RippleVision: Unobtrusive Gaze-Dependent Guidance via Directed Wave Motion in Virtual Reality

Jérôme Kudnick, Daniel T. Mayer, Colin Groth, Bipul Mohanto, Ralf Dörner, and Martin Weier

Fig. 1: *RippleVision* guides the gaze (highlighted as eye symbol) using waves toward the point of interest (object inside red outline). In stereoscopic vision, the wave is shifted by half a cycle in one eye relative to the other.

Abstract— The inherent freedom of exploration in virtual reality poses a challenge for directing user attention toward relevant points or objects of interest, which are potentially located outside the user's field of view or are temporarily obscured. Accordingly, guidance must sustain users' perceptual focus while preserving the sense of presence. This paper presents *RippleVision*, a subtle gaze guidance technique using wave-like ripples that modulate brightness with inverted polarity between the eyes and appear only within a cone from the point of interest to the gaze location. A calibration study defined detectability and acceptability thresholds for the cues, followed by a comparative user study against four state-of-the-art techniques. Results show *RippleVision* achieves similar search times to clearly visible techniques while significantly reducing visual dominance. Moreover, its effectiveness scales with cue visibility, improving search performance with minimal impact on perceived visual obstruction and presence.

Index Terms— Gaze guidance, subtle cue, immersion preserving, concentric ripple slice, stereo inverse, screen-space shader, eye tracking, extended reality.

1 Introduction

Directing audience attention is central to storytelling and training scenarios in immersive virtual environments. In head-mounted displays (HMDs), both interactive virtual scenes and 360° narrative video present more content than fits within the user's instantaneous field-of-view (FoV) or camera frustum [8]. As a result, plot- or task-critical events often occur outside the line of sight or, in the case of 360° video, outside the current head orientation, leading to missed narrative beats and increased time-to-orient duration. In order to control the user's attention, their gaze must be guided to the relevant points of interest (POIs).

- *Jérôme Kudnick is with RheinMain University of Applied Sciences. E-mail: jerome.kudnick@hs-rm.de*
- *Daniel T. Mayer is with RheinMain University of Applied Sciences. E-mail: daniel.mayer@student.hs-rm.de*
- *Colin Groth is with New York University. E-mail: c.groth@nyu.edu*
- *Bipul Mohanto is with University of Rostock. E-mail: bipul.mohanto@uni-rostock.de*
- *Ralf Dörner is with RheinMain University of Applied Sciences. E-mail: ralf.doerner@hs-rm.de*
- *Martin Weier is with RheinMain University of Applied Sciences. E-mail: martin.weier@hs-rm.de*

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

Previous work approaches gaze guidance techniques using (i) *overt* diegetic/ non-diegetic indicators, e.g., arrows, halos, wedges, or overlays, and (ii) *subtle*, content-preserving modulations that leverage properties of the human visual system (HVS) [2–4, 7, 20, 24, 27]. Overt cues are easily interpretable and effective; however, they can occlude the content and break immersion. More subtle cues typically require careful calibration – especially in complex, realistic scenes and 360° video – as these can interact with background motion, colors, and contrast, or leverage physiological triggers such as triggered eye blinks to mask redirection [9].

Due to the high rod density in the periphery, motion is a powerful peripheral stimulus: onsets and low-frequency temporal changes bias saccades pre-attentively [22]. However, naïve motion cues are plainly visible and, on their own, do not encode directionality. Therefore, we seek a guidance metaphor that (i) resides primarily in the periphery, (ii) communicates direction to the POI, and (iii) preserves the appearance of the scene, being non-obstructive and non-obtrusive.

RippleVision addresses this need with gaze-contingent *spatio-temporal waves* in screen space. The wave pattern can be used to guide to POIs throughout a 360° FoV, for example, when rendering for an HMD. The ripples originate at the current gaze point (GP) and propagate toward a POI while remaining largely confined to the periphery. The ripples modulate luminance with short spatial wavelength and moderate temporal frequency and are rendered as a phase-inverse stereo pairs (stereo-inverse brightness between eyes) with the goal of increasing salience without altering perceived structure.

In order to minimize obtrusiveness, *RippleVision* uses opposing

amplitude-modulated luminance in both eyes to maintain the overall appearance of the color when viewed binocularly, but at the same time trigger perceptual reflexes. The binocular stereo-inverse brightness increases peripheral detectability while keeping the effect near threshold at the fovea, inspired by prior stereo-inverse brightness modulation filters [7]. Parameters are selected via a short calibration to achieve just-noticeable and acceptable yet non-dominant settings that generalize across scenes.

This design allows *RippleVision* to be used in various scenarios, such as directional guidance in training, educational exploration, narrative 360° videos, and other scenarios benefiting from directional guidance of the gaze.

With this methodology, perceptual studies were conducted to answer the following research questions:

RQ1: Does *RippleVision* significantly reduce search times compared to existing guidance techniques?

RQ2: Is *RippleVision* less obtrusive than state-of-the-art?

In summary, our novel contributions in this work are as follows:

- The introduction of *RippleVision*, a novel screen-space technique employing stereo-inverse spherical wave patterns with modulated amplitude and brightness to subtly direct visual attention.
- A calibration study to establish parameter configurations for *RippleVision*, enabling its application at minimally perceptible or clearly detectable yet non-dominant levels.
- An extensive comparative evaluation of *RippleVision* against multiple established guidance techniques, demonstrating its effectiveness and unique advantages.

2 Background and Related Work

In this section, we first provide conceptual foundations on gaze guidance before exploring notable examples of previous work.

2.1 Conceptual Foundations

Gaze guidance techniques can be broadly understood in terms of four categories of cues: **Directness**, **Subtlety**, **Coverage**, and **Diegesis** [19, 21, 27].

Directness: Direct cues capture gaze immediately and require no instruction, typically highlighting the POI or its vicinity; indirect cues, by contrast, must be interpreted to be understood, such as the direction indicated by an arrow.

Subtlety: Overt guidance techniques use clearly visible cues to capture the attention of the user to guide their gaze toward a POI. These techniques are often very effective but can distract the user from their original goal and potentially grab attention in unexpected ways, thus disrupting the user experience of the virtual worlds. Subtle guidance techniques, on the other hand, attempt to address this challenge by being less noticeable. An often implemented method to obtain this goal is to be less visible through high transparency, slight changes in color modulations, or using inherent properties of VR.

Coverage: Coverage categorizes techniques based on their coverage of the visual FoV. Since the POI can be outside the user' FoV, many techniques first attempt to shift the gaze in the general direction of the POI, before subsequently assisting the user locate the exact target. These techniques often use indirect cues, as direct highlighting of the POI is only possible if the target is visible in the first place. In-view guidance, by contrast, places direct cues near the target to maximize effectiveness.

Diegesis: Diegetic techniques incorporate cues as part of the virtual worlds itself. As part of the world, these cues can direct the user to the POI without disrupting the user's sense of presence while having the advantage of being clearly visible. Lange et al. [15], for example, use bee-inspired swarm motions as cues to guide attention, while Norouzi et al. [17] utilize animals for diegetic guidance toward a POI. However, due to their direct integration, these cues need to fit the narrative of the virtual worlds and are thus difficult to generalized. Therefore, these techniques will not be considered for further evaluation.

Table 1: Notable examples of gaze guidance techniques for extended reality (XR) from related work. Illustrated are three inherent properties of the cues used by the techniques. Directness refers to whether a cue works direct or indirect; subtly describes whether a cue is overt or subtlety; coverage pertains to the techniques ability to guide to POIs outside of the user's view.

Name	Directness	Subtlety	Converage
Arrow [3, 16, 27]	Indirect	Overt	out-of-view
DynSWAVE [20]	Indirect	Overt	out-of-view
WedgeVR [10]	Indirect	Overt	out-of-view
HaloVR [10]	Indirect	Overt	out-of-view
Attention funnel [2]	Indirect	Overt	out-of-view
3D Radar [4]	Indirect	Overt	out-of-view
EyeSee360 [11]	Indirect	Overt	out-of-view
ChromaGazer [24]	Direct	Subtle	in-view
SIBM [7]	Direct	Subtle	out-of-view
Deadeye [13]	Direct	Subtle	in-view
WarpVision [14]	Direct	Subtle	semi out-of-view

2.2 Gaze Guidance Techniques

Prior research spans a broad spectrum of gaze-guidance techniques; representative examples are summarized in Tab. 1. Arrows are one of the simplest and most common examples [3, 26, 27]. With strategic placement, arrows allow for out-of-view guidance and immediate recognition. Even though arrows are indirect cues, they are understood instantly and considered to be highly effective. However, arrows are also exceedingly overt and distracting.

Techniques such as *WedgeVR* and *HaloVR* by Gruenefeld et al. [10] use overt geometric shapes placed directly in the user's FoV and are primarily designed to guide to out-of-view POIs. *HaloVR* renders circles in the user's peripheral field, centered on the direction of out-of-view POIs; consequently, it serves as an indirect cue. The gaze guidance technique *DynSWAVE* [20] projects waves onto a sphere around the user, which are directed toward the position of the POI. The distance to the target is emphasized by altering the speed of the waves. Although we also use waves, our technique does not overlay the entire screen using waves on a sphere, but only a section between the POI and GP inside the screen space. Furthermore, the waves are presented stereo-inverse for each eye and adapted using eye tracking. The overt out-of-view technique *Attention Funnel* [2] uses repeating patterns (such as squares or rings) as cues along a curve connecting the camera position and the POI. These indirect cues are supposed to act as a funnel toward the indicated location. *3D Radar* [4] and *EyeSee360* [11] act in a radar-like way and aim to show the user the location of multiple POIs at the same time. These techniques are especially useful in very limited FoV setups to assist in understanding the layout of the room. Nonetheless, they may require additional mental effort [19]. A comprehensive comparison by Woodworth and Borst [27] reaffirmed the efficiency of overt cues, offering a significant advantage over unassisted search tasks, where *DynSWAVE* was categorized as one of the best-performing cues. Quinn and Gabbard [19] concluded that visual feedback provided by techniques, such as gaze-dependent waves in *DynSWAVE*, is an effective characteristic in gaze guidance.

ChromaGazer [24] is an example of a subtle guidance technique, which uses color vibrations as direct cues to guide toward in-view POIs. This technique is designed to be unobtrusive and maintain low cognitive demand. The subtle technique *SIBM* by Grogorick et al. [6, 7] exploits the binocular rivalry through stereo-inverse brightness modulations. To guide attention to in-view POIs, a spherical cue is placed directly on the target. For out-of-view POIs, the cue moves in a repeating motion outward. While we also use the rivalry effect of SIBM, instead of spherical cues, our technique uses gaze-dependent wave motions that are not placed on top of the POI, but direct toward it. This gaze dependency further allows our technique to never overlay or alter an area that the user is currently looking at, and only act within their peripheral FoV. Another technique that uses the stereoscopic properties

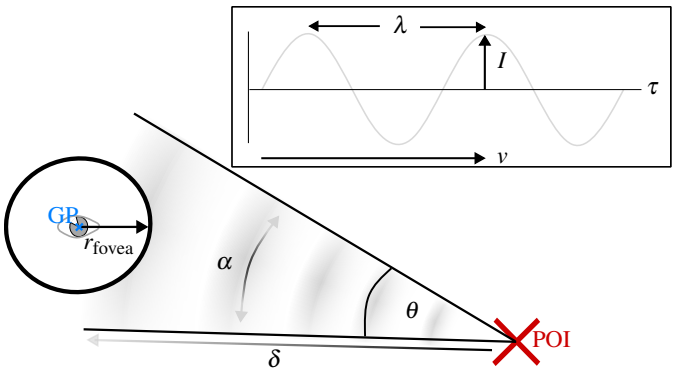

Fig. 2: *RippleVision's* parameters to configure spatial wavelength (λ), intensity (I), velocity (v), and falloffs (α, δ) of the effect. *RippleVision* originates from the POI displayed only inside the spatial extent (θ) and outside the spherical region (r_{fovea}) around the GP.

of VR is *Deadeye* [13], which renders the target object in only one eye. The authors argue that the preattentive action of the cue subtly attracts the user's gaze. Out-of-view guidance is not feasible with *Deadeye*. Kudnick et al. introduced the subtle guidance technique *WarpVision* [14]. This technique applies spatial curvature to enlarge the area of the POI, similar to a fisheye lens. While this direct approach has been shown to provide effective performance, it is primarily designed for subtle and immersion-preserving guidance. *WarpVision's* cues are gaze-dependent and decrease in intensity with a smaller distance to the POI. The technique was evaluated as efficient for in-view guidance, while out-of-view POIs may be addressed by large curvatures that extend into the viewport. However, *WarpVision's* efficiency is highly dependent on its intensity. While higher intensities attracted the gaze significantly faster, they also had a higher negative impact on the user's immersion. Hence, we consider *WarpVision* to allow only for a semi out-of-view coverage.

These examples show that efficient gaze guidance techniques for out-of-view POIs predominantly use overt indirect cues, which not only require an understanding of the cue but are also clearly visible, thus impacting the user's immersion in the scene. On the contrary, we propose *RippleVision*, a subtle direct out-of-view guidance technique. Our primary goal is to guide the user to out-of-view POIs without interrupting their sense of presence in the virtual environment.

3 RippleVision

The technique *RippleVision* uses wave-like temporal modulations to guide attention. The wave patterns propagate from the eye-tracked GP towards the POI. They are expressed as modulations in local brightness with a phase shift for each eye. A visual example is given in Fig. 1. The waves are only visible inside a circular slice that is positioned at the POI and extends to the distance of the GP. As a result only a limited portion of the screen is modulated by *RippleVision's* waves. This portion is maximized when the POI lies opposite to the GP (effectively directly behind the user), resulting in the modulation over a broad fraction of the screen space in the direction of the POI. For the purpose of ensuring that the user's foveal vision stays unobstructed and prevent distractions, the waves are not displayed in the area directly around the fovea. Thereby, *RippleVision* is designed to use a cue that provides the user with direct feedback based on their GP. Moreover, it indicates the direction and distance to the POI, as recommended in the systematic review by Quinn and Gabbard [19].

RippleVision can be calibrated using multiple parameters (see Fig. 2). The spatial wavelength and intensity (amplitude) of the waves are controlled by $\lambda, I \in \mathbb{R}^+$, respectively. The temporal progression of the wave is determined by the velocity parameter $v \in \mathbb{R}^+$. The angular span of the circular slice in which the waves are visible is defined by $\theta \in [0, 2\pi]$ in normalized screen space and adjusted according to the aspect ratio. To avoid distracting central vision and ensure visibility only in peripheral vision, the parameter $r_{fovea} \in \mathbb{R}^+$ can be used to adjust the radius of the area around the gaze position that remains unaffected by *RippleVision*. The final configuration parameters are $\alpha \in \mathbb{R}^+$, which controls the falloff of the wave modulation toward the edges of the circular slice, and $\delta \in \mathbb{R}^+$, which controls the falloff based on the distance to the POI. Both parameters α and δ are meant to ensure a smooth transition at the slice's boundary.

The waves from *RippleVision* are generated by a post-processing fragment shader that implements a circular wave function. The screen-space dimensions are denoted by w (width) and h (height). The position of the POI in screen space is denoted as $C \in \mathbb{R}^2$, the position of the GP in screen space is denoted as $T \in \mathbb{R}^2$, and the screen-space UV coordinates of each pixel as $U \in \mathbb{R}^2$. Using these definitions, the aspect-ratio–corrected distance to the target can be computed as follows:

$$d_c = \sqrt{(C_x - U_x)^2 + \left((C_y - U_y) \cdot \frac{h}{w}\right)^2}$$

With this, the ripple strength can be computed for the left

$$\text{ripple}_l(t) = 1 + I \cdot \sin\left(\frac{-d_c - t}{\lambda}\right)$$

and phase-shifted for the right eye by half a wavelength to implement the stereo-inverse brightness modulation

$$\text{ripple}_r(t) = 1 + I \cdot \sin\left(\frac{-d_c - t + \lambda \cdot \pi}{\lambda}\right)$$

Here, $t = \tau \cdot v$ denotes the user-defined velocity-adapted system time τ in seconds.

These ripple functions are only executed if the UV coordinates are inside the circular slice defined by θ and outside r_{fovea}. To this end, the relative angle ϕ of each UV coordinate is calculated using the directional vector $\vec{d}_{\text{rel}} = T - C$ and the relative direction to each UV coordinate $\vec{u}_{\text{rel}} = U - C$ as

$$\phi = \text{clamp}(\hat{u}_{\text{rel}} \cdot \hat{d}_{\text{rel}}, -1, 1)$$

The normalized vectors for $\vec{d}_{\text{rel}}$ and $\vec{u}_{\text{rel}}$ are denoted as $\hat{d}_{\text{rel}}$ and $\hat{u}_{\text{rel}}$, respectively. A pixel at the UV coordinate is only modulated if this ϕ is within the limits defined by the configurable parameter θ. In addition, to hide the transitions along the edges of the slices, an additional α_{final} value is used to lower the intensity of the ripple effect toward the edges.

$$\alpha_{\text{final}} = \left(\text{clamp}\left(\frac{\theta/2 - acos(\phi) - \alpha\theta/2}{\alpha\theta/2}, 0, 1\right)\right)^{\alpha}$$

Apart from UV coordinates that fall outside the slice angle, *RippleVision's* cue is also disabled beyond a certain distance and within the central vision area, which is defined by the parameter r_{fovea}. To this end, we project the vector $\vec{d}_{\text{rel}}$ onto the direction $\hat{u}_{\text{rel}}$

$$\vec{d}_{\text{proj}} = (\vec{d}_{\text{rel}} \cdot \hat{u}_{\text{rel}})\hat{u}_{\text{rel}}$$

and compute the maximal distance $d_{\max}$ of the ripple effect. The distance is either given by the length of $\vec{d}_{\text{rel}}$ or the first intersection with the foveal area.

$$d_{\max} = \begin{cases} \|\vec{d}_{\text{rel}}\| & , \text{if } d > r_{fovea} \\ \min\left(\|\vec{d}_{\text{proj}} + \vec{g}\|, \|\vec{d}_{\text{proj}} - \vec{g}\|\right) & , \text{else} \end{cases}$$

with $d = \|\vec{d}_{\text{rel}} - \vec{d}_{\text{proj}}\|$ and $\vec{g} = \sqrt{r_{fovea}^2 - d^2}\,\hat{u}_{\text{rel}}$.

Having determined $d_{\max}$, a δ_{final} value is computed to further tone down the ripple effect based on distance.

$$\delta_{\text{final}} = \left(1 - \text{clamp}\left(\frac{d_c}{d_{\max}}, 0, 1\right)\right)^{\delta}$$

Finally, the strength of the ripple effect is computed using

$$\text{ripple}_{\text{final}}(t) = 1 + \alpha_{\text{final}}\,\delta_{\text{final}}\left(\text{ripple}_{l|r}(t) - 1\right)$$

This function $\text{ripple}_{\text{final}}(t)$ is used to adjust the brightness of each pixel within the CIELAB color space. Specifically, the color of each pixel is first converted to the CIELAB space, where the final ripple factor, $\text{ripple}_{\text{final}}$, is multiplied with the L (lightness) component. The modified color is then converted back to sRGB. The CIELAB color space was chosen because it is perceptually uniform; that is, equal distances within this space correspond to approximately equal perceived differences in color. This property is crucial because the visual effect produced by *RippleVision* on corresponding regions in both eyes is designed to be stereo-inverse, so that the color perceived by both eyes conveys the average. A sample Unity project, which includes the shader code and other relevant files, is publicly available[1].

4 Perceptual Studies

Two studies were conducted in VR to determine the optimal properties and effectiveness of *RippleVision*. Prior to both studies, participants were informed about the procedure and provided their written consent. Based on previous experiments, the risks for both user studies were assessed as minimal. For this reason and in accordance with the legal framework of the institution, no ethical review board was involved in the review of the study. Participation was voluntary, data was anonymized, and potential risks associated with the use of VR systems were discussed individually and summarized in the declaration of consent. The first study is a calibration study that was performed to explore the parameter space of *Ripplevision*. Based on these insights, an application study was conducted to investigate the effectiveness of *RippleVision* and compare it with other gaze guidance techniques. This study aims to provide insights into our research questions, for which three hypotheses are formulated.

- **(H1)** *RippleVision* reduces search times by as much or more than state-of-the-art.
- **(H2)** *RippleVision* does not lose its effectiveness even when the POI is out-of-view.
- **(H3)** *RippleVision* is less obtrusive than state-of-the-art.

4.1 Apparatus and Stimuli

Both studies were conducted using the Unity game engine (version 6000.0.44f1) and the VIVE Focus 3 HMD with eye-tracking module, which was individually calibrated. To evaluate *RippleVision*, two virtual scenes (i) artificial and (ii) realistic were created (Fig. 3). The artificial scene is based on the environment used by Kudnick et al. [14] (Fig. 3a). The user is positioned in the center of 250 geometric objects of varying shapes and colors (Fig. 3c), which circle around the user in random directions and constant speeds with enabled collisions and act as distractors. The object acting as POI is a moving sphere that is only instantiated once in a random location and color. The realistic condition uses a dark, Cyberpunk-styled environment (Fig. 3b) adapted from [12], intended to approximate real-world VR applications such as games. Here, the user is located near the center of the platform and is surrounded by multiple different polygonal annuli (Fig. 3d). In contrast to the previous scene, these objects are static and always presented to the user in a front-facing orientation with a random rotation around their center on their local z-axis. In this scene, motion is conveyed through a dynamic background of flying ships and moving cables. The object used as POI is a uniquely instantiated hexagon.

4.2 Calibration Study

Prior to evaluating *RippleVision* against other techniques, suitable parameter settings had to be determined. Based on the formulation in Sec. 3, the four main parameters were calibrated: Intensity I, wavelength λ, velocity v, and angle θ. The artificial scene was used for the calibration because of minimal confounding factors. Participants viewed the artificial scene with a static sphere as a POI in front of them and *RippleVision*'s directional cue. Six static objects were placed in the scene to provide reference points.

Participants and Procedure: The calibration study involved eight participants (7M, 1F; age 21–40, M = 26.38, SD = 3.38), balanced in VR experience (4 low, 4 high). The study was divided into two identical parts: first, the calibration of the parameters to the detectability threshold, and second, calibration to the acceptability threshold. Both parts had 8 trials, each with the task to adjust the given parameter by button presses to either be visible but not distracting (first part) or barely visible (second part), with no time limit. The first three trials adjusted the intensity parameter, which in vision experiments is usually most important [1, 5, 18, 23]. All parameters (including intensity) were randomly initialized in $I \in [0, 0.5]$, $\lambda \in [0.0005, 0.0045]$, $v \in [0.01, 0.05]$, and $\theta \in [25, 90]$. Those min/max values were chosen in equal intervals around the mean of previous empirical testing, with a variance that still ensures the method is visible in any case. The next three trials adjusted the wavelength and velocity parameters simultaneously, as we found a strong correlation between both parameters. Participants had two pairs of up/down buttons and could freely explore their parameter weighting. Here, the intensity was fixed to the average of the former trials, while the other three parameters were again randomly initialized. The last two trials adjusted the angle parameter, with all others fixed to the formerly chosen average values.

Results and Discussion: The results depicted in Fig. 4 present the results of the parameter calibration. For intensity, a clear distinction emerges between detectability and acceptability thresholds, providing a direct indicator for parameter adjustment. In contrast, this separation was less pronounced for the remaining parameters. Spatial wavelength showed the greatest variability, with broad distributions across both thresholds, suggesting either a limited functional relevance or a particularly strong perceptual influence, warranting further investigation. Velocity demonstrated minimal differences between thresholds and a narrow distribution, indicating general agreement among participants and reducing the need for additional study. The angle parameter exhibited somewhat greater variability, but comparable means and distributions across thresholds. This suggests it can be set independently of cue visibility. Since angular variation grows substantially when the point of interest lies in peripheral view, this parameter was not considered critical for the current study. Based on these findings, six *RippleVision* variants were selected for further evaluation of subjective perception and objective efficiency (2 intensities × 3 spatial wavelengths × 1 velocity × 1 angle; see Tab. 2). The intensity values correspond to the median detectability and acceptability thresholds, while the three spatial wavelengths were derived from the 25th, 50th, and 75th percentiles across both thresholds. Velocity and angle were set to the mean values obtained from this calibration study. The two additional fall-off parameters, α and δ, were defined in consultation with an expert, as we deemed an individual calibration to be unnecessary. The foveal radius r_{fovea} was chosen to reflect the 5° foveal visual angle [25].

4.3 Application Study

The second study aims to evaluate the effectiveness of *RippleVision*. It was carried out using a within-subjects design.

Participants: The application study included 21 participants (14M, 7F; age: 19-58, M = 31.2, SD = 10.6) with normal or corrected to normal vision. The participants' experience with VR was distributed as follows: 10 low, 6 medium, and 5 high. Data from one participant were excluded from further analysis as they did not correctly follow the study procedure.

Techniques: We compare the six *RippleVision* settings established in the calibration study to four common gaze guidance techniques. A visual example of all techniques is presented in Fig. 5. We consider two baselines for the search task: A baseline for slow searching is given by testing subjects without assistance. A second baseline for fast search is given by the technique *ARROW* (Fig. 5a). *ARROW* uses two simple red arrows to direct the subject toward the POI: The first arrow is positioned slightly above the user to indicate the direction, and the other arrow is

[1] https://github.com/XRG-HSRM/RippleVision

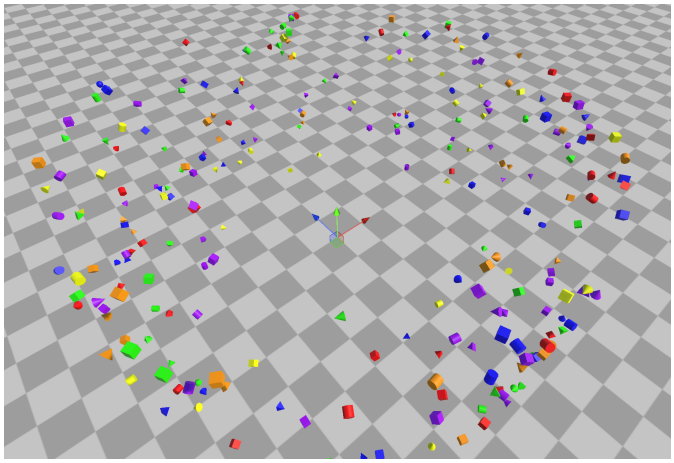

(a) Artificial scene

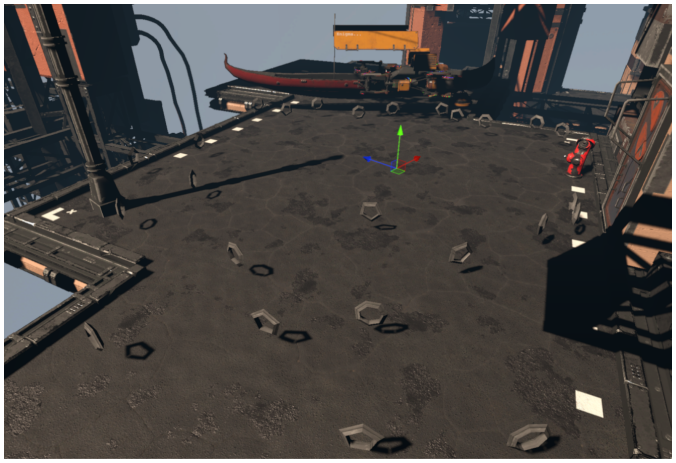

(b) Realistic scene

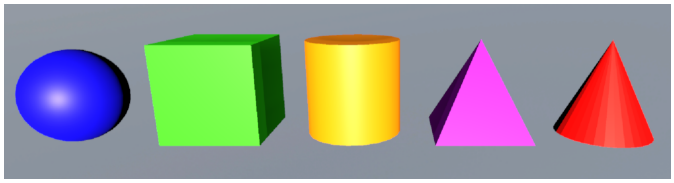

(c) Shapes used in the artificial scene

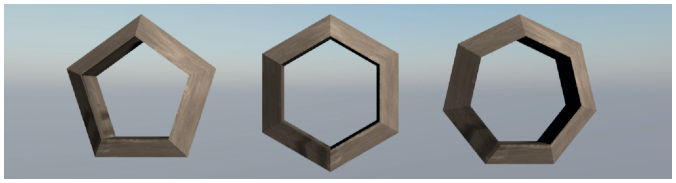

(d) Shapes used in the realistic scene

Fig. 3: Both scenes were used in the study design. In the artificial scene (a), the search task was to find a sphere while being circled by geometric shapes. These shapes (c) were always shown with a random color selected from blue, red, green, yellow, orange, and purple. In the realistic scene (c), the task was to find a hexagon between penta- and heptagons (d). In both scenes, the participant is placed at the coordinate cross visible in the center.

Table 2: *RippleVision* variants used in the application study. To simplify interpretation, the spatial wavelength λ is also given in degrees of visual angle and the velocity v as pixels per second (pps). The angle θ can be interpreted directly as degrees.

Variant	I	λ (in deg.)	v (in pps)	θ in deg.
RV-L1	0.09	0.0015 (0.93°)	0.035 (538)	49°
RV-L2	0.09	0.00243 (1.50°)	0.035 (538)	49°
RV-L3	0.09	0.0036 (2.22°)	0.035 (538)	49°
RV-H1	0.36	0.0015 (0.93°)	0.035 (538)	49°
RV-H2	0.36	0.00243 (1.50°)	0.035 (538)	49°
RV-H3	0.36	0.0036 (2.22°)	0.035 (538)	49°

directly above the POI, pointing downward. The downward arrow was added to directly connect the gaze to the POI, thereby compensating for the general disadvantages of arrows, as discussed by Woodworth and Borst [27]. This technique is hypothesized to be the fastest possible visual guidance technique based on the general familiarity with arrows and their proven usefulness. Further comparisons with other shader-based guidance techniques include *DynSWAVE* [20] (Fig. 5b), *WarpVision* [14] (Fig. 5c), and *Deadeye* [13] (Fig. 5d). *DynSWAVE* and *WarpVision* were chosen for their conceptual similarity of using gaze-dependent cues and proven effectiveness, whereby *DynSWAVE* additionally utilizes a similar wave-based cue as *RippleVision*, while *WarpVision* was shown to have low negative impact on immersion. *Deadeye* was selected for comparison with POI highlighting based cue. Examples for two intensities of the employed *RippleVision* settings are shown in Fig. 5e and Fig. 5f.

Procedure: Throughout the study, participants were allowed unrestricted head and body rotation, while other movements were restricted. All techniques were briefly presented in an initial demonstrative search task (a simplified version of the artificial scene with only 10 distractors), which ensured that differences in performance can be attributed to the techniques' efficiencies rather than the speed of initial comprehension. For the task, participants were told to find the POI as quickly as possible. A total of 44 search tasks were run in randomized order, 22 in each scene: 11 (ARROW + DynSWAVE + WarpVision + Deadeye + 6 *RippleVision* settings + unassisted search) x 2 (repetitions) x 2 (scenes). The participants were given a short break between scenes.

To gain insight into subjective receptions of the techniques, the participants were asked to rate three statements after each run:

(**S1**) The technique obstructed my vision.

(**S2**) The technique disrupted my sense of presence.

(**S3**) The technique was visually dominant.

The first statement was chosen to determine in what way the techniques impact and interact with the user's vision and view of the scene, while the second statement assessed whether the technique has any impact on immersion. The last statement estimated the subtlety of the techniques. All questions were to be rated on a Likert scale ranging from -2 (disagree) to +2 (agree).

Upon finding the POI in each search task, participants were instructed to press the trigger button on the VR hand controller. Their current gaze was used as an indicator for the correct identification of the POI. If the selection was incorrect, they could continue searching; otherwise, the subjective questions were presented, and ratings had to be selected before continuing with the next search task. If the POI was not found within three minutes, the search task was automatically terminated and the subjective questions were displayed immediately, before the next run started. This ensured that participants did not get stuck in a singular search task. The entire procedure was carried out in around 40 minutes, depending on the participant's proficiency with the task.

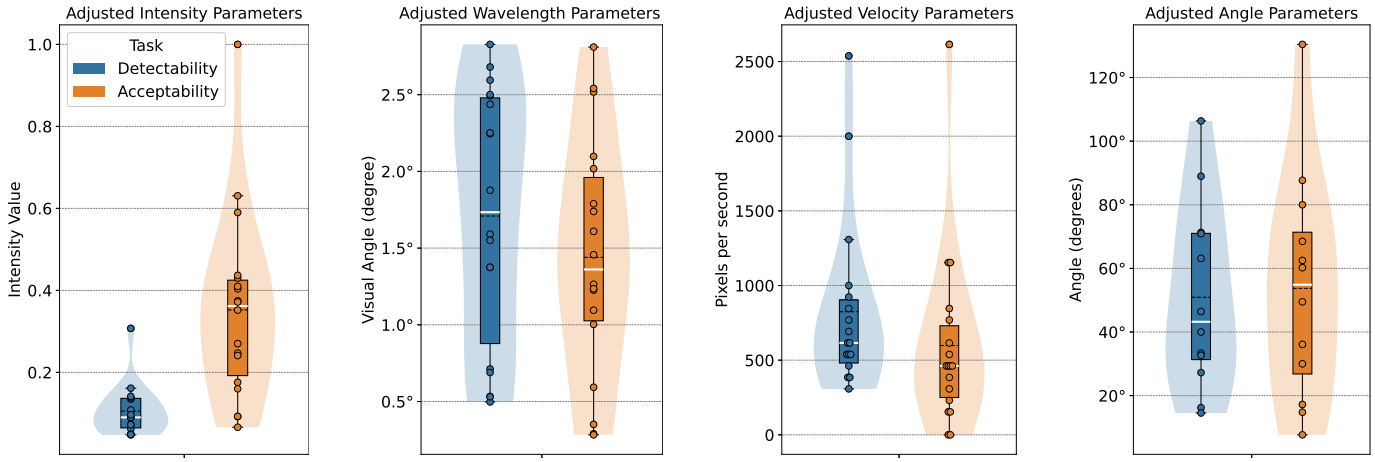

(a) Results for I

(b) Results for λ, given in degrees of visual angle

(c) Results for v, given in pixels per second

(d) Results for θ in degrees

Fig. 4: Calibration study results for the tuned parameters of intensity I (a), wavelength λ (b), velocity v (c), and angle θ (d) between the detectability and acceptability threshold. The average is depicted as a dotted black line, the median as a white line.

5 Results

The results of the application study will be separated into objective and subjective results.

5.1 Objective Results

In order to gain insight into the effectiveness of the employed techniques and answer **H1**, we explore the times until the participants successfully found the POI. Figure 7 shows the mean and median search times and the median first fixation on the POI. As expected, the two baselines achieved the lowest and highest median search times, respectively. *ARROW* (fast baseline) is closely followed in efficiency by *DynSWAVE*, *RV-H2*, and *RV-H3*. *WarpVision* places roughly between these four and the remaining techniques. Especially noticeable is the time variation of the remaining techniques between scenes. *Deadeye*, for example, shows barely any effect in the artificial scene, while at the same time offering a noticeable support in the other scene. Vice versa, all low-intensity *RippleVision* settings offered less help in the realistic scene.

An analysis of saccadic movement showed that there are neither significant differences in the number of saccades per second nor the average saccade angles between techniques; with the exception of *RV-H2* in the artificial scene, which led to significantly less saccades than *Deadeye* (Wilcoxon signed-rank: $p_{\text{holm}} < 0.001$, $W = 108$) and the unassisted search ($p_{\text{holm}} < 0.036$, $W = 188$).

For further evaluation of the effectiveness of techniques and to take the right-censored data into account (due to the time limit of the search), the Kaplan-Meier estimator is used, which approximates the probability of an event (finding the POI) not yet occurring at a specific point in time (time since the search task started). For analysis of significant differences between these curves, the pairwise cox proportional hazard regression is used with Bonferroni-Holm corrected p-values for each individual scene (Fig. 6c). Figure 6a shows the results in the artificial scene, where the strong effect of *ARROW* is confirmed through its stark downward curve and significantly more effective guidance than any other technique. *RV-H3* and *DynSWAVE* show similar patterns, whereby *DynSWAVE* falls off in efficiency toward the end, indicating that not all participants successfully found the target within the first 10 seconds. All other techniques except for *Deadeye* closely follow *RV-H3* in efficiency, whereby *RV-L1* performs the worst of all *RippleVision* settings and is significantly worse than all high-intensity settings. The different spatial wavelengths used by *RippleVision* do not show significant differences between each other (when looking at the same intensity). However, as can be observed from the p-values reported in Fig. 6c, there are many significant differences between the spatial wavelengths compared to other techniques. *RV-L1*, for example, is the only *RippleVision* setting that is significantly worse compared to *DynSWAVE* ($z = -4.06, p = 0.001$), while *RV-L2* with its slightly higher spatial wavelength is comparable ($z = -2.06, p = 0.624$). The least effective technique was *Deadeye*, which was the only technique to show no significant difference to the unassisted search, with a probability of 80% to successfully find the POI after one minute.

In the realistic scene (Fig. 6b), the techniques show different effectiveness. While *ARROW* is still the fastest to guide toward the POI and closely followed by *DynSWAVE*, it is no longer significantly better than *DynSWAVE*. All *Ripplevision* versions fall off compared to their capability in the artificial scene, with the exception of *RV-H2*. This is especially noticeable with *RV-L1*, which does not provide any visible assistance. Unlike in the artificial scene, *Deadeye's* curve falls clearly below the unassisted search, and the provided assistance is similar to *RV-L2* and *RV-L3*, but not significantly better. The only *RippleVision* setting that is not significantly worse than *DynSWAVE* is *RV-H2* ($z = 2.55, p = 0.214$). *WarpVision* falls between all other techniques and is similarly effective to *RV-H3* and not significantly worse than any *RippleVision* setting. Therefore, we accept **H1** and conclude that *RippleVision* can be as effective in guiding users to a POI as state-of-the-art guidance techniques.

Since there were no significant differences in search times between *RippleVision* settings and *WarpVision* in the artificial scene, a further comparison is employed to gain insights into **H2**. In Fig. 8, the time of a search task for these two techniques is depicted depending on the initial angle of the POI to the GP in the artificial scene. It can be observed that *WarpVision* lies between the *RippleVision* variants with the lowest and highest spatial wavelength and intensity settings. However, with an increase in initial angle, *RippleVision's* effectiveness does not decrease similarly as *WarpVision* does. Even if the average search time using *RV-L1* is longer, and its confidence interval is larger, this variant has a similar tendency to *RV-H3*, seemingly only offset by around eight seconds. Therefore, we accept **H2** and conclude that *RippleVision* does not lose its effectiveness even when the POI is out-of-view.

5.2 Subjective Results

In order to answer **H3**, subjective results will be analyzed using Likert plots and the Wilcoxon signed-rank test for significant differences, with Bonferroni-Holm adjusted p-values applied to both scenes independently. This test was chosen because Likert-scale responses are ordinal and usually not normally distributed. The subjective rating on the obstruction of vision caused by techniques (Fig. 6d) shows a clear tendency across both scenes. The greatest obstruction was caused by *ARROW*, with more than $3/4$ of the participants perceiving it as

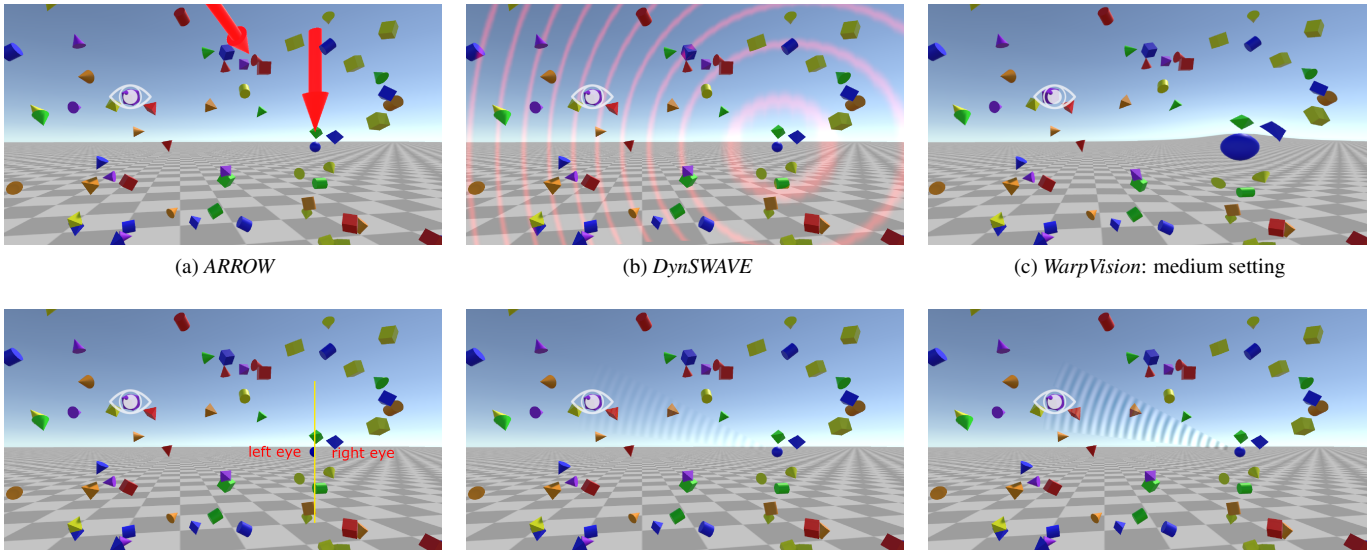

(a) *ARROW* (b) *DynSWAVE* (c) *WarpVision*: medium setting

(d) *Deadeye* (e) *RippleVision*: low intensity, intermediate wavelength (f) *RippleVision*: high intensity, intermediate wavelength

Fig. 5: Exemplary depiction of search tasks using the employed techniques *ARROW* (a),*DynSWAVE* (b), *WarpVision* (c), *Deadeye* (d), and *RippleVision* (e, f). The POI is the blue sphere on the right side. The gaze position is on a purple cube, marked by the white eye symbol.

obstructive, closely followed by *DynSWAVE*. All other techniques are perceived as significantly less obstructive in both scenes, while also not showing significant differences between each other in the artificial scene. In the realistic scene, there was a notable difference between *RV-L1* and *WarpVision* ($z = -3.62$, $p < 0.01$). This *RippleVision* version was also significantly less obstructive than all versions with a long spatial wavelength.

The subjective rating on the techniques' impact on the sense of presence (Fig. 6e) shows a greater difference, *ARROW* and *DynSWAVE* have a great negative impact on the sense of presence, with the majority of participants perceiving them as invasive. Notably, there is a large difference between the two scenes, with almost twice the people now fully agreeing with the statement. There was no consensus among participants regarding *Deadeye*. Although the sense of presence was not affected according to roughly $2/3$ of the participants, the other $1/3$ perceived such a disruption. *WarpVision* proved to be less disruptive in the artificial scene, which is also validated by having no significant differences to any techniques besides *ARROW* ($z = 5.09, p < 0.01$) and *Deadeye* ($z = 5.06, p < 0.01$). In contrast, the differences between all *RippleVision* versions and *WarpVision* are significant in the realistic scene (all with $p < 0.01$). *RF-L1* further proved to be significantly less harmful to the sense of presence than any *RippleVision* setting configured to have a long spatial wavelength.

The subjective rating on visual dominance shows the most varying distribution (Fig. 6f). *ARROW* and *DynSWAVE* are almost entirely perceived as visually dominant. The next-greatest dominance shows *WarpVision*, with almost equally distributed ratings in both scenes, but significantly less dominance than *ARROW* and *DynSWAVE* (all comparisons $p < 0.01$). The *RippleVision* variants show a nearly steady increase in dominance depending on their intensity and spatial wavelength employed. All low-intensity variants are significantly less dominant than their high-intensity counterparts in the realistic scene, while being similarly dominant in the artificial scene. *Deadeye* lies between these variants. It was perceived to be significantly less dominant than *RV-H3* in the artificial ($z = -3.13, p = 0.04$) scene and more dominant than *RV-L1* ($z = -3.85, p < 0.01$) and *RV-L2* ($z = -3.17, p = 0.02$) in the realistic scene. Overall, we accept **H3** and conclude that *RippleVision* is less obtrusive than the compared state-of-the-art guidance technique, especially due to a lowered visual dominance without a greater impact on the sense of presence or visual obstructiveness.

6 Discussion

The results show that *RippleVision* is an effective alternative to existing gaze guidance techniques, and the initial calibration of the parameters proved to be relevant, as there is a clear difference between the individual settings. As expected, an increase in intensity improved the objective effectiveness, which allowed *RippleVision* to be as effective as the overt technique *DynSWAVE* in both scenes. Notable is the impact of the employed spatial wavelengths, which had a clear effect on *RippleVision's* efficiency, suggesting that intermediate spatial wavelengths provide favorable results. As already considered in the calibration study, the angle parameter should be re-examined. Currently, the angle increases drastically with an increase in distance/angle of the GP to the POI, thereby resulting in full coverage of one side of the screen if the POI is further out-of-view. A smaller angle or an alternative implementation, such as a logarithmic increase, could assist in not only guiding the user in the direction of the POI while it is out-of-view, but also directly toward it. This adjustment could further reduce *RippleVision's* obtrusive properties, especially regarding visual obstruction. Nevertheless, we can still confirm **RQ1** as our results strongly indicate that *RippleVision* guides the gaze as effectively as state-of-the-art techniques.

When comparing the subjective perception of *RippleVision*, it can be established that it is less overt than the employed *DynSWAVE* and *ARROW*, even with settings that noticeably change the brightness. With the least impact on the sense of presence, *RippleVision* also preserved user immersion best. Consequently, we can answer **RQ2** and confirm that *RippleVision* is less obtrusive than state-of-the-art. It should be noted that due to the small sample size of participants, further validation is required to strengthen the evidence addressing both research questions.

The clear differences in the techniques, perceptions, and resulting effectiveness between scenes are likely due to two reasons: the search task and the lighting. In the artificial scene, a multitude of moving objects greatly increased the unassisted complexity of the search task compared to the realistic scene. This increase is presumably the key difference in the raw search times with the assistance of certain techniques. Especially *Deadeye* seems to struggle as a cue in such a 360° scenario and its number of distractors. Although *WarpVision* does not offer unrestricted out-of-view guidance either, its guiding capability does not appear to be affected similarly. *WarpVision's* effectiveness

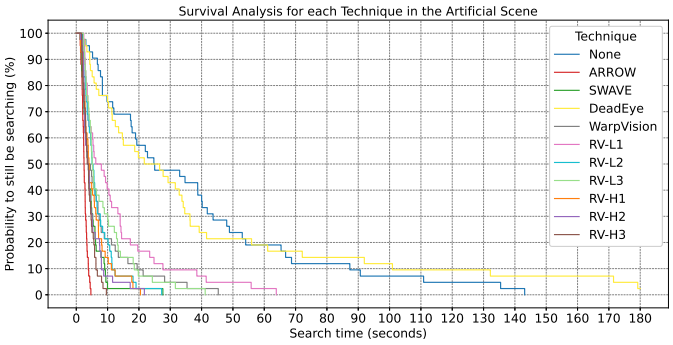

(a) Artificial scene

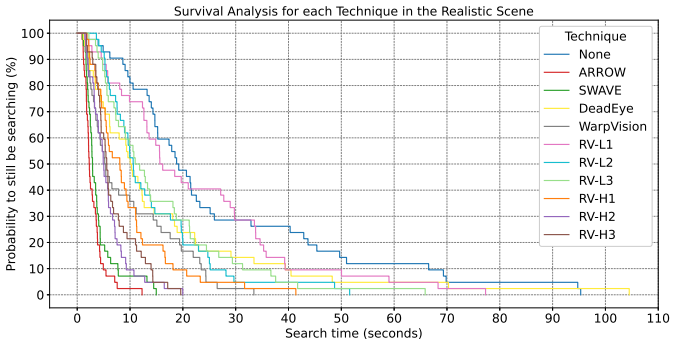

(b) Realistic scene

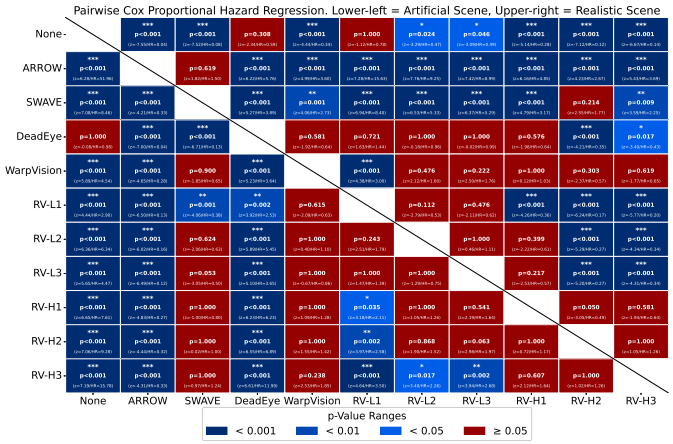

(c) Pairwise comparison of techniques. Each cell shows the z-value, hazard ratio (HR), and Bonferroni-Holm corrected p-value within scenes. The hazard ratio indicates the probability to find the POI faster with the column technique compared to the row technique (a hazard ratio of 0.5 would imply that the column technique is 50% less effective).

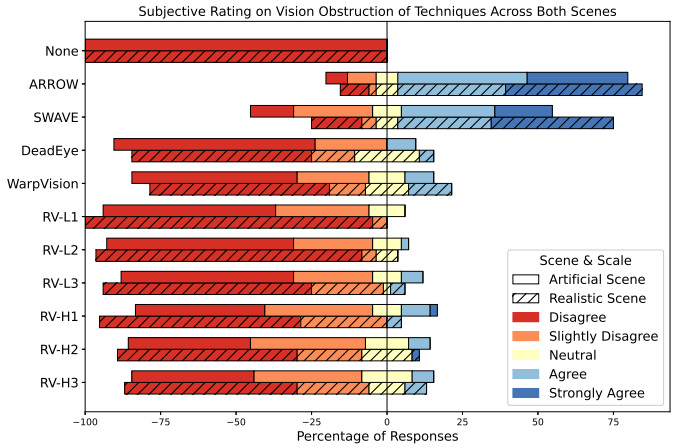

(d) Participant ratings to the statement **S1** "The technique obstructed my vision".

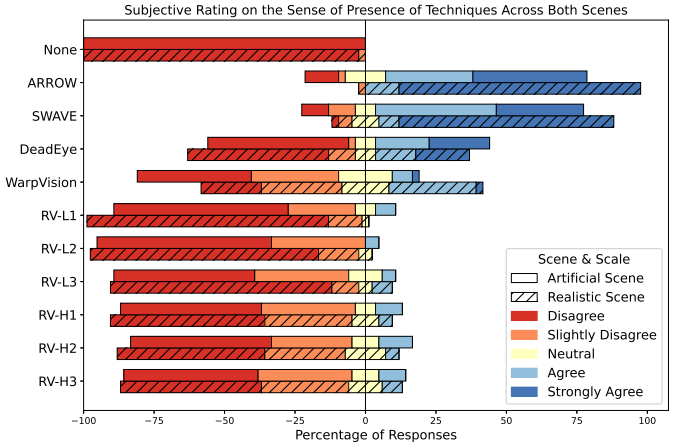

(e) Participant ratings to the statement **S2** "The technique disrupted my sense of presence."

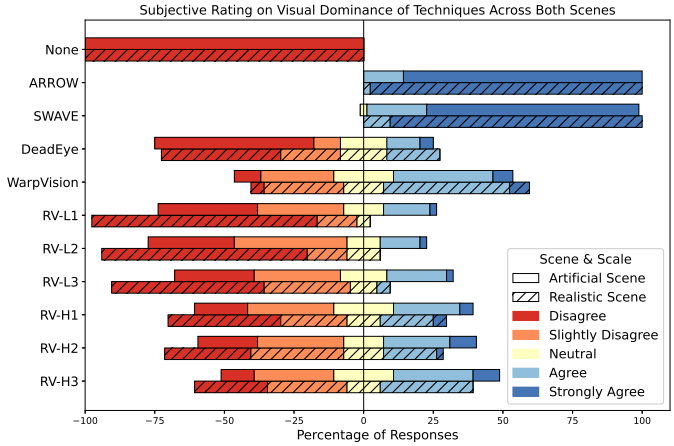

(f) Participant ratings to the statement **S3** "The technique was visually dominant."

Fig. 6: Results of the application study. Survival functions without confidence intervals in the artificial (a) and realistic (b) scene, followed by a pairwise cox proportional hazard regression (c), show the objective results. Likert plots show the subjective results for the statements on visual obstruction (a), sense of presence (b), and visual dominance (c) of techniques which were asked after each search task.

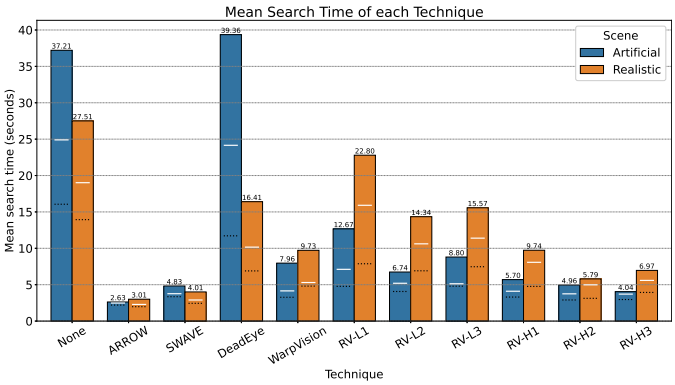

Fig. 7: Mean and median search times of trial runs with employed techniques. The mean is displayed as bar, the median as white line. The black dotted line shows the median first fixation time of the POI.

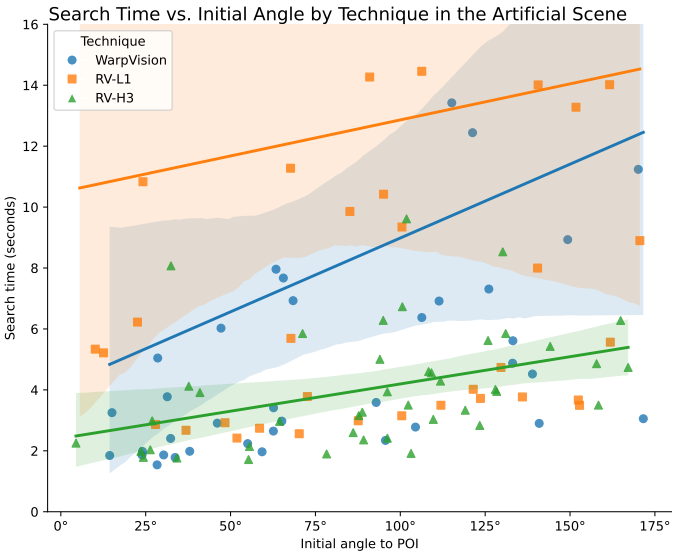

Fig. 8: Scatterplot of total search times of trial runs against the initial angles of the POI to the GP at the start of the task, with a fitted linear regression line and confidence intervals. Shown in the artificial scene for the in-view technique *WarpVision* and the lowest intensity/shortest spatial wavelength and highest intensity/largest spatial wavelength *RippleVision* variants.

depends on the distance to the POI and the number of distractors. At the same time, *Deadeye* clearly improves the search task in an environment with fewer distractors.

Although *RippleVision* offers out-of-view guidance, its effectiveness seems to be lower in the realistic scene. This loss is likely not attributable to the amount of distractors but to the darker lighting and textures used in the cyberpunk-like realistic scene, which reduced the visibility of *RippleVision's* cue. It might, furthermore, be a result of the way the waves adjust brightness in stereoscopic inverse, as scenes that are moderately lit can easily be adjusted to be brighter or darker, while scenes that are very dark are likely to reduce the perceived brightness of darker waves compared to their bright counterpart, thus reducing the effect of stereoscopic disparity. The same mechanisms can be speculated to reduce efficiency in very bright scenes. Moreover, though MSAA and SMAA was enabled, the realistic scene showed aliasing artifacts. Here, some users reported that the cue was interacting with the aliasing artifacts, which could reduce its perceived effectiveness. In addition, since *Ripplevision* was calibrated for the detectability and acceptability threshold in a scene comparable to the artificial search task, it can be concluded that the darker lighting also affected these thresholds and may need additional calibration. The remaining two techniques *ARROW* and *DynSWAVE* were unaffected by notable changes in the scene.

Overall, participants found the perceptual experience of *RippleVision* itself to be difficult to articulate, while the waves themselves could be seen more clearly in higher intensity settings, some participants described the effect of lower settings as almost subconscious yet intuitive cues for guidance. Although no participants spontaneously reported visual fatigue in their post-study comments, further research should evaluate how *RippleVision* impacts this aspect, especially in search tasks that require more time.

7 Limitations and Future Work

Several constraints arise from the HMD and study design. The HMD operated wirelessly with the application running on the experimenter's workstation to allow unrestricted head rotation. Occasional minor connectivity drops in the local network caused brief interruptions in the procedure; these delays were rare and did not measurably affect the collected data.

Although the Vive Focus 3 is specified at a 120° FoV, the display edges exhibit pronounced geometric distortion, likely attenuating techniques that rely on peripheral warping, particularly *WarpVision*, which uses a similar peripheral distortion to attract gaze toward out-of-view POIs. Moreover, the panel resolution constitutes an additional limitation: at low-intensity settings, *RippleVision* could be misinterpreted as screen noise or rendering artifacts. This is a hardware limitation specific to the Vive Focus 3.

The calibration study was conducted only in an artificial scene; consequently, *RippleVision* was not re-tuned for the realistic scene. Moreover, parameters for comparison techniques were adopted from prior work rather than established via a dedicated calibration. As a result, comparisons of subjective obtrusiveness are limited to the efficiency in parameter calibration of previous work.

In the realistic scene, residual aliasing could not be eliminated, which may have led participants to perceive guidance where none was present and to misinterpret subtle *RippleVision* waves.

The search task itself poses a limitation. Even techniques lacking full out-of-view guidance effectively aided participants, likely because they rapidly scanned the visible region before rotating to a new area. In tasks where users are not primed to redirect their view, it is plausible that in-view guidance techniques would perform substantially worse.

It could be shown that the parameters for *RippleVision* were of significant importance and led to clear differences between scenes. Therefore, and due to the limited number of participants in the calibration study, the calibration of the parameters should be further investigated to dynamically adjust brightness based on the background to remove the necessity for individual calibrations for each virtual environment. Furthermore, to grant the user more insights, the waves could be adjusted to change their characteristics, such as their wavelength or speed, based on the distance and depth of the POI.

Future work may also look into alternative cues, as it remains unclear whether the waves themselves were the main directional indicator or whether another type of movement inside the circular slice could improve guidance further. Alternatives might explore wave-like curvature of the space toward the POI to avoid changes in brightness or barely perceivable flickering. Lastly, the combination of subtle out-of-view and in-view techniques should be investigated, like a combination of *RippleVision* as a directional cue and *WarpVision* as a direct cue once the POI is inside the FoV.

8 Conclusion

In this study, we introduce a novel gaze-dependent guidance technique, *RippleVision*, using stereo-inverse brightness modulated waves as a cue. Two studies were conducted, one to calibrate the parameters for *RippleVision* and the other to show its effectiveness over state-of-the-art. Our results demonstrated that *RippleVision* achieves comparable effectiveness with significantly less perceptible disruption. *RippleVision* also offers subtle out-of-view gaze guidance.

REFERENCES

[1] X. P. amd Colin Groth, D. J. Navarro, Z. Zou, Y. Zhu, A. Serrano, K. Myszkowski, Q. Sun, and P. Chakravarthula. Why slow feels fast and fast feels slow: Evaluating and predicting speed misperception. In *IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, 2025. 4

[2] F. Biocca, A. Tang, C. Owen, and F. Xiao. Attention funnel: omnidirectional 3d cursor for mobile augmented reality platforms. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '06, 8 pages, p. 1115–1122. Association for Computing Machinery, New York, NY, USA, 2006. doi: 10.1145/1124772.1124939 1, 2

[3] N. Biswas, A. Singh, and S. Bhattacharya. Augmented 3D arrows for visualizing off-screen Points of Interest without clutter. *Displays*, 79:102502, Sept. 2023. doi: 10.1016/j.displa.2023.102502 1, 2

[4] F. Bork, C. Schnelzer, U. Eck, and N. Navab. Towards Efficient Visual Guidance in Limited Field-of-View Head-Mounted Displays. *IEEE Transactions on Visualization and Computer Graphics*, 24(11):2983–2992, Nov. 2018. doi: 10.1109/TVCG.2018.2868584 1, 2

[5] J. M. Foley, S. Varadharajan, C. C. Koh, and M. C. Farias. Detection of gabor patterns of different sizes, shapes, phases and eccentricities. *Vision Research*, 47(1):85–107, 2007. 4

[6] S. Grogorick, M. Stengel, E. Eisemann, and M. Magnor. Subtle gaze guidance for immersive environments. In *Proceedings of the ACM Symposium on Applied Perception*, SAP '17, pp. 1–7. Association for Computing Machinery, 2017. doi: 10.1145/3119881.3119890 2

[7] S. Grogorick, J.-P. Tauscher, N. Heesen, S. Castillo, and M. Magnor. Stereo inverse brightness modulation for guidance in dynamic panorama videos in virtual reality. *Computer Graphics Forum*, 39, 08 2020. doi: 10.1111/cgf.14091 1, 2

[8] C. Groth, S. Fricke, S. Castillo, and M. Magnor. Wavelet-based fast decoding of 360° videos. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, pp. 1–9, 2023. doi: 10.1109/TVCG.2023.3247080 1

[9] C. Groth, T. Scholz, S. Castillo, J.-P. Tauscher, and M. Magnor. Instant hand redirection in virtual reality through electrical muscle stimulation-triggered eye blinks. In *ACM Symposium on Virtual Reality Software and Technology (VRST)*, number 37, pp. 1–11, 2023. doi: 10.1145/3611659.3615717 1

[10] U. Gruenefeld, A. E. Ali, S. Boll, and W. Heuten. Beyond halo and wedge: visualizing out-of-view objects on head-mounted virtual and augmented reality devices. In *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services*, pp. 1–11. ACM, 2018. doi: 10.1145/3229434.3229438 2

[11] U. Gruenefeld, D. Ennenga, A. E. Ali, W. Heuten, and S. Boll. Eyesee360: designing a visualization technique for out-of-view objects in head-mounted augmented reality. In *Proceedings of the 5th Symposium on Spatial User Interaction*, SUI '17, 10 pages, p. 109–118. Association for Computing Machinery, New York, NY, USA, 2017. doi: 10.1145/3131277.3132175 2

[12] IL.ranch. cyberpunk - tall city (bip). Unity Asset Store, 2023. [Unity asset], https://assetstore.unity.com/packages/3d/environments/sci-fi/cyberpunk-tall-city-bip-263146, Accessed: September 9, 2025. 4

[13] A. Krekhov, S. Cmentowski, A. Waschk, and J. Kruger. Deadeye Visualization Revisited: Investigation of Preattentiveness and Applicability in Virtual Environments . *IEEE Transactions on Visualization & Computer Graphics*, 26(01):547–557, Jan. 2020. doi: 10.1109/TVCG.2019.2934370 2, 3, 5

[14] J. Kudnick, M. Weier, C. Groth, B. Fu, and R. Horst. Warpvision: Using spatial curvature to guide attention in virtual reality. In *IEEE International Symposium on Mixed and Augmented Reality (ISMAR), TVCG*, 2025. 2, 3, 4, 5

[15] D. Lange, T. C. Stratmann, U. Gruenefeld, and S. Boll. HiveFive: Immersion preserving attention guidance in virtual reality. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI '20, pp. 1–13. Association for Computing Machinery, 2020. doi: 10.1145/3313831.3376803 2

[16] Y.-C. Lin, Y.-J. Chang, H.-N. Hu, H.-T. Cheng, C.-W. Huang, and M. Sun. Tell me where to look: Investigating ways for assisting focus in 360° video. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, pp. 2535–2545. Association for Computing Machinery, 2017. doi: 10.1145/3025453.3025757 2

[17] N. Norouzi, G. Bruder, A. Erickson, K. Kim, J. Bailenson, P. Wisniewski, C. Hughes, and G. Welch. Virtual animals as diegetic attention guidance mechanisms in 360-degree experiences. *IEEE Transactions on Visualization and Computer Graphics*, 27(11):4321–4331, 2021. doi: 10.1109/TVCG.2021.3106490 2

[18] E. Peli. In search of a contrast metric: Matching the perceived contrast of gabor patches at different phases and bandwidths. *Vision Research*, 37(23):3217–3224, 1997. 4

[19] K. Quinn and J. L. Gabbard. Augmented Reality Visualization Techniques for Attention Guidance to Out-of-View Objects: A Systematic Review. In *2024 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 826–835, Oct. 2024. ISSN: 2473-0726. doi: 10.1109/ISMAR62088.2024.00098 2, 3

[20] P. Renner and T. Pfeiffer. Attention guiding techniques using peripheral vision and eye tracking for feedback in augmented-reality-based assistance systems. In *2017 IEEE Symposium on 3D User Interfaces (3DUI)*, pp. 186–194, Mar. 2017. doi: 10.1109/3DUI.2017.7893338 1, 2, 5

[21] S. Rothe, D. Buschek, and H. Hußmann. Guidance in cinematic virtual reality-taxonomy, research status and challenges. 3(1):19, 2019. Number: 1 Publisher: Multidisciplinary Digital Publishing Institute. doi: 10.3390/mti3010019 2

[22] G. I. Rozhkova, A. V. Belokopytov, M. A. Gracheva, E. I. Ershov, and P. P. Nikolaev. A simple method for comparing peripheral and central color vision by means of two smartphones. *bioRxiv*, 2021. doi: 10.1101/2021.01.12.426150 1

[23] T. Scholz, C. Groth, S. Castillo, M. Eisemann, and M. Magnor. Measuring velocity perception regarding stimulus eccentricity. In *ACM Symposium on Applied Perception (SAP)*, pp. 1–9. ACM, August 2024. 4

[24] R. Tosa, S. Hattori, Y. Hiroi, Y. Itoh, and T. Hiraki. ChromaGazer: Unobtrusive Visual Modulation using Imperceptible Color Vibration for Visual Guidance. *IEEE Transactions on Visualization and Computer Graphics*, 31(5):3450–3458, May 2025. doi: 10.1109/TVCG.2025.3549173 1, 2

[25] M. Weier, M. Stengel, T. Roth, P. Didyk, E. Eisemann, M. Eisemann, S. Grogorick, A. Hinkenjann, E. Kruijff, M. Magnor, K. Myszkowski, and P. Slusallek. Perception-driven accelerated rendering. *Computer Graphics Forum (Proc. of Eurographics EG)*, 36(2):611–643, Apr 2017. doi: 10.1111/cgf.13150 4

[26] P. Wen, F. Lu, and A. Z. Mohamad Ali. Using attentional guidance methods in virtual reality laboratories reduces students' cognitive load and improves their academic performance. *Virtual Reality*, 28(2):110, May 2024. doi: 10.1007/s10055-024-01012-0 2

[27] J. W. Woodworth and C. W. Borst. Visual cues in VR for guiding attention vs. restoring attention after a short distraction. *Computers & Graphics*, 118:194–209, Feb. 2024. doi: 10.1016/j.cag.2023.12.008 1, 2, 5