

SeamlessVR: Bridging the Immersive to Non-Immersive Visualization Divide

Shuqi Liao , Sparsh Chaudhri , Maanas K. Karwa , and Voicu Popescu 




Fig. 1: *SeamlessVR* transition from immersive visualization in the headset (left) to non-immersive visualization on screen (right): illustration frames acquired with an XR headset in passthrough mode (top row), and actual frames seen by the user in the headset (bottom row). Once the user engages the transition (a), the 3D visualization morphs to a 2D visualization (a to b to c) that matches what the user sees once they remove the headset (c to d).

Abstract—The paper describes *SeamlessVR*, a method for switching effectively from immersive visualization, in a virtual reality (VR) headset, to non-immersive visualization, on screen. *SeamlessVR* implements a continuous morph of the 3D visualization to a 2D visualization that matches what the user will see on screen after removing the headset. This visualization continuity reduces the cognitive effort of connecting the immersive to the non-immersive visualization, helping the user continue on screen a visualization task started in the headset. We have compared *SeamlessVR* to the conventional approach of directly removing the headset in an IRB-approved user study with $N = 30$ participants. *SeamlessVR* had a significant advantage over the conventional approach in terms of time and accuracy for target tracking in complex abstract and realistic scenes and in terms of participants' perception of the switch from immersive to non-immersive visualization, as well as in terms of usability. *SeamlessVR* did not pose cybersickness concerns.

Index Terms—Transitional interface, cross-reality, VR headset removal.

1 INTRODUCTION

Immersive visualization, such as that provided by virtual reality (VR) headsets, is a uniquely powerful form of interactive visualization. Indeed, immersive visualization provides the user with a natural interface for view selection through head motions, as well as with depth cues, allowing the user to perceive directly, and not just to deduce, the distance to and in between various parts of the scene. Immersive visualization strengthens the user's sense of being present in the 3D scene being visualized, increasing the effectiveness of the interactive visualization application. Recent years have brought leap-forward progress of VR headset technology. We now have all-in-one VR headsets that are completely untethered, with on-board rendering, networking, tracking, and power, and that provide a compelling visual experience, with acceptable

brightness, resolution, field of view, and ergonomics. Multiple trillion-dollar companies are firmly invested in the VR technology space, so the continued, if not accelerated, progress of immersive visualization technology is likely.

We put forth that, at least for the foreseeable future, immersive visualization *will not* completely supplant non-immersive visualization. There are tasks such as reading and writing [15], or such as scrutinizing a 2D image that shows a diagram [43], a photograph, or a rendering of a 3D scene or dataset, which can be performed well using a conventional keyboard+mouse+screen computer interface. The aim is not for immersive visualization to completely *replace* conventional non-immersive visualization, but rather for immersive visualization to be used *in conjunction* with non-immersive visualization.

It is our aim to make the VR headset another essential interface peripheral on a user's desk, alongside the screen, keyboard, and mouse, which the user can grab and use when needed, and only for as long as needed. Unfortunately, immersive visualization is presently used more often as an exceptional event, which the user schedules in advance and to which the user allots an exclusive block of time. For example, a university researcher might walk over to a centralized visualization facility on campus to see their data in VR. A student might use immersive visualization in a VR lab that complements the lecture. Whereas these events dedicated exclusively to immersive visualization will continue

• Shuqi Liao is with Purdue University. E-mail: liao201@purdue.edu.
• Sparsh Chaudhri is with Purdue University. E-mail: chaudhri@purdue.edu.
• Maanas K. Karwa is with Purdue University. E-mail: mkarwa@purdue.edu.
• Voicu Popescu is with Purdue University. E-mail: popescu@purdue.edu.

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

to have their role, it is also important to allow for a fine interleaving between immersive and non-immersive visualization, in support of impromptu, short-term, and repeated uses of immersive visualization.

Recent work explores combining immersive VR headsets with traditional computing devices to enhance data analytics efficiency [9, 18, 23, 26, 27]. These systems leverage VR’s spatial capabilities while retaining familiar desktop interfaces for deeper, aggregated analysis. Such dual-environment setups aim to improve workflow flexibility, efficiency, and insight generation across diverse domains. Numerous studies point out that precise editing in immersive environments remains challenging [43]. Although prior work points out the strengths and weaknesses of immersive and non-immersive visualization, calling for the use of both hybrid visualization systems do not presently provide for an effective way of switching from immersive to non-immersive visualization. Consider the scenario of a user seated at their desk, visualizing a 3D scene immersively, using a VR headset. Removing the headset is a disruptive event, and, continuing the visualization task on the 2D screen incurs the delay and cognitive effort of having to mentally connect the immersive visualization seen in the headset to the non-immersive visualization seen on screen.

In this paper we describe *SeamlessVR*, a method for switching effectively from immersive to non-immersive visualization. The user, who started a visualization task immersively, in the headset, can choose to switch at any time to a non-immersive visualization, on screen. The user indicates the desired switch through the immersive interface, for example, by pressing a button on the handheld controller, which initiates a continuous morph of the 3D visualization to a 2D visualization that matches what the user will see on screen after removing the headset (Fig. 1). The morph moves the scene vertices gradually from their 3D position in the immersive visualization to their location on the 2D screen in the non-immersive visualization. Once the morph completes, the user removes the headset, with the last visualization seen in the headset closely matching the first visualization seen on screen. This visualization continuity promises to reduce the cognitive effort of connecting the immersive to the non-immersive visualization, helping the user continue on screen the visualization task started in the headset. We also refer the reader to the video accompanying our paper.

We have compared *SeamlessVR* to the conventional approach of directly removing the headset in an IRB-approved user study. N = 30 participants were asked to start visualization tasks immersively, in the headset, and then to remove the headset and continue the tasks non-immersively, on screen. The tasks included counting objects and keeping track of stationary and dynamic targets in simple/complex, and abstract/realistic 3D scenes. Task performance was quantified objectively, in terms of accuracy and time, while user experience was evaluated subjectively, in terms of user preference, usability, and simulator sickness. The results show that *SeamlessVR* has a significant advantage over the conventional approach for target tracking in complex abstract scenes. In such scenes, participants found it difficult to map the 3D visualization to the 2D visualization, from memory, without the aid of the morph. *SeamlessVR* had a smaller, but still statistically significant, advantage for target tracking in complex realistic scenes, where salient landmarks were more readily available to aid with the mapping from 3D to 2D. There were no statistically significant differences between the two conditions for counting, where the morph brought both the benefit of visualization continuity, as well as the added complexity of scene motion. *SeamlessVR* had significant advantages over the control condition in terms of participants’ perception of the switch from immersive to non-immersive visualization, as well as in terms of usability. Finally, *SeamlessVR* did not pose cybersickness concerns.

In summary, our paper contributes (1) *SeamlessVR*, a method for switching from immersive visualization, in a VR headset, to non-immersive visualization, on screen, with visual continuity, and (2) a controlled user study that investigates and confirms the scenarios where *SeamlessVR* has an advantage over the conventional approach of removing the headset directly.

2 PRIOR WORK

2.1 Transition between Virtuality and Reality

We review prior work on cross-reality systems, focusing on visual continuity. The Reality-Virtuality Continuum (RVC) proposed by Milgram et al. [35] provides a foundational framework for understanding transitions across varying degrees of virtuality, from fully immersive virtual environments to augmented and real-world settings. The importance of designing interfaces that allow the user to move fluidly between the virtual and real worlds was noted in the seminal *MagicBook* [7] work that coined the phrase *Transitional Interface*. However, building transitional interfaces for seamless transitions along the RVC, particularly in maintaining visual continuity, remains a significant research challenge and an area of ongoing exploration.

Managing transitions between virtuality and reality is a prominent research area with a substantial body of prior work. A comprehensive review [4] categorizes cross-reality systems into three types: *transitional*, involving a gradual shift along the RVC; *substitutional*, where physical objects can be accessed for interaction from within the virtual environment; and *multi-user*, allowing multiple users to interact together, possibly with each user experiencing their own reality. Our *SeamlessVR* method focuses on the continuous transition from immersive to non-immersive, placing it within the *transitional* category.

Several works aim to build virtual replicas of the physical world to facilitate the transition between the real and virtual worlds. For instance, Valkov et al. [48] highlight four essential elements for smooth immersion: the real world, a virtual replica, a seamless transition, and the fully virtual environment. This approach gradually transitions between physical and virtual objects by using timers to manage object visibility, ensuring users do not notice changes within their current field of view. Similarly, Pointecker et al. [39] introduce a virtual replica as an intermediate layer between reality and virtuality, comparing transition techniques such as fade, dissolve, translate, and combinations thereof. Users found the replica helpful for reducing spatial disorientation and for aiding cognitive processing, although constructing high-fidelity replicas requires extensive pre-processing and fine-tuning. Our *SeamlessVR* approach builds a digital twin of the 2D screen, which is both simple, and effective, as *SeamlessVR* is not preoccupied with connecting the user to the real world, but rather to keep the user in the same virtual scene or dataset across visualization modalities.

Researchers have also investigated aiding the transition between real and virtual through visualization techniques. These methods often incorporate in the virtual world some familiar elements from the real world, such as phones, portals, or doors, in order to make the transition between environments more intuitive [16, 19, 24, 38, 46]. Another way to provide a continuous transition experience across realities is to use transition animations, such as fading [16, 22, 38, 39, 46], dissolving [16, 38, 39], or translation [16, 39], similar to the transitions between slides most presentation editing software provide. Evaluations show that users tend to prefer simple and efficient transitions over elaborate animations or time-consuming interactions [17, 22]. The transition *SeamlessVR* has to achieve is both constrained and aided by well-defined starting and ending visualizations. Unlike the aforementioned approaches, which focus on ensuring smooth transitions between different scenes along the RVC, *SeamlessVR* instead aims to preserve the visual continuity of the same scene across the RVC. In other words, the visualizations are not disparate, but rather strongly connected, as they show the same 3D scene. The goal is *not* to provide the user with a visual palate-cleanser as they move from one visualization course to the next, but rather to help them connect as well as possible the two visualizations. Furthermore, the same 3D scene underlying both visualizations provides a dense set of reliably correspondences between the two visualizations, which can and have to be used when designing the transition.

Some research focuses on minimizing the sense of fragmentation when transitioning from immersive virtual environments back to the real world [22, 29, 37]. For example, one study [29] investigates the user experience as they exit from virtual reality, offering design guidelines to improve this process. The guidelines consider the abruptness of the transition, whether the headset should switch to passthrough mode

or not, how to align the scales of the two worlds, how to reconfigure the virtual world in preparation for re-emergence to reality, and what to do about the social setting change. There is no single answer that fits all users, scenarios, and applications. Similarly, another work [22] introduces “outro-transitions”, techniques designed to smoothly guide users from virtual environments back to reality. This study highlights the importance of simple, effective transitions to avoid abrupt breaks in immersion, particularly in scenarios where users frequently switch between virtual and physical worlds. *SeamlessVR* focuses on switching from one visualization modality to another, bypassing the real world. Nonetheless, some of the design constraints identified by this body prior work, such as a smooth transition, aligning the scales of the worlds of the two visualizations, and reconfiguring one world in preparation for switching to the next world are relevant and followed by *SeamlessVR*. As discussed in the Conclusions Sec. 5, it is not yet useful for *SeamlessVR* to switch to passthrough mode before removing the headset as the passthrough video does not yet show the 2D screen with sufficient quality. Switching from a virtual 2D screen to the video passthrough screen and to the directly seen screen introduces an unwarranted drop in visual quality during the video passthrough stage.

2.2 Data Visualization Transformation in Mixed Reality

Lee et al. [31] introduced a design space for transforming visualizations between immersive (3D) and non-immersive (2D) environments, providing a useful taxonomy for categorizing *SeamlessVR*. Accordingly, *SeamlessVR* aligns with a perspective projection transformation method that transitions between a 3D visualization and a corresponding 2D view. Following these guidelines, *SeamlessVR* supports scenarios where an immersive 3D visualization can be temporarily suspended and then resumed on a non-immersive 2D display, or vice versa. The transformations to which *SeamlessVR* resorts are *continuous*, they occur over a *fixed* duration, and they ultimately result in a *permanent* change of the visualization modality.

The extension of 2D visualizations into 3D is an active research field [30, 33, 41, 42]. For example, a recent study focused specifically on the visualization of graphs [45]. Users were given the ability to switch between immersive visualization in an XR headset and non-immersive visualization on a desktop screen. The study notes the benefits of switching between the two modalities. The user is expected to wear the headset at all times, even during the non-immersive, on screen visualization, and the study notes the fact that current video seethrough technology (i.e., Varjo XR-3 [3]) does not show the desktop screen sufficiently well so it had to be virtualized.

Schwajda et al. [44] focused on transforming graph data visualizations specifically from an interactive 2D display into an augmented 3D environment. Their method takes advantage of specific graph data characteristics—such as nodes and edges—and involves carefully defining initial and final visualization states, as well as applying interpolation to achieve a smooth transformation. The approach aims to convert optimally the 2D graph into a 3D representation on a sphere’s surface, thereby altering the dataset’s visual modality. By comparison, *SeamlessVR* aims to maintain visual continuity without modifying the user’s view of the underlying dataset, and it is also more general, being applicable to a broad range of 3D datasets.

2.3 Visual Continuity in Cross-Reality Systems

Mergereality [52] proposes high-level design principles for achieving seamless transitions between virtual and real environments, demonstrating the feasibility of transferring and synchronizing content across multiple devices, such as tablets, computer screens, and mixed reality (MR) displays, yet these principles remain largely unvalidated through concrete evaluations. Building on these insights, we focus on assisting users in initiating and switching between immersive and non-immersive visualizations for a variety of 3D datasets. This includes handling both simple and complex, as well as static and dynamic scenarios, and we validate our approach through four distinct user tasks.

Research in cross-reality interaction has investigated bi-directional integrations between XR headsets and conventional displays (e.g., desktop screens [12] or smartphones [54]). These systems aim to syn-

chronize both interaction and visualization between immersive and non-immersive modalities [49]. For example, a recent interface [53] employs projective texture mapping to enable seamless transitions between 2D and 3D visualizations of complex biological structures. In this setup, a drawing tablet and an optical see-through headset (HoloLens2 [1]) are integrated with hand-gesture input, ensuring both visual and interaction continuity. Notably, while the see-through headset preserves the user’s direct view of the real world at the highest possible visual fidelity, its limited field of view constrains the volume of the 3D visualization when compared to video passthrough XR headsets.

In summary, our review of prior work confirms that transitional interfaces are useful in a variety of domains, that they can be implemented with current XR headset technology, and that transitioning between visualization modalities is best done gradually, with the preservation of visual continuity. *SeamlessVR* builds upon this prior work, focusing on the scenario where the user completely removes the headset as opposed to merely switching to passthrough mode. *SeamlessVR* also excels at maintaining visual continuity as the visualization transitions from 3D to 2D. We morph 3D scene vertices to their 2D visualization destination, as opposed to simply hiding a 3D model beyond the screen by pushing it across the invisible clipping plane of the screen [52], or to projecting the 3D visualization onto the screen plane [53]. In other words, we change the 3D visualization into 2D visualization as opposed to showing both and progressively removing one. In addition to geometric visual continuity, our approach also has the ability to reduce depth perception gradually from 3D to 2D visualization. We validate *SeamlessVR* on visualization tasks that are general and low level, i.e., counting and tracking static and dynamic targets in simple and complex, and abstract and realistic virtual environments, contributing to the body of empirical evidence supporting transitional interfaces.

3 CONTINUOUS 3D TO 2D VISUALIZATION MORPHING

It is our goal to allow a user to switch from immersive visualization in a VR headset to non-immersive visualization on a 2D screen with good visual continuity, such that the user can continue on screen a visualization task started in the headset. In this section we first detail the design concerns that have to be met for switching between immersive and non-immersive visualization (Sec.3.1), and then we describe how our *SeamlessVR* approach meets these concerns (Sec.3.2).

3.1 Design Concerns

The ideal transition from immersive to non-immersive visualization is one where what the user sees does not change at all. This ideal transition is not possible for general scenes, because a 3D scene seen in the headset with depth perception cannot be seen by the user the same way on the 2D screen. As such, the transition has to gradually change from one visualization to the other, allowing the user to maintain scene awareness without too high of a cognitive effort. We have developed *SeamlessVR* according to the following design concerns:

Visual continuity. Removing the headset should not cause a significant visualization discontinuity. For this, what the user sees in the headset before removing it should match closely what the user sees on screen once they remove the headset. Any visual discontinuity translates to a higher cognitive load for the user to map the headset visualization to the screen visualization from memory. Visual continuity allows the user to track the entities visualized as the visualization changes, allowing the user to leverage and continue the visual analysis they started before the transition. Visual continuity is a high-level concern that implies several low level design concerns.

View parameter continuity. The user view used last to visualize the 3D scene in the headset has to be the view used first to visualize the 3D scene on screen. Different viewpoints, different view directions, or different scale factors would yield incongruent visualizations in the headset and on screen, showing different parts of the scene.

Appearance continuity. The two visualizations should be as close as possible not only in what they show, but also how they show it. Differences in brightness, dynamic range, and color space should be reduced gradually as the user removes the headset.




Fig. 2: Transition initialization and morphing. The 3D bounding box of the scene is translated and scaled from BB_{3D} to BB_{2D} , and, at the same time, BB_{2D} is morphed to the 2D screen S . A scene vertex P corresponds to a BB_{2D} point Q , which is used to define the final screen position T along the 2D visualization ray OQ . P is then gradually displaced to T . O is defined as the user's head position U when the transition is initiated.

Continuous transition from stereoscopic to monoscopic visualization. Once the user removes the headset, they will see the visualization of the 3D scene monoscopically, on a 2D screen. For this, the 3D visualization depth cues have to be suppressed gradually such that the user is not asked to perform an abrupt change of the depth where they focus. This is achieved by flattening the 3D visualization to a 2D visualization.

Virtual to physical 2D screen alignment. The plane on which the 2D visualization is flattened has to match the plane of the physical screen. Furthermore, the frames of the virtual and physical screens have to align. This is achieved by positioning the virtual screen in alignment with the physical screen.

Temporal continuity. In order to support not only static but also dynamic 3D scenes, that change based on a recorded or real-time simulation or based on user interactions, the 2D visualization on screen has to be in sync with the 3D visualization. This is achieved by communicating from the headset to the computer all changes to scene.

3.2 SeamlessVR

Pre-processing. Before the visualization session, the virtual world coordinates of the physical screen have to be calibrated to align the virtual and physical screens. In our implementation we ask the user to click the corners of the physical screen with a handheld controller.

Before the transition. The user visualizes a 3D scene in the headset. For every frame, the scene is animated and rendered conventionally as the transition has not yet begun.

Initiating the transition. When desired, the user initiates the transition from immersive to non-immersive visualization, for example, by pressing a button on the handheld controller.

Initializing the transition. Before the transition can begin, transition control has to compute (Fig. 2): (a) the 3D axis-aligned bounding box BB_{3D} of the 3D scene as visualized in 3D, in the headset; (b) the 3D axis-aligned bounding box BB_{2D} of the 3D scene as it will be visualized in 2D, on screen; (c) the view V_{2D} used to render the 3D scene on the virtual and physical screens. BB_{3D} is computed straightforwardly to encompass the 3D scene as it is currently displayed in the headset. Since BB_{3D} could be much larger than the 2D screen, and since it might not be aligned with the screen, BB_{3D} has to be first transformed to an axis aligned bounding box BB_{2D} in between the user and the 2D screen. In Fig. 2, BB_{3D} is to the left of and larger than the screen so it has to be scaled down and translated to BB_{2D} . The viewpoint O of the 2D visualization is set to the user viewpoint U when the transition is initiated, and it is *not* updated to account for subsequent user head motion. Updating O continually by tethering it to the user viewpoint would create an unstable visualization that changes based on user head motions, and is likely to lead to disorientation and even cybersickness. Indeed, prior work has widely observed that moving visual stimuli incongruent with the user's head motion leads a person to experiencevection, i.e., an imagined self-motion that is associated with cybersickness [10]. The frustum of the 2D visualization view V_{2D} is defined using viewpoint O and the screen frame S . The transition is initialized quickly, over the span of one frame.




Fig. 3: Side view (first row) and top view (second row) illustration of the flattening of the 3D visualization from Fig. 1.

During the transition. The 3D visualization morphs gradually to a 2D visualization on screen (Fig. 3). The scene is rendered with a morph. A scene vertex P (Fig. 2) is displaced gradually to its screen position T , which is computed in two steps: (1) compute the BB_{2D} point Q that corresponds to P using the transformation from BB_{3D} to BB_{2D} , and (2) find T as the intersection of ray OQ with the screen plane S . The pseudo-code algorithm for the vertex shader that implements the morph is given in Alg. 1. For correct visibility sorting in the final, flattened visualization, the morph ends before a vertex reaches the plane of screen S , as controlled through the shader parameter ϵ . This way, the visualization is not flattened to a 2D rectangle, but rather to a thin box.

End of transition. When the transition ends, the headset communicates to the computer the 2D visualization view V_{2D} . For dynamic scenes, the headset also communicates the current state of the 3D scene and the current animation step, to synchronize the screen visualization to that in the headset (Fig. 4). This way the computer can begin to animate and render the scene to provide the 2D visualization on screen.

After the transition. Once the transition concludes, the user can remove the headset and continue the visualization task on screen. If the headset has a video passthrough mode, the headset could switch to showing the real world, including the screen, before the user removes

Algorithm 1 Vertex shader morphing algorithm

Input: With the notations from Fig. 2: 3D scene vertex P , transformation X from BB_{3D} to BB_{2D} , 3D rectangle screen S , 2D visualization viewpoint O , current morph step i , total morph steps n , early stopping threshold ϵ .

Output: morphed vertex P_i

- 1: $Q = XP$
- 2: $T = OQ \cap S$
- 3: $P_i = P + \frac{(T-P)}{\|T-P\|} \cdot \frac{i}{(n+\epsilon)}$



Fig. 4: Illustration of the synchronization of a dynamic 3D scene between the headset and the computer. The image is a second-person view frame acquired with an XR headset in passthrough mode to show the user, the screen, and what the user sees in the headset.

the headset. The benefit is that this switch to passthrough mode provides one more step in the transition from the headset to the screen visualization. One disadvantage is that the passthrough mode might introduce distortions, as the physical screen is close to the user and resolving the distance between the eye cameras and the eyes distorts the video frame. Another disadvantage is that the passthrough mode shows the user the real world with lower resolution and lower dynamic range than when the user sees it directly.

The morphing process can be reversed to support switching from non-immersive visualization, on screen, to immersive visualization, in the headset. Indeed, the morph is reversible to the BB_{2D} box in between the user and the screen, using Fig. 2. The first step is to compute a transformation of the 3D scene bounding box to BB_{2D} , which is used to recover the Q position of a scene vertex P in BB_{2D} . Then T is pushed away gradually from the screen on ray TP . This way, the user can interleave immersive and non-immersive visualization with any granularity.

In summary, the *SeamlessVR* approach satisfies the design concerns from Sec. 3.1. *View parameter continuity* is enforced by rendering on screen with the same camera as for the last headset frame. *Appearance continuity* is partially satisfied by using the same shading in the two visualizations. The three dimensionality of the visualization dissolves progressively, with the user perceiving less and less left-right eye disparity, as the 3D visualization folds into the 2D visualization, which provides a *continuous transition from stereoscopic to monoscopic visualization*. *Virtual to physical 2D screen alignment* is provided by the XR headset ability of acquiring the 3D coordinates of the corners of the physical screen. We conclude that *SeamlessVR* satisfies the overarching design concern of **visual continuity**.

4 USER STUDY

We conducted a controlled user study to compare our *SeamlessVR* approach for switching from immersive to non-immersive visualization with **visual continuity** to the conventional approach of switching abruptly by directly removing the headset. Our study was conducted with the approval of our Institutional Review Board. We describe the methods (Sec. 4.1), we present the results (Sec. 4.2), and we conclude with a summative discussion of the study results (Sec. 4.3).

4.1 Methods

Participants. $N = 30$ participants were recruited from the undergraduate and graduate student population of our department. The participants' average age was 22 years ($SD = 3.6$), ranging from 18 to 28, with 4 participants self-identified as female and 26 as male. Unfortunately, this severe gender imbalance reflects the overall representation in our department. No participant indicated never having used a VR application before, 10 had used VR once, 15 occasionally, and 5 frequently. All participants had normal or corrected-to-normal vision. The experimenter, who wears vision glasses and never removes them when putting on a headset, helped participants wearing vision glasses practice putting on the headset on top of their glasses, and removing the headset without removing their glasses. The headset is equipped with a head strap which facilitates this task.

Study design. The study compares *SeamlessVR* (experimental condition) to the conventional approach of switching directly from the 3D visualization inside the headset to the 2D screen-based visualization (control condition). We use a within-subject design, with each participant performing tasks in each condition, in counterbalanced fashion to avoid that learning effects skew our results.

Implementation and experimental procedure. The tasks were developed and implemented on Unity (version 2022.3.17f1) along with the Meta XR All-in-One SDK (v68.0.0) and the Oculus XR Plugin (v4.1.2). The prototype is available as open source¹. All development and testing were conducted on a high-performance machine equipped with an Intel Core i9-13900 processor, 32 GB of RAM, and an NVIDIA GeForce RTX 4080 graphics card. A Meta Quest 3 headset [2] with two Meta Quest Touch Plus controllers was used run to all the experiments.




Fig. 5: User study tasks $T1$ (left), $T2$ (middle), and $T3$ (right). The top row shows headset frames where the participant begins the task, and the bottom row shows the participant completing the tasks on screen.

The headset was plugged into the computer. Participants received instructions before each task, in the headset, in textual form. The participant was seated about 1 m from the screen, which was 59.7 cm \times 33.6 cm in size and had a resolution of 3,840 px \times 2,160 px. A participant was involved in a single 30 min session and was compensated with a 30 USD gift card.

4.1.1 Study Hypotheses

We expect that, due to the visually continuous transition from immersive to non-immersive visualization, *SeamlessVR* outperforms the control condition in a variety of scenarios. Specifically, we expect *SeamlessVR* to allow participants to start the task immersively and then to continue non-immersively, whereas in the control condition participants have to start over, or to continue from memory. We expect that this will translate to a task completion time or correctness advantage for *SeamlessVR*. Furthermore, we expect that participants will find the abrupt switch in the middle of the task frustrating, giving *SeamlessVR* an advantage in terms of personal user preference and usability. We expect that these *SeamlessVR* advantages hold for a variety of scenes, static or dynamic, abstract or realistic. Finally, we do not expect *SeamlessVR* to induce cybersickness since it is akin to a scene animation, and not to a view modification incongruent with headset motions. We have formulated the following hypotheses, which guided task and data collection design.

- **H1** *SeamlessVR* outperforms control in terms of objective metrics such as task completion time and correctness.
- **H2** *SeamlessVR* outperforms control in terms of subjective metrics such as user subjective preference and usability.
- **H3** *SeamlessVR* does not pose cybersickness concerns.

4.1.2 Tasks

Each participant performed four tasks (Figs. 5 and 1), seated, repeating each task three times. Upon hearing a beep, participants took off their headsets mid-task and continued the task on screen. The timing of the beep was specific to each task and was tuned through a pilot study.

The *first task*, $T1$, involved counting 3D shapes. The task aims to investigate whether a participant can leverage having started counting with the headset on as they continue counting on the screen, or whether they need to start from the beginning after the switch to non-immersive visualization. The *second task*, $T2$, required participants to remember and identify a highlighted sphere out of many identical spheres. The task aims to investigate whether participants can map the immersive visualization of the 3D scene, where the highlight was shown, to the non-immersive visualization, such that they can transfer the highlight to identify the target. The *third task*, $T3$, involved playing a shell game, where the goal is to keep track visually of two shells that each hide a ball as a total of five shells are shuffled. The task aims to investigate whether participants can map the 3D visualization to the 2D visualization in the case of a dynamic scene. The *fourth task*, $T4$, was similar to the second task but involved identifying an element within a complex and realistic 3D model of a factory (Fig. 1). The task aims to investigate the quality of visualization transition in the context of a realistic 3D scene.

¹<https://github.com/VRSeamlessTransition/SeamlessVR>

T1: counting. In the first task, participants were asked to count floating objects. A total of 125 objects, varying in shape (pyramids, spheres, and cubes) and color (red, green, and blue), were displayed in 3D in front of the participant. The participant was asked to count objects of a specific shape and color, with correct answers in the range of 11 to 17. The objects were stationary. After a certain time interval, participants were instructed to switch to non-immersive visualization, either by removing the headset directly in the control condition or by engaging our method for transitioning and then removing the headset in the experimental condition. The duration in the control condition was 9 s inside the headset with the original view, after which a beep was heard, and participants were asked to remove their headsets. The duration of the experimental condition was 3 s inside the headset with the original view, followed by a 6 s morph into 2D, after which a beep was heard and the participants were asked to remove their headsets. Therefore, participants saw the 3D scene within the headset for the same total amount of time of 9 s. Once in the non-immersive mode, participants continued counting the objects on screen in both conditions. After completion of the count, the participants were asked to enter the number of objects. There was no time limit for counting on the screen.

T2: static target tracking. In the second task, participants were asked to track a sphere across the transition between visualization modes. A participant was first instructed that they will see a red target sphere and that they will be asked to keep track of. Initially, multiple floating spheres of identical size were displayed in 3D in front of the participant for 3 s. All spheres were of the same color except for the target sphere which was red. After the 3 s, the highlight disappeared, with the target sphere assuming the color of the other spheres. In the control condition, the participant had an additional 3 s to observe the target, then was asked to remove the headset and click the target sphere on the screen using the computer mouse. In the experimental condition, the morph started automatically after the 3 s and lasted for 3 s, after which the participant was asked to remove the headset and to click the target on screen using the mouse.

T3: dynamic target tracking. In the third task, participants were presented with five shells arranged in a row. Initially, two balls were slowly and visibly placed under two of the shells. Once the balls were concealed, the cups were shuffled with the participants attempting to track the shells where the balls are hidden. In the control condition, the shells were shuffled for 6 s, after which a beep prompted the participant to remove the headset, while the shuffling continued. Then the shuffling continued for another 6 s on the screen after which the participant was asked to indicate the two shells concealing the balls. In the experimental condition, the shells were shuffled while morphing from 3D to 2D for 6 s, after which the beep prompted the participant to remove the headset and finish the task on screen like for the control condition. Therefore, participants see the shells for a total of 12 s under both conditions. The shuffle continued after participants removed the headset. Participants were instructed to remove the headset as quickly as possible while maintaining their gaze on the target shells.

T4: static target tracking in complex scene. In the fourth task, participants were required to track specific components across the immersive/non-immersive visualization boundary within a complex, realistic 3D model of a factory. The task is similar to *T2*, except that the spheres were replaced with factory model elements, such as pipes, gears, and machinery. In the experimental condition, participants saw the 3D visualization for 3 s, and the morph took 5 s, whereas in the control condition the 3D visualization lasted for 8 s.

4.1.3 Data Collection and Analysis

In terms of objective metrics, we have recorded the time each participant needed to perform each trial of each task in each condition, as well as the correctness of the task execution.

For *T1* and *T3*, correctness is obvious. For *T2*, the task is considered to be executed correctly if the user clicks on any pixel that belongs to the target sphere. In addition to correctness, the *T2* error is quantified as the Euclidian distance between the pixel on which the user clicks and the projection of the center of the target sphere. Similarly, *T4* is executed correctly when the user clicks on the target element of the

factory scene. No error is computed for *T4* as the target elements have varying size, making the error inconsistent.

In terms of subjective metrics, we have collected responses to our custom user preference questionnaire. There were six questions on it, each with responses on a five-point Likert scale [32]. The response options were: “strongly disagree”, “disagree”, “neither agree nor disagree”, “agree” and “strongly agree”. The questions were:

- Q1:** Removing the headset made me lose track of the task.
- Q2:** For counting, I had to start over after removing the headset.
- Q3:** The transition from headset to screen was smooth.
- Q4:** The transition from headset to screen was seamless.
- Q5:** The transition from headset to screen was too fast.
- Q6:** The transition from headset to screen was clear.

A participant answered all questions after the experimental condition, and only the first four (relevant) questions after the control condition. Finally, we have also investigated our *SeamlessVR* approach in terms of usability, using the standard SUS questionnaire [8], and in terms of cybersickness, using the standard SSQ questionnaire [28]. We analyzed time (a continuous dependent variable) and correctness (a dichotomous dependent variable) across the four tasks. We also analyzed subjective dependent variables derived from the SUS and the custom user preference questionnaire, both using Likert-like scales.

Descriptive statistics are presented through box plots for time and bar plots for correctness. The boxplot shows the interquartile range (IQR) as a thick bar, the entire range as whiskers, the mean as a small triangle, the median as a white horizontal line, and outliers as small circles. Outliers are defined as data points that fall below $Q1 - 1.5IQR$ or above $Q3 + 1.5IQR$, where $Q1$ and $Q3$ represent the first and third quartiles (25% and 75%), respectively. We also compare conditions using inferential statistics. Time and correctness were not normally distributed, so we used the non-parametric Wilcoxon’s Signed-Rank test [51] for time and McNemar’s test [34] for correctness. All statistical tests were conducted at a significance level of $\alpha = 0.05$. A paired t-test [47] was applied to the SUS scores, leveraging their normal distribution. In contrast, the Wilcoxon Signed-Rank test with a normal approximation was used for the custom user experience questionnaire, which were not normally distributed and included cases with zero differences. For significant results, we quantified effect size using Cohen’s r for Wilcoxon’s tests [50], Cohen’s h for McNemar tests [36], and Cohen’s d for paired t-tests [11], and we used it to estimate the statistical power provided by our $N = 30$ participants.

4.2 Results

We first discuss the results on the objective time (Sec. 4.2.1) and correctness (Sec. 4.2.2) metrics for each task, and then we discuss the results on the subjective metrics over all tasks (Sec. 4.2.3).

4.2.1 Time

The task completion time is shown through box-plots in Fig. 6 (left) and numerically in Tab. 1 (row 1).

T1: counting. The left panel of Fig. 6 shows that the *T1* task completion time is shorter for *SeamlessVR* (20.5 s) than for the control condition (21.5 s), but the difference is not significant (Tab. 1). Since in both conditions participants saw the 3D scene in the headset for the same amount of time, we attribute the time advantage for *SeamlessVR* to the counting that participants were able to perform in the headset and then were able to carry over to the screen to complete the count, and to the fact that counting might become easier as the 3D scene flattened in preparation for the switch to the non-immersive visualization. However, the morph animates the scene, which can complicate counting.

T2: static target tracking. For *T2* the average *SeamlessVR* time of 15.5 s is significantly shorter than the 17.8 s time for control. The difference is significant even for a conservative Bonferroni adjusted significance level $\alpha = 0.05/4 = 0.0125$ that accounts for the four time comparisons (one per task). We explain this significant advantage by the fact that keeping track of the target during the morph is a relatively simple visualization task, with the participant having to only attend to a single object. On the other hand, the counting in *T1* required keeping track of the objects that were already counted, which is more

Table 1: Time (Wilcoxon’s signed rank test) and correctness (McNemar’s test) comparison between the two conditions for the four tasks. Negative Z values indicate shorter *SeamlessVR* times compared to control. Significant differences are indicated with an asterisk (*).

Metric	T1: counting	T2: static target tracking	T3: dynamic target tracking	T4: static tracking in factory
Time	Z = -1.115, p = 0.265	Z = -3.374, p < 0.001*	Z = -0.084, p = 0.933	Z = -3.897, p < 0.001*
Correctness	p = 0.324	p < 0.001*	p = 0.815	p = 0.041*




Fig. 6: Task completion times (left) and correctness (right) across tasks and conditions.

challenging. Furthermore, for *T2*, once the participant switches to the screen, the participant has to spend time searching for the target from memory, which includes building a mental mapping between the 3D and 2D visualizations. The effect sizes are “large” (row 1 of Tab. 2), so our $N = 30$ participants provide sufficient statistical power.

T3: dynamic target tracking. For this task the shells move for the same amount of time, so any differences in time between the two conditions stem exclusively from the thinking time participants take to provide their answer after the shells have stopped moving. Since the shell game cannot be solved by starting over, there are no great differences in time whether the participant knows or does not know the answer, so time is not a particularly discerning metric for *T3*. Indeed, the times for the two conditions are similar.

T4: static target tracking in a complex scene. The time advantage of *SeamlessVR* over control found in *T2* is confirmed for the realistic factory model used in *T4*. The average time for *SeamlessVR*, i.e., 16.7 s, is significantly shorter than the 18.9 s for control, even for a conservative Bonferroni adjusted significance level of $\alpha = 0.05/4 = 0.0125$. It is easy to keep track of the target during the morph, the effect is large, so the analysis is adequately powered (row 2 of Tab. 2).

4.2.2 Correctness

Task completion correctness is shown in Fig. 6 (right) and in Tab. 1 (row 2).

T1: counting. The counting accuracy is slightly higher for the control condition compared to *SeamlessVR*. The difference is not significant. Furthermore, both numbers are below 50%, which indicates that the task of counting between 11 and 17 objects out of 125 might have been too challenging. As such, the motion of the objects during the morph made counting in the *SeamlessVR* condition even more challenging. Participants made spontaneous observations about the counting task: some found counting in 3D very hard and essentially were waiting to count on screen, some found the counting task easy in 3D and would finish counting before the switch to the screen, and some mentioned that it was hard for them to keep track of the objects that were already counted as they were moving. Fig. 7, left, shows that counting errors were either 0 or 1, with a median error of 1, for both conditions.

T2: static target tracking. *SeamlessVR* has a significant advantage over control in terms of target tracking correctness (77% vs 19%), even for a conservative Bonferroni adjusted significance level $\alpha = 0.05/4 = 0.0125$. The visualization discontinuity between immersive and non-immersive makes it hard to track the target in the control condition, whereas *SeamlessVR* shows the target location before the participant removes the headset. In the control condition, participants found it hard to memorize the scene from the 3D visualization and then to map the 3D scene, from memory, to the 2D visualization on screen. Fig. 7, right, shows that in the control condition participants simply lost track of the target, clicking hundreds of pixels away from the target’s correct screen

Table 2: Effect sizes and statistical power for *T2* and *T4*.

Metric	Task	Effect Size	Power
Time	T2	$r = 0.62$ (large)	95%
	T4	$r = 0.71$ (large)	99%
Correctness	T2	$h = 1.23$ (large)	99%
	T4	$h = 0.30$ (small)	53%

location. The effect size is “large”, and $N = 30$ participants provide ample statistical power (row 3 of Tab. 2).

T3: dynamic target tracking. In our pilot study we used three shells with one hidden ball, but the task proved to be too simple. As such, for the actual study, we used five shells and two balls. Even so, participants performed the task correctly in both conditions, with no significant differences. The shells are aligned, there is only a few of them, and their motion is continuous, therefore participants succeeded in keeping track of the two shells with the hidden balls even in the control condition. Although a dynamic scene, the shell game lacks the complexity needed to probe the difference between the two conditions.

T4: static target tracking in complex scene. The *SeamlessVR* condition has a significant advantage over control, but the p value is close to the significance level $\alpha = 0.05$, and it far exceeds Bonferroni adjusted significance level $\alpha = 0.05/4 = 0.0125$. However, the advantage is much smaller than in the case of *T2*, where the static target tracking was performed in an abstract scene with many similar objects, i.e., spheres. The effect size is small, and the test is under-powered (row 4 of Tab. 2). The complexity of the factory scene played to the advantage of the control condition, making it easier for the participant to memorize the relative position of the target within the scene, leveraging nearby landmarks. In other words, participants will memorize what the target is, and what the target neighborhood looks like, which helped them overcome the visualization discontinuity in the control condition more reliably. Some participants mentioned that in the control condition, when the target was one out of a group of identical objects, they resorted to counting the objects to remember the relative position of the target within the group. While this may increase correctness, it also increases time, as shown in Fig. 6, left.

4.2.3 Subjective Metrics

User preference questionnaire. The five-point Likert scale answers to the first four questions of the questionnaire are provided in Fig. 8. The statistical comparison between the two conditions for each of the four questions is given in Tab. 3. *SeamlessVR* has a significant advantage over control in terms of participants not losing track of




Fig. 7: Counting error magnitude for *T1* (left) and target tracking error magnitude for *T2* (right), for each of the two conditions.




Fig. 8: User preference questionnaire scores for **Q1** (task continuity after headset removal), **Q2** (no need to restart count upon headset removal), **Q3** (smooth headset-to-screen transition) and **Q4** (seamless transition). For each pair, the control condition is at the top (red) and the *SeamlessVR* condition is at the bottom (blue). For each bar, the five segments indicate the number of times each answer value was recorded over the 30 participants. **Q1** and **Q2** are negative questions, and their score was flipped, i.e., 6 - a instead of a, for more to always mean better.

Table 3: Comparison between the two conditions for each of the first four questions of the user preference questionnaire. *SeamlessVR* has a statistically significant advantage for each question (see asterisk *).

Q.	Control vs. <i>SeamlessVR</i>	Effect size r	Power
Q1	$Z = -4.371$, $p < 0.001^*$	0.80 (large)	99%
Q2	$Z = -4.206$, $p = 0.011^*$	0.77 (large)	99%
Q3	$Z = -3.887$, $p < 0.001^*$	0.71 (large)	99%
Q4	$Z = -4.381$, $p < 0.001^*$	0.80 (large)	99%

the task when removing the headset **Q1**, of not having to start over the count when removing the headset **Q2**, of participants judging the headset to screen transition as being smooth **Q3** and seamless **Q4**. The differences are significant even for a conservatively Bonferroni adjusted level of $\alpha = 0.05/4 = 0.0125$, which accounts for the four tests, one per question. All effect sizes (Cohen's r) are "large", and the tests are adequately powered (Tab. 3). Q5 and Q6 only pertain to the *SeamlessVR* condition. As shown in Fig. 10, participants rated the transition neither too fast (Q5: 3.7/5.0) nor unclear (Q6: 3.9/5.0).

System usability (SUS). The overall average SUS scores are 80.8 for *SeamlessVR*, which corresponds to the adjective label of "Excellent", to the 90th percentile, to the letter grade of "A", to the highest acceptability rating of "Acceptable", and to a highest net promoter score (NPS) rating of "Promoter" [5, 8]. The SUS score for the control condition is 69.1, which corresponds to an adjective label between "Good" and "OK", to a percentile lower than 59, to the letter grade of "C", to the acceptability rating of "Marginal", and to an NPS rating of "Passive". This difference in SUS scores is significant (t-test results of $t(29) = -2.86$, $p = 0.008$), with a Cohen's effect size $d = 0.523$, and 80% power.

Cybersickness (SSQ). The nausea (N), oculomotor (O), disorientation (D), and total (TS) SSQ scores for the control condition were $N = 5.4$, $O = 10.6$, $D = 13.5$, and $TS = 11.0$ (Fig. 9). For *SeamlessVR*, the SSQ scores were slightly lower at $N = 3.2$, $O = 5.6$, $D = 7.4$, and $TS = 6.0$. There are no significant differences between the two conditions. The




Fig. 9: SSQ scores for the two conditions.




Fig. 10: Scores for questions **Q5** (transition speed being too fast) and **Q6** (clear transition) of the user preference questionnaire, for the *SeamlessVR* condition. The scores for the negative question **Q5** were flipped.

scores are low for both conditions, with most VR experiments reporting scores above 20 [13, 40]. Indeed, our study exposed participants to VR for short amounts of time, as participants had to remove the headset frequently. The morph between the 3D and 2D visualizations is brief in nature and the participant is anchored by the horizon and ground plane, which do not change during the morph. It is plausible that the control condition ($D = 13.5$) is more disorienting than the *SeamlessVR* condition ($D = 7.4$) as the visual discontinuity between the headset and the screen is more salient. Consequently, *SeamlessVR* lowers the potential for disorientation compared to repeatedly removing the headset abruptly.

4.3 Discussion

Results Overview. Tab. 4 gives an overview of the results for each task and for each metric, noting the attributes of the scene used in the task. A statistical significant advantage of *SeamlessVR* (EC) over the conventional approach of removing the headset abruptly (CC) is indicated through green shading; yellow shading means that the difference is not significant; we have grayed out the time cell for *T3* because all trials end at the same time, by design, i.e., when the shells stop moving. User preference here only encompasses the first four questions which have been administered for both conditions. In terms of objective metrics, *T2* shows a *SeamlessVR* advantage, whereas *T1* does not. In terms of subjective metrics, *SeamlessVR* has an advantage for all tasks. Finally, cybersickness is not a concern for *SeamlessVR*, as it isn't a concern for the control condition.

Research Hypotheses Support. Based on the results summarized in Tab. 4, the study lends support to our research hypotheses (Sec. 4.1.1) as follows. **H1** is partially supported. *SeamlessVR* has an advantage over control in terms of task completion time and correctness when the scene has no landmarks *or* when it is complex, but not when the scene has landmarks *and* it is simple. **H2** is supported. *SeamlessVR* has an advantage over control in terms of subjective user preference and usability for all tasks and all types of scenes investigated. **H3** is supported. *SeamlessVR* does not pose cybersickness concerns.

General Findings. The study results indicate that *SeamlessVR* meets our primary design concern of providing **visual continuity** when switching from immersive to non-immersive visualization. This implies that our implementation achieves a sufficient degree of *view parameter continuity*, of *appearance continuity*, of *stereoscopic to monoscopic transition continuity*, of *virtual to physical screen alignment*, and of *temporal continuity*. The tasks of our study bracket the range of tasks where *SeamlessVR* is most useful. The connection between the headset and screen visualizations that *SeamlessVR* fosters is most markedly

		Scene			Metrics				
		Static / Dynamic	Abstract / Realistic	Landmark	Time	Correctness	User preference	Usability	Cybersickness
Tasks	T1	S	A	Yes	NA	NA	NA	NA	NA
	T2	S	A	No					
	T3	D	A	Yes					
	T4	S	R	Yes					

EC > CC
EC ≈ CC
NA

Table 4: Overview of study findings: EC—experimental condition (*SeamlessVR*), CC—control condition.

useful when the alternative of building the connection from memory is challenging. This is the case when the scene has no landmarks, like for the blue spheres of *T2*. This is also the case when the scene *does* have landmarks but it is complex, e.g., factory in *T4*, where participants solved the task correctly, but took longer to do so.

SeamlessVR does not have an advantage for simple scenes or for scenes with a low level of depth complexity, such as the shell game of *T3*, where the user can easily map from short term memory the game in the headset to the game on screen. It is also the case that *SeamlessVR* has less of an advantage for very challenging tasks where the added visualization complexity brought by the morph distracts the user from continuing the task during the morph, before removing the headset, as was the case for the counting task of *T1*. Interestingly, for *SeamlessVR*, participants responded that they did not start over the count after removing the headset, and that they did so for control (significant advantage for **Q2**), even though *SeamlessVR* did not have an advantage for task completion time and accuracy. For such tasks, the morph speed should be reduced to be less intrusive. We tuned the morphing speed in a pilot study, but finding a single speed that suits all users is challenging.

SeamlessVR relies on a 3D morph that converts the 3D dataset to a thin (almost 2D) version that mimics the image on screen the user sees after removing the headset. This morph is most useful **(1)** when the 3D dataset much larger than its 2D image on screen [14, 25], case in which the morph also has the role of scaling it down, **(2)** when the user's focus is on a region not aligned with the screen [9, 23], case in which the morph also has the role of rotating the 3D dataset, **(3)** when the scene is complex, dynamic, and without landmarks [6, 25], case in which connecting the 3D visualization to the 2D visualization from memory is challenging, and **(4)** when continuity across the two visualization modalities is important [20, 21], such as when the task started immersively is interrupted and has to be continued in 2D, as is the case, for example, when the user is interrupted by a real world interlocutor, or when the user wants to alternate note taking with interactive visual exploration.

Use Cases. In our study the transition was initiated automatically for each trial to make sure a participant sees each task in each condition for approximately the same amount of time, to avoid confounding factors. For examples, allowing a participant to remove the headset right away, or to not remove it at all, defeats the purpose of the study. However, in practice, the transition is initiated by the user as desired, leveraging any one of a variety of input modalities such as keyboard, mouse, handheld controller, hand gesture, or voice. Furthermore, in practice, the user should be allowed to select the duration of the transition based on their individual preference. In our study, participants were asked to assume a view direction perpendicular to the screen and at its center. In practice, the user should be allowed to look at the screen from a skewed direction, as needed, for example, in a collaborative scenario with two users sitting in front of the same monitor. This is implemented by using the headset pose when the transition is initiated.

SeamlessVR is ready to be integrated with any hybrid visualization system that allows the user to toggle between immersive and non-immersive visualization, such as VRception [20] and ReLive [23]. One particularly promising use case of *SeamlessVR* is to bring to (3D) life images in a textbook or instruction manual, and, once explored in 3D, to fold the visualization again back to 2D, to continue to read or write with a conventional computer interface. Another use case is to provide the same functionality but with a physical document, e.g., a paper brochure, leveraging an XR implementation of *SeamlessVR*.

5 CONCLUSIONS, LIMITATIONS, AND FUTURE WORK

Conclusions. We have presented *SeamlessVR*, an approach for explicitly connecting an immersive visualization, shown in a VR headset, to a non-immersive visualization of the same 3D scene, shown on a 2D screen. *SeamlessVR* brings continuity between the two visualizations, allowing the user to start a visualization task in the headset and then to remove the headset and continue the visualization on screen. A user study compared *SeamlessVR* to the conventional approach of removing the headset directly on four visualization tasks, using objective and

subjective metrics. *SeamlessVR* has significant advantages over the conventional approach in terms of task completion time and task completion correctness when the 3D scene visualized is complex and scarce in visual landmarks. The advantages and qualities of *SeamlessVR* were also confirmed subjectively, by the study participants, through a user preference questionnaire. *SeamlessVR* received a significantly higher usability score than the conventional approach, and it did not pose cybersickness concerns. *SeamlessVR* is simple—with a small computational cost that can be easily handled by the VR headset, robust—with a morphing operation that is well defined even for complex scenes and that does not rely on optimizations, and general—without making any assumptions about the scene.

Limitations and Future Work. One possible improvement to our *SeamlessVR* implementation is to automate the real-world screen acquisition, such that the user doesn't have to calibrate the screen manually. For desktop screens that are unlikely to move between sessions this could also be achieved by saving the calibration and the spatial anchor data that are needed to reuse an earlier calibration. Appearance continuity could also be improved by taking into account display properties such as brightness, contrast, and color space.

This initial study does not attempt to quantify differences between *SeamlessVR* and control as a function of scene properties, and we have instead opted for testing as many scenes and tasks as possible. Future studies should test the same task on two scenes that are different only along one dimension, for example in terms of animation. We foresee that *T2* run on a dynamic version of the blue spheres scene (Fig. 5, middle) will prove to be close to impossible, resulting in even more sizeable advantages for *SeamlessVR*.

Our study confirms benefits of *SeamlessVR* in a limited context. Future work is needed to investigate any potential benefits of *SeamlessVR* in the context of other visualization tasks, over multiple sessions, and with a more diverse group of participants, e.g., with a more balanced gender representation and a wider range of ages. Future studies should also investigate whether tuning the morph speed based on the task, the user individual characteristics, and the user preferences can further increase the benefits of *SeamlessVR*.

Our work aims to make removing the headset a less disruptive event. Although the form factor of headsets has improved, putting on and removing a headset repeatedly remains burdensome. One aspect is that the headset is not designed for such use, and therefore supporting a fine interleaving between immersive and non-immersive visualization sessions requires buy-in from headset manufacturers. As VR headsets are now also XR headsets providing video passthrough modes, the user could conceivably just switch to passthrough mode at the end of the transition and *not* remove the headset. However, at least for now, the brightness, the dynamic range, resolution, and field of view of video see through XR headsets are not sufficient for comfortable viewing of detailed visualizations on real-world 2D screens. This is not a concern for optical see-through XR headsets, but such headsets suffer from other limitations such as limited field of view, brightness, and visualization transparency. It is also the case that, at least for now, headsets remain sufficiently uncomfortable for the user to be motivated to remove the headset and watch the visualization on the real world screen without the encumbrance of the headset. As headset form factors continue to improve, it is conceivable that the user might decide not to remove the headset when alternating between immersive and non-immersive visualization.

In our work we have only examined the transition from immersive to non-immersive visualization, but the morph is reversible, and future work should examine the interleaving of multiple immersive and non-immersive visualization sessions. *SeamlessVR* is ready to be integrated into applications to allow the user to cross the immersive/non-immersive divide at will.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grants No. 2219842, 2309564, and 2212200.

REFERENCES

[1] Hololens 2, mixed reality headset. <https://www.microsoft.com/en-us/hololens>. Accessed: 2024-09-14. 3

[2] Meta quest 3, mixed reality vr headset. <https://www.meta.com/quest/quest-3/>. Accessed: 2024-09-14. 5

[3] Varjo xr-3, mixed reality headset. <https://varjo.com/products/varjo-xr-3/>. Accessed: 2024-09-14. 3

[4] J. Auda, U. Gruenefeld, S. Faltaous, S. Mayer, and S. Schneegass. A scoping survey on cross-reality systems. *ACM Computing Surveys*, 56(4):1–38, 2023. 2

[5] A. Bangor, P. T. Kortum, and J. T. Miller. An empirical evaluation of the system usability scale. *Intl. Journal of Human–Computer Interaction*, 24(6):574–594, 2008. 8

[6] A. Bhardwaj, J. Chae, R. H. Noeske, and J. R. Kim. Tangibledata: Interactive data visualization with mid-air haptics. In *Proceedings of the 27th ACM Symposium on Virtual Reality Software and Technology*, pp. 1–11, 2021. 9

[7] M. Billinghurst, H. Kato, and I. Poupyrev. The magicbook-moving seamlessly between reality and virtuality. *IEEE Computer Graphics and applications*, 21(3):6–8, 2001. 2

[8] J. Brooke. Sus: A quick and dirty usability scale. *Usability Evaluation in Industry*, 1996. 6, 8

[9] W. Büschel, A. Lehmann, and R. Dachselt. Miria: A mixed reality toolkit for the in-situ visualization and analysis of spatio-temporal interaction data. In *Proceedings of the 2021 CHI conference on human factors in computing systems*, pp. 1–15, 2021. 2, 9

[10] E. Chang, H. T. Kim, and B. Yoo. Virtual reality sickness: a review of causes and measurements. *International Journal of Human–Computer Interaction*, 36(17):1658–1682, 2020. 4

[11] J. Cohen. *Statistical power analysis for the behavioral sciences*. routledge, 2013. 6

[12] R. Cools, M. Gottsacker, A. Simeone, G. Bruder, G. Welch, and S. Feiner. Towards a desktop-ar prototyping framework: Prototyping cross-reality between desktops and augmented reality. In *2022 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*, pp. 175–182. IEEE, 2022. 3

[13] M. Dennison and M. D’Zmura. Effects of unexpected visual motion on postural sway and motion sickness. *Applied ergonomics*, 71:9–16, 2018. 8

[14] C. Donalek, S. G. Djorgovski, A. Cioc, A. Wang, J. Zhang, E. Lawler, S. Yeh, A. Mahabal, M. Graham, A. Drake, et al. Immersive and collaborative data visualization using virtual reality platforms. In *2014 IEEE International Conference on Big Data (Big Data)*, pp. 609–614. IEEE, 2014. 9

[15] S. Feiner and A. Shamash. Hybrid user interfaces: Breeding virtually bigger interfaces for physically smaller computers. In *Proceedings of the 4th annual ACM symposium on User interface software and technology*, pp. 9–17, 1991. 1

[16] N. Feld, P. Bimberg, B. Weyers, and D. Zielasko. Keep it simple? evaluation of transitions in virtual reality. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, pp. 1–7, 2023. 2

[17] N. Feld, P. Bimberg, B. Weyers, and D. Zielasko. Simple and efficient? evaluation of transitions for task-driven cross-reality experiences. *IEEE Transactions on Visualization and Computer Graphics*, 2024. 2

[18] B. Fröhler, C. Anthes, F. Pointecker, J. Friedl, D. Schwajda, A. Riegler, S. Tripathi, C. Holzmann, M. Brunner, H. Jodlbauer, et al. A survey on cross-virtuality analytics. In *Computer Graphics Forum*, vol. 41, pp. 465–494. Wiley Online Library, 2022. 2

[19] C. George, A. N. Tien, and H. Hussmann. Seamless, bi-directional transitions along the reality-virtuality continuum: A conceptualization and prototype exploration. In *2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 412–424. IEEE, 2020. 2

[20] U. Gruenefeld, J. Auda, F. Mathis, S. Schneegass, M. Khamis, J. Gugenheimer, and S. Mayer. Vrception: Rapid prototyping of cross-reality systems in virtual reality. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pp. 1–15, 2022. 9

[21] J. Gugenheimer, E. Stemasov, J. Frommel, and E. Rukzio. Sharevr: Enabling co-located experiences for virtual reality between hmd and non-hmd users. In *Proceedings of the 2017 CHI conference on human factors in computing systems*, pp. 4021–4033, 2017. 9

[22] R. Horst, R. Naraghi-Taghi-Off, L. Rau, and R. Dörner. Back to reality: Transition techniques from short hmd-based virtual experiences to the physical world. *Multimedia Tools and Applications*, 83(15):46683–46706, 2024. 2, 3

[23] S. Hubenschmid, J. Wieland, D. I. Fink, A. Batch, J. Zagermann, N. Elmqvist, and H. Reiterer. Relive: Bridging in-situ and ex-situ visual analytics for analyzing mixed reality user studies. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pp. 1–20, 2022. 2, 9

[24] M. Husung and E. Langbehn. Of portals and orbs: An evaluation of scene transition techniques for virtual reality. In *Proceedings of Mensch Und Computer 2019*, pp. 245–254. 2019. 2

[25] S. Jalayer, Y. Xiao, and M. Kersten-Oertel. Vrnconnect: Towards more intuitive interaction of 3d brain connectivity data in virtual environments. In *2024 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*, pp. 24–29. IEEE, 2024. 9

[26] N. Janaka, R. Cai, A. Ram, L. Zhu, S. Zhao, and K. Q. Yong. Pilotar: Streamlining pilot studies with ohmds from concept to insight. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 8(3):1–35, 2024. 2

[27] P. Jansen, J. Britten, A. Häusele, T. Segschneider, M. Colley, and E. Rukzio. Autovis: Enabling mixed-immersive analysis of automotive user interface interaction studies. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pp. 1–23, 2023. 2

[28] R. S. Kennedy, N. E. Lane, K. S. Berbaum, and M. G. Lilienthal. Simulator sickness questionnaire: An enhanced method for quantifying simulator sickness. *The international journal of aviation psychology*, 3(3):203–220, 1993. 6

[29] J. Knibbe, J. Schjerlund, M. Petraeus, and K. Hornbæk. The dream is collapsing: the experience of exiting vr. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pp. 1–13, 2018. 2

[30] R. Langner, M. Satkowski, W. Büschel, and R. Dachselt. Marvis: Combining mobile devices and augmented reality for visual data analysis. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pp. 1–17, 2021. 3

[31] B. Lee, M. Cordeil, A. Prouzeau, B. Jenny, and T. Dwyer. A design space for data visualisation transformations between 2d and 3d in mixed-reality environments. In *Proceedings of the 2022 CHI conference on human factors in computing systems*, pp. 1–14, 2022. 3

[32] R. Likert. A technique for the measurement of attitudes. *Archives of psychology*, 1932. 6

[33] T. Mahmood, E. Butler, N. Davis, J. Huang, and A. Lu. Building multiple coordinated spaces for effective immersive analytics through distributed cognition. In *2018 International Symposium on Big Data Visual and Immersive Analytics (BDVA)*, pp. 1–11. IEEE, 2018. 3

[34] Q. McNemar. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika*, 12(2):153–157, 1947. 6

[35] P. Milgram, H. Takemura, A. Utsumi, and F. Kishino. Augmented reality: A class of displays on the reality-virtuality continuum. In *Telemanipulator and telepresence technologies*, vol. 2351, pp. 282–292. Spie, 1995. 2

[36] J. Olivier, W. L. May, and M. L. Bell. Relative effect sizes for measures of risk. *Communications in statistics-theory and methods*, 46(14):6774–6781, 2017. 6

[37] T. Piumsomboon, G. Ong, C. Urban, B. Ens, X. Bai, and S. Hoermann. Ex-cit xr: Expert-elicitation of xr techniques for disengaging from ives. In *2022 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*, pp. 710–711. IEEE, 2022. 2

[38] F. Pointecker, J. Friedl, D. Schwajda, H.-C. Jetter, and C. Anthes. Bridging the gap across realities: Visual transitions between virtual and augmented reality. In *2022 IEEE international symposium on mixed and augmented reality (ISMAR)*, pp. 827–836. IEEE, 2022. 2

[39] F. Pointecker, J. Friedl-Knirsch, H.-C. Jetter, and C. Anthes. From real to virtual: Exploring replica-enhanced environment transitions along the reality-virtuality continuum. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pp. 1–13, 2024. 2

[40] E. D. Ragan, S. Scerbo, F. Bacim, and D. A. Bowman. Amplified head rotation in virtual reality and the effects on 3d search, training transfer, and spatial orientation. *IEEE transactions on visualization and computer graphics*, 23(8):1880–1895, 2016. 8

[41] P. Reipschläger and R. Dachselt. Designar: Immersive 3d-modeling combining augmented reality with interactive displays. In *Proceedings of the 2019 ACM international conference on interactive surfaces and spaces*, pp. 29–41, 2019. 3

[42] P. Reipschläger, T. Flemisch, and R. Dachselt. Personal augmented reality for information visualization on large interactive displays. *IEEE Transac-*

tions on Visualization and Computer Graphics, 27(2):1182–1192, 2020.

3

- [43] M. Satkowski, W. Luo, and R. Dachselt. Towards in-situ authoring of ar visualizations with mobile devices. In *2021 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*, pp. 324–325. IEEE, 2021. 1, 2
- [44] D. Schwajda, J. Friedl, F. Pointecker, H.-C. Jetter, and C. Anthes. Transforming graph data visualisations from 2d displays into augmented reality 3d space: A quantitative study. *Frontiers in Virtual Reality*, 4:1155628, 2023. 3
- [45] M. R. Seraji, P. Piray, V. Zahednejad, and W. Stuerzlinge. Analyzing user behaviour patterns in a cross-virtuality immersive analytics system. *IEEE Transactions on Visualization and Computer Graphics*, 2024. 3
- [46] R. Soret, A.-M. Montes-Solano, C. Manzini, V. Peysakhovich, and E. F. Fabre. Pushing open the door to reality: On facilitating the transitions from virtual to real environments. *Applied ergonomics*, 97:103535, 2021. 2
- [47] Student. The probable error of a mean. *Biometrika*, pp. 1–25, 1908. 6
- [48] D. Valkov and S. Flagge. Smooth immersion: The benefits of making the transition to virtual environments a continuous process. In *Proceedings of the 5th Symposium on Spatial User Interaction*, pp. 12–19, 2017. 2
- [49] X. Wang, L. Besançon, D. Rousseau, M. Sereno, M. Ammi, and T. Isenberg. Towards an understanding of augmented reality extensions for existing 3d data analysis tools. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–13, 2020. 3
- [50] R. Wilcox. A robust nonparametric measure of effect size based on an analog of cohen's d, plus inferences about the median of the typical difference. *Journal of Modern Applied Statistical Methods*, 17, 2018. 6
- [51] F. Wilcoxon. Individual comparisons by ranking methods. In *Breakthroughs in statistics: Methodology and distribution*, pp. 196–202. Springer, 1992. 6
- [52] S. Wu, D. Byrne, and M. W. Steenson. " megereality": Leveraging physical affordances for multi-device gestural interaction in augmented reality. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–4, 2020. 3
- [53] L. Zhao, T. Isenberg, F. Xie, H.-N. Liang, and L. Yu. Spatialtouch: Exploring spatial data visualizations in cross-reality. *IEEE Transactions on Visualization and Computer Graphics*, 2025. 3
- [54] F. Zhu and T. Grossman. Bishare: Exploring bidirectional interactions between smartphones and head-mounted augmented reality. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–14, 2020. 3