Foveated Rendering: Motivation, Taxonomy, and Research Directions

SUSMIJA JABBIREDDY, University of Maryland, College Park, USA XUETONG SUN, University of Maryland, College Park, USA XIAOXU MENG, University of Maryland, College Park, USA AMITABH VARSHNEY, University of Maryland, College Park, USA

With the recent interest in virtual reality and augmented reality, there is a newfound demand for displays that can provide high resolution with a wide field of view (FOV). However, such displays incur significantly higher costs for rendering the larger number of pixels. This poses the challenge of rendering realistic real-time images that have a wide FOV and high resolution using limited computing resources. The human visual system does not need every pixel to be rendered at a uniformly high quality. Foveated rendering methods provide perceptually high-quality images while reducing computational workload and are becoming a crucial component for large-scale rendering. In this paper, we present key motivations, research directions, and challenges for leveraging the limitations of the human visual system as they relate to foveated rendering. We provide a taxonomy to compare and contrast various foveated techniques based on key factors. We also review aliasing artifacts arising due to foveation methods and discuss several approaches that attempt to mitigate such effects. Finally, we present several open problems and possible future research directions that can further reduce computational costs while generating perceptually high-quality renderings.

CCS Concepts: • General and reference \rightarrow Surveys and overviews; • Computing methodologies \rightarrow Mixed / augmented reality; Rendering; Antialiasing; Perception; Ray tracing.

Additional Key Words and Phrases: foveated rendering, foveated displays, eye-tracking, gaze-contingent rendering, variable-rate sampling, head mounted displays

1 INTRODUCTION

Virtual and augmented reality (VR and AR) are poised to transform computer-mediated communication spanning an exciting range of applications in science, engineering, medicine, arts, entertainment, commerce, and several other areas. While impressive, the current-generation VR and AR systems are unable to match the visual fidelity of our real-world experiences along several dimensions, including resolution, the field of view, and latency. Delivering such a high-quality visual experience in real-time requires enormous computational resources beyond the recent improvements in the hardware and software systems for rendering.

A very small fraction of the scene that is projected on the fovea, the center of the human retina, is perceived by the human visual system (HVS) at its finest details. Visual acuity is the ability to observe detail and is measured as the grating resolution in cycles/degrees. Studies show that though HVS has a wide field of view (FOV), the region that has the highest visual acuity, also known as the foveal region, covers only the central 2.5° of the visual field [Levin et al. 2011]. Foveation refers to a decrease in the acuity with angular distance in the human visual system [Guenter et al. 2012]. Foveated rendering leverages this feature of the human visual system to selectively render a small fraction of the graphics frame in fine detail.

MOST VR PIXELS ARE PERIPHERAL



Fig. 1. Percentage of pixels that lie in the fovea region when viewed on an iPhone 7 plus screen, 27" desktop monitor and a VR head-mounted display. Image adapted from Patney *et al.* [2016b].

Foveated rendering allows us to reconcile the mutually conflicting goals of high visual realism, interactive frame rates, and low power consumption on modern VR and AR devices. As foveated rendering becomes an integral part of future AR/VR devices, it is important to understand the current state and the observed trends in the field.

The concept of foveated displays is not new. Early implementations of the gaze-contingent displays find applications in building flight simulators [Murphy and Duchowski 2001]. However, early methods were developed mainly for desktop monitors, where the FOV was approximately $43^{\circ} \times 33^{\circ}$ [Murphy and Duchowski 2001]. This is a fraction of modern wide FOV displays. Patney et al. [2016b] give a comparison of the percentage of pixels that lie in the peripheral region across various devices ranging from a small FOV smartphone screen to a wide FOV VR HMD, as shown in Figure 1. Thus, foveated rendering becomes even more important for VR and AR headsets compared with traditional displays. In a 100° wide FOV HMD, only 4% of the screen pixels lie in the foveal region while the rest lie in the peripheral region [Patney et al. 2016b]. When such a large percentage of pixels lie in the low acuity peripheral region, efficient rendering systems can obtain significant computational gains by allocating fewer resources to render the peripheral region.

Until recently, real-time eye-tracking was not commonly available in VR and AR headsets. However, with the commoditization of eyetrackers in VR and AR headsets, foveated rendering techniques can use the gaze information to dynamically render the foveal region at a higher quality and the rest at a lower quality without a drop in the overall perceived quality of the image. As the FOV of future HMDs increases to match that of the human vision, foveated rendering will emerge as an essential component of the real-time rendering 3D graphics pipeline for large-scale, wide FOV and high-resolution VR and AR displays. 4D light fields is another area that can benefit from

Authors' addresses: Susmija Jabbireddy, University of Maryland, College Park, College Park, USA, jsreddy@umd.com; Xuetong Sun, University of Maryland, College Park, College Park, USA; Xiaoxu Meng, University of Maryland, College Park, College Park, USA; Amitabh Varshney, University of Maryland, College Park, College Park, USA.

2 . Susmija Jabbireddy, Xuetong Sun, Xiaoxu Meng, and Amitabh Varshney

the foveated rendering techniques. 4D light fields represent a scene from multiple camera positions, which incur a huge rendering cost. Sun *et al.* [2017] and Meng *et al.* [2020b] present methods that use foveated rendering techniques to accelerate the rendering process of light fields, thus enabling real-time light field rendering.

Recently, a classification of foveated displays has been presented by Spjut *et al.* [2020]. Their classification categorizes foveated rendering systems based on the resolution distribution function and gaze contingency. The resolution distribution function measures how well the system's acuity distribution matches the pre-measured visual acuity for the user (as measured by the Snellen eye chart). Gaze contingency measures how the system adapts to changes in the gaze direction. This paper focuses on foveated rendering methods for VR/AR systems. We consider different attributes that characterize foveated rendering, including several visual factors, resolution distribution, foveation space, gaze information, and anti-aliasing response. Further, we have classified several influential research papers along these dimensions. A similar review is also provided in a concurrent work by Mohanto *et al.* [Mohanto *et al.* 2022] with a different taxonomy of foveated rendering methods.

We organize our paper as follows. Section 2 provides a brief overview of the human visual system and its limitations that enable foveated rendering. A taxonomy of various foveated rendering methods based on their salient characteristics is presented in Section 3. Section 4 discusses the methods used to evaluate the perceived quality of the foveated image. Finally, some concluding remarks along with the challenges and possible future research directions are presented in Section 5.

2 A BRIEF OVERVIEW OF HUMAN VISUAL SYSTEM

Understanding human perception and its limitations is beneficial to improve the rendering quality and efficiency. The eye and the brain make up the human visual system. While the eye serves as a camera, the brain is responsible for processing the information. When a light ray hits the eye, it first passes through the cornea and undergoes refraction. The refracted ray passes through the aqueous humor, the iris, the lens, the vitreous humor, and finally, reaches the retina - the image sensor of the human eye. Figure 2 shows the complete anatomy of the human eye. The light photons reaching the retina are then detected and converted to electrical signals by photoreceptors. The information from the photoreceptors is transmitted to the central neural system by ganglion cells for further processing.

It is well known that not all the information that enters the eye reaches the brain. A significant amount of compression occurs at the retinal level. Several components are responsible for this preprocessing. Firstly, the distribution of photoreceptors and the ganglion cells is not uniform across the retina. There are two categories of photoreceptors, namely rods and cones. The rods and cones differ in characteristics such as shape, distribution across the retina, patterns of synaptic connections, among many others [Purves et al. 2001]. We have 100's of millions of rods, while the number of cones is around six-seven million [Taylor and Batey 2012]. Moreover, a large portion of the cones is concentrated near the central part of the retina, known as the fovea, and their density greatly decreases



Fig. 2. Anatomy of the human eye. Image from Urone *et al.* [2012], reproduced under Creative Commons Attribution License.

as the distance from the center increases. On the other hand, rods are spread out across the entire retina and are absent at the fovea. Figure 3 shows the density distribution of rods and cones across the visual field. In addition, rods are highly sensitive to light, whereas cones provide excellent spatial discrimination and are responsible for color vision. The differences in the rod and cone systems characterize three modes of vision based on the amount of light that reaches the human eye. The first one is scotopic or night vision, which facilitates perception in dim lighting. In this type of vision, only rods are active. Rods cannot discriminate colors, and therefore, scotopic vision is a grayscale vision. The second type of vision is photopic vision, or daylight vision, which provides for color perception mediated by the cone cells. Mesopic vision, the third type of vision, is a combination of photopic and scotopic vision and is perceived in low but not quite dark lighting conditions. Both rods and cones are simultaneously active for this vision. Most indoor and outdoor lighting situations use photopic vision, where cones are the most active ones. However, if one wishes to render dim lighting, the properties of rods are to be considered. Therefore, the perceived level of detail at any point in time also depends on the active mode of vision.

The ganglion cells in the inner part of the retina transmit the information from the photoreceptors to the neural processing system. While we have 100's of millions of photoreceptors, the number of ganglion cells present in the retina is only around 1.2 million [Kim et al. 2021]. This many-to-one connection further compresses the information reaching the brain. Studies [Wässle et al. 1990] observed that the ganglion cell density is also non-uniform across the retina. The number of photoreceptors connected to a single ganglion cell increases with an increase in the distance from the fovea.

The spatially varying photoreceptor and the ganglion cell densities, with a higher density at the center, lead to foveated vision in the human visual system. As a result, the visual sensitivity is maximum at the fovea and reduces towards the periphery. The exact transition point from the foveal to the peripheral region is not defined uniformly across various disciplines. The eccentricity of a point in the retina is defined as the angular distance of that point from the



Fig. 3. Density distribution of photoreceptors (rods and cones) with eccentricity in the human eye. Image adapted from Mustafi *et al.* [2012]

fovea. The fovea is circumscribed by para-fovea, which is circumscribed by peri-fovea. Based on these regions, the human visual field can be broadly divided into central and peripheral vision. Neuropsychologists typically define central vision as the region from 0° to $2.6^{\circ} - 3.6^{\circ}$ eccentricity, with everything beyond that considered as peripheral vision [Loschky et al. 2019]. Visual cognition researchers generally consider the region from 0° to 5° eccentricity as the central vision, comprising the foveal region from 0° to approximately $0.5^{\circ} - 1^{\circ}$ eccentricity and para-fovea from approximately $0.5^{\circ} - 1^{\circ}$ to 5° and anything beyond 5° is considered peripheral vision. Vision science researchers consider the central vision to extend from 0° to 10° with the peripheral vision beyond that [Loschky et al. 2019]. Whatever be the exact definition, the vast majority of the human visual field falls in peripheral vision.

Foveated rendering exploits the foveation in the human visual system and renders the peripheral region at a relatively lower quality. For a more comprehensive survey on the physiological and perceptual aspects of human vision and limitations of the human visual perception, readers are referred to the work by Weier *et al.* [2017].

3 A TAXONOMY OF FOVEATED RENDERING SYSTEMS

The non-uniform distribution of rods and cones across the retina, along with their varying responses to the incoming light, leads to several interesting properties that change from the foveal to the peripheral vision. For instance, perception of detail reduces with an increase in eccentricity and the time to perceive a same-sized object increases with eccentricity [Sun and Varshney 2018]. Rods have limited spatial discrimination while cones have excellent spatial localization and color discrimination [Purves et al. 2001]. However, it is interesting that the ability to detect motion in the peripheral region is higher than what would be expected. Moreover, the ability to differentiate an object's relative velocity and flicker sensitivity stays uniform across the visual field. In addition, Patney et al. [2016b] show that preserving contrast in the peripheral region improves image perception for filtered images. Studies also show that while color perception decreases with an increase in eccentricity, the decrease is gradual and color discrimination survives at high eccentricities [Hansen et al. 2009a]. Rosenholtz [Rosenholtz 2016] argues that peripheral vision is not a mere low-resolution version of

foveal vision, but instead is responsible for a texture like representation that preserves the summary statistics of a scene. The different properties that are needed to be preserved in the peripheral region depends on the task at hand. Foveated rendering can leverage these human vision characteristics that vary with eccentricity.

While most real-time rendering uses the rasterization pipeline, ray-tracing provides more realistic images with accurate reflections, refractions, and shadows and is gradually eliciting greater interest [Ludvigsen and Elster 2010]. Foveation can be applied to both rasterization and ray-tracing pipelines. The graphics pipeline for real-time rasterization is often dominated by the computational complexity of shading operations. Shading costs can be reduced by simplifying the operation, pre-computing parts of the operation, sharing results across a set of fragments, and/or reducing the number of operations. Foveated rendering attempts to reduce the average number of operations per pixel by leveraging the non-uniform visual sensitivity across the visual field. We next categorize and compare foveated rendering techniques based on their characteristics and discuss their use in efficient real-time rendering. Figure 4 gives an overview of the presented taxonomy.

Section 3.1 presents various visual attributes of the human visual system that spatially vary across the visual field. Section 3.2 discusses various models representing the spatial variance of visual sensitivity, and Section 3.3 considers the stage at which foveation is implemented in foveated rendering systems. Section 3.4 gives an overview of the techniques based on their relation to gaze tracking. Finally, Section 3.5 discusses anti-aliasing techniques for foveated rendering.

3.1 Visual Factors

Visual perception depends on various factors including, but not limited to, eccentricity, contrast, luminance, brightness, color, shape, motion, fixations, and saliency. In this section, we survey the foveated rendering methods based on the different visual factors considered.

3.1.1 Content-independent factors. One widely used visual measure is visual acuity, defined as the reciprocal of the angle (in minutes) subtended by a just resolvable region at a given eccentricity. Visual acuity varies across the visual field. Early works [Anstis 1974] show that acuity drops with an increase in eccentricity. Therefore, reducing the resolution as a function of angular distance from the gaze point has always characterized foveated rendering. One of the distinguishing features among the various foveated rendering approaches has been the way this reduction is accomplished. One common way is to render multiple discrete layers independently; each sampled at a different rate [2012; 2016]. These layers, called eccentricity layers, are later upsampled to the screen resolution and composed together. As shading is the most time-consuming operation in the rendering pipeline, certain works [2014; 2014; [n. d.]] vary the shading rate across the visual field, shading only once per several pixels in the peripheral region. Some recent works [Koskela et al. 2019b; Meng et al. 2018; Tursun et al. 2019] also demonstrate varying resolution in a more continuous manner. These methods map the visual field to a non-linear space where a uniform rendering matches the visual acuity of human vision. While the above methods are mainly developed using a rasterization pipeline, ray 4 . Susmija Jabbireddy, Xuetong Sun, Xiaoxu Meng, and Amitabh Varshney



Fig. 4. Taxonomy of the foveated rendering methods based on some of the key factors. A full foveated rendering pipeline might choose any sequence of options from these categories.

tracing methods naturally allow for smooth non-uniform sampling in the screen space. Weier *et al.* [2016] and Fujita and Harada [2014] incorporate foveated rendering into ray-tracing and achieve sparse sampling by varying the sampling probability with an increase in distance from the gaze point. For more discussion on the trade-offs between discrete and continuous distributions, we refer the readers to the recent work on foveation displays [Spjut and Boudaoud 2019]. The idea behind these methods is to mimic a distribution that closely aligns with the visual acuity of the HVS. Different methods approximate the non-uniform sampling distribution function differently to best match acuity across the visual field; we discuss this in detail in Section 3.2.

Furthermore, recent studies [Abrams et al. 2012; Barbot et al. 2021] show that the visual acuity is higher across the horizontal axis than the vertical axis at the same eccentricity values. Studies further show that the lower visual field has better visual acuity. Most of the current foveated rendering techniques do not factor this asymmetry during rendering and assume uniform acuity across horizontal and vertical directions. One exception is the recent method by Ye *et al.* [2022] that processes the horizontal and vertical visual fields independently, achieving a superior image quality.

3.1.2 Content dependent factors.

Contrast. Image contrast is another major factor impacting visual perception. The minimum amount of contrast (contrast threshold) required to detect a pattern depends on its spatial frequency, in addition to the content-independent eccentricity. The reciprocal of contrast threshold is called contrast sensitivity, and it is characterized by the contrast sensitivity function (CSF) across the human visual field. The function describes the capacity of the HVS to recognize differences in patterns as a function of spatial frequencies [Bull 2014]. Similar to spatial acuity, contrast sensitivity for a given spatial frequency is highest at the fovea and decreases with increasing eccentricity. The contrast sensitivity depends on several factors such as size, contrast, and the viewing angle of the target pattern,

and several functions for modeling CSF have been proposed in the literature [Barten 2003; Pointer and Hess 1989].



Fig. 5. An illustration of the presence of aliasing zone in human peripheral vision. Bandlimiting filter response is plotted on the y-axis against the angular frequency on the x-axis, normalized to the detection cut-off frequency. Aliasing zone represents the range of frequencies that lie between the discrimination and detection cut-off frequencies. Image adapted from Patney *et al.* [2016b].

The relationship between contrast sensitivity and spatial frequency varies across the visual field [Thibos et al. 1996]. In general, the foveal CSF shows a gradual loss in contrast sensitivity with increasing spatial frequency, while peripheral CSF shows that contrast sensitivity drops abruptly. The contrast sensitivity function has a cut-off spatial frequency beyond which a higher frequency is not discernible. The cut-off spatial frequency depends on the eccentricity of the image patch location in the visual field. Further, the contrast sensitivity function is task-dependent. The peripheral CSFs for detection and discrimination tasks indicate different limits to performance at high spatial frequency [Anderson et al. 2002]. The peripheral contrast sensitivity for the discrimination task is observed to degrade at a much faster rate than for the detection task, as demonstrated in Figure 5. The difference between the cutoff frequencies for the detection task and the discrimination task indicates that the human visual system can perceive higher frequencies in the periphery but cannot resolve details. The range of frequencies between the detection and discrimination cut-off frequencies is known as the aliasing zone. The foveated renderers that determine the sampling rate based on the discrimination cut-off frequency provide faster performance than those based on the detection cut-off frequency. However, rendering the peripheral region based on the discrimination cut-off frequency reduces contrast for regions whose frequencies are in the aliasing zone. To maintain the detection ability, Patney et al. [2016a; 2016b] suggest enhancing peripheral contrast in the aliasing zone, which was degraded by filtering to maintain the perceptual quality of a non-foveated image. Tursun et al. [2019] exploits the influence of luminance contrast on visual perception. They observe that while a given reduction of spatial resolution in high-contrast regions reduces the perceived quality, the same reduction in low-contrast regions maintains it. The minimal required resolution for each region is estimated using the local luminance contrast and the angular distance from the gaze point. In addition to luminance, the sensitivity to color also reduces towards the peripheral region. The hue resolution or the number of gray levels within each RGB channel that can be perceived reduces with eccentricity [Hansen et al. 2009b]. Liu et al. [2008] show that the number of bits representing each color channel can be monotonically reduced from 8 to 4 as the eccentricity increases from 0° to 30°.

Saliency. Visual saliency, often a good indicator of visual attention, is another factor influencing the perceived level of detail. Visual saliency is a distinct subjective quality where certain regions of a scene are more likely to draw the viewer's attention when compared to their surrounding regions. The concept of foveal and attentional spotlights states that the attention and gaze do not necessarily coincide [Levin et al. 2011]. As the user's attention is automatically attracted to the visually salient stimuli, identifying such salient regions can redistribute the rendering cost by allocating more resources to the visually important regions. Many computational models have been introduced to estimate visual saliency [Kim et al. 2010; Lee et al. 2005; Song et al. 2014; Yohanandan et al. 2018]. In addition to the low-level features like local color and brightness contrast, high-level features such as the position and identity of the objects and the scene context have a significant influence on the visual attention model. The set of visually salient regions is not uniquely considered across various research works. For example, the adaptive multi-rate shading method by He et al. [2014] considers regions near object silhouettes, shadow edges, and regions of potential specular highlights as visually salient. A higher sampling rate is used to render these regions so that they are perceived at relatively higher acuity. On the other hand, Stengel et al. [2016] consider the regions of high spatial and temporal contrast and saturated colors as visually significant to develop a foveated rendering system.

3.2 Eccentricity-based analytical visual models

Foveated rendering systems vary resolution or sampling rate across the visual field based on one or more visual factors discussed above. Based on the scene content and task, these systems approximate the human visual field using analytical expressions. These approximations allow different kinds of pixel distributions. In this section, we discuss the similarities and differences among the different distributions.

3.2.1 *Hyperbolic model.* Visual acuity can be quantitatively represented as the reciprocal of the minimum angle of resolution [Weymouth 1958], which is the smallest angle at which two points are perceived as different. The minimum angle of resolution increases linearly with an increase in eccentricity. One possible mathematical characterization of visual acuity over eccentricity is given by

Acuity =
$$\frac{1}{me + \omega_0}$$
 (1)

where, ω_0 corresponds to the smallest resolvable angular resolution that occurs at the fovea (e = 0) and m represents the slope. We refer to this distribution as the hyperbolic model where the acuity changes as a function of 1/eccentricity. This hyperbolic model serves as a good approximation for visual acuity at low eccentricities (less than 8° angular radii) [Guenter et al. 2012], after which the acuity drops more steeply. Many foveated renderers [Guenter et al. 2012; He et al. 2014; Patney et al. 2016b; Stengel et al. 2016; Swafford et al. 2016; Vaidyanathan et al. 2014] base their sampling distribution on this hyperbolic fall-off of acuity.

Guenter et al.'s [2012] pioneering work shows that significant performance improvement is possible with foveated rendering for rasterization. Assuming that a few discrete layers are sufficient to model the acuity fall-off in the human visual system, they use three nested and overlapping rectangular layers rendered at different resolutions, as shown in Figure 6. These three layers, known as eccentricity layers, are centered at the gaze point or the fovea. The innermost, foveal layer, is rendered at the highest resolution (which is the native display resolution). The middle layer is larger than the inner layer and is rendered at a relatively lower resolution. The outermost layer covers the entire screen and is rendered at a much lower resolution. All three layers are then interpolated to the display resolution and blended smoothly. The hyperbolic function in Equation 1 is used to compute the size and resolution for the eccentricity layers. The parameters depend on the minimal angular resolution slope *m*, which is estimated from the user studies. The approach assumes symmetric radial acuity fall-off, ignoring the differences in the horizontal and vertical axes of the visual field. They limit their method to three layers to approximate the human visual system, although more layers can give a better approximation at the cost of increased complexity. Thus, the the total number of rendered pixels is reduced compared to rendering the entire frame at a uniformly high resolution.

Stengel *et al.* [2016] extend the above model to incorporate the effect of smooth-pursuit eye movements. This type of eye movement is triggered unconsciously when a moving object attracts attention. They model the region of focus as a straight line, obtained as integration of gaze positions over consecutive frames. The eccentricity is now measured based on the distance to this line. Consequently, the rendering at the highest-resolution is over a larger elliptical region enveloping the straight line, instead of a circular region around a

6 • Susmija Jabbireddy, Xuetong Sun, Xiaoxu Meng, and Amitabh Varshney



Fig. 6. Three discrete layers are used to match the human acuity across the visual field. The three layers are rendered at different resolutions based on the distance to the gaze point (red dot). Image adapted from Guenter *et al.* [2012].

single point. This method gives superior performance on higher latency displays.

Spjut *et al.* [2020] characterize the acuity distribution of HVS using a hyperbolic model and evaluate a foveated display based on how well it matches this distribution. For natural viewing of images, the hyperbolic model has proven to be accurate for small eccentricity values [Guenter et al. 2012; Stengel et al. 2016]. However, often, the region beyond 30° eccentricity is rendered at a uniform lower resolution, which potentially limits the extent of speedup that can be achieved through foveated rendering.

3.2.2 Linear model. Weier *et al.* [2016] consider a linear fall-off of acuity with increasing eccentricity as opposed to the above hyperbolic fall-off. By modeling a linear model into the ray-tracing pipeline, the sampling probability reduces with an increase in distance from the gaze point. The visual field is divided into three regions, similar to Guenter *et al.* [2012]. Each region is characterized by three parameters: r_0, r_1, p_{min} as shown in Figure 7. The inner region (the region which falls within r_0 degrees) is the foveal area, which is rendered at full resolution and so is sampled with a probability of one. The pixels in the outermost peripheral region (the region that lies beyond r_1 degrees from the gaze point) are sampled with a minimum probability of p_{min} . The pixels in the region between the two layers (between r_0 and r_1) are sampled according to

the linear equation:

$$p(x) = 1 - (1 - p_{min})\frac{d(x) - r_0}{r_1 - r_0},$$
(2)

where d(x) is the distance of the pixel x from the center of the visual field in degrees. r_0 , r_1 , and p_{min} are user-defined parameters. The linear approximation model is more relaxed and holds good only for small eccentricity values, beyond which the region is rendered at a uniform lower resolution. However, the linear fall-off maintains motion perception in the periphery better than the model with hyperbolic fall-off due to the increasing sampling rate.



Fig. 7. Linear acuity model used for sampling pixels in the region between r_0 and r_1 . Image adapted from Weier *et al.* [2016].

3.2.3 Logarithmic model. In the human visual system, a log-polar mapping of the eye's retinal image approximates the excitation process of the visual cortex [Araujo and Dias 1996]. This log-polar mapping ensures that the sensitivity to perceive fine details is high at the center of the visual field and decreases logarithmically with an increase in distance from the fovea. Contrasting the above models that use multi-resolution rates during rendering, a mapping from Cartesian space to log-polar space that matches human visual acuity can use a single uniform resolution rendering, applied directly to the transformed space. Log-polar mapping has been proven useful in many areas like computer vision, robotics, computer graphics, and image processing because of its power to provide enough visual detail using limited computational resources [Antonelli et al. 2015].

Meng *et al.* [2018] provide a foveated rendering method for meshes using the log-polar mapping of the human visual system. Their system introduces a kernel log-polar mapping technique that offers the flexibility to model acuity fall-off that matches HVS. Figure 8 shows the mapping from the Cartesian space to the log-polar space. The acuity decreases logarithmically and depends on the kernel. The rendering acceleration in this technique is achieved through deferred shading [Deering et al. 1988], a widely-used technique in real-time rendering. The information about the positions, normals, textures, and materials for each surface necessary for shading computations is rendered into the geometry buffer (G-buffer). The contents of the G-buffer are transformed from the Cartesian space to the log-polar space. The direct and indirect lighting at each pixel is computed and rendered to the reduced resolution log-polar buffer. Lighting calculations are performed in the reduced log-polar space. Inverse kernel log-polar mapping is applied to map the shading back to the Cartesian screen space. This method is able to systematically vary the sampling rate and sampling distribution continuously in the log-polar space.



Fig. 8. Kernel Log-polar mapping to match the human acuity fall-off. The gaze point is at the center of the image. (a) is the image in Cartesian coordinates. (b) is the corresponding image in log-polar coordinates. The region close to the gaze point occupies the major portion in the log-polar space. Image adapted from Meng *et al.* [2018].

Koskela et al. [2019b] introduce a similar mapping from Cartesian space to visual-polar space. The path tracing with one sample per pixel is performed in the visual-polar space. The polar-coordinate space is modified so that the sampling distribution aligns with the human visual acuity distribution. To match HVS, the number of samples along the angular and the radial axes can be adjusted. They observe that varying the number of samples along the angular axis results in peripheral region artifacts while varying the samples along the radial axis leads to foveal region artifacts. Based on these findings, their optimization technique varies the resolution along the angular axis in the fovea and rescales the radial axis in the peripheral region. Denoising is applied to the noisy path-traced visual-polar space image [Koskela et al. 2019a,b]. The reconstructed visual-polar space image is transformed back to the Cartesian coordinates using the inverse mapping. They report that the visual-polar mapping reduces distracting artifacts compared to the log-polar mapping.

3.2.4 Other Approximations : Conformal rendering [Bastani et al. 2017] models visual acuity as a non-linear function of eccentricity. The technique aims to mimic the smooth transition from the foveal to the peripheral region using a non-linear mapping of the screen distance from the eye-gaze point. The projected vertices of the virtual scene are warped into a non-linear space that matches the retinal acuity and the HMD lens characteristics. The warped image is then rasterized at a reduced resolution and unwarped back into Cartesian space. The complexity of conformal foveated rendering depends on the complexity of the scene. As the number of vertices in a scene increases, the performance reduces. In contrast to the methods that use discrete layers to model the acuity fall-off [Guenter et al. 2012; Weier et al. 2016], the methods by Meng *et al.* [2018] and Bastani *et al.* [2017] simulate the acuity fall-off from foveal to the peripheral region using a continuous smooth function with



Fig. 9. Various approximate analytical models are used to vary the spatial resolution across the visual field. We plot the resolution factor as a function of distance from the center of the visual field following (a) hyperbolic fall-off [Guenter et al. 2012], (b) kernel log-polar fall-off [Meng et al. 2018], (c) inverse quadratic fall-off [Reddy 2001] and (d) linear fall-off [Weier et al. 2016].

non-linear mapping. This smooth transition helps in reducing visual artifacts.

Reddy [2001] proposes a visual acuity model as a function of the angular velocity of a stimulus projected onto the retina and the eccentricity. The visual acuity varies as an inverse quadratic function of eccentricity [Reddy 1998]. Zheng et al. [2018] develop a foveated rendering method based on the visual acuity model by Reddy [2001] and adjust the tessellation levels accordingly. Friston et al. [2019] use a simple radial power-falloff function $p(x) = f(x^2)$ that maps the distance from the gaze point to a foveated distance. The pixel locations are then scaled based on this foveated distance. The overall effect magnifies the region close to the gaze point, giving more importance to the foveal region. They assume that the foveation function is arbitrary and free to change every frame, but require the function to be invertible to map the foveated image back to an unfoveated one for display. Fujita and Harada et al. [2014] assume that the acuity drops off as a function of $d(x)^{-2/3}$, where d(x) is the distance from the gaze point and define a sampling distribution function correspondingly. Recently, Li et al. [2021] presented a logrectilinear mapping-based foveated rendering to model exponential decay in resolution with an increase in eccentricity.

The various perceptual models described above mimic the visual sensitivity of the HVS across the visual field. They differ mainly in their motivation, complexity, and distribute samples with subtle differences. The most commonly used hyperbolic model mimics the very early operation of retinal projection in the eye, whereas logarithmic model represents the visual cortex part of HVS. Computational complexity is a major factor for practical purposes, especially when real-time rendering is required. The least complex, linear model, results in a higher resolution level in the peripheral region compared to the other models. The logarithmic model offers a more tunable fall-off distribution.

8 • Susmija Jabbireddy, Xuetong Sun, Xiaoxu Meng, and Amitabh Varshney

Distribution model	Function governing change of resolution
Hyperbolic	1/e
Linear	-е
Logarithmic	1/log(e)
Inverse Quadratic	$1/e^{-2/3}$

Table 1. Various models characterizing the degradation of visual sensitivity as a function of eccentricity e

Table 1 shows how the visual acuity degrades as a function of eccentricity e for each model. Futhermore, Figure 9 shows the resolution factor distributions across the visual field. The parameters are chosen from the corresponding papers. The models mainly differ in the resolution requirements in the near-peripheral region. The hyperbolic function offers a drastic resolution reduction, followed by kernel log-polar, inverse quadratic, and linear fall-offs. The kernel log-polar mapping demands higher resolution in the near periphery compared to the hyperbolic fall-off. In the far periphery, all the models converge to similar resolutions. One can choose a resolution distribution function from many of these approximations based on the scene complexity and the task to be performed. Moreover, we have seen that several visual factors impact the perceptual quality of a scene. Most of the existing models consider only a few of these visual factors. There has not been much research thus far in weighing these factors and their inter-relationships to develop an overall perceptual model. Furthermore, the visual sensitivity can change dynamically based on the external stimuli presented to the user, further increasing the complexity of the required model. Techniques for simplifying the complex models that consider the various inter-relationships between the visual factors without affecting the perceived visual quality or performance are therefore important.

3.3 Foveation Space

We classify foveated rendering techniques into screen-based, objectbased, or optics-based methods based on the stage of the rendering pipeline that incorporates the concept of foveation: the screen space, the object space, or the optical space (Figure 10).

3.3.1 Screen-based Methods. Most foveated rendering techniques vary the sampling rate in the screen space based on the distance from the point of focus or the gaze point. Screen-based foveated rendering approaches [Guenter et al. 2012; He et al. 2014; Meng et al. 2018; Patney et al. 2016b; Vaidyanathan et al. 2014] involve manipulating the frame-buffer contents just prior to the display to reduce the overall shading rate. Vaidyanathan et al. [2014] provide an architecture called coarse pixel shading that enables sparse shading operations for the peripheral region in screen space. Swafford et al. [2016] study the effect of foveation on ambient occlusion which is associated with a high computational cost on modern real-time rendering pipelines. Specifically, they present the effect of varying per-pixel depth samples in foveal and peripheral regions for Screen-Space Ambient Occlusion, a technique which approximates ambient occlusion. They show that the banding effect arising from a low number of per-pixel samples is imperceptible in the peripheral region due to the reduced acuity, thus improving performance. Another foveated renderer that makes changes in the screen space is conformal rendering [Bastani et al. 2017] in which the projected vertices are warped to a non-linear space before rasterization. The image is then rasterized at a reduced resolution and finally unwarped into the Cartesian space.

3.3.2 Object-based Methods. Object-based foveated rendering methods involve manipulating the model geometry prior to rendering. One of the very early works on foveated rendering by Levoy and Whitaker [1990] is an object-based approach. Following a raytracing approach for volume rendering, the number of rays cast per unit area and the number of samples per unit length along each ray are changed based on the pixel's angular distance from the gaze direction. The system prefilters the 3D volume using a 3D MIPmap and uses fewer samples in the peripheral region. Several early works [Danforth et al. 2000; Luebke and Hallen 2001; Murphy and Duchowski 2001; Ohshima et al. 1996] also model objects based on the gaze direction. The desired level of detail for the object is determined based on its distance from the gaze point. Swafford et al. [2016] have developed another object-space technique. They vary tessellation levels in the foveal and peripheral regions. A higher tessellation factor is used for the tiles that will fall within the foveal area after screen-space projection. For those that fall in the peripheral region, a lower tessellation factor is used. Linear interpolation determines the level of tessellation for the tiles that fall in the region between the foveal and peripheral areas. Their results show performance gain compared to rendering the entire scene at a uniform tesselation. Along similar lines to Swafford et al. [2016], Zheng et al. [2018] propose a method to tessellate based on visual sensitivity at a given eccentricity. They suggest removing all the imperceivable polygons and then dynamically tessellating the model. The polygons whose edge length in the screen space is less than the user's minimum perceptible length are not further sub-divided. The tessellation level for the perceivable polygons is given by the ratio of the edge length of the polygon in the screen space to the minimum perceptible length.



Fig. 10. Foveation can be applied at any of the three stages - (a) Optics space, (b) Screen space and (c) Object space.

3.3.3 Optics-based Methods. Optics-based foveated displays involve optically manipulating light from one or more displays [Kim et al. 2019a; Lee et al. 2020; Rolland et al. 1998; Wu and Kim 2020; Yoo

et al. 2020; Yoshida et al. 1995]. Such methods employ additional optical components in the device design, and so the overall form factor and the manufacturing cost of the device increases. Foveation effect at optical level can be achieved by either using a single display [Yoo et al. 2020], or two separate displays [Kim et al. 2019a; Lee et al. 2020; Tan et al. 2018a; Yoshida et al. 1995]. Yoo *et al.* [2020] propose a single-display-based near-eye foveated system based on temporal polarization multiplexing and provides two operating modes, whereas Kim *et al.* [2019a] and Lee *et al.* [2020] use two displays – one for the foveal region with a narrow FOV and the other for the peripheral region with a wide FOV.

The concept of double displays dates back to the late 1990s. Doubledisplay systems optically combine the light from two display modules, each with a different resolution and field of view. Yoshida et al. [1995] propose a high-resolution-inset HMD, which optically superimposes a high-resolution image over a low-resolution wide FOV image. The part of the scene around the gaze point is generated at a high resolution. The high-resolution-inset is optically duplicated into a grid of non-overlapping copies to fill the entire display [Rolland et al. 1998; Yoshida et al. 1995]. An LCD array selects one element of this grid based on the gaze direction and transports it to the eyes through optical fibers [Yoshida et al. 1995]. The background image covers the entire FOV and is generated at a lower resolution. The foveal inset is then combined with the background image using opto-electrical components alone. However, these devices were too heavy and expensive when they were proposed. With recent developments in the optical elements, Kim et al. [2019a] and Lee et al. [2020] present prototypes for foveated near-eye devices using two displays with a small form factor. Kim et al. [2019a] use a micro OLED display for the foveal display system and a projector-based Maxwellian-view display for the peripheral display system. The two displays are mechanically steered towards the gaze direction based on the built-in eye tracker's information. Lee et al. [2020] use a holographic near-eye display for the foveal display and elements based on polarization optics for the peripheral display system. The foveal display is steered using a micro-electro-mechanical system (MEMS) mirror and a switchable Pancharatnam-Berry phase (S-PBP) grating module. The peripheral display is ensured to support enough eye box, thereby avoiding the need to steer the display. The light coming from these two displays are then combined optically before reaching the eye.

The concept of foveation is also increasingly used in holographic displays [Maimone et al. 2017; Yaraş et al. 2010], a highly promising direction to true 3D displays with powerful features such as variable focal control, optical aberration correction, and non-conflicting depth cues, to reduce the rendering cost of computer-generated holograms (CGH) [Chang et al. 2020; Ju and Park 2019; Wei and Sakamoto 2019]. For example, Chang *et al.* [2020] use the concept of foveated rendering to speed-up the computational cost of CGH for 3D volumes based on a layered approach. The volume is divided into several parallel layers, and the image at each layer is segmented into a foveal region and a peripheral region. The foveal region is considered at a high resolution for the hologram calculation, whereas the peripheral region is down-sampled to a lower resolution. As a result, the reconstructed images from the hologram appear with a higher quality in the foveal region than the peripheral region.



Fig. 11. Dynamic and Static Foveated Rendering. Dynamic foveated rendering depends on the information from the eye-tracking hardware. Static foveated rendering does not rely on the eye-tracking hardware.

3.4 Gaze Point Information

This section discusses the use of eye-tracking knowledge for foveated rendering systems (Figure 11) and classifies the methods based on the configuration in which they were originally proposed.

3.4.1 Static Foveated Rendering. Static foveated rendering techniques do not rely on the eye-tracking devices and assume that a user focuses on specific positions in the image. The likelihood of a user glancing at a particular area estimates the desired acuity level in that region. Many early works exploited static foveation, in which they either assume that the user looks at the center of the screen [Funkhouser and Séquin 1993] or use a content-based model of visual attention [Horvitz and Lengyel 1997; Yee et al. 2001]. The direction of the head was considered a good approximation for the gaze direction before the eye-tracking devices were developed. Most of the early work on foveated rendering is based on this approximation and assumes that the gaze point is at the center of the screen.

One example of static foveated rendering is the fixed foveated rendering (FFR) [Oculus 2018] developed by Oculus, which assumes that most users look towards the center of the display. FFR uses a tile-based approach. The image is sub-divided into tiles of varying acuity levels, based on the distance from the center. The tile corresponding to the higher resolution region lies at the center while the tiles towards the edges correspond to lower resolution. Each tile is rendered at a uniform spatial resolution. Static foveated rendering requires the users to maintain their focus at predefined fixed zones, mostly the center of the screen. A significant advantage of the static foveated rendering techniques is that they do not depend on the quality of the eye-tracking devices and are compatible with all existing devices. Moreover, static foveated rendering approaches can be further improved with the recent improvements in the visual attention and saliency models [Itti et al. 1998; Kim et al. 2010; Lee et al. 2005]. The area around the salient regions is always rendered at a higher resolution. However, static techniques require a larger high-resolution region than the dynamic approaches. So the average rendering cost for a static foveated rendering technique could be higher than the dynamic method.

3.4.2 Dynamic Foveated Rendering. Traditional foveated rendering is predicated on eye-tracking technology [Reddy 2001]. Eye-tracking devices can pinpoint the user's gaze position in real-time. Research on head-mounted eye trackers dates back to the 1960s [Clay et al. 2019]. However, the recent advancements in computational power, massively parallel image processing, low cost, and small-sized hardware have made it possible to use real-time eye-tracking with VR and AR. By actively tracking the user's gaze using eye-tracking tools, a small region around the gaze point is rendered at a higher resolution and the peripheral region at a lower resolution. Most of the foveated rendering approaches [Guenter et al. 2012; Patney et al. 2016b; Stengel et al. 2016; Swafford et al. 2016; Tursun et al. 2019] are dynamic approaches and depend on the performance of the eye-tracking devices to obtain the gaze point. In addition to the hardware-based eye-trackers, there has been a growing interest in developing methods to predict future gaze positions using deeplearning methods [Hu et al. 2020; Xu et al. 2018]. As the accuracy and efficiency of the eve-tracking solutions improves further, dynamic foveated rendering is likely to greatly enhance the rendering performance and quality.

3.5 Anti-aliasing

Aliasing is prominent in foveated rendering as the peripheral region is rendered at a lower resolution. The aliasing artifacts stemming from foveated rendering can be either spatial or temporal.

Spatial aliasing artifacts occur when the level of detail of the virtual world is higher than the rendered resolution [Hoffman et al. 2018b]. They arise at the object level and appear irrespective of any motion. Aliasing during scene motion generates temporal artifacts. These artifacts are aligned to the output display pixel grid rather than the virtual world coordinates [Hoffman et al. 2018b]. As the user view changes, these artifacts result in flickering and scintillation effects, disrupting the user experience in the virtual world.

Studies show that the human visual system is sensitive to temporal aliasing artifacts even at higher eccentricities [Hoffman et al. 2018a,b]. McKee and Nakayama [1984] show that motion acuity falls dramatically from 0° to 10° and then drops more subtly from 10° to 40°. The peripheral motion sensitivity of the human visual system poses a critical challenge to all foveated rendering techniques. Patney *et al.* [2016b] show that minimizing temporal aliasing in foveated rendering is necessary for an effective user experience.

3.5.1 Avoid Artifacts. One approach to ensure temporal and spatial stability is to design foveated rendering techniques that prevent artifacts from occurring. These techniques generally identify the feature that might cause a detectable artifact and remove that feature or nullify its effect in the rendering pipeline [Bastani et al. 2017]. The foveal and peripheral regions are constantly updated according to the gaze direction. In general, the position of the rendered pixels in the high-resolution foveal region and the low-resolution peripheral region aligns with the display-coordinate system. The low-resolution rendered pixels are later upsampled to match the native display resolution. As the user's head rotates, the value of each rendered pixel changes irrespective of the scene content, causing the pixel color to shift and flicker. This generates the time-varying aliasing artifacts in the upsampled display pixels, thus disrupting

the user experience. Turner et al. [2018] observe that proper angular alignment of the rendered frustums minimizes such frame-to-frame flickering effects with head rotation. They present a phase-aligned foveated rendering system that aligns the low-resolution region to the world-coordinate system rather than the display-coordinate system as shown in Figure 12. The low-resolution region, which is sampled in the world space, is then upsampled and re-projected to align with the display-coordinate system. The artifacts in the peripheral region now move along with the content of the virtual scene, and so the temporal variation of the artifacts is minimized as the head rotates. Phase-aligned foveated rendering significantly reduces the perceivability of motion artifacts in the peripheral region, making aggressive foveation possible. However, spatial aliasing and a distinct boundary between the foveal and peripheral regions are noticeable in the rendered frame. The discern-ability of the fovealperipheral boundary can be reduced by ensuring a smooth transition in the resolution levels from foveal to the peripheral region [Bastani et al. 2017]. Franke et al. [2021] and Mueller et al. [2021] use temporal coherence to reduce the number of shading operations. Franke et al. [2021] reproject the peripheral region of the previous frame to the current frame, and only the foveal pixels and the less coherent peripheral pixels are rendered. They propose to use the exact world space pixel positions of the fragments rather than the depth values to avoid reprojection artifacts.



Fig. 12. Phase-aligned Foveation. The foveal region (green) is rendered in display coordinates, while the peripheral regions (red, orange, pink and blue) are rotationally fixed to the world coordinate system. Image adapted from Turner *et al.* [2018].

3.5.2 Mitigate Artifacts. Most of the foveated rendering approaches reduce the perceptibility of the artifacts in the post-rendering phase. The choice of interpolation techniques used to up-sample the low-resolution peripheral region also affects the detection of the artifacts in the peripheral region. For example, the simple nearest-neighbor interpolation magnifies the temporally unstable artifacts [Hoffman et al. 2018a], though it can offer good contrast retention when compared to bi-linear interpolation.

Guenter *et al.* [2012] use a combination of multi-sample antialiasing (MSAA), temporal reverse reprojection, and temporal jitter of the spatial sampling grid to reduce the spatial and temporal artifacts. MSAA reduces spatial aliasing along silhouette edges by increasing the effective sampling resolution in those regions. Temporal reverse projection with frame jitter reduces the aliasing throughout the image. Temporal anti-aliasing strategies are better at reducing aliasing arising from sampling highly specular materials at a lower sampling rate. Reprojection-based temporal anti-aliasing is a common technique that uses the previous frame's information to mitigate the temporal artifacts. The current frame is reprojected onto the previous frame, and the information from the two frames is used to compute the final color. However, as details from the previous frame may continue to exist beyond when they can be correctly reprojected, temporal anti-aliasing may produce high-frequency artifacts or ghosting [Patney et al. 2016b]. Karis [2014] reduces such ghosting artifacts by conditioning samples from previous frames that are consistent with samples in the current frame. Meng *et al.* [2018] apply temporal anti-aliasing with Halton sampling in the screen space as a post-processing method to mitigate the artifacts appearing in the peripheral region.

Patney et al. [2016b] use pre-filters in addition to temporal antialiasing methods to mitigate temporal artifacts. They introduce variance sampling as a post-process image-enhancement technique. Around each pixel, a variable-size axis-aligned bounding box is constructed based on local color distribution. The back-projected and resampled information from the previous frame that lies within the defined bounding box is integrated with the current frame information. The authors show that by explicitly using the local color information, the ghosting artifacts are reduced. This method is further extended to account for saccadic eye movements. Due to the saccadic movement of the eye, the previous frame's peripheral region can become the foveal region for the current frame. Such a situation requires information from several frames to converge to the level of detail required for the foveal region. Patney et al. [2016b] accelerate the rate of convergence based on the shading rates in the two consecutive frames. This reduces the blurring artifacts caused by eye saccades.

Weier et al. [2018] use depth-of-field information to design a post-process anti-aliasing technique. The depth-of-field effect that occurs when focusing on objects can be used as a low-pass filter to minimize the high-frequency artifacts in the peripheral region. Temporal anti-aliasing is first applied to the foveated rendered image to obtain temporally smooth samples. These samples are used to reconstruct the full image using push-pull interpolation. The aliasing artifacts are then mitigated using a low-pass depth-of-field filter. The approximate depth of the focal point is estimated using a support vector machine-based gaze depth estimator that takes various depth measurements as input. Given the gaze-depth and estimation inaccuracy, two circles of confusion are computed, which determine the depth-of-field. A multi-layer filter is designed based on the depth-of-field model. The filtering of the image occurs in layers, which are eventually blended with different weights to give the final output.

In addition to the above methods, denoising approaches based on deep neural networks have been developed to reduce the detection of artifacts. DeepFovea [Kaplanyan et al. 2019] takes inspiration from the internal model of the human visual system that infers content from the sparse peripheral information. A manifold representing the distribution of samples from a large collection of natural videos is learned using generative adversarial networks [Goodfellow et al. 2014]. For a given sparse foveated input video stream, DeepFovea reconstructs the peripheral region by finding the closest natural video that corresponds to the sparse input on the learned manifold. As the reconstructed output is close to the realistic videos, the temporal aliasing artifacts are minimized.

4 EVALUATION

One can measure the performance of a foveated system across two dimensions - the quality of the rendered image and the average rendering time. However, quality evaluation is a challenging task as a ground-truth reference foveated image does not exist. The peripheral region is sensitive to flickering artifacts, spatial and temporal motion, contrast, and salient features [Patney et al. 2016b] and it is not easy to quantify these perceptual metrics. Further research is required to understand the important characteristics of human perception. Also, the optical components of the HMD limit the perceived resolution. The spatial resolution of optical systems is commonly expressed in terms of the modulation transfer function (MTF), which gives the normalized frequency response of the system. Beams et al. [2020] show that the maximum resolution provided by the optical system depends on the eccentricity in the FOV. This dependence on the optical components also has to be considered while evaluating the foveated rendering system. This section gives an overview of the common qualitative and quantitative measures used to evaluate a foveated system.

4.1 User study-based evaluation

Empirical user studies are the most common and reliable way to evaluate foveated rendering systems. These user studies can be broadly classified into three categories [Hsu et al. 2017]:

- single-stimulus absolute category rating or pairwise comparisons;
- double-stimulus quality comparison;
- slider methods or adjustment methods

Single-Stimulus Absolute Category Rating: The participant is presented with a single stimulus, a foveally-rendered image, on the screen for a certain time. The participant is then asked to make a judgment on the quality of the image, ranging from acceptable to unacceptable.

Double-Stimulus Quality Comparison: This type of experiment requires the availability of a full-resolution image. Two renderings of the scene, a full-resolution (unfoveated) image, and a foveated image are shown to the participant in a random order. The participant is then asked to compare the two images and determine the perceptually superior image.

Slider Methods: These methods are more helpful in finding the optimal parameters for the foveated-rendering system. There are two types of slider experiments - descending and ascending. In descending methods, the participant is presented with a full-resolution high-quality image and a slider. The slider is used to control the parameters of foveation. By moving the slider position gradually, the participant is asked to determine the point at which the quality of the image degrades. The ascending method is similar to the descending method, but the participant is initially provided with a reference full-resolution image for a short period. The quality of the image gradually increases from a fully foveated image as the slider moves, and the user is asked to determine the point at which the quality of the image is perceptually similar to the reference image.

12 • Susmija Jabbireddy, Xuetong Sun, Xiaoxu Meng, and Amitabh Varshney

Paper	Visual Factors			Distribution				Space			Gaze		Aliasing	
	Acuity	Contrast	Saliency	Hyperbolic	Linear	Logarithmic	Quadratic	Object	Screen	Optics	Static	Dynamic	Avoid	Mitigate
Funkhouser and Séquin [1993]	1	-	1	1	-	-	-	1	-	-	1	-	-	-
Leubke and Hallen [2001]	1	1	1	-	-	-	-	1	-	-	-	1	-	-
Murphy and Duchowski [2001]	1	-	-	-	-	-	-	1	-	-	-	1	-	-
Reddy [2001]	1	-	-	-	-	-	1	1	-	-	1	-	-	-
Guenter <i>et al.</i> [2012]	1	-	-	1	-	-	-	-	1	-	-	1	-	1
Fujita and Harada [2014]	1	-	-	1	-	-	-	-	1	-	1	-	-	1
Stengel <i>et al.</i> [2016]	1	1	1	1	-	-	-	-	1	-	-	1	1	1
Swafford et al. [2016]	1	-	-	1	-	-	-	1	1	-	-	1	-	-
Patney et al. [2016b]	1	1	-	-	1	-	-	-	1	-	-	1	-	1
Weier <i>et al.</i> [2017]	1	-	-	-	1	-	-	-	1	-	-	1	1	1
Turner <i>et al.</i> [2018]	1	-	-	-	-	-	-	-	1	-	-	1	1	-
Zheng <i>et al.</i> [2018]	1	-	-	-	-	-	1	1	-	-	-	1	-	-
Meng et al. [2018]	1	-	-	-	-	1	-	-	1	-	-	~	-	1
Weier et al. [2018]	1	-	-	-	1	-	-	-	1	-	-	~	-	1
Koskela <i>et al.</i> [2019b]	1	-	-	-	-	1	-	-	1	-	-	1	1	1
Tursun <i>et al.</i> [2019]	1	1	-	1	-	-	-	-	1	-	-	1	1	-
Kim <i>et al.</i> [2019a]	1	-	-	-	-	-	-	-	-	1	-	1	-	1
Lee et al. [2020]	1	-	-	-	-	-	-	-	-	1	-	1	-	1
Yoo et al. [2020]	1	-	-	-	-	-	-	-	-	1	-	1	-	1

Table 2. Classification of existing prominent foveated rendering works based on the presented taxonomy

The user studies also provide personalized calibration of optimal foveation parameters since the level of acceptance of foveated imagery may vary from person to person. However, user studies are generally very time-consuming and expensive. Therefore, perception and foveation-based evaluation metrics are required to help us design and compare various methods. The user-study experiments can help us evaluate methods that have a high probability of success based on the initial evaluation metrics.

4.2 Computational metric-based evaluation

Traditional image quality metrics are designed for uniform quality images. Structural Similarity Index (SSIM) and HDR Visual Difference Predictor (HDR-VDP2) are well-known perceptually-informed metrics [Swafford et al. 2016]. SSIM [Wang et al. 2004] uses structural distortion in an image as an estimate of perceived visual distortion. It is based on the assumption that HVS is highly adapted for the extraction of structural information from the visual field. SSIM uses mean, variance, and covariance of the original high-resolution image and the foveated image to measure their structural differences. HDR-VDP2 [Mantiuk et al. 2011] considers the contrast sensitivity measurements. These metrics were modeled with the assumption of a uniform sensitivity across the visual field.

Several recent papers [Guo et al. 2018; Rimac-Drije et al. 2010; Sanghoon Lee et al. 2002; Swafford et al. 2016; Tsai and Liu 2014; Wang et al. 2001; You et al. 2014] extend the uniform image quality metrics to consider foveation. Rimac et al. [2011] combine SSIM with a foveation-based sensitivity function to account for the nonuniform perceived structural differences across the visual field. Swafford et al. [2016] extend the HDR-VDP2 metric by including the contrast-sensitivity degradation factor. Tsai et al. [2014] develop a window-based metric and assign weights based on the distance from the predicted salient regions. Guo et al. [2018] vary the spatial resolution based on eccentricity to calculate the quality of the image. Recently, Mantiuk et al. [2021] developed a novel metric named FovVideoVDP that incorporates several factors such as spatial resolution variance with eccentricity, spatial and temporal contrast sensitivity, and scene content. Most of these metrics have been developed for foveated video quality assessment and rely on a fullresolution reference image or video. However, complete perceptually informed and computationally simple metrics to evaluate the perceived quality of wide FOV images and videos and level of immersion are relatively unexplored.

In terms of rendering time, we provide the results reported for each of the discussed foveated rendering methods. This, however, should not be interpreted to be a relative comparison among the different methods unless explicitly noted otherwise; the overall rendering time varies with the complexity of the scene, shading models, and the hardware used. Guenter et al. [2012] report an overall speedup of $5 \times - 6 \times$ in rendering time and $10 \times - 15 \times$ in the number of pixels rendered. Multi-rate shading [He et al. 2014] provides an average speedup of about $3 \times - 5 \times$ and the multi-resolution approach [Reed 2015] provides a speedup of $1.3 \times - 2 \times$. Zheng et al. [2018] reduce the rendering time by a factor of $1.2 \times$ compared to Guenter et al. [2012]. Patney et al. [2016b] report 2× more efficiency in reducing shading cost compared to Guenter et al. [2012]. Stengel et al. [2016] report a speedup of 1.34× in overall rendering time while reducing the shading time by 1.7×. Swafford et al. [2016] reduce the rendering time in half by carefully reducing the resolution in the periphery. The kernel-foveated rendering method by Meng et al. [2018] achieves $2.8 \times - 3.2 \times$ speedup by using log-polar transformations. By estimating the maximum resolution required for each pixel, Tursun et al. [2019] report 3.08× speedup in overall rendering time. The ray-tracing-based foveated rendering method by Weier et al. [2016] provides an average speedup of 2.55×, reducing the number of sampled pixels by 79% on benchmark scenes. By using depth-of-field filtering as a post-processing step, Weier et al. [2018] reduce the number of sampled pixels by 69% without any visible artifacts.

5 CONCLUSION

In conclusion, we have described the limitations of the human visual system and how foveated rendering leverages such limitations to

reduce the computational resources required for real-time rendering. We compare and discuss various foveated rendering works by providing a taxonomy based on the key factors in the rendering pipeline. Table 2 provides a summary of our taxonomy. Our first classification differentiates methods based on the visual factors considered while developing the rendering system: content-independent (acuity) and content-dependent (image contrast, saliency). As can be seen from table 2, all the papers consider acuity degradation with eccentricity in their rendering systems. In addition, a few works [Patney et al. 2016a; Stengel et al. 2016; Tursun et al. 2019] leverage the contentdependent information to further improve the foveation effect. The second category differentiates the methods based on the underlying analytical models for peripheral degradation. While all the methods reduce the sampling rate with an increase in distance from the focus point, the reduction rate is different across methods. The best-suited model is selected based on the target task and scene complexity. We next differentiate methods according to the space where the foveation effect is employed: object-space, screen-space, and opticsspace. With the evolution of the graphics pipeline and dedicated hardware, there is an increasing trend to incorporate foveation in the screen or the image space. More recently, several works have modified the display at the lens level to incorporate foveation. These optics-based foveated methods mitigate the hardware limitation in maintaining high pixel densities to match the angular resolution of the HVS. We also separate the methods based on the amount of gaze information used in their initial proposed solution. With improvement in eye-tracking technology, more and more works tend to develop and study foveated systems using gaze point information. As the eye-tracking technology improves in both accuracy and latency, dynamic foveated rendering seems to be the way to achieve real-time rendering with high perceptual quality. We also discuss anti-aliasing approaches that complement foveated rendering. These include deep-learning-based approaches that produce realistic, anti-aliased, foveated output images from sparse inputs. It is to be noted that the provided categorization is not disjoint, and a solution can combine several options in each category, for example, including several visual factors or combining screen-space and optics-space foveation methods.

We also provided a brief overview of the human visual system. However, the way the human brain processes information from the images formed on the retina is not yet fully understood. As the neuroscience community advances our understanding of visual perception and cognition, their studies can guide our efforts towards minimal and sufficient rendering without compromising human perception and cognition.

The ultimate goal is to build head-mounted displays that can simultaneously provide wide FOV and high resolution that match human perceptual capabilities. Visual acuity is also limited by diffraction, aberrations of the eye lens, in addition to the density of the photoreceptor cells [Smith 1997]. In addition, even factors such as refractive error, illumination, contrast, and location of the retina being stimulated affect the perceived image quality. Studies also show that the photoreceptor density decreases with age [Panda-Jonas et al. 1995]. So, these factors are to be considered while studying the effectiveness of a foveated rendering system. Interestingly, the foveation characteristics differ even between the two eyes. Meng *et* *al.* [2020a] leverages the concept of eye dominance, which states that the human visual system prefers visual stimuli from one eye more than the other. This suggests that the non-dominant eye can permit greater foveation than the dominant eye without any perceptual difference. Similarly, given that only rods are active for night vision can be used to render scenes under dim lighting differently than those with bright lighting and at a much lower computational cost. Another possible direction of research is to leverage the color discrimination ability of the human eye across the visual field, which decreases with an increase in eccentricity. Further, much of the current effort in building foveated rendering system is for a personalized single user system; foveated rendering systems for shared multi-user displays is another exciting area for future research.

A few recent works [Kim et al. 2019a; Lee et al. 2019; Lee et al. 2018; Tan et al. 2018b] use the concept of foveation to develop neareye foveated displays. Kim *et al.* [2019a] use a deep-learning-based gaze-tracking system [Kim et al. 2019b] for pupil center estimation to move the small foveal display appropriately. Such systems can also benefit from estimating the future gaze direction based on the current frame and gaze information. Improving the accuracy of the saliency models based on long-term memory and the information from the current frame can guide the renderer to render appropriate sections at a high resolution. As deep-learning models are increasingly able to provide better results in several tasks, they can also be employed to improve the gaze and saliency estimation tasks.

Although there has been significant progress in foveated rendering, it is still challenging to realize the full potential of foveated rendering, especially in complex scenes with photo-realistic lighting. Further, there is an urgent need to develop a uniform evaluation metric that uses perceptual criteria to quantify the quality of the foveated rendering and allows us to compare various foveation methods. We believe that the combination of developments in both software and hardware will make foveated rendering more dominant in every commodity headset.

ACKNOWLEDGMENTS

This work has been supported in part by the NSF Grants 15-64212, 18-23321, and the State of Maryland's MPower initiative. The authors would like to thank Alexandar Rowden, Anshul Shah, Sweta Agrawal and Yogesh Balaji for their helpful suggestions.

REFERENCES

- Jared Abrams, Aaron Nizam, and Marisa Carrasco. 2012. Isoeccentric locations are not equivalent: The extent of the vertical meridian asymmetry. Vision Research 52, 1 (2012), 70–78. https://doi.org/10.1016/j.visres.2011.10.016
- R. S. Anderson, M. B. Zlatkova, and R. O. Beirne. 2002. The contrast sensitivity function for detection and resolution of blue-on-yellow gratings in foveal and peripheral vision. *Ophthalmic and Physiological Optics* 22, 5 (2002), 420–426. https://doi.org/10.1046/j.1475-1313.2002.00068.x arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1046/j.1475-1313.2002.00068.x
- S.M. Anstis. 1974. A chart demonstrating variations in acuity with retinal position. Vision Research 14, 7 (1974), 589–592. https://doi.org/10.1016/0042-6989(74)90049-2
- Marco Antonelli, Francisco D. Igual, Francisco Ramos, and V. Javier Traver. 2015. Speeding up the log-polar transform with inexpensive parallel hardware: graphics units and multi-core architectures. *Journal of Real-Time Image Processing* 10, 3 (01 Sep 2015), 533–550. https://doi.org/10.1007/s11554-012-0281-6
- H. Araujo and J. M. Dias. 1996. An introduction to the log-polar mapping. In *Proceedings* II Workshop on Cybernetic Vision. 139–144.
- Antoine Barbot, Shutian Xue, and Marisa Carrasco. 2021. Asymmetries in visual acuity around the visual field. Journal of Vision 21, 1 (01 2021), 2–2. https://doi.org/10.1167/jov.21.1.2

arXiv:https://arvojournals.org/arvo/content_public/journal/jov/938512/i1534-7362-21-1-2_1609762263.01982.pdf

- Peter G. J. Barten. 2003. Formula for the contrast sensitivity of the human eye. In Image Quality and System Performance, Yoichi Miyake and D. Rene Rasmussen (Eds.), Vol. 5294. International Society for Optics and Photonics, SPIE, 231 – 238. https://doi.org/10.1117/12.537476
- Behnam Bastani, Eric Turner, Carlin Vieri, Haomiao Jiang, Brian Funt, and Nikhil Balram. 2017. Foveated Pipeline for AR/VR Head-Mounted Displays. *Information Display* 33 (11 2017), 14–19 and 35. https://doi.org/10.1002/j.2637-496x.2017.tb01040.x
- R. Beams, B. Collins, A. S. Kim, and A. Badano. 2020. Angular Dependence of the Spatial Resolution in Virtual Reality Displays. In 2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR). 836–841. https://doi.org/10.1109/VR46266.2020.00108
- David R. Bull. 2014. Chapter 2 The Human Visual System. In Communicating Pictures, David R. Bull (Ed.). Academic Press, Oxford, 17 – 61. https://doi.org/10.1016/B978-0-12-405906-1.00002-7
- Chenliang Chang, Wei Cui, and Liang Gao. 2020. Foveated holographic near-eye 3D display. Opt. Express 28, 2 (Jan 2020), 1345–1356. https://doi.org/10.1364/OE.384421
- Viviane Clay, Peter König, and Sabine Koenig. 2019. Eye Tracking in Virtual Reality. Journal of Eye Movement Research 12 (04 2019). https://doi.org/10.16910/jemr.12.1.3 Robert Danforth, Andrew Duchowski, Robert Geist, and Elizabeth McAliley. 2000. A
- platform for gaze-contingent virtual environments. In Smart Graphics (Papers from the 2000 AAAI Spring Symposium, Technical Report SS-00-04). 66-70.
- Michael Deering, Stephanie Winner, Bic Schediwy, Chris Duffy, and Neil Hunt. 1988. The Triangle Processor and Normal Vector Shader: A VLSI System for High Performance Graphics. SIGGRAPH Comput. Graph. 22, 4 (June 1988), 21–30. https://doi.org/10. 1145/378456.378468
- Linus Franke, Laura Fink, Jana Martschinke, Kai Selgrad, and Marc Stamminger. 2021. Time-Warped Foveated Rendering for Virtual Reality Headsets. Computer Graphics Forum 40, 1 (2021), 110–123. https://doi.org/10.1111/cgf.14176 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14176
- Sebastian Friston, Tobias Ritschel, and Anthony Steed. 2019. Perceptual Rasterization for Head-mounted Display Image Synthesis. ACM Trans. Graph. 38, 4, Article 97 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3323033
- Masahiro Fujita and Takahiro Harada. 2014. Foveated real-time ray tracing for virtual reality headset. Light Transport Entertainment Research (2014).
- Thomas A. Funkhouser and Carlo H. Séquin. 1993. Adaptive Display Algorithm for Interactive Frame Rates During Visualization of Complex Virtual Environments. In Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques (Anaheim, CA) (SIGGRAPH '93). ACM, New York, NY, USA, 247–254. https://doi.org/10.1145/166117.166149
- Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. In Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2 (Montreal, Canada) (NIPS'14). MIT Press, Cambridge, MA, USA, 2672–2680. http://dl.acm.org/citation.cfm?id=2969033.2969125
- 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. https://doi.org/10.1145/2366145.2366183
- Peiyao Guo, Qiu Shen, Zhan Ma, David J Brady, and Yao Wang. 2018. Perceptual quality assessment of immersive images considering peripheral vision impact. arXiv preprint arXiv:1802.09065 (2018).
- Thorsten Hansen, Lars Pracejus, and Karl R. Gegenfurtner. 2009a. Color perception in the intermediate periphery of the visual field. *Journal of Vision* 9, 4 (04 2009), 26–26. https://doi.org/10.1167/9.4.26 arXiv:https://arvojournals.org/arvo/content_public/journal/jov/933534/jov-9-4-26.pdf
- Thorsten Hansen, Lars Pracejus, and Karl R. Gegenfurtner. 2009b. Color perception in the intermediate periphery of the visual field. *Journal of Vision* 9, 4 (04 2009), 26–26. https://doi.org/10.1167/9.4.26 arXiv:https://arvojournals.org/arvo/content_public/journal/jov/933534/jov-9-4-26.pdf
- Yong He, Yan Gu, and Kayvon Fatahalian. 2014. Extending the Graphics Pipeline with Adaptive, Multi-rate Shading. ACM Trans. Graph. 33, 4, Article 142 (July 2014), 12 pages. https://doi.org/10.1145/2601097.2601105
- David Hoffman, Zoe Meraz, and Eric Turner. 2018a. Limits of peripheral acuity and implications for VR system design. Journal of the Society for Information Display 26, 8 (2018), 483-495. https://doi.org/10.1002/jsid.730 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/jsid.730
- David M. Hoffman, Zoe Meraz, and Eric Turner. 2018b. 65-2: Sensitivity to Peripheral Artifacts in VR Display Systems. SID Symposium Digest of Technical Papers 49, 1 (2018), 858–861. https://doi.org/10.1002/sdtp.12261
- Eric Horvitz and Jed Lengyel. 1997. Perception, Attention, and Resources: A Decisiontheoretic Approach to Graphics Rendering. In Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence (Providence, Rhode Island) (UAI'97). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 238–249. http: //dl.acm.org/citation.cfm?id=2074226.2074255

- Chih-Fan Hsu, Anthony Chen, Cheng-Hsin Hsu, Chun-Ying Huang, Chin-Laung Lei, and Kuan-Ta Chen. 2017. Is Foveated Rendering Perceivable in Virtual Reality?: Exploring the Efficiency and Consistency of Quality Assessment Methods. In Proceedings of the 25th ACM International Conference on Multimedia (Mountain View, California, USA) (MM '17). ACM, New York, NY, USA, 55–63. https: //doi.org/10.1145/3123266.3123434
- Z. Hu, S. Li, C. Zhang, K. Yi, G. Wang, and D. Manocha. 2020. DGaze: CNN-Based Gaze Prediction in Dynamic Scenes. *IEEE Transactions on Visualization and Computer Graphics* 26, 5 (2020), 1902–1911.
- L. Itti, C. Koch, and E. Niebur. 1998. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20, 11 (1998), 1254–1259.
- Yeon-Gyeong Ju and Jae-Hyeung Park. 2019. Foveated computer-generated hologram and its progressive update using triangular mesh scene model for near-eye displays. Opt. Express 27, 17 (Aug 2019), 23725–23738. https://doi.org/10.1364/OE.27.023725
- Anton Kaplanyan, Anton Sochenov, Thomas Leimkühler, Mikhail Okunev, Todd Goodall, and Gizem Rufo. 2019. Deepfovea: Neural Reconstruction for Foveated Rendering and Video Compression Using Learned Natural Video Statistics. In ACM SIGGRAPH 2019 Talks (Los Angeles, California) (SIGGRAPH '19). ACM, New York, NY, USA, Article 58, 2 pages. https://doi.org/10.1145/3306307.3328186
- Brian Karis. 2014. High-quality temporal supersampling. Advances in Real-Time Rendering in Games, SIGGRAPH Courses 1 (2014), 1–55.
- Jonghyun Kim, Youngmo Jeong, Michael Stengel, Kaan Akşit, Rachel Albert, Ben Boudaoud, Trey Greer, Joohwan Kim, Ward Lopes, Zander Majercik, Peter Shirley, Josef Spjut, Morgan McGuire, and David Luebke. 2019a. Foveated AR: Dynamicallyfoveated Augmented Reality Display. ACM Trans. Graph. 38, 4, Article 99 (July 2019), 15 pages. https://doi.org/10.1145/3306346.3322987
- Joohwan Kim, Michael Stengel, Alexander Majercik, Shalini De Mello, David Dunn, Samuli Laine, Morgan McGuire, and David Luebke. 2019b. NVGaze: An Anatomically-Informed Dataset for Low-Latency, Near-Eye Gaze Estimation. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (Glasgow, Scotland Uk) (CHI '19). ACM, New York, NY, USA, Article 550, 12 pages. https://doi.org/10.1145/3290605.3300780
- Ungsoo Samuel Kim, Omar A. Mahroo, John D. Mollon, and Patrick Yu-Wai-Man. 2021. Retinal Ganglion Cells—Diversity of Cell Types and Clinical Relevance. *Frontiers in Neurology* 12 (2021), 635. https://doi.org/10.3389/fneur.2021.661938
- Youngmin Kim, Amitabh Varshney, David W. Jacobs, and François Guimbretière. 2010. Mesh Saliency and Human Eye Fixations. ACM Trans. Appl. Percept. 7, 2, Article 12 (Feb. 2010), 13 pages. https://doi.org/10.1145/1670671.1670676
- Matias Koskela, Kalle Immonen, Markku Mäkitalo, Alessandro Foi, Timo Viitanen, Pekka Jääskeläinen, Heikki Kultala, and Jarmo Takala. 2019a. Blockwise Multi-Order Feature Regression for Real-Time Path Tracing Reconstruction. ACM Transactions on Graphics (TOG) 38, 5 (June 2019). https://doi.org/10.1145/3269978
- Matias Koskela, Atro Lotvonen, Markku Mäkitalo, Petrus Kivi, Timo Viitanen, and Pekka Jääskeläinen. 2019b. Foveated Real-Time Path Tracing in Visual-Polar Space. In Eurographics Symposium on Rendering - DL-only and Industry Track, Tamy Boubekeur and Pradeep Sen (Eds.). The Eurographics Association. https://doi.org/10.2312/sr. 20191219
- Chang Ha Lee, Amitabh Varshney, and David W. Jacobs. 2005. Mesh Saliency. In ACM SIGGRAPH 2005 Papers (Los Angeles, California) (SIGGRAPH '05). Association for Computing Machinery, New York, NY, USA, 659–666. https://doi.org/10.1145/ 1186822.1073244
- Jin Su Lee, Yoo Kwang Kim, Mu Young Lee, and Yong Hyub Won. 2019. Enhanced see-through near-eye display using time-division multiplexing of a Maxwellian-view and holographic display. *Opt. Express* 27, 2 (Jan 2019), 689–701. https://doi.org/10. 1364/OE.27.000689
- S. Lee, J. Cho, B. Lee, Y. Jo, C. Jang, D. Kim, and B. Lee. 2018. Foveated Retinal Optimization for See-Through Near-Eye Multi-Layer Displays. *IEEE Access* 6 (2018), 2170–2180. https://doi.org/10.1109/ACCESS.2017.2782219
- Seungjae Lee, Mengfei Wang, Gang Li, Lu Lu, Yusufu Sulai, Changwon Jang, and Barry Silverstein. 2020. Foveated near-eye display for mixed reality using liquid crystal photonics. *Scientific Reports* 10, 1 (2020), 1–11.
- Leonard A Levin, Siv FE Nilsson, James Ver Hoeve, Samuel Wu, Paul L Kaufman, and Albert Alm. 2011. *Adler's Physiology of the Eye E-Book: Expert Consult-Online and Print.* Elsevier Health Sciences.
- Marc Levoy and Ross Whitaker. 1990. Gaze-directed Volume Rendering. SIGGRAPH Comput. Graph. 24, 2 (Feb. 1990), 217–223. https://doi.org/10.1145/91394.91449
- David Li, Ruofei Du, Adharsh Babu, Camelia D. Brumar, and Amitabh Varshney. 2021. A Log-Rectilinear Transformation for Foveated 360-degree Video Streaming. IEEE Transactions on Visualization and Computer Graphics (2021).
- Sheng Liu and Hong Hua. 2008. Spatialchromatic Foveation for Gaze Contingent Displays. In Proceedings of the 2008 Symposium on Eye Tracking Research & Amp; Applications (Savannah, Georgia) (ETRA '08). Association for Computing Machinery, New York, NY, USA, 139–142. https://doi.org/10.1145/1344471.1344507

- Lester C. Loschky, Sebastien Szaffarczyk, Clement Beugnet, Michael E. Young, and Muriel Boucart. 2019. The contributions of central and peripheral vision to scene-gist recognition with a 180° visual field. *Journal of Vision* 19, 5 (05 2019), 15–15. https://doi.org/10.1167/19.5.15 arXiv:https://jov.arvojournals.org/arvo/content_public/journal/jov/937990/i1534-7362-19-5-15.pdf
- Holger Ludvigsen and Anne Cathrine Elster. 2010. Real-Time Ray Tracing Using Nvidia OptiX. In Eurographics 2010 - Short Papers, H. P. A. Lensch and S. Seipel (Eds.). The Eurographics Association. https://doi.org/10.2312/egsh.20101049
- David Luebke and Benjamin Hallen. 2001. Perceptually Driven Simplification for Interactive Rendering. In Proceedings of the 12th Eurographics Conference on Rendering (London, UK) (EGWR'01). Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 223–234. https://doi.org/10.2312/EGWR/EGWR01/223-234
- Andrew Maimone, Andreas Georgiou, and Joel S. Kollin. 2017. Holographic Near-Eye Displays for Virtual and Augmented Reality. ACM Trans. Graph. 36, 4, Article 85 (July 2017), 16 pages. https://doi.org/10.1145/3072959.3073624
- Rafat Mantiuk, Kil Joong Kim, Allan G. Rempel, and Wolfgang Heidrich. 2011. HDR-VDP-2: A Calibrated Visual Metric for Visibility and Quality Predictions in All Luminance Conditions. ACM Trans. Graph. 30, 4, Article 40 (July 2011), 14 pages. https://doi.org/10.1145/2010324.1964935
- Rafał K. Mantiuk, Gyorgy Denes, Alexandre Chapiro, Anton Kaplanyan, Gizem Rufo, Romain Bachy, Trisha Lian, and Anjul Patney. 2021. FovVideoVDP: A Visible Difference Predictor for Wide Field-of-View Video. ACM Trans. Graph. 40, 4, Article 49 (July 2021), 19 pages. https://doi.org/10.1145/3450626.3459831
- Suzanne P. Mckee and Ken Nakayama. 1984. The detection of motion in the peripheral visual field. Vision Research 24, 1 (1984), 25 - 32. https://doi.org/10.1016/0042-6989(84)90140-8
- X. Meng, R. Du, J. F. JaJa, and A. Varshney. 2020b. 3D-Kernel Foveated Rendering for Light Fields. *IEEE Transactions on Visualization and Computer Graphics* (2020), 1–1. https://doi.org/10.1109/TVCG.2020.2975801
- X. Meng, R. Du, and A. Varshney. 2020a. Eye-dominance-guided Foveated Rendering. IEEE Transactions on Visualization and Computer Graphics (2020), 1–1. https://doi. org/10.1109/TVCG.2020.2973442
- 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. https://doi.org/10.1145/3203199
- Bipul Mohanto, ABM Tariqul Islam, Enrico Gobbetti, and Oliver Staadt. 2022. An integrative view of foveated rendering. *Computers & Graphics* 102 (2022), 474–501. https://doi.org/10.1016/j.cag.2021.10.010
- Joerg H. Mueller, Thomas Neff, Philip Voglreiter, Markus Steinberger, and Dieter Schmalstieg. 2021. Temporally Adaptive Shading Reuse for Real-Time Rendering and Virtual Reality. ACM Trans. Graph. 40, 2, Article 11 (April 2021), 14 pages. https://doi.org/10.1145/3446790
- Hunter Murphy and Andrew Duchowski. 2001. Gaze-Contingent Level Of Detail Rendering. EuroGraphics 2001 (01 2001).
- Oculus. 2018. Oculus Go: Fixed Foveated Rendering. https://developer.oculus.com/ documentation/unreal/latest/concepts/unreal-ffr
- T. Ohshima, H. Yamamoto, and H. Tamura. 1996. Gaze-directed adaptive rendering for interacting with virtual space. In *Proceedings of the IEEE 1996 Virtual Reality Annual International Symposium*. 103–110. https://doi.org/10.1109/VRAIS.1996.490517
- Songhomitra Panda-Jonas, Jost B. Jonas, and Martha Jakobczyk-Zmija. 1995. Retinal Photoreceptor Density Decreases with Age. Ophthalmology 102, 12 (1995), 1853– 1859. https://doi.org/10.1016/S0161-6420(95)30784-1
- Anjul Patney, Joohwan Kim, Marco Salvi, Anton Kaplanyan, Chris Wyman, Nir Benty, Aaron Lefohn, and David Luebke. 2016a. Perceptually-based Foveated Virtual Reality. In ACM SIGGRAPH 2016 Emerging Technologies (Anaheim, California) (SIGGRAPH '16). ACM, New York, NY, USA, Article 17, 2 pages. https://doi.org/10.1145/2929464. 2929472
- Anjul Patney, Marco Salvi, Joohwan Kim, Anton Kaplanyan, Chris Wyman, Nir Benty, David Luebke, and Aaron Lefohn. 2016b. Towards Foveated Rendering for Gazetracked Virtual Reality. ACM Trans. Graph. 35, 6, Article 179 (Nov. 2016), 12 pages. https://doi.org/10.1145/2980179.2980246
- J.S. Pointer and R.F. Hess. 1989. The contrast sensitivity gradient across the human visual field: With emphasis on the low spatial frequency range. Vision Research 29, 9 (1989), 1133 – 1151. https://doi.org/10.1016/0042-6989(89)90061-8
- Dale Purves, George J Augustine, David. Fitzpatrick, Lawrence C Katz, Anthony-Samuel LaMantia, James O McNamara, and S Mark Williams. 2001. Neuroscience. 2nd edition. Sunderland (MA): Sinauer Associates. https://www.ncbi.nlm.nih.gov/ books/NBK10799/
- Martin Reddy. 2001. Perceptually Optimized 3D Graphics. IEEE Comput. Graph. Appl. 21, 5 (Sept. 2001), 68–75. https://doi.org/10.1109/38.946633
- Nathan Reed. 2015. NVIDIA Multi Resolution Rendering. (2015). http://www.reedbeta. com/talks/gameworks-vr/
- Snježana Rimac-Drije, Mario Vranješ, and Drago Žagar. 2010. Foveated Mean Squared Error-a Novel Video Quality Metric. *Multimedia Tools Appl.* 49, 3 (Sept. 2010), 425-445. https://doi.org/10.1007/s11042-009-0442-1

16 . Susmija Jabbireddy, Xuetong Sun, Xiaoxu Meng, and Amitabh Varshney

- Snježana Rimac-Drlje, Goran Martinović, and Branka Zovko-Cihlar. 2011. Foveationbased content Adaptive Structural Similarity index. In 2011 18th International Conference on Systems, Signals and Image Processing. IEEE, 1–4.
- Jannick P. Rolland, Akitoshi Yoshida, Larry D. Davis, and John H. Reif. 1998. Highresolution inset head-mounted display. *Appl. Opt.* 37, 19 (Jul 1998), 4183–4193. https://doi.org/10.1364/AO.37.004183
- Ruth Rosenholtz. 2016. Capabilities and Limitations of Peripheral Vision. Annual Review of Vision Science 2, 1 (2016), 437–457. https://doi.org/10.1146/annurevvision-082114-035733 arXiv:https://doi.org/10.1146/annurev-vision-082114-035733 PMID: 28532349.
- Sanghoon Lee, M. S. Pattichis, and A. C. Bovik. 2002. Foveated video quality assessment. IEEE Transactions on Multimedia 4, 1 (2002), 129–132. https://doi.org/10.1109/6046. 985561
- George Smith. 1997. The eye and visual optical instruments / George Smith, David A. Atchison. Cambridge University Press, Cambridge ;.
- Ran Song, Yonghuai Liu, Ralph R. Martin, and Paul L. Rosin. 2014. Mesh Saliency via Spectral Processing. ACM Trans. Graph. 33, 1, Article 6 (Feb. 2014), 17 pages. https://doi.org/10.1145/2530691
- Josef Spjut and Ben Boudaoud. 2019. Foveated Displays: Toward Classification of the Emerging Field. In ACM SIGGRAPH 2019 Talks (Los Angeles, California) (SIGGRAPH '19). ACM, New York, NY, USA, Article 57, 2 pages. https://doi.org/10.1145/3306307. 3328145
- J. Spjut, B. Boudaoud, J. Kim, T. Greer, R. Albert, M. Stengel, K. Akşit, and D. Luebke. 2020. Toward Standardized Classification of Foveated Displays. *IEEE Transactions* on Visualization and Computer Graphics 26, 5 (2020), 2126–2134. https://doi.org/10. 1109/TVCG.2020.2973053
- Michael Stengel, Steve Grogorick, Martin Eisemann, and Marcus Magnor. 2016. Adaptive Image-Space Sampling for Gaze-Contingent Real-time Rendering. Comput. Graph. Forum 35, 4 (July 2016), 129–139. https://doi.org/10.1111/cgf.12956
- Qi Sun, Fu-Chung Huang, Joohwan Kim, Li-Yi Wei, David Luebke, and Arie Kaufman. 2017. Perceptually-guided Foveation for Light Field Displays. ACM Trans. Graph. 36, 6, Article 192 (Nov. 2017), 13 pages. https://doi.org/10.1145/3130800.3130807
- Xuetong Sun and Amitabh Varshney. 2018. Investigating Perception Time in the Far Peripheral Vision for Virtual and Augmented Reality. In Proceedings of the 15th ACM Symposium on Applied Perception (Vancouver, British Columbia, Canada) (SAP '18). ACM, New York, NY, USA, Article 13, 8 pages. https://doi.org/10.1145/3225153. 3225160
- Nicholas T. Swafford, José A. Iglesias-Guitian, Charalampos Koniaris, Bochang Moon, Darren Cosker, and Kenny Mitchell. 2016. User, Metric, and Computational Evaluation of Foveated Rendering Methods. In Proceedings of the ACM Symposium on Applied Perception (Anaheim, California) (SAP '16). ACM, New York, NY, USA, 7–14. https://doi.org/10.1145/2931002.2931011
- Guanjun Tan, Yun-Han Lee, Tao Zhan, Jilin Yang, Sheng Liu, Dongfeng Zhao, and Shin-Tson Wu. 2018a. Foveated imaging for near-eye displays. *Opt. Express* 26, 19 (Sep 2018), 25076–25085. https://doi.org/10.1364/OE.26.025076
- Guanjun Tan, Yun-Han Lee, Tao Zhan, Jilin Yang, Sheng Liu, Dongfeng Zhao, and Shin-Tson Wu. 2018b. Foveated imaging for near-eye displays. *Opt. Express* 26, 19 (Sep 2018), 25076–25085. https://doi.org/10.1364/OE.26.025076
- R. Taylor and D. Batey. 2012. Handbook of Retinal Screening in Diabetes: Diagnosis and Management. Wiley. https://books.google.com/books?id=836CqvTVKwkC
- Larry N. Thibos, David L. Still, and Arthur Bradley. 1996. Characterization of spatial aliasing and contrast sensitivity in peripheral vision. *Vision Research* 36, 2 (1996), 249 – 258. https://doi.org/10.1016/0042-6989(95)00109-D
- W. Tsai and Y. Liu. 2014. Foveation-based image quality assessment. In 2014 IEEE Visual Communications and Image Processing Conference. 25–28. https://doi.org/10.1109/ VCIP.2014.7051495
- E. Turner, H. Jiang, D. Saint-Macary, and B. Bastani. 2018. Phase-Aligned Foveated Rendering for Virtual Reality Headsets. In 2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR). IEEE Computer Society, Los Alamitos, CA, USA, 1–2. https://doi.org/10.1109/VR.2018.8446142
- Okan Tarhan Tursun, Elena Arabadzhiyska-Koleva, Marek Wernikowski, Radosław Mantiuk, Hans-Peter Seidel, Karol Myszkowski, and Piotr Didyk. 2019. Luminancecontrast-aware Foveated Rendering. ACM Trans. Graph. 38, 4, Article 98 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3322985
- P.P. Urone and Hinrichs R. 2012. College Physics. OpenStax. https://openstax.org/ books/college-physics/pages/26-1-physics-of-the-eye
- K. Vaidyanathan, M. Salvi, R. Toth, T. Foley, T. Akenine-Möller, J. Nilsson, J. Munkberg, J. Hasselgren, M. Sugihara, P. Clarberg, T. Janczak, and A. Lefohn. 2014. Coarse Pixel Shading. In Proceedings of High Performance Graphics (Lyon, France) (HPG '14). Eurographics Association, Goslar Germany, Germany, 9–18. http://dl.acm.org/ citation.cfm?id=2980009.2980011
- VRWorks. [n. d.]. Multi-Res Shading. ([n. d.]). https://developer.nvidia.com/vrworks/ graphics/multiresshading "[Online; accessed 6-December-2021]".
- Martin Reddy. 1998. Specification and evaluation of level of detail selection criteria. Virtual Reality 3, 2 (1998), 132–143.

- Zhou Wang, Alan Conrad Bovik, Ligang Lu, and Jack L. Kouloheris. 2001. Foveated wavelet image quality index. In *Applications of Digital Image Processing XXIV*, Andrew G. Tescher (Ed.), Vol. 4472. International Society for Optics and Photonics, SPIE, 42 – 52. https://doi.org/10.1117/12.449797
- Zhou Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. 2004. Image Quality Assessment: From Error Visibility to Structural Similarity. *Trans. Img. Proc.* 13, 4 (April 2004), 600–612. https://doi.org/10.1109/TIP.2003.819861
- Lingjie Wei and Yuji Sakamoto. 2019. Fast calculation method with foveated rendering for computer-generated holograms using an angle-changeable ray-tracing method. *Appl. Opt.* 58, 5 (Feb 2019), A258–A266. https://doi.org/10.1364/AO.58.00A258
- Martin Weier, Thorsten Roth, André Hinkenjann, and Philipp Slusallek. 2018. Foveated Depth-of-Field Filtering in Head-Mounted Displays. ACM Trans. Appl. Percept. 15, 4, Article 26 (Sept. 2018), 14 pages. https://doi.org/10.1145/3238301
- Martin Weier, Thorsten Roth, Ernst Kruijff, André Hinkenjann, Arsène Pérard-Gayot, Philipp Slusallek, and Yongmin Li. 2016. Foveated Real-Time Ray Tracing for Head-Mounted Displays. Comput. Graph. Forum 35, 7 (Oct. 2016), 289–298. https: //doi.org/10.1111/cgf.13026
- M. Weier, M. Stengel, T. Roth, P. Didyk, E. Eisemann, M. Eisemann, S. Grogorick, A. Hinkenjann, E. Kruijff, M. Magnor, K. Myszkowski, and P. Slusallek. 2017. Perception-driven Accelerated Rendering. *Comput. Graph. Forum* 36, 2 (May 2017), 611–643. https://doi.org/10.1111/cgf.13150
- Frank W. Weymouth. 1958. Visual Sensory Units and the Minimal Angle of Resolution*. American Journal of Ophthalmology 46, 1, Part 2 (1958), 102 – 113. https://doi.org/ 10.1016/0002-9394(58)90042-4
- Jui-Yi Wu and Jonghyun Kim. 2020. Prescription AR: a fully-customized prescriptionembedded augmented reality display. Opt. Express 28, 5 (Mar 2020), 6225-6241. https://doi.org/10.1364/OE.380945
- Heinz Wässle, Ulrike Grünert, Jürgen Röhrenbeck, and Brian B. Boycott. 1990. Retinal ganglion cell density and cortical magnification factor in the primate. Vision Research 30, 11 (1990), 1897–1911. https://doi.org/10.1016/0042-6989(90)90166-I Optics Physiology and Vision.
- Yanyu Xu, Yanbing Dong, Junru Wu, Zhengzhong Sun, Zhiru Shi, Jingyi Yu, and Shenghua Gao. 2018. Gaze Prediction in Dynamic 360° Immersive Videos. In The IEEE Conference on Computer Vision and Pattern Recognition (CVPR).
- Fahri Yaraş, Hoonjong Kang, and Levent Onural. 2010. State of the Art in Holographic Displays: A Survey. J. Display Technol. 6, 10 (Oct 2010), 443–454. http://jdt.osa.org/ abstract.cfm?URI=jdt-6-10-443
- Jiannan Ye, Anqi Xie, Susmija Jabbireddy, Yunchuan Li, Xubo Yang, and Xiaoxu Meng. 2022. Rectangular Mapping-based Foveated Rendering. In 2022 IEEE Conference on Virtual Reality and 3D User Interfaces (VR). 756–764. https://doi.org/10.1109/ VR51125.2022.00097
- Hector Yee, Sumanita Pattanaik, and Donald P. Greenberg. 2001. Spatiotemporal Sensitivity and Visual Attention for Efficient Rendering of Dynamic Environments. ACM Trans. Graph. 20, 1 (Jan. 2001), 39–65. https://doi.org/10.1145/383745.383748
- Shivanthan Yohanandan, Andy Song, Adrian G. Dyer, and Dacheng Tao. 2018. Saliency Preservation in Low-Resolution Grayscale Images. In *Computer Vision – ECCV* 2018, Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss (Eds.). Springer International Publishing, Cham, 237–254.
- Chanhyung Yoo, Jianghao Xiong, Seokil Moon, Dongheon Yoo, Chang-Kun Lee, Shin-Tson Wu, and Byoungho Lee. 2020. Foveated display system based on a doublet geometric phase lens. *Opt. Express* 28, 16 (Aug 2020), 23690–23702. https://doi.org/ 10.1364/OE.399808
- A. Yoshida, J. P. Rolland, and J. H. Reif. 1995. Design and applications of a high-resolution insert head-mounted-display. In *Proceedings Virtual Reality Annual International Symposium* '95. 84–93. https://doi.org/10.1109/VRAIS.1995.512483
- Akitoshi Yoshida, Jannick P. Rolland, and John H. Reif. 1995. Optical design and analysis of a head-mounted display with a high-resolution insert. In *Novel Optical Systems Design and Optimization*, Jose M. Sasian (Ed.), Vol. 2537. International Society for Optics and Photonics, SPIE, 71 – 82. https://doi.org/10.1117/12.216372
- J. You, T. Ebrahimi, and A. Perkis. 2014. Attention Driven Foveated Video Quality Assessment. *IEEE Transactions on Image Processing* 23, 1 (2014), 200–213. https: //doi.org/10.1109/TIP.2013.2287611
- Zipeng Zheng, Zhuo Yang, Yinwei Zhan, Yuqing Li, and Wenxin Yu. 2018. Perceptual Model Optimized Efficient Foveated Rendering. In *Proceedings of the 24th ACM Symposium on Virtual Reality Software and Technology* (Tokyo, Japan) (VRST '18). ACM, New York, NY, USA, Article 109, 2 pages. https://doi.org/10.1145/3281505. 3281588