CheckerVision: Grid-Based Interleaving of Concurrent 3D Environments in VR

Kevin Linne
Technische Hochschule Mittelhessen
Gießen, Germany
kevin.linne@mni.thm.de

Amir Hossein Ebrahimi
Technische Hochschule Mittelhessen
Gießen, Germany
amir.ebrahimi@mni.thm.de

Jan Gugenheimer
TU Darmstadt
Darmstadt, Germany
jan.gugenheimer@tu-darmstadt.de

Martin Weigel
Technische Hochschule Mittelhessen
Gießen, Germany
martin.weigel@mni.thm.de

Figure 1: CheckerVision uses a checkerboard segmentation pattern to interleave two virtual environments (A/B) across the entire field of view while preserving stereoscopic depth perception. We implemented three grid resolutions: coarse (8×8), medium (64×64), and fine (512×512).

Abstract

CheckerVision enables the simultaneous visualization of two virtual environments by segmenting the user's field of view into a checkerboard pattern. Alternating cells are assigned to one of two environments, distributing both worlds across the entire viewport rather than allocating contiguous regions to each. This work contributes (1) a concept for FoV-segmented multi-environment visualization, (2) a Unity-based stereoscopic implementation using layer-based rendering and stencil masking with a runtime-generated checkerboard mask, and (3) an empirical evaluation of how segmentation granularity relates to subjective experience. In a within-subjects study (N=12), participants explored two distinct scenes under three grid resolutions and reported presence (IPQ), usability (SUS), and preference rankings. Across conditions, presence and usability ratings were consistent and did not differ reliably between resolutions; preference rankings likewise showed no significant differences, with a descriptive tendency toward finer segmentation.

DIS Companion '26, Singapore, Singapore

ACM ISBN 979-8-4007-2632-3/2026/06
https://doi.org/10.1145/3802974.3809440

CCS Concepts

• **Human-centered computing** → *Virtual reality.*

Keywords

Virtual Reality, Multi-Environment Interaction

ACM Reference Format:
Kevin Linne, Amir Hossein Ebrahimi, Jan Gugenheimer, and Martin Weigel. 2026. CheckerVision: Grid-Based Interleaving of Concurrent 3D Environments in VR. In *Designing Interactive Systems Conference (DIS Companion '26), June 13–17, 2026, Singapore, Singapore.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3802974.3809440

1 Introduction

Most VR applications present a single coherent virtual environment (VE) at a time. At the same time, many scenarios would benefit from maintaining awareness of more than one VE in parallel, for example when cross-checking reference information, comparing alternative design states, coordinating parallel activities, or moving between task and support spaces. Prior systems typically enable such multi-environment access through explicit transitions, menus, and portal-like mechanisms that reveal another VE on demand [1–3, 5]. These approaches preserve visual fidelity by presenting one VE as a coherent view, but they make cross-environment awareness an *interaction problem*: users must switch, open, or navigate to access the other VE. A complementary line of work explores *simultaneous* visibility by placing multiple VEs within the user's view. This design space can be roughly characterized along two dimensions: (1) whether visibility is provided through *contiguous* FoV regions (e.g., splits) or through *distributed* allocations across the full FoV, and (2) whether the content is presented as a lightweight overview or at a level of visual richness comparable to a full VE. Schjerlund et al.'s Ovrlap exemplifies distributed multi-view visualization by showing several perspectives at once, enabling multi-view awareness but with limited per-view detail [6]. In contrast, FoV partitioning approaches such as split-based layouts allocate contiguous regions to different VEs [4], which can support richer detail within each region but may concentrate each VE into specific FoV areas.

CheckerVision explores an alternative point in this space by combining two ideas from prior work: (i) distributed allocation across the full FoV, and (ii) rendering two visually rich, stereoscopic VEs in parallel. Instead of assigning each VE to a contiguous region, CheckerVision segments the FoV into a checkerboard and alternates cell ownership between two VEs. This yields an interleaved view in which both VEs remain present across the entire viewport. The design builds on the visual system's ability to perceive coherent scenes from incomplete spatial samples, allowing users to maintain a stable impression of each VE despite missing cells. To support different preferences and perceptual trade-offs, CheckerVision exposes segmentation granularity as a controllable parameter. Interaction remains VE-specific through controller-scoped manipulation, enabling object actions in both worlds without an explicit mode switch.

We implement CheckerVision as a Unity-based VR prototype using stencil masking, which allows us to render two VEs in one scene. As a first step, we report an initial within-subject user study (N=12) that compares three segmentation granularities of CheckerVision. Participants explored two VEs and reported presence (IPQ), usability (SUS), and preference rankings. Across conditions, subjective ratings remained consistently high, with no reliable differences between resolutions; preference rankings likewise showed no significant differences, with descriptive tendencies favoring finer segmentation.

Contributions. We contribute (1) a FoV-interleaving concept for simultaneously viewing two stereoscopic, visually rich VEs via checkerboard segmentation, (2) a Unity/URP implementation using runtime-generated mask textures and stencil-based compositing with controller-scoped interaction, and (3) an initial within-subject evaluation across grid granularities using IPQ, SUS, and preference ranking (incl. assumption checks and effect sizes).

2 CheckerVision: Full FoV Visualization of Two VEs

2.1 Concept

CheckerVision targets parallel viewing of multiple VEs while maintaining stereoscopic depth perception. Many multi-environment systems partition the FoV into contiguous regions, for example by splitting it horizontally or vertically [4]. Such layouts make multiple scenes visible at once, but they limit how much of each VE is visible and require users to direct their gaze toward a specific region to inspect one VE in detail. As a result, increasing the visible area of one VE reduces the visible area of the other.

CheckerVision follows the principle of using the FoV in both horizontal and vertical directions simultaneously. It segments the FoV into a regular grid and assigns alternating cells to different VEs, distributing each VE across the entire FoV. This keeps both VEs partially visible independent of gaze direction and avoids a default notion of a primary versus secondary VE. Pilot-style comparisons considered alternative segmentation geometries (squares, circles, triangles and stripes), with triangles and stripes being visually distracting and less structured, and circles leading to inefficient use of screen space due to unused regions. A checkerboard layout was selected because it tiles the FoV without unused space and supports a direct alternation between VEs. The segmentation is parameterized by grid resolution, which controls cell size and therefore the perceived balance between separation and visual interleaving. CheckerVision exposes this parameter as a user-adjustable setting with discrete presets spanning coarse to fine segmentations (e.g., 8×8, 64×64, 512×512).

Interaction follows standard VR grabbing while remaining VE-specific: each controller is mapped to one VE and can only manipulate objects in that VE. Users can also adjust grid resolution to adapt the visualization.

2.2 Implementation

CheckerVision was implemented as a Unity-based VR application for Meta Quest 3 with two separate virtual environments (A and B). All scene objects are assigned to VE-specific rendering layers. Interactable objects additionally use interaction layers to route input per VE.

Stereoscopic visualization is achieved with stencil masking in the Unity Universal Render Pipeline (URP). An initial prototype represented the checkerboard as many individual quads, each writing a stencil ID for its cell. However, performance degraded as the number of quads increased. We therefore implemented a texture-driven stencil shader that encodes the segmentation in a single camera-aligned quad. At runtime, the system generates a mask texture in the headset's display resolution that encodes a checkerboard pattern with alternating red and green cells. The texture is regenerated whenever the user changes the grid resolution, updating segment size and count while keeping rendering cost stable. The mask is applied on a camera-aligned quad parented to the VR camera so it follows head motion. The quad is placed 60 cm in front of the camera and scaled to fill the field of view. During rendering, the mask writes two stencil values per pixel: pixels corresponding to green cells are marked with one value and red cells with another. In subsequent passes, the two VEs are rendered conditionally via stencil tests, such that one VE is drawn only in the "green" cells and the other only in the "red" cells. Because masking and compositing happen within the stereo rendering pipeline, both environments remain stereoscopic.

Interaction is implemented using XR controllers. Grid resolution is adjusted via button bindings, which triggers regeneration of the checkerboard mask. Object manipulation follows standard grab interactions mapped to the grip input, but remains VE-specific by associating each controller with one environment. This mapping enforces a clear separation of input between environments, but may limit more fluid cross-environment interactions or require coordinated use of both hands. It also introduces an additional cognitive demand, as the mapping between controllers and environments is not inherently self-explanatory and must be learned. Locomotion uses joystick-based teleportation; teleport actions update the user position consistently across both environments, while object manipulation remains restricted to the controller's assigned VE.

3 User Study

We conducted an initial user study to examine how CheckerVision's field-of-view segmentation influences subjective experience. The study focuses on **presence**, **usability**, and **preference** across grid resolutions, targeting whether an intermediate granularity provides a balanced experience between perceiving both environments in parallel and maintaining a coherent impression of each environment.

3.1 Research Hypotheses

We expected grid granularity to affect the balance between perceptual separation and visual interleaving of the two environments. Based on pilot-style testing, we considered an intermediate grid (64×64) a plausible compromise: coarse segmentation makes the two VEs appear strongly fragmented, whereas very fine segmentation increases visual interleaving and may reduce perceptual separation. We therefore hypothesized:

- H1 Perceived presence (IPQ) will be highest for 64×64 compared to 8×8 and 512×512.
- H2 Participants will report the highest usability (SUS) in the 64×64 grid condition compared to 8×8 and 512×512.
- H3 64×64 will be ranked first more often than 8×8 and 512×512.

3.2 Method

Twelve individuals participated ($N = 12$; 9 male, 3 female), aged 21–28 years ($M = 24.83$, $SD = 2.29$). Four participants reported using visual correction (two glasses, two contact lenses); eight reported no visual aid. Prior VR experience was limited: four participants reported no prior use and eight reported rare use. The study followed a within-subjects design with three grid-resolution conditions and fully counterbalanced order. Participants explored each condition freely. Afterwards, they completed the Igroup Presence Questionnaire (IPQ), a standardized measure of spatial presence, involvement, and experienced realism in virtual environments. The IPQ is rated on a Likert scale ranging from -3 to 3. Usability was assessed using the System Usability Scale (SUS), a widely used 10-item questionnaire providing a global measure of perceived usability. Scores range from 0 to 100, with higher values indicating better usability. After the final condition, participants provided a ranking and open-ended feedback.

3.3 Data Analysis

We analyzed subjective ratings using repeated-measures statistical tests, as all participants experienced each condition. Prior to inferential analysis, we checked assumptions of normality (Shapiro–Wilk) and sphericity (Mauchly's test) where applicable. For presence (IPQ), we conducted a two-way repeated-measures ANOVA with the within-subject factors *Resolution* (8×8, 64×64, 512×512) and *Environment* (Nature, City). For usability (SUS), we applied a one-way repeated-measures ANOVA with *Resolution* as the within-subject factor. For preference rankings, we used a Friedman test to assess differences between conditions, as ranking data is ordinal. Across all analyses, we report effect sizes alongside test statistics, using a significance threshold of $\alpha = .05$.

4 Results

4.1 Presence (IPQ)

Table 1 reports IPQ overall scores by resolution and environment (including the per-resolution combined score). A two-way repeated-measures ANOVA showed no main effect of Resolution on IPQ scores, and no main effect of Environment. The Resolution × Environment interaction was also not significant. Accordingly, no post-hoc comparisons are reported.

Descriptively, the highest combined mean occurred in the 512×512 condition ($M = 0.402, SD = 1.376$), followed by 64×64 ($M = 0.292, SD = 1.025$) and 8×8 ($M = 0.128, SD = 0.943$; Table 1). Thus, H1 (highest presence in 64×64) was not supported by the inferential analysis, and the descriptive trend did not align with H1. To contextualize absolute score levels, we relate the observed means to published IPQ ranking thresholds on the $[-3, 3]$ scale [7]. Across conditions, means primarily fell into the Low to Moderate range. However, these thresholds were derived from conventional single-environment VR experiences and are only partially applicable to CheckerVision, which deliberately interleaves two environments within the same FoV. Interpreted in this context, presence values near the Low/Moderate boundary indicate that participants still reported a coherent sense of "being there" in each environment despite continuous visual competition and reduced pixel coverage per environment.

This interpretation is also reflected in participants' qualitative feedback. Several participants noted that they could "interact with both worlds" (P-1) and "watch both worlds at the same time" (P-4), suggesting that a sense of presence can be maintained across both environments despite reduced visual completeness.

4.2 User Experience (SUS)

SUS ratings were high across all three resolutions (8×8: $M = 81.9$, $SD = 8.80$; 64×64: $M = 83.8$, $SD = 8.76$; 512×512: $M = 85.8$, $SD = 10.0$). A one-way repeated-measures ANOVA with the within-subject factor *Resolution* (three levels) found no effect of resolution on SUS. H2 (highest usability for 64×64) was not supported. Descriptively, SUS increased slightly with finer segmentation, but differences were small relative to variability, suggesting that grid resolution did not systematically influence perceived usability in this study. Qualitative feedback further supports this interpretation. Participants reported that "the interaction worked" (P-3) and highlighted the ability to manipulate objects across environments (P-7), while also noting challenges such as "disorientation during teleportation" (P-2) or difficulty in perceiving vertical movement (P-2).

Res.	Env.	Overall	SP	INV	REAL	GP
8×8	Nature	0.17 (0.95)	1.05 (1.19)	0.50 (1.59)	-1.35 (0.99)	0.58 (1.38)
8×8	City	0.08 (0.98)	0.92 (1.23)	0.52 (1.58)	-1.54 (1.16)	0.67 (1.23)
8×8	Combined	0.13 (0.94)	0.98 (1.18)	0.51 (1.55)	-1.45 (1.02)	0.62 (1.13)
64×64	Nature	0.32 (1.05)	1.12 (1.04)	0.71 (1.68)	-1.17 (1.53)	0.67 (1.50)
64×64	City	0.27 (1.09)	1.12 (1.07)	0.23 (1.78)	-0.92 (1.61)	0.92 (1.51)
64×64	Combined	0.29 (1.03)	1.12 (0.95)	0.47 (1.58)	-1.04 (1.47)	0.79 (1.28)
512×512	Nature	0.27 (1.39)	0.93 (1.25)	0.65 (2.03)	-0.92 (1.69)	0.17 (1.95)
512×512	City	0.54 (1.43)	1.43 (1.19)	0.79 (2.06)	-0.96 (1.86)	1.00 (1.54)
512×512	Combined	0.40 (1.38)	1.18 (1.15)	0.72 (1.99)	-0.94 (1.72)	0.58 (1.42)
Overall	Nature	0.25 (1.05)	1.03 (0.93)	0.62 (1.73)	-1.14 (1.31)	0.47 (1.35)
Overall	City	0.30 (1.08)	1.15 (1.01)	0.51 (1.71)	-1.14 (1.35)	0.86 (1.20)
Overall	Combined	0.27 (0.99)	1.09 (0.91)	0.56 (1.64)	-1.14 (1.28)	0.67 (1.10)

Table 1: IPQ descriptive statistics on a $[-3, 3]$ scale, reported as mean (SD). *Overall* is the mean of all 14 IPQ items. SP, INV, and REAL denote the IPQ subscales; GP is the single-item general presence rating. *Combined* denotes the per-participant mean of Nature and City within a resolution (and analogously for the overall row).

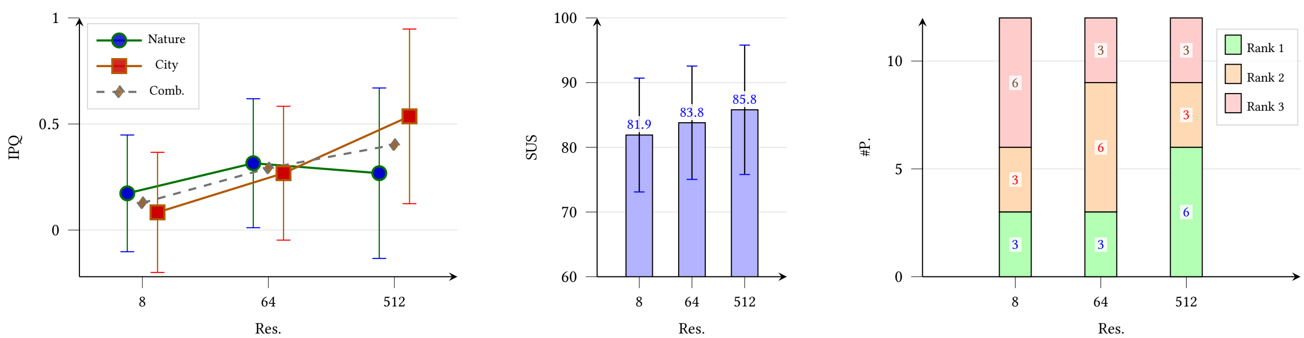

(a) IPQ (mean ± SE). **(b) SUS (mean ± SD).** **(c) Ranking counts.**

Figure 2: Overview of subjective study outcomes across grid resolutions. (a) IPQ overall scores (mean ± SE, $N = 12$). (b) SUS scores (mean ± SD). (c) Preference ranking distribution.

4.3 User Preference

Participants ranked the three grid resolutions from best (1) to worst (3). Descriptively, 512×512 was most frequently ranked first (50%), whereas 8×8 was most frequently ranked last (50%). Mean ranks were 2.25 ($SD = 0.866$) for 8×8, 2.00 ($SD = 0.739$) for 64×64, and 1.75 ($SD = 0.866$) for 512×512. The Friedman test indicated no difference in ranks between conditions. Thus, H3 (64×64 is most often ranked first) was not supported. Overall, rankings suggest a descriptive tendency toward preferring the finest grid, but this tendency was not consistent enough across participants to yield a reliable effect in this sample. Participants' comments also reflect these tendencies. Finer grids were described as making the environments "really overlap, almost both at the same time" (P-6), whereas intermediate resolutions were sometimes perceived as "too many tiles for clear separation, but too few for true blending" (P-6). This highlights the trade-off between perceptual integration and separation across grid resolutions.

5 Limitations and Future Work

Evaluation setting and measures. Participants explored the system without a performance-oriented task. This design aligns with the subjective nature of presence and usability ratings, but it does not reveal how segmentation affects objective behavior (e.g., search, selection, error rates) or longer-term use. Moreover, the hypothesis that an intermediate grid (64×64) provides an optimal balance was primarily motivated by pilot observations rather than perceptual theory. In addition, our evaluation does not include comparisons against alternative multi-environment techniques (e.g., switching or split-view), limiting conclusions about relative advantages over existing approaches. Future studies should include task-driven protocols, baseline comparisons, and combine questionnaires with performance and behavioral measures.

Perceptual asymmetry and content-agnostic segmentation. CheckerVision distributes both environments uniformly across the field of view using a fixed grid that is independent of scene content and task context. However, perceptual balance is influenced not only by spatial allocation but also by scene characteristics such as contrast, spatial frequency, and structural density, as well as task relevance. In our study, the two environments differed in geometry and visual detail, yet we did not observe a significant effect of environment on presence, suggesting a degree of robustness. Nevertheless, environments with stronger visual features or higher task relevance may become perceptually dominant despite identical segmentation. In practice, this raises the question whether grid allocation should adapt to content or task demands, for example by assigning more cells to a task-relevant environment or redistributing cells based on scene statistics. Future work should systematically investigate such adaptive or content-aware segmentation strategies and their impact on perceptual balance.

System constraints and interaction design. The prototype visualizes two VEs with a fixed checkerboard assignment and controller-scoped interaction, and locomotion is mirrored across environments. These choices simplify VE separation, but they also constrain interaction patterns (e.g., cross-VE actions, asymmetric locomotion, or more than two VEs). Future work could extend CheckerVision to adaptive or content-aware segmentation, support additional VEs, and explore interaction techniques that enable controlled cross-VE operations while preserving separation.

6 Conclusion

CheckerVision explores an alternative to contiguous viewport splits by interleaving two stereoscopic virtual environments across the full field of view using a checkerboard segmentation. In a within-subject study comparing three segmentation granularities, subjective presence, usability, and preference showed no reliable differences between conditions, although descriptively participants tended to prefer the finest grid. These results suggest that interleaved FoV segmentation can maintain usable and moderately consistent experiences while enabling parallel visibility of multiple environments, and motivate future work on task-driven evaluations and more flexible segmentation and interaction designs.

Acknowledgments

We thank Kilian Stock for assistance with the stencil shader implementation.

References

[1] Matt Gottsacker, Yahya Hmaiti, Mykola Maslych, Hiroshi Furuya, Gerd Bruder, Gregory F. Welch, and Joseph J. LaViola. 2025. From Alt-Tab to World-Snap: Exploring Different Metaphors for Swift and Seamless VR World Switching. In *2025 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*. 943–944. doi:10.1109/ISMAR-Adjunct68609.2025.00260

[2] Malte Husung and Eike Langbehn. 2019. Of Portals and Orbs: An Evaluation of Scene Transition Techniques for Virtual Reality. In *Proceedings of Mensch Und Computer 2019* (Hamburg, Germany) *(MuC ’19)*. Association for Computing Machinery, New York, NY, USA, 245–254. doi:10.1145/3340764.3340779

[3] Kevin Linne, Sven Thomas, Jennifer Roth, and Martin Weigel. 2025. SliVR: A 360° VR-Hub for Fast Selections in Multiple Virtual Environments. In *2025 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 551–559. doi:10.1109/ISMAR67309.2025.00065

[4] Reiji Miura, Shunichi Kasahara, Michiteru Kitazaki, Adrien Verhulst, Masahiko Inami, and Maki Sugimoto. 2021. MultiSoma: Distributed Embodiment with Synchronized Behavior and Perception. In *Proceedings of the Augmented Humans International Conference 2021* (Rovaniemi, Finland) *(AHs ’21)*. Association for Computing Machinery, New York, NY, USA, 1–9. doi:10.1145/3458709.3458878

[5] Henning Pohl, Klemen Lilija, Jess McIntosh, and Kasper Hornbæk. 2021. Poros: Configurable Proxies for Distant Interactions in VR. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI ’21)*. Association for Computing Machinery, New York, NY, USA, Article 532, 12 pages. doi:10.1145/3411764.3445685

[6] Jonas Schjerlund, Kasper Hornbæk, and Joanna Bergström. 2022. OVRlap: Perceiving Multiple Locations Simultaneously to Improve Interaction in VR. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI ’22)*. Association for Computing Machinery, New York, NY, USA, Article 355, 13 pages. doi:10.1145/3491102.3501873

[7] Tanh Quang Tran, Tobias Langlotz, Jacob Young, Thomas W. Schubert, and Holger Regenbrecht. 2024. Classifying Presence Scores: Insights and Analysis from Two Decades of the Igroup Presence Questionnaire (IPQ). *ACM Trans. Comput.-Hum. Interact.* 31, 5, Article 61 (Nov. 2024), 26 pages. doi:10.1145/3689046