# Towards Understanding Depth Perception in Foveated Rendering

SOPHIE KERGASSNER, Università della Svizzera italiana, Switzerland TAIMOOR TARIQ, Università della Svizzera italiana, Switzerland PIOTR DIDYK, Università della Svizzera italiana, Switzerland



Fig. 1. We design a broadband stimulus to accurately estimate disparity thresholds across various eccentricities and blur levels. We fit a custom surface through the retrieved threshold estimations. It shows the lowest resolvable disparity *T* as a continuous function of eccentricity  $\theta$  and Gaussian filter sigma  $\sigma$ . The green line denotes the minimum *T*-value per eccentricity  $\theta$ , encoding the amount of blur that can be induced by foveated rendering without degrading stereoscopic depth perception. We validate our findings in a validation study, by showing participants unaltered renderings vs. strongly foveated renderings.

The true vision for real-time virtual and augmented reality is reproducing our visual reality in its entirety on immersive displays. To this end, foveated rendering leverages the limitations of spatial acuity in human peripheral vision to allocate computational resources to the fovea while reducing quality in the periphery. Such methods are often derived from studies on the spatial resolution of the human visual system and its ability to perceive blur in the periphery, enabling the potential for high spatial quality in real-time. However, the effects of blur on other visual cues that depend on luminance contrast, such as depth, remain largely unexplored. It is critical to understand this interplay, as accurate depth representation is a fundamental aspect of visual realism. In this paper, we present the first evaluation exploring the effects of foveated rendering on stereoscopic depth perception. We design a psychovisual experiment to quantitatively study the effects of peripheral blur on depth perception. Our analysis demonstrates that stereoscopic acuity remains unaffected (or even improves) by high levels of peripheral blur. Based on our studies, we derive a simple perceptual model that determines the amount of foveation that does not affect stereoacuity. Furthermore, we analyze the model in the context of common foveation practices reported in literature. The findings indicate that foveated rendering does not impact stereoscopic depth perception, and stereoacuity remains unaffected with up to 2× stronger foveation than commonly used. Finally, we conduct a validation experiment and show that our findings hold for complex natural stimuli.

Authors' Contact Information: Sophie Kergaßner, sophie.kergassner@usi.ch, Università della Svizzera italiana, Switzerland; Taimoor Tariq, taimoor.tariq@usi.ch, Università della Svizzera italiana, Switzerland; Piotr Didyk, piotr.didyk@usi.ch, Università della Svizzera italiana, Switzerland.

#### 

This work is licensed under a Creative Commons Attribution 4.0 International License. *SIGGRAPH Conference Papers '25, Vancouver, BC, Canada* © 2025 Copyright held by the owner/author(s). ACM ISBN 979-8-4007-1540-2/2025/08 https://doi.org/10.1145/3721238.3730609 CCS Concepts: • Computing methodologies  $\rightarrow$  Perception; Virtual reality.

Additional Key Words and Phrases: Human Perception, Stereoacuity, Perceptual Model, Foveated Rendering

#### **ACM Reference Format:**

Sophie Kergaßner, Taimoor Tariq, and Piotr Didyk. 2025. Towards Understanding Depth Perception in Foveated Rendering. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '25), August 10–14, 2025, Vancouver, BC, Canada.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3721238. 3730609

#### 1 Introduction

The human visual system (HVS) exhibits a non-uniform sensitivity towards different visual cues across the visual field. While visual acuity, shape discrimination, and sensitivity to color and depth corrugations are highest in the central fovea, there is a substantial decline in sensitivity to these factors with increasing eccentricity [Abramov et al. 1991; Baldwin et al. 2016; Moreland and Cruz 1959; Prince and Rogers 1998; Rosenholtz 2016; Whitaker et al. 1993]. This is the prime motivation behind foveated rendering techniques, that allocate computational resources to the fovea while reducing fidelity in the periphery. With the advent of gaze-tracking systems in augmented reality (AR) and virtual reality (VR), foveated rendering has emerged as a key enabler of real-time high-quality rendering for virtual realities. Consequently, it is being incorporated by graphics APIs, such as Unity and Nvidia VRS, and used by commercial AR and VR devices, such as Meta Quest and Apple Vision Pro.

The basis for foveated rendering methods are studies and psychophysical models of peripheral vision [Strasburger et al. 2011], which form the foundation for perceptually optimized rendering methods [Guenter et al. 2012; Murphy and Duchowski 2001; Patney et al. 2016; Surace et al. 2023; Tursun et al. 2019]. However, while they focus on different aspects of spatial quality, they neglect the potential influence of foveated rendering on visual cues such as motion [Tariq and Didyk 2024] and depth [Sun et al. 2019]. Consequently, the effects of foveated rendering on more general human perception, which governs the interaction and navigation in three-dimensional environments, remain largely unexplored. This is particularly relevant given the availability and importance of depth cues in novel displays, which inherently utilize foveated rendering.

In this work, we address the problem of reproducing and preserving depth cues on novel displays. While innovative display designs have already highlighted their importance [Akşit et al. 2017; Akeley et al. 2004; Narain et al. 2015; Qin et al. 2023], our work addresses the problem from the rendering perspective. We take first steps towards understanding how foveated rendering influences stereopsis, one of the most fundamental depth cues [Cutting and Vishton 1995]. While past research has already demonstrated a decrease in stereoacuity due to eccentricity [Prince and Rogers 1998], as well as the significant impact of blur on stereopsis [Cormack et al. 1991], these effects have not yet been investigated in conjunction, as occurring in foveated rendering. To close this gap, we design a broadband stimulus that serves as a basis for our psychophysical experiment (Section 3) investigating the effect of blur intensity on stereopsis at different eccentricities. Based on the measurements, we derive a model (Section 4) that describes the maximum blur intensity which does not compromise stereopsis. We discuss the findings in the context of common practices in foveated rendering literature. Our findings suggest that stereopsis in peripheral vision is remarkably resilient to loss of spatial resolution, offering new possibilities for resource-efficient rendering strategies without sacrificing spatial awareness. Finally, we validate our findings by applying our model to realistic scenes (Section 5).

#### 2 Background and Previous Work

Decoding human vision has received the efforts of many. In this work, we focus in particular on the perception of stereoscopic depth, and the effects of peripheral blur on stereoacuity.

## 2.1 Peripheral Vision and Foveated Rendering

Peripheral vision exhibits significantly reduced spatial resolution compared to foveal vision [Rosenholtz 2016]. This is primarily due to the lower density of cone photoreceptors and ganglion cells in the retina's periphery [Curcio et al. 1990]. While the fovea supports high-acuity tasks, the periphery excels at motion detection [Mckee and Nakayama 1984] and spatial learning [Alfano and Michel 1990; Fortenbaugh et al. 2008]. Foveated rendering utilizes these perceptual limitations by reducing rendering quality in peripheral regions. While researchers initially focused on reducing spatial resolution in the periphery [Guenter et al. 2012], subsequent research investigated pre- and post-processing steps which allowed for more aggressive foveation. Patney et al. [2016] post-applied contrast enhancement, allowing higher blur rates. Similarly, Tariq et al. [2022] recovered detectable frequencies in a cost-efficient post-processing step to re-enhance foveated content. Tursun et al. [2019] modeled the masking effects of the underlying content, enabling contentaware foveation. Surace et al. [2023] examined substantially reduced geometric complexity in the periphery. Another approach featured inducing mental tasks in the fovea to allow for an unnoticed increase in foveation. [Krajancich et al. 2023]. However, recent work on motion perception revealed a damage of speed estimations induced by foveation and resolved this by recovering spatio-temporal energy in a post-processing step [Tariq and Didyk 2024].

#### 2.2 Depth Perception

To navigate 3D space, humans rely on a number of different depth cues that convey either relative or absolute distance information [Cutting and Vishton 1995; Sweet and Kaiser 2012]. Cues such as accommodation, convergence, or familiar size convey information about the absolute distance between the observer and an object. Other cues provide solely relative information of ordinal or metric kind, allowing for the estimation of relative depth differences between objects. For example, occlusion conveys only ordinal information, whereas cues such as stereoscopic disparity or relative size provide metric information regarding the magnitude of depth differences.

With the advent of stereoscopic displays, the recreation of binocular vision became feasible. This enabled the display of stereo-

scopic depth cues such as binocular disparity, resulting in stereopsis. Stereopsis is created by perceiving a difference in vergence angles between a focus point  $p_f$  and a reference point  $p_r$ , which naturally occurs when those points are at different depths with respect to the eyes



Fig. 2. Different vergence angles occurring for different points.

(Figure 2). The vergence angles can be described by  $D(p) = \angle LpR$ , denoting the angles between the two locations of the eyes in space (*L* and *R*), and the respective points  $p_f$  and  $p_r$ . Stereopsis is quantified by the disparity d [°], denoting the difference of these vergence angles:

$$d = D(p_1) - D(p_2).$$
 (1)

The sign of *d* denotes whether objects are relatively closer, or farther. In case both vergence angles are the same, d = 0 and no disparity is present. This can happen when both points are at the same position in space, the image is rendered with an interpupillary distance (IPD) of 0mm, or both points lie on the horopter [Blakemore 1970].

Large population studies found average disparity thresholds of 30-60'' disparity, with lowest thresholds as low as 2'' [Bosten et al. 2015; Coutant and Westheimer 1993]. However, stereoacuity depends strongly on various properties exhibited by the luminance pattern. Siderov and Harwerth [1995] studied band-limited luminance patterns, finding peak sensitivities at 8 cycles per degree (cpd) in the fovea, and 2cpd at  $10^{\circ}$  eccentricity. Didyk et al. [2012] manipulated disparity based on the underlying luminance contrast. Their perceptual model is based on a finding of Marr and Poggio [1979], who claim that disparity is mostly supported by the most sensitive channel. Additionally, overall stereopsis declines significantly for increasing eccentricities [Mochizuki et al. 2012; Prince and Rogers

Table 1. Tested permutations of the independent variables  $\theta$  and  $\sigma$ .  $\sigma$  is constant in pixel-scale, thus corrected for  $\theta > 0$ .

$\rightarrow$ increasing $\sigma$ ['] $\rightarrow$							
$\theta = 0^{\circ}$	0.0'	3.0'	6.0'	9.0'	12.0'	15.0'	×
$\theta = 10^{\circ}$	0.0'	2.9′	5.8′	8.7 <b>′</b>	11.6′	14.6′	×
$\theta = 20^{\circ}$	0.0'	2.6'	5.3'	8.0'	10.6'	13.3'	26.6'

1998]. In addition to the factors mentioned above, stereoacuity is affected by the characteristics of the displayed depth modulation. Bradshaw and Rogers [1999] found that humans are more sensitive to horizontally than to vertically oriented corrugations, and are particularly sensitivity to depth modulations of 0.3–0.5cpd. This finding was reproduced by Didyk et al. [2011]. Prince and Rogers [1998] measured the peak sensitivity to spatial corrugation frequencies in the periphery and found that the sensitivity M-scales.

The endeavor of accurately reproducing cues of depth and focus in AR and VR has also been a subject of strong efforts from a display design perspective. An early multi-focal display was presented by Akeley et al. [2004], which delivered the important depth cue of accommodation. Zhao et al. [2023] introduced the concept of varifocal VR, which has a movable display to accurately deliver vergence cues to users in addition to disparity. Recently, Split-Lohmann multifocal displays demonstrated the ability to deliver a dense set of focal planes simultaneously using a single exposure [Qin et al. 2023].

While all of the aforementioned studies and efforts address various aspects of human depth perception at the fovea, only few of them explore the peripheral vision, which forms ~95% of our visual field. Although we rely on foveal vision for many tasks, reconstructing a realistic depth impression across the entire visual field is essential, particularly as the periphery plays a critical role in accurately understanding spatial layout [Alfano and Michel 1990; Fortenbaugh et al. 2008]. To the best of our knowledge, human disparity perception under the influence of peripheral blur is an open question. Given the advent of gaze-tracked VR, it is critical to understand how techniques such as foveated rendering affect depth perception.

## 3 Methods

We aim to design a perceptual model that quantifies the stereoscopic acuity threshold  $T(\theta, \sigma)$  (minimum disparity that is resolvable) as a function of blur intensity  $\sigma$  and eccentricity  $\theta$ . For this, we design an experiment displaying a textured ring with a sinusoidal depth corrugation of varying amplitude. The texture is broadband, to quantitatively study the effect of the loss of spatial frequencies on depth perception. We adjust the depth amplitude, and thus the disparity d between peaks and troughs, in a 2-Alternative Forced Choice staircase procedure to determine the resolvability threshold. We measure the stimulus at three eccentricities: 0° (fovea), 10°, and  $20^{\circ}$  eccentricity (periphery). The upper measurable eccentricity is limited by headset components, which allow binocular vision only up to 25° eccentricity. At each eccentricity, we blur the stimulus with a Gaussian filter of varying size. The filter's  $\sigma$  varies from 0'-15'. For  $20^{\circ}$  eccentricity, we include an additional measurement for  $\sigma = 26.6'$  (see Table 1 for an overview).



Fig. 3. Overview of the controlled stimulus parameters: eccentricity  $\theta$ , the ring's width, blur intensity  $\sigma$ , texture scale based on  $S_F(\theta)$ , depth map  $I_D$  based on  $w(S_C(\theta))$ , and the dithered border.

The stimulus design minimizes the influence of all depth cues except binocular disparity. While natural content presents a variety of depth cues such as motion parallax, shading, texture density, and deformation, the purpose of this study is to particularly measure disparity sensitivity in isolation. To this end, we ensure that monocular observation of the stimulus does not convey depth information, except for minor texture deformation. However, the visibility of the texture deformation is well below the resolvable disparity limit and, therefore, negligible [Bradshaw and Rogers 1999].

#### 3.1 Stimulus Design

We design our stimulus in the shape of a ring. Using a ring has the advantage of measuring the whole visual field uniformly, independent of the distinctive characteristics and strengths exhibited by different visual directions. The ring exhibits a variety of properties, including luminance pattern, and modulated depth (Figure 3). We control every parameter in a way that optimizes it for human vision. This is an effort to try to measure the lowest possible stereoacuity threshold. Given the numerous additional factors that influence depth perception that we do not control, such as presentation time, participant training, eye movement, or overall luminance level [Fendick and Westheimer 1983; Rady and Ishak 1955; Wolski et al. 2022], it is important to acknowledge that our measured thresholds may not accurately reflect the absolute threshold. Instead, they provide a general trend and demonstrate the behavior of stereoacuity under different conditions. The respective design choices are motivated in the following sections.

3.1.1 Texture. 50% Random Dot Patterns are widely used in vision science, particularly for researching depth perception [Julesz 1960]. They serve as a simple baseline pattern for our metric, exhibiting controllable features such as scale, frequency range, and dynamic range. As we aim to measure the lowest possible disparity threshold, we scale the random dot pattern to the optimally perceptible size. This scale equals the peak sensitivity value of the contrast sensitivity function (CSF) for the given eccentricity  $\theta$ . The stelaCSF<sup>1</sup> [Mantiuk et al. 2022] yields a peak sensitivity  $S_F(\theta)$  of  $S_F(0^\circ) = 4.1$ cpd,  $S_F(10^\circ) = 1.8$ cpd, and  $S_F(20^\circ) = 1.3$ cpd. We scale the dot pattern

<sup>&</sup>lt;sup>1</sup>https://github.com/gfxdisp/stelaCSF, retrieved 18. Oct 2024

SIGGRAPH Conference Papers '25, August 10-14, 2025, Vancouver, BC, Canada.

#### 4 · Sophie Kergaßner, Taimoor Tariq, and Piotr Didyk

to match  $S_F(\theta)$ , by scaling the dots such that two dots equal one cycle.

Our setup allows us to accurately pre-blur the ring's texture to precisely measure each  $\sigma$  level. We pre-render all required blur intensities by using  $\sigma$  as the standard deviation of a Gaussian filter kernel. After blurring, we recover the full dynamic range, as a reduced contrast might make it more challenging to stereo match the images and thus perceive the depth. This would result in a loss of comparability between stimuli with varying  $\sigma$ .

Additionally, the ring must be of a sufficient minimum width in the visual field to ensure good visibility. However, to reduce crosstalk between eccentricities and ensure accurate measurements, the width is minimized. To account for the varying sensitivity per eccentricity, we scale the ring wider with higher eccentricities. We set the ring's width based on the CSF, choosing a width of 13 dots for both peripheral stimuli. This is approximately equal to  $3.8^{\circ}$  in the visual field for  $10^{\circ}$ , and  $5^{\circ}$  for  $20^{\circ}$  eccentricity. For  $0^{\circ}$  eccentricity we set a width of  $6.7^{\circ}$ , so that always 2 full cycles are visible. Additionally, the ring has a randomly dithered border (Figure 3). This attenuates the monoscopic depth cue of geometric deformation, which occurs due to the pixel shifting.

3.1.2 *Depth.* Similarly to the spatial frequency of the texture, we display the ring with a depth corrugation frequency that humans are most sensitive to. In the fovea, humans have the highest sensitivity to depth corrugations with a spatial frequency of 0.2-0.5cpd [Bradshaw and Rogers 1999; Didyk et al. 2011]. However, the ideal spatial frequency for depth corrugations shifts with eccentricity. According to Prince and Rogers [1998], the sensitivity M-scales with 0.08 cycles per mm of cortex. This value is an average of the sensitivities across all directions of the visual field, and is therefore evenly applicable to our ring. We extract the M-values as an average across the four directions provided from [Rovamo and Virsu 1979], and retrieve the peak sensitivity  $S_c(\theta)$ , with  $S_c(10^\circ) = 0.133$  cpd, and  $S_c(20^\circ) = 0.073$  cpd. For  $S_c(0^\circ)$  we set a value of 0.3 cpd. Based on the circumference of the ring, we determine the number of depth corrugation cycles  $w(\theta)$  per eccentricity. This yields  $w(10^\circ) = 8$ , and  $w(20^\circ) = 9$ . The depth maps  $I_D(w(\theta))$  are pre-rendered sinusoidal patterns (as shown in Figure 3). For  $0^{\circ}$  they are horizontally oriented [Bradshaw and Rogers 1999], for the other cases they are radially oriented. For each eccentricity, we pre-render five depth maps with a phase shift of  $+0.4\pi$  each for later randomization.

To induce stereopsis, we artificially shift a pixel  $p_R$  in the right

eye's image, while leaving it untouched in the left eye's image ( $p_L$ ). This procedure follows implementations of previous work [Bradshaw and Rogers 1999; Didyk et al. 2011; Siderov and Harwerth 1995] and guarantees precise control of the presented disparity. The pixel shift introduces



Fig. 4. Schematic representation of the induced disparity between reference point  $p_Z$  and artificially created point  $p_I$ .

an intersection point  $p_I$  which our brain perceives as being in

SIGGRAPH Conference Papers '25, August 10–14, 2025, Vancouver, BC, Canada.



Fig. 5. Our three final stimuli for  $\theta \in \{0^\circ, 10^\circ, 20^\circ\}$ . The dithered border, as well as a representative depth map  $I_D$  are shown in the right halves. The lower left shows the distortion in  $I_R$  (blue) based on  $I_D$ .

front of the zero-plane (Figure 4). The disparity is measured between the point  $p_Z$ , which serves as a zero point reference (trough,  $I_D(p_Z) = 0$ ), and  $p_I$ , the artificially created illusion point (peak,  $I_D(p_I) = 1$ ). We generate the right eye's image  $I_R$  by resampling the left eye's image  $I_L$  at the respective image coordinates:

$$I_{R}(u,v) = I_{L}(u+s(d) \cdot I_{D}(u,v),v),$$
(2)

with s(d) being the shift amount in UV-coordinates based on the desired disparity d. The shift s is weighed by the depth map  $I_D$  to create the corrugations. This procedure results in a wavily warped image  $I_R$ , as each pixel is sampled slightly shifted from  $I_L$  if  $I_D > 0$ . The final stimuli are shown in Figure 5. An exemplary stereogram is provided in the Supplementary Material.

## 3.2 Psychophysical Experiment

This section details the experimental protocol. One complete session covered the judgment of all 19 conditions listed in Table 1.

3.2.1 Procedure. For each condition, the threshold estimation was conducted as follows: Subjects viewed the stimulus with a highlight either marking the peaks, or the troughs (Figure 6). Each stimulus presentation was triggered by the participant and was visible for 1.5 seconds. This duration was chosen to have adequate stimuli display time for the  $20^{\circ}$  eccentricity stimulus, as shorter presentation times

led to higher variability in measurements across different trials, as shown in pilot experiments. Pilot testing also showed that unlimited presentation times lead to visual artifacts, such as strong afterimages, during observation. Participants were asked whether the highlight marks the peaks or the troughs. After the presentation, they logged their answers via keyboard input. The stimulus intensity (amount of disparity *d*) was determined by the "Best PEST" procedure [Lieberman and Pentland 1982;



Fig. 6. The highlights mark the peaks.

Pentland 1980]. After a maximum of 60 stimulus presentations per condition, the PEST procedure was stopped and the data was extracted.

3.2.2 Apparatus. The study was conducted on a Varjo XR-3 headset. We deactivated the focus displays in the Varjo Base Application and only presented the stimulus on the context displays (2880×2720px per eye). The application was created in Unity Version 2022.3.9 using OpenGL shader. Since all resource-intensive calculations are



Fig. 7. Plot of estimated thresholds per eccentricity, as well as the polynomial fits. Each blue point denotes one measured threshold, the halo of each point denotes its weight *w*. The orange scatter plots denote the mean and standard error per  $\sigma$ . The line plots show the retrieved parabolas, fitted through the weighted threshold data.

precalculated, we operated in real-time at 90 frames per second constantly.

3.2.3 Participants. The threshold estimation of  $\theta = 10^{\circ}$  and  $\theta = 20^{\circ}$  was conducted on 11 voluntary participants (1F, 10M), with an average age of  $28 \pm 4.7$  years. All participants had normal or corrected-to-normal vision and were tested negative for stereo blindness. Nine participants reported prior experience in conducting perceptual experiments. The threshold estimation of  $\theta = 0^{\circ}$  was conducted on 10 of the 11 previous participants (1F, 9M; mean age  $28.1 \pm 4.9$  years). Eight participants reported prior experience. With the exception of the two participating authors, all participants were naïve to the project. Additionally, all participants signed their previous consent and were compensated for their time. The study was approved by the institutional ethics committee.

#### 4 Results and Perceptual Model

This section covers the evaluation of the psychophysical experiment to retrieve all thresholds  $T_p(\theta, \sigma)$  per participant and condition, as well as the design and fitting of our perceptual model.

#### 4.1 Threshold Estimation

The PEST procedure yields one dataset per measured condition which is evaluated independently to extract the estimated threshold  $T_p(\theta, \sigma)$ . We generate a data point for each measured disparity intensity *d*, such that the value denotes the probability of detection as  $P(d) = \frac{c_d}{n_d}$ , where  $n_d$  is the number of samples per disparity, and  $c_d$  is the number of correctly detected stimuli. While a detection probability of 50% means that the participant failed to resolve the stimulus and answered randomly, a 100% score means that the participant correctly resolved the depth in each trial. We fit a psychometric curve and determine the 75% point. The 75% point is commonly used to describe an above-random score, which means that participants resolved the stimulus correctly. We fit a Weibull cumulative distribution function *F*, and estimate the threshold *T* 

$$x = T_p(\theta, \sigma)$$
, such that  $F(x) = 0.75$ . (3)

Since our probability values range from P = 0.5 (random guess) to P = 1 (always correct answer), we fit a modified Weibull curve that outputs values  $F \in [0.5, 1]$ :

$$F(x) = 1 - \frac{1}{2} \cdot e^{-(\frac{x}{\lambda})^k},$$
(4)

where  $\lambda$  and k are the parameters we fit. For the fit, we weigh each data point P(d) with the weight  $w(d) = n_d$ . This ensures that



Fig. 8. Fit of the three parabola parameters  $p_1$ ,  $p_2$ , and  $p_3$ . The fitted curves are exponential curves of shape  $p_i(\theta) = ae^{b\theta} + c$ .

frequently measured disparity levels become a larger weight in the fit, as it is more certain that P(d) describes the real value.

Since the PEST procedure can lead to sparse sampling of important regions and random answers can render data unreliable, the Weibull fit can get unstable. We obtain the final threshold *T* as well as its standard deviation  $T_{\sigma}$  by bootstrapping the data 100 times. This also yields a relative uncertainty of  $u = \frac{T_{\sigma}}{T}$  per *T*. Measurements with u > 0.3 are treated as outliers and excluded from further evaluation, as we consider the PEST procedure to be too unstable and therefore unreliable. This affects ~20% of our measurements. A plot of the excluded data points is provided in the Supplementary Material.

## 4.2 Perceptual Model

We aim to fit a continuous surface M through our acquired thresholds  $T_p(\theta, \sigma)$ . For this, we take a look at the characteristics exhibited by our measured thresholds (Figure 7): Instead of increasing monotonically, the mean values of T first decrease with increasing  $\sigma$ , before eventually rising again. To match this characteristic, we fit a parabola of the form  $p_1(\theta) \cdot (\sigma - p_2(\theta))^2 + p_3(\theta)$  to each eccentricity  $\theta$  (Figure 7), where  $p_1$  defines the steepness,  $p_2$  the shift of the minimum in  $\sigma$ -direction, and  $p_3$  denotes the T-value of the minimum of the respective parabola. Each data point is weighted by an inverse of its uncertainty u with  $w(u) = \frac{1}{u^2}$  to ensure that more reliable threshold estimations take more influence on the fit. From this, we interpolate between the three parabolas to complete the model and retrieve the surface M of shape

$$M = p_1(\theta) \cdot (\sigma - p_2(\theta))^2 + p_3(\theta), \quad \theta \in [0, 20].$$
(5)

To estimate the parameters' interpolation characteristics for varying  $\theta$ , we examine the characteristics of our fitted parabolas (Figure 7): The minimum threshold, associated with  $p_2$ , shifts to a higher  $\sigma$  with increasing eccentricity. This indicates that the HVS tolerates higher blur levels in the periphery. Furthermore, the overall threshold *T* 

#### 6 · Sophie Kergaßner, Taimoor Tariq, and Piotr Didyk



Fig. 9. Our blur rates set into comparison with various  $\sigma$ -values reported in literature [Mantiuk et al. 2022; Patney et al. 2016; Thibos et al. 1987; Tursun et al. 2019]. If a cut-off frequency f[cpd] is reported, we retrieve  $\sigma$ by assuming that the Gaussian filter has a kernel of size  $\sigma = \frac{3}{2\pi f}$ .

shifts to higher values with increasing eccentricity, which resembles the y-offset encoded in  $p_3$ . The slope  $p_1$  remains mostly unaltered. Figure 8 displays the parameter-values  $p_i$  per eccentricity. As expected, they resemble the intuitive understanding discussed above. To model the parabolas for a continuous  $\theta$ , we fit an exponential curve of the form  $p_i(\theta) = ae^{b\theta} + c$ . An exponential curves increase strictly, which resembles our data. The fitting yields:

$$p_{1}(\theta) = 2.07 \cdot 10^{-11} \cdot e^{0.87 \cdot \theta} + 0.003$$

$$p_{2}(\theta) = 8.85 \cdot e^{0.04 \cdot \theta} - 7.5$$

$$p_{3}(\theta) = 0.04 \cdot e^{0.15 \cdot \theta} + 0.12$$
(6)

The final surface fit M is displayed in Figure 1. It is described by inserting the fitted parameters  $p_i$  in Equation 5, which results in:

$$M = \left[ 2.07 \cdot 10^{-11} \cdot e^{0.87 \cdot \theta} + 0.003 \right] \cdot \left( \sigma - \left[ 8.85 \cdot e^{0.04 \cdot \theta} - 7.5 \right] \right)^2 + \left[ 0.04 \cdot e^{0.15 \cdot \theta} + 0.12 \right]$$
(7)

This model holds for our measured scope of  $\sigma \in [0, 15]$ , and  $\theta \in [0, 20]$ . However, we note that  $p_1$  exhibits a steep slope for higher eccentricities, resulting in extreme values for  $\theta > 20$ . If such conditions are expected we suggest replacing  $p_1$  with a constant fit of  $p_1 = 0.0034$ , which is the mean across  $p_1(0^\circ)$ ,  $p_1(10^\circ)$  and  $p_1(20^\circ)$ . Replacing  $p_1$  with this value only results in a minor difference to the current model, as the surface is still well-defined by  $p_2$  and  $p_3$ , and  $p_1$ 's variability across the three measured eccentricities is not extreme.

#### 4.3 Discussion

We retrieve  $p_2(\theta)$  as the blur level at which participants presented optimal stereoacuity. This yields exemplary  $\sigma$ -values of  $p_2(10^\circ) =$ 5.41', and  $p_2(20^\circ) = 11.33'$ . Figure 9 sets these values in comparison to foveation values reported in literature. Our model exceeds all reported blur intensities, showing that stereoacuity is unaffected by commonly used levels of foveation. However, the shape of our model, especially the high tolerance for blur beyond the limit of visibility, is surprising. This means that stereopsis does not depend on perceivable high frequencies.



Fig. 10. Illustration of how aliased frequencies influence stereo matching.



Fig. 11. Renderings of the two natural scenes ( $FOREST^2$  and  $KITCHEN^3$ ) used in the validation experiment.

Moreover, there seems to be a constant trend that the removal of high frequencies ( $\sigma < p_2$ ) improves the stereoacuity slightly. We hypothesize that this is due to the removal of aliased frequencies: When sparsely sampling a continuous signal, only frequencies up to  $0.5 \times$ sampling rate can be correctly reconstructed (Nyquist-Shannon sampling theorem). All frequencies above this threshold lead to aliasing. This means that the HVS may interpolate the aliased samples incorrectly, leading to ambiguities (Figure 10). This false recreation leads the HVS to stereo match those signals incorrectly. Since correct stereo matching is the basis for correctly resolving disparity, this phenomenon leads to a distorted depth perception. Removing all confounding frequencies beforehand (low-pass filtering) prevents this mismatch. Figure 9 marks the two equivalent  $\sigma$ -values that remove all aliased frequencies according to Thibos et al. [1987]. However, these  $\sigma$ -values are ~50% lower than our threshold. This contradicts the assumption that the trend of T is solely influenced by the presence of aliased frequencies, since T seems to be persistent beyond the removal of just the aliased frequencies. Instead, we hypothesize that there must be high frequencies that disturb stereo matching beyond the limit given by Thibos et al. [1987].

## 5 Validation

The experiment described in the previous section was conducted on simple, isolated stimuli. To verify whether the model holds for complex stimuli, we conducted a validation experiment on natural scenes and verify that the foveation derived from our model does not affect depth perception. The experiment is set in a FOREST and a KITCHEN (Figure 11).

## 5.1 Methods

We show participants half-split images of the natural scenes. The central part is blacked out up to  $15^{\circ}$  eccentricity. We validate our model on high eccentricities, as the fovea remains almost unaffected

<sup>&</sup>lt;sup>2</sup>Forest Environment from NatureManufacture (Unity)

<sup>&</sup>lt;sup>3</sup>Retro Farmhouse from NOT\_Lonely (Unity)

by our model, yielding no measurable changes. As we allow the gaze to move within a margin of  $3-5^{\circ}$  from the fixation point, the effective blacking is ~10°. Furthermore, an examination of > 15° eccentricity retains enough binocular image space, as the headset only supports binocular vision up to 25° eccentricity. The left and right halves of the image display different disparity levels: one half is rendered with 0mm IPD (no disparity), while the other half is rendered with a varying IPD (0–20mm) to control the overall magnitude of disparities in the scene (see Supplementary Material for vergence angle histograms). The IPD, and thus the overall scene depth, was adjusted in a staircase procedure. We simulate foveation by locally interpolating between pre-blurred images according to the per-pixel eccentricity value. Afterwards, we linearly map the values of the input image.

Participants were shown both scenes in full-resolution (ORG) and foveated (Fov) rendering styles. Each participant first viewed the FOREST scene (F), and then the KITCHEN scene (K). The order of presented rendering styles was randomized, and the style sequence was counterbalanced across scenes (e.g. F-Fov, F-ORG, K-ORG, K-Fov). Before threshold estimation began, participants were allowed to freely explore each scene using their central vision to familiarize themselves with prominent depth cues. This step was added to reduce measurement noise caused by initial variability in scene familiarity, as identified in pilot testing. Participants were asked to identify the side of the image which exhibited disparity. To estimate the IPD threshold, the displayed disparity was adjusted using the "Best PEST" procedure, as described in Section 3.2.1. The experiment was conducted using the same apparatus as described in Section 3.2.2. We conducted the study on 15 participants (13M, 1F, 1X; mean age 26.1  $\pm$  2.5 years). 11 participants reported prior experience with perceptual studies. All had normal or corrected-to-normal vision, signed their previous consent, and were compensated for their time. The study was approved by the institutional ethics committee.

## 5.2 Results

The PEST procedures are evaluated and weighed as stated in Section 4.1. Figure 12 shows the estimated IPD thresholds, as well as a representation of the change of the participant's IPD threshold between the ORG and Fov stimuli. The change rate is determined by dividing a participant's threshold presented for an ORG stimulus by the threshold of the respective Fov stimulus. This results in a positive change value in case the participant had a lower threshold in the Fov condition.

Neither the means nor the medians differ substantially for both the thresholds and change rates. The overlapping box plot notches yield a 95% certainty that the medians do not differ significantly, which is the case for all reported values. This suggests that participants exhibited similar stereoacuity for both rendering styles, and that the foveation did not compromise it. This finding is reflected by the constant change rate of ~0%, which shows that the participants' performance did not change in either of the settings. Although the change rates vary up to  $\pm 20\%$ , we did not find a significant trend in personal preferences, meaning that participants did not have a



Fig. 12. Left: Estimated IPD thresholds for both scenes and rendering styles. The area of each dot denotes its weight. The scatter marks the mean and standard error per condition, while the box plots show the median, 25% and 75% percentiles, 1.5 IQR whiskers, and 95% confidence interval. Right: Change of the log-threshold per participant and scene. Only participants with two valid measurements per scene were considered. A positive value denotes a smaller threshold for the Fov style.

similar change rate across conditions. In addition, the directions of change were not related to expert vs. inexperienced viewer.

## 5.3 Discussion

The results demonstrate that the removal of high frequencies does not affect stereoacuity. This corresponds to our hypothesis, that high frequencies do not carry essential information for stereo matching. Additionally, the results confirm that our model holds for complex scenes that differ significantly from the simple stimulus used in our first experiment. However, for complex stimuli, we were not able to replicate the increased stereoacuity anticipated with our suggested blur intensity. Two factors may contribute to this: Firstly, the observed decline in *T* towards  $p_2$  was subtle. Secondly, the measurement process for the IPD thresholds is inherently more variable due to the abstract nature of the task and the complex stimuli, which may have posed a challenge for participants to resolve consistently. However, the applied foveation did not detrimentally affect overall stereoacuity.

## 6 Future Work

Recreating a correct spatial understanding is crucial for various real-time applications, such as simulations, medical training, entertainment, and design interfaces. Our findings indicate that strong peripheral blur does not compromise depth perception derived from stereopsis. Furthermore, the results of our psychophysical experiment suggest that an optimal amount of peripheral blur can actually improve stereoacuity. We hypothesize that this improvement may be due to the reduction of spatial aliasing in peripheral vision. However, this argument does not fully account for our observations, as the limits presented by Thibos et al. [1987] do not align with the blur levels at which stereoacuity was highest in our experiments.

We believe that further investigation of this effect is a fascinating avenue for future research, which could enhance our understanding of the mechanisms behind depth perception. This may lead to interesting applications where careful manipulation of image content improves spatial awareness. Additionally, potential findings can lead to new conclusions about the quality and reproduction of visual cues, particularly in light of recent efforts to enhance and recreate

#### 8 • Sophie Kergaßner, Taimoor Tariq, and Piotr Didyk

high-frequency details in foveated rendering [Patney et al. 2016; Tariq et al. 2022; Walton et al. 2021].

It has been demonstrated that stereoacuity decreases in peripheral vision [Blakemore 1970; Mochizuki et al. 2012; Prince and Rogers 1998]. Our results align with these findings and complement them by showing that this trend is not further enhanced by peripheral blur. This highlights the need for careful treatment of depth information in foveated content: Our findings indicate that reducing the spatial quality of an image does not entail that the quality of depth reproduction can be reduced as well. This observation is particularly important in applications that require depth processing. For example, stereoscopic and multiscopic content can be transmitted as a combination of rendered images and depth maps to recreate missing views during the decoding step [Merkle et al. 2009]. When developing gaze-contingent compression methods, the compression of depth maps should be handled carefully and ideally separately from the compression of the image content. Similarly, any depth-related rendering process, such as displacement mapping or level-of-detail methods [Murphy and Duchowski 2001; Surace et al. 2023], should be handled carefully and without implicitly assuming that foveation masks naïve compression of depth information. Exploring these issues in the future will broaden our understanding of the critical factors that must be replicated to achieve real-time realism in virtual and augmented realities.

## 7 Conclusion

Techniques such as foveated rendering are an essential tool for reproducing real-time visual realism in its entirety. However, it is paramount that other dimensions of realism, such as the perception of depth, are maintained in conjunction with perceived spatial realism. The goal of this work is to understand how foveated rendering affects our perception of disparity in immersive environments. To this end, we measure disparity thresholds across various eccentricities  $(0-20^{\circ})$  under varying blur intensities (0-26.6' Gaussian filter  $\sigma$ ). Based on our measurements, we derive a perceptual model which demonstrates that the blur intensities applied in common foveation procedures do not affect stereoacuity. Moreover, the threshold for stable stereoacuity is much higher ( $\sim 2 \times$ ) than the blur applied in traditional foveated rendering. We validate these findings on complex scenes and show that the overall perceived depth in a scene is not affected by strong peripheral blur. Additionally, our findings suggest that it is important to maintain high depth quality even in strongly foveated stereoscopic content.

#### Acknowledgments

This project has received funding from the European Research Council under the European Union's Horizon 2020 research and innovation program (Grant 804226, PERDY).

#### References

- Israel Abramov, James Gordon, and Hoover Chan. 1991. Color Appearance in the Peripheral Retina: Effects of Stimulus Size. *Jour. Optical Society America* 8, 2 (1991), 404–414.
- Kaan Akşit, Ward Lopes, Jonghyun Kim, Peter Shirley, and David Luebke. 2017. Near-Eye Varifocal Augmented Reality Display Using See-Through Screens. ACM Trans. Graph. 36, 6 (2017), 1–13.

SIGGRAPH Conference Papers '25, August 10-14, 2025, Vancouver, BC, Canada.

- Kurt Akeley, Simon J. Watt, Ahna Reza Girshick, and Martin S. Banks. 2004. A Stereo Display Prototype with Multiple Focal Distances. ACM Trans. Graph. 23, 3 (2004), 804–813.
- Patricia L. Alfano and George F. Michel. 1990. Restricting the Field of View: Perceptual and Performance Effects. *Perceptual and Motor Skills* 70, 1 (1990), 35–45.
- Joseph Baldwin, Alistair Burleigh, Robert Pepperell, and Nicole Ruta. 2016. The Perceived Size and Shape of Objects in Peripheral Vision. *I-Perception* 7, 4 (2016), 1–23.
- Colin Blakemore. 1970. The Range and Scope of Binocular Depth Discrimination in Man. Jour. Physiology 211, 3 (1970), 599–622.
- J. M. Bosten, P. T. Goodbourn, A. J. Lawrance-Owen, G. Bargary, R. E. Hogg, and J. D. Mollon. 2015. A Population Study of Binocular Function. *Vision Res.* 110 (2015), 34–50.
- Mark F. Bradshaw and Brian J. Rogers. 1999. Sensitivity to Horizontal and Vertical Corrugations Defined by Binocular Disparity. Vision Res. 39, 18 (1999), 3049–3056.
- Lawrence K. Cormack, Scott B. Stevenson, and Clifton M. Schor. 1991. Interocular Correlation, Luminance Contrast and Cyclopean Processing. Vision Res. 31, 12 (1991), 2195–2207.
- B. E. Coutant and G. Westheimer. 1993. Population Distribution of Stereoscopic Ability. Ophthalmic and Physiological Optics 13, 1 (1993), 3–7.
- Christine A. Curcio, Kenneth R. Sloan, Robert E. Kalina, and Anita E. Hendrickson. 1990. Human Photoreceptor Topography. *Jour. Comparative Neurology* 292, 4 (1990), 497–523.
- James E. Cutting and Peter M. Vishton. 1995. Perceiving Layout and Knowing Distances: The Integration, Relative Potency, and Contextual Use of Different Information about Depth. In *Perception of Space and Motion*. Academic Press, San Diego, USA, 69–117.
- Piotr Didyk, Tobias Ritschel, Elmar Eisemann, Karol Myszkowski, and Hans-Peter Seidel. 2011. A Perceptual Model for Disparity. ACM Trans. Graph. 30, 4 (2011), 1–10.
- Piotr Didyk, Tobias Ritschel, Elmar Eisemann, Karol Myszkowski, Hans-Peter Seidel, and Wojciech Matusik. 2012. A Luminance-Contrast-Aware Disparity Model and Applications. ACM Trans. Graph. 31, 6 (2012), 1–10.
- Miacheal Fendick and Gerald Westheimer. 1983. Effects of Practice and the Separation of Test Targets on Foveal and Peripheral Stereoacuity. Vision Res. 23, 2 (1983), 145–150.
- Francesca C. Fortenbaugh, John C. Hicks, and Kathleen A. Turano. 2008. The Effect of Peripheral Visual Field Loss on Representations of Space: Evidence for Distortion and Adaptation. Investigative Opthalmology & Visual Science 49, 6 (2008), 2765–2772.
- Brian Guenter, Mark Finch, Steven Drucker, Desney Tan, and John Snyder. 2012. Foveated 3D Graphics. ACM Trans. Graph. 31, 6 (2012), 1–10.
- Bela Julesz. 1960. Binocular Depth Perception of Computer-Generated Patterns. *The Bell System Technical Journal* 39, 5 (1960), 1125–1162.
- Brooke Krajancich, Petr Kellnhofer, and Gordon Wetzstein. 2023. Towards Attention–Aware Foveated Rendering. ACM Trans. Graph. 42, 4 (2023), 1–10.
- Harris R. Lieberman and Alex P. Pentland. 1982. Microcomputer-Based Estimation of Psychophysical Thresholds: The Best PEST. Behavior Research Methods & Instrumentation 14, 1 (1982), 21–25.
- Rafał K. Mantiuk, Maliha Ashraf, and Alexandre Chapiro. 2022. stelaCSF: A Unified Model of Contrast Sensitivity as the Function of Spatio-Temporal Frequency, Eccentricity, Luminance and Area. ACM Trans. Graph. 41, 4 (2022), 1–16.
- D. Marr and T. Poggio. 1979. A Computational Theory of Human Stereo Vision. Proc. Royal Society London 204, 1156 (1979), 301–328.
- Suzanne P. Mckee and Ken Nakayama. 1984. The Detection of Motion in the Peripheral Visual Field. Vision Res. 24, 1 (1984), 25–32.
- P. Merkle, Y. Wang, K. Muller, A. Smolic, and T. Wiegand. 2009. Video Plus Depth Compression for Mobile 3D Services. In 3DTV Conf.: The True Vision-Capture, Transmission and Display of 3D Video. IEEE, 1–4.
- Hiroshi Mochizuki, Nobuyuki Shoji, Eriko Ando, Maiko Otsuka, Kenichiro Takahashi, Tomoya Handa, et al. 2012. The Magnitude of Stereopsis in Peripheral Visual Fields. *Kitasato Medical Journel* 41 (2012), 1–5.
- J.D. Moreland and A. Cruz. 1959. Colour Perception with the Peripheral Retina. Optica Acta: Int. Jour. Optics 6, 2 (1959), 117–151.
- Hunter A. Murphy and Andrew T. Duchowski. 2001. Gaze-Contingent Level Of Detail Rendering. In Eurographics Short Presentations. Eurographics Assoc., 1–10.
- Rahul Narain, Rachel A. Albert, Abdullah Bulbul, Gregory J. Ward, Martin S. Banks, and James F. O'Brien. 2015. Optimal Presentation of Imagery with Focus Cues on Multi-Plane Displays. ACM Trans. Graph. 34, 4 (2015), 1–12.
- 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 Trans. Graph. 35, 6 (2016), 1–12.
- Alex Pentland. 1980. Maximum Likelihood Estimation: The Best PEST. Perception & Psychophysics 28, 4 (1980), 377–379.
- Simon J. D. Prince and Brian J. Rogers. 1998. Sensitivity to Disparity Corrugations in Peripheral Vision. Vision Res. 38, 17 (1998), 2533–2537.
- Yingsi Qin, Wei-Yu Chen, Matthew O'Toole, and Aswin C. Sankaranarayanan. 2023. Split-Lohmann Multifocal Displays. ACM Trans. Graph. 42, 4 (2023), 1–18.

- Towards Understanding Depth Perception in Foveated Rendering 9
- A. A. Rady and I. G. H. Ishak. 1955. Relative Contributions of Disparity and Convergence to Stereoscopic Acuity. *Jour. Optical Society America* 45, 7 (1955), 530–534.
- Ruth Rosenholtz. 2016. Capabilities and Limitations of Peripheral Vision. Annual Review of Vision Science 2, 1 (2016), 437–457.
- J. Rovamo and V. Virsu. 1979. An Estimation and Application of the Human Cortical Magnification Factor. *Experimental Brain Research* 37, 3 (1979), 495–510.
- John Siderov and Ronald S. Harwerth. 1995. Stereopsis, Spatial Frequency and Retinal Eccentricity. Vision Res. 35, 16 (1995), 2329–2337.
- Hans Strasburger, Ingo Rentschler, and Martin Jüttner. 2011. Peripheral Vision and Pattern Recognition: A Review. *Jour. Vision* 11, 5 (2011), 1–84.
- Qi Sun, Fu-Chung Huang, Li-Yi Wei, David Luebke, Arie Kaufman, and Joohwan Kim. 2019. Eccentricity Effects on Blur and Depth Perception. arXiv:1708.06034
- Luca Surace, Cara Tursun, Ufuk Celikcan, and Piotr Didyk. 2023. Gaze-Contingent Perceptual Level of Detail Prediction. In *Eurographics Symp. Rendering*. Eurographics Assoc., 1–12.
- Barbara Sweet and Mary Kaiser. 2012. Depth Perception, Cueing, and Control. 1-10.
- Taimoor Tariq and Piotr Didyk. 2024. Towards Motion Metamers for Foveated Rendering. ACM Trans. Graph. 43, 4 (2024), 1–10.

- Taimoor Tariq, Cara Tursun, and Piotr Didyk. 2022. Noise-Based Enhancement for Foveated Rendering. ACM Trans. Graph. 41, 4 (2022), 1–14.
- Larry N. Thibos, D. J. Walsh, and F. E. Cheney. 1987. Vision beyond the resolution limit: aliasing in the periphery. Vision Res. 27, 12 (1987), 2193–2197.
- Okan Tarhan Tursun, Elena Arabadzhiyska-Koleva, Marek Wernikowski, Radosław Mantiuk, Hans-Peter Seidel, Karol Myszkowski, and Piotr Didyk. 2019. Luminance-Contrast-Aware Foveated Rendering. *ACM Trans. Graph.* 38, 4 (2019), 1–14.
- David R. Walton, Rafael Kuffner dos Anjos, Sebastian Friston, David Swapp, K. Akşit, Anthony Steed, and Tobias Ritschel. 2021. Beyond Blur: Real-Time Ventral Metamers for Foveated Rendering. *ACM Trans. Graph.* 40, 4 (2021), 1–14.
- David Whitaker, Keziah Latham, Pia Mäkelä, and Jyrki Rovamo. 1993. Detection and Discrimination of Curvature in Foveal and Peripheral Vision. *Vision Res.* 33, 16 (1993), 2215–2224.
- Krzysztof Wolski, Fangcheng Zhong, Karol Myszkowski, and Rafał K. Mantiuk. 2022. Dark Stereo: Improving Depth Perception under Low Luminance. ACM Trans. Graph. 41, 4 (2022), 1–12.
- Yang Zhao, Dave Lindberg, Bruce Cleary, Olivier Mercier, Ryan Mcclelland, Eric Penner, Yu-Jen Lin, Julia Majors, and Douglas Lanman. 2023. Retinal-Resolution Varifocal VR. In ACM SIGGRAPH '23 Emerging Technologies. ACM, New York, USA, 1–3.