The Impact of Blur on Motion Parallax and Binocular Disparity

Sophie Kergaßner; Università della Svizzera italiana, Switzerland
Piotr Didyk; Università della Svizzera italiana, Switzerland

Abstract

Foveated rendering is a key technique for reducing computational load in immersive display systems by lowering rendered image quality in the peripheral visual field while preserving high fidelity in the fovea. While the impact of foveation on perceived spatial detail is well understood, its influence on other visual qualities, such as depth from motion parallax, remains unclear. In this work, we investigate how foveated rendering affects motion-based depth perception across the visual field. Building on previous work on binocular disparity, we use a comparable experimental setup to isolate motion parallax as the sole depth cue and measure depth discrimination thresholds under varying levels of foveation, modeled as varying intensities of spatial blur, and eccentricity. Our results show that depth from motion is immediately impaired by visible foveation, with stronger impairments at higher levels of blur. These findings suggest that motion-based depth cues may be more sensitive to foveated rendering than disparity cues, which were previously found to be largely unaffected.

Introduction

Rendering perceptually convincing images is computationally expensive, particularly when rendering for immersive displays such as virtual reality (VR) and augmented reality (AR) headsets, where high-quality images must be generated in real time for a pleasant experience. The need to maintain high frame rates, combined with limited hardware resources of VR/AR devices, significantly restricts the available computational budget. To meet these performance requirements, the rendering process must be simplified in a way that is imperceptible to the viewer. This can be achieved by exploiting the compressive characteristics of the human visual system (HVS), whose sensitivity is highly non-uniform across the visual field. In particular, visual acuity decreases rapidly with increasing eccentricity from the fovea. A widely used technique that leverages this property is *foveated rendering* [6, 17], where image resolution is reduced with increasing distance from the point of gaze, allowing computational resources to be concentrated in the foveal region. By adjusting the rate at which resolution decreases across eccentricity, it is possible to achieve a perceptually identical image at significantly lower computational cost. This makes foveated rendering a key enabling technology for real-time, high-quality immersive systems.

While the perceptual consequences of reduced peripheral resolution are relatively well understood with respect to spatial detail, its influence on other visual qualities is less clear. Recent work suggests that foveated rendering can affect perceptual processes beyond spatial acuity, including the perception of ego motion [23] and scene layout [17]. Previous work on depth perception, however, found that stereoacuity, i.e., the precision of depth perception from binocular disparity, remains largely robust under foveated rendering [9]. This raises the question of whether other depth cues exhibit similar resilience, particularly with respect to the observed sensitivity of motion perception to peripheral blur [23], which suggests that motion-dependent depth cues may be more strongly affected. In this context, the motion-dependent depth cue is *motion parallax*, which arises from relative image motion during observer movement.

In this work, we investigate the impact of foveated rendering on depth perception from motion parallax. To this end, we extend the experimental paradigm introduced in [9] by adapting the stimulus to isolate motion-parallax-based depth perception. We conduct a psychophysical user study and fit a perceptual model to the collected data in order to describe how depth sensitivity changes across eccentricity under various conditions of foveation intensity. We model the reduced resolution from foveation by applying spatial Gaussian blur to our images [1]. Because the stimulus is a direct adaptation of the one used by Kergaßner et al. [9], the resulting perceptual model is numerically comparable to the previously obtained model for stereoacuity. This allows us to directly contrast the qualitative and quantitative effects of foveated rendering on two fundamental relative depth cues, and to assess how peripheral resolution loss influences depth perception in dynamic immersive scenes.

Background

Human vision is highly non-uniform, with sharp detail at the center and coarser perception in the periphery. This characteristic has implications for how we experience immersive displays.

Foveated Rendering Rendering high-quality images for immersive displays such as VR and AR headsets is computationally expensive, particularly when targeting high frame rates. Foveated rendering provides a practical solution by leveraging the non-uniform sensitivity to spatial detail of the HVS. Visual acuity is highest in the fovea and declines toward the periphery due to the lower density of cone photoreceptors and retinal ganglion cells [2, 20]. This allows the visual system to tolerate substantial loss of detail in peripheral regions without affecting perceived image quality. This loss of detail can manifest in various forms [15], such as reduced spatial [17] and temporal resolution [5], decreased geometric complexity [13, 21], and chromatic degradation [4, 12]. In this work, we focus in particular on reduced spatial resolution in the periphery, which has received considerable attention in previous work. Early approaches to foveated rendering focused primarily on reducing spatial resolution in peripheral regions based solely on eccentricity [6, 17]. Post-processing methods such as contrast enhancement or frequency reintroduction allow more aggressive foveation [8, 17, 24, 27]. Other approaches leverage content-aware masking to further decrease resolution in areas where it is unlikely to be noticed [26], or exploit higher-level attentional or task-

based models to foveate based on the degree of viewer distraction [10]. These strategies collectively enable modern foveated rendering systems that maintain perceptually convincing images while reducing computational cost. While most foveation methods are designed and evaluated to preserve spatial image quality, the impact on other perceptual aspects beyond spatial resolution remains underexplored. Indeed, recent work shows that the perception of motion speed [23] and global layout [17] are affected by moderate foveation. These findings suggest that perceptual equivalence cannot be assessed based on spatial resolution alone, and that other cues may be degraded even when image quality appears largely unchanged.

Depth Perception One of such possibly impaired cues is depth perception. Immersive displays make it feasible to present rich depth cues, supporting realistic perception of 3D structure. While many depth cues convey relative or absolute distance information solely through pictorial cues [3, 22], interactive scenarios particularly rely on dynamic cues that respond to observer motion. Two of the most relevant cues in VR are stereoscopic disparity, inherent to VR headset design, and motion parallax, which arises from relative image motion as the observer moves. Both provide relative information about depth differences [19]. Previous work investigated how peripheral blur influences depth perception from binocular disparity, showing that stereopsis remains remarkably robust even under substantial foveation [9]. However, the effect of peripheral blur on motion parallax has not been systematically examined. In the present study, we focus on motion-parallax-based depth perception in interactive viewing scenarios. To relate these measurements to previous work on stereopsis, we assume that depth estimates from disparity and motion are comparable by treating the endpoints of a monoscopic observer's lateral movement as equivalent to the positions of the eyes in a stereoscopic system [7, 19]. Formally, the perceived disparity d is defined as the difference in vergence angles between two points A and B:

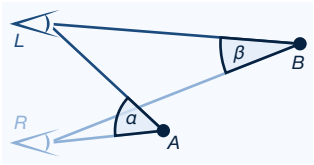

Figure 1. *Geometry of perceived binocular disparity (Equation 1) induced by having two eyes L and R. An equivalent measure of disparity magnitude from motion parallax can be found when the two eyes are treated as endpoints of the movement of a monoscopic observation.*

$$d = D(A) - D(B), \tag{1}$$

where D denotes the vergence angle, which is the angle between the two visual axes and the observed point (α and β in Figure 1). This equivalence allows us to express motion-parallax thresholds in units of "disparity" and directly compare them to the stereoacuity measurements reported by Kergaßner et al. [9].

Method

We aim to examine the influence of spatial blur on the perceived depth from motion parallax across a range of eccentricities. We design our study to extend the perceptual model proposed by Kergaßner et al. [9], who examined the impact of foveation on stereoacuity. In that interest, our goal is to stay as close as possible to the original experimental paradigm, to achieve numerical comparability between the two models. Throughout this manuscript, we also follow the naming conventions of [9] for consistency.

Stimulus Design Participants are presented with a textured ring featuring sinusoidal depth corrugations (Figure 2). The ring spans the field of view evenly across one of the three examined eccentricities $\theta \in \{0°, 10°, 20°\}$. We examine a range of spatial blur intensities $\sigma \in [0', 26.6']$ applied to the luminance pattern. Different blur intensities are achieved by applying a Gaussian blur with varying standard deviation σ to a full-resolution image. This emulates foveation from reduced resolution [1]. The luminance pattern and depth modulation are both designed to be optimally visible for a human observer. This is achieved by tuning both patterns to a spatial frequency where the contrast sensitivity function (CSF) and disparity sensitivity function (DSF) exhibit peak sensitivity. This procedure yields a luminance frequency of 4.1 cycles per degree (cpd) at 0°, 1.8 cpd at 10°, and 1.3 cpd at 20° eccentricity. The depth pattern consists of vertical corrugations with a spatial frequency of 0.3 cpd at 0°. For 10° and 20° eccentricity, the stimulus contains 8 and 9 full depth cycles, respectively.

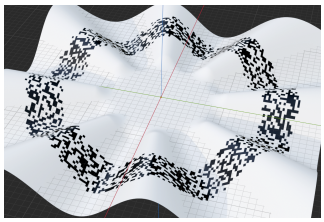

Figure 2. *Side view of the stimulus geometry, displaying its depth corrugations and an exemplary luminance pattern for 20° eccentricity.*

Stimulus Presentation Participants observe a monoscopic representation of the ring. We realize this by rendering an orthographically tilted view of the stimulus and displaying the resulting 2D image to the observer. The observer's movement is directly linked to the perspective from which the stimulus is rendered. Thus, by moving their head, participants perceive different monoscopic views of the stimulus that are consistent with their position in space (Figure 3 and 4). This results in a depth impression solely induced by motion parallax. All other depth cues are suppressed. Participant movement was restricted to a maximum angular deviation of 4° relative to the surface normal of the stimulus plane, corresponding to a lateral movement range of approximately 40 cm (Figure 3). This constraint ensures a well-defined range of motion, which is necessary to calculate the perceived depth in units of disparity (Equation 1). Lastly, to give further perspective cues, a matte ellipsoid and a thin rod aligned with the plane normal were placed at the center of the stimulus (Figure 4). These elements provide shading cues that help observers maintain orientation during movement. During testing, we found that correct processing of motion parallax critically depends on participants physically moving their head. In conditions where participants instead manipulated a digital object with their hand to

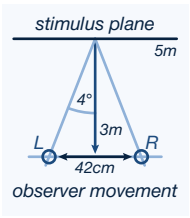

Figure 3. *Schematic of the restriction of observer movement.*

Figure 4. *Participant movement leads to a perspectively different rendering of the stimulus. The depth modulation is strongly amplified for illustrating purpose. The distortion of the texture was not the primary depth cue.*

induce relative motion, depth perception became ambiguous with regard to depth polarity, meaning that the perception of peaks and troughs randomly flipped during observation [16].

User Study Setup and Participants For each condition (combination of eccentricity θ and blur intensity σ), we estimate the minimum depth amplitude that participants can reliably resolve using a 2AFC threshold estimation procedure. The resulting thresholds T serve as a direct measure of depth sensitivity, where lower thresholds indicate more precise depth perception from motion parallax. We examined 9 participants with an average age of 29 ± 6 years. All had normal or corrected to normal vision, previous experience with perceptual experiments, were compensated for their time, and signed previous consent. The study was approved by the institutional ethics committee. The total duration of the experiment was 180 minutes, split into multiple individual sessions over multiple days.

Procedure and Task Each participant observed each condition once in randomized order, resulting in 19 2AFC estimation procedures per participant. Per condition, participants were prompted to move in front of the display to induce motion parallax. They were instructed to use the full allowed range of motion to perceive the correct amount of motion parallax. Participants were asked whether the marked regions corresponded to depth peaks (hills) or troughs (valleys). Based on the participant's response, the exhibited depth for the next stimulus display was determined following the bestPEST procedure [11, 18]. In total, participants judged 30 stimuli per condition, resulting in a set of 30 binary responses (correct/incorrect) over a range of tested blur intensities.

Results

To generalize our findings, we fit a model to our measured data. With our data preprocessing, model design and fitting, we aim to stay closely to the procedure described by Kergaßner et al. [9] (Section 4.1 and 4.2) to maintain comparability.

Data Preprocessing Per participant and condition, we retrieve a set of 30 binary responses from which we aim to retrieve a reliable 75% threshold as the perceptibility threshold. Following Kergaßner et al. [9], we bootstrap the data 1000 times by fitting a Weibull function, to retrieve the mean 75% threshold point $T_p(\theta,\sigma)$, as well as its standard deviation T_σ. We interpret the ratio between the estimated mean threshold T and its standard deviation T_σ as a measure of precision, i.e., the reliability of the 2AFC procedure. Therefore, we retrieve a relative uncertainty $u = \frac{T_\sigma}{T}$ per estimated threshold, describing this reliability property. We convert the uncertainty u to a weight $w = \frac{1}{u}, w \leq 10$ to determine each threshold's influence on the final fit. Ultimately, w is large if T_σ is small, giving more weight to reliable, meaning less noisy, threshold estimations. Figure 5 displays the measured thresholds T, as well their weight w displayed in area. The diamonds indicate the weighted mean per blur intensity, retrieved from T and w.

Model Design and Fit We fit a function with four free parameters to our estimated means. The fitted function is a modified softplus function, where the parameters control displacement in 2D space (d_x, d_y), steepness (s), and softness of the knee $(k,\ k \leq 3)$:

$$M = s(\theta) \cdot \log\left[1 + e^{k(\theta)[\sigma - d_x(\theta)]}\right] + d_y(\theta). \tag{2}$$

We determine all four parameters for each of the three eccentricities. The resulting fits are displayed in Figure 5. We complete the model by interpolating the parameters across eccentricities:

$$\begin{aligned} d_x(\theta) &= 2.517 \cdot (\theta + 1)^{0.2501} \\ d_y(\theta) &= 0.08571 \cdot e^{0.1743\theta} + 0.9465 \\ s(\theta) &= 6.502 \cdot 10^{-05}\theta^2 - 2.367 \cdot 10^{-03}\theta + 0.07908 \\ k(\theta) &= 3 \end{aligned} \tag{3}$$

The interpolated parameters are shown in Figure 7. Our final model is described by substituting Equations 3 into Equation 2. A 3D representation of the surface is displayed in Figure 6.

Binocular Disparity vs. Motion Parallax

The direct comparability of the experimental paradigms enables a systematic analysis of how spatial blur affects depth perception from binocular disparity and motion parallax.

Quantitative Comparison Overall, depth from motion parallax requires substantially larger depth amplitudes to be resolved compared to binocular disparity. While stereoacuity thresholds are in the order of less than $1.5'$, the corresponding thresholds for motion parallax are around $2'$–$6'$, indicating that motion parallax provides a less precise relative depth signal under the tested conditions. These results suggest that binocular disparity provides more precise depth information, while motion parallax requires larger depth differences before being perceivable.

Robustness of Binocular Disparity Previous work has demonstrated a remarkable resilience of stereoacuity to blur, and therefore foveation [9]. As shown in Figure 6 (bottom), stereoacuity thresholds remain largely unaffected even at high blur levels. Notably, depth perception from binocular disparity remains stable well beyond the visibility threshold predicted by the CSF. In fact, blur can be increased to approximately $3\times$ the level at which it becomes perceptible before stereoacuity begins to degrade. This corresponds to a clearly visible reduction in spatial image quality, yet participants are still able to reliably resolve depth from disparity. This high tolerance is illustrated by the red line in Figure 6, which denotes the maximum blur intensity that does not significantly impair performance at each eccentricity. The curve extends to substantially higher blur levels than the CSF limit, particularly

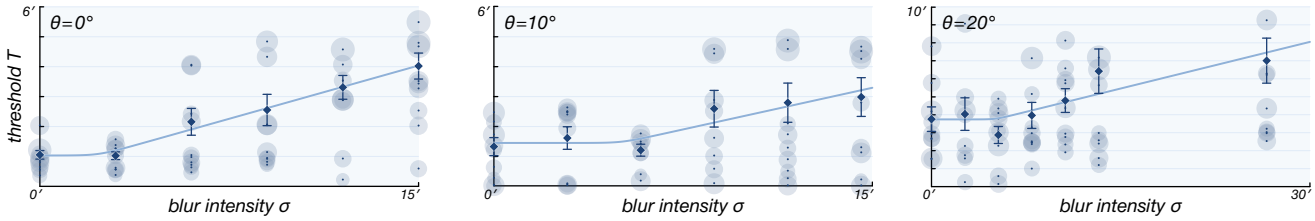

Figure 5. *Measured thresholds for the three examined eccentricities 0°, 10°, and 20°. The individual points display the bootstrapped mean thresholds per participant and condition, with the area denoting the weight w. The diamonds display the weighted average per blur intensity, as well as the standard error. The curves display the fitted functions as defined in Equation 2.*

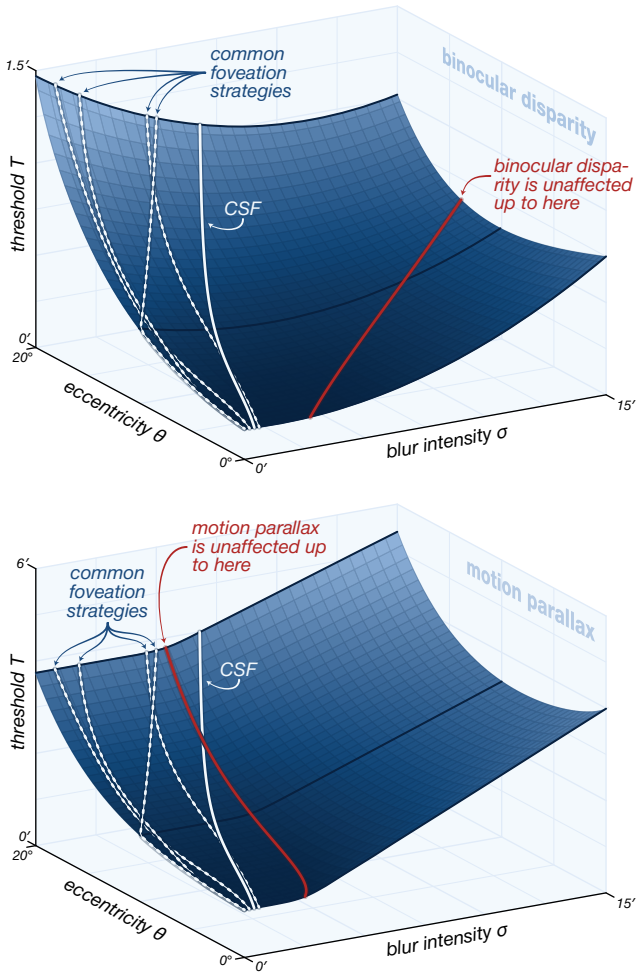

Figure 6. *Retrieved perceptual models of binocular disparity (top, [9]) and motion parallax (bottom). The red lines denote the maximum blur intensity per eccentricity, which does not harm depth perception from the particular cue. The solid white lines display the maximum blur per eccentricity, which is imperceivable as described by stelaCSF [14]. The dashed lines depict foveation strengths of common foveation techniques per eccentricity. Along the upper boundary of the surface, from left to right, the lines correspond to Thibos et al. (detectability limit) [25], Tursun et al. [26], Thibos et al. (resolvability limit) [25], Patney et al. [17]. The further the red line is from the white lines, the more robust the respective depth cue is to spatial blur.*

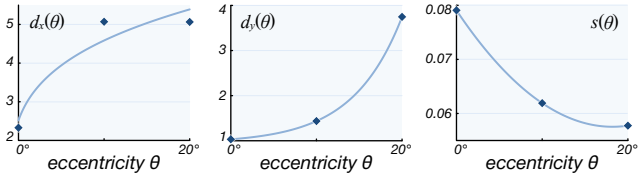

Figure 7. *Interpolation of the three function parameters of Equation 3 across eccentricities $\theta \in [0, 20]$. $k(\theta)$ is a constant with $k(\theta) = 3$.*

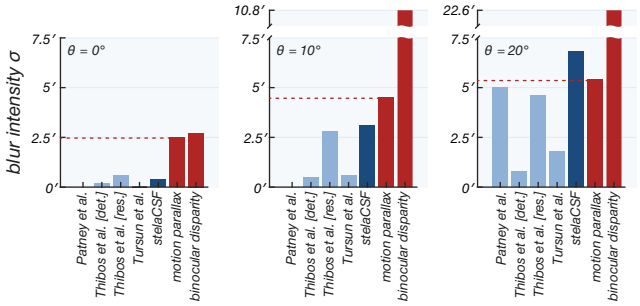

Figure 8. *Overview of blur intensities marked at the three eccentricities in Figure 6. The light blue bars depict foveation strengths of common foveation techniques per eccentricity [17, 25, 26]. The dark blue bars highlight the sensitivity cutoff as described by stelaCSF [14]. The red bars depict our estimated cutoff intensities for motion parallax, as well as the ones for binocular disparity [9]. Particularly at 10° and 20° eccentricity, motion-parallax thresholds are close to the CSF-defined visibility limits, whereas stereoacuity thresholds remain well above, highlighting the greater sensitivity of motion-parallax depth perception to peripheral blur.*

in the periphery, highlighting the robustness of disparity-based depth perception under strong foveation. In addition, Kergaßner et al. [9] observed that stereoacuity thresholds can initially decrease at low levels of blur before increasing at higher blur intensities. This behavior suggests that small amounts of blur may improve depth perception from binocular disparity. One possible explanation is that mild blur attenuates high-frequency components that can give rise to spatial aliasing, even beyond the aliasing limits described by Thibos et al. [25]. Such aliasing artifacts may particularly impair the stereo matching process, and their reduction could therefore facilitate more reliable disparity estimation. Although this effect is subtle, it highlights the complex interaction between spatial resolution and disparity processing.

Sensitivity of Motion Parallax In contrast, the thresholds measured for motion parallax do not exhibit such behavior: they show

neither an initial improvement at low blur levels nor a comparable robustness to visible blur. Instead, performance remains relatively stable for small to moderate blur levels and then deteriorates once the blur exceeds a certain magnitude. As shown in Figure 6, this transition occurs at blur levels that closely align with the visibility limit defined by the CSF. This indicates that the threshold for reliable depth perception from motion parallax is closely tied to the perceptual visibility of blur. In other words, as soon as the blur becomes perceptible, it begins to impair depth perception from motion parallax. Figure 8 further shows that the measured thresholds for motion parallax lie very close to blur levels typically employed by foveated rendering techniques, which are designed to remain below the perceptual visibility threshold. This leaves only a narrow margin: even a slight increase in blur beyond these levels would push the system into a regime where motion-parallax-based depth perception is degraded. In contrast to binocular disparity, which remains robust well beyond the visibility threshold, motion parallax therefore operates close to its functional limit under common foveation settings. This suggests that increasing blur beyond imperceptible levels can quickly lead to a breakdown of motion-based depth cues.

Conclusion

In this work, we investigate how reduced peripheral spatial resolution, as introduced by foveated rendering, affects the resolvability of depth cues. While previous approaches have largely focused on preserving spatial fidelity under eccentricity-based foveation, we show that perceptual consequences extend beyond spatial quality alone and can influence higher-level visual judgments. In particular, we measure depth perception from motion parallax under various conditions of peripheral image resolution, modeled by spatial blur, across the visual field. Subsequently, we relate our findings to previously reported measurements for binocular disparity. While previous work has shown that stereoacuity remains relatively robust under foveation, our results indicate that motion-dependent depth perception is more sensitive to similar degradations. We show that peripheral blur has little impact as long as it remains below the perceptual visibility threshold. Once this threshold is exceeded, however, motion-dependent depth cues degrade, highlighting the need for caution when applying more aggressive foveation.

Acknowledgments

We thank all voluntary participants for taking part in the user study. 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

[1] Rachel Albert, Anjul Patney, David Luebke, and Joohwan Kim. Latency Requirements for Foveated Rendering in Virtual Reality. *ACM Trans. Appl. Percept.*, 14(4), 2017.

[2] Christine A. Curcio, Kenneth R. Sloan, Robert E. Kalina, and Anita E. Hendrickson. Human Photoreceptor Topography. *Jour. Comparative Neurology*, 292(4):497–523, 1990.

[3] James E. Cutting and Peter M. Vishton. Perceiving Layout and Knowing Distances: The Integration, Relative Potency, and Contextual Use of Different Information about Depth. In *Perception of Space and Motion*, Handbook of Perception and Cognition, pages 69–117. Academic Press, 1995.

[4] Budmonde Duinkharjav, Kenneth Chen, Abhishek Tyagi, Jiayi He, Yuhao Zhu, and Qi Sun. Color-Perception-Guided Display Power Reduction for Virtual Reality. *ACM Trans. Graph.*, 41(6), 2022.

[5] Linus Franke, Laura Fink, Jana Martschinke, Kai Selgrad, and Marc Stamminger. Time-Warped Foveated Rendering for Virtual Reality Headsets. *Computer Graphics Forum*, 40(1):110–123, 2021.

[6] Brian Guenter, Mark Finch, Steven Drucker, Desney Tan, and John Snyder. Foveated 3D Graphics. *ACM Trans. Graph.*, 31(6):1–10, 2012.

[7] Petr Kellnhofer, Piotr Didyk, Tobias Ritschel, Belen Masia, Karol Myszkowski, and Hans-Peter Seidel. Motion Parallax in Stereo 3D: Model and Applications. *ACM Trans. Graph.*, 35(6):1–12, 2016.

[8] Sophie Kergaßner and Piotr Didyk. It's Not Just a Phase: Creating Phase-Aligned Peripheral Metamers, 2026. arXiv:2601.19425.

[9] Sophie Kergaßner, Taimoor Tariq, and Piotr Didyk. Towards Understanding Depth Perception in Foveated Rendering. In *Proc. SIGGRAPH Conference Papers*, pages 1–9, 2025.

[10] Brooke Krajancich, Petr Kellnhofer, and Gordon Wetzstein. Towards Attention–Aware Foveated Rendering. *ACM Trans. Graph.*, 42(4):1–10, 2023.

[11] Harris R. Lieberman and Alex P. Pentland. Microcomputer-Based Estimation of Psychophysical Thresholds: The Best PEST. *Behavior Res. Methods & Instrumentation*, 14(1):21–25, 1982.

[12] Sheng Liu and Hong Hua. Spatialchromatic Foveation for Gaze Contingent Displays. In *Proc. Symp. Eye Tracking Research & Applications*, ETRA, page 139–142. ACM, 2008.

[13] David Luebke and Benjamin Hallen. Perceptually Driven Simplification for Interactive Rendering. In Steven J. Gortler and Karol Myszkowski, editors, *Rendering Techniques*, pages 223–234. Springer Vienna, 2001.

[14] Rafał K. Mantiuk, Maliha Ashraf, and Alexandre Chapiro. stelaCSF: A Unified Model of Contrast Sensitivity as the Function of Spatio-Temporal Frequency, Eccentricity, Luminance and Area. *ACM Trans. Graph.*, 41(4):1–16, 2022.

[15] Bipul Mohanto, ABM Tariqul Islam, Enrico Gobbetti, and Oliver Staadt. An Integrative View of Foveated Rendering. *Computers & Graphics*, 102:474–501, 2022.

[16] Mark Nawrot and Lindsey Joyce. The Pursuit Theory of Motion Parallax. *Vision Res.*, 46(28):4709–4725, 2006.

[17] Anjul Patney, Marco Salvi, Joohwan Kim, Anton Kaplanyan, Chris Wyman, Nir Benty, David Luebke, and Aaron Lefohn. Towards Foveated Rendering for Gaze-Tracked Virtual Reality. *ACM Trans. Graph.*, 35(6):1–12, 2016.

[18] Alex Pentland. Maximum Likelihood Estimation: The Best PEST. *Perception & Psychophysics*, 28(4):377–379, 1980.

[19] Brian Rogers and Maureen Graham. Similarities Between Motion Parallax and Stereopsis in Human Depth Perception. *Vision Res.*, 22(2):1–10, 1982.

[20] Ruth Rosenholtz. Capabilities and Limitations of Peripheral Vision. *Ann. Rev. Vision Science*, 2(1):437–457, 2016.

[21] Luca Surace, Cara Tursun, Ufuk Celikcan, and Piotr Didyk. Gaze-Contingent Perceptual Level of Detail Prediction. In *Eurographics Symp. Rendering*, pages 1–12. Eurographics Assoc., 2023.

[22] Barbara Sweet and Mary Kaiser. *Depth Perception, Cueing, and Control*, pages 1–10. AIAA, 2012.

[23] Taimoor Tariq and Piotr Didyk. Towards Motion Metamers for Foveated Rendering. *ACM Trans. Graph.*, 43(4):1–10, 2024.

[24] Taimoor Tariq, Cara Tursun, and Piotr Didyk. Noise-Based Enhancement for Foveated Rendering. *ACM Trans. Graph.*, 41(4):1–14, 2022.

[25] Larry N. Thibos, D. J. Walsh, and F. E. Cheney. Vision Beyond the Resolution Limit: Aliasing in the Periphery. *Vision Res.*, 27(12):2193–2197, 1987.

[26] Okan Tarhan Tursun, Elena Arabadzhiyska-Koleva, Marek Wernikowski, Radosław Mantiuk, Hans-Peter Seidel, Karol Myszkowski, and Piotr Didyk. Luminance-Contrast-Aware Foveated Rendering. *ACM Trans. Graph.*, 38(4):1–14, 2019.

[27] David R. Walton, Rafael Kuffner Dos Anjos, Sebastian Friston, David Swapp, Kaan Akşit, Anthony Steed, and Tobias Ritschel. Beyond Blur: Real-Time Ventral Metamers for Foveated Rendering. *ACM Trans. Graph.*, 40(4):1–14, 2021.