kim, and Y. J. Kim. Dynamic deep octree for high-resolution volumetric painting in virtual reality. *Computer Graphics Forum*, 37(7):179–190, 2018.

[33] R. Komiyama, T. Miyaki, and J. Rekimoto. JackIn space: designing a seamless transition between first and third person view for effective telepresence collaborations. In *Proceedings of the 8th Augmented Human International Conference on - AH '17*, pp. 1–9. ACM Press, 2017. doi: 10.1145/3041164.3041183

[34] R. Kopper, T. Ni, D. A. Bowman, and M. Pinho. Design and evaluation of navigation techniques for multiscale virtual environments. In *Ieee virtual reality conference (vr 2006)*, pp. 175–182. Ieee, 2006.

[35] R. Kopper, Tao Ni, D. A. Bowman, and M. Pinho. Design and Evaluation of Navigation Techniques for Multiscale Virtual Environments. In *IEEE Virtual Reality Conference (VR 2006)*, pp. 175–182, Mar. 2006. doi: 10.1109/VR.2006.47

[36] A. Krekhov, S. Cmentowski, K. Emmerich, M. Masuch, and J. Krüger. Gullivr: A walking-oriented technique for navigation in virtual reality games based on virtual body resizing. In *Proceedings of the 2018 Annual Symposium on Computer-Human Interaction in Play*, pp. 243–256, 2018.

[37] J. J. LaViola Jr, D. A. Feliz, D. F. Keefe, and R. C. Zeleznik. Hands-free multi-scale navigation in virtual environments. In *Proceedings of the 2001 symposium on Interactive 3D graphics*, pp. 9–15, 2001.

[38] J. J. LaViola Jr, E. Kruijff, R. P. McMahan, D. Bowman, and I. P. Poupyrev. *3D user interfaces: theory and practice*. Addison-Wesley Professional, 2017.

[39] 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*, pp. 1–10, 2020.

[40] J.-I. Lee, P. Asente, and W. Stuerzlinger. A comparison of zoom-in transition methods for multiscale vr. In *ACM SIGGRAPH 2022 Posters*, pp. 1–2. ACM, 2022.

[41] T. B. V. LLC. The body vr. [steamVR], 2016.
[42] J. D. Mackinlay, S. K. Card, and G. G. Robertson. Rapid controlled movement through a virtual 3D workspace. *ACM SIGGRAPH Computer Graphics*, 24(4):171–176, Sept. 1990. doi: 10.1145/97880.97898
[43] J. Matejka, T. Grossman, and G. Fitzmaurice. Swift: Reducing the effects of latency in online video scrubbing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 637–646, 2012.
[44] A. Matviienko, F. Müller, M. Schmitz, M. Fendrich, and M. Mühlhäuser. Skyport: Investigating 3d teleportation methods in virtual environments. In *CHI Conference on Human Factors in Computing Systems*, pp. 1–11, 2022.
[45] M. May, P. Péruch, and A. Savoyant. Navigating in a virtual environment with map-acquired knowledge: Encoding and alignment effects. *Ecological Psychology*, 7(1):21–36, 1995.
[46] J. McCrae, I. Mordatch, M. Glueck, and A. Khan. Multiscale 3d navigation. *Proceedings of the 2009 symposium on Interactive 3D graphics and games - I3D '09*, pp. 7–14, 2009. doi: 10.1145/1507149.1507151
[47] D. Medeiros, E. Cordeiro, D. Mendes, M. Sousa, A. Raposo, A. Ferreira, and J. Jorge. Effects of speed and transitions on target-based travel techniques. In *Proceedings of the 22nd ACM Conference on Virtual Reality Software and Technology - VRST '16*, pp. 327–328. ACM Press, 2016. doi: 10.1145/2993369.2996348
[48] 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, 2019.
[49] E. Mellet, L. Laou, L. Petit, L. Zago, B. Mazoyer, and N. Tzourio-Mazoyer. Impact of the virtual reality on the neural representation of an environment. *Human brain mapping*, 31(7):1065–1075, 2010.
[50] D. Miau and S. Feiner. Spacetokens: Interactive map widgets for location-centric interactions. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pp. 1–12, 2018.
[51] M. R. Mine. Virtual environment interaction techniques. *UNC Chapel Hill CS Dept*, 1995.
[52] K. R. Moghadam and E. D. Ragan. Towards understanding scene transition techniques in immersive 360 movies and cinematic experiences. In *2017 IEEE Virtual Reality (VR)*, pp. 375–376. IEEE, 2017.
[53] Oculus. Oculus quill. [Oculus Store], 2016.
[54] K. Rahimi, C. Banigan, and E. D. Ragan. Scene transitions and teleportation in virtual reality and the implications for spatial awareness and sickness. *IEEE transactions on visualization and computer graphics*, 26(6):2273–2287, 2018.
[55] V. Rudakova, N. Lin, N. Trayan, T. M. Sezgin, J. Dorsey, and H. E. Rushmeier. Cher-ish: A sketch-and image-based system for 3d representation and documentation of cultural heritage sites. In *GCH*, pp. 195–199, 2017.
[56] K. A. Satriadi, B. Ens, M. Cordeil, B. Jenny, T. Czauderna, and W. Willett. Augmented reality map navigation with freehand gestures. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 593–603. IEEE, 2019.
[57] S. Studios. In death. [steamVR], 2019.
[58] M. Sukan and S. Feiner. SnapAR: Storing snapshots for quick viewpoint switching in hand-held augmented reality. In *2010 IEEE International Symposium on Mixed and Augmented Reality*, pp. 273–274, 2010. doi: 10.1109/ISMAR.2010.5643603
[59] M. Sukan, S. Feiner, B. Tversky, and S. Energin. Quick viewpoint switching for manipulating virtual objects in hand-held augmented reality using stored snapshots. In *2012 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 217–226, 2012. doi: 10.1109/ISMAR.2012.6402560
[60] M. Usoh, K. Arthur, M. C. Whitton, R. Bastos, A. Steed, M. Slater, and F. P. Brooks Jr. 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.
[61] Valve. The broken seal: Arena. [steamVR], 2019.
[62] Valve. Half-life: Alyx. [steamVR], 2020.
[63] E. Veas, A. Mulloni, E. Kruijff, H. Regenbrecht, and D. Schmalstieg. Techniques for view transition in multi-camera outdoor environments. In *Graphics Interface*, p. 9, 2010.
[64] O. Views. Overview. [steamVR], 2018.
[65] D. Waller and N. Greenauer. The role of body-based sensory information in the acquisition of enduring spatial representations. *Psychological research*, 71(3):322–332, 2007.
[66] D. Waller, J. M. Loomis, and D. Haun. Body-based senses enhance knowledge of directions in large-scale environments. *Psychonomic bulletin & review*, 11(1):157–163, 2004.
[67] C. Ware and D. Fleet. Context sensitive flying interface. In *Proceedings of the 1997 Symposium on Interactive 3D graphics*, pp. 127–ff, 1997.
[68] Wikipedia. Alice's Adventures in Wonderland — Wikipedia, the free encyclopedia. `https://en.wikipedia.org/wiki/Alice%27s_Adventures_in_Wonderland`, 2021. [Online; accessed 23-July-2021].
[69] J. O. Wobbrock, L. Findlater, D. Gergle, and J. J. Higgins. The aligned rank transform for nonparametric factorial analyses using only ANOVA procedures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pp. 143–146. ACM, New York, NY, USA, 2011. doi: 10.1145/1978942.1978963
[70] X. Zhang and G. W. Furnas. Social interactions in multiscale cves. In *Proceedings of the 4th international conference on Collaborative virtual environments*, pp. 31–38, 2002.
[71] X. L. Zhang. Multiscale traveling: crossing the boundary between space and scale. *Virtual reality*, 13(2):101, 2009.
[72] J. Zhao, D. Wigdor, and R. Balakrishnan. Trailmap: facilitating information seeking in a multi-scale digital map via implicit bookmarking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 3009–3018, 2013.