

ScaleCue-WIM: Boundary-Controlled WIM Scaling for Scale-Aware Navigation in Multiscale Virtual Environments

Junkun Long
Computational Media and Arts Thrust
The Hong Kong University of Science and Technology
(Guangzhou)
Guangzhou, Guangdong, China
jlong892@connect.hkust-gz.edu.cn

Wolfgang Stuerzlinger
School of Interactive Arts + Technology (SIAT)
Simon Fraser University
Vancouver, British Columbia, Canada
w.s@sfu.ca

Pan Hui
Computational Media and Arts Thrust
The Hong Kong University of Science and Technology
(Guangzhou)
Guangzhou, Guangdong, China
panhui@hkust-gz.edu.cn

Hai-Ning Liang
Computational Media and Arts Thrust
The Hong Kong University of Science and Technology
(Guangzhou)
Guangzhou, Guangdong, China
hainingliang@hkust-gz.edu.cn




Figure 1: ScaleCue-WIM scaling workflow, illustrated with Earth. The user selects Earth with the ray to place it at the scale center (left). Scale-in (top): shrinking the Scale-Frame to Earth's Semantic Box enlarges the WIM content until that box aligns with the WIM frame, after which a smaller region, such as the Arctic, can be selected as the next focus. Scale-out (bottom): expanding the Scale-Frame to the fixed WIM frame shrinks the preview content until the next outer Numerical Box fits that frame. The illustrated boundary-triggered transitions are delimited by semantic or numerical scale references, and the Metric Extent Readout reports the currently selected real-world side length.

Abstract

Navigating multiscale virtual environments (MSVEs) requires users to understand their current scale and follow scale changes during navigation. World-in-Miniature (WIM) interfaces provide an exocentric preview for inspecting views across scales. A remaining design question is how to connect semantic and numerical scale references to preview adjustment so that users can follow these



This work is licensed under a Creative Commons Attribution 4.0 International License.
VRST '26, Sendai, Japan
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2811-2/2026/11
<https://doi.org/10.1145/3822517.3848745>

changes. We present ScaleCue-WIM, which combines these references with a central Scale-Frame. Users resize the frame to select an extent; reaching a boundary triggers automatic preview scaling, during which the frame remains visible as a reference. We conducted an exploratory within-subjects study with 24 participants comparing ScaleCue-WIM with Cue-Augmented Scaled and Scrolling WIM (CASSWIM), an author-constructed continuous-scaling workflow with comparable scale cues. Across the guided cross-scale routes tested, participants completed routes 7.37 s (8.7%) faster and used 4.66 (20.0%) fewer scale-control operations with ScaleCue-WIM. Participants also reported stronger perceived scale support. Participants described how the workflow helped them interpret scale changes, while also noting learning and control trade-offs. These findings offer insights into how WIM interfaces can support users' understanding of scale as they inspect and adjust previews for multiscale navigation.

CCS Concepts

• **Human-centered computing** → **Virtual reality**; *Interaction techniques*; *Empirical studies in HCI*.

Keywords

Multiscale virtual environments, world-in-miniature, virtual reality navigation, scale cues, interaction techniques

ACM Reference Format:

Junkun Long, Wolfgang Stuerzlinger, Pan Hui, and Hai-Ning Liang. 2026. ScaleCue-WIM: Boundary-Controlled WIM Scaling for Scale-Aware Navigation in Multiscale Virtual Environments. In *32nd ACM Symposium on Virtual Reality Software and Technology (VRST '26)*, November 16–18, 2026, Sendai, Japan. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3822517.3848745>

1 Introduction

Multiscale virtual environments (MSVEs) allow users to explore vastly different levels of scale (LoS) spanning from microscopic biological structures to astronomical systems. To navigate such environments, users need to know which LoS they are in, how the current view relates to larger and smaller structures, which object or region should become the next scale center, and how large the next scale change should be. Therefore, MSVE navigation combines locomotion, wayfinding, and scale reasoning [2, 20, 26, 33].

Prior multiscale navigation techniques have primarily addressed the mechanics of cross-scale locomotion. These approaches adapt travel speed, control-display ratio, user scale, or view parameters to nearby structures, allowing users to move more effectively across large spatial ranges [1, 6, 31, 34, 43, 46]. Our focus is the complementary problem of supporting users' understanding of scale while they inspect and adjust an exocentric preview before changing the immersive view.

Exocentric approaches, especially World-in-Miniature (WIM) interfaces, provide an established basis for browsing LoS and planning cross-scale routes. A WIM gives users an exocentric representation that can support overview, object selection, manipulation, route planning, preview, and travel [35, 42]. Later WIM systems extended this idea with scaling, scrolling, hierarchical maps, power-law representations, landmarks, occlusion-management techniques, and

integrated navigation workflows for multiscale environments [2, 29, 44, 45, 48]. These systems show that WIMs can support both navigation and scale adjustment in MSVEs. Scalable WIM [29] combined continuous spatial scaling with landmarks and cues for scale, position, and direction. The interface design question we address is how semantic and numerical scale references can assist users in understanding their current scale and following scale changes throughout WIM-preview adjustment. ScaleCue-WIM connects these references to extent adjustment and automatic preview transitions, with the Scale-Frame remaining visible as a reference.

The available scale references vary across an MSVE. Dense biological scenes contain repeated, nested structures, while sparse astronomical scenes contain large intervals with few visible landmarks. These differences motivate combining semantic extents with numerical scale references.

ScaleCue-WIM combines a *Scale-Frame* with semantic and numerical references to support users as they inspect and adjust the WIM preview. The Scale-Frame is a cubic frame at the WIM's geometric center, which also serves as the scale center. Users place an object or region at this center and resize the frame to select an extent. Reaching a Semantic Box, Numerical Box, or the WIM frame starts an automatic transition to the corresponding semantic or power-of-ten metric extent (Figure 1; detailed in Section 3.3). The resulting WIM state previews the candidate view's LoS, position, and viewing direction before the user teleports.

We conducted an exploratory within-subjects study with 24 participants comparing ScaleCue-WIM and Cue-Augmented Scaled and Scrolling WIM (CASSWIM) as complete WIM preview-and-teleport workflows. To ensure a fair comparison, we constructed CASSWIM by extending SSWIM [48] with the same types of semantic and numerical cues used in ScaleCue-WIM. Both conditions provided comparable scale cues but used different scale-control workflows. In the guided routes tested, participants completed tasks faster and used fewer scale-control operations with ScaleCue-WIM, and they reported stronger perceived scale support.

This paper contributes: (1) the design and implementation of ScaleCue-WIM, a workflow that connects semantic and numerical scale references with extent selection and preview transitions to support scale judgments during WIM adjustment; and (2) an exploratory comparison with a cue-augmented continuous-scaling workflow, providing evidence on guided route execution, perceived scale support, and learning and control trade-offs.

2 Related Work

2.1 Egocentric and Exocentric Multiscale Navigation

Egocentric MSVE techniques mainly support travel across large spatial ranges. Pan-and-zoom and space-scale models established movement across orders of magnitude [3, 11, 38], while foundational VR research examined how viewpoint motion should be controlled [4]. Later systems adapted travel speed, control-display ratio, user scale, or view parameters to nearby structures [1, 6, 25, 31, 34, 43, 46]. Research also shows that transition design affects how users understand nested spatial relationships [27], while a third-person bird's-eye view can support orientation during long-distance travel [7]. This work primarily concerns control of the

immersive view. WIM interfaces provide a complementary space for inspecting scale and configuring a target view.

World-in-Miniature (WIM) interfaces provide an exocentric planning space that complements the immersive detail view. The original metaphor supported overview, selection, manipulation, and travel [35, 42]. SSWIM later added scaling and scrolling to overcome fixed-scale coverage limits [48], while other systems addressed object access, occlusion, and integrated navigation workflows [44, 45]. MeTACAST also uses a hand-held WIM to support target- and context-aware spatial selection in particle data [49].

Using Danyluk et al.’s design space [8], both conditions used a single headset-VR WIM with geometry clipped to a cube and a fixed physical size with variable scope and scale. The WIM frame stayed at a fixed position and orientation in the virtual environment; users translated, rotated, and scaled the content inside it. This provided a stable display area, while the cubic frame defined the visible extent of the WIM content. We held these WIM properties constant across conditions. The implemented size and scale range are reported in Section 4.4.

Several systems address scale support more directly. Bacim et al. compared Multi-Scale WIM, hierarchical maps, and combined techniques for MSVE wayfinding [2]. Scalable WIM combined continuous power-law scaling with landmark labels and nested power-of-ten cubes, cylinders, or spheres that provided spatial and scale context [29]. Lee and Stuerzlinger examined input methods and scaling centers for exocentric scale control [28]. Weissker et al. compared ways to specify target scale during teleportation and provided a WIM preview of the surroundings at the selected position and scale [47]. PRIMO used progressive selection of subvolumes to inspect dense, homogeneous objects, comparing predefined subdivisions with freely positioned selection boxes [36]. ScaleFree uses dynamically updated density fields to support progressive navigation in multiscale point clouds, moving the viewpoint to a selected region and adjusting the viewing scale to its extent [50]. These studies provide scale controls, region-selection methods, contextual references, and navigation previews. We examine a workflow that connects semantic and numerical extents to boundary-triggered WIM-preview transitions and keeps the selected extent visible during adjustment, comparing it with cue-augmented continuous WIM scaling.

2.2 Scale Cues, Labels, and Operational Boundaries

Semantic and numerical references provide complementary forms of scale support in MSVEs. Semantic information conveys what an object or region is, how it is contained within larger structures, and which spatial extent might be meaningful to inspect next; numerical information conveys physical size and order-of-magnitude change, which can help when semantic structure is sparse or ambiguous. The distinction is important because a single MSVE can contain semantic-rich scenes, dense biological structures, and sparse space across its LoS.

Semantic support is grounded in spatial cognition: people navigate not only by metric coordinates but by landmarks, regions, routes, and hierarchical organization [9, 14, 15, 30, 32, 37]. In multiscale VR, perceptual cues can influence how users interpret their

scale relative to the environment [24]. This motivates providing explicit scale references during interaction. In semantic-rich scenes—rooms, buildings, streets, map-like regions—and for bounded biological structures such as cells or organelles, visible objects and meaningful boundaries offer references for choosing the next scale center and extent. Dense biological scenes are harder: repeated instances and nested structures make exhaustive labeling impractical, which Labels on Levels and HyperLabels address through hierarchy-aware label management, representative-instance labeling, parent-level context, and clutter reduction [21, 22]. These works inform the semantic side of ScaleCue-WIM: labels identify objects and regions, while Semantic Boxes represent meaningful extents that can serve as operational boundaries for WIM-preview transitions.

Numerical references complement semantic cues where meaningful objects or regions are sparse. In sparse space, targets may be too small or distant to remain visible or selectable, so persistent landmarks and labels keep them available as anchors [39]; but labels alone do not specify a region’s physical size or a scale change’s magnitude. Logarithmic spaces, power-law representations, and scale indicators provide ways to represent and navigate very large spatial ranges [10–12, 29]. ScaleCue-WIM supplies this through Numerical Boxes, which mark power-of-ten metric extents, and a Metric Extent Readout, which reports the real-world side length currently selected by the Scale-Frame.

Direct-manipulation and external-cognition research motivates keeping relevant action information visible during interaction [17, 41]. We apply this rationale to WIM scaling by showing the selected extent and the next transition boundary throughout adjustment. The design aims to support scale judgments throughout preview adjustment by combining semantic and numerical references, visible transition boundaries, and preview inspection before teleportation.

3 ScaleCue-WIM

ScaleCue-WIM keeps semantic and numerical scale references available as users adjust and inspect a WIM preview. Users configure a candidate immersive view, including its LoS, position, and viewing direction, then enter it through teleportation. The current prototype environment spans DNA-level biological structures, human-scale spaces, and solar-system scenes.

3.1 Design Overview and Workflow

The *Scale-Frame* is a cubic frame at the WIM’s geometric center, which also serves as the scale center. Users place an object or region at this center and resize the Scale-Frame to choose the extent to inspect. The Metric Extent Readout shows the selected real-world side length, while Semantic and Numerical Boxes mark possible extents. Shrinking the Scale-Frame to an internal box starts a scale-in transition; expanding it to the fixed WIM frame starts a scale-out transition.

The workflow has four stages. First, users center an object or region by selecting its visible object or label with a ray, or by translating and rotating the WIM content. Second, they resize the Scale-Frame to select an extent. Third, a boundary contact triggers automatic scaling of the WIM preview, following the rules in Section 3.3. Finally, users inspect the resulting preview and either continue adjusting it or enter it through teleportation.

3.2 Scale Cues in the WIM

ScaleCue-WIM combines manually authored Labels and Semantic Boxes with automatically generated Numerical Boxes and a Metric Extent Readout (Figure 2). The Semantic Boxes form a manually defined parent–child hierarchy. Labels identify objects or regions, Semantic Boxes mark meaningful extents, Numerical Boxes mark power-of-ten metric extents, and the readout reports the Scale-Frame’s current real-world side length.

Labels identify objects, regions, or representative instances. In map-like human-scale scenes, peer-level labels can remain visible, as on a map. For bounded objects with clear visual extents, labels appear mainly during ray preselection to confirm object identity. In dense biological scenes, ScaleCue-WIM uses hierarchy-aware label management inspired by Labels on Levels [21] and Hyper-Labels [22], including parent-level context and representative instances. In sparse space, persistent labels or landmarks keep small or distant targets available as navigation anchors.

Semantic Boxes enclose meaningful objects or regions, including rooms, buildings, cells, organelles, planets, orbital regions, and the solar system. ScaleCue-WIM constructs a default scaling chain through the semantic hierarchy by selecting, at each level, the child whose physical center is closest to that of its parent. To limit clutter, it displays Semantic Boxes for the current scale center, its parent, the next child along this chain, and any ray-preselected object. These boxes provide local hierarchical context and semantic extents for scaling. The relevant part of the chain is updated when the object or region at the scale center changes. Boxes outside the current view can be inspected by recentering or changing scale.

Numerical Boxes provide metric extents around the scale center. In the prototype, their side lengths follow powers of ten, 10^k meters, so adjacent boxes differ by one order of magnitude. They provide scale references in intervals where a semantic parent or child is unavailable or less useful.

The *Metric Extent Readout* is attached to the Scale-Frame and continuously reports its real-world side length. Unlike Numerical Boxes, which provide discrete order-of-magnitude references, the readout provides continuous feedback for the currently selected extent.

3.3 Scale-Frame Scaling

Figure 1 illustrates scale-in and scale-out. During manual selection, the right thumbstick resizes the Scale-Frame while the WIM preview stays at its current scale. When the Scale-Frame reaches a trigger boundary, manual resizing pauses and the system automatically enlarges or shrinks the preview content. The WIM’s outer frame stays fixed. During the transition, the Scale-Frame remains visible and continues to represent the same selected real-world side length, as shown by its unchanged Metric Extent Readout.

For scale-in, users shrink the Scale-Frame until it reaches an internal Semantic Box or Numerical Box. The system then enlarges the WIM content until the target box matches the size of the fixed WIM frame. For a Semantic Box, the system also places the corresponding object or region at the WIM center. For example, when the Scale-Frame reaches Earth’s Semantic Box, Earth is centered and enlarged to fill the WIM.

For scale-out, users expand the Scale-Frame until it reaches the fixed WIM frame. The system then shrinks the WIM content, bringing a larger surrounding extent into view, until the parent Semantic Box or next outer Numerical Box matches the size of the WIM frame. For a parent Semantic Box, the system also places the corresponding parent object or region at the WIM center.

After alignment, the preview is held briefly before a continuing thumbstick input resumes Scale-Frame resizing (Section 4.4). Users can also release the thumbstick between predefined boundaries. On the next inward input, the system first enlarges the preview to make this intermediate extent fill the WIM frame, then resumes resizing. This allows users to continue from an extent they selected themselves.

3.4 Preview and Teleport

ScaleCue-WIM supports multiscale navigation through a preview-and-commit workflow. Users configure a candidate immersive view in the WIM, including its level of scale (LoS), position, and viewing direction. They then commit the previewed configuration by teleporting to the corresponding immersive view. Teleportation changes the viewpoint discretely, reducing continuous optic flow during travel, but it can disrupt spatial orientation and updating [40]. ScaleCue-WIM lets users inspect the destination configuration before teleporting, providing a preview of the upcoming transition.

4 User Study

We conducted an exploratory within-subjects study comparing ScaleCue-WIM and CASSWIM during guided cross-scale route execution. We examined participant-level completion time, scale-operation count, and perceived scale support; both conditions are described in Section 4.3.

The study addressed two research questions:

- **RQ1:** How do completion time and scale-operation count differ between ScaleCue-WIM and CASSWIM during guided cross-scale route execution?
- **RQ2:** How does perceived scale support during WIM-preview adjustment differ between ScaleCue-WIM and CASSWIM?

4.1 Participants

We recruited 24 participants (14 male, 10 female; mean age 25.5, $SD = 1.8$, range 22–30) from a local university community. All participants had normal or corrected-to-normal vision and reported no vestibular disorders. Before the experiment, participants completed a background questionnaire about prior VR experience (7-point; $M = 4.1$, $SD = 1.5$) and self-rated familiarity with biological structures ($M = 2.0$, $SD = 1.0$), geographic and human-scale maps ($M = 3.4$, $SD = 1.3$), and astronomical and solar-system content ($M = 2.3$, $SD = 1.2$). The sample size of 24 supported a fully crossed counterbalancing of condition order and task-set assignment (four order-by-set sequences, six participants each). The study was approved by the university’s ethics review process, and all participants provided informed consent.

4.2 Apparatus and Controls

The study used a Meta Quest 3 head-mounted display with two Meta Quest Touch Plus controllers (2064×2208 pixels per eye, ~110°




Figure 2: Main visual components of ScaleCue-WIM. The WIM frame defines the boundary of the miniature world, while the Scale-Frame is used to select an extent around the scale center. Labels and Semantic Boxes provide semantic references for objects and regions. Numerical Boxes mark power-of-ten metric extents, and the Metric Extent Readout reports the real-world side length currently selected by the Scale-Frame.

horizontal field of view, up to 120 Hz). The application ran via a PC-tethered link from a desktop (Intel Core i9-14900K, NVIDIA RTX 4090D, Windows 11) and was implemented in Unity (2022.3.62f1). The headset view was mirrored to an external monitor so that the experimenter could observe task progress.

Participants used the left controller to manipulate the WIM content. Pressing the X button activated rotation: the left thumbstick controlled rotation about the x - and y -axes, while holding the left trigger and moving the thumbstick controlled rotation about the z -axis. Pressing the Y button activated translation: the thumbstick moved the content forward, backward, left, or right, while holding the trigger and moving it up or down moved the content vertically. Rotation and translation rates were $30^\circ/\text{s}$ and 2 Unity units/s, respectively (2 m/s under the scene’s 1-unit-to-1-m convention). Each active directional input moved or rotated the WIM content at the fixed rate for that axis; increasing thumbstick deflection did not increase the rate. The right thumbstick scaled the WIM content in CASSWIM or the Scale-Frame in ScaleCue-WIM. Pressing the right trigger selected the ray target; selecting a label also recentered its associated object or region. The A button committed the configured preview and teleported the participant. A trial ended when the ray intersected the active red checkpoint, with no separate confirmation input.

4.3 Conditions

We compared two matched WIM preview-and-teleport workflows. Both used the same MSVE, types of scale cues, ray selection, user-controlled recentering, manual 6-DoF manipulation of the WIM content, and teleport command. Participants could translate and rotate the WIM content to center an unlabeled region or refine the preview.

CASSWIM. CASSWIM extended SSWM [48] with the same types of scale cues used in ScaleCue-WIM. Participants directly scaled the WIM content. Labels, Semantic Boxes, Numerical Boxes, and the Metric Extent Readout attached to the WIM frame served as visual references. The displayed Semantic Boxes corresponded to the object closest to the scale center and its related parent and child along the scaling chain. Continuous scaling updated this display without automatically recentering the WIM content.

ScaleCue-WIM. Participants used the shared selection and manual recentering controls but adjusted scale by resizing the Scale-Frame. The Metric Extent Readout showed the selected extent. Contact with an internal box or the fixed WIM frame triggered an automatic preview transition, including object or region recentering when a Semantic Box was the target. Labels remained available for selection and recentering.

4.4 Scaling Speed and Animation

Both techniques used the same manual scale mapping: each right-thumbstick update changed the controlled extent by a factor of $s = 1.1$, with a minimum interval of 0.05 s between updates. We selected these values during prototyping to provide a controllable scaling rate.

The physical WIM remained fixed at 1 m per side. With the implemented scale-index range $-213 \leq n \leq 320$ and $L(n) = 1.1^n$ m, it represented extents from approximately 1.53×10^{-9} m (1.53 nm) to 1.76×10^{13} m (about 118 AU).

CASSWIM applied each accepted manual update directly to the WIM preview, whereas ScaleCue-WIM applied it to the Scale-Frame. During automatic alignment, ScaleCue-WIM multiplied the size of the preview content by 1.1 per step for scale-in and divided it by 1.1 per step for scale-out, with at most one step per rendered frame. The transition stopped when the target box matched the size of the fixed WIM frame. Its duration depended on the number of steps needed to reach the target and the frame timing. The aligned preview was then held for 0.2 s before a continuing thumbstick input resumed Scale-Frame resizing. Completion time included the entire transition and this hold.

4.5 Tasks and Environments

Participants completed guided routes in a single MSVE spanning biological, human-scale, and astronomical scenes. An always-visible text panel above the WIM listed the destination and intermediate route points in order. Figure 3 shows the ScaleCue-WIM workspace and the technique-neutral route guidance used in the study. After participants visited these points, the final red-sphere checkpoint became active. A trial ended when the ray intersected it.

We used three route types: cell-to-planet routes scaled outward, planet-to-cell routes scaled inward, and cell-to-cell routes scaled




Figure 3: Unity application view in the ScaleCue-WIM condition. The text panel above the WIM lists the route points in order. Both conditions used this form of technique-neutral route guidance.

outward to a geographic level and then inward to another cell. The last type also required a user-selected change of scale center.

We created two matched six-trial task sets, each containing two routes of each type. Trial order was randomized for each participant. The sets matched the required label-based center selections, scale span, and scale-in and scale-out phases, but used different routes to reduce learning effects.

Condition order and task-set assignment were fully crossed and counterbalanced across the 24 participants using four order-by-set sequences. This scheme balanced which technique was experienced first and ensured that each task set appeared equally often with each technique and in each presentation position.

4.6 Objective Measures

Completion time was measured from the first WIM interaction after the participant finished reading the route description until the ray intersected the final checkpoint. Route-reading time was excluded.

Scale-operation count recorded separate user-initiated scale inputs: direct WIM-scaling inputs in CASSWIM and Scale-Frame resizing inputs in ScaleCue-WIM. A continuous thumbstick hold counted as one operation, including when it continued through automatic alignment. Releasing and resuming the input counted as a new operation. The two measures therefore captured route duration and the number of separate scale-control actions under the instruction to finish as quickly as possible.

4.7 Subjective Measures

After each condition block, participants completed four subjective instruments. We used seven study-specific 7-point items to assess perceived scale support: awareness of the current LoS, judgment of real-world extent, understanding of scale changes, use of scale cues, confidence in WIM scale control, judgment of when to stop scaling, and avoidance of unintended LoS. Higher ratings indicate stronger perceived scale support.

We additionally administered the System Usability Scale (SUS) [5] for general usability and NASA-TLX [13] for perceived workload. NASA-TLX items were rated on a 7-point scale and averaged without weighting. We also administered an adapted 16-item SSQ

symptom checklist [18] before the study and after each condition block to monitor simulator sickness. Participants rated each SSQ symptom on a 7-point scale. We analyzed the unweighted symptom sum (range 16–112) as a monitoring measure. At the end of the study, a forced-preference questionnaire asked which interface participants preferred overall and which interface better supported scale awareness, extent judgment, WIM scale control, judging when to stop scaling, and route-execution efficiency. The semi-structured interview then asked participants to explain their preferences and describe how they used the semantic and numerical cues in each technique. The perceived scale-support items and forced-preference questions are provided in the supplementary material.

4.8 Procedure

Participants completed the background questionnaire and pre-study SSQ before the experimenter introduced LoS, the route panel, the controls, and both WIM workflows. They then practiced in a separate training scene.

For the main study, participants were instructed to follow the displayed route points in order and complete each trial as quickly as possible. They completed two six-trial condition blocks. After each block, they completed the scale-support questionnaire, SUS, NASA-TLX, and SSQ. The forced-preference questionnaire and semi-structured interview followed the second block. The session lasted approximately 75–90 minutes.

4.9 Analysis

For each condition, we averaged each participant's six heterogeneous but balanced routes, giving every participant equal weight in the condition comparison. We compared these participant-level means using two-sided paired-samples *t*-tests and report mean differences, 95% confidence intervals, and standardized within-participant effect sizes (d_z) [23]. Here, $d_z = \bar{D}/s_D$, where D is each participant's ScaleCue-WIM-minus-CASSWIM difference. The paired differences showed no clear departure from normality (completion time: Shapiro-Wilk $W = .961$, $p = .461$; scale-operation count: $W = .956$, $p = .355$). We also examined the Condition $\times$ Route-type interaction using a two-way repeated-measures ANOVA and applied Greenhouse-Geisser correction to both outcome analyses. The Greenhouse-Geisser epsilons were .775 for completion time and .848 for scale-operation count. Route-type paired-samples *t*-tests were exploratory and uncorrected for multiplicity.

Questionnaire outcomes were analyzed as within-participant condition-block measures. SUS was scored using the standard procedure, and the study-specific scale-support items were averaged into one composite per condition. The paired differences showed no clear departure from normality (Shapiro-Wilk: NASA-TLX, $W = .953$, $p = .320$; SUS, $W = .945$, $p = .210$; perceived scale support, $W = .979$, $p = .869$). We compared these three outcomes using two-sided paired-samples *t*-tests with Holm correction [16]. SSQ symptom sums were analyzed separately as a monitoring measure using a two-sided Wilcoxon signed-rank test because paired differences were strongly non-normal, with the rank-biserial correlation reported as the accompanying effect size [19]. Forced-preference responses were tested against chance using two-sided binomial tests. The aspect preference count summarized the number of the five

Table 1: Participant-level objective and subjective outcomes ($N = 24$). Values are mean (SD), except for overall preference (participants choosing each technique); differences are ScaleCue-WIM minus CASSWIM. NASA-TLX and scale support use 1–7 scales; SSQ symptom sum is an unweighted 16-item adapted score (16–112). †: Holm-adjusted across NASA-TLX, SUS, and scale support. SSQ uses a two-sided Wilcoxon signed-rank test as a simulator-sickness monitoring measure; the remaining non-preference outcomes use paired-samples t -tests. Overall-preference p uses a two-sided binomial test.

Measure	CASSWIM	ScaleCue-WIM	Diff. [95% CI]	d_z	p
<i>Objective outcomes</i>					
Completion time (s)	84.82 (13.65)	77.44 (9.46)	-7.37 [-11.20, -3.54]	-0.81	< .001
Scale operations	23.30 (7.80)	18.64 (6.45)	-4.66 [-6.47, -2.85]	-1.09	< .001
<i>Subjective outcomes and preference</i>					
NASA-TLX	2.31 (0.62)	1.51 (0.53)	-0.79 [-1.02, -0.56]	-1.45	< .001†
SUS	77.81 (9.76)	86.67 (8.68)	8.85 [2.96, 14.74]	0.63	.005†
Scale support	3.23 (0.75)	5.32 (0.61)	2.10 [1.75, 2.44]	2.57	< .001†
SSQ symptom sum	21.58 (7.53)	20.13 (6.05)	—	—	.037
Overall preference	1/24	23/24	—	—	< .001

aspect-specific items favoring each technique, excluding overall preference.

One author transcribed and coded the interview recordings using a hybrid deductive-inductive process. The initial coding framework followed the six topics covered by the interview guide: scale cues and awareness, Scale-Frame anchoring and rescaling, interruptions during continuous scaling, navigation and task support, learning and ease of use, and freedom and control. Complete responses were the coding units, and one response could receive multiple codes. During analysis, the initial codes were refined using more specific patterns in the responses and then grouped into the four themes reported below. The resulting themes were checked against the interview summaries. The complete semi-structured interview guide is provided in the supplementary material.

5 Results

Figure 4 provides an overview of the objective and subjective outcomes at the participant level.

5.1 Objective Results

Table 1 summarizes the objective and subjective outcomes aggregated at the participant level. Negative differences indicate lower values for ScaleCue-WIM than CASSWIM. ScaleCue-WIM reduced the mean completion time from 84.82 s to 77.44 s. A two-sided paired-samples t -test on participant-level block means showed a significant reduction, $t(23) = -3.98$, $p < .001$, $d_z = -0.81$, with a mean paired difference of -7.37 s, 95% CI $[-11.20, -3.54]$. The paired-difference SD was 9.07 s, and the within-participant condition correlation was $r = .750$.

ScaleCue-WIM also reduced scale-operation count from 23.30 to 18.64 operations per trial. A two-sided paired-samples t -test showed a significant reduction, $t(23) = -5.32$, $p < .001$, $d_z = -1.09$, with a mean paired difference of -4.66 operations, 95% CI $[-6.47, -2.85]$. The paired-difference SD was 4.29 operations, and the within-participant condition correlation was $r = .835$.




Figure 4: Participant-level outcomes ($N = 24$). Lines in (a) and (b) connect each participant's means across six routes per condition. The other panels show individual questionnaire scores or aspect preference counts. Boxes show quartiles and medians, whiskers extend to the most extreme observations within 1.5 interquartile ranges, and open diamonds show means. All observations are plotted. Brackets indicate paired-samples t -tests ($p < .01$, *** $p < .001$), with Holm correction for the three questionnaire comparisons. Aspect preference count excludes overall preference and is descriptive only.**

Route-type paired-samples t -tests (Table 2, part (a)) were exploratory and uncorrected for multiplicity. They showed lower scale-operation counts with ScaleCue-WIM in all three route types (d_z between -0.48 and -0.65 ; all $p < .05$), whereas the completion-time reduction reached significance only for cell-to-planet routes ($t(23) = -4.88$, $p < .001$, $d_z = -1.00$), not for cell-to-cell ($d_z = -0.37$, $p = .081$) or planet-to-cell routes ($d_z = -0.12$, $p = .557$). The Condition $\times$ Route-type interaction was not significant for either measure (Table 2, part (b)).

5.2 Subjective Results

The seven-item perceived scale-support composite showed acceptable internal consistency for CASSWIM (Cronbach's $\alpha = .73$) and

Table 2: Objective outcomes by route type at the participant level ($N = 24$, two routes per type per condition). Part (a) reports exploratory route-type paired-samples t -tests (ScaleCue-WIM minus CASSWIM; negative values indicate shorter completion times or fewer scale operations under ScaleCue-WIM; uncorrected for multiplicity). Part (b) reports the Condition $\times$ Route-type interaction with Greenhouse-Geisser-corrected degrees of freedom and p values.

Measure	Route type	CASSWIM	ScaleCue-WIM	Diff. [95% CI]	d_z	p
<i>(a) Route-type paired-samples t-tests</i>						
Time (s)	Cell-to-planet	76.93 (34.67)	64.05 (28.40)	-12.88 [-18.33, -7.42]	-1.00	< .001
	Planet-to-cell	87.59 (21.19)	85.38 (15.44)	-2.21 [-9.89, 5.47]	-0.12	.557
	Cell-to-cell	89.93 (14.48)	82.90 (15.54)	-7.03 [-14.99, 0.94]	-0.37	.081
Scale ops	Cell-to-planet	19.92 (14.91)	16.04 (16.02)	-3.88 [-6.41, -1.34]	-0.65	.004
	Planet-to-cell	25.81 (7.49)	22.40 (7.06)	-3.42 [-6.41, -0.42]	-0.48	.027
	Cell-to-cell	24.17 (8.18)	17.48 (5.94)	-6.69 [-11.02, -2.35]	-0.65	.004
<i>(b) Condition $\times$ Route type interaction (two-way RM-ANOVA)</i>						
Time (s)	$F(1.55, 35.65) = 2.26, \eta_p^2 = .09$					.130
Scale ops	$F(1.70, 38.99) = 1.10, \eta_p^2 = .05$					.333

high internal consistency for ScaleCue-WIM ($\alpha = .91$). Participants reported substantially stronger perceived scale support with ScaleCue-WIM: the composite increased from 3.23 with CASSWIM to 5.32 on the 7-point scale, $t(23) = 12.61$, $p_{\text{Holm}} < .001$, $d_z = 2.57$. The paired-difference SD was 0.81, and the within-participant condition correlation was $r = .293$.

SUS scores also favored ScaleCue-WIM (86.67 vs. 77.81), $t(23) = 3.11$, $p_{\text{Holm}} = .005$, $d_z = 0.63$, and the unweighted NASA-TLX score was lower with ScaleCue-WIM (1.51 vs. 2.31 on the 7-point scale), $t(23) = -7.11$, $p_{\text{Holm}} < .001$, $d_z = -1.45$. Paired-difference SDs were 13.95 for SUS and 0.55 for NASA-TLX; the corresponding within-participant condition correlations were $r = -.141$ and $r = .561$. SUS and NASA-TLX provided complementary measures of usability and workload. In this comparison, stronger perceived scale support with ScaleCue-WIM was accompanied by higher usability ratings and lower workload ratings.

In the forced-preference questionnaire, 23 of 24 participants preferred ScaleCue-WIM overall ($p < .001$, two-sided binomial test). ScaleCue-WIM was also selected more often for the preference items on scale awareness, extent judgment, WIM scale control, judging when to stop scaling, and route-execution efficiency. Figure 4 summarizes the five aspect-specific responses as aspect preference counts.

The adapted SSQ symptom sums were lower after ScaleCue-WIM (CASSWIM: $M = 21.58$, $SD = 7.53$; ScaleCue-WIM: $M = 20.13$, $SD = 6.05$). Because the paired differences were non-normal and 11 participant pairs showed no difference, we evaluated this monitoring measure using a two-sided Wilcoxon signed-rank test. The test indicated lower symptom sums after ScaleCue-WIM, $W = 16$, $p = .037$, with a rank-biserial correlation of $r_{\text{rb}} = -.65$.

5.3 Interview Findings

The interviews contextualized the quantitative results through four themes: operationalized scale cues, Scale-Frame anchoring, monitoring during continuous scaling, and trade-offs in learning and freedom.

Operationalized scale cues. Participants described linking automatic preview changes to the Semantic and Numerical Boxes defining the resulting extent. They interpreted transitions to semantic extents as movement to another level in the semantic hierarchy. Numerical Boxes provided repeated reference points at regular order-of-magnitude intervals, which participants described as helping them follow scale changes as they continued scaling. In CASSWIM, participants found it harder to connect the boxes with changes in the preview when the boxes served only as visual references. P2, P4, P7, P10, P13, and P24 noted that they focused on completing the route and often paid little attention to the visible boxes.

Scale-Frame anchoring. P1, P5, P8, P10, P14, and P24 described the Scale-Frame as a stable reference during automatic preview scaling. The frame continued to represent the same selected real-world side length, with an unchanged value in the Metric Extent Readout, while the WIM content changed. Participants reported using this reference to relate the current preview to the preceding scale range and to judge how much and in which direction the scale had changed.

Monitoring during continuous scaling. P7, P18, and P20 reported that continuous scaling through unfamiliar ranges in CASSWIM made it difficult to tell how much the scale had changed or whether they had reached the intended level. They described pausing to assess the current scale, compare it with their target, and decide whether to continue scaling. With ScaleCue-WIM, they felt better supported in understanding their current scale and following changes as scaling continued. They reported less confusion and fewer pauses to reassess the scale.

Learning and freedom trade-offs. P3, P5, P9, P16, P17, P19, and P20 reported needing time to understand the Scale-Frame and boundary-controlled workflow. They described CASSWIM as more direct because it allowed continuous adjustment, although they perceived it as offering less scale support. P12 considered ScaleCue-WIM unnecessarily complex for the study tasks. P14 and P19 also viewed predefined boundaries as less flexible; P14 suggested a customizable Numerical Box density. Participants who discussed the Metric Extent Readout treated it as secondary to the boxes and Scale-Frame. P7 noted that the tasks did not require exact metric judgments.

6 Discussion

The central design question is how semantic and numerical references can support users in understanding scale as they inspect and adjust a WIM preview. ScaleCue-WIM connects these references to scaling and keeps the Scale-Frame visible during preview transitions. Participants completed routes faster and used fewer scale-control operations with this workflow, while reporting stronger perceived scale support.

The mean time saving was 7.37 s per route (8.7%), alongside 4.66 fewer scale-control operations (20.0%). Participants therefore started scale-control inputs fewer times while completing routes more quickly. The value of scale support may also depend on the range a user is navigating. Users may move confidently through familiar scales and need more guidance in unfamiliar ones. Support

may therefore be useful at different moments and to different degrees for different users. Subjective ratings and interview accounts help explain this value alongside the route-level efficiency gains.

Participants reported stronger perceived scale support, higher usability, and lower workload with ScaleCue-WIM. Both workflows provided the same types of semantic and numerical cues, but participants described using them differently. In ScaleCue-WIM, they found the boxes easier to notice and use when the boxes guided scaling, and used the Scale-Frame to relate the preceding and resulting scale ranges. In CASSWIM, some reported overlooking the boxes while following the route or pausing to inspect unfamiliar ranges. These accounts suggest that connecting scale references to preview adjustment helped participants interpret scale changes. For WIM design, this points to making the relationship between the selected extent, the scaling action, and the resulting preview easy to follow. References that users can consult across these steps may support both choosing a view and understanding how it changes.

The interviews also point to a trade-off between guidance and direct control. Some participants needed time to learn the Scale-Frame, and some preferred the directness of continuous scaling. This suggests a design that lets users switch between boundary-controlled transitions and continuous preview scaling for local adjustments. Such a design would retain guidance between meaningful extents while giving users more control when the desired view falls between boundaries.

7 Limitations and Future Work

First, we used prescribed routes to give participants comparable task requirements and focus the comparison on scale control. The always-visible route description reduced route-discovery and memory demands. Future work will extend the evaluation to open-ended tasks by removing the prescribed route guidance and letting users choose destinations and routes through the existing label selection, recentring, and preview controls. Exact extent-matching tasks would also help assess the value of the Metric Extent Readout.

Second, the study compared integrated workflows rather than isolating the contribution of the Scale-Frame, operational boundaries, automatic alignment, or their visual feedback. The comparison includes differences in automatic transition timing as well as participants' actions. Factorial comparisons are needed to determine which elements are necessary and how they interact.

Third, scale-operation count records separate scale-control inputs but does not identify why an input ended. Planned adjustments, corrections, inspection pauses, and individual strategies can therefore produce similar counts. Completion time summarizes the whole route, and the questionnaires summarize each condition block. We did not measure when support was useful within individual scale ranges. Combining event logs with brief reports at selected scale ranges would help identify these moments and relate them to users' familiarity with each range.

The perceived scale-support composite measures self-reported support rather than objective scale awareness and requires further validation. Interview themes were developed by one author without independent coding. The interview prompts elicited comparisons of scale support, which may also have shaped the emphasis of participants' responses.

Finally, each route type contained only two routes per condition, and neither Condition $\times$ Route-type interaction was significant. A wider range of routes is needed to examine variation across task demands. Although condition order and task-set assignment were counterbalanced, each participant encountered each condition and task set only once; possible order and task-set effects cannot be fully excluded. The study also used one prototype MSVE and a young university sample with moderate VR experience. Future work should test broader populations and environments.

8 Conclusion

We presented ScaleCue-WIM, a workflow designed to support users' scale judgments as they inspect and adjust a WIM preview before teleportation. It combines a central Scale-Frame with semantic and numerical references for extent adjustment and automatic preview transitions. In an exploratory within-subjects study with 24 participants, ScaleCue-WIM supported faster guided cross-scale route execution with fewer scale-control operations than CASSWIM. Participants reported stronger perceived scale support, higher usability, and lower workload. The findings motivate further evaluation of the workflow's benefits and control trade-offs in open-ended multiscale navigation.

References

- [1] Ferran Argelaguet and Morgant Maignant. 2016. GiAnt: stereoscopic-compliant multi-scale navigation in VEs. In *Proceedings of the 22nd ACM Conference on Virtual Reality Software and Technology*. Association for Computing Machinery, New York, NY, USA, 269–277. doi:10.1145/2993369.2993391
- [2] Felipe Bacim, Doug Bowman, and Marcio Pinho. 2009. Wayfinding Techniques for MultiScale Virtual Environments. In *2009 IEEE Symposium on 3D User Interfaces*. IEEE, 67–74. doi:10.1109/3DUI.2009.4811207
- [3] Benjamin B Bederson and James D Hollan. 1994. Pad++: A Zooming Graphical Interface for Exploring Alternate Interface Physics. In *Proceedings of the 7th annual ACM symposium on User interface software and technology*. ACM Press, New York, New York, USA, 17–26. doi:10.1145/192426.192435
- [4] Doug A. Bowman, David Koller, and Larry F. Hodges. 1997. Travel in Immersive Virtual Environments: An Evaluation of Viewpoint Motion Control Techniques. In *Proceedings of IEEE 1997 Annual International Symposium on Virtual Reality*. IEEE Comput. Soc. Press, 45–52. doi:10.1109/VRAIS.1997.583043
- [5] John Brooke. 1996. SUS: A 'Quick and Dirty' Usability Scale. In *Usability Evaluation in Industry*, Patrick W. Jordan, Bruce Thomas, Bernard A. Weerdmeester, and Ian L. McClelland (Eds.). Taylor & Francis, 189–194. https://www.researchgate.net/publication/319394819_SUS_-_a_quick_and_dirty_usability_scale
- [6] Isaac Cho, Jialei Li, and Zachary Wartell. 2018. Multi-Scale 7DOF View Adjustment. *IEEE Transactions on Visualization and Computer Graphics* 24, 3 (2018), 1331–1344. doi:10.1109/TVCG.2017.2668405
- [7] Sebastian Cmentowski, Andrey Krekhov, and Jens Krüger. 2019. Outstanding: A Multi-Perspective Travel Approach for Virtual Reality Games. In *Proceedings of the Annual Symposium on Computer-Human Interaction in Play*. Association for Computing Machinery, New York, NY, USA, 287–299. doi:10.1145/3311350.3347183
- [8] Kurtis Danyluk, Barrett Ens, Bernhard Jenny, and Wesley Willett. 2021. A Design Space Exploration of Worlds in Miniature. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–15. doi:10.1145/3411764.3445098
- [9] Markus Dräger and Alexander Koller. 2012. Generation of landmark-based navigation instructions from open-source data. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, Avignon, France, 757–766. <https://aclanthology.org/E12-1077/>
- [10] Chi-Wing Fu and Andrew J Hanson. 2007. A transparently scalable visualization architecture for exploring the universe. *IEEE Transactions on Visualization and Computer Graphics* 13, 1 (2007), 108–121. doi:10.1109/TVCG.2007.2
- [11] George W Furnas and Benjamin B Bederson. 1995. Space-scale diagrams: Understanding multiscale interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM Press, New York, New York, USA, 234–241. doi:10.1145/223904.223934

- [12] Andrew J Hanson, Chi-Wing Fu, and Eric A Wernert. 2000. Very large scale visualization methods for astrophysical data. In *Data Visualization 2000 (Eurographics)*. Springer, Vienna, 115–124. doi:10.1007/978-3-7091-6783-0_12
- [13] Sandra G Hart and Lowell E Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. In *Human Mental Workload*, Peter A. Hancock and Najmedin Meshkati (Eds.). Advances in Psychology, Vol. 52. Elsevier, 139–183. doi:10.1016/S0166-4115(08)62386-9
- [14] Stephen C. Hirtle and John Jonides. 1985. Evidence of hierarchies in cognitive maps. *Memory & Cognition* 13, 3 (1985), 208–217. doi:10.3758/BF03197683
- [15] Carol S. Holding. 1992. Clusters and reference points in cognitive representations of the environment. *Journal of Environmental Psychology* 12, 1 (1992), 45–55. doi:10.1016/S0272-4944(05)80296-8
- [16] Sture Holm. 1979. A Simple Sequentially Rejective Multiple Test Procedure. *Scandinavian Journal of Statistics* 6, 2 (1979), 65–70. <https://www.jstor.org/stable/4615733>
- [17] Edwin L. Hutchins, James D. Hollan, and Donald A. Norman. 1986. Direct Manipulation Interfaces. In *User Centered System Design: New Perspectives on Human-Computer Interaction*, Donald A. Norman and Stephen W. Draper (Eds.). Lawrence Erlbaum Associates, Hillsdale, NJ, 87–124.
- [18] Robert S. Kennedy, Norman E. Lane, Kevin S. Berbaum, and Michael G. Lilienthal. 1993. Simulator sickness questionnaire: An enhanced method for quantifying simulator sickness. *The International Journal of Aviation Psychology* 3, 3 (1993), 203–220. doi:10.1207/s15327108ijap0303_3
- [19] Dave S. Kerby. 2014. The Simple Difference Formula: An Approach to Teaching Nonparametric Correlation. *Comprehensive Psychology* 3, Article 1 (2014), 9 pages. doi:10.2466/11.IT.3.1
- [20] Regis Kopper, Tao Ni, Doug A. Bowman, and Marcio Pinho. 2006. Design and Evaluation of Navigation Techniques for Multiscale Virtual Environments. In *IEEE Virtual Reality Conference (VR 2006)*. IEEE, 175–182. doi:10.1109/VR.2006.47
- [21] David Kouřil, Ladislav Čmólik, Barbora Kozlíková, Hsiang-Yun Wu, Graham Johnson, David S. Goodsell, Arthur Olson, M. Eduard Gröller, and Ivan Viola. 2019. Labels on levels: Labeling of multi-scale multi-instance and crowded 3D biological environments. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2019), 977–986. doi:10.1109/TVCG.2018.2864491
- [22] David Kouřil, Tobias Isenberger, Barbora Kozlíková, Miriah Meyer, M. Eduard Gröller, and Ivan Viola. 2021. HyperLabels: Browsing of dense and hierarchical molecular 3D models. *IEEE Transactions on Visualization and Computer Graphics* 27, 8 (2021), 3493–3504. doi:10.1109/TVCG.2020.2975583
- [23] Daniel Lakens. 2013. Calculating and Reporting Effect Sizes to Facilitate Cumulative Science: A Practical Primer for *t*-Tests and ANOVAs. *Frontiers in Psychology* 4, Article 863 (2013), 12 pages. doi:10.3389/fpsyg.2013.00863
- [24] Eike Langbehn, Gerd Bruder, and Frank Steinicke. 2016. Scale matters! Analysis of dominant scale estimation in the presence of conflicting cues in multi-scale collaborative virtual environments. In *2016 IEEE Symposium on 3D User Interfaces (3DUI)*. IEEE, 211–220. doi:10.1109/3DUI.2016.7460054
- [25] Joseph J. LaViola, Jr., Daniel Acevedo Feliz, Daniel F. Keefe, and Robert C. Zeleznik. 2001. Hands-free multi-scale navigation in virtual environments. In *Proceedings of the 2001 Symposium on Interactive 3D Graphics*. Association for Computing Machinery, New York, NY, USA, 9–15. doi:10.1145/364338.364339
- [26] Joseph J. LaViola, Jr., Ernst Kruijff, Ryan P. McMahan, Doug A. Bowman, and Ivan Poupyrev. 2017. *3D User Interfaces: Theory and Practice* (2 ed.). Addison-Wesley Professional.
- [27] Jong-In Lee, Paul Asente, and Wolfgang Stuerzlinger. 2023. Designing Viewpoint Transition Techniques in Multiscale Virtual Environments. In *2023 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*. IEEE, 680–690. doi:10.1109/VR55154.2023.00083
- [28] Jong-In Lee and Wolfgang Stuerzlinger. 2025. Scaling Techniques for Exocentric Navigation Interfaces in Multiscale Virtual Environments. *IEEE Transactions on Visualization and Computer Graphics* 31, 5 (2025), 2704–2714. doi:10.1109/TVCG.2025.3549535
- [29] Yinggang Li, Chi-Wing Fu, and Andrew J. Hanson. 2006. Scalable WIM: Effective Exploration in Large-scale Astrophysical Environments. *IEEE Transactions on Visualization and Computer Graphics* 12, 5 (2006), 1005–1012. doi:10.1109/TVCG.2006.176
- [30] Kristin L. Lovelace, Mary Hegarty, and Daniel R. Montello. 1999. Elements of Good Route Directions in Familiar and Unfamiliar Environments. In *Spatial Information Theory: Cognitive and Computational Foundations of Geographic Information Science (Lecture Notes in Computer Science, Vol. 1661)*, Christian Freksa and David M. Mark (Eds.). Springer, Berlin, Heidelberg, 65–82. doi:10.1007/3-540-48384-5_5
- [31] James McCrae, Igor Mordatch, Michael Glueck, and Azam Khan. 2009. Multiscale 3D navigation. In *Proceedings of the 2009 Symposium on Interactive 3D Graphics and Games*. Association for Computing Machinery, New York, NY, USA, 7–14. doi:10.1145/1507149.1507151
- [32] Timothy P. McNamara, James K. Hardy, and Stephen C. Hirtle. 1989. Subjective hierarchies in spatial memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 15, 2 (1989), 211–227. doi:10.1037/0278-7393.15.2.211
- [33] Daniel R. Montello. 2005. Navigation. In *The Cambridge Handbook of Visuospatial Thinking*, Priti Shah and Akira Miyake (Eds.). Cambridge University Press, Chapter 7, 257–294. doi:10.1017/CBO9780511610448.008
- [34] Domi Papoi and Wolfgang Stuerzlinger. 2019. Improved Automatic Speed Control for 3D Navigation. In *Advances in Computer Graphics (Lecture Notes in Computer Science, Vol. 11542)*. Springer, Cham, 278–290. doi:10.1007/978-3-030-22514-8_23
- [35] Randy Pausch, Tommy Burnette, Dan Brockway, and Michael E. Weiblen. 1995. Navigation and locomotion in virtual worlds via flight into hand-held miniatures. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*. ACM Press, New York, New York, USA, 399–400. doi:10.1145/218380.218495
- [36] Leonardo Pavanatto, Alexander Giovannelli, Brian Giera, Timo Bremer, Haichao Miao, and Doug A. Bowman. 2025. Exploring Multiscale Navigation of Homogeneous and Dense Objects with Progressive Refinement in Virtual Reality. In *2025 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*. IEEE, 228–237. doi:10.1109/VR59515.2025.00047
- [37] Michael Peer, Iva K. Brunec, Nora S. Newcombe, and Russell A. Epstein. 2021. Structuring Knowledge with Cognitive Maps and Cognitive Graphs. *Trends in Cognitive Sciences* 25, 1 (2021), 37–54. doi:10.1016/j.tics.2020.10.004
- [38] Ken Perlin and David Fox. 1993. Pad: an alternative approach to the computer interface. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*. Association for Computing Machinery, New York, NY, USA, 57–64. doi:10.1145/166117.166125
- [39] Jeffrey S. Pierce and Randy Pausch. 2004. Navigation with Place Representations and Visible Landmarks. In *IEEE Virtual Reality 2004*. IEEE, 173–180. doi:10.1109/VR.2004.1310071
- [40] Aniruddha Prithul, Isayas Berhe Adhanom, and Eelke Folmer. 2021. Teleportation in Virtual Reality: A Mini-Review. *Frontiers in Virtual Reality* 2, Article 730792 (2021), 7 pages. doi:10.3389/frvir.2021.730792
- [41] Mike Scaife and Yvonne Rogers. 1996. External cognition: how do graphical representations work? *International Journal of Human-Computer Studies* 45, 2 (1996), 185–213. doi:10.1006/ijhc.1996.0048
- [42] Richard Stoakley, Matthew J. Conway, and Randy Pausch. 1995. Virtual Reality on a WIM: Interactive Worlds in Miniature. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM Press, New York, New York, USA, 265–272. doi:10.1145/223904.223938
- [43] Daniel R. Trindade and Alberto B. Raposo. 2011. Improving 3D navigation in multiscale environments using cubemap-based techniques. In *Proceedings of the 2011 ACM Symposium on Applied Computing*. Association for Computing Machinery, New York, NY, USA, 1215–1221. doi:10.1145/1982185.1982454
- [44] Ramon Trueba, Carlos Andujar, and Ferran Argelaguet. 2009. Multi-scale Manipulation in Indoor Scenes with the World in Miniature Metaphor. In *Proceedings of the 15th Joint virtual reality Eurographics conference on Virtual Environments*. The Eurographics Association, 93–100. doi:10.2312/EGVE/JVRC09/093-100
- [45] Samuel Truman and Sebastian von Mammen. 2020. An Integrated Design of World-in-Miniature Navigation in Virtual Reality. In *Proceedings of the 15th International Conference on the Foundations of Digital Games*. Association for Computing Machinery, New York, NY, USA, 1–9. doi:10.1145/3402942.3402994
- [46] Colin Ware and Daniel Fleet. 1997. Context sensitive flying interface. In *Proceedings of the 1997 Symposium on Interactive 3D Graphics*. ACM Press, New York, New York, USA, 127–130. doi:10.1145/253284.253319
- [47] Tim Weissker, Matthias Franzgrote, and Torsten Kuhlen. 2024. Try This for Size: Multi-Scale Teleportation in Immersive Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics* 30, 5 (2024), 2298–2308. doi:10.1109/TVCG.2024.3372043
- [48] Chadwick A. Wingrave, Yonca Hacıahmetoglu, and Doug A. Bowman. 2006. Overcoming World in Miniature Limitations by a Scaled and Scrolling WIM. In *3D User Interfaces (3DUI'06)*. IEEE, 11–16. doi:10.1109/VR.2006.106
- [49] Lixiang Zhao, Tobias Isenberger, Fuqi Xie, Hai-Ning Liang, and Lingyun Yu. 2024. MeTACAST: Target- and Context-Aware Spatial Selection in VR. *IEEE Transactions on Visualization and Computer Graphics* 30, 1 (2024), 480–494. doi:10.1109/TVCG.2023.3326517
- [50] Lixiang Zhao, Fuqi Xie, Tobias Isenberger, Hai-Ning Liang, and Lingyun Yu. 2026. ScaleFree: Dynamic KDE for Multiscale Point Cloud Exploration in VR. In *2026 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 358–368. doi:10.1109/VR67842.2026.00057