Retinotopic Foveated Rendering

Yan Zhang*
Shanghai Jiao Tong University

Keyao You†
Shanghai Jiao Tong University

Xiaodan Hu‡
Nara Institute of Science and Technology

Hangyu Zhou§
Shanghai Jiao Tong University

Kiyoshi Kiyokawa¶
Nara Institute of Science and Technology

Xubo Yang‖
Shanghai Jiao Tong University

Figure 1: We propose retinotopic foveated rendering (RFR) based on the horizontal-vertical asymmetry (HVA) and vertical-meridian asymmetry (VMA) that are evident from advancements in retinotopy research. We first derived the asymmetric model for cortical magnification factor (CMF), which is proportional to the visual acuity. The comparison of the radially symmetric CMF model in typical FR and the radially asymmetric CMF model in the proposed RFR is illustrated in the left column, indicating the overestimated visual acuity near the vertical meridian in existing FR methods. The asymmetric CMF model is derived from the fMRI data of the early cortex, which defines the asymmetric rendering precision along the horizontal, lower-vertical, and upper-vertical meridian at ratios of 0.72:0.57:0.32. The across-meridian asymmetry is achieved through interpolation in terms of the quasi-linear degradation nature of the visual acuity from the horizontal meridian to the vertical meridian. As a result, RFR keeps consistent rendering precision in the horizontal meridian while eliminating excess rendering in the upper and lower FOV, as shown in the zoom-in images enclosed in red, orange, and blue boxes in the middle column. Notice that the degradation of image quality in RFR is unperceivable since it follows the radially asymmetric regression of CMF. We validated the proposed RFR method by integrating the asymmetric computation model with both polar and rectangular mapping-based foveated rendering. The results show that RFR methods reduce shading in the fragmented shader by 27.2% on average compared with typical FR methods, resulting in the increase in frame rate by 17.45%.

Abstract

Foveated rendering (FR) improves the rendering performance of virtual reality (VR) by allocating fewer computational loads in the peripheral field of view (FOV). Existing FR techniques are built based on the radially symmetric regression model of human visual acuity. However, horizontal-vertical asymmetry (HVA) and vertical meridian asymmetry (VMA) in the cortical magnification factor (CMF) of the human visual system have been evidenced by retinotopy research of neuroscience, suggesting the radially asymmetric regression of visual acuity. In this paper, we begin with functional magnetic resonance imaging (fMRI) data, construct an anisotropic CMF model of the human visual system, and then introduce the first radially asymmetric regression model of the rendering precision for FR applications. We conducted a pilot experiment to adapt the proposed model to VR head-mounted displays (HMDs). A user study demonstrates that retinotopic foveated rendering (RFR) provides participants with perceptually equal image quality compared to typical FR methods while reducing fragments shading by 27.2% averagely, leading to the acceleration of 1/6 for graphics rendering. We anticipate that our study will enhance the rendering performance of VR by bridging the gap between retinotopy research in neuroscience and computer graphics in VR.

*e-mail: yan-zh@sjtu.edu.cn
†e-mail: youkeyao@sjtu.edu.cn
‡e-mail: hu.xiaodan.ht1@is.naist.jp
§e-mail: zhouhy0571@sjtu.edu.cn
¶e-mail: kiyo@is.naist.jp
‖e-mail: yangxubo@sjtu.edu.cn (corresponding author)

Index Terms: Computing methodologies—Computer graphics—Graphics systems and interfaces—Virtual reality; Computing methodologies—Computer graphics—Rendering

1 Introduction

Virtual reality (VR) offers an immersive experience by providing users with high-quality and low-latency graphics. Although the tremendous advances in hardware have increased the computation load of VR platforms, delivering realistic virtual content in real-time still faces challenges from graphics rendering and streaming. Recently, foveated rendering has demonstrated the advantage of displaying perceptually equal images while reducing computation load and streaming bandwidth dramatically. Thus, it is considered to be one of the essential techniques of VR rendering techniques [10]. So far, foveated rendering has become more powerful by utilizing better rendering pipelines [25], integrating with cutting-edge imaging techniques [6, 33], optimizing the peripheral blur [20, 38], and introducing a more comprehensive perception model [22, 34]. Very recently, a few VR head-mounted displays (HMDs), such as Meta Quest Pro[1] and Apple Vision Pro[2], were released with the support for the foveated rendering feature, showing appealing graphics performance.

Although the implementation of foveated rendering could be technically complicated due to the graphics techniques adopted, the basic idea is simple - arranging the computation load in terms of the variance of visual performance in the user's field of view (FOV). Correspondingly, existing foveated rendering methods are built on the well-known concept that the human visual acuity isotropically deteriorates with the visual eccentricity increase. Thus, the computation load follows a radially symmetric falloff from the user's gaze to the periphery [10].

However, the latest neuroscience research on retinotopy has revealed the human visual system possesses horizontal-vertical asymmetry (HVA) and vertical meridian asymmetry (VMA) of the cortical magnification factor (CMF) on the early visual cortex. Since CMF has been proven to be proportional to a variety of visual performances, including visual acuity [7] and contrast sensitivity [29], the HVA and VMA indicate the distribution of visual performance follows an anisotropic instead of isotropic regression across human vision [16]. Fig. 2 illustrates the retinotopy in the V1 area of the visual cortex, which is proven to strongly relate with the visual acuity [7], within the FOV of $\pm 8^\circ$ based on the functional magnetic resonance imaging (fMRI) data given by Human Connectome Project (HCP) [35]. In addition to the aggressive regression from the fovea to the parafovea (see Fig. 2 (a)), which forms the physiological basis of typical foveated graphics, the retinotopy in the V1 visual field shows a noticeable radially asymmetric distribution. Specifically, the horizontal FOV extends more than the vertical FOV, and the lower FOV extends more than the upper FOV.

Therefore, it is essential to optimize foveated rendering in light of the retinotopy of the human visual system. In this paper, we propose retinotopic foveated rendering (RFR). Fig. 1 illustrates an overview of the present study. A radially asymmetric model of visual acuity is derived from fMRI data of HCP, with the regression that follows the retinotopy pattern. Then, we implement retinotopic foveated rendering by integrating the fMRI-driven model into conventional foveated rendering methods, using both log-polar mapping [25] and rectangular mapping [41]. The radial asymmetrical allocation of computation load is enabled by using asymmetric factors for the horizontal, lower-vertical, and upper-vertical meridians, respectively. A pilot experiment is conducted to validate the proposed method by evaluating the maximum rendering precision in the fovea and the global falloff to the periphery. A final user study with 24 participants proves that the retinotopic foveated rendering supports perceptually equal graphics to typical foveated rendering, while the shading performance is accelerated by 27.2% by avoiding the redundant rendering in the upper and lower FOVs.

[1] https://www.meta.com/quest/quest-pro

[2] https://www.apple.com/apple-vision-pro

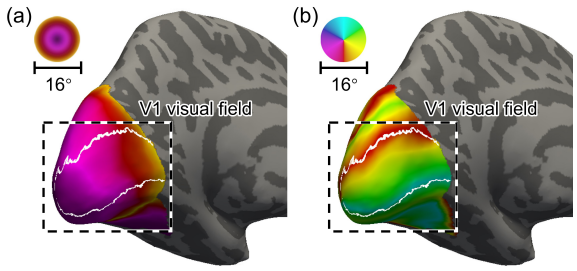

Figure 2: The retinotopy of the FOV within $\pm 8^\circ$ in the V1 visual field of the left hemisphere. The same FOV is mapped to the V1 area according to the eccentricities (a) and polar angle (b), indicating the radially asymmetric regression of the human visual performance. Our retinotopic foveated rendering takes both the eccentricity-dependent and the polar-angle-dependent variation into account.

The primary contributions of the present paper are highlighted as follows:

1) We introduce the first fMRI-driven regression model of rendering precision for foveated rendering methods, allocating the computation load more efficiently by taking the HVA and VMA of the human visual system into account.

2) We introduce the RFR algorithm by formulating closed-form mapping between screen pixels and shading fragments based on the fMRI-driven model. A pilot experiment is conducted for validating the RFR algorithm.

3) We conduct a user study to evaluate the graphics produced by RFR and typical foveated rendering. The participative results indicate all the evaluated images are perceptually equal, while RFR saving 27.2% shading cost, hereby, performs $1/6$ faster than typical foveated rendering methods approxiamtely.

2 Related Works

2.1 Foveated Rendering

The human visual system (HVS) supports a monocular FOV around 180° horizontally and 120° vertically, where the performance maintains a non-uniform distribution from the central to the peripheral [32]. In general, the visual acuity of adults could be as high as 60 cycles per degree (cpd) in the fovea ($< 3^\circ$) but degrades to lower than 5cpd in the periphery ($> 30^\circ$) [40]. Therefore, it is reasonable to minimize the image quality while keeping users unnoticeable by allocating the computation load corresponding to attributes of the HVS, which forms the basis of foveated rendering.

Real-time foveated rendering was demonstrated with the multiresolution video system by Geisler et al. [9]. Then, Guenter et al. applied the concept to modern rendering pipelines on GPUs, showing a promising performance of generating perceptually equal graphics while saving the computation load considerably [10]. The following works on foveated rendering include improving the image quality by optimizing artifacts in peripheral vision [20, 27, 41] and generating perceptually better blurry effects [38], as well as validating more aggressive computation models by counting users' contrast sensitivity [34] and attention [22]. In addition, foveated rendering also benefits the cutting-edge techniques, such as light field displays [33], computational holography [6], and ray-tracing graphics [31, 39], by reducing the computation efficiently.

As foveated rendering attenuates computational load by avoiding the unresolved image details to human eyes, constructing the computation model that meets the visual acuity distribution of the HVS primarily impacts the rendering performance. Guenter et al. propose a hyperbolic model that decomposes an image into three layers

with degraded sampling rates against the eccentricity [10], while it overestimates the visual acuity in the periphery and thus causes redundant computation. A derived model of the hyperbolic falloff is utilized by Weier et al., for achieving better motion perception and fewer artifacts in the periphery [39]. Meng et al. [25] propose a log-polar falloff that further reduces the computation load in foveated rendering by reviewing the visual acuity model derived from the physiological and anatomical evidence [12]. Very recently, Ye et al. exploited a rectangular falloff for foveated rendering, optimizing the artifacts for horizontal and vertical image patterns [41].

However, research on retinotopy has proven the noticeable HVA and VMA in HVS, suggesting the anisotropic model of visual acuity [17]. Existing foveated rendering methods leverage the computation falloff by following the isotropic regression of visual acuity, leading to a global computation redundancy in the upper and lower FOVs. In this paper, we review the advances in retinotopy research, exploiting the open-source fMRI dataset to predict the precise visual acuity distribution, and then deriving the first fMRI-driven model for foveated rendering. The results expect to benefit foveated rendering by introducing a more efficient computation model and bridging foveated graphics with fMRI-based neuroscience, facilitating the development of rendering techniques in the future.

2.2 Retinotopy Based on fMRI

Retinotopy is the mapping of the visual stimulus from the retina to the visual cortex via the optic nerve. The retinotopy in the primary visual cortex of stimulus within a unit degree of vision is defined as the cortical magnification factor (CMF). In general, the higher the CMF is, the better the visual performance will be [32]. Specifically, psychophysical and anatomical evidence has suggested the linear proportion between the visual acuity and CMF [7, 36, 37], allowing the indirect measurement of retinotopy, and vice versa. As a result, isotropic retinotopy models were established according to tremendous psychophysical experiments, where CMF progressively reduces with the eccentricity increase [12]. However, with the advancements in fMRI techniques in the 21st century, neuroscience has directly measured retinotopy. This is achieved by presenting visual stimuli to participants and subsequently recording the activated areas within the visual cortex [35], revealing the more precise model of retinotopy with the anisotropic distribution [30].

Current fMRI-based retinotopy demonstrates that the performance of HVS not only depends on the vision area's eccentricity but also is influenced by the radial position in the FOV [17]. Specific to the CMF, research based on the open-source fMRI dataset indicates that the primary cortex areas corresponding to the horizontal meridian are approximately 60% greater than the areas for the vertical meridian within the measured FOV of $\pm 8^\circ$, and the lower-vertical meridian vision is approximately 25% greater than the symmetrical upper part [16]. The radial asymmetry of CMF is demonstrated to explain the relationship between visual acuity and contrast sensitivity with the polar angle well [1, 2, 4, 5].

Foveated rendering aims to minimize the computation load by omitting graphics details users cannot resolve, yet existing methods work on the radial symmetrical computation model of the psychophysical or empirical models of visual acuity, neglecting the advances in retinotopy. The proposed work drives foveated rendering based on fMRI-based retinotopy, thus, improving the performance by validating the HVA and VMA in the computation model.

3 fMRI-Driven CMF Model

3.1 CMF and Foveated Rendering

The analytic model for CMF can be categorized into local mapping and global mapping. The local mapping of CMF is defined as the stimulated area in the primary cortex by a visual input within a unit degree of the FOV. Since visual acuity is proportional to the local mapping of CMF [7], it is reasonable to leverage the computation load regression in foveated rendering as the local mapping varies across the FOV. Although the variation of CMF by the polar angle is less explored in the retinotopy based on psychophysical experiments, the eccentricity-dependent regression of CMF is well formulated and also demonstrated by fMRI experiments [23]. The local mapping of CMF along the four meridians in the FOV can be predicted by the following equation [24]:

$$M = \frac{M_0}{1 + E/E_2} \tag{1}$$

Where M_0 gives the maximum CMF in the fovea, and E_2 determines how fast the CMF M regresses with an increasing eccentricity E. Although there are a few analytic functions for predicting the eccentricity-dependent CMF, Eq. 1 is demonstrated to be simple and efficient since it provides a good prediction for both the fovea and the periphery [32].

Compared to leveraging computation load by using the local mapping of CMF, a global mapping is more straightforward for implementing foveated rendering in rendering pipelines. The global mapping of CMF can be derived by integrating Eq. 1 along the visual meridian, which is expressed as [32]:

$$\delta = \int_0^E M(E)dE = M_0 E_2 \cdot \ln \frac{E_2 + E}{E_2} \tag{2}$$

Where δ represents the cortical distance, in millimeters, from the center of the fovea to the visual area at the eccentricity of E on the early cortex, thus, is proportional to the amount of resolved visual input within the corresponding FOV. Since the visual input presented to the users in foveated rendering is generated in the shading stage, the projection between the screen pixels and the shading fragments should follow δ in Eq. 2. The simplified expression of Eq. 2, which is $\delta \propto \ln E$, has been used to drive the rendering precision falloff by Meng et al. in their kernel-based foveated rendering method, showing a promising computation load attenuation [25].

3.2 Asymmetric CMF Prediction Based on fMRI

The prediction of CMF by using Eq. 1 or Eq. 2 neglects the radial asymmetry in the FOV, but it performs well when the visual field is close to the meridians. Benefiting from the progress of fMRI techniques, neuroscience directly records the mapping of visual input on the primary visual cortex via the retina and optic nerve. However, the precise measurement of retinotopy is still challenging due to the limited spatial resolution of fMRI [19]. Furthermore, the visual stimulus to participants in fMRI experiments is in a small FOV since the display is remotely located due to the strong magnetic field generated by fMRI devices. Consequently, recording CMF for an arbitrary FOV by fMRI is infeasible at this point [17].

Nevertheless, recent works on the open-source dataset, HCP, provide reasonable fMRI data for fitting the CMF model in visual meridians [3]. We adopted the processed fMRI data in the FOV of $1^\circ - 7^\circ$ by Benson et al. to build the asymmetric CMF regression along the horizontal, lower-vertical, and upper-vertical meridians [4]. Meanwhile, Kuper et al. suggest the prediction of CMF based on Eq. 1 with Horton et al.' M_0 and E_2 matches close to the fMRI results in the horizontal meridian [18, 23]. As a consequence, we decided to derive the analytic function of global mapping of CMF for each meridian by utilizing Eq. 2 with a constant $M_0 = 23.07\text{mm}/^\circ$ to fit the processed fMRI data. The asymmetric regression of CMF at different meridians is achieved by using distinct E_2, as shown in Table 1. Local mapping of CMF is, hereby, predicted by substituting the derived E_2 into Eq. 1. Fig. 3 plots the curves of CMF versus eccentricity of global mapping (left) and local mapping (right). The red, yellow, and blue lines represent the predicted CMF in the horizontal, lower-vertical, and upper-vertical meridians, respectively, and the measured global mapping of CMF by fMRI are labeled by circle,

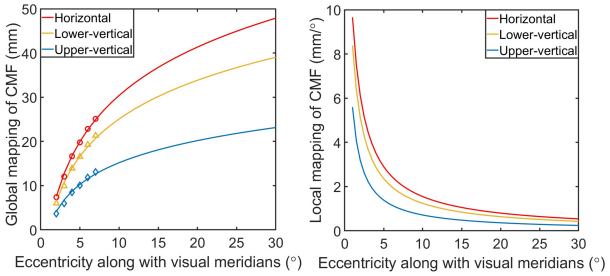

Figure 3: The asymmetry of CMF within a FOV of 30° is illustrated by the fitted curves of global mapping (left) and local mapping (right). Source fMRI data is labeled as the circle, triangle, and diamond markers in the figure of global mapping. Human visual acuity expects to show the same HVA and VMA as CMF due to the linear proportion nature [7], suggesting foveated rendering should possess an asymmetric computation load.

triangle, and diamond markers in the left figure. Please check the supplementary material for the detailed derivation of fMRI-driven CMF model.

Table 1: E_2 for each meridian are derived by utilizing Eq. 2 to fit the processed fMRI data of HCP, where M_0 is set to be a constant measured by Horton et al. [18].

Meridians	Horizontal	Lower-vertical	Upper-vertical
E_2	0.72	0.57	0.32
R square	0.9967	0.9899	0.9798

3.3 Interpolation Model With K(θ)

So far, the CMF functions are built for each meridian with different E_2 derived, while the existing dataset does not support the precise measurement of retinotopy for visions away from visual meridians, which hinders building a model to predict CMF at arbitrary FOV via simply fitting fMRI data. Fortunately, psychophysical experiments have shown a quasi-linear regression for both visual acuity [2] and contrast sensitivity [1] at iso-eccentricity visions from the horizontal meridian to the vertical meridian as the polar angle varies. Thus, it suggests that the CMF also follows the quasi-linear variation between meridians [4, 7]. Consequently, we build the interpolation model for predicting the local mapping of CMF at arbitrary FOV with an eccentricity of E and polar angle of θ:

$$M(E,\theta) = K(\theta)\cdot M_V(E) + (1-K(\theta))\cdot M_H(E) \tag{3}$$

Where θ is the polar angle measured counter-clockwise from the horizontal meridian, and $K(\theta)$ is a monotonically increasing function within a quadrant. $M_v(E)$ and $M_H(E)$ are the asymmetric CMFs in the vertical and horizontal meridian with the same eccentricity E, which can be computed by:

$$M_V(E) = \begin{cases} \dfrac{M_0}{1+E/E_{2VU}}, & 0 \le \theta < \pi \\ \dfrac{M_0}{1+E/E_{2VL}}, & \pi \le \theta < 2\pi \end{cases} \qquad M_H(E) = \frac{M_0}{1+E/E_{2H}} \tag{4}$$

Where E_{2H}, E_{2VL}, and E_{2VU} are the derived E_2 from the fMRI data for predicting the asymmetric CMF in the horizontal, lower-vertical, and upper-vertical meridians in Table 1, and $M_0 = 23.07\text{mm}/^\circ$ is the CMF at the fovea given by Horton et al. [18].

Although the exact variation of CMF with the polar angle is not formulated by neuroscience yet, a few studies suggest that the HVA is most pronounced at the vertical meridian and goes to be marginal at the diagonal vision [1,15]. Therefore, we pick the logistic function with the midpoint shifts to $\pi/4$ as the interpolation kernel $K(\theta)$ in Eq. 3, as it performs a quasi-linear increase in a quadrant and keeps a gentle slope reaching the meridians. The image in the top-left of Fig. 1 illustrates the radial asymmetric contours of CMF (as well as the human visual acuity due to the linear proportion) within a FOV of 35° based on Eq. 3. Typical foveated rendering follows the symmetric contours indicated by the white concentric circles, leading to the considerable redundancy of computation load in the upper and lower FOV.

4 Algorithm for Retinotopic Foveated Rendering

4.1 Asymmetric Rendering Precision From CMF Models

The regression model of CMF can be utilized to drive foveated graphics in many aspects, such as optimizing the streaming of graphics assests [33] and the rasterization of primitives [8]. Since shading becomes the dominant computation cost in modern rendering pipelines [11], a simple but efficient way to optimize rendering performance is implementing the foveation-based non-injective projection between the screen space and the fragment shader space [10, 25, 41].

Therefore, we define the rendering precision P as the number of fragments mapped to an image pixel. The lower P of pixels is, the less shading in the fragment shader will be, thus, reducing the computation load. Let $P \propto M$, the rendering precision in our RPR method can be derived from the asymmetric CMF model in Eq. 3 and Eq. 4:

$$P = \frac{\partial u \partial v}{\partial x \partial y} = \begin{cases} \dfrac{K(\theta)\cdot P_0}{1+\frac{r}{C_{VU}D}} + \dfrac{(1-K(\theta))\cdot P_0}{1+\frac{r}{C_H D}}, & 0 \le \theta < \pi \\ \dfrac{K(\theta)\cdot P_0}{1+\frac{r}{C_{VL}D}} + \dfrac{(1-K(\theta))\cdot P_0}{1+\frac{r}{C_H D}}, & \pi \le \theta < 2\pi \end{cases} \tag{5}$$

Where (u,v) is the coordinate of fragments and (x,y) is the coordinate of the screen pixel. $r = \|x-x_0, y-y_0\|_2$ is the pixel distance of the pixel at (x,y) to the user's gaze point (x_0,y_0). The eccentricity distance of a pixel to (x_0,y_0) is derived by r/D, where D is the average angular resolution of the HMD, and $\theta = \arctan(\frac{x-x_0}{y-y_0})$ is the polar angle of the pixel. The maximum rendering precision for the gaze area is P_0, and it is conducted to perform a radial asymmetric regression to (x,y) by introducing the asymmetric coefficient C_H, C_{VL}, and C_{VU}, where $C_H : C_{VL} : C_{VU} = E_{2H} : E_{2VL} : E_{2VU}$.

4.2 Model Simplification for Rendering Pipeline Adaption

For implementing Eq. 5 to operate foveated rendering, we refer to the two-pass kernel-based method by Meng et al. [25]. Their rendering pipeline consists of two passes, where the first pass maps the screen pixels to the fragments according to the varied rendering precision before entering the shading stage, and the second pass outputs the displayed pixels by mapping the shaded fragments to the screen space reversely. However, the asymmetric regression model defined by Eq. 5 faces challenges when integrating into the two-pass approach, since we did not find an analytical solution for the coordinate transform with the two-step mapping. As a consequence, we simplify Eq. 5 by removing the constant part in the denominator:

$$P_{approx} = \frac{\partial u \partial v}{\partial x \partial y} = S(\theta)\cdot \frac{P_0}{r^\alpha} =$$

$$\begin{cases} \frac{K(\theta)\cdot P_0 C_{VU} D}{r^{\alpha}} + \frac{(1-K(\theta))\cdot P_0 C_H D}{r^{\alpha}}, & 0 \le \theta < \pi \\ \frac{K(\theta)\cdot P_0 C_{VL} D}{r^{\alpha}} + \frac{(1-K(\theta))\cdot P_0 C_H D}{r^{\alpha}}, & \pi \le \theta < 2\pi \end{cases} \tag{6}$$

Note that we mapped the pixel distance r to r^{α} since it allows the global adjustment of the regression from the fovea to the periphery while keeping the asymmetric computation load among different quadrants. Moreover, the simplification separates Eq. 6 into polar-angle-dependent asymmetric coefficient $S(\theta)$, the maximum rendering precision P_0 at the fovea, and eccentricity-dependent regression coefficient r^{α}, benefiting the leverage of the computation load in foveated rendering.

Fig. 4 depicts the visualization of the asymmetric rendering precision based on Eq. 6. The simulated FOV is H60°×V60° for the sampling map (left column). The angular resolution D is 20cpd for all configurations. In general, all the configurations show a heart-shaped distribution, which is similar to the CMF contours in Fig. 1. The maximum rendering precision in the fovea is set to be 10, 30, and 50, which increases the global rendering precision in the visual fields. The α shifts from 1 to 1.8 with an interval of 0.4, leading to the increasing reduction of computation load from the fovea to the periphery. The sparsest shading rate is found in visual fields close to the upper-vertical meridian, which results from the smallest asymmetric coefficient C_{VU}. Besides each rendering precision map, images with the FOV of H90°×V96° are rendered by using our PM-RFR methods (see Section 4 for details), where the same P_0 and α are adopted and the gaze point is set to be the door handle in the center. For all configurations, the zoom-in images demonstrate the distinct image quality among meridians due to asymmetric rendering precision. Apparently, raising the computation load with a bigger P_0 results in better image quality for all visual fields. The image quality in the central vision (white box) is efficiently improved by increasing α from 1 to 1.4, which is reasonable since more computation load is allocated for this area. The $\alpha = 1.8$ does not provide significant improvement for the central but causes the deterioration of image quality and noticeable artifacts in the left-horizontal (red box), lower-vertical (orange box), and upper-vertical (blue box) visual fields.

The bottom row depicts in Fig. 4 the rendering precision in each visual meridian by utilizing Eq. 5 (solid lines) and Eq. 6 (dot lines) with $P_0 = 30$. With all regression velocity by α, the average difference in rendering precision at the eccentricity of 1° between Eq. 5 and Eq. 6 is 28.9%. The difference reduces to be 11% at the eccentricity of 5° and 5.7% at the eccentricity of 10° with $\alpha = 1$, thus, the original and simplified models perform similarly with $\alpha = 1$. All configurations possess HVA and VMA of 47.7% and 56.9%, leading to the aggressive degradation of image quality from the horizontal meridian to the vertical meridian.

4.3 RFR Algorithms

Then, we explored the foveated rendering by adopting the radial asymmetric rendering precision defined by Eq. 6. We reviewed two representative mapping methods for foveated rendering, which are the log-polar mapping by Meng et al. [25] and the rectangular mapping by Ye et al. [41], and validated the proposed asymmetric distribution. The two mapping methods project pixels in screen space into the fragment shader space defined by the Cartesian coordinate system and the polar coordinate system, leading to the distinct artifacts in the periphery [41], thus, requiring different regression of rendering precision in practice. Therefore, we proposed rectangular-mapping-based RFR (RM-RFR) and polar-mapping-based RFR (PM-RFR). The following pilot experiment in Sec. 5.2 suggests the different slopes of the regression models in RM-RFR and PM-RFR. Moreover, this section focuses on optimizing the existing foveated rendering methods by integrating the proposed fMRI-driven asymmetric computation model, so that we omit the details of rendering pipelines here.

4.3.1 Rectangular-Mapping-Based RFR

RM-RFR performs the projection between two Cartesian coordinate systems. Hence, it allows the naive usage of Eq. 6 to formulate the transformation between the screen pixel coordinate (x,y) and the fragment coordinate (u,v), by equally separating P_{approx} into the horizontal and vertical dimension:

$$P_{approx} = \frac{\partial u}{\partial x}\cdot\frac{\partial v}{\partial y} = \left(\frac{\sqrt{S(\theta)\cdot P_0}}{r^{\frac{\alpha}{2}}}\right)^2 \tag{7}$$

Then, the mapping from (x,y) to (u,v) can be derived from the integral of Eq. 7:

$$\begin{aligned} u &= \sqrt{\frac{2}{2-\alpha}}\cdot\sqrt{S(\theta)\cdot P_0}\cdot r^{1-\frac{\alpha}{2}}\cdot\cos(\theta) + u_0 \\ v &= \sqrt{\frac{2}{2-\alpha}}\cdot\sqrt{S(\theta)\cdot P_0}\cdot r^{1-\frac{\alpha}{2}}\cdot\sin(\theta) + v_0 \end{aligned} \tag{8}$$

Where (u_0, v_0) is the coordinate of the central fragment. Similarly, the inverse transformation from the shaded fragment to the displayed pixel can be solved as:

$$\begin{aligned} x &= \left(\frac{r'^2}{S(\theta)\cdot P_0}\cdot\frac{2-\alpha}{2}\right)^{\frac{1}{2-\alpha}}\cdot\frac{u-u_0}{r'} + x_0 \\ y &= \left(\frac{r'^2}{S(\theta)\cdot P_0}\cdot\frac{2-\alpha}{2}\right)^{\frac{1}{2-\alpha}}\cdot\frac{v-v_0}{r'} + y_0 \end{aligned} \tag{9}$$

Where $r' = \|u-u_0, v-v_0\|_2$ is the pixel distance of the fragment at (u,v) to the origin, and (x_0, y_0) is gaze point of the user. Note that α should be strictly smaller than 2 for validating the closed-form transformation between (x,y) and (u,v) in the RM-RFR method.

4.3.2 Polar-Mapping-Based RFR

Different from the RM-RFR method, PM-RFR builds the mapping between the Cartesian coordinate of the screen and the Polar coordinate of the fragment space, which prevents the leverage of the computation by using Eq. 6 directly. In order to operate foveated rendering based on the asymmetric computation model, we exploit Jacobian determinant to Eq. 6. Thus, Eq. 6 in the polar coordinate system can be derived as:

$$P_{approx} = \frac{\partial u \partial v}{\partial r \partial \theta} = \frac{r\partial u \partial v}{\partial x \partial y} = \frac{S(\theta)\cdot P_0}{r^{\alpha-1}} \tag{10}$$

Different from the separation of P_{approx} in Eq. 7, the rendering precision in PM-RFR is unequally allocated in the radial and tangential dimensions:

$$\frac{\partial u}{\partial r} = \frac{S(\theta)\cdot P_0}{r^{\alpha-1}}, \quad \frac{\partial v}{\partial \theta} = 1 \tag{11}$$

Then, the mapping of screen pixel (x,y) to the fragments (u,v) is formulated:

$$u = \frac{S(\theta)\cdot P_0}{2-\alpha}\cdot r^{2-\alpha}, \quad v = \theta \tag{12}$$

Similarly, the inverse mapping is:

$$x = \left(\frac{u\cdot(2-\alpha)}{S(\theta)\cdot P_0}\right)^{\frac{1}{2-\alpha}}\cdot cos(v), \quad y = \left(\frac{u\cdot(2-\alpha)}{S(\theta)\cdot P_0}\right)^{\frac{1}{2-\alpha}}\cdot sin(v) \tag{13}$$

Note that the PM-RFR method also follows the constraint of $\alpha < 2$ in Eq. 12 and Eq. 13.

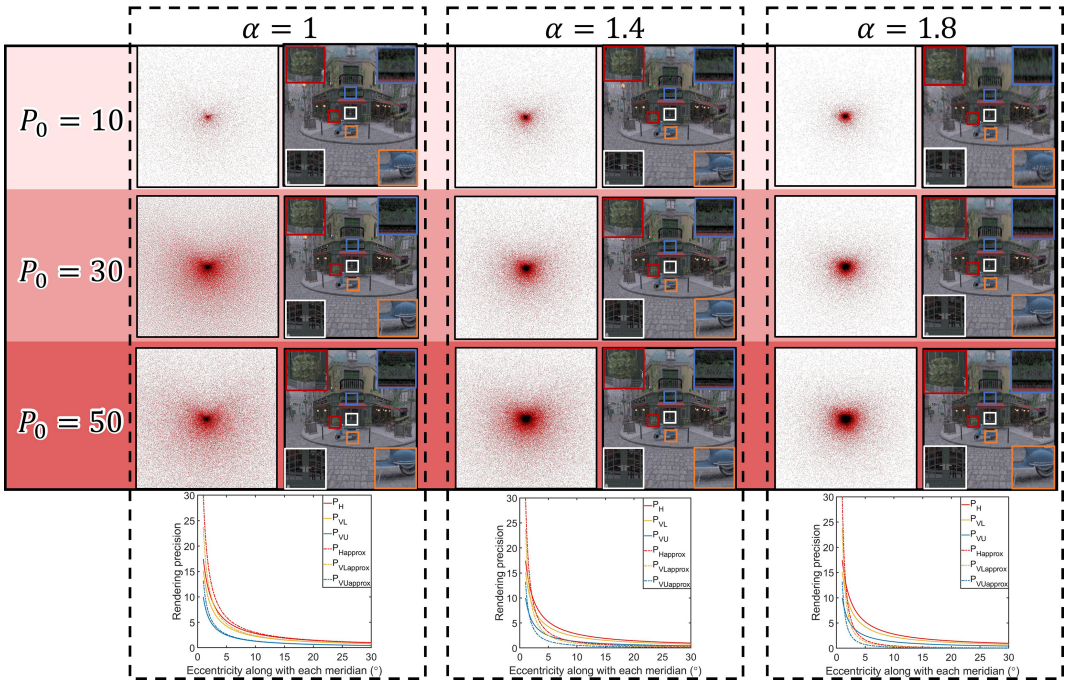

Figure 4: The visualization of radial asymmetric rendering precision of the RFR method based on Eq. 6. The simulated FOV for the rendering precision map is H60°×V60°. The global computation load is leveraged by raising the maximum rendering precision P_0. The regression velocity from the fovea to the periphery is modulated by α. The figures beside each precision map are the image rendered by our polar-mapping-based RFR method. The same configuration of the P_0 and α is utilized for rendering, and the FOV of images is H90°×V96°. The bottom row shows the asymmetric reduction of rendering precision along each visual meridian by utilizing the original model Eq. 5 and the simplified model Eq. 6, showing the HVA and VMA of 47.7% and 56.9% for RFR with all configurations.

4.4 Implementation of RFR

Although NVIDIA Variable Rate Shading (VRS) provides a robust framework for implementing foveated rendering, it only supports several pre-defined shading rates [26, 27, 34]. The RFR method defines the continuous variation of rendering precision both along and across visual meridians. Therefore, we determined to implement RFR with customized rendering pipelines that modulate the shading rate continuously. Moreover, foveated rendering is demonstrated to accelerate both rasterization-based methods and ray-tracing methods [10, 21], while the advantage of saving computation costs by adopting the proposed asymmetric model is more crucial in the latter. As a consequence, we validate RM-RFR and PM-RFR with ray-tracing in our customized pipeline. Please refer to the pseudo-code in the supplementary material for the realization of RFR.

Fig. 5 shows the comparison between FR methods and the proposed RFR methods using Nvidia's Bistro and Sponza scenes. Filtering is omitted in the rendering pipelines to make the characteristics of RM-RFR and PM-RFR most pronounced. The configurations for RFR methods are obtained from the pilot experiment in Section 5.2. The same $P_0 = 33$ is utilized, and the different α of 1.2 and 1.6 are adopted for the RM-RFR and PM-RFR methods, respectively. All the foveated methods have shading textures with an approximately equal size, which is shown at the bottom. The zoom-in images at the visual fields of central, left-horizontal, upper-vertical, and lower-vertical visual fields show the HVA and VMA of the rendering precision in RFR methods compared to the typical FR methods with symmetric rendering precision. Due to HVA and VMA in the shading stage, RFR methods avoid the redundant computation of fragments, which causes the larger black areas in the texture. As a result, RM-RFR and PM-RFR reduce 30.6% and 23.8% of shading than typical rectangular and polar mapping-based FR methods.

We evaluated the rendered Bistro scenes by using FR and RFR methods with PSNR, FPSNR, and SSIM, which is presented in Table. 2. Not surprisingly, images rendered by RFR are lower scored than FR with the same mapping strategy, since less computation power is arranged in the upper and lower visual fields. While SSIM indicates the considerable difference between images rendered with FR and RFR methods due to the increased peripheral noise, the user study in Sec 5.3 proves images are perceptually equivalent, implying that the spatial noise is less noticed in the periphery.

Table 2: The evaluated PSNR, FPSNR, and SSIM for the Bistro scene rendered with FR and RFR methods in Fig. 5.

Metrics	FR_{rect}	FR_{polar}	RM-RFR	PM-RFR
PSNR	26.2058	25.8521	26.199	25.723
FPSNR	27.5596	27.4992	27.2992	27.3505
SSIM	0.59356	0.59068	0.5877	0.58022

5 User Studies

As shown in Fig. 4, the RFR method leverages regression of computation load from the fovea to the periphery by tuning α. Although a configuration with $\alpha = 1$ shows a regression similar to the fMRI-driven CMF model, both hardware (e.g., the display performance of HMDs) and software (e.g., artifacts in the periphery) may cause deviations in the perceived image quality from the target regression. Therefore, a practical α needs to be evaluated for VR application scenarios. Meanwhile, existing HMDs are still struggling to support a satisfactory resolution in the user's central vision, and foveated rendering often results in redundant computation in the central visual field [25, 41]. Hence, the RFR methods also require the selection

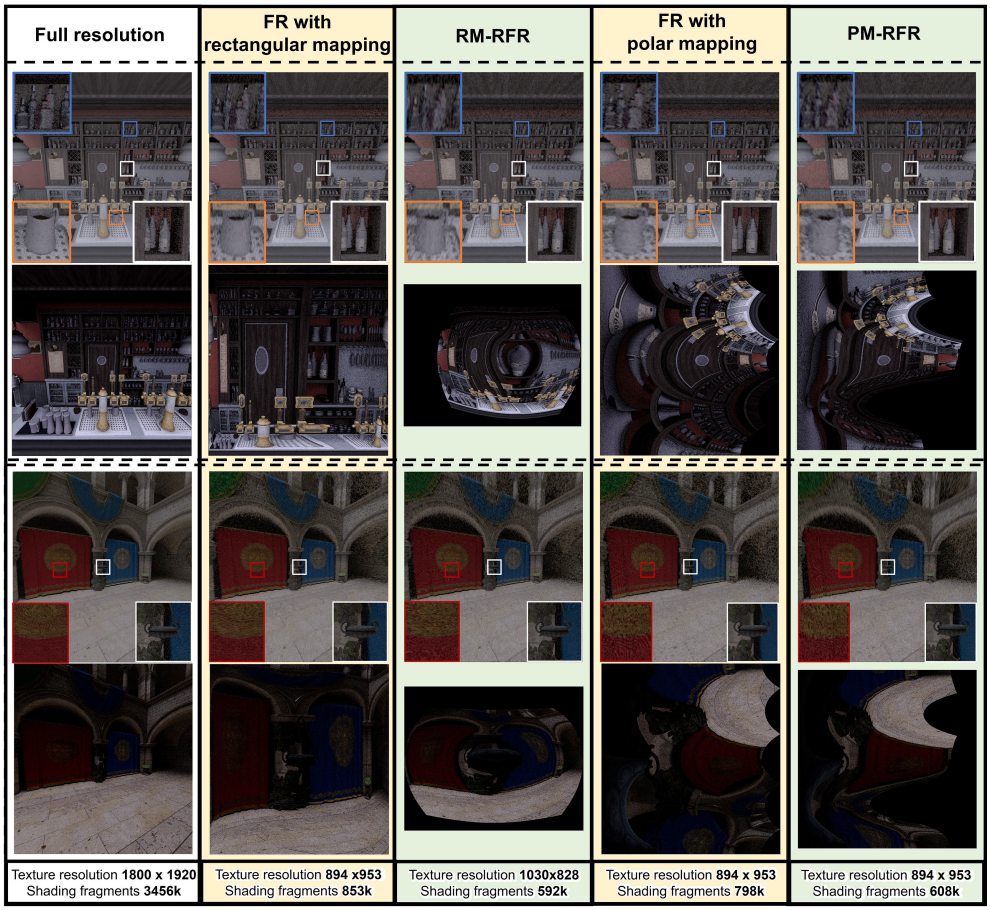

Figure 5: The comparison between RFR methods and typical foveated rendering (FR). We integrate the proposed radial asymmetric computation model into FR with rectangular and polar mapping, leading to the RM-RFR and PM-RFR methods. Overall, all the foveated methods display a sharper image in the central vision (white box) than the full resolution case and show degraded quality in the periphery. Zoom-in images in the upper-vertical (blue box) and lower-vertical (orange box) visual fields given by RM-RFR and PM-RFR methods are in worse image quality compared with the FR methods, and the images in the red box show similar quality between FR and RFR with the same mapping methods, hence, demonstrating the asymmetry of rendering precision in RFR. The downsized texture in the fragment shader of each method is placed beneath the rendered images. All the foveated methods have a roughly equal resolution for the warped shader, but the shaded fragments are considerably reduced for validating the radial asymmetry of computation load in RFR methods. As a result, an average reduction of 27.2% for shading is achieved by using RFR compared with typical FR.

of P_0 to provide the desired image quality. To address this, we conducted a pilot experiment to measure the essential coefficients P_0 and α for validating RM-RFR and PM-RFR in VR HMDs. Then, RFR with the derived configuration was included in the final user study that evaluated the perceived image quality using the proposed methods and typical FR methods.

5.1 Apparatus

We used a Meta Quest Pro (OS ver. 56.0) for displaying all the images in the user study. The Quest Pro headset offers a resolution of 1800×1920 per eye and supports a binocular FOV of H106°$\times$V96°. It has a high refresh rate of 90 Hz and provides an eye-tracking feature with integrated cameras. The PC used for running the rendering graphics is equipped with an Nvidia GTX 3080 graphics card and an Intel Core i9-12900K CPU. We developed the rendering pipeline based on the Scriptable Rendering Pipeline (SRP) feature, combined with the ray-tracing shader in Unity (ver.2021.3.29f1). The experimental Graphics API Direct3D12 is set in Unity, and we selected multi-pass for enabling the customized rendering pipeline. We utilized SDKs, such as Core RP Library (ver. 12.1.12) and Oculus Integration (ver. 56.0), to validate the ray-tracing feature and to enable eye-tracking, respectively. Although all the experiment scenes run dynamically in Unity, the Meta Quest Pro does not support the graphics API Direct3D12. As a consequence, the scenes presented to participants were pre-captured in Unity and displayed as static images in VR. Participants used a keyboard to select their preferred image. After completing the experiments, all participants were asked to fill out a questionnaire using a smartphone.

5.2 Pilot Experiment

5.2.1 Participants

For the pilot experiment, 12 participants with normal or corrected-to-normal vision (6 males and 6 females, aged 21 to 28 years (M = 23.75, SD = 2.05)) from a local university were recruited. Each participant received an amount equivalent to 11 US Dollars for their participation. All participants provided informed consent, and the study was approved by the University's Ethics Committee.

5.2.2 Procedure

To begin the study, participants were asked to wear the headsets and adjust them for the best vision after sitting comfortably. The

visual stimulus consists of a full-resolution image and foveated images rendered by RM-RFR and PM-RFR methods with α values ranging from 1 to 1.8 in steps of 0.2 and P_0 values ranging from 10 to 50 in steps of 10. The Bistro scene shown in Fig. 4 was used to generate the stimulus. Each foveated image is compared with the full-resolution image six times. The pair order is randomized, ensuring that each foveated image appears both before and after the full-resolution image three times. The study utilized a modified two-alternative forced-choice (2AFC) task with an "unsure" option. Participants were asked to choose the image with better quality using the left and right arrows on the keyboard or click the up arrow if they considered both images to be of equal quality. Before displaying the images, a red cross appears at the central vision for 0.75s to indicate the foveal area, then it disappears. The compared images are both shown for 2.5s, with an interval of 0.75s. The red cross is displayed on the interval screen to remind participants of the central vision. The participants were asked to look at the indicated area by the red cross, and choose the better image or make an "unsure" option after each image pair was displayed. The next image pair was shown after the participants clicked the keyboard. The study includes 312 trials in total, including 12 control conditions where both images are in full resolution, and lasts around one hour. To prevent fatigue, the RM-RFR and PM-RFR tests were separated with a break time, and participants were also told to have a break as long as he or she felt tired. Half of the participants took the RM-RFR test before the PM-RFR test, while the order was reversed for the other half. During the study, participants' gaze positions were tracked using the eye-tracking feature of the Meta Quest Pro.

5.2.3 Results and Analysis

The mean proportion of the participants choosing the full-resolution image, in addition to the "unsure" option, is shown in Fig. 6. Each bar represents the proportion of choices for full-resolution (solid colors) and the "unsure" option (light colors). Overall, the proportion of choosing the full-resolution image decreases with the P_0 increasing to 50, indicating that the perceived image quality is primarily influenced by the global rendering precision. Moreover, the influence of α on the choice for full-resolution progressively increases as P_0 ranges from 10 to 50. Notably, the PM-RFR method outperforms with fewer choices for full-resolution during the test (PM-RFR: 53%, RM-RFR: 62%). We determined thresholds of P_0 to realize an unnoticeable regression in computation load with the RM-RFR and PM-RFR methods. The threshold fitting is conducted by adding half of the selection on the "unsure" option to the full-resolution option. Then, this data was used to fit the probability curves to P_0 using a cumulative normal distribution function for each α condition, utilizing the Palamedes Toolbox [28]. We identified the threshold 0.5 as the target P_0 for rendering an image perceptually equivalent to the full-resolution case. The derived P_0 for both methods across α values ranging from 1 to 1.8 is listed in Table. 3. Notably, the minimal P_0 values for both RM-RFR and PM-RFR methods are found at different α values. We opted for configurations with identical $P_0 = 33$ but varied α values of 1.2 and 1.6 for RM-RFR and PM-RFR methods, respectively, for the final user study.

Table 3: The evaluated P_0 for RM-RFR and PM-RFR with α from 1 to 1.8 to display perceptually equal images as full-resolution rendering.

α	1	1.2	1.4	1.6	1.8
P_0 for RM-RFR	51.0	34.5	43.0	83.3	151.5
R square	0.9761	0.9967	0.9452	0.8771	0.8609
P_0 for PM-RFR	55.6	38.2	34.2	33.1	82.0
R square	0.9672	0.8932	0.9529	0.9830	0.9491

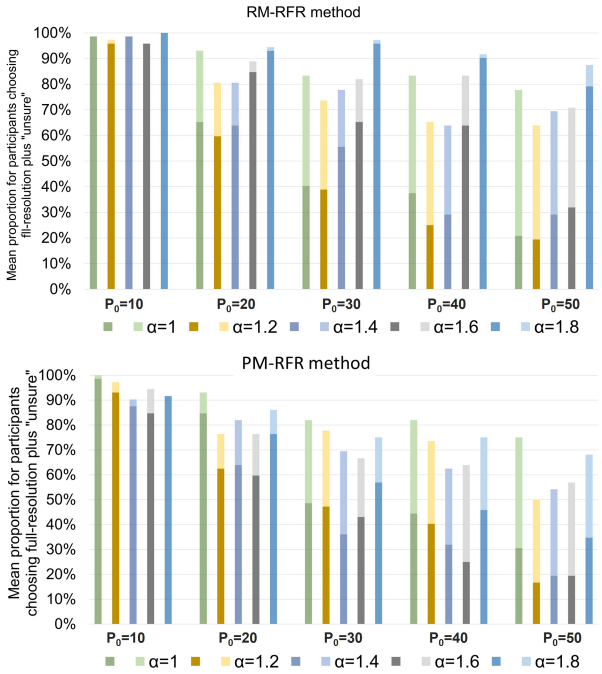

Figure 6: The mean proportion of participants choosing the full-resolution image (solid colors) plus the "unsure" option (light colors) in the pilot experiment.

5.3 Final User Study

5.3.1 Participants

24 participants with normal or corrected-to-normal vision (12 males and 12 females, aged 21 to 28 years (M = 23.58, SD = 1.56)) were recruited for this study. Each participant received an amount equivalent to 11 US Dollars for their participation. All participants provided informed consent, and the study was approved by the University's Ethics Committee.

5.3.2 Procedure

We employed a procedure similar to that used in the pilot experiment. Images for the stimuli were rendered based on full-resolution, polar FR, rectangular FR, RM-RFR, and PM-RFR using the Bistro and Sponza scenes, as shown in Fig. 5. Beyond the foveal field, we additionally selected two gaze areas in the upper-vertical and lower-vertical visual fields for Bistro and two gaze areas in the left-horizontal and right-horizontal visual fields for Sponza. Standard 2AFC tasks were included in the final user study. The images rendered by all five methods were randomly paired, and the participants were asked to choose the one with better quality. Each trial was repeated four times, resulting in 240 trials in total. The study lasted around 40 minutes. To prevent fatigue, we scheduled a break after every 60 trials.

5.3.3 Results and Analysis

Fig. 7 presents the results of the final user study. Overall, we did not observe a significant difference between images rendered by the full-resolution method and those rendered by other methods, indicating that all the foveated methods produce images that are perceptually equivalent to full-resolution rendering but require much less computation load, as shown in Fig. 5. Furthermore, we did not find any significant difference between the two typical FR methods and our RFR methods. Therefore, utilizing a radial asymmetrical distribution for rendering precision does not cause any degradation

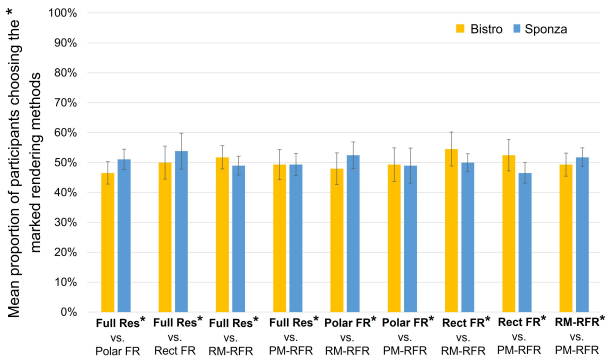

Figure 7: The mean proportion of participants choosing the $*$ marked rendering methods in the final user study for Bistro and Sponza scenes. The bar on each histogram represents $\pm$95%CI. The "Full Res" indicates the full-resolution method. No significant difference was found between groups using typical FR methods and the proposed RFR methods, or between groups using full-resolution and the RFR methods based on the results of the t-test. This proves the reduced rendering precision near the upper and lower meridian by RFR does not cause perceivable degradation in image quality.

Table 4: The FPS for rendering the Bistro scene using typical FR methods and the proposed RFR methods is presented. The hardware here is the same as that used in the apparatus for conducting the user study. The proposed RFR methods achieve an average FPS that is 17.45% higher than that of typical FR methods.

Methods	Full Res	FR_{rect}	FR_{polar}	RM-RFR	PM-RFR
Shading fragments	3,456,000	852,876	797,876	592,104	608,181
FPS	1.66	5.9	5.14	6.71	6.23

in perceivable image quality. As a result, the proposed RFR is able to push the performance of existing FR techniques one step further.

6 Discussion and Future Works

Rendering acceleration by RFR. Fig. 5 depicts the texture in the fragment shader stage of FR and RFR methods. While the texture resolution is nearly identical, the total number of shading fragments differs between the FR and RFR rendering pipelines. Table 4 lists the exact numbers of shading fragments for different methods. Using the Bistro scene from Fig.5, FPS values were tested at a full resolution of 1800×1920 pixels. All the foveated techniques here are realized by our customized rendering pipeline with a ray tracing sampling rate of 30. As with Fig. 5, we set $P_0 = 33$ to ensure all foveated methods have a consistent texture size. The listed FPS values were averaged over 300 frames, and each method was measured three times. On the whole, conventional FR methods accelerated the rendering to approximately 3.3x. In contrast, RFR methods achieved an even higher acceleration rate, reaching 3.9x. By comparing FR and RFR with the same mapping strategy, we observed a higher FPS by 13.7% when using RM-RFR and 21.2% when using PM-RFR. The asymmetric allocation of computation load when using RFR can be easily transferred to any foveated graphics technique, we believe that the proposed model will benefit the development of FR in many aspects.

Anti-aliasing for RFR. Unlike typical FR methods that leverage the regression in a single dimension, the proposed RFR method introduced the regression in the tangential direction, which causes more artifacts in the periphery. As shown in Fig. 4, pronounced radial artifacts are visible, especially in scenarios with $\alpha = 1.8$ in the upper FOV, and notably in the gaze area when $P_0 = 10$ and $\alpha = 1$. While we experimented with temporal anti-aliasing (TAA) to counteract these artifacts, the results were marginal. Additionally, camera jitter emerged as a new challenge, especially pronounced when wearing HMDs. Increasing the computation load will attenuate the artifacts efficiently, as evidenced in the image with $P_0 = 50$. Adopting a retinotopy-based sampling probability map may reduce the artifacts considerably [31]. Further RFR research should take the anti-aliasing feature into account.

RFR-based ray tracing. We implemented the RFR method and the typical FR method with ray tracing. Although the same sampling rate is utilized for fragment shading in the full-resolution rendering, it shows worse image quality than FR methods in the central vision in Fig. 5. The advantage of FR-based ray tracing is attributed to the equivalently increased sampling rate for pixels in the central vision via the non-uniform projection between the screen space and the fragment shader space [8]. We believe the proposed RFR method will further enhance ray tracing methods by optimally allocating the sampling rate based on the asymmetric nature of HVS.

Retinotopic foveated graphics. Radial asymmetry has been confirmed in the primary cortex. This not only affects visual acuity but also impacts various low-level visual features such as contrast sensitivity and color recognition. These features have been shown to strongly correlate with CMF, and are expected to similarly exhibit the radial asymmetric characteristic. Given this, graphics techniques in VR stand to benefit from the fMRI-driven computation model, going beyond merely reducing shading costs.

User studies in static scenes. As indicated in Sec 5.1, the RFR method is evaluated in static scenes due to the hardware restriction, while the temporal artifacts in dynamic scenes must be taken into account for practical implementation [39]. Nevertheless, the results from the static scenes reveal the optimal regression of the sampling rate in FR-based ray tracing. Since amortized strategies, such as TAA, are utilized to increase the sampling rate in ray-tracing graphics but cause ghost images due to the collected temporal information, an optimization in terms of the asymmetric regression of the sampling rate may provide the optimal trade-off between image quality and ghost images.

Individual-preferred RFR. Although the size of the early cortex varies with individuals, fMRI evidence suggests similar HVA and VMA in group-level [13, 16]. However, a minor VMA in children is confirmed [14]. The fMRI-driven model in the presented paper is derived from both the fMRI data and psychophysical evidence due to the restriction of fMRI techniques. We believe the advancement of fMRI will reveal the asymmetry nature in HVS more precisely, and then facilitates the development of individual-preferred RFR methods.

7 Conclusion

In this paper, we proposed retinotopic foveated rendering that links advancements in neuroscience with graphics rendering techniques. We introduced the first fMRI-driven computation model for implementing foveated rendering. RFR with this radially asymmetric computation model achieves a reduction of 27.2% for fragment shading compared to typical FR methods. We conducted a pilot experiment to validate the proposed asymmetric computation model in the FR algorithm. The final user study demonstrates that images rendered by typical FR methods and the proposed RM-RFR and PM-RFR methods are perceptually equal. The proposed method also performs 17.45% faster than the typical FR methods on average. In summary, we anticipate a global optimization with the proposed RFR method in foveated graphics in the future.

ACKNOWLEDGMENTS

This work was funded by the Shanghai Pujiang Program (23PJ1406800) and the China Postdoctoral Science Foundation (2023M742260).

REFERENCES

[1] J. Abrams, A. Nizam, and M. Carrasco. Isoeccentric locations are not equivalent: The extent of the vertical meridian asymmetry. *Vision research*, 52(1):70–78, 2012.

[2] A. Barbot, S. Xue, and M. Carrasco. Asymmetries in visual acuity around the visual field. *Journal of Vision*, 21(1):2–2, 2021.

[3] N. C. Benson, K. W. Jamison, M. J. Arcaro, A. T. Vu, M. F. Glasser, T. S. Coalson, D. C. Van Essen, E. Yacoub, K. Ugurbil, J. Winawer, et al. The human connectome project 7 tesla retinotopy dataset: Description and population receptive field analysis. *Journal of vision*, 18(13):23–23, 2018.

[4] N. C. Benson, E. R. Kupers, A. Barbot, M. Carrasco, and J. Winawer. Cortical magnification in human visual cortex parallels task performance around the visual field. *Elife*, 10:e67685, 2021.

[5] M. Carrasco, C. P. Talgar, and E. L. Cameron. Characterizing visual performance fields: Effects of transient covert attention, spatial frequency, eccentricity, task and set size. *Spatial vision*, 15(1):61, 2001.

[6] P. Chakravarthula, Z. Zhang, O. Tursun, P. Didyk, Q. Sun, and H. Fuchs. Gaze-contingent retinal speckle suppression for perceptually-matched foveated holographic displays. *IEEE Transactions on Visualization and Computer Graphics*, 27(11):4194–4203, 2021.

[7] R. O. Duncan and G. M. Boynton. Cortical magnification within human primary visual cortex correlates with acuity thresholds. *Neuron*, 38(4):659–671, 2003.

[8] S. Friston, T. Ritschel, and A. Steed. Perceptual rasterization for head-mounted display image synthesis. *ACM Trans. Graph.*, 38(4):97–1, 2019.

[9] W. S. Geisler and J. S. Perry. Real-time foveated multiresolution system for low-bandwidth video communication. In *Human vision and electronic imaging III*, vol. 3299, pp. 294–305. SPIE, 1998.

[10] B. Guenter, M. Finch, S. Drucker, D. Tan, and J. Snyder. Foveated 3d graphics. *ACM transactions on Graphics (tOG)*, 31(6):1–10, 2012.

[11] Y. He, Y. Gu, and K. Fatahalian. Extending the graphics pipeline with adaptive, multi-rate shading. *ACM Transactions on Graphics (TOG)*, 33(4):1–12, 2014.

[12] A. Helder. An introduction to the log-polar mapping. In *II Workshop on Cibernetic Vision, December, 1996*, 1996.

[13] M. M. Himmelberg, J. W. Kurzawski, N. C. Benson, D. G. Pelli, M. Carrasco, and J. Winawer. Cross-dataset reproducibility of human retinotopic maps. *Neuroimage*, 244:118609, 2021.

[14] M. M. Himmelberg, E. Tünçok, J. Gomez, K. Grill-Spector, M. Carrasco, and J. Winawer. Comparing retinotopic maps of children and adults reveals a late-stage change in how v1 samples the visual field. *Nature communications*, 14(1):1561, 2023.

[15] M. M. Himmelberg, J. Winawer, and M. Carrasco. Stimulus-dependent contrast sensitivity asymmetries around the visual field. *Journal of vision*, 20(9):18–18, 2020.

[16] M. M. Himmelberg, J. Winawer, and M. Carrasco. Linking individual differences in human primary visual cortex to contrast sensitivity around the visual field. *Nature communications*, 13(1):3309, 2022.

[17] M. M. Himmelberg, J. Winawer, and M. Carrasco. Polar angle asymmetries in visual perception and neural architecture. *Trends in Neurosciences*, 2023.

[18] J. C. Horton and W. F. Hoyt. The representation of the visual field in human striate cortex: a revision of the classic holmes map. *Archives of ophthalmology*, 109(6):816–824, 1991.

[19] L. Huber, H. Desmond, C. J. Wiggins, K. Uludağ, S. Kashyap, D. C. Jangraw, P. A. Bandettini, B. A. Poser, and D. Ivanov. Ultra-high resolution blood volume fmri and bold fmri in humans at 9.4 t: capabilities and challenges. *Neuroimage*, 178:769–779, 2018.

[20] A. S. Kaplanyan, A. Sochenov, T. Leimkühler, M. Okunev, T. Goodall, and G. Rufo. Deepfovea: Neural reconstruction for foveated rendering and video compression using learned statistics of natural videos. *ACM Transactions on Graphics (TOG)*, 38(6):1–13, 2019.

[21] M. Koskela, A. Lotvonen, M. Mäkitalo, P. Kivi, T. Viitanen, and P. Jääskeläinen. Foveated real-time path tracing in visual-polar space. In *Proceedings of 30th Eurographics Symposium on Rendering*. The Eurographics Association, 2019.

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

[23] E. R. Kupers, N. C. Benson, M. Carrasco, and J. Winawer. Asymmetries around the visual field: From retina to cortex to behavior. *PLoS computational biology*, 18(1):e1009771, 2022.

[24] D. M. Levi, S. A. Klein, and A. Aitsebaomo. Vernier acuity, crowding and cortical magnification. *Vision research*, 25(7):963–977, 1985.

[25] X. Meng, R. Du, M. Zwicker, and A. Varshney. Kernel foveated rendering. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 1(1):1–20, 2018.

[26] A. Patney, J. Kim, M. Salvi, A. Kaplanyan, C. Wyman, N. Benty, A. Lefohn, and D. Luebke. Perceptually-based foveated virtual reality. In *ACM SIGGRAPH 2016 emerging technologies*, pp. 1–2. 2016.

[27] A. Patney, M. Salvi, J. Kim, A. Kaplanyan, C. Wyman, N. Benty, D. Luebke, and A. Lefohn. Towards foveated rendering for gaze-tracked virtual reality. *ACM Transactions on Graphics (TOG)*, 35(6):1–12, 2016.

[28] N. Prins and F. A. Kingdom. Applying the model-comparison approach to test specific research hypotheses in psychophysical research using the palamedes toolbox. *Frontiers in psychology*, 9:1250, 2018.

[29] J. Rovamo, V. Virsu, and R. Näsänen. Cortical magnification factor predicts the photopic contrast sensitivity of peripheral vision. *Nature*, 271(5640):54–56, 1978.

[30] M. F. Silva, J. W. Brascamp, S. Ferreira, M. Castelo-Branco, S. O. Dumoulin, and B. M. Harvey. Radial asymmetries in population receptive field size and cortical magnification factor in early visual cortex. *NeuroImage*, 167:41–52, 2018.

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

[32] H. Strasburger, I. Rentschler, and M. Jüttner. Peripheral vision and pattern recognition: A review. *Journal of vision*, 11(5):13–13, 2011.

[33] Q. Sun, F.-C. Huang, J. Kim, L.-Y. Wei, D. Luebke, and A. Kaufman. Perceptually-guided foveation for light field displays. *ACM Transactions on Graphics (TOG)*, 36(6):1–13, 2017.

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

[35] K. Uğurbil, J. Xu, E. J. Auerbach, S. Moeller, A. T. Vu, J. M. Duarte-Carvajalino, C. Lenglet, X. Wu, S. Schmitter, P. F. Van de Moortele, et al. Pushing spatial and temporal resolution for functional and diffusion mri in the human connectome project. *Neuroimage*, 80:80–104, 2013.

[36] V. Virsu, R. Näsänen, and K. Osmoviita. Cortical magnification and peripheral vision. *JOSA A*, 4(8):1568–1578, 1987.

[37] V. Virsu and J. Rovamo. Visual resolution, contrast sensitivity, and the cortical magnification factor. *Experimental brain research*, 37:475–494, 1979.

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

[39] M. Weier, T. Roth, E. Kruijff, A. Hinkenjann, A. Pérard-Gayot, P. Slusallek, and Y. Li. Foveated real-time ray tracing for head-mounted displays. In *Computer Graphics Forum*, vol. 35, pp. 289–298. Wiley Online Library, 2016.

[40] G. Westheimer. Visual acuity. *Annual Review of Psychology*, 16(1):359–380, 1965.

[41] J. Ye, A. Xie, S. Jabbireddy, Y. Li, X. Yang, and X. Meng. Rectangular mapping-based foveated rendering. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 756–764. IEEE, 2022.