Dichoptic Foveation

HENRY KAM, New York University, USA
COLIN GROTH, New York University, USA
JENNA KANG, New York University, USA
PRATHAM SARAF, New York University, USA
QI SUN, New York University, USA
KENNETH CHEN, New York University, USA

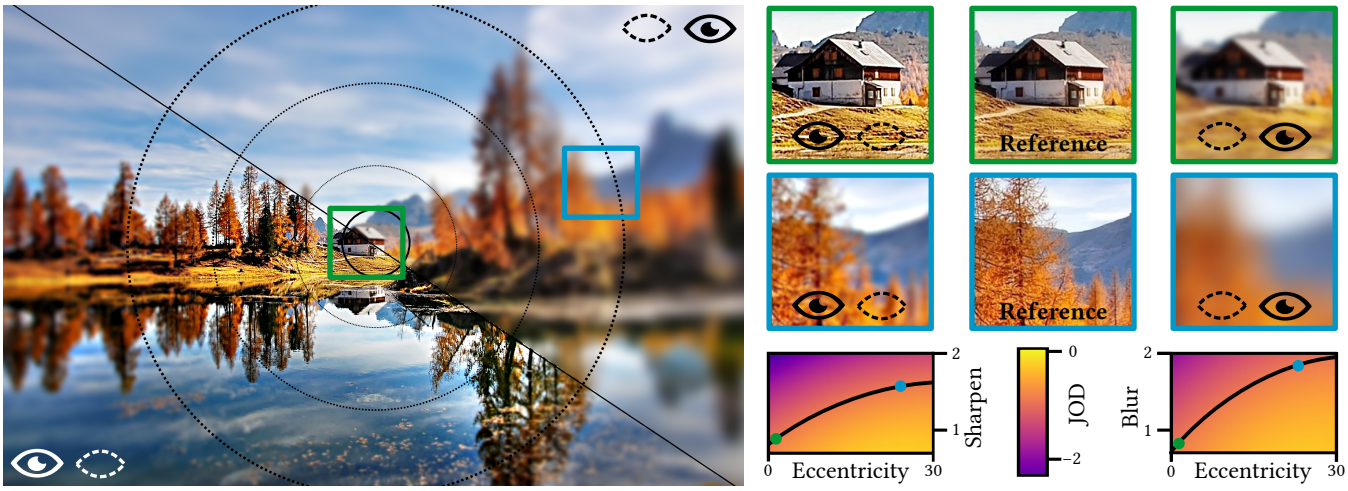


Fig. 1. We present dichoptic foveation, a new paradigm to parameterize quality degradation of dichoptically-rendered stimuli in stereoscopic displays like virtual reality. Particularly, we apply Gaussian blur on the stimulus of one eye (right inset images here), while the other eye's stimulus is processed by a simultaneous sharpening routine (left insets). This is guided by our computational model (see heatmaps), which is built on psychophysical data. Dichoptic foveation has direct implications for better quality-to-compute ratios in foveated rendering. Image by kordi_vahle from Pixabay.

Interocular differences in visual perception can induce a variety of effects when fused by the brain. For example, prior works have found that carefully crafted binocular differences in local detail can improve contrast. It has also been found that when the frequency content of two stimuli are slightly different, blur suppression leads to a fused percept that is typically dominated by the sharper image. In this paper, we develop a psychophysical framework to measure the perception of natural image stimuli with interocular frequency differences across the visual field. To this end, we study the effect of *dichoptic foveation*, which we define as the application of blur to one eye and a simultaneous sharpening filter to the other. Stimuli were viewed in a virtual reality (VR) head-mounted display (HMD) and placed at different retinal eccentricities. Study data were scaled to a perceptual just objectionable difference (JOD) scale, and a 4D model was fit to it; our results suggest that interocular frequency differences can be well described by a simple computational model. We applied the model in a realistic VR scenario with free exploration of 360° videos to improve a base foveated rendering system by enhancing high frequency information dichoptically.

CCS Concepts: • **Computing methodologies** → **Perception**; *Visibility*.

Additional Key Words and Phrases: dichoptic perception, foveation, VR, rendering

ACM Reference Format:
Henry Kam, Colin Groth, Jenna Kang, Pratham Saraf, Qi Sun, and Kenneth Chen. 2026. Dichoptic Foveation. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '26), July 19–23, 2026, Los Angeles, CA, USA.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3799902.3811108

1 Introduction

Understanding how we integrate binocular information is a long-standing problem with direct implications for designing stereoscopic displays like VR HMDs. For example, this principle has been shown to boost contrast [Zhong et al. 2019]. A large body of visual psychophysics work has also shown that when the two eyes receive information with unequal spatial fidelity, the fused percept is often dominated by the sharper stimulus. However, no previous work has explored this fusion effect for stimuli located across the visual field.

Authors' Contact Information: Henry Kam, henrykam@nyu.edu, New York University, New York, USA; Colin Groth, c.groth@nyu.edu, New York University, New York, USA; Jenna Kang, jennakang@nyu.edu, New York University, New York, USA; Pratham Saraf, ps5218@nyu.edu, New York University, New York, USA; Qi Sun, qisun@nyu.edu, New York University, New York, USA; Kenneth Chen, kennychen@nyu.edu, New York University, New York, USA.

This work is licensed under a Creative Commons Attribution 4.0 International License.
SIGGRAPH Conference Papers '26, Los Angeles, CA, USA
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2554-8/2026/07
https://doi.org/10.1145/3799902.3811108

Significant blur in one eye can remain undetected when perceived stereoscopically, as shown in studies of interocular blur suppression [Collins and Goode 1994; Georgeson and Wallis 2014; Schor et al. 1987] or binocular contrast differences [Wang et al. 2023a]. This finding suggests that the visual system may preserve high-frequency information even when interocular frequency differences occur. Prior work on dichoptic rendering has shown that carefully crafting differences in binocular imagery can improve depth perception [Wang and Cooper 2021] or contrast perception [Zhong et al. 2019]. However, sensitivity to these interocular differences across retinal eccentricities has not been studied, especially regarding how blur and sharpening across the two eyes interact dichoptically.

In this paper, we study *dichoptic foveation*, which we define as the deliberate presentation of a blurred stimulus to one eye and the sharpened image to the other, and its variation with retinal eccentricity, defined as angular distance from the fovea, as exemplified in Figure 1. We ask the question: at a given eccentricity, can fusing blurred and sharpened images across the two eyes yield equivalent, or even improved perceived quality compared to symmetric presentation?

Based on our psychophysical data, we fit a simple perceptual model that predicts a quality correlate given blur–sharpen parameters and eccentricity. We then demonstrate how this model can be applied to improve foveated rendering, and evaluate this system via a subjective evaluation study. Using blur magnitude as a proxy to calculate system-agnostic computational cost, we show that our application of dichoptic foveation in VR can reduce sampling rates by about half that of the baseline foveated rendering technique. For this comparison, both foveated rendering and dichoptic foveation were calibrated to have minimal perceptual impact compared to the reference full-resolution content. Please refer to our supplementary webpage for all data, video examples, and interactive comparisons. To summarize, we

- develop a psychophysical framework and collected a large-scale dataset measuring perception of dichoptic stimuli with 12,989 human subject study trials;
- fit a model over three variables (blur, sharpen, and eccentricity) that predicts quality in just-objectionable-differences (JODs);
- leverage the model as an application to improve foveated rendering, which demonstrates the potential for ~50% compute savings at similar perceptual quality.

2 Background & Related Work

First, we describe the scientific background of binocular perception, then we address prior literature on dichoptic and foveated rendering.

2.1 Binocular and Dichoptic Perception

Stereoscopic perception, or stereopsis, emerges when the human visual system (HVS) integrates retinal images from the left and right eyes into a unified percept [Blake and Wilson 2011; Cang et al. 2023]. The neuropsychological mechanism that enables the combination of two monocular images into a single percept is referred to as binocular fusion, and occurs when light stimulates corresponding regions on the two retinas [Sperling 1970].

Fusion occurs in the forward-facing binocular visual field, the region where fields of view (FOVs) overlap between the eyes, which spans approximately 110° in the horizontal direction [Stidwill and Fletcher 2017]. Outside of this region, each eye has a 40° monocular periphery where no overlap occurs between the two eyes [Stidwill and Fletcher 2017]. It is critical to consider this characteristic of the HVS in dichoptic perception because the input of both eyes is required to form a fused stimulus.

Within the binocular visual field, the HVS can tolerate small interocular differences in sharpness without disrupting fusion. As these differences increase, blur suppression may occur, in which the sharper retinal image dominates [Georgeson and Wallis 2014]. When interocular discrepancies become sufficiently large, binocular rivalry may occur. Binocular rivalry is often characterized by alternating perceptual dominance between the two monocular images [Arnold et al. 2007]. Riesen et al. [2019] found that rivalry and fusion are not necessarily mutually exclusive states; a "tristable" zone spans 25°–35° where left or right eye-dominance could occur.

Dichoptic perception refers to the perception of differing visual inputs between the two eyes. The described processes define the possible regimes of dichoptic perception [Jiang and Meng 2023; Qian et al. 2019]. When dichoptic inputs are sufficiently similar, binocular fusion occurs and yields stereopsis [Jiang and Meng 2023]; differences primarily in sharpness with similar structure can result in interocular blur suppression, producing a fused percept dominated by one eye [Lew et al. 2021]. Applying blur dichoptically to high-spatial-frequency content has a stronger likelihood of fusion breaks than luminance-contrast changes, especially under divergence stress [Dostalek et al. 2019]. When dichoptic differences become too large, fusion fails and binocular rivalry emerges, which may lead to alternating monocular dominance [Law et al. 2013].

2.2 Dichoptic Rendering in Computer Graphics

An example application of dichoptic perception to computer graphics is the work of Zhong et al. [2019], which presents images tone-mapped differently for each eye. The authors found that the fused image has improved perceived contrast. Similar results were observed for unintended dichoptic contrast differences in AR [Wang et al. 2023a]. Hoffman and Banks [2010] investigated how focus cues, like blur and depth-dependent information, are used by the HVS to solve the binocular correspondence problem. For guidance tasks in VR, dichoptic stimuli are often used to attract gaze [Kudnick et al. 2026, 2025]. Recent related work has found that applying different compression rates [Fezza and Larabi 2017; Lee et al. 2025; Shao et al. 2010] to each eye can reduce computation at similar perceptual quality. No prior work, however, has studied the perception of dichoptic frequency differences across the visual field.

2.3 Foveated Rendering Techniques

Human visual acuity decreases with increasing retinal eccentricity due to the variation in density of rods, cones, and retinal ganglion cells across the visual field. Foveated rendering is a set of techniques taking advantage of this relationship by reducing rendering samples with eccentricity, which is typically characterized as gaze-contingent blur [Ashraf et al. 2025; Fan et al. 2024; Groth et al. 2023a;

Guenter et al. 2012; Kaplanyan et al. 2019; Krajancich et al. 2023; Meng et al. 2018; Mohanto et al. 2025; Patney et al. 2016; Stengel et al. 2016; Tariq and Didyk 2024; Tariq et al. 2022; Zhang et al. 2024]. Gaze-contingent rendering has also been applied for display power reduction [Chen et al. 2024; Duinkharjav et al. 2022], for boosting stereo perception [Jiménez-Navarro et al. 2026; Kergaßner et al. 2025], for reducing cybersickness [Groth et al. 2024, 2021a,b], or for improving neural rendering efficiency [Deng et al. 2022; Lin et al. 2025]. Another set of papers explores how asymmetric levels of blur in both eyes enhance foveated rendering [Meng et al. 2020; Wang et al. 2023b, 2025]. While this asymmetry effectively creates a dichoptic percept, these works predominantly take advantage of eye dominance. Nevertheless, their results clearly show that stimuli with slight differences in blur can be fused. These findings motivate our work, but our approach is significantly different. Independent of eye dominance, we explore combinations of blur and sharpening operations on either eye that yield the highest level of modulation at the threshold of detectability, relying on binocular fusion and interocular blur suppression.

3 Psychophysical Framework

The core hypothesis of this experiment is that binocular fusion due to asymmetric blur–sharpen operations in each eye may preserve more visual detail than asymmetric dichoptic blur. Similar to dichoptic tone-mapping techniques [Zhong et al. 2019], if one eye retains higher-frequency structure (sharp) while the other eye provides a low-pass version of the same content (blur), the visual system may combine these asymmetric inputs into a fused percept that could appear closer to the reference. Because visual acuity decreases with retinal eccentricity, we hypothesize that this blur–sharpen relationship is also eccentricity-dependent.

3.1 Method

We parameterize both blur and sharpen operations by convolving images with Gaussian kernels. For a given input image I, the blurred image is produced as

$$I_{\text{blur}} = G_\sigma * I, \tag{1}$$

where G_σ is a Gaussian kernel with standard deviation σ, and $*$ is the convolution operator. In our implementation, we use a finite-support kernel whose size s is chosen to capture almost all of the Gaussian's mass,

$$s = 2 \cdot \lceil 3\sigma \rceil + 1. \tag{2}$$

This definition is a common choice that truncates the filter at approximately $\pm 3\sigma$. Sharpening is implemented via an unsharp masking-style operator. We first compute a blurred version of the original image using the same σ, then compute the high-frequency residual R, and subsequently add this residual back to the original image with a multiplicative gain ϕ,

$$\begin{aligned} I_{\text{sharp}} &= I + \phi \cdot R \\ &= I + \phi \cdot (I - G_\sigma * I). \end{aligned} \tag{3}$$

This formulation demonstrates how unsharp masking boosts the image components attenuated by the Gaussian low-pass filter. In practice, outputs are clamped to the valid display range. However, clipping only occurs when the residual R is large enough, which in practice happens for sharpening gains of $\phi > 1$. This constraint was accounted for when sampling parameters in our study.



Fig. 2. *Study stimuli.* We showed twelve distinct image patches featuring a variety of image content, color profiles, and frequency of the content. Images from Pixabay.

We intentionally couple the sharpening scale with the blur scale by using the same Gaussian standard deviation, σ, for both operations. If sharpening were the inverse operation of blurring, using the same σ could potentially offset the effect during binocular fusion. In practice, unsharp masking is not an exact inverse, so this hypothetical cancellation is only approximate. An example of these operations is shown in Figure 4.

3.2 Experimental Setup

We used the same experimental setup for the pilot studies and the main perceptual study, which is first described here.

Hardware apparatus. Stimuli were presented on a Meta Quest Pro HMD. The headset has a refresh rate of 90 Hz, 1,800×1,920 resolution per eye, 106° horizontal FOV, and 80° binocular overlap.

Stimuli. Our study was conducted on natural image patches sourced from Pixabay. The selected patches are shown in Figure 2. While perceptual studies often rely on controlled stimuli like Gabor patches, our goal was to study quality in a realistic setting. We assigned each patch to a low, medium, or high spatial-frequency bin based on the relative amount of high-frequency energy in its Fourier power spectrum. Each set of test stimuli contained two image patches from each frequency bin. Patches were scaled to 141×141 resolution for display, which was the smallest stimulus size that retains the highest spatial frequencies above the Nyquist limit of the headset. Patches were displayed on a solid gray background and had a size of approximately 7° of visual angle. Patches were presented spanning equally-spaced steps across the headset's binocular field of view (eccentricities of 0°, 15°, or 30°) along the horizontal meridian, and the placement of stimuli on the left or right side was randomized across trials. Observers were instructed to fixate on a persistent central crosshair; stimuli were hidden if gaze deviated. The crosshair was removed when stimuli were presented at the foveal location (0°).

3.3 Pilot Experiments

One critical caveat is that binocular fusion is not strictly guaranteed. Sufficiently large interocular differences can lead to binocular rivalry or other forms of perceptual instability. We therefore expect the fusion of dichoptic pairs to fail for some parameter ranges, depending on the blur–sharpen strength, the stimulus spatial frequency content, and retinal eccentricity. For the main experiment,

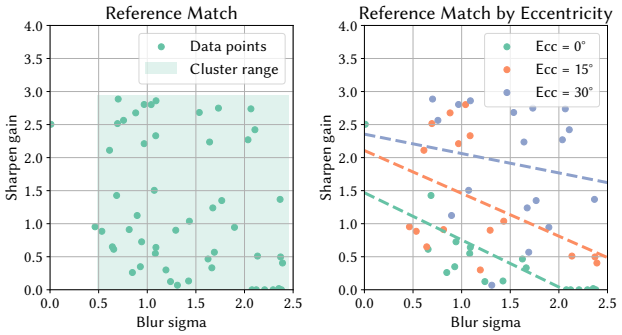


Fig. 3. *Pilot study results.* The left plot shows method of adjustment results of pilot 1. The results show that the selected parameters cluster well in the range outlined by the green background. The right plot shows the linear fits to per-eccentricity data, which were used to sample conditions for the 2nd pilot.

we conducted two pilot studies to identify a reasonable parameter range as well as find blur–sharpen levels with uniform perceptual spacing.

3.3.1 Pilot 1: Calibrating Blur–Sharpen Range. We conducted a method-of-adjustment [Wichmann and Hill 2001] study to identify a suitable parameter range for blur (σ) and sharpening gain (ϕ). We recruited 4 participants (1 female, 3 male, aged 22-31) for this pilot. In each trial, they viewed a test dichoptic image pair placed at the three tested eccentricities, and were able to flip between this test stimulus and an unmodified reference image. The test dichoptic pair was initialized with a random combination of $\sigma \in [0, 3]$ and $\phi \in [0, 4]$, and observers could adjust these two parameters until they found the fused percept to closely match the reference. Across participants and images, the resulting adjustments clustered within a well-defined range of blur and sharpen values (see Figure 3). Based on these observations, we selected $\sigma \in [0, 2.5]$ and $\phi \in [0, 2.5]$ as the range for subsequent testing.

3.3.2 Pilot 2: Determining Main Study Parameters. The first pilot established the parameter range, but left the question open on how best to sample this range to achieve perceptually-uniform spacing. While simple linear spacing would be the most straightforward parametrization, it might not result in a *perceptually* uniform test set [Wichmann and Hill 2001], leading to weak model approximation in the main study.

We ran a second pilot study to select specific stimulus levels that vary in approximately uniform perceived quality. Using the same stimulus pipeline, 5 observers (4 male, 1 female, aged 22–31 years), completed a two-alternative forced-choice (2AFC) similarity task with an unaltered reference. Each trial presented two test dichoptic pairs (parameterized by (σ, ϕ)), where participants selected the test pair that appeared closer to the reference. We evaluated a grid of parameter values by uniformly sampling five blur levels and five sharpening gains over the range determined in the first pilot $\sigma, \phi \in [0, 2.5]$, forming a Latin grid for pairwise comparison. The



Fig. 4. *Stimulus processing.* We show a reference patch (left) after applying sharpening (middle) and blurring (right) of medium strength. Image by terski from Pixabay

pairwise preference data were converted into a perceptual JOD scale via *pwcmp* [Perez-Ortiz and Mantiuk 2017].

To prioritize accurate sampling in foveal conditions, we used the 0° condition and examined JOD scores as a function of σ and ϕ separately. We fit a linear model to these data to estimate the parameter value corresponding to 0 JOD, which corresponds to the reference condition (i.e., $\sigma = 0$ and $\phi = 0$). We then selected four additional levels at 0.5 JOD decrements from this reference point, yielding five blur and five sharpening values that are approximately evenly spaced in perceived quality,

Blur (σ): [0, 0.2704, 0.5408, 0.8113, 1.0817],
Sharpen (ϕ): [0, 0.2235, 0.4471, 0.6706, 0.8941].

3.4 Main Experiment

We conducted a within-subjects experiment to measure the perceived image quality of dichoptic pairs as a function of eccentricity and image content. The experiment used the same general experimental setup described in Section 3.2.

Task and procedure. The experiment compared different dichoptic test stimuli against a binocular unaltered reference image in a 2AFC task. Dichoptic test stimuli considered combinations of eccentricity, σ, ϕ, and image content. Each trial first presented the binocular reference and two dichoptic test stimuli. Participants could freely toggle between the reference and two test stimuli, with a 500 ms mid-gray blank screen inserted between views to prevent direct comparisons. Pairwise comparisons were scheduled using ASAP [Mikhailiuk et al. 2021], which samples stimuli by maximizing information gain. This reduces the number of trials significantly compared to a full study design. During each trial, the participant indicated which of the test stimuli appeared closer to the reference and made selections with a standard keyboard. The experiment was divided into 15 min blocks with mandatory breaks between blocks due to the study duration (approximately one hour per participant). Each participant completed 419 trials.

Participants. A total of 31 participants (22 male, 9 female, aged 18-29 years) completed the experiment. Participants reported normal or corrected-to-normal vision. The study protocol was approved by the university's institutional review board (IRB). Participants gave informed consent before starting the experiment.

3.5 Results

We scaled the pairwise comparison responses to JODs using Thurstone's Case V model [Thurstone 2017], as implemented by the *pwcmp* algorithm [Perez-Ortiz and Mantiuk 2017]. Confidence intervals were estimated by simulating 500 bootstrap samples using

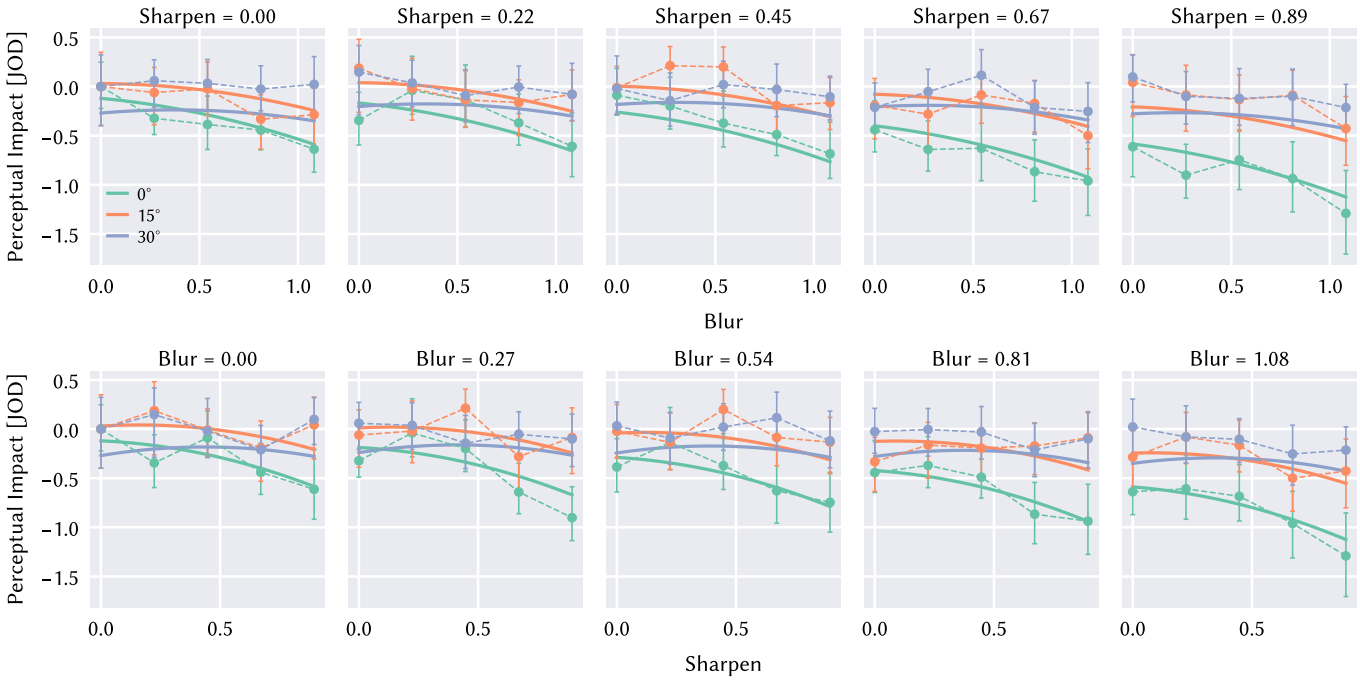


Fig. 5. *Main study results.* Plots visualizing the results of our perceptual study as a function of JOD. The plots show varying blur (top row), sharpening (bottom row), and eccentricity (color). The error on the datapoints represents the standard deviation, and the fitted quadratic functions are represented by solid lines.

the same *pwcmp* algorithm. The Kolmogorov-Smirnov test was run on the bootstrap distributions, where 64 of 75 conditions did not reject normality, suggesting that the normality assumptions are reasonable for most conditions.

We anchored JOD scores to the baseline condition (JOD=0 for $\sigma = 0$, $\phi = 0$) so that they reflected perceptual impact at the given eccentricity. We ran a three-way factorial analysis of variance (ANOVA) on anchored JOD scores across the main factors blur, sharpen, and eccentricity, including all two-way interactions. The analysis showed strong main effects of blur ($p = 5.44 \cdot 10^{-8}$, $F(4, 32) = 18.61$), sharpen ($p = 6.32 \cdot 10^{-8}$, $F(4, 32) = 18.35$), and eccentricity ($p = 5.77 \cdot 10^{-16}$, $F(2, 32) = 127.39$). The interaction effect between blur and sharpen was not significant ($p = 0.707$, $F(16, 32) = 0.77$). However, the interaction between blur and eccentricity ($p = 0.03$, $F(8, 32) = 2.45$), and between sharpen and eccentricity was significant ($p = 5.38 \cdot 10^{-5}$, $F(8, 32) = 6.45$), indicating that eccentricity modulates the impact of the blur–sharpen fusion (see Figure 5). These interaction effects are consistent with the known fall-off of visual acuity with eccentricity, creating non-uniform dichoptic effects across the visual field.

4 Computational Model

To interpolate our data, we fit 4D perceptual models that estimate perceptual impact (JOD scores) given blur (σ), sharpen (ϕ), and eccentricity (θ) as input (see Figure 6). We considered four parameterizations: linear, linear with pairwise interactions, exponential, and quadratic. Among these, the quadratic model $\mathcal{M}_{\text{quadratic}}(\cdot)$ yielded the best numerical fit ($R^2 = 0.82$). It is expressed as

$$\begin{aligned}\mathcal{M}_{\text{quadratic}}(\sigma, \phi, \theta) = &k_1 + k_2\sigma + k_3\phi + k_4\theta + k_5\sigma^2 \\ &+k_6\sigma\phi + k_7\sigma\theta + k_8\phi^2 + k_9\phi\theta + k_{10}\theta^2.\end{aligned} \tag{4}$$

The parameters, $k_i \in \mathbf{k}$, were fitted to the empirical results of the perceptual experiment (see Section 3.4). The optimized parameters are $\mathbf{k} = [-.12, -.189, -.109, .025, -.226, -.074, .012, -.459, .017, -.001]$.

Practical application of a perceptual model also relies on the ability to extrapolate beyond the tested regime. In our analysis, the linear model ($R^2 = 0.65$), $\mathcal{M}_{\text{linear}}(\cdot)$,

$$\mathcal{M}_{\text{linear}}(\sigma, \phi, \theta) = k_1 + k_2\sigma + k_3\phi + k_4\theta, \tag{5}$$

extrapolated most accurately to values outside the experimental parameter space. This was most pronounced when exploring higher eccentricities (cf. Section 5). The fitted parameters were

$$\mathbf{k} = [-.179, -.284, -.301, .016].$$

5 Application to Foveated Rendering

While our model offers fundamental insights into the neurological processing of stereoscopic information, it also enables practical applications in computer graphics. A direct application of our model is the optimization of foveated rendering. By rendering the foveated image using dichoptic blur–sharpen image processing, we argue that visual quality can be better preserved even when fewer samples are used. These two objectives of equal perceptual quality and lower sampling requirements are validated separately.

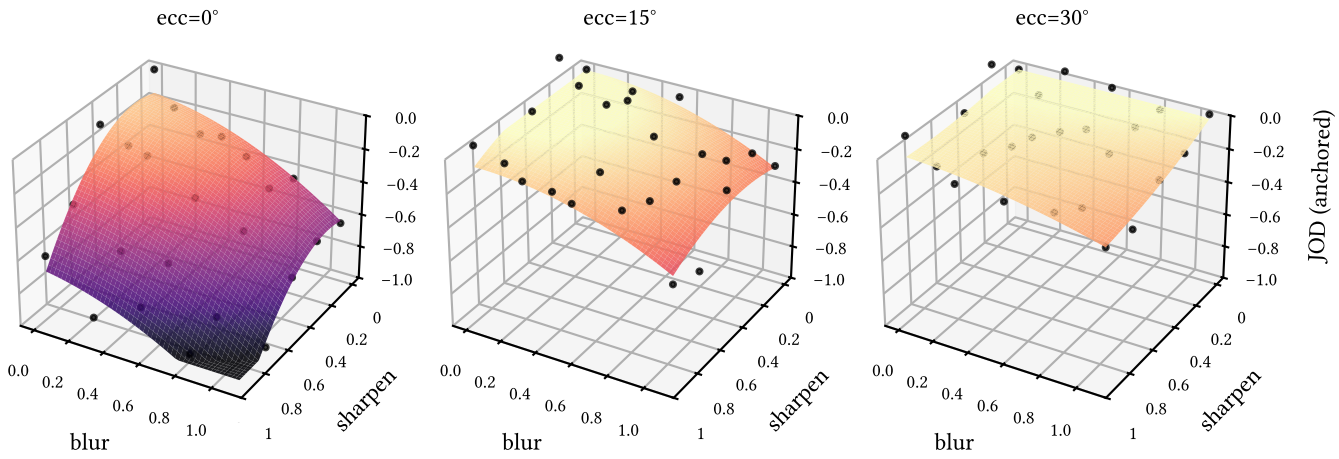


Fig. 6. *Perceptual model.* Visualization of the quadratic model surface for each eccentricity, respectively. The results show how sensitivity to dichoptic effects decreases with eccentricity.

5.1 Implementation & Sampling

To validate the quality objective, we implemented dichoptic foveated rendering through post-processing shaders in Unity. By applying blurring and sharpening in post-processing, we directly simulated the visual output of an undersampled-then-reconstructed frame, providing a rigorous testbed for evaluating our model's ability to preserve fidelity in high-resolution, dynamic environments.

To determine the optimal blur (σ) and sharpening (ϕ) parameters for a given pixel at eccentricity θ, we used the linear model $\mathcal{M}_{\text{linear}}(\cdot)$, as it was found to be more robust with respect to extrapolation (see Section 4). This property is essential for covering the full range of eccentricities in wide-FOV HMDs. Because the linear model yields a planar objective surface, the optimization space lacks local minima, precluding the need for iterative optimization solvers; instead, we adopted an analytical approach to sampling.

Our primary objective was to maximize the blur magnitude σ, as it directly influences the reduction in required samples and thus yields computational savings during rendering (see Section 5.5). In contrast, the computational overhead of sharpening is constant. As shown in Equation 3, sharpening is implemented via a single addition and multiplication. The $O(1)$ complexity renders its cost negligible relative to rendering costs (see Section 3.1), ensuring that the computational benefits gained from increased blur are not offset by additional pipeline complexity.

Our model revealed an inverse relationship between sharpening strength and blur magnitude. Since human sensitivity to high-frequency sharpening artifacts diminishes significantly at higher eccentricities, we hypothesized that sharpening gain could be decreased with eccentricity to permit greater peripheral blur. However, in the fovea, visual quality must be better preserved due to higher acuity and to prevent binocular rivalry. We therefore assigned the highest sharpening gain validated in our experiments ($\phi = 0.89$) to the fovea. This high sharpening minimized the required blur at the gaze center while still providing meaningful computational savings compared to standard rendering.

In our experiment, we tested two conditions: a sharpening gain that is constant across the visual field, and one that decreases linearly to the lowest validated sharpen value ($\phi = 0.22$) between 0° to 60° eccentricity. The corresponding blur values are derived from sampling the model at −1 JOD.

Recent stereo-aware rendering work has shown that exploiting inter-view coherence can improve stereo consistency while keeping computational overhead low [Philippi et al. 2025]. Our approach is designed to complement, rather than replace, existing foveated rendering frameworks. In this application, dichoptic sharpening operates exclusively on the low-pass information that remains after classic foveation has been applied. High-frequency information beyond the Nyquist limit of foveated rendering is not visible to the viewer and may even introduce aliasing. By only sharpening on the foveated image, the computational cost of dichoptic rendering never exceeds that of foveated rendering.

5.2 Baseline Calibration

Prior foveated rendering methods are usually calibrated by a central, fixed gaze point [Kergaßner et al. 2025; Patney et al. 2016; Tariq and Didyk 2024; Tariq et al. 2022; Walton et al. 2021]. In contrast, we apply a dynamic gaze setting, which requires recalibration to account for artifacts due to saccadic eye movements. We calibrated the baseline foveated rendering parameters following the formulation of Tariq et al. [2022],

$$\sigma = c \cdot \max(\theta - 8, 0). \tag{6}$$

To calibrate the foveation intensity, c, we conducted a 2AFC study (N=14, 9 male, 5 female, aged 20-31) with three 360° stereoscopic videos from the LIVE Database [Jin et al. 2021]. We chose diverse videos with natural content, people, and a street with cars. Each trial compared stimuli with foveated rendering applied with one of four intensities ($c \in [0, 1.5]$) to their unmodified reference. Each comparison was repeated three times per video, resulting in 36 trials per participant.

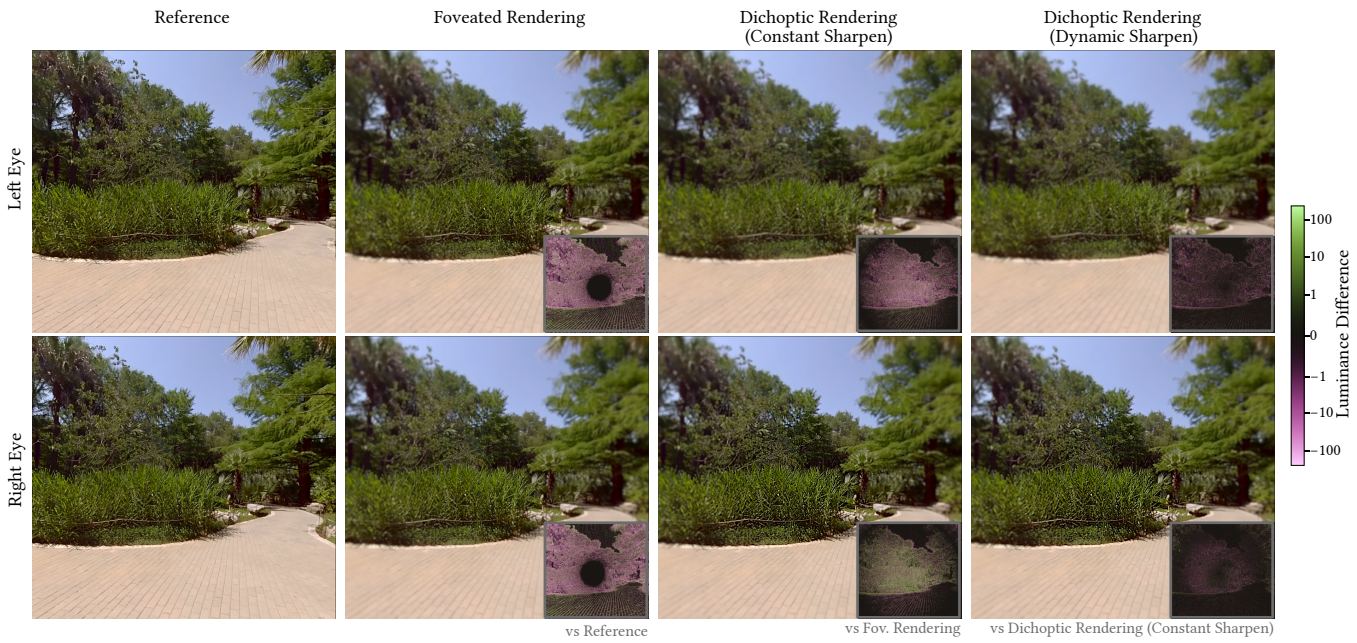


Fig. 7. *Application renderings.* Visual stimuli of reference, foveated rendering, and the two dichoptic rendering approaches as displayed in the validation study. The insets highlight luminance differences between methods for better visibility. ©The University of Texas at Austin.

We analyzed participants' accuracy (% of reference chosen) as a function of foveation intensity c and fit a logistic psychometric function to the data. Performance increased reliably with blur; the psychometric fit performed better than a constant model (bootstrap $p = 3.33 \cdot 10^{-4}$) and showed no evidence of poor fit (bootstrap deviance $p = 0.197$). We obtained the perceptual threshold (-1 JOD) from the fitted curve, which was $c = 0.11$ with a parametric-bootstrap 95% confidence interval of (0.09, 0.14). The data and fitted curve are visualized in the supplementary webpage.

5.3 Subjective Quality Evaluation

We conducted a within-subjects repeated-measures experiment to rate the perceptual quality of the baseline reference, the calibrated foveated rendering, and two conditions for dichoptic foveation on high-resolution stereoscopic 360° videos in VR (see Figure 7). The two dichoptic conditions applied constant and dynamic sharpening, respectively, as described in Section 5.1. Immersive video content allowed us to test the model's ability to generalize to full-screen, dynamic scenes.

Apparatus and stimuli. The Meta Quest Pro HMD was used. Stimuli consisted of ten 360° stereoscopic videos at a resolution of 7680 × 3840, obtained from the LIVE Database [Jin et al. 2021]. Each participant viewed all videos under all rendering modalities in randomized order.

Participants and procedure. A total of 14 participants (8 male, 5 female, 1 other, aged 18–26 years) completed the experiment. Similar to the main study, participants were screened for normal vision and compensated for their time. Each trial presented a ten-second 360° video in which participants were allowed to freely explore the scene.

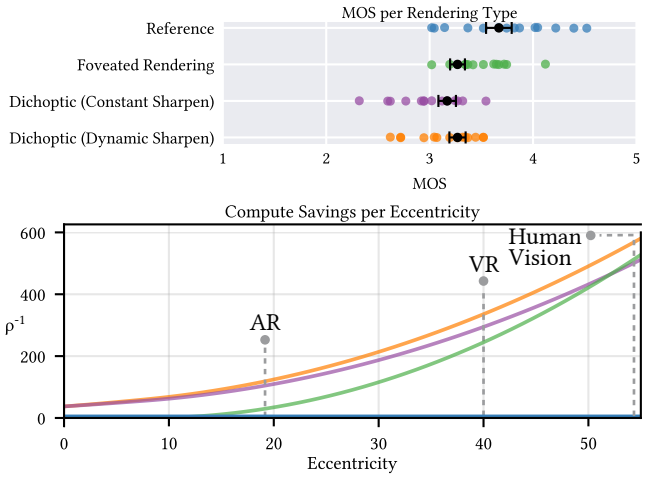


Fig. 8. *Validation result.* The top graph visualizes perceptual quality for the investigated techniques in MOS. Higher values indicate better quality. The dots depict subject-level ratings, while the weighted mean with standard deviation is shown in black. The bottom graph shows computational benefits measured by the inverse of sampling density (gain factor ρ^{-1}, y-axis) as a function of eccentricity (x-axis). Colors correlate with the methods of the top graph. The binocular FOV limits of common devices are also shown.

After each video finished playing, participants rated video content using a 5-point Likert scale which conformed to the ITU-T P.910 standard for video quality assessment [ITU-T 2023].

5.4 Results

The user ratings were scaled to mean opinion scores (MOS) using the Sureal method [Netflix 2016]. The results can be seen in Figure 8. The mean MOS values for dichoptic rendering with constant sharpening (3.16), dynamic sharpening (3.26), and standard foveated rendering (3.26) were similar, whereas the reference condition received a slightly higher score (3.66). To analyze significance of these results, we first tested the subject-level MOS for normality using Shapiro-Wilk ($W = 0.9856$, $p = 0.738$), which affirmed this assumption. A repeated-measures ANOVA was then conducted to determine if the difference in MOS between the render types was significant ($p < 0.0001$, $F = 9.59$). Given the confirmation of the significance hypothesis, we followed with post-hoc pairwise comparisons using t-tests with Bonferroni correction. We found that only dichoptic foveated rendering with constant sharpening was significantly different from the reference ($p = 0.015$). Both dynamic sharpening and foveated rendering showed no significant difference with the reference video in terms of visual quality. This result suggests that our model generalizes well beyond simple image patch stimuli to dynamic, full-screen content.

5.5 Computational Benefit

We demonstrated that dichoptic foveation achieves similar visual quality compared to full-resolution rendering, despite significant application of blur. The computational benefit of our approach is measured using a theoretical, hardware-agnostic framework based on blur kernel strength, σ. We consider foveated raytracing, where the primary objective is to maximize the local sample spacing $\Delta x(\theta)$ relative to retinal eccentricity θ. In this framework, any $\Delta x(\theta) \geq 1$ represents a reduction in local sampling density compared to full-resolution rendering; therefore, larger values of Δx directly correspond to greater computational savings.

First, we define the relationship between blur σ and the spatial sampling interval Δx. In alignment with standard practices in the rendering literature, we assume that sparse sampling followed by an ideal reconstruction filter behaves as a low-pass operation [Albert et al. 2017; Hoffman et al. 2018]. Specifically, we approximate the resulting point-spread function (PSF) as a Gaussian kernel. The frequency domain expression of the PSF [Goodman 2005] with standard deviation σ is defined as

$$H(f) = e^{-2\pi^2\sigma^2 f^2}. \tag{7}$$

In rendering, the practical limit of frequency f is given by the Nyquist theorem where the spatial frequency drops below 0.5 cycles per pixel. Since a Gaussian has no true cutoff, we set $f_c = \frac{1}{2\pi\sigma}$ as the natural scale frequency, where the Gaussian's energy is concentrated, given its spatial width, σ. To reproduce this blur without additional loss, the sampling rate must satisfy

$$f_s \geq 2f_c = \frac{1}{\pi\sigma}. \tag{8}$$

Given that a frequency in pixel space is naturally related to the sampling spacing by $f_s = \frac{1}{\Delta x}$, we then have

$$\Delta x = \pi\sigma. \tag{9}$$

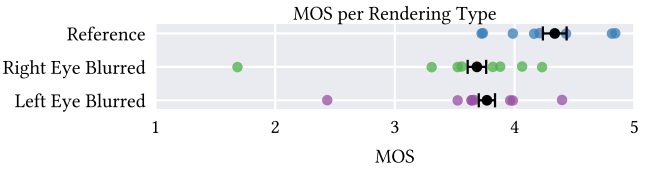


Fig. 9. *Ablation result.* The jitter plot compares the perceived quality (MOS) of reference videos to uniform dichoptic rendering. For the analysis, we split the results into stimuli that blurred the right and left eye respectively. Greater MOS indicates better quality.

For 2D images, the sample density relative to full resolution is

$$\rho = \frac{1}{\Delta x^2} = \frac{1}{(\pi\sigma)^2}. \tag{10}$$

Figure 8 shows the theoretical computational savings that can be achieved by foveated and dichoptic rendering relative to uniform sampling as a function of eccentricity.

Computational savings for the full display are computed by aggregating local eccentricity-dependent savings across the entire FOV. We assume central gaze and approximate the display with a circular area. With these assumptions, savings are integrated across the spherical domain. Because the solid angle of each concentric ring increases with eccentricity, the base saving function must be weighted by the area of the corresponding spherical cap. We express the total theoretical computational cost, C, using the angular sampling formulation,

$$C = 2\pi \int_0^{\theta_{\max}} f(\theta)\sin(\theta)\, d\theta. \tag{11}$$

In this expression, $f(\theta)$ represents the local sampling density at eccentricity θ. We evaluate the cost for both standard foveated rendering (C_{fr}) and our proposed dichoptic foveation (C_{dr}) to derive a ratio of their respective workloads, $\frac{C_{\text{dr}}}{C_{\text{fr}}}$. For the VR HMD that was used in our study, the analysis yields a benefit of approximately 1.9, suggesting that dichoptic foveation requires only about half the samples of traditional foveated rendering to maintain the same perceived quality. For smaller stereoscopic displays, such as those used in augmented reality (AR), dichoptic foveation yielded theoretical savings of up to 11×. Additional device-specific results are provided in the supplementary material.

5.6 Ablation: Dichoptic Rendering Without Eyetracking

Many modern devices lack eye tracking. As a simple adaptation, we evaluated the application of dichoptic parameters derived at zero eccentricity uniformly across the image, eliminating the need for gaze-contingent rendering. Our results showed that even in the contrast-sensitive fovea, substantial dichoptic modulation was tolerated. Under uniform dichoptic rendering, theoretical sampling requirements were reduced by approximately 97% relative to native-resolution rendering.

To validate the perceptual quality of this approach, we conducted a small-scale ablation study (9 participants, 6 male, 3 female, aged 19–26 years), where participants rated the video quality of the uniform dichoptic rendering method against the unaltered reference. Similar

to Section 5.3, users conducted a Likert-based rating task across ten 360° videos. Results are shown in Figure 9. To explore potential ocular dominance effects, we analyze our results in two conditions: blurring the left and right eye respectively. A repeated-measures ANOVA indicated a significant effect of condition ($p = 0.003$). Post-hoc comparisons showed that when the left eye was blurred, the dichoptic condition ($MOS = 3.63$) was not significantly different from the reference ($MOS = 4.19$) after Bonferroni correction ($p = 0.04$). These results indicate that uniform dichoptic rendering can indeed achieve a perceptual quality comparable to native rendering. In contrast, when the right eye was blurred, results indicated a significant difference from the reference ($p = 0.06$). This asymmetry may be related to eye dominance, as nearly 70% of people have been found to be right-eye dominant [Seijas et al. 2005]. Preserving sharpness in the dominant eye may therefore yield higher perceived quality, highlighting an important direction for future work.

6 Limitations & Future Work

The objective of this paper is to define a general approach to dichoptic foveation that applies to a typical observer. As such, we computed average JOD scores across all users in our psychophysical study; individual differences in sensitivity could lead to variation in perceived quality or the potential for binocular rivalry. A per-observer calibration procedure could therefore improve subject-level performance and represents a promising direction for future work.

Our model does not explicitly account for individual eye dominance or dynamic effects like change blindness [Groth et al. 2023b], which may offer additional opportunities for optimization. In this study, the assignment of blur and sharpening filters was randomized across the left and right eyes to ensure the resulting computational model remains broadly applicable. However, as suggested by our ablation study, the effectiveness of dichoptic foveation may depend on which eye receives the sharper signal. Future work could incorporate eye-dominance-aware rendering strategies to enable more personalized dichoptic foveation.

All experiments were conducted using a commercial HMD (Meta Quest Pro), which introduces limitations related to eye-tracking latency, display resolution, optical distortions, and related factors. Such limitations may introduce auxiliary artifacts that are not directly related to the phenomena studied in this work. Nevertheless, our model is calibrated under conditions that are representative of a realistic VR viewing scenario.

Dichoptic effects, furthermore, require binocular vision. However, typical VR displays have a monocular region in the outer edges of the visual field, where there is no overlap between the two eyes. Our technique, however, only works in this overlapping region; headsets with very high FOV may have a smaller overlapping region relative to the overall display area, leading to less savings. In this high-FOV monocular regime, normal foveated rendering can be applied to supplement dichoptic foveation.

Our application-focused evaluation assessed perceptual quality relative to a baseline reference condition. Therefore, long-term effects like discomfort due to eye strain or rivalry are unexplored, and should be investigated as future work when used in daily use cases.

The aim of our computational analysis was to be system-agnostic. However, real-world raytracing systems come with system-level optimizations that may lead to deviations when compared with our theoretical savings. Future work should therefore implement our approach on such a system to evaluate actual hardware-specific savings. Other techniques, such as those targeting bandwidth-reduction, could supplement our method for real-time stereo transmission [Denes et al. 2019].

Additional future directions include exploring a wider range of eccentricities or blur–sharpen parameters, as well as additional interaction effects due to contrast or color. Furthermore, alternative display modalities like AR may realize even more gains with our technique, due to low FOV (see Figure 8) and reduced contrasts that could mask distortions.

7 Conclusion

This work introduces *dichoptic foveation*, a paradigm that leverages binocular fusion to optimize stereoscopic rendering. Psychophysical measurements show that appropriate combinations of blur and sharpening, parameterized by retinal eccentricity, preserve or improve perceived quality relative to symmetric presentations. The resulting 4D perceptual model provides a compact and interpretable quality predictor that directly links rendering parameters to human perception. Applied to foveated rendering in VR, the model enables ~50% sampling-rate reductions compared to traditional foveated rendering, while maintaining perceptually equivalent quality to full-resolution renderings.

Acknowledgments

This work is partially supported by National Science Foundation grant #2225861, and the DARPA Intrinsic Cognitive Security (ICS) program.

References

Rachel Albert, Anjul Patney, David Luebke, and Joohwan Kim. 2017. Latency Requirements for Foveated Rendering in Virtual Reality. *ACM Trans. Appl. Percept.* 14, 4, Article 25 (Sept. 2017), 13 pages. doi:10.1145/3127589

Derek H Arnold, Philip M Grove, and Thomas SA Wallis. 2007. Staying focused: A functional account of perceptual suppression during binocular rivalry. *Journal of Vision* 7, 7 (2007), 7–7.

Maliha Ashraf, Alexandre Chapiro, and Rafał K Mantiuk. 2025. Resolution limit of the eye—how many pixels can we see? *Nature Communications* 16, 1 (2025), 9086.

Randolph Blake and Hugh Wilson. 2011. Binocular vision. *Vision research* 51, 7 (2011), 754–770.

Jianhua Cang, Jieming Fu, and Seiji Tanabe. 2023. Neural circuits for binocular vision: Ocular dominance, interocular matching, and disparity selectivity. *Frontiers in Neural Circuits* 17 (2023), 1084027.

Kenneth Chen, Thomas Wan, Nathan Matsuda, Ajit Ninan, Alexandre Chapiro, and Qi Sun. 2024. PEA-PODs: Perceptual Evaluation of Algorithms for Power Optimization in XR Displays. *ACM Trans. Graph.* 43, 4, Article 67 (jul 2024), 17 pages. doi:10.1145/3658126

Michael J Collins and Andrew Goode. 1994. Interocular blur suppression and monovision. *Acta ophthalmologica* 72, 3 (1994), 376–380.

Gyorgy Denes, Kuba Maruszczyk, George Ash, and Rafał K Mantiuk. 2019. Temporal Resolution Multiplexing: Exploiting the limitations of spatio-temporal vision for more efficient VR rendering. *IEEE transactions on visualization and computer graphics* 25, 5 (2019), 2072–2082.

Nianchen Deng, Zhenyi He, Jiannan Ye, Budmonde Duinkharjav, Praneeth Chakravarthula, Xubo Yang, and Qi Sun. 2022. FoV-NeRF: Foveated Neural Radiance Fields for Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics* 28, 11 (2022), 3854–3864. doi:10.1109/TVCG.2022.3203102

Miroslav Dostalek, Jan Hejda, Karel Fliegel, Michaela Duchackova, Ladislav Dusek, Jiri Hozman, Tomas Lukes, and Rudolf Autrata. 2019. Influence of artificially generated

interocular blur difference on fusion stability under vergence stress. *Journal of Eye Movement Research* 12, 4 (2019), 27.

Budmonde Duinkharjav, Kenneth Chen, Abhishek Tyagi, Jiayi He, Yuhao Zhu, and Qi Sun. 2022. Color-Perception-Guided Display Power Reduction for Virtual Reality. *ACM Trans. Graph.* 41, 6, Article 210 (Nov. 2022), 16 pages. doi:10.1145/3550454.3555473

Runze Fan, Xuehuai Shi, Kangyu Wang, Qixiang Ma, and Lili Wang. 2024. Scene-aware foveated rendering. *IEEE Transactions on Visualization and Computer Graphics* (2024).

Sid Ahmed Fezza and Mohamed-Chaker Larabi. 2017. Perceptually Driven Nonuniform Asymmetric Coding of Stereoscopic 3D Video. *IEEE Transactions on Circuits and Systems for Video Technology* 27, 10 (2017), 2231–2245. doi:10.1109/TCSVT.2016.2583618

Mark A Georgeson and Stuart A Wallis. 2014. Binocular fusion, suppression and diplopia for blurred edges. *Ophthalmic and Physiological Optics* 34, 2 (2014), 163–185.

Joseph W Goodman. 2005. *Introduction to Fourier optics*. Roberts and Company publishers.

Colin Groth, Sascha Fricke, Susana Castillo, and Marcus Magnor. 2023a. Wavelet-Based Fast Decoding of 360°Videos. *IEEE Transactions on Visualization and Computer Graphics (TVCG, Proc. IEEE VR)* (2023), 1–9. doi:10.1109/{TVCG}.2023.3247080

Colin Groth, Marcus Magnor, Steve Grogorick, Martin Eisemann, and Piotr Didyk. 2024. Cybersickness Reduction via Gaze-Contingent Image Deformation. *ACM Transactions on Graphics (Proc. of Siggraph)* 43, 4 (2024), 1–14. doi:10.1145/3658138

Colin Groth, Timon Scholz, Susana Castillo, Jan-Philipp Tauscher, and Marcus Magnor. 2023b. Instant Hand Redirection in Virtual Reality Through Electrical Muscle Stimulation-Triggered Eye Blinks. In *Proceedings of the 29th ACM Symposium on Virtual Reality Software and Technology*. 1–11.

Colin Groth, Jan-Philipp Tauscher, Nikkel Heesen, Susana Castillo, and Marcus Magnor. 2021a. Visual techniques to reduce cybersickness in virtual reality. In *2021 IEEE conference on virtual reality and 3D user interfaces abstracts and workshops (VRW)*. 486–487.

Colin Groth, Jan-Philipp Tauscher, Nikkel Heesen, Steve Grogorick, Susana Castillo, and Marcus Magnor. 2021b. Mitigation of Cybersickness in Immersive 360°Videos. In *IEEE Virtual Reality Workshop on Immersive Sickness Prevention (WISP)*. 169–177. doi:10.1109/{VRW}52623.2021.00039

Brian Guenter, Mark Finch, Steven Drucker, Desney Tan, and John Snyder. 2012. Foveated 3D graphics. *ACM Trans. Graph.* 31, 6, Article 164 (Nov. 2012), 10 pages. doi:10.1145/2366145.2366183

David Hoffman, Zoe Meraz, and Eric Turner. 2018. Limits of peripheral acuity and implications for VR system design. *Journal of the Society for Information Display* 26, 8 (2018), 483–495. arXiv:https://sid.onlinelibrary.wiley.com/doi/pdf/10.1002/jsid.730 doi:10.1002/jsid.730

David M Hoffman and Martin S Banks. 2010. Focus information is used to interpret binocular images. *Journal of vision* 10, 5 (2010), 13–13.

ITU-T. 2023. *Recommendation ITU-T P.910: Subjective video quality assessment methods for multimedia applications*. ITU-T Recommendation P.910. International Telecommunication Union. https://www.itu.int/rec/T-REC-P.910

Rong Jiang and Ming Meng. 2023. Integration and suppression interact in binocular vision. *Journal of Vision* 23, 10 (2023), 17–17.

Daniel Jiménez-Navarro, Colin Groth, Xi Peng, Jorge Pina, Qi Sun, Praneeth Chakravarthula, Karol Myszkowski, Hans-Peter Seidel, and Ana Serrano. 2026. Overdriving Visual Depth Perception via Sound Modulation in VR. *IEEE Transactions on Visualization and Computer Graphics* (2026).

Yize Jin, Meixu Chen, Todd Goodall, Anjul Patney, and Alan C. Bovik. 2021. Subjective and Objective Quality Assessment of 2D and 3D Foveated Video Compression in Virtual Reality. *Trans. Img. Proc.* 30 (Jan. 2021), 5905–5919. doi:10.1109/TIP.2021.3087322

Anton S. Kaplanyan, Anton Sochenov, Thomas Leimkühler, Mikhail Okunev, Todd Goodall, and Gizem Rufo. 2019. DeepFovea: neural reconstruction for foveated rendering and video compression using learned statistics of natural videos. *ACM Trans. Graph.* 38, 6, Article 212 (Nov. 2019), 13 pages. doi:10.1145/3355089.3356557

Sophie Kergaßner, Taimoor Tariq, and Piotr Didyk. 2025. Towards Understanding Depth Perception in Foveated Rendering. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '25)*. Association for Computing Machinery, New York, NY, USA, Article 42, 9 pages. doi:10.1145/3721238.3730609

Brooke Krajancich, Petr Kellnhofer, and Gordon Wetzstein. 2023. Towards attention–aware foveated rendering. *ACM Transactions on Graphics (TOG)* 42, 4 (2023), 1–10.

Jérôme Kudnick, Daniel T. Mayer, Colin Groth, Bipul Mohanto, Ralf Dörner, and Martin Weier. 2026. RippleVision: Unobtrusive Gaze-Dependent Guidance via Directed Wave Motion in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics (TVCG, Proc. IEEE VR)* (2026).

Jérôme Kudnick, Martin Weier, Colin Groth, Biying Fu, and Robin Horst. 2025. WarpVision: Using Spatial Curvature to Guide Attention in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics (TVCG, Proc. IEEE ISMAR)* (2025).

Phillip CF Law, Bryan K Paton, Richard H Thomson, Guang B Liu, Steven M Miller, and Trung T Ngo. 2013. Dichoptic viewing methods for binocular rivalry research: prospects for large-scale clinical and genetic studies. *Twin Research and Human Genetics* 16, 6 (2013), 1033–1078.

Haechan Lee, Namil Kim, Hoe Sung Ryu, and Uijong Ju. 2025. When One Eye Sees Less: Uncovering Perceptual Quality Thresholds of Asymmetric Quality Degradation in 4K XR Displays. In *2025 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 998–1008. doi:10.1109/ISMAR67309.2025.00107

Wei Hau Lew, Scott B Stevenson, and Daniel R Coates. 2021. Stimulus dependence of interocular suppression. *Scientific Reports* 11, 1 (2021), 9309.

Weikai Lin, Yu Feng, and Yuhao Zhu. 2025. MetaSapiens: Real-Time Neural Rendering with Efficiency-Aware Pruning and Accelerated Foveated Rendering. In *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 1* (Rotterdam, Netherlands) *(ASPLOS '25)*. Association for Computing Machinery, New York, NY, USA, 669–682. doi:10.1145/3669940.3707227

Xiaoxu Meng, Ruofei Du, and Amitabh Varshney. 2020. Eye-dominance-guided foveated rendering. *IEEE transactions on visualization and computer graphics* 26, 5 (2020), 1972–1980.

Xiaoxu Meng, Ruofei Du, Matthias Zwicker, and Amitabh Varshney. 2018. Kernel Foveated Rendering. *Proc. ACM Comput. Graph. Interact. Tech.* 1, 1, Article 5 (July 2018), 20 pages. doi:10.1145/3203199

Aliaksei Mikhailiuk, Clifford Wilmot, Maria Perez-Ortiz, Dingcheng Yue, and Rafal Mantiuk. 2021. Active Sampling for Pairwise Comparisons via Approximate Message Passing and Information Gain Maximization. In *2020 IEEE International Conference on Pattern Recognition (ICPR)*.

Bipul Mohanto, Martin Weier, Colin Groth, and Oliver Staadt. 2025. A Unified Real-Time Spatial and Chromatic Foveated Rendering in VR. In *GI Workshop on Virtual and Augmented Reality*. 10–18420.

Netflix. 2016. SUREAL: Subjective quality scores recovery from noisy measurements. https://github.com/Netflix/sureal. GitHub repository.

Anjul Patney, Marco Salvi, Joohwan Kim, Anton Kaplanyan, Chris Wyman, Nir Benty, David Luebke, and Aaron Lefohn. 2016. Towards foveated rendering for gaze-tracked virtual reality. *ACM Transactions On Graphics (TOG)* 35, 6 (2016), 1–12.

Maria Perez-Ortiz and Rafal K. Mantiuk. 2017. A practical guide and software for analysing pairwise comparison experiments. arXiv:1712.03686 [stat.AP] https://arxiv.org/abs/1712.03686

Henrik Philippi, Henrik Wann Jensen, and Jeppe Revall Frisvad. 2025. Stable Sample Caching for Interactive Stereoscopic Ray Tracing. In *33rd Pacific Conference on Computer Graphics and Applications*. Eurographics Association.

Cheng Stella Qian, Sam Ling, and Jan W Brascamp. 2019. Dichoptic vision in the absence of attention: neither fusion nor rivalry. *Scientific Reports* 9, 1 (2019), 12904.

Guillaume Riesen, Anthony M Norcia, and Justin L Gardner. 2019. Humans perceive binocular rivalry and fusion in a tristable dynamic state. *Journal of Neuroscience* 39, 43 (2019), 8527–8537.

Clifton Schor, Lori Landsman, and Paul Erickson. 1987. Ocular dominance and the interocular suppression of blur in monovision. *Optometry and Vision Science* 64, 10 (1987), 723–730.

Olga Seijas, Patricia Gómez de Liaño, Raquel Gómez de Liaño, Colin J. Roberts, Elkin Piedrahita, and Eugenio Diaz. 2005. Ocular dominance: eye preference within the context of binocular functions. *Graefe's Archive for Clinical and Experimental Ophthalmology* 243, 9 (2005), 926–932. doi:10.1007/s00417-005-1128-7

Feng Shao, Gangyi Jiang, Xu Wang, Mei Yu, and Ken Chen. 2010. Stereoscopic video coding with asymmetric luminance and chrominance qualities. *IEEE Transactions on Consumer Electronics* 56, 4 (2010), 2460–2468.

George Sperling. 1970. Binocular vision: A physical and a neural theory. *The American Journal of Psychology* (1970), 461–534.

Michael Stengel, Steve Grogorick, Martin Eisemann, and Marcus Magnor. 2016. Adaptive image-space sampling for gaze-contingent real-time rendering. In *Computer graphics forum*, Vol. 35. Wiley Online Library, 129–139.

David Stidwill and Robert Fletcher. 2017. *Normal binocular vision: Theory, investigation and practical aspects*. John Wiley & Sons.

Taimoor Tariq and Piotr Didyk. 2024. Towards motion metamers for foveated rendering. *ACM Transactions on Graphics (TOG)* 43, 4 (2024), 1–10.

Taimoor Tariq, Cara Tursun, and Piotr Didyk. 2022. Noise-based enhancement for foveated rendering. *ACM Transactions on Graphics (TOG)* 41, 4 (2022), 1–14.

Louis L Thurstone. 2017. A law of comparative judgment. In *Scaling*. Routledge, 81–92.

David R. Walton, Rafael Kuffner Dos Anjos, Sebastian Friston, David Swapp, Kaan Akşit, Anthony Steed, and Tobias Ritschel. 2021. Beyond blur: real-time ventral metamers for foveated rendering. *ACM Trans. Graph.* 40, 4, Article 48 (July 2021), 14 pages. doi:10.1145/3450626.3459943

Minqi Wang and Emily A. Cooper. 2021. A Re-examination of Dichoptic Tone Mapping. *ACM Trans. Graph.* 40, 2, Article 13 (April 2021), 15 pages. doi:10.1145/3443702

Minqi Wang, Jian Ding, Dennis M Levi, and Emily A Cooper. 2023a. The effect of interocular contrast differences on the appearance of augmented reality imagery. *ACM Transactions on Applied Perception* 21, 1 (2023), 1–23.

Zhimin Wang, Xiangyuan Gu, and Feng Lu. 2023b. Deamp: Dominant-eye-aware foveated rendering with multi-parameter optimization. In *2023 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 632–641.

Zhimin Wang, Xiangyuan Gu, and Feng Lu. 2025. Dominant-Eye-Aware Asymmetric Foveated Rendering for Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics* 31, 10 (2025), 9225–9236. doi:10.1109/TVCG.2025.3593899

Felix A. Wichmann and Nicholas J. Hill. 2001. The psychometric function: I. Fitting, sampling, and goodness of fit. *Perception & Psychophysics* 63, 8 (Nov. 2001), 1293–1313. doi:10.3758/BF03194544

Yan Zhang, Keyao You, Xiaodan Hu, Hangyu Zhou, Kiyoshi Kiyokawa, and Xubo Yang. 2024. Retinotopic Foveated Rendering. In *2024 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*. 903–912. doi:10.1109/VR58804.2024.00109

Fangcheng Zhong, George Alex Koulieris, George Drettakis, Martin S Banks, Mathieu Chambe, Frédo Durand, and Rafał K Mantiuk. 2019. DiCE: dichoptic contrast enhancement for VR and stereo displays. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–13.