Gaze-Contingent Counter-Vection Noise for Cybersickness Reduction

Colin Groth* New York University
Henry Kam† New York University
Piotr Didyk‡ Università della Svizzera italiana
Qi Sun§ New York University

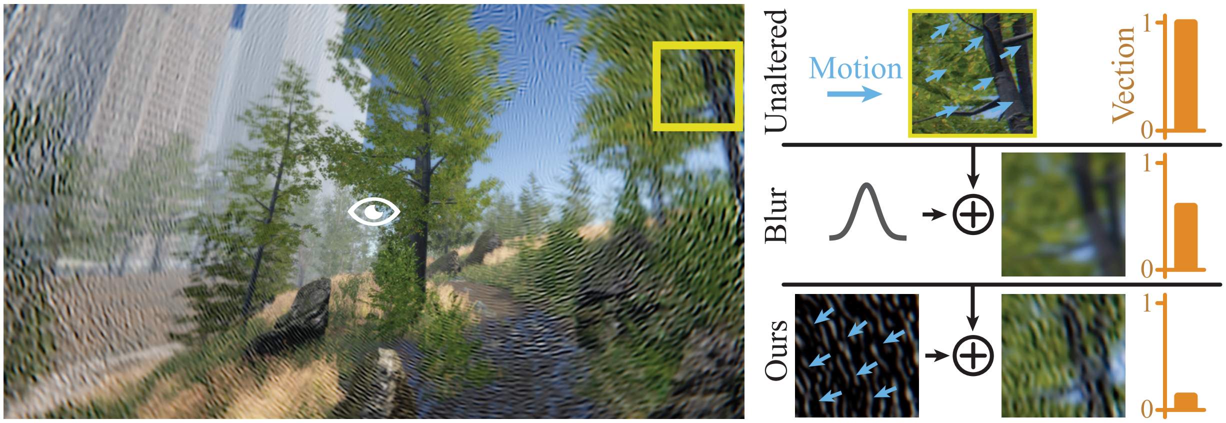

Figure 1: We present a novel method that reduces cybersickness by band-limited Gabor noise. The Gabors are added to the luminance channel of a foveated output (Blur), while the phase of the noise is temporally shifted to induce the impression of motion. By adjusting the orientation and phase shift opposite to the scene motion, we induce the impression of *counter-vection*. Given the natural spatial pooling of visual information in the periphery, the counter-vection signals are integrated with the scene motion, creating a percept with reduced vection. In an experiment, we show that this reduced vection results in significantly less cybersickness.

Abstract

Cybersickness remains a common side effect of many immersive experiences, occurring particularly when fast, virtual-only movements induce high levels of optical flow. These visual movements can induce the perception of self-motion, denoted vection, which leads to a neurological sensory mismatch with the balance system, causing the body to react with sickness. Many techniques have been proposed to reduce cybersickness, aiming either to align the signals of both sensory systems or to suppress the motion perception of one system. One of the simplest, unobtrusive, and popular techniques is foveated blurring, which applies a gaze-contingent low-pass filter to the visual field. However, while the filtering removes high-frequency motion information responsible for high vection, the preserved low-frequency information of the content still carries significant motion cues. Therefore, the effectiveness of the approach varies depending on the visual content. In this paper, we introduce a novel method to enhance foveated blurring by integrating Gabor noise to actively counteract optical flow. Specifically, by leveraging the wave-like characteristics of Gabor functions, we dynamically shift their phase to generate localized motion signals. When these functions are oriented against the scene's optical flow, they induce a "counter-vection" effect. Through spatial pooling, this conflicting motion information cancels out primary sickness-inducing stimuli without degrading the user's overall motion perception. In our real-time implementation, the noise is fine-tuned to induce most counter-vection without being actively perceptible to the average user. In a naturalistic VR experiment, we validate the effectiveness of the approach and show that it can significantly improve the performance of regular foveated blurring.

Index Terms: VR, cybersickness, noise.

1 Introduction

By providing highly realistic and interactive environments, virtual reality (VR) offers experiences far more engaging than traditional two-dimensional displays. However, VR experiences are often accompanied by unwanted physiological reactions, including nausea, headache, and disorientation [31] that compromise comfort and usability [18]. While various theories and factors exist, a major cause of cybersickness is commonly attributed to a sensory conflict: a mismatch between the visual system, which actively perceives motion, and the vestibular system, which remains mostly stationary. This discrepancy is largely driven by vection — the visually induced illusion of self-motion [11].

Current mitigation strategies generally fall into two categories: stimulating vestibular cues via external hardware (e.g., GVS) or reducing the visual cues that trigger vection. Among the latter, foveated blurring is a common rendering-based solution that removes high-frequency peripheral detail [17]. Scenes with less high frequency information are generally perceived to be slower [53, 60]. However, this method faces a fundamental trade-off: aggressive blurring can suppress vection yet simultaneously degrade visual quality [19].

To overcome this barrier, we propose moving beyond simple signal attenuation. We introduce the concept of **counter-vection**: the integration of subtle visual signals that move in the direction opposite to the perceived scene motion. By leveraging the visual system's natural pooling process, whereby the periphery integrates nearby motion signals into a single percept, we can rebalance the perceived motion magnitude without removing scene content.

Specifically, we present a novel, real-time technique that gen-

*e-mail: c.groth@nyu.edu
†e-mail: henrykam@nyu.edu
‡e-mail: piotr.didyk@usi.ch
§e-mail: qisun@nyu.edu

erates these counter-vection signals using dynamically synthesized Gabor noise. Unlike "ghosting" or wave-based methods that are visually obtrusive or spatially rigid, our approach uses non-uniform noise patterns that are perceptually subtle yet locally controllable. By implementing these patterns within a foveated rendering pipeline, we additionally modulate luminance in the frequency bands removed by blurring. This ensures the counter-vection signal is both effective and visually integrated. Our user experiment demonstrates that this method significantly reduces cybersickness while preserving the immersion essential to the VR experience.

2 Related Work

In this section, we describe how vection leads to cybersickness and explore previous research on mitigating cybersickness in VR.

Vection and Cybersickness. Cybersickness in VR is commonly attributed to conflicts between visually induced motion cues and vestibular signals, also known as visual-vestibular conflict [55]. Visual information plays a particularly important role in this sensory conflict in VR, since immersive head-mounted displays (HMDs) isolate users from the real world, while providing compelling stereoscopic motion cues [3, 45]. Such motion cues can be the result of controller-based locomotion in virtual scenes or camera motion in immersive video content. Visually induced sensations of motion, known as vection, have therefore been widely discussed as an important factor of cybersickness [11, 26, 40]. Many works suggest that peripheral vision is influential for perceived vection due to its high sensitivity to temporal changes and global motion [68, 22, 50, 48, 66]. Other studies have shown that perceived vection is also dependent on stimulus properties, including luminance, contrast, spatial frequency, and motion speed [57, 22, 8, 14]. Dizmen et al. further showed that compelling self-motion can be induced without even stimulating the entire visual field [10], however, only when paying active attention to the virtual motion [28]. These findings inspire our work: through a foveated method, we aim to induce subtle motion cues that induce vection against the scene motion (counter-vection).

Visual Mitigations for Cybersickness Reduction. A common strategy for mitigating cybersickness is the attenuation of visual motion cues as a more practical approach than, for example, active vestibular stimulation [20, 24]. Early approaches apply opaque occlusions, either to the outer visual field [5, 37, 61] or dynamically via gaze-contingent masking [1, 19, 18]. While effective at reducing vection, these occlusions severely limit peripheral awareness.

To provide a less intrusive experience, other methods have experimented with peripheral blurring and gaze-contingent filtering [52, 25, 38, 46, 47, 17, 27]. By applying Gaussian filters to the periphery, these techniques reduce contrast and eliminate high spatial frequencies, lowering the visibility of motion-inducing structures. Blurring effectively removes high-frequency spatial details while leaving low-frequency signals intact. However, the periphery continues to derive vection from low-frequency motion signals [50]. Therefore, foveated blurring often yields varying success in suppressing cybersickness.

Counter-vection and Reverse Optic Flow. Alternative mitigation strategies actively disrupt optic flow by introducing opposing visual cues, which we denote as counter-vection. Early foundational work explored augmenting the periphery outside the display viewport with artificial motion patterns to counteract visually induced motion [76]. More recent approaches have integrated these counter-vection patterns into the display, using the scene's optical flow or camera motion [51].

For example, Park et al. overlaid reverse optic flow cues using distinct directional arrows [51]. While their method reduced cybersickness, the highly noticeable arrows extended into peripheral and foveal vision, causing severe visual distraction. There have been attempts at more subtle counter-vection patterns, such as adding them to more constrained areas [10, 33, 32]. However, these approaches either still rely on introducing overtly visible geometric patterns that disrupt the user experience or have not successfully reduced cybersickness.

Content-based Motion Manipulation. The physiological limitations of the peripheral human visual system (HVS) have long been exploited in graphics, most prominently through foveated rendering [12, 21, 30, 34, 42, 52, 63, 77]. Some approaches exploit reduced peripheral sensitivity to alter the visual field in ways that remain imperceptible. One such example is the use of ventral metamers [13], physically distinct images that become perceptually indistinguishable when they produce the same pooled peripheral statistics [16]. This arises from the fact that peripheral vision represents input in terms of local summary statistics rather than preserving precise spatial layout [56, 58].

This framework also helps to explain why noise can act as an effective metameric substitute. Prior texture-synthesis work showed that images generated by starting from white noise and matching multiscale filter-response statistics can generate new images with similar texture appearance [23]. This logic can also extend to foveated rendering, as spatial frequencies beyond the peripheral resolution limit may still be detectable even when they are no longer individually resolvable [67, 66]. In such works, metameric rendering has been used in a variety of applications, ranging from optimizing subjective image quality [74], imperceptibly embedding data [54], and enhancing motion perception [64]. For example, Tariq et al. [65] demonstrated that procedurally generated noise can significantly enhance the spatial perception of a blurred periphery. Crucially, these works demonstrate how carefully controlled noise can be blended into rendered content without compromising viewers' perception of the original scene.

However, this same reasoning extends to motion. Recent work shows that removing high-frequency peripheral detail can impair motion judgements; adding controlled spatiotemporal energy can restore motion cues without objectionable quality loss [64], Thus, we expand on prior work and blend metameric noise patterns with the rendered virtual content. In contrast, our primary objective is to use this noise to subtly suppress motion cues and thereby reduce cybersickness.

3 Counter-Vection Noise

The primary goal of the proposed method is to mitigate cybersickness by modulating the user's perception of vection by subtle, foveated Gabor noise. By introducing motion energy that opposes the optic flow of the virtual scene (counter-vection), we aim to balance the visual-vestibular conflict arising from visual-only motion cues (i.e., visual motion without corresponding physical movements). Our method's counter-vection is superimposed on foveated frames, which can be the result of either a computation-focused foveated rendering implementation or a foveated blurring post-process shader application. In both cases, the low-pass filter operation of the foveation naturally removes high-frequency information. Our counter-vection cues are introduced in the *aliasing band*, a spatial frequency band above the foveated region where content is detectable but not resolvable [66]. Thus, acting on the foveated frames resolves interference problems of the vection cues from the rendered content and our synthesized signals. The primary challenge in implementing such a system is ensuring minimal visibility while maintaining sufficient motion energy to influence the vestibular system. Furthermore, the noise must respond adaptively to the motion in different regions of the visual field.

3.1 Noise Pattern

The counter-vection signal must introduce motion energy without producing visible structure to preserve the VR experience. Therefore, the noise pattern requires several properties. First, the signal must be spatiotemporally broadband so that it does not form stable edges or shapes while still inducing sufficient motion energy in the HVS. Second, the contrast must remain low and spatial correlations must remain local to avoid conscious detection. Third, the motion direction and spatial frequency must be controllable so that the induced motion can be aligned with or opposed to the optic flow of the scene.

We evaluated several procedural noise formulations. *Dynamic Perlin noise* was initially considered due to its smooth spatial structure. However, Perlin noise provides limited control over the spectral characteristics of the signal. Spatial frequency, orientation, and temporal evolution are coupled, which makes it difficult to adjust motion energy independently of the spatial appearance.

Instead, we employ *Dynamic Gabor Noise*. Gabor functions form localized band-pass filters with explicit control over spatial frequency and orientation. A dynamic noise field is generated by sampling multiple Gabors with randomized spatial positions and phases while maintaining controlled frequency and orientation distributions. Temporal evolution is introduced through phase drift over time, which determines the perceived motion direction and magnitude. This formulation allows the spatial frequency to be constrained to the aliasing band above the foveated region and provides direct control over the dominant motion direction. By aligning the phase drift opposite to the optic flow of the virtual scene, the noise induces counter-vection motion energy. The random spatial placement of the Gabors produces a perceptually texture-like signal without coherent structure, which keeps the pattern visually subtle while still generating sufficient motion energy in the periphery to influence vection.

3.2 Initialization

We generate the noise using sparse Gabor convolution [4]. This method convolves Gabor kernels with sparse random impulses, effectively placing localized Gabor patches at random image locations. To obtain a uniform yet stochastic distribution of motion primitives, we divide the image into a grid of cells. Each cell corresponds to approximately 1° of visual angle (35 pixels). For simplicity, we assume a uniform pixel distribution across the field of view (FOV) for the grid definition. Within each cell, a fixed number of impulse locations are sampled using a stable random seed. The number of impulses controls the trade-off between the quality of the noise and the computational overhead. We empirically selected eight impulses per cell to provide the best trade-off. The impulse locations remain constant over time to avoid temporal instabilities, while the associated Gabor kernel parameters vary smoothly according to the properties of the underlying content.

Each Gabor kernel defines the local characteristics of the noise. A Gabor function consists of a sinusoidal carrier modulated by a Gaussian envelope. The kernel is defined by four main parameters: amplitude A, spatial frequency f, phase ϕ, and orientation θ. The Gabor kernel G at a local coordinate $\mathbf{p}$ is expressed as

$$G(\mathbf{p}) = A \exp\left(-\pi\sigma_G^2\|\mathbf{p}\|^2\right)\cos(2\pi f(\cos\theta,\sin\theta)\cdot\mathbf{p}+\phi),$$

where $\mathbf{p}=(x,y)$ denotes the local coordinate relative to the impulse location, and σ_G defines the spatial spread factor which controls the width of the Gaussian envelope. The final noise value N at pixel $\mathbf{p}$ is computed as a weighted sum of nearby Gabor contributions

$$N(\mathbf{p}) = \sum_{i\in I} w_i\, G(\mathbf{p}-\mathbf{p}_i),$$

for a set of impulses I. Here, $\mathbf{p}_i$ denotes the impulse location of the i-th Gabor and w_i is a random weight sampled from $[-1,1]$. The symmetric weight distribution enforces a zero-centered mean, preventing changes in the average brightness of the underlying image when the noise is added.

For the accumulation of contributions, we evaluate Gabors from a 3×3 neighborhood of cells around each pixel. This ensures C^0 continuity and prevents visible grid artifacts. Each Gabor kernel is dynamically controlled through the parameters A, f, ϕ, and θ. These parameters determine the visibility and temporal behavior of the generated counter-vection signal. The following sections describe how each parameter is defined.

3.3 Orientation

The orientation θ of each Gabor kernel determines the direction of the sinusoidal carrier and therefore the direction of the perceived wave propagation when the phase changes over time. To generate a counter-vection signal, the carrier orientation is aligned with the local motion direction, but defined such that a positive phase shift ($\Delta\phi > 0$) results in a propagation direction diametrically opposed to the local motion.

Let $\mathbf{m}(\mathbf{p})$ denote the 2D motion vector at pixel $\mathbf{p}$, obtained from the motion field of the rendered scene. The orientation is defined by the direction of this motion vector

$$\theta_{\text{flow}}(\mathbf{p}) = \operatorname{atan2}\big(\mathbf{m}_y(\mathbf{p}),\mathbf{m}_x(\mathbf{p})\big),$$

where $\operatorname{atan2}(\cdot)$ returns the polar angle of the motion vector.

In regions where the motion magnitude is very small, the motion direction becomes unstable and can introduce temporal fluctuations in the noise pattern. To avoid such instabilities, the orientation is smoothly interpolated between a stable reference θ_{ref} and the motion-derived θ_{flow} based on a confidence value $\tau\in[0,1]$. We derive τ by mapping the motion magnitude $\|\mathbf{m}(\mathbf{p})\|$ through a smooth step function between $\tau_{\min}$ and $\tau_{\max}$; the mapping ensures that τ remains at zero in near-static regions and transitions to one only when the motion magnitude is sufficient to provide a reliable directional signal. In other words, the orientation follows the optic flow in regions with reliable motion while remaining temporally stable in nearly static areas. The final orientation is defined as

$$\theta(\mathbf{p}) = (1-\tau)\,\theta_{\text{ref}} + \tau\,\theta_{\text{flow}}(\mathbf{p}).$$

When combined with the temporal phase evolution described next, the resulting wave propagation occurs opposite to the scene motion and therefore produces the intended counter-vection signal.

3.4 Phase

Shifting the Gabor phase ϕ over time controls the temporal evolution of the sinusoidal carrier and therefore determines the perceived motion of the Gabor pattern. Psychophysical evidence, such as the double-drift illusion, demonstrates that such local phase shifts can induce the perception of global motion [9, 69].

Let ϕ_t denote the phase of a kernel at time step t. The phase is updated according to the magnitude of the local motion vector $\mathbf{m}(\mathbf{p})$ and the spatial frequency f of the kernel

$$\phi_{t+1} = \operatorname{mod}(\phi_t + \Delta\phi,\, 2\pi),$$

where the phase increment $\Delta\phi$ is defined as

$$\Delta\phi = 2\pi f\,\|\mathbf{m}(\mathbf{p})\|.$$

The motion vector magnitude determines the local temporal displacement of the carrier wave. The resulting phase drift propagates opposite to the optic flow, because the carrier orientation is aligned with the motion direction (see Section 3.3)

Large phase increments that do not satisfy $\Delta\phi < \frac{\pi}{2}$ can introduce temporal aliasing artifacts because the motion direction of the phase shift becomes visually ambiguous. To avoid phase wrapping artifacts, the phase shift is constrained such that the instantaneous displacement remains below half a period of the sinusoidal carrier.

3.5 Spatial Frequency

The spatial frequency f determines the wavelength of the sinusoidal carrier. Lower spatial frequencies produce broader patterns with large spatial support, whereas higher spatial frequencies generate fine structures. Consequently, spatial frequency strongly influences the detectability of the signal by the HVS.

Our method targets signals that remain detectable but not resolvable by the HVS. Such signals generate motion energy that can influence perception without introducing continuously perceived structures [66]. The range of spatial frequencies that satisfies this condition is commonly referred to as the *aliasing band* [65, 66].

Thibos et al. studied the characterization of this band as a function of a lower and an upper spatial frequency limit at different retinal eccentricities [66]. Based on the reported measurements, we approximate the lower boundary of the aliasing band using a power-law fit

$$f(e) = 193.79\,(e+2.011)^{-1.105},$$

where e denotes the eccentricity in degrees of visual angle.

We intentionally select the lower bound of the aliasing band. This conservative design choice ensures that the signal remains effective across users. Spatial frequencies near the upper bound may have limited perceptual impact because the patterns can go undetected by insensitive observers.

Two additional constraints are defined to further restrict the frequency range. First, sufficiently low spatial frequencies would produce structures larger than the spatial support of the Gaussian envelope. To ensure that each Gabor kernel remains properly localized, we limit the minimum frequency such that the wavelength does not exceed twice the cell size (equivalent to 0.5 cycles per degree (CPD)). Second, we ensure that high spatial frequencies do not exceed the limit defined by the Nyquist sampling theorem of 0.5 cycles per pixel (CPP).

3.6 Amplitude

The amplitude A controls the strength of the luminance modulation of each Gabor kernel and therefore determines the deviation from the local mean luminance of the underlying image. Higher amplitudes increase the contrast of the generated signal and increase detectability, while lower amplitudes reduce visibility. Perceptual studies show that the sensitivity to luminance contrast depends strongly on retinal eccentricity and temporal change [41, 2, 70]. Consequently, we seek to determine the relationship of the luminance contrast of the synthesized noise as a function of scene speed and eccentricity. The optimization criterion is driven by our overarching goal to identify luminance settings that maximize counter-vection while minimizing noise visibility.

3.6.1 Modulation by Speed

In this paper, we determine the relationships among luminance, visibility, and vection from the psychophysical datasets reported by Guo et al. [22]. Through multiple psychophysical experiments, they independently measured the effects of luminance, as defined by Michelson contrast, on perceived visibility and vection magnitude. We combine those results into a unified model that provides a more global understanding of how speed influences the noise amplitude.

Step 1: Modeling Luminance as a Function of Vection and Speed. The first dataset of Guo et al. reports perceived vection magnitude for different luminance contrasts of a moving stimulus [22]. The measurements were collected at two motion speeds of 20°/s and 60°/s. Given this data, we model luminance contrast L as a function of vection magnitude v and stimulus speed s. The relationship follows an exponential form

$$L(v,s) = \alpha \exp\big(a(s)\,v\big), \tag{1}$$

where α is a shared intercept parameter and $a(s)$ controls the growth rate of the exponential. The fit is achieved using least squares for the two discrete speeds 20°/s and 60°/s, resulting in

$$\alpha = 0.05378, \qquad a(20) = 4.92, \qquad a(60) = 3.437,$$

with a high coefficient of determination.

To evaluate the function for arbitrary speeds s, we require a continuous formulation of $a(s)$. A linear interpolation would introduce a negative slope that eventually produces negative values of $a(s)$, which would invalidate the exponential form of $L(v,s)$. We therefore adopt a hyperbolic formulation that preserves strictly positive values of $a(s)$. For any positive speed s, we define

$$a(s) = \gamma + \frac{\beta}{s},$$

where the constants β and γ are determined from the fitted values at 20°/s and 60°/s:

$$\beta = 44.4952, \qquad \gamma = 2.69546.$$

Step 2: Modeling Visibility as a Function of Luminance and Speed. The second dataset reports the visibility of stimuli at different luminance contrasts [22]. The measurements were again collected for motion speeds of 20°/s and 60°/s. We model visibility as a logarithmic function of luminance contrast. This choice follows the Weber–Fechner law, which states that perceptual responses to physical stimulus intensity often follow a logarithmic relationship. We therefore define visibility $\mathcal{V}$ as

$$\mathcal{V}(L,s) = b(s)\,\log\big(1 + c(s)L\big), \tag{2}$$

where $b(s)$ and $c(s)$ denote speed-dependent parameters.

The model incorporates the constraint $(L,\mathcal{V}) = (0,0)$, since zero luminance modulation implies zero visibility.

We first fit the function independently to the two measured speeds, yielding discrete parameter estimates for $b(s)$ and $c(s)$. To obtain a closed-form model for arbitrary speeds, we interpolate both parameters linearly between the two data points. The resulting functions are

$$b(s) = -0.002135\,s + 0.7298, \qquad c(s) = 0.018589\,s + 2.0254.$$

Step 3: Optimization. The perceptual models from Equation (1) and Equation (2) link luminance contrast to two opposing effects: the induced amount of vection and the visibility of the signal. Our objective is to determine a luminance level that maximizes the counter-vection effect while minimizing the visibility of the added noise.

To express vection as a function of luminance and speed, we permute the model from Equation (1). This yields

$$v(L,s) = \frac{\log(L/\alpha)}{a(s)}. \tag{3}$$

Using this formulation together with the visibility model from Equation (2), we define the utility function

$$U(L,s) = v(L,s) - \lambda\,\mathcal{V}(L,s),$$

where $\lambda > 0$ controls the relative penalty assigned to visibility. Substituting the two models results in

$$U(L,s) = \frac{\log(L/\alpha)}{a(s)} - \lambda\, b(s)\log\big(1 + c(s)L\big).$$

The optimal luminance $L^*(s)$ is obtained by determining the stationary point with respect to L

$$\frac{\partial U}{\partial L} = \frac{1}{a(s)L} - \lambda\,\frac{b(s)c(s)}{1 + c(s)L} = 0,$$

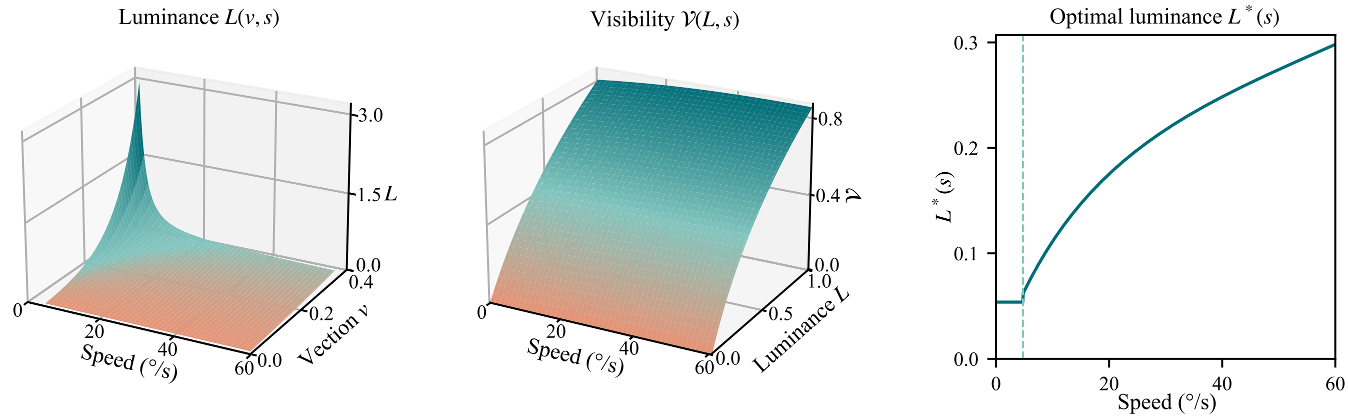

Figure 2: Mathematical models describing luminance L as a function of motion-related parameters. The **left plot** shows the model obtained in Step 1 (in Section 3.6), which expresses luminance as a function of vection strength and motion speed. The **middle plot** shows the visibility function derived in Step 2, which relates luminance to motion speed. In the final step, both models were combined in an optimization procedure to determine the optimal luminance for a given speed; the resulting function is visualized in the **right plot**.

which yields the closed-form solution

$$L^*(s) = \frac{1}{c(s)\big(\lambda b(s)a(s) - 1\big)}.$$

In practice, the formulation must satisfy the constraint $L \geq \alpha$. If $L < \alpha$, the logarithm in $v(L,s)$ becomes negative (cf. Equation (3)), which leads to undefined vection values in the model. Therefore, the luminance function is clipped at a lower transition speed s_L. Due to the nonlinearity of $L^*(s)$ and the constraint $L \geq \alpha$, the optimal luminance function is expressed by a piecewise definition

$$L^*(s) = \begin{cases} \alpha, & s \leq s_L, \\ \dfrac{1}{c(s)\big(\lambda b(s)a(s) - 1\big)}, & s > s_L, \end{cases}$$

where $s_L = 4.76$ and $\alpha = 0.05378$. In our application, we set $\lambda = 1$ to obtain equal weight between visibility and vection.

3.6.2 Modulation by Eccentricity

Perceptual sensitivity depends strongly on retinal eccentricity [41, 2, 70, 17]. Signals presented near the fovea are easily detectable and can therefore become distracting. Signals in the visual periphery are less visible due to the falloff of contrast sensitivity, as described by the contrast sensitivity function [41, 2]. At the same time, peripheral motion contributes strongly to the perception of vection [68, 22]. These properties motivate an eccentricity-dependent modulation of the noise amplitude.

We therefore scale the amplitude using a factor that increases with retinal eccentricity e. This modulation factor E is defined by a sigmoid function

$$E(e) = \frac{1}{1 + \exp\big(-n(e - \delta)\big)},$$

where n controls the slope of the transition and δ defines the midpoint of the function. This formulation produces a smooth increase in signal strength from the central visual field towards the periphery.

Previous cybersickness mitigation techniques often employ a linear increase after a predefined eccentricity threshold (here defined by e_f) to preserve the foveal region. A linear formulation introduces a discontinuity in the derivative at the transition point and does not distribute the signal optimally across the visual field. Instead, the sigmoid formulation produces a smooth transition and allocates more motion energy to peripheral regions, where the influence on counter-vection is stronger [68, 22].

We set $n = 0.38$ and $\delta = 18°$. These parameters satisfy two practical constraints commonly used in perceptual rendering approaches [18]: first, the central foveal region of approximately 6° remains largely unaffected (see fovea definition in Section 4.3). Second, the function reaches its maximum at half the display's FOV, corresponding to the maximum eccentricity for a gaze at the center of the screen. Users typically focus most of their visual attention on the middle region, leading to this assumption. For numerical stability, the exponent is clamped for eccentricities above 60°.

We compared the sigmoid formulation to a linear amplitude increase in an informal study with six participants. The linear baseline was adjusted to produce the same integrated motion energy across the visual field. All participants reported that the sigmoid-based modulation produced noticeably lower visibility while preserving the intended perceptual effect.

The final amplitude A for the Gabor kernels is modulated by the product of the speed-dependent luminance model of Section 3.6.1 and the eccentricity modulation factor:

$$E(e) \cdot L^*(s).$$

4 Real-Time Application

We implement the method as a post-processing system in *Unity* using the *Universal Render Pipeline (URP)*. The implementation relies on custom render features and shaders. The approach operates entirely on the foveated render target and therefore does not require an additional full-resolution rendering pass.

4.1 Pipeline Overview

Figure 3 provides a high-level overview of the rendering pipeline. The pipeline consists of the following stages, which are executed in sequence:

1. Projection of the gaze vectors of both eyes into screen space. These coordinates define the center of the gaze-contingent foveation.

2. Computation of a full-resolution motion vector texture from the depth buffer and the camera view-projection matrices. The formulation removes head motion as a factor.

3. Construction of a motion strength MipMap. The highest level stores the magnitude of the motion vector for each pixel.

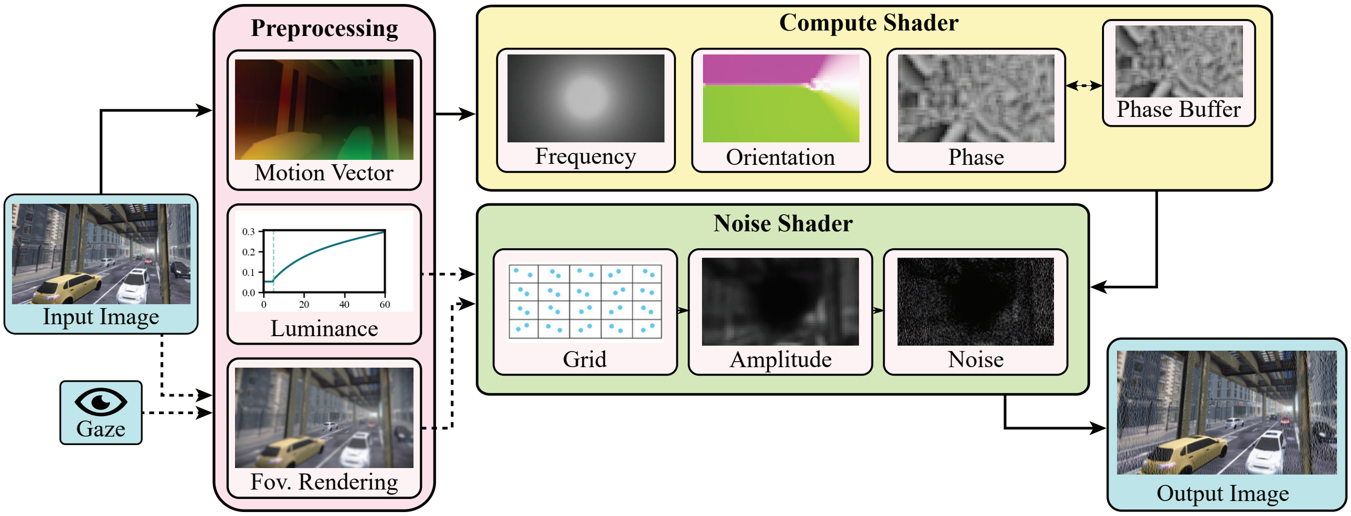

Figure 3: Illustration of our rendering pipeline of the real-time application. Solid arrows show the process flow. Dotted arrows indicate inputs. The Phase Buffer holds information about the last iteration.

Lower-resolution levels are generated through bilinear filtering.

4. Computation of a luminance image from the rendered color buffer.
5. Application of gaze-contingent foveated blurring using a separable Gaussian filter.
6. Computation of Gabor parameters (frequency, orientation, phase) in a compute shader. The results are stored in a structured buffer.
7. Generation of Gabor noise in a fragment shader. The noise is added to the luminance channel of the foveated image to produce the final output.

4.2 Virtual-Only Camera Motion Vectors

Unity provides a view matrix for each eye. The provided matrix corresponds to the final view transform

$$\mathbf{V}_{\text{full}} = \mathbf{V}_{\text{eye}} \cdot \mathbf{V}_{\text{HMD}} \cdot \mathbf{V}_{\text{rig}},$$

where $\mathbf{V}_{\text{eye}}$ represents the eye offset from the HMD center, $\mathbf{V}_{\text{HMD}}$ represents the tracked pose of the head-mounted display relative to the rig, and $\mathbf{V}_{\text{rig}}$ represents the event-driven virtual camera motion in the scene.

Standard motion vector computation uses the current and previous view matrices $\mathbf{V}$ and projection matrices $\mathbf{P}$. For a world position $\mathbf{x}$, the clip-space positions $\mathbf{p}$ are derived by

$$\mathbf{p}_t = \mathbf{P}_t \mathbf{V}_{\text{full},t} \mathbf{x}, \quad \mathbf{p}_{t-1} = \mathbf{P}_{t-1} \mathbf{V}_{\text{full},t-1} \mathbf{x},$$

and the screen-space motion vector $\mathbf{m}$ is

$$\mathbf{m} = \Pi(\mathbf{p}_t) - \Pi(\mathbf{p}_{t-1}),$$

where $\Pi(\cdot)$ denotes perspective division and mapping to normalized screen coordinates.

Using $\mathbf{V}_{\text{full}}$ directly results in motion vectors that contain both virtual scene motion and user head motion. Head motion visually follows corresponding vestibular cues and therefore does not cause cybersickness. Instead, counteracting these motions would introduce an artificial visual–vestibular mismatch.

To isolate motion originating from the virtual scene, we construct a modified previous-frame view matrix that assumes zero HMD movement. First, we isolate the combined eye and HMD transformation

$$\mathbf{V}_{\text{HMD/eye}} = \mathbf{V}_{\text{full}} \cdot \mathbf{V}_{\text{rig}}^{-1} = \mathbf{V}_{\text{eye}} \cdot \mathbf{V}_{\text{HMD}}.$$

The corrected previous view matrix becomes

$$\mathbf{V}_{\text{full-prev}} = \mathbf{V}_{\text{HMD/eye}} \cdot \mathbf{V}_{\text{rig-prev}}.$$

This formulation replaces the previous HMD pose with the current pose while preserving the previous virtual camera motion. Motion vectors computed with this matrix represent only virtual scene motion.

4.3 Foveated Blurring

Foveated blurring is implemented as a separable Gaussian filter executed in two passes. The first pass performs horizontal convolution, and the second pass performs vertical convolution.

The blur radius depends on the pixel eccentricity relative to the gaze position. The standard deviation of the Gaussian kernel follows a linear eccentricity model used in prior work [17, 19, 18]:

$$\sigma = \max(e - e_f, 0) \cdot m.$$

The parameter e denotes the visual eccentricity of the pixel. The slope parameter is set to $m = 0.45$ to have a comparable blur as in prior work [17]. The blur formulation leaves a central region of size $2e_f$ unaffected. The human fovea spans approximately 5° of visual angle [75]. Visual sensitivity to blur is highest in this region, while the influence on cybersickness remains minimal [17]. We introduce a 1° buffer and set $e_f = 6^\circ$.

The Gaussian support is truncated at 3σ, and the kernel size is limited to 61 samples to maintain real-time performance. Blur strength scales approximately linearly with the average optical flow magnitude [19, 18]. We follow this design and scale the blur by motion strength sampled from the highest level of the motion strength MipMap. Movements above 30°/s are clamped. The blurred image forms the basis for the subsequent noise overlay.

4.4 Noise Parameter Computation

The parameters of the Gabor noise signal are computed in a low-dimensional compute shader. Each thread corresponds to a logical block of the image. The block size equals the initialization grid used for the sparse Gabor impulse generation (Section 3.2), which corresponds to 35×35 pixels in our implementation.

This design reduces the number of compute threads by approximately a factor of 1200 compared to per-pixel computation and minimizes texture fetches.

The parameters follow the formulation described in Section 3. Orientation estimation uses the motion vectors within each block. The orientation of the former frame defines the stable reference θ_{ref}. We set

$$\tau_{\min} = 0.05 \text{ pixels/frame}, \quad \tau_{\max} = 1 \text{ pixel/frame}.$$

Motion vectors below $\tau_{\min}$ are treated as random noise and ignored. Vectors above $\tau_{\max}$ are treated as fully reliable. The upper bound reflects the fact that pixel displacements of one pixel or larger provide stable orientation estimates. The lower threshold was determined empirically for the scenes used in our experiments.

The frequency parameter is restricted to remain below the Nyquist limit

$$f \leq \frac{1}{2\Delta x},$$

where Δx denotes the pixel spacing in screen space.

4.5 Noise Generation

Noise generation occurs in the fragment shader. For each pixel with screen coordinate $\mathbf{p}$, the position within the logical cell is normalized to cell coordinates.

The amplitude parameter, on the other hand, is computed per pixel instead of per cell. This formulation allows spatial variation below the cell size and enables sensitivity to high-frequency contrast variations. The amplitude depends on the luminance value of the pixel. The luminance image $L(\mathbf{p})$ is computed using the standard RGB-to-luminance conversion:

$$L(\mathbf{p}) = 0.299\,R(\mathbf{p}) + 0.587\,G(\mathbf{p}) + 0.114\,B(\mathbf{p}),$$

where $R(\mathbf{p})$, $G(\mathbf{p})$, and $B(\mathbf{p})$ denote the color channels at pixel position $\mathbf{p}$. The luminance serves as an approximation of the local signal power. The Gabor envelope uses a spread parameter $\sigma_G = 0.05$, which corresponds to a 3σ Gaussian envelope [64].

Before adding the noise, the image is converted from RGB to YUV color space. The noise is added only to the Y luminance channel. This procedure reduces color artifacts and increases perceptual subtlety while maintaining effectiveness.

Scene motion magnitude also influences the amplitude. Higher motion speeds increase cybersickness risk and therefore require stronger mitigation. The phase shift already scales with motion speed through the formulation described in Section 3.4. In addition, the amplitude increases with motion magnitude. The relationship between Gabor intensity and perceived vection follows a logarithmic model [22]. To maintain a proportional relationship between scene motion speed and counter-vection strength, we apply a logarithmic scaling to the amplitude as described in Section 3.6.

5 Validation Experiment

We validated the effectiveness of our application in a within-subject experiment using a realistic VR scenario. In three sessions per participant, we tested the effect of our method against regular foveated blurring and a control condition. A recovery time of at least 24 hours was maintained between sessions to avoid carry-over effects. Conditions were fully counterbalanced in order. The experiment was approved by the corresponding IRB.

Figure 4: The virtual environment of our experiment displayed two distinct regions: a forest region with natural vegetation (left) and a city region with buildings and cars (right).

5.1 Experimental Design

5.1.1 Stimuli

Conditions. We investigated three different conditions, all presented to each participant (one per session): (1) a control condition that showed the full-resolution rendering output without any post-processing modifications. (2) our method, which applied subtle, gaze-contingent Gabor noise to the rendered frames (see Section 4). (3) foveated blurring, which removed high-frequency details in the periphery by a separable Gaussian blur. The same blur is applied for foveated blurring and as a basis for our method, following the implementation of Section 4.

Virtual Environment. In all sessions, we presented a photorealistic virtual environment consisting of two distinct regions: urban and nature. The urban region primarily contained geometric structures and straight lines, whereas the nature region featured a more heterogeneous shape composition. Both regions were merged seamlessly to create a smooth transition. The camera path in the experiment was predefined and cannot be altered by the participants. It traversed both regions twice in a fixed order. We carefully designed the camera path to feature a balanced mixture of movements. For linear movements, the camera constantly accelerated and decelerated sinusoidally with speeds varying between a minimum of 1 m/s and a maximum of 7 m/s. Such non-linear movements are generally known to induce cybersickness [49, 20]. Furthermore, we included sections with unpredictable movements containing both linear and angular components. Full 360° rotations were performed around each of the three principal axes at 30°/s.

5.1.2 Apparatus

The experiment used a Varjo Aero HMD with a visible FOV of approximately 115° and a refresh rate of 90 Hz. The frames were rendered on the glasses with a resolution of 3840×2160 pixels, corresponding to a peak angular resolution of about 35 pixels per degree (PPD). The built-in device-level foveated rendering of the headset was disabled to ensure full control over our experimental conditions. Rendering was performed on a commodity computer with an NVIDIA GeForce RTX 5080 graphics card at about 70 fps.

5.1.3 Participants

We conducted an a priori power analysis for a within-participant design with three conditions. The analysis assumed a medium effect size ($f = 0.25$), a significance level of $\alpha = 0.05$, and a statistical power of 0.80. We ran the analysis under the conservative assumption of no prior knowledge (0.5 correlation) between repeated measures. The analysis indicated a required sample size of *28 participants*. Accordingly, we recruited 33 participants for the experiment. A total of 29 participants completed all three sessions (7 females, 22 males; age range = 20–31; mean age = 24.24; SD = 1.94). One participant was excluded from the analysis due to corrupted data. Participants were compensated with $30.

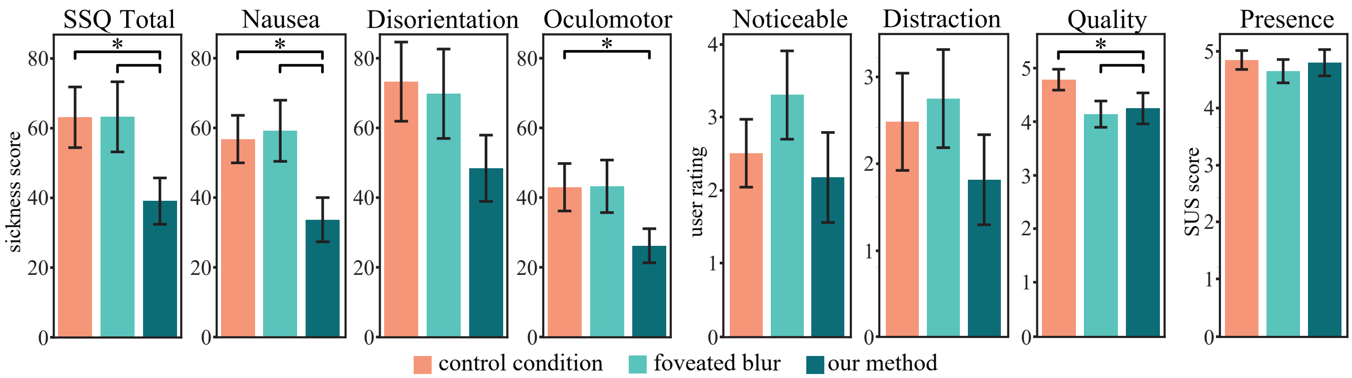

Figure 5: Results of the validation experiment (N=28) comparing the control condition, foveated blur, and counter-vection noise (ours). The SSQ scores are shown with the total score and the three subscales, nausea, disorientation, and oculomotor, where lower values indicate reduced cybersickness. Subjective ratings for noticeability and distraction of the distortions (only when any were noticed; Q2), and quality were rated on a 7-point rating scale. Presence is measured with the SUS questionnaire. Error bars denote the standard error of the mean. Asterisks (*) indicate statistically significant differences.

5.1.4 Measurement

Our primary tool for quantifying cybersickness was the SSQ [31]. Following common procedure, participants completed the SSQ twice, before and after each session of the experiment, to counteract variations in their daily condition. Total sickness scores and subscores were calculated according to the original procedure by Kennedy et al. [31]. The participants' feeling of presence was self-assessed through the Slater-Usoh-Steed (SUS) after each session [62, 72, 73].

Additionally, we evaluated the subtlety of all methods through subjective ratings. We let participants fill in the four questions:
[Q1] Rate the visual quality of the VR experience, disregarding any specific details of the scene.
[Q2] Did you notice visual distortions? (Yes/No)
[Q3] How noticeable were the distortions?
[Q4] How distracting were the distortions?

Except for the second question, all responses were recorded on a 7-point rating scale ranging from "very low" (1) to "very high" (7). Questions three and four were only presented if a participant answered "Yes" to the second question. Here, distortion refers to perceptual artifacts that are not part of the underlying scene content, similar to IQA definitions [7].

5.1.5 Procedure

The experiment consisted of three sessions per participant, with one condition (control, blur, or noise) per session. Each session followed an identical protocol. At the start of the study, participants were informed about the experimental procedure and side effects, and provided written consent. Demographic data were collected in this phase.

At the beginning of each session, participants completed the pre-exposure SSQ to establish a baseline. They then adjusted the VR headset and performed a built-in eye-tracker calibration. During the VR experience, participants remained seated but were allowed to look around freely as the camera followed the predefined path.

We instructed participants to report any severe symptoms immediately, which would result in the session being terminated. Otherwise, the VR experience lasted for 12 minutes. Following each session, participants completed the post-exposure SSQ, the SUS, and the subtlety-related questions.

5.2 Analysis and Results

Statistical analysis was conducted to evaluate the impact of the three within-subject conditions (control, foveation, and noise) on cybersickness ($N = 28$). Shapiro-Wilk tests for normality were non-significant ($p > 0.05$) for the differences between the control and noise conditions (SSQ scores), though a significant deviation was observed in comparisons involving the foveation method. However, visual inspection of Q-Q plots revealed a predominantly linear distribution with minor deviations limited to the extreme tails, suggesting that the data structure remained suitable for parametric testing. Due to the inherent resilience of the F-test against moderate deviations from normality, the repeated measures ANOVA was retained as the primary analytical tool. Following the primary analysis, post-hoc comparisons were executed using dependent t-tests. All resulting p-values were then corrected with the Holm-Bonferroni method to control the family-wise error rate.

The repeated measures ANOVA revealed a significant main effect of condition on total SSQ scores ($F(2,54) = 4.32$, $p = 0.0182$). Post-hoc analysis demonstrated that the noise condition significantly reduced total sickness compared to the control ($p = 0.0333$) and foveation ($p = 0.0470$), while foveated blurring showed no improvement over control ($p = 0.9889$). Analysis of the nausea subscale confirmed a significant main effect ($F(2,54) = 5.53$, $p = 0.0065$), with the noise condition performing significantly better than both control ($p = 0.0207$) and foveation ($p = 0.0207$). For the oculomotor subscale, the main effect was significant ($F(2,54) = 4.17$, $p = 0.0207$); here, post-hoc tests showed the noise condition only differed significantly from the control ($p = 0.0245$) and not from foveation ($p = 0.0625$). The disorientation subscale did not yield a significant main effect ($F(2,54) = 2.30$, $p = 0.1104$).

Analysis of the total time spent within the virtual environment yielded no significant differences across conditions ($\mu_{control} = 641, \mu_{blur} = 639, \mu_{noise} = 634$; $F(2,54) = 0.0476, p = 0.9537$), though 10 sessions in each condition were terminated prematurely due to severe sickness. Presence scores remained high across all groups, with averages of 4.85 for control, 4.65 for foveation, and 4.80 for noise (see Figure 5), showing no statistically significant variation. Subjective evaluations of visual quality (**Q1**) were generally high, though the control condition ($M = 4.79$) was rated significantly better than foveation ($M = 4.14, p = 0.0098$) and noise ($M = 4.25, p = 0.0222$). Interestingly, the proportion of participants noticing visual distortions (**Q2**) was identical for control and foveation (43%) but lower for the noise condition (29%). Further analysis suggests that these reports of distortion were not arbitrary, as specific participants consistently reported distortion across sessions. However, a subset of these participants no longer reported distortions specifically during the noise condition. The followup questions on noticeability (**Q3**) and distraction (**Q4**) are intended to provide further insight into this behavior. Analyzing the results of both scales, we observe that indeed the mean scores for noticeability and distraction indicated that foveated blurring was perceived as more prominent ($M_{notice} = 3.33, M_{distract} = 2.73$) than both the control ($M_{notice} = 2.53, M_{distract} = 2.47$) and noise ($M_{notice} = 2.2, M_{distract} = 1.8$) conditions. However, we have to reject the statistical significance between conditions due to high inter-participant variability and a low sample size (only presented after positive response in **Q2**). These results suggest that while the noise method may effectively mask perceived artifacts for some users, individual differences in perception remain a dominant factor.

6 Discussion & Limitations

In this section, we discuss the results and further implications of our findings, as well as their limitations.

Effectiveness and Subtlety of Counter-vection Noise. The results of our experiment indicate that our counter-vection method with Gabor noise can significantly mitigate cybersickness compared to both an unmodified presentation and foveated blurring. Because foveated blurring alone did not demonstrate a significant improvement over the control condition, these findings suggest that actively introducing opposing motion energy in the peripheral aliasing band, rather than simply stripping away high-frequency detail, is more effective at reducing visual-vestibular conflict.

Additionally, the subjective metrics of visual quality and presence highlight the subtlety of our implementation. While visual quality ratings for both foveation and noise conditions were slightly lower than the control, the participants' sense of presence remained high across all sessions. Interestingly, the proportion of participants who reported consciously noticing visual distortions was substantially lower in the noise condition (29%) than in the foveated blurring condition (43%). This result suggests that peripheral vision is tolerant to random luminance perturbations when the global multi-scale structure of the scene is preserved. This observation is consistent with the concept of peripheral metamers, where images that share similar pooled multiscale statistics appear perceptually indistinguishable despite local differences [13]. Since the additive signal has a texture-like frequency spectrum, is confined to a range of luminances, and is tuned to local motion, it can perceptually recover blur-induced losses without disturbing chromatic appearance or overall scene structure [15, 39, 59]

Our current implementation is calibrated to prioritize subtlety, ensuring the noise stays barely noticeable for the average user. However, cybersickness mitigation is inherently balanced by a trade-off between visibility and mitigation strength. It is likely that sensitive users might benefit from, or even prefer, a more intensive counter-vection signal. Altering the optimization parameter in Section 3 enables individual calibration of the mitigation technique to account for inter-personal sensitivities. In this work, we aimed for a generalizable solution that works for most users and requires no further calibration. The noise generation alone is computationally lightweight and likely applicable to standalone VR devices, where the simulated blur is replaced with real foveated rendering. Also, the method may be further improved when paired with gaze guidance method [35, 36] or chroma modulations [43]. For headsets without eye tracking, gaze prediction methods promise to be a powerful alternative [29].

Results of Foveated Blur In contrast to what many prior works suggest, foveated blurring did not reduce cybersickness in our experiment. foveated blurring is highly dependent on scene composition and motion structure, since the method essentially acts as a low-pass filter. While high-frequency motion information is always removed, scenes with large-scale low-frequency optic flow can still strongly drive vection, including particularly naturalistic scenes like ours [71]. In a former work, we observed the same behaviour: blur alone did not significantly reduce cybersickness [17].

Phase shift vs. Physical Displacement. In our approach, we achieve localized movement by temporally shifting the phase of the Gabors. While motion could alternatively be achieved by physically translating the Gabor kernels in image space according to the motion vector $\mathbf{m}(\mathbf{p})$, such an approach introduces significant challenges. Physical displacement causes kernels to cluster in motion sinks and leave voids in source regions, necessitating constant spawning, culling, and tracking of samples. Instead, we keep the kernel centers stationary and modulate the internal phase ϕ, which keeps the motion energy local. This decouples the perceived carrier motion from the spatial sampling distribution, ensuring a stable, high-quality motion percept without the computational overhead or sampling artifacts associated with dynamic particle systems.

Generalization and Binocular Integration. Our counter-vection technique operates as a screen-space post-processing shader and therefore remains independent of scene-specific rendering pipelines. This property enables straightforward integration into existing XR systems and makes the approach suitable for deployment at the device level without modifications to application content. Such a design supports broad applicability across different scenes and rendering engines.

The current implementation synthesizes the Gabor noise independently for each eye. The resulting luminance signals are not stereoscopically correlated, which could induce conflicting binocular signals in certain areas and may increase visual fatigue during long sessions [6]. Future work should therefore investigate stereoscopically consistent noise generation. One possible extension to the method is integrating depth information and aligning the sinusoidal phase shifts with the spatial depth of the underlying virtual content. Such depth-consistent modulation could strengthen the perceptual binding of the signal to the scene, potentially reducing noise resolvability while further increasing the effectiveness of the counter-vection stimulus.

Limitations. We further acknowledge that potential flicker-related effects by the Gabor noise may induce visual discomfort in sensitive users during prolonged exposure. Also, we emphasize that the amplitude model (see Section 3.6) is an engineering approximation rather than a validated one. Since we designed the experiment to be passively explored for better comparability, direct generalization to fully interactive VR is not guaranteed. However, the proposed method operates on scene motion itself and is implemented application-independent. We therefore expect the approach to transfer to other settings. Any perceptual manipulation introduces a trade-off between mitigation strength and visual fidelity. Our method therefore might not be ideal for applications requiring accurate peripheral scene interpretation. The results demonstrate a clear positive trend for our mitigation strategy, yet we acknowledge that these results might be biased by the gender imbalance of the participant pool. Prior work has reported higher cybersickness susceptibility in female users [44]. Future studies with larger and more balanced samples could therefore investigate potential sex-related differences to obtain a more complete understanding of the mitigation effects.

7 Conclusion

We introduced a real-time counter-vection technique that induces motion-opposing Gabor noise into the periphery to rebalance perceived motion without removing scene content. Our user experiment demonstrated that cybersickness is reduced while presence is maintained, indicating perceptual signal integration as a promising future direction for VR mitigation techniques.

Acknowledgments

The authors gratefully acknowledge funding by the German Science Foundation (DFG project 557564533) and the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 804226 PERDY).

References

[1] I. B. Adhanom, N. N. Griffin, P. MacNeilage, and E. Folmer. The effect of a foveated field-of-view restrictor on vr sickness. In *IEEE Conf. Virtual Reality and 3D User Interfaces*, pp. 645–652. IEEE, 2020. 2

[2] M. Ashraf, R. K. Mantiuk, A. Chapiro, and S. Wuerger. castlecsf—a contrast sensitivity function of color, area, spatiotemporal frequency, luminance and eccentricity. *Journal of vision*, 24(4):5–5, 2024. 4, 5

[3] K. R. Bankieris, V. R. Bejjanki, and R. N. Aslin. Sensory cue-combination in the context of newly learned categories. *Scientific Reports*, 7(1), 2017. 2

[4] P. Bénard, A. Lagae, P. Vangorp, S. Lefebvre, G. Drettakis, and J. Thollot. A dynamic noise primitive for coherent stylization. In *Computer Graphics Forum*, vol. 29, pp. 1497–1506. Wiley Online Library, 2010. 3

[5] J. E. Bos, S. C. de Vries, M. L. van Emmerik, and E. L. Groen. The effect of internal and external fields of view on visually induced motion sickness. *Applied ergonomics*, 41(4):516–521, 2010. 2

[6] T. Cai, H. Zhu, J. Xu, S. Wu, X. Li, and S. He. Human cortical neural correlates of visual fatigue during binocular depth perception: an fnirs study. *PLoS One*, 12(2):e0172426, 2017. 9

[7] U. Cogalan, M. Bemana, K. Myszkowski, H.-P. Seidel, and C. Groth. Milo: A lightweight perceptual quality metric for image and latent-space optimization. *ACM Transactions on Graphics (TOG)*, 44(6):1–11, 2025. 8

[8] K. K. De Valois, R. L. De Valois, and E. W. Yund. Responses of striate cortex cells to grating and checkerboard patterns. *The Journal of Physiology*, 291:483–505, June 1979. doi: 10.1113/jphysiol.1979.sp012827 2

[9] R. L. De Valois and K. K. De Valois. Vernier acuity with stationary moving gabors. *Vision research*, 31(9):1619–1626, 1991. 3

[10] C. J. Dizmen and R. H. So. Vection induced by a pair of patches of synchronized visual motion stimuli covering total field of views as small as 10 square-degrees. *i-Perception*, 14(5):20416695231201463, Sept. 2023. doi: 10.1177/20416695231201463 2

[11] S. D'Amour, L. R. Harris, S. Berti, and B. Keshavarz. The role of cognitive factors and personality traits in the perception of illusory self-motion (vection). *Attention, Perception, & Psychophysics*, 83:1804–1817, 2021. 1, 2

[12] R. Fan, X. Shi, K. Wang, Q. Ma, and L. Wang. Scene-aware foveated rendering. *IEEE Transactions on Visualization and Computer Graphics*, 2024. 2

[13] J. Freeman and E. P. Simoncelli. Metamers of the ventral stream. *Nature neuroscience*, 14(9):1195–1201, 2011. 2, 9

[14] Y. Fujii, O. G. Kio, D. Au, L. M. Wilcox, and R. S. Allison. Effects of frame rate on vection and postural sway. *Displays*, 58:33–43, July 2019. doi: 10.1016/j.displa.2019.03.002 2

[15] M. A. Georgeson. Contrast overconstancy. *Journal of the Optical Society of America A*, 8(3):579–586, 1991. 9

[16] J. A. Greenwood, P. J. Bex, and S. C. Dakin. Positional averaging explains crowding with letter-like stimuli. *Proceedings of the National Academy of Sciences*, 106(31):13130–13135, 2009. doi: 10.1073/pnas.0901352106 2

[17] C. Groth, M. Magnor, S. Grogorick, M. Eisemann, and P. Didyk. Cybersickness reduction via gaze-contingent image deformation. *ACM Transactions on Graphics (Proc. of Siggraph)*, 43(4):1–14, Jul 2024. doi: 10.1145/3658138 1, 2, 5, 6, 9

[18] C. Groth, J.-P. Tauscher, N. Heesen, S. Castillo, and M. Magnor. Visual techniques to reduce cybersickness in virtual reality. In *IEEE Conf. Virtual Reality and 3D User Interfaces*, pp. 486–487, 2021. 1, 2, 5, 6

[19] C. Groth, J.-P. Tauscher, N. Heesen, S. Grogorick, S. Castillo, and M. Magnor. Mitigation of cybersickness in immersive 360°videos. In *IEEE VR Workshop on Immersive Sickness Prevention*, pp. 169–177, 2021. 1, 2, 6

[20] C. Groth, J.-P. Tauscher, N. Heesen, M. Hattenbach, S. Castillo, and M. Magnor. Omnidirectional galvanic vestibular stimulation in virtual reality. *IEEE Trans. Vis. Comput. Graph.*, 28(5):2234–2244, 2022. 2, 7

[21] B. Guenter, M. Finch, S. Drucker, D. Tan, and J. Snyder. Foveated 3d graphics. *ACM Trans. Graph.*, 31(6), Nov. 2012. doi: 10.1145/2366145.2366183 2

[22] X. Guo, S. Nakamura, Y. Fujii, T. Seno, and S. Palmisano. Effects of luminance contrast, averaged luminance and spatial frequency on vection. *Experimental Brain Research*, 239:3507–3525, 2021. 2, 4, 5, 7

[23] D. J. Heeger and J. R. Bergen. Pyramid-based texture analysis/synthesis. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pp. 229–238, 1995. 2

[24] R. Hendrichs, J. H. Israel, M. Magnor, and C. Groth. Towards general motion sickness reduction via semi-random galvanic vestibular stimulation. *IEEE Transactions on Visualization and Computer Graphics*, 2026. 2

[25] S. Hillaire, A. Lécuyer, R. Cozot, and G. Casiez. Depth-of-field blur effects for first-person navigation in virtual environments. *IEEE computer graphics and applications*, 28(6):47–55, 2008. 2

[26] P. Hu, Q. Sun, P. Didyk, L.-Y. Wei, and A. E. Kaufman. Reducing simulator sickness with perceptual camera control. *ACM Trans. Graph.*, 38(6), Nov. 2019. doi: 10.1145/3355089.3356490 2

[27] H. Kam, C. Groth, J. Kang, P. Saraf, Q. Sun, and K. Chen. Dichoptic foveation. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers*, SIGGRAPH Conference Papers '26. Association for Computing Machinery, New York, NY, USA, 2026. doi: 10.1145/3799902.3811108 2

[28] H. Kam, N. L. Williams, J. Lubin, J. R. Bergen, M. Sizintsev, Q. Sun, and C. Groth. Attention drives sickness in augmented reality. *IEEE Transactions on Visualization and Computer Graphics*, 2026. 2

[29] J. Kang, C. Groth, T. Wu, F. Torrens, P. Sangkloy, G. Wetzstein, and Q. Sun. Infinite gaze generation for videos with autoregressive diffusion. *arXiv preprint arXiv:2603.24938*, 2026. 9

[30] A. Kaplanyan, A. Sochenov, T. Leimkuehler, M. Okunev, T. Goodall, and G. Rufo. DeepFovea: Neural reconstruction for foveated rendering and video compression using learned statistics of natural videos. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)*, 38(4):212:1–212:13, 2019. 2

[31] R. S. Kennedy, N. E. Lane, K. S. Berbaum, and M. G. Lilienthal. Simulator sickness questionnaire: An enhanced method for quantifying simulator sickness. *The Int. Journal of Aviation Psychology*, 3(3):203–220, 1993. 1, 8

[32] S. Kim and G. J. Kim. Dynamically Adjusted and Peripheral Visualization of Reverse Optical Flow for VR Sickness Reduction. *Electronics*, 12(4), Jan. 2023. doi: 10.3390/electronics12040861 2

[33] S. M. Kim and G. J. Kim. Sparse peripheral counter-vection flow visualization with reverse optical flow for vr sickness reduction. In *International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*, pp. 798–799, 2022. 2

[34] B. Krajancich, P. Kellnhofer, and G. Wetzstein. Towards attention–aware foveated rendering. *ACM Transactions on Graphics (TOG)*, 42(4):1–10, 2023. 2

[35] J. Kudnick, D. T. Mayer, C. Groth, B. Mohanto, R. Dörner, and M. Weier. Ripplevision: Unobtrusive gaze-dependent guidance via directed wave motion in virtual reality. *IEEE Transactions on Visualization and Computer Graphics*, 2026. 9

[36] J. Kudnick, M. Weier, C. Groth, B. Fu, and R. Horst. Warpvision: Using spatial curvature to guide attention in virtual reality. *IEEE Transactions on Visualization and Computer Graphics*, 2025. 9

[37] J.-W. Lin, H. B.-L. Duh, D. E. Parker, H. Abi-Rached, and T. A. Furness. Effects of field of view on presence, enjoyment, memory, and simulator sickness in a virtual environment. In *Proc. IEEE Virtual Reality*, pp. 164–171. IEEE, 2002. 2

[38] Y.-X. Lin, R. Venkatakrishnan, R. Venkatakrishnan, E. Ebrahimi, W.-C. Lin, and S. V. Babu. How the presence and size of static peripheral

blur affects cybersickness in virtual reality. *ACM Trans. Appl. Percept.*, 17(4), 2020. 2

[39] B. B. Lopes, T. M. G. Loureiro, F. A. d. C. Brito, L. Miquilini, A. Brasil, M. F. Costa, R. C. Salomão, D. F. Ventura, A. L. Brino, and G. d. S. Souza. Luminance noise impacts putative luminance components of visual evoked potentials but not chromatic components. *Psicologia: Reflexão e Crítica*, 38:22, 2025. 9

[40] W. Luu, B. Zangerl, M. Kalloniatis, and J. Kim. Effects of stereopsis on vection, presence and cybersickness in head-mounted display (hmd) virtual reality. *Scientific reports*, 11(1):12373, 2021. 2

[41] R. K. Mantiuk, M. Ashraf, and A. Chapiro. stelacsf: a unified model of contrast sensitivity as the function of spatio-temporal frequency, eccentricity, luminance and area. *ACM Transactions on Graphics (TOG)*, 41(4):1–16, 2022. 4, 5

[42] X. Meng, R. Du, and A. Varshney. Eye-dominance-guided foveated rendering. *IEEE transactions on visualization and computer graphics*, 26(5):1972–1980, 2020. 2

[43] B. Mohanto, M. Weier, C. Groth, J. Kudnick, and O. Staadt. Perceptually calibrated foveated chroma attenuation for gaze-contingent vr displays. In *2026 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*, pp. 1191–1192. IEEE, 2026. 9

[44] J. Munafo, M. Diedrick, and T. A. Stoffregen. The virtual reality head-mounted display oculus rift induces motion sickness and is sexist in its effects. *Experimental brain research*, 235(3):889–901, 2017. 9

[45] B. Murovec, J. Spaniol, J. L. Campos, and B. Keshavarz. Multisensory effects on illusory self-motion (vection): The role of visual, auditory, and tactile cues. *Multisensory Research*, 34(8):869–890, 2021. 2

[46] G. Nie, Y. Liu, and Y. Wang. Prevention of visually induced motion sickness based on dynamic real-time content-aware non-salient area blurring. In *IEEE Int. Symp. Mixed and Augmented Reality Adjunct*, pp. 75–78, 2017. 2

[47] G.-Y. Nie, H. B.-L. Duh, Y. Liu, and Y. Wang. Analysis on mitigation of visually induced motion sickness by applying dynamical blurring on a user's retina. *TVCG*, 2019. 2

[48] T. Nie, C. H. Pospick, V. Cantory, D. Zhang, J. J. DeGuzman, V. Interrante, I. B. Adhanom, and E. S. Rosenberg. Peripheral teleportation: A rest frame design to mitigate cybersickness during virtual locomotion. *IEEE Transactions on Visualization and Computer Graphics*, 31(5):2891–2900, May 2025. doi: 10.1109/TVCG.2025.3549568 2

[49] J. F. O'Hanlon, M. E. McCauley, et al. Motion sickness incidence as a function of the frequency and acceleration of vertical sinusoidal motion. *Aerospace medicine*, 45(4):366–369, 1974. 7

[50] S. Palmisano and B. Gillam. Stimulus Eccentricity and Spatial Frequency Interact to Determine Circular Vection. *Perception*, 27(9):1067–1077, Sept. 1998. doi: 10.1068/p271067 2

[51] S. H. Park, B. Han, and G. J. Kim. Mixing in reverse optical flow to mitigate vection and simulation sickness in virtual reality. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*, 2022. 2

[52] A. Patney, M. Salvi, J. Kim, A. Kaplanyan, C. Wyman, N. Benty, D. Luebke, and A. Lefohn. Towards foveated rendering for gaze-tracked virtual reality. *ACM Trans. Graph.*, 35(6), Nov. 2016. doi: 10.1145/2980179.2980246 2

[53] X. Peng, C. 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 *2025 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 1075–1084. IEEE, 2025. 1

[54] L. Qiu Lin, H. Kam, Q. Sun, and K. Akşit. *Foveation Improves Payload Capacity in Steganography*. Association for Computing Machinery, New York, NY, USA, 2025. 2

[55] J. T. Reason and J. J. Brand. *Motion sickness.* Academic press, 1975. 2

[56] R. Rosenholtz, J. Huang, A. Raj, B. J. Balas, and L. Ilie. A summary statistic representation in peripheral vision explains visual search. *Journal of Vision*, 12(4):14, 2012. doi: 10.1167/12.4.14 2

[57] X. M. Sauvan and C. Bonnet. Spatiotemporal boundaries of linear vection. *Perception & Psychophysics*, 57(6):898–904, Jan. 1995. doi: 10.3758/BF03206804 2

[58] A. M. Schmid, K. P. Purpura, I. E. Ohiorhenuan, F. Mechler, and J. D. Victor. Subpopulations of neurons in visual area v2 perform differentiation and integration operations in space and time. *Frontiers in systems neuroscience*, 3:915, 2009. 2

[59] A. J. Schofield and M. A. Georgeson. Sensitivity to modulations of luminance and contrast in visual white noise: Separate mechanisms with similar behaviour. *Vision Research*, 39(16):2697–2716, 1999. 9

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

[61] A. F. Seay, D. M. Krum, L. Hodges, and W. Ribarsky. Simulator sickness and presence in a high fov virtual environment. In *Proc. IEEE Virtual Reality*, pp. 299–300. IEEE, 2001. 2

[62] M. Slater, J. McCarthy, and F. Maringelli. The influence of body movement on subjective presence in virtual environments. *Human factors*, 40(3):469–477, 1998. 8

[63] M. Stengel, S. Grogorick, M. Eisemann, and M. Magnor. Adaptive image-space sampling for gaze-contingent real-time rendering. In *Computer graphics forum*, vol. 35, pp. 129–139. Wiley Online Library, 2016. 2

[64] T. Tariq and P. Didyk. Towards motion metamers for foveated rendering. *ACM Transactions on Graphics (TOG)*, 43(4):1–10, 2024. 2, 7

[65] T. Tariq, C. Tursun, and P. Didyk. Noise-based enhancement for foveated rendering. *ACM Transactions on Graphics (TOG)*, 41(4):1–14, 2022. 2, 4

[66] L. N. Thibos, D. L. Still, and A. Bradley. Characterization of spatial aliasing and contrast sensitivity in peripheral vision. *Vision Research*, 36(2):249–258, 1996. doi: 10.1016/0042-6989(95)00109-D 2, 4

[67] L. N. Thibos, D. Walsh, and F. Cheney. Vision beyond the resolution limit: aliasing in the periphery. *Vision Research*, 27(12):2193–2197, 1987. 2

[68] B. Thompson, B. C. Hansen, R. F. Hess, and N. F. Troje. Peripheral vision: Good for biological motion, bad for signal noise segregation? *Journal of Vision*, 7(10):1–12, 2007. 2, 5

[69] P. Tse and P.-J. Hsieh. The infinite regress illusion reveals faulty integration of local and global motion signals. *Vision research*, 46(22):3881–3885, 2006. 3

[70] O. T. Tursun, E. Arabadzhiyska-Koleva, M. Wernikowski, R. Mantiuk, H.-P. Seidel, K. Myszkowski, and P. Didyk. Luminance-contrast-aware foveated rendering. *ACM Trans. Graph.*, 38(4):1–14, 2019. 4, 5

[71] M. Unser and T. Blu. Mathematical properties of the jpeg2000 wavelet filters. *IEEE Trans. Image Process.*, 12(9):1080–1090, 2003. doi: 10.1109/TIP.2003.812329 9

[72] M. Usoh, K. Arthur, M. C. Whitton, R. Bastos, A. Steed, M. Slater, and F. P. Brooks Jr. Walking¿ walking-in-place¿ flying, in virtual environments. In *Annual Conf. Computer graphics and interactive techniques*, pp. 359–364, 1999. 8

[73] M. Usoh, E. Catena, S. Arman, and M. Slater. Using presence questionnaires in reality. *Presence (Camb)*, 9(5):497–503, 2000. 8

[74] D. R. Walton, R. K. Dos Anjos, S. Friston, D. Swapp, K. Akşit, A. Steed, and T. Ritschel. Beyond blur: Real-time ventral metamers for foveated rendering. *ACM Trans. Graph.*, 40(4):1–14, 2021. 2

[75] M. Weier, T. Roth, E. Kruijff, A. Hinkenjann, A. Pérard-Gayot, P. Slusallek, and Y. Li. Foveated real-time ray tracing for head-mounted displays. *Computer Graphics Forum*, 35(7):289–298, 2016. doi: 10.1111/cgf.13026 6

[76] R. Xiao and H. Benko. Augmenting the Field-of-View of Head-Mounted Displays with Sparse Peripheral Displays. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, pp. 1221–1232. Association for Computing Machinery, May 2016. doi: 10.1145/2858036.2858212 2

[77] Y. Zhang, K. You, X. Hu, H. Zhou, K. Kiyokawa, and X. Yang. Retinotopic foveated rendering. In *2024 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*, pp. 903–912, 2024. doi: 10.1109/VR58804.2024.00109 2