

Gaze-Contingent Ocular Parallax Rendering for Virtual Reality

ROBERT KONRAD, Stanford University

ANASTASIOS ANGELOPOULOS, Stanford University

GORDON WETZSTEIN, Stanford University




Fig. 1. The centers of rotation and projection of the eyes are not the same. As a consequence, small amounts of parallax are created in the retinal image as we fixate on different objects in the scene. The nodal points of the eye, representing the centers of projection, are shown as small blue circles on the left along with a ray diagram illustrating the optical mechanism of ocular parallax. Simulated retinal images that include the falloff of acuity in the periphery of the visual field are shown on the right. As a user fixates on the candle in the center of the scene (center, red circle indicates fixation point), the bottle is partly occluded by the candle. As their gaze moves to the left, ocular parallax reveals the bottle behind the candle in the center (right). Ocular parallax is a gaze-contingent effect exhibiting the strongest effect size in near to mid peripheral vision, where visual acuity is lower than in the fovea. In this paper, we introduce ocular parallax rendering for eye-tracking-enabled virtual reality displays, and study the complex interplay between micro parallax, occlusion, visual acuity, disparity distortions, and other perceptual aspects of this technology in simulation and with a series of user experiments.

Immersive computer graphics systems strive to generate perceptually realistic user experiences. Current-generation virtual reality (VR) displays are successful in accurately rendering many perceptually important effects, including perspective, disparity, motion parallax, and other depth cues. In this paper we introduce ocular parallax rendering, a technology that accurately renders small amounts of gaze-contingent parallax capable of improving depth perception and realism in VR. Ocular parallax describes the small amounts of depth-dependent image shifts on the retina that are created as the eye rotates. The effect occurs because the centers of rotation and projection of the eye are not the same. We study the perceptual implications of ocular parallax rendering by designing and conducting a series of user experiments. Specifically, we estimate perceptual detection and discrimination thresholds for this effect and demonstrate that it is clearly visible in most VR applications. Additionally, we show that ocular parallax rendering provides an effective ordinal depth cue and it improves the impression of realistic depth in VR.

1 INTRODUCTION

Immersive computer graphics systems, such as virtual reality (VR) displays, aim at synthesizing a perceptually realistic user experience. To achieve this goal, several components are required: interactive, photorealistic rendering; a high-resolution, low-persistence, stereoscopic display; and low-latency head tracking. Modern VR systems provide all of these capabilities and create experiences that support many, but not all, of the monocular and binocular depth cues of the human visual system, including occlusions, shading, binocular

disparity, and motion parallax. The support of focus cues (e.g. accommodation and retinal blur) has also received attention in research and industry over the last few years. In this paper, we study a depth cue of human vision that has not been discussed in the context of virtual reality and that may help further improve depth perception and perceptual realism: ocular parallax.

The centers of rotation and projection in the human eye are not the same. Therefore, changes in gaze direction create small amounts of depth-dependent image shifts on our retina—an effect known as ocular parallax. This depth cue was first described by Brewster [1845] and it has been demonstrated to produce parallax well within the range of human visual acuity [Bingham 1993; Hadani et al. 1980; Mapp and Ono 1986]. Interestingly, species as diverse as the chameleon and the sandlance critically rely on this depth cue to judge distance [Land 1995; Pettigrew et al. 1999].

To render ocular parallax into a VR/AR experience, eye tracking is required. Conveniently, many emerging wearable display systems already have eye tracking integrated, either to support foveated rendering [Guenter et al. 2012; Patney et al. 2016], accurate registration of physical and digital images in AR, or other gaze-contingent display modes. With eye tracking available, there is no additional computational cost to integrate ocular parallax. The perspective of the rendered image simply changes depending on the gaze direction. However, the magnitude of depth-dependent motion induced by ocular parallax rendering increases in the periphery of the visual field, where visual acuity is lower than in the fovea. Moreover, the resolution offered by current-generation VR displays is well below

Authors' addresses: Robert Konrad, Stanford University, rkkonrad@stanford.edu; Anastasios Angelopoulos, Stanford University, nikolasa@stanford.edu; Gordon Wetzstein, Stanford University, gordon.wetzstein@stanford.edu.

the visual acuity of human vision and it is not clear if the subtle ocular parallax effect is perceivable in VR at all.

To further our understanding of ocular parallax and its perceptual effects in VR, we thoroughly analyze the tradeoffs between perceived parallax, visual acuity, and disparity for near-eye displays. We build a prototype gaze-tracked VR display, conduct a series of user experiments that quantify effect sizes of ocular parallax rendering, and measure its impact on depth perception and the user experience in general. We find that detection thresholds for ocular parallax rendering are almost an order of magnitude lower than the visual acuity at the same extrafoveal locus, which verifies that our sensitivity to small amounts of differential motion is well below the acuity limit, especially in the periphery of the visual field [McKee and Nakayama 1984]. We also show that the relative ocular parallax of objects with respect to a background target can be discriminated accurately even for relatively small object distances that fall well within the depth ranges of most virtual environments. Furthermore, we show that ocular parallax rendering provides an effective ordinal depth cue, helping users better distinguish the relative depth ordering of a scene, but that it does not benefit absolute distance estimates to objects. Finally, we show that ocular parallax rendering improves the impression of realistic depth in a 3D scene. For zero additional computational cost, ocular parallax rendering has the potential to improve both depth perception and perceptual realism of eye-tracked AR/VR systems.

Specifically, we make the following contributions:

- We introduce gaze-contingent ocular parallax rendering for VR display systems.
- We design and conduct user experiments to quantify detection and discrimination thresholds of ocular parallax rendering.
- We design and conduct user experiments to quantify the effectiveness of ocular parallax rendering as both an ordinal and absolute depth cue.
- We conduct a user experiment that demonstrates improved perceptual realism using ocular parallax.

2 RELATED WORK

Depth Cues. Human depth perception relies on a variety of cues [Howard and Rogers 2002; Palmer 1999]. Many of these cues are pictorial and can be synthesized using photorealistic rendering techniques, including occlusions, perspective foreshortening, texture and shading gradients, as well as relative and familiar object size. Unlike conventional 2D displays, head-mounted displays (HMDs) use stereoscopic displays and head tracking and can thus support two additional visual depth cues, disparity and motion parallax, as well as one oculomotor cue, vergence. Emerging near-eye displays also support visual focus cues like retinal blur and chromatic aberrations, which in turn drive accommodation, another oculomotor cue (see discussion below). All of these cues are important for human depth perception to varying degrees depending on the fixation distance [Cutting and Vishton 1995]. Studying visual cues, such as disparity [Didyk et al. 2011] or motion parallax [Kellnhofer et al. 2016], and their impact on computational display applications has been an integral part of graphics research. In this work, we explore

ocular parallax as another visual cue that may improve the user experience in immersive computer graphics applications.

Ocular Parallax. Ocular parallax describes the change in perspective as the eye rotates, primarily due to the user fixating on different parts of the scene. This visual cue is well known [Brewster 1845] and has a measurable effect on depth perception [Bingham 1993; Kudo and Ohnishi 1998; Kudo et al. 1999; Mapp and Ono 1986]. Similar to other monocular visual cues, such as retinal blur and chromatic aberration, the change of the retinal image caused by ocular parallax may be small. Nonetheless, supporting all of these cues with an HMD can improve visual comfort [Hoffman et al. 2008], perceived realism, and the user experience as a whole.

Kudo and Ohnishi [2000] discuss gaze-contingent optical distortions in head-mounted displays (HMDs) and attribute them in part to ocular parallax. This effect is commonly known as “pupil swim”. However, they did not propose ocular parallax rendering for HMDs or study its perceptual effects with HMDs. Building on emerging head-mounted displays with eye-tracking technology, to our knowledge we are the first to propose ocular parallax as a gaze-contingent rendering mode for VR/AR and evaluate its perceptual implications with a series of user experiments.

Gaze-contingent and Computational Displays. Eye tracking enables gaze-contingent rendering techniques that adapt effects like magnification, stylization, or geometric level-of-detail to the user’s viewpoint [Duchowski et al. 2004]. Gaze-contingent rendering is becoming an integral part of modern wearable display systems, enabling techniques such as foveated rendering [Guenther et al. 2012; Patney et al. 2016] and gaze-contingent varifocal [Dunn et al. 2017; Johnson et al. 2016; Konrad et al. 2015; Liu et al. 2008; Padmanaban et al. 2017] or multifocal [Akeley et al. 2004; Mercier et al. 2017; Rolland et al. 2000] image displays. Although rendering accommodation-dependent effects, such as chromatic aberrations [Cholewiak et al. 2017] and blur at depth edges [Marshall et al. 1996; Zannoli et al. 2014], have not been directly evaluated with eye-tracked displays, these techniques could be optimized by tracking the user’s gaze or accommodation. Ocular parallax rendering is complimentary to these techniques and could be integrated, without computational overhead, into conventional HMDs with eye tracking and optionally combined with any other gaze-contingent rendering algorithms.

Other than the proposed method, the only techniques described in the literature that could inherently provide ocular parallax cues are near-eye light field [Hua and Javidi 2014; Huang et al. 2015; Lanman and Luebke 2013] and holographic near eye displays [Maimone et al. 2017]. However, none of these works mention or study the effect of ocular parallax. Mercier et al. [2017] proposed a multifocal plane display that shifts the decomposed layers according to the tracked pupil position, but because ocular parallax was not explicitly rendered into the displayed images, it was omitted as a visual cue.

3 OCULAR PARALLAX

In this section, we discuss schematic models of the human eye with a specific focus on how they model the centers of projection and rotation. Moreover, we discuss and analyze the trade-off between




Fig. 2. Illustration of a schematic eye, including the front and rear nodal points N , N' , the center of rotation C , and the anterior vertex of the cornea V . The nodal points are two parameters of a thick lens model that refracts light rays as depicted. The exact locations of these points depend on the schematic eye model used.

	Gull. 1 Relaxed	Gull. 1 Acc.	Gull.-Ems. Relaxed	Gull.-Ems. Acc.	Ems.
VN	7.078	6.533	7.062	6.562	5.556
VN'	7.331	6.847	7.363	6.909	5.556

Table 1. Overview of parameters of popular schematic eye models: Gullstrand number 1 (Gull. 1), Gullstrand-Emsley (Gull.-Ems.), and Emsley reduced (Ems.). The distances of front and rear nodal points N , N' are listed in mm with respect to the anterior vertex of the cornea V for the relaxed and accommodated (acc.) state.

ocular parallax and image sharpness, accounting for the decrease in peripheral visual acuity as well as retinal blur due to accommodation.

Eye Models. Complex optical systems with multiple refractive surfaces can be reduced to six cardinal points that fully define the Gaussian imaging and magnification properties of the system. These cardinal points include the front and rear focal points, the front and rear principle points, and the front and rear nodal points, N and N' . For the purpose of modeling ocular parallax, we only require the front nodal point N , which is the center of projection of the eye, as well as the center of rotation C (see Figure 2). Several schematic eye models have been proposed in the literature, each listing slightly different values for the cardinal points [Atchison 2017]. Some of the most popular models include the Gullstrand number 1, the Gullstrand-Emsley, and the Emsley reduced eyes, which are models of decreasing complexity. We outline the location of their respective nodal points and centers of rotation in Table 1. The center of rotation of the eye was measured to be 14.7536 mm from the cornea, on average, for emmetropic subjects [Fry and Hill 1962].

Although the specific locations of the cardinal points are slightly different for each eye model, the distance between center of rotation C and center of projection N is 7–8 mm in all cases. Note that nodal points for the Gullstrand number 1 and Gullstrand-Emsley models are accommodation dependent, i.e. the nodal points move slightly toward the cornea when the eye accommodates to close distances. However, in most current-generation VR/AR systems, the focal plane of the display is fixed. For example, the focal planes of the




Fig. 3. The amount of ocular parallax, measured in degrees of visual angle, increases with increasing eccentricity and relative distance between objects. Relative object distances of 3 diopters and greater are above the minimum angle of resolution (MAR, red dashed line) and may therefore be detectable by a human observer. However, the amount of parallax on the foveola may be too small to be detected for static stimuli (left). The resolution available by modern VR displays, like the HTC Vive Pro (dashed blue lines), is slightly lower than the MAR of human vision in the foveola, but it exceeds the MAR for 2–3° of visual angle—precisely where we expect to see the effects of ocular parallax.

Oculus Rift and Microsoft Hololens are approximately 1.3 m¹ and 2 m² in front of the user, respectively. Since users will accommodate at that fixed focal distance [Padmanaban et al. 2017], we use the relaxed setting of the popular Gullstrand-Emsley eye model for all experiments in this paper, i.e. $NC = 7.6916$ mm.

Parallax and Acuity. We can model the amount of perceived ocular parallax expected in various viewing conditions. Similar to the illustration in Figure 1, we simulate two points that are directly in front of the eye at some relative distance to one another. As the eye rotates, the retinal images of these points will be perceived at an increasing eccentricity, or distance from the fovea, measured in degrees of visual angle. The larger the eccentricity, the larger the parallax, or relative distance, of the points on the retina. However, the density of photoreceptors, in particular the cones, decreases rapidly with increasing eccentricity. Thus, while one would expect a larger amount of parallax in the periphery, the threshold to perceive it there is higher due to the falloff in visual acuity.

Figure 3 illustrates the tradeoff between ocular parallax and visual acuity. Here, the minimum angle of resolution (MAR) represents a measure for visual acuity and approximates the resolution of human vision for a *static* target. To model the falloff of acuity in the periphery of the visual field, we use the linear model proposed by Guenter et al. [2012]: $\omega = m e + \omega_0$. Here, ω is the minimum angle of resolution (dashed red lines in Fig. 3), e is the eccentricity in degrees, m is the slope modeling the falloff of acuity, and ω_0 is the MAR at the fovea. We set $\omega_0 = 1/60$ to model 20/20 vision and $m = 0.022$ as proposed by Guenter et al. [2012]. As seen in Figure 3, a relative distance of 3 D (diopters) should theoretically be detectable by a human observer for eccentricity angles smaller or equal to 40°.

The left part of Figure 3 also shows a semi-logarithmic plot zooming into the foveal region. We see that the magnitude of ocular parallax expected in the foveola, i.e. $e < 1^\circ$, may not be sufficient to be perceivable. Yet for eccentricities larger than 1°, relative object distances of 2–3 D may make ocular parallax a useful depth cue.

¹Oculus Best Practices Guide

²<https://docs.microsoft.com/en-us/windows/mixed-reality/comfort>

Projection Matrix. Vertices in eye space are transformed into clip space using the projection matrix. A perspective projection in stereo rendering is usually represented as an asymmetric off-axis view frustum defined by a near and far clipping plane, z_{near} and

z_{far} , as well as the left (l), right (r), top (t) and bottom (b) boundary values on the near clipping plane [Shirley et al. 2009]. Using a right-handed coordinate system, the corresponding projection matrix has the general structure outlined in Equation 5. For clarity, we only show the projection transform of the right eye, but a matrix of similar form is applied for the left eye:

$$\mathbf{v}_{L/R}^{(clip)} = \mathbf{P}_{L/R} \cdot \mathbf{v}_{L/R}^{(eye)}, \quad (4)$$

$$\mathbf{P}_R = \begin{bmatrix} \frac{2 \cdot z_{near}}{r_R - l_R} & 0 & \frac{r_R + l_R}{r_R - l_R} & 0 \\ 0 & \frac{2 \cdot z_{near}}{t_R - b_R} & \frac{t_R + b_R}{t_R - b_R} & 0 \\ 0 & 0 & \frac{-(z_{far} + z_{near})}{z_{far} - z_{near}} & \frac{-2 \cdot z_{far} \cdot z_{near}}{z_{far} - z_{near}} \\ 0 & 0 & -1 & 0 \end{bmatrix}. \quad (5)$$

The frustum boundary values are determined from parameters of the physical setup, including the field of view of the display, the distance to the virtual image³ d and the position of the front nodal points $N_{L/R}$ defined with respect to the corresponding center of rotation C . Assuming that we know the fields of view $\alpha^{\{l,r,t,b\}}$ defining the asymmetric frustum of the conventional stereo rendering mode (see Fig. 4), we can compute the asymmetric view frustum of ocular parallax rendering as

$$\{l_R, r_R\} = \frac{z_{near} + N_R^{(z)}}{d + N_R^{(z)}} \left(d \cdot \tan(\alpha_R^{\{l,r\}}) + N_R^{(x)} \right), \quad (6)$$

$$\{t_R, b_R\} = \frac{z_{near} + N_R^{(z)}}{d + N_R^{(z)}} \left(d \cdot \tan(\alpha_R^{\{t,b\}}) + N_R^{(y)} \right). \quad (7)$$

The projection matrix is updated on a per-frame basis using the tracked nodal points. Applying the above modifications to the view and projection transforms renders perceptually accurate ocular parallax using slight modifications of the graphics pipeline, under the assumption that there are no optical distortions. Currently, we use the manufacturer-supplied optical distortion correction to account for optical aberrations of the HMD lenses across the visual field.

4.2 Perceptual Effects of Ocular Parallax Rendering

Ocular parallax rendering is expected to have implications on several depth cues. We briefly discuss these here to motivate the user experiments that specifically evaluate each of these effects in the following sections.

Micro Parallax. Parallax describes both magnitude and direction of the retinal velocity of objects at different depths while either the scene is in motion or the user’s head moves. We expect a similar effect to occur, albeit at a significantly smaller magnitude, due to changes in gaze direction when ocular parallax rendering is enabled. Such changes in gaze direction rotate the nodal point about the center of rotation, inducing tiny amounts of depth-dependent “micro parallax” into the retinal image. The relative magnitudes of the velocities of objects at different depths could provide an ordinal depth cue, helping users better understand the relative depth ordering of a scene. This has, for example, been shown to be the

case for head motion parallax [Yonas et al. 1987], but it is not clear whether ocular parallax is an effective ordinal depth cue as well. Furthermore, the absolute magnitude of retinal velocity induced by ocular parallax could also provide an absolute, metric depth cue, although it is not clear whether the small velocity magnitudes created by micro parallax are sufficient for the user to robustly estimate absolute distance. While micro parallax has an effect on all user experiments conducted in the following, we specifically quantify its perceptual effect size in the discrimination threshold experiment in Section 5.2.

Gaze-contingent Occlusion. Micro parallax near occlusion boundaries is particularly interesting because there we observe gaze-contingent occlusion (see Fig. 1). When objects at different depths overlap, the observed parallax due to eye rotations causes the accretion-deletion of only the farther object’s texture. While occlusion can only provide ordinal depth information, it is considered one of the strongest depth cues in static environments [Cutting and Vishton 1995]. Particularly relevant to gaze-contingent occlusion is the fact that time-varying accretion-deletion of texture due to head-motion-induced parallax has been shown to be an effective ordinal depth cue [Yonas et al. 1987]. Yet, it is unknown whether the same is true for the small amounts of accretion-deletion of texture observed with ocular parallax rendering. We evaluate the perceptual effect size of gaze-induced occlusion by estimating the detection thresholds for ocular parallax in Section 5.1 and its effectiveness as an ordinal depth cue in Section 6.1.

Disparity Distortion. Conventional stereoscopic rendering techniques assume that the centers of projection and rotation of the eyes are equivalent. This is generally an adequate assumption to correctly render the disparities for *foveal* vision. As illustrated in Figure 5, when a (blue) point is stereoscopically rendered, its position is projected onto the virtual image from the center of rotations of left and right eyes. The projected positions are indicated by the blue L/R for left and right eye, respectively. When viewed by a user who fixates on this point, the user’s front nodal points optically align through vergence with the rendered perspective and they perceive the correct disparity. With conventional rendering, the purple point is also rendered with the same perspective as the blue point, i.e. from the centers of rotation of the eyes (dashed purple line). As the user fixates on the blue point and the peripheral point is projected through the front nodal points, however, the rendered disparity and the perceived disparity are different. Therefore, we expect the disparities computed with conventional stereo rendering to be distorted in the *periphery* of the visual field.

We plot the amount of expected disparity distortion for increasing eccentricity angles for several fixation distances in Figure 5. While these distortions increase in the periphery, they may not be perceptible because stereoacuity also decreases in the periphery [McKee 1983; Rawlings and Shipley 1969]. For example, stereoacuity at 10° eccentricity was measured to be 80 arcsec [Fendick and Westheimer 1983]. In comparison, the expected amount of disparity distortion due to a lack of appropriate ocular parallax rendering at 10° eccentricity is 73.76 arcsec (fixating at 4 D). This implies that the disparity distortions of conventional stereo rendering in peripheral vision may be too small to have a measurable effect on depth perception.

³We make the assumption that the virtual image is at optical infinity, i.e. $d = \infty$, such that distant objects do not move during eye motion but objects at closer distances shift relative to the background.




Fig. 5. Conventional stereoscopic rendering distorts disparities in the periphery. Here, two points are rendered by projecting them through the centers of rotation of both left (L) and right (R) eyes onto a virtual image (dashed lines); this process represents conventional stereo rendering. The centers of projection of a user's eyes, however, are the front nodal points. This small discrepancy results in distortions of the perceived disparity (top right) and depth (bottom right) in the periphery of the visual field, which is corrected with ocular parallax rendering.

We validate this intuition with a user experiment that studies the effect of ocular parallax on absolute depth perception in Section 6.2. Further discussions on this topic can be found in the supplement.

4.3 Implementation

Hardware. We implement ocular parallax rendering with a prototype virtual reality system. The VR system is an HTC Vive Pro connected to the open source binocular Pupil Labs eye tracker. The HTC Vive Pro has a field of view of 110° , a refresh rate of 90 Hz, and a 1440×1600 pixel organic light-emitting diode display, resulting in a theoretical resolution of 4.58 arcmin/pixel. The HTC Vive Pro supports built in ipd adjustment. The Pupil Labs eye tracker snaps into the the HTC Vive Pro and estimates a global fixation point (i.e. F) at 120 Hz, supporting about 1° of gaze accuracy and 0.08° of gaze precision (see supplement for additional accuracy characterization). In the egocentric depth perception study, an HTC Vive Tracker estimates the position and orientation of the users' hand at 120 Hz with an accuracy of 1.9 mm root-mean-square-error (RMSE) and a precision of 1.5 mm RMSE⁴.

Software and Calibration. Unity was used as the rendering engine for both the ocular parallax rendering and the user experiments. Pixel-precise rendering and anti-aliasing were used for all rendered stimuli. Pupil Labs provides a Unity plugin that interfaces with their Python library, providing eye-tracking calibration and gaze tracking. All software related to acquiring subject data was written as C# scripts in Unity. The data were then analyzed in Python.

5 PERCEPTUAL THRESHOLDS FOR OCULAR PARALLAX IN VR

The primary goal of this section is to establish depth-dependent detection and discrimination thresholds for ocular parallax in VR. Although ocular parallax has been shown to be well within the range

of human visual acuity for natural viewing conditions (e.g., Bingham [1993]), we are not aware of any work that actually verified that this small effect size is even perceivable with the limited resolution offered by current-generation VR displays or, if it is, what the specific thresholds are. These thresholds are crucial for understanding in which conditions ocular parallax rendering is visible and how reliable of a depth cue it may be. To estimate these thresholds, we perform two psychophysical experiments that are discussed in the following sections with the apparatus described in Section 4.3.

5.1 Detection Thresholds for Ocular Parallax

With this experiment, we aim to estimate a depth-dependent detection threshold at which ocular parallax is perceivable.

Stimuli. As seen in Figure 6 (left), we presented two circular surfaces at varying distances to the subject. The surfaces were scaled to subtend 2° of visual angle irrespective of depth. The farther surface was solid red and the front one was textured with white noise. Without ocular parallax enabled, the front surface exactly occluded the back one.

An additional small red and white fixation target was then rendered, circling around the scene at a distance of 16° from the center, or 15° from the edges of the front and back surfaces. This was the largest eccentricity that kept the target within the aberration-free viewing zone of the Vive Pro. This fixation target rotated with an angular velocity of $90^\circ/\text{s}$, resulting in a perceived retinal motion of $24.81^\circ/\text{s}$, which falls within smooth pursuit rates [Westheimer 1954]. The starting position and rotation direction of the fixation target were randomized per trial.

Conditions. All stimuli were presented monocularly to the right eye. In each trial, the absolute distance between the viewer and back surface was randomly chosen as 1, 2, or 3 D. The relative distance between the front and back surface was also randomly chosen as 0, 0.25, 0.5, 0.75, or 1 D. In each trial, subjects viewed the stimulus with ocular parallax rendering enabled and disabled. These two conditions were presented in different 2-second intervals in random order, separated by a 0.5 second blank frame.

Subjects. Six adults participated (age range 26–38, 1 female). Due to the demanding nature of our psychophysical experiment, only a few subjects were recruited, which is common for low-level psychophysics (see e.g. [Patney et al. 2016]).

All subjects in this and all following experiments had normal or corrected to normal vision, no history of visual deficiency, and no color blindness. All subjects gave informed consent. The research protocol was approved by the Institutional Review Board at the participating university.

Procedure. To start the session, each subject performed a 7-point eye-tracker calibration that is provided by the manufacturer. To minimize eye-tracker error over the course of the experiment, each trial began with a single-point re-calibration of the eye trackers; subjects were instructed to fixate on a cross target centered on the screen, and the measured gaze direction was used to compensate for possible drift error. Subjects then viewed the stimulus rendered with one of the two ocular parallax rendering conditions for 2 seconds,

⁴<http://doc-ok.org/?p=1478>




Fig. 6. Detection and discrimination thresholds for ocular parallax in VR. Two experiments were conducted to estimate perceptual thresholds using an HTC Vive Pro head mounted display presenting stimuli of the form shown on the left. A red surface subtending 2° of visual angle was completely occluded by a noisy gray surface in front of it (left). The relative distances of these surfaces were varied with conditions described in the text. Detection thresholds (center left) and discrimination thresholds (center right) were estimated from the psychometric functions fitted to the recorded subject data (examples shown right).

then a blank screen for 0.5 seconds, and then again the stimulus with the other rendering condition from the first interval for another 2 seconds. Each trial constituted a two-alternative forced choice test, and subjects were asked to choose the time interval which exhibited more relative motion between the two surfaces with a keyboard. This concluded the trial. No feedback was provided. Subjects were instructed to fixate only on the moving fixation target, and never to the center surfaces.

There were 15 distance configurations and 15 trials for each configuration for a total of 225 trials per subject. The experiment took about 25 minutes to complete, including instruction and eye tracking calibration per subject.

Analysis. For each of the 15 distance configurations, we computed the proportion of correct responses. Using Bayesian inference methods [Schütt et al. 2016; Wichmann and Hill 2001a,b], we fit a psychometric function to each subject’s performance at each of the three absolute depths of the back surface. Each psychometric function gives us a detection threshold, measured in diopters of relative distance from the absolute distance of the back surface. The thresholds represent where the psychometric function exceeded a 75% chance for a correct response. An example of one of these psychometric functions is shown in Figure 6 (top right) and all measured psychometric functions are available in the supplement.

Results. The detection thresholds, averaged across subjects, are plotted in Figure 6 (center left). We see that these thresholds are invariant to the absolute distance of the back surface. This is expected because the conditions were equally spaced in diopters and ocular parallax, like other types of parallax cues, is perceptually linear in dioptric space (see supplement for an analysis). The estimated detection threshold is approximately 0.36 D. This implies that, for an eccentricity of 15° and a relative object distance as low as 0.36 D, ocular parallax may be perceivable in VR.

This result is surprising because visual acuity alone (see Fig. 3) predicts detection thresholds that are an order of magnitude higher

than what we measured. Yet, our results are consistent with data reported for physical viewing conditions (i.e. non-VR settings) [Bingham 1993] and with differential velocity thresholds [McKee and Nakayama 1984]. We expect the detection thresholds to increase linearly with eccentricity in the extra-foveal region because both the magnitude of micro parallax (see Fig. 3) and differential velocity thresholds increase roughly linearly there [McKee 1983]. Also, we expect detection thresholds to decrease with faster eye movements because differential velocity thresholds inversely vary with retinal velocity [McKee 1983].

Our results emphasize that the human visual system is much more sensitive to small amounts of motion, even in peripheral vision, than naively expected. Even the small amount of parallax induced by ocular parallax rendering may be visible in many VR applications. An important question that arises from this insight is whether ocular parallax rendering can improve depth perception or the realism of a 3D scene. We perform several experiments in Sections 6 and 7 that aim at answering this question.

5.2 Discrimination Thresholds for Ocular Parallax

A discrimination threshold tells us what the smallest amount of perceivable change in ocular parallax is. For a fixed eye eccentricity, this threshold depends on both the absolute distance of the reference surface (i.e., the back surface) and also the relative offset from that surface. Conceptually, one would have to estimate discrimination thresholds for each combination of absolute and relative distance, which would be an arduous task. However, due to the fact that the detection thresholds are depth independent, we assume that the discrimination thresholds are also independent of the absolute distance to the back surface. This assumption makes it easier to set up an experiment to estimate discrimination thresholds for ocular parallax in VR, which we did by performing a second experiment that uses the same apparatus, stimuli, and analysis as the first experiment and a very similar procedure, but with slightly different conditions. Six adults participated (age range 26–32, 1 female).

Conditions. The back surface was fixed to 0 D for all conditions. Instead of presenting the stimulus with ocular parallax rendering

enabled for only one interval, as done in the previous experiment, we enabled ocular parallax rendering for both intervals. The relative offset from the back surface was randomly chosen as 1, 2, or 3 D and assigned as the depth for one of the front surfaces shown in each trial. The other appeared at one of the following distances from the previously chosen relative offset: 0, 0.45, 0.9, 1.35, and 1.8 D for the 1 D and 2 D initial offsets; and 0, 0.7, 1.4, 2.1, and 2.8 D for the 3 D initial offset. Again, there were 15 distance configurations and 15 trials for each configuration for a total of 225 trials per subject.

Results. We used a similar analysis as the one described in the previous subsection to estimate the discrimination thresholds, which are plotted in Figure 6 (center right). As expected, the discrimination thresholds increase linearly with increasing amounts of ocular parallax, per Weber’s law, but due to the proximity to the detection threshold magnitude, the slope is less than 1. A linear fit between the discrimination thresholds and relative offset from 0 D has a slope of 0.11 and an intercept of 0.38 D, further verifying the measured 0.36 D detection threshold from the previous experiment.

In conclusion, the estimated discrimination thresholds are well within the range of the depth ranges of natural scenes commonly rendered for immersive virtual reality applications. This result motivates further studies to investigate if or when ocular parallax could be used as a reliable depth cue in these applications. We take a first step at answering this question by conducting experiments that study whether ocular parallax has a measurable effect on depth perception in the following sections.

6 OCULAR PARALLAX AND DEPTH PERCEPTION

Motivated by the discussion in Section 4.2 and the surprisingly low detection and discrimination thresholds measured in the previous section, we proceed to investigate the importance of ocular parallax as a depth cue. Here, we distinguish between ordinal depth perception, which provides information about the *relative* depth ordering of a scene (i.e., object A is closer than object B), and metric depth cues, which also provide *absolute* distance estimates (i.e., object A is 1 m away and object B is 2 m away). It seems intuitive that gaze-contingent occlusion can provide ordinal depth information due to the deletion and accretion of the farther object’s texture near occlusion boundaries. Moreover, gaze-induced micro parallax also provides ordinal information in the relative magnitudes of retinal velocities of objects at different depths. However, the effectiveness of ocular parallax as a metric depth cue is questionable; occlusion only provides ordinal depth information and the absolute motion velocity magnitudes necessary to determine metric depth may be unreliable. We investigate the effect of ocular parallax rendering on both ordinal and metric depth perception in the following experiments.

6.1 Effect of Ocular Parallax on Ordinal Depth Estimation

In this section, we study the benefits of ocular parallax rendering on ordinal depth estimation. We also investigate whether depth perception gained from ocular parallax is purely a visual process, or requires some non-visual, extra-retinal signal like the magnitude or direction of eye rotation. The perception of depth from head-motion parallax, for example, has been shown to rely on an extra-retinal




Fig. 7. Effect of ocular parallax on ordinal depth perception. Subjects viewed the monocular stimuli consisting of two distinctly textured surfaces separated by 1 D or 2 D (top) and were asked which one was closer. The proportions of correct responses, averaged across subjects per condition, are plotted on the bottom. Subjects performed significantly better with ocular parallax rendering enabled compared to conventional rendering. However, they also performed slightly better than conventional rendering with reversed ocular parallax rendering, indicating that the extra-retinal signal of eye rotation may not be crucial for depth perception. Significance is indicated at the $p \leq 0.05$, 0.01, and 0.001 levels with *, **, and ***, respectively.

signal in the form of optokinetic response eye movements [Nawrot 2003]. This experiment is modeled after a related experiment investigating retinal blur as an ordinal depth cue [Zannoli et al. 2016]. Twenty-one adults participated (age range 22–43, 4 females), of which two were excluded for failing to follow instructions.

Stimuli. The monocular stimuli consisted of two differently textured, frontoparallel surfaces at different depths (Fig. 7, top). The textures, the same those in the aforementioned depth order study [Zannoli et al. 2016], had the same space-average luminance and contrast energy, and exhibited similar amplitude spectra. The rear surface was fixed at 0.5 D and the front surface appeared at either 1.5 D or 2.5 D. The border between the two surfaces had a sinusoidal shape. The surfaces were scaled to subtend the same visual angle irrespective of depth. Subjects viewed the surfaces through a 20° circular aperture that was unaffected by ocular parallax rendering.

Procedure. First, all subjects performed a 7-point eye-tracker calibration that is provided by the manufacturer to start the session. As in the psychophysical experiments, each individual trial began with a single-point re-calibration of the eye trackers. Following the re-calibration, the stimulus was presented monocularly to the right eye for 3 seconds with one of the three rendering conditions. Subjects were instructed to gaze between the two surfaces. The stimulus was then replaced with a prompt asking which surface, left or right, appeared nearer, and a keyboard response concluded the trial. No feedback was provided.

Conditions. We evaluated three different rendering conditions (conventional, ocular parallax, and reversed ocular parallax) and two near surface distances (1.5 D and 2.5 D). With the monocular stimuli devoid of perspective and accommodation/retinal blur cues, the only expected sources of depth information are gaze-contingent micro parallax and occlusion. We included both correct and reversed ocular parallax rendering to understand whether depth perception in this scenario is purely a visual process, requiring only the retinal image with micro parallax, or whether it also requires an extra-retinal signal in the form of eye rotation direction. With reversed ocular parallax rendering, the sign of the visual signal is negated, causing occluded surface textures to accrete when they would normally be deleted. Subjects expecting the visual and extra-retinal signals to be consistent could misinterpret depth orderings. Reversed ocular parallax was implemented by negating the x and y components of the estimated nodal points of the eyes, $N_{L/R}$.

The near surface distances create a 1 D and 2 D separation from the rear surface, resulting in different magnitudes of the gaze-contingent effects. Following the threshold experiments, we expect the proportion of correct responses to increase with surface separation. Overall, there were 6 conditions and each condition was evaluated with 15 trials for a total of 90 trials per subject.

Results. The proportion of correct responses, averaged across subjects per condition, is plotted in Figure 7 (bottom). As expected, subjects in the conventional static rendering condition performed close to random, correctly identifying the nearer surface in 48.8% and 52.6% of trials for 1 D and 2 D of surface separation, respectively. Enabling ocular parallax rendering clearly improved ordinal depth judgment with subjects performing at 66.7% and 75.8% correct identification for the 1 D and 2 D separations, respectively. Subjects in the reversed ocular parallax condition fell in-between, performing at 60.4% and 67.4% correct identification for the two separation distances.

We conducted a 2×3 repeated-measures ANOVA on the proportion of correct responses with independent variables of rendering mode (conventional, ocular parallax, reversed ocular parallax) and separation distance (1 D or 2 D). Greenhouse-Geisser sphericity correction was applied. The ANOVA shows a very significant effect of rendering mode ($F(1.7, 30.65) = 17.98, p < 0.0001$) as well as a significant effect of distance ($F(1, 18) = 8.14, p < 0.05$). The ANOVA does not reveal a significant interaction between rendering mode and distance ($F(1.78, 31.96) = 0.42, p = 0.64$).

Post-hoc tests were conducted as pairwise t -tests between rendering modes at each separation distance, with Bonferroni correction applied to the p -values. The post-hoc tests found that, at 1 D of separation, ocular parallax rendering shows a significant improvement over conventional rendering ($p < 0.001$), but not over reversed ocular parallax rendering. Reversed ocular parallax rendering does not show a significant improvement over conventional rendering. At 2 D of separation, ocular parallax rendering shows a significant improvement over conventional rendering ($p < 0.001$) as well as reversed ocular parallax rendering ($p < 0.05$). Reversed ocular parallax rendering also shows a significant improvement over conventional rendering ($p < 0.05$) at this separation.




Fig. 8. Effect of ocular parallax on absolute, egocentric depth perception. Subjects viewed a target (pencil) and then, with the display turned off, were asked to reach for it (top right). We did not find a statistically significant difference between conventional stereo rendering and stereo rendering with ocular parallax in this experiment (bottom plot).

In summary, this experiment demonstrates that ocular parallax significantly improves ordinal depth perception over conventional rendering. However, reversed ocular parallax also improves ordinal depth perception over conventional rendering, but not nearly as much as when rendered correctly. The reduced performance in the reversed ocular parallax condition compared to the correct ocular parallax condition suggests that extra-retinal signals, like eye rotation, play an important role in the perception of depth. However, this effect seems to be weaker than for motion parallax, where a reversal in eye movement causes a reversal in the sign of the perceived depth [Nawrot 2003]. Still, subjects performed better, in some conditions significantly so, when the directions of the retinal image motion and eye rotation were consistent. Therefore, enabling correct ocular parallax rendering can benefit ordinal depth estimation which can be particularly useful for viewing 3D scenes that often contain many occlusion boundaries.

6.2 Effect of Ocular Parallax on Metric Depth Estimation

Next, we study whether ocular parallax rendering can benefit egocentric distance (distance from one's self) estimation in a stereo environment. While unlikely, gaze-induced micro parallax and the reduction in peripheral disparity errors could affect metric distance estimation.

In the experiment, subjects performed a blind reaching task introduced by Napieralski et al. [2011] to estimate distances to objects because verbal estimates have been shown to be inaccurate [Renner et al. 2013]. In a photo-realistic dining room scene, subjects viewed a stereoscopically rendered pencil for a minimum of 5 seconds after which the screen was blanked and they reached their hand to

where they had last seen the pencil (Figure 8, top left). Each subject performed the task with ocular parallax enabled and disabled, and the reach target distances were set proportionally—50, 58, 67, 75, 82, and 90%—to each subjects' maximum arm reach. Each of the 12 conditions were evaluated with five blind reaching tasks for a total of 60 randomly-ordered trials per subject. During the session, 6-DOF head pose tracking was enabled, and the hand's position was tracked via an HTC Vive tracker mounted on an optical post (Figure 8, top right). Sixteen young adults participated (age range 22–32, 5 females), of which one was excluded for not passing a standard Randot stereo vision test, and two others were excluded due to the eye tracker failing to track their pupils.

Subjects saw little to no improvement in their egocentric distance estimates with ocular parallax rendering enabled. Figure 8 shows each subject's reach, as a percentage of their maximum arm reach, to a presented target distance for the ocular parallax enabled and disabled conditions. Linear models were fit to the two sets of reaches; the slopes for the ocular parallax enabled and disabled modes were 1.141 and 1.128, respectively, while the intercepts were 0.167 and 0.165, respectively. A multiple regression analysis did not show a significant difference between the two conditions. It is therefore unlikely that the systematic underestimation of distances to virtual objects compared to real ones [Renner et al. 2013] can be explained by the omission of ocular parallax rendering. For more details regarding this experiment, please refer to the supplement.

7 OCULAR PARALLAX AND PERCEPTUAL REALISM

We next measured the effect of ocular parallax on perceptual realism. Twenty-one adults participated in the experiment (age range 22–43, 4 females), of which two were excluded for failing to follow instructions.

Stimuli and Conditions. A 3D scene (Fig. 9, left) was presented to the right eye with conventional, ocular parallax, or reversed ocular parallax rendering. The scene was composed of four grid planes extending into the distance and a field of frontoparallel targets. These targets were randomly distributed in depth between 0.5 D and 3.5 D and were scaled to subtend 2.6° of visual angle regardless of depth. The targets' lateral positions were randomly chosen such that they did not occlude one another and the entire target field subtended 60° . Subjects viewed the stimulus in three rendering conditions: conventional, ocular parallax, and reversed ocular parallax.

Procedure. Each trial began with the single-point eye-tracking re-calibration used in the psychophysical experiments. Then, the target positions used for the subsequent stimuli were generated. The first stimulus was presented for 3 seconds with one of the three rendering conditions. The targets then disappeared for 1 second, leaving only the grid pattern shown with conventional rendering. The second stimulus was then presented for 3 seconds using one of the three rendering conditions, but not the same as the one used for the first stimulus. After the second stimulus, a prompt appeared asking the subjects to determine, in a forced-choice judgment, the stimulus that portrayed a stronger impression of realistic depth. No feedback was provided. Each trial therefore consisted of three




Fig. 9. Evaluating perceptual realism. Subjects viewed a 3D scene consisting of targets that are randomly distributed in depth but that do not occlude one another (left). This stimulus was presented with either conventional, ocular parallax, or reversed ocular parallax rendering and we asked subjects to indicate which rendering mode provided a stronger impression of realistic depth. Results of pairwise comparisons between these rendering modes show the percent of times the first member of the pair was chosen over the second (right). Rendering with correct and reversed ocular parallax conveyed a stronger impression of depth compared to conventional rendering, but when compared against one another no difference was observed. This result indicates that the relative magnitudes of depth-dependent motion velocities are most important for perceptual realism but not necessarily their direction. Significance is indicated at the $p \leq 0.05$, 0.01 , and 0.001 levels with *, **, and ***, respectively.

pairwise comparisons: *conventional vs. ocular parallax*, *conventional vs. reversed ocular parallax*, and *ocular parallax vs. reversed ocular parallax*.

Results. The results of the pairwise comparisons are averaged across users and trials and plotted in Figure 9 (right). In 76.8% of trials, subjects reported a stronger sense of realistic depth when viewing ocular parallax rendering over conventional rendering. Interestingly, subjects also reported more realistic depth for reversed ocular parallax rendering over conventional rendering in 75.1% of trials. Each of these percentages is significantly greater than 50% ($p < 0.001$ and $p < 0.001$ respectively, one-tailed binomial test). Users had more difficulty differentiating between the two ocular parallax modes and reported more realistic depth in ocular parallax rendering in only 49.1% of trials, which was not significantly greater than 50%.

These results suggest that gaze-induced micro parallax, offered by ocular parallax rendering, has a significant effect on perceptual realism when compared to conventional rendering. Moreover, we see that the relative motion magnitudes of the depth-dependent retinal velocities are more important than their direction. This insight further emphasizes that, similar to the ordinal depth perception task (Sec. 6.1), extra-retinal signals were likely not notably factored into the perception of depth on this subjective task. Unlike the experiment in Section 6.1, where both gaze-contingent occlusion and micro parallax were presented as cues, in this experiment, micro parallax was the primary indicator of depth because the targets did not overlap.

8 DISCUSSION

In summary, our primary contribution is to introduce a new technology for virtual and augmented reality: ocular parallax rendering. This technique is enabled by eye tracking systems which can already be found in some of the latest headsets, like the Microsoft HoloLens 2, Magic Leap One, Varjo, Fove, and HTC Vive Pro Eye. Ocular parallax rendering could be jointly implemented with other gaze-contingent rendering methods, such as foveated rendering, and it requires no additional computational cost compared to conventional stereo rendering.

To evaluate ocular parallax rendering, we designed and conducted a series of user experiments. First, we measured detection thresholds and show that the effect is perceivable when the relative distance between two partially occluding objects is as low as 0.36 diopters. We also measured discrimination thresholds and confirm that the just noticeable difference between two objects at different relative depths in front of a background stimulus is directly proportional to their absolute depth from the background. These thresholds confirm that ocular parallax is an important *visual cue* in VR. Our third and fourth experiments show that ocular parallax acts as a reliable ordinal depth cue but that it is not a reliable metric depth cue. Finally, our fifth experiment demonstrates that ocular parallax rendering significantly improves the impression of realistic depth when viewing a 3D scene over conventional rendering.

Ocular Parallax and Retinal Blur. Retinal blur describes the depth-dependent defocus blur of objects on the retina, relative to the accommodation distance. In physical environments, retinal blur discrimination thresholds are comparable to those of ocular parallax—approximately ± 0.35 D [Ogle and Schwartz 1959]—and have been shown to serve as both a reliable ordinal and metric depth cue [Vishwanath and Blaser 2010]. While recent research efforts in VR/AR have studied focus cues such as retinal blur, chromatic aberrations, and accommodation, none have studied ocular parallax, even though their effect sizes are very similar.

Gaze-contingent retinal blur and ocular parallax rendering can certainly be implemented simultaneously as they are complementary cues. While depth-of-field rendering blurs objects at different depths, we do not expect it to alter our measurements or insights significantly. This is because the perception of motion is understood to use low spatial frequencies [Smith and Snowden 1994], so the loss of the high spatial frequencies due to depth-of-field rendering should not impair motion, or ocular parallax, perception. Indeed, our measured detection thresholds are consistent with those measured in a physical environment [Bingham 1993], where retinal blur was apparent.

Limitations. Although our system uses some of the highest-end components available, including a high-resolution wide-field-of-view VR display and a 120 Hz eye tracker, the latency of 20 ms for gaze-contingent ocular parallax rendering is high. Faster and more accurate eye tracking would certainly help improve the user experience for all gaze-contingent rendering schemes, including ocular parallax.

Applications. We envision ocular parallax rendering to be a standard part of the graphics pipeline of eye-tracking-enabled near-eye displays. It improves perceptual realism, ordinal depth perception, and it may offer other perceptual benefits. In particular, optical see-through augmented reality systems may benefit from ocular parallax rendering as a user usually sees a digitally rendered stimulus overlaid on a reference stimulus (i.e. the physical world); visual cue consistency between these stimuli may be even more important than in VR. Finally, one could also imagine that an amplified version of ocular parallax rendering could be an effective gaze-contingent user interface that allows users to transform objects, navigate through virtual environments, or perform other tasks. Tasks that require hands-free operation could particularly benefit from this type of gaze-contingent interaction mode.

Future Work. As the fields of view of emerging VR and AR systems keep increasing, understanding perceptual effects in peripheral vision becomes ever more important. With this work, we thoroughly evaluate the perceptual implications of one technique, ocular parallax rendering, which shows the strongest effects in near-mid peripheral vision. However, many other technologies, such as foveated rendering, reducing motion sickness, or other means to improve depth perception or perceptual realism, could also benefit from studying their perceptual effects on peripheral vision. Additionally, as many of these effects vary from person to person, performing a per-user calibration of the distance between the centers of rotation and projection could further increase perceptual realism.

Conclusions. Virtual and augmented reality systems have focused on improving resolution, field of view, device form factor, and other characteristics. With this work, we hope to stimulate new directions for gaze-contingent rendering and improve perceptual realism and depth perception with next-generation near-eye displays.

REFERENCES

- Kurt Akeley, Simon J. Watt, Ahna Reza Girshick, and Martin S. Banks. 2004. A Stereo Display Prototype with Multiple Focal Distances. *ACM Trans. Graph. (SIGGRAPH)* 23, 3 (2004), 804–813.
- David A. Atchison. 2017. Schematic Eyes. In *Handbook of Visual Optics, Volume I - Fundamentals and Eye Optics*, Pablo Artal (Ed.). CRC Press, Chapter 16.
- Geoffrey P. Bingham. 1993. Optical flow from eye movement with head immobilized: “Ocular occlusion” beyond the nose. *Vision Research* 33, 5 (1993), 777 – 789.
- David Brewster. 1845. On the Law of Visible Position in Single and Binocular Vision, and on the representation of Solid Figures by the Union of dissimilar Plane Pictures on the Retina. *Proc. Royal Society of Edinburgh* 1 (1845).
- Steven A. Cholewiak, Gordon D. Love, Pratul P. Srinivasan, Ren Ng, and Martin S. Banks. 2017. Chromabur: Rendering Chromatic Eye Aberration Improves Accommodation and Realism. *ACM Trans. Graph. (SIGGRAPH Asia)* 36, 6 (2017), 210:1–210:12.
- James Cutting and Peter Vishton. 1995. Perceiving layout and knowing distances: The interaction, relative potency, and contextual use of different information about depth. In *Perception of Space and Motion*, William Epstein and Sheena Rogers (Eds.). Academic Press, Chapter 3, 69–117.
- Piotr Didyk, Tobias Ritschel, Elmar Eisemann, Karol Myszkowski, and Hans-Peter Seidel. 2011. A Perceptual Model for Disparity. *ACM Trans. Graph. (SIGGRAPH)* 30, 4 (2011), 96:1–96:10.
- Andrew T. Duchowski, Nathan Cournia, and Hunter A. Murphy. 2004. Gaze-Contingent Displays: A Review. *Cyberpsychology & behavior* 7 (2004), 621–34.
- A. T. Duchowski, D. H. House, J. Gestring, R. I. Wang, K. Krejtz, I. Krejtz, R. Mantiuk, and B. Bazyluk. 2014. Reducing Visual Discomfort of 3D Stereoscopic Displays with Gaze-contingent Depth-of-field. In *ACM Symposium on Applied Perception*. 39–46.
- David Dunn, Cary Tippets, Kent Torell, Petr Kellnhofer, Kaan Aksit, Piotr Didyk, Karol Myszkowski, David Luebke, and Henry Fuchs. 2017. Wide Field Of View Varifocal

- Near-Eye Display Using See-Through Deformable Membrane Mirrors. *IEEE TVCG* 23, 4 (2017), 1322–1331.
- M. Fendick and G. Westheimer. 1983. Effects of practice and the separation of test targets on foveal and peripheral stereoacuity. *Vision Research* 23, 2 (1983), 145–150.
- G. A. Fry and W. W. Hill. 1962. The Center of Rotation of the Eye. *Optometry and Vision Science* (1962), 581–595.
- Brian Guenter, Mark Finch, Steven Drucker, Desney Tan, and John Snyder. 2012. Foveated 3D Graphics. *ACM Trans. Graph. (SIGGRAPH Asia)* 31, 6 (2012), 164:1–164:10.
- I. Hadani, G. Ishai, and M. Gur. 1980. Visual stability and space perception in monocular vision: mathematical model. *OSA J. Opt. Soc. Am.* 70, 1 (1980), 60–65.
- David M. Hoffman, Ahna R. Girshick, Kurt Akeley, and Martin S. Banks. 2008. Vergence–Accommodation conflicts hinder visual performance and cause visual fatigue. *Journal of Vision* 8, 3 (2008), 33.
- Ian P. Howard and Brian J. Rogers. 2002. *Seeing in Depth*. Oxford University Press.
- Hong Hua and Bahram Javidi. 2014. A 3D integral imaging optical see-through head-mounted display. *Optics Express* 22, 11 (2014), 13484–13491.
- Fu-Chung Huang, Kevin Chen, and Gordon Wetzstein. 2015. The Light Field Stereoscope: Immersive Computer Graphics via Factored Near-Eye Light Field Display with Focus Cues. *ACM Trans. Graph. (SIGGRAPH)* 34, 4 (2015).
- Paul V. Johnson, Jared A.Q. Parnell, Joohwan Kim, Christopher D. Saunter, Gordon D. Love, and Martin S. Banks. 2016. Dynamic lens and monovision 3D displays to improve viewer comfort. *OSA Opt. Express* 24, 11 (2016), 11808–11827.
- Petr Kellnhofer, Piotr Didyk, Tobias Ritschel, Belen Masia, Karol Myszkowski, and Hans-Peter Seidel. 2016. Motion Parallax in Stereo 3D: Model and Applications. *ACM Trans. Graph. (SIGGRAPH)* 35, 6 (2016), 176:1–176:12.
- Robert Konrad, Emily Cooper, and Gordon Wetzstein. 2015. Novel Optical Configurations for Virtual Reality: Evaluating User Preference and Performance with Focus-tunable and Monovision Near-eye Displays. In *Proc. SIGCHI*.
- H. Kudo and N. Ohnishi. 1998. Study on the ocular parallax as a monocular depth cue induced by small eye movements during a gaze. In *Proc. IEEE Engineering in Medicine and Biology Society*, Vol. 6. 3180–3183.
- H Kudo and N Ohnishi. 2000. Effect of the sight line shift when a head-mounted display is used. In *Proc. EMBS International Conference*, Vol. 1. 548–550.
- H. Kudo, M. Saito, T. Yamamura, and N. Ohnishi. 1999. Measurement of the ability in monocular depth perception during gazing at near visual target-effect of the ocular parallax cue. In *Proc. IEEE International Conference on Systems, Man, and Cybernetics*, Vol. 2. 34–37.
- Michael F. Land. 1995. Fast-focus telephoto eye. *Nature* 373 (1995), 658–659.
- Douglas Lanman and David Luebke. 2013. Near-eye Light Field Displays. *ACM Trans. Graph. (SIGGRAPH Asia)* 32, 6 (2013), 220:1–220:10.
- Sheng Liu, Dewen Cheng, and Hong Hua. 2008. An Optical See-through Head Mounted Display with Addressable Focal Planes. In *Proc. ISMAR*. 33–42.
- Andrew Maimone, Andreas Georgiou, and Joel S. Kollin. 2017. Holographic Near-eye Displays for Virtual and Augmented Reality. *ACM Trans. Graph. (SIGGRAPH)* 36, 4 (2017), 85:1–85:16.
- Alistair P. Mapp and Hiroshi Ono. 1986. The rhino-optical phenomenon: Ocular parallax and the visible field beyond the nose. *Vision Research* 26, 7 (1986), 1163–1165.
- Jonathan A. Marshall, Christina A. Burbeck, Dan Ariely, Jannick P. Rolland, and Kevin E. Martin. 1996. Occlusion edge blur: a cue to relative visual depth. *J. Opt. Soc. Am. A* 13, 4 (1996), 681–688.
- George Mather and David R R Smith. 2002. Blur Discrimination and its Relation to Blur-Mediated Depth Perception. *Perception* 31, 10 (2002), 1211–1219.
- M. Mauderer, S. Conte, M. A. Nacenta, and D. Vishwanath. 2014. Depth Perception with Gaze-contingent Depth of Field. *ACM SIGCHI* (2014).
- Suzanne P. McKee. 1983. The spatial requirements for fine stereoacuity. *Vision Research* 23, 2 (1983), 191–198.
- Suzanne P. McKee and Ken Nakayama. 1984. The detection of motion in the peripheral visual field. *Vision Research* 24, 1 (1984), 25–32.
- Olivier Mercier, Yusuf Sulai, Kevin Mackenzie, Marina Zannoli, James Hillis, Derek Nowrouzezahrai, and Douglas Lanman. 2017. Fast Gaze-contingent Optimal Decompositions for Multifocal Displays. *ACM Trans. Graph. (SIGGRAPH Asia)* 36, 6 (2017).
- Phillip E. Napieralski, Bliss M. Altenhoff, Jeffrey W. Bertrand, Lindsay O. Long, Sabarish V. Babu, Christopher C. Pagano, Justin Kern, and Timothy A. Davis. 2011. Near-field Distance Perception in Real and Virtual Environments Using Both Verbal and Action Responses. *ACM Trans. Appl. Percept.* 8, 3, Article 18 (Aug. 2011), 19 pages.
- Mark Nawrot. 2003. Eye movements provide the extra-retinal signal required for the perception of depth from motion parallax. *Vision research* 43, 14 (2003), 1553–1562.
- Kenneth N Ogle and J Theodore Schwartz. 1959. Depth of focus of the human eye. *JOSA* 49, 3 (1959), 273–280.
- Nitish Padmanaban, Robert Konrad, Tal Stramer, Emily A. Cooper, and Gordon Wetzstein. 2017. Optimizing virtual reality for all users through gaze-contingent and adaptive focus displays. *PNAS* 114, 9 (2017), 2183–2188.
- Stephen E. Palmer. 1999. *Vision Science - Photons to Phenomenology*. MIT Press.
- Stephen E Palmer and Joseph L Brooks. 2008. Edge-region grouping in figure-ground organization and depth perception. *Journal of Experimental Psychology: Human Perception and Performance* 34, 6 (2008), 1353.
- Anjul Patney, Marco Salvi, Joohwan Kim, Anton Kaplanyan, Chris Wyman, Nir Benty, David Luebke, and Aaron Lefohn. 2016. Towards Foveated Rendering for Gaze-tracked Virtual Reality. *ACM Trans. Graph. (SIGGRAPH Asia)* 35, 6 (2016), 179:1–179:12.
- John D. Pettigrew, Shaun P. Collin, and Matthias Ott. 1999. Convergence of specialised behaviour, eye movements and visual optics in the sandlance (Teleostei) and the chameleon (Reptilia). *Current Biology* 9, 8 (1999), 421–424.
- Samuel C. Rawlings and T. Shipley. 1969. Stereoscopic Acuity and Horizontal Angular Distance from Fixation*. *J. Opt. Soc. Am.* 59, 8 (Aug 1969), 991–993.
- Rebekka S. Renner, Boris M. Velichkovsky, and Jens R. Helmert. 2013. The Perception of Egocentric Distances in Virtual Environments - A Review. *ACM Comput. Surv.* 46, 2 (2013), 23:1–23:40.
- Jannick P. Rolland, Myron W. Krueger, and Alexei Goon. 2000. Multifocal planes head-mounted displays. *OSA Appl. Opt.* 39, 19 (2000), 3209–3215.
- Heiko H. Schütt, Stefan Harmeling, Jakob H. Macke, and Felix A. Wichmann. 2016. Painfree and accurate Bayesian estimation of psychometric functions for (potentially) overdispersed data. *Vision Research* 122 (2016), 105–123.
- Peter Shirley, Michael Ashikhmin, and Steve Marschner. 2009. *Fundamentals of Computer Graphics*. A K Peters/CRC Press.
- A.T. Smith and R.J. Snowden. 1994. *Visual Detection of Motion*. Academic Press. <https://books.google.com/books?id=p11qAAAAAAAJ>
- Dhanraj Vishwanath and Erik Blaser. 2010. Retinal blur and the perception of egocentric distance. *Journal of Vision* 10, 10 (08 2010), 26–26.
- G. Westheimer. 1954. Eye movement responses to a horizontally moving visual stimulus. *A.M.A. Archives of Ophthalmology* 52, 6 (1954), 932–941.
- Felix A. Wichmann and N. Jeremy Hill. 2001a. The psychometric function: I. Fitting, sampling, and goodness of fit. *Perception & Psychophysics* 63, 8 (01 Nov 2001), 1293–1313.
- Felix A. Wichmann and N. Jeremy Hill. 2001b. The psychometric function: II. Bootstrap-based confidence intervals and sampling. *Perception & Psychophysics* 63, 8 (01 Nov 2001), 1314–1329.
- Albert Yonas, Lincoln G. Craton, and William B. Thompson. 1987. Relative motion: Kinetic information for the order of depth at an edge. *Perception & Psychophysics* 41, 1 (01 Jan 1987), 53–59.
- Marina Zannoli, Rachel A. Albert, Abdullah Bulbul, Rahul Narain, James F. O'Brien, and Martin Banks. 2014. Correct blur and accommodation information is a reliable cue to depth ordering. *Journal of Vision* 14, 10 (2014), 138.
- Marina Zannoli, Gordon D. Love, Rahul Narain, and Martin S. Banks. 2016. Blur and the perception of depth at occlusions. *Journal of Vision* 16, 6 (04 2016), 17–17.

Supplementary Information: Gaze-Contingent Ocular Parallax Rendering for Virtual Reality

In this document we provide additional discussion and results in support of the primary text.

1 ADDITIONAL DETAILS ON DISPARITY DISTORTION

In Figure 1, we show additional details for calculating the disparity distortions of objects in the periphery of the visual field. On the right, we illustrate the settings we simulated: a user fixating on the optical axis on one of four different depths, either 1 D (1 m), 2 D (50 cm), 3 D (33.3 cm), or 4 D (25 cm). For each of these fixation distances, we then traced out points with the same respective distance from the center point between the eyes along an arc (red points). We then calculated what the disparity angles to left and right eye would be when rays go directly from the individual points through the center of rotation of the eyes and when they go through the front nodal point. The discrepancy between these disparity angles is plotted for all four distance settings and for a range of eccentricity angles (i.e., points moving along the arc) in the left. We see that the disparity distortion is zero for foveal vision, but steadily increases the farther we go into peripheral vision. Although the disparity distortion measured in angles of visual degrees is more pronounced for closer distances, when converting these distortion to meters, they are fairly independent of depth (center). However, the maximum depth distortion is only about 4 mm at an eccentricity of 60° , which is negligible. The perceived location of the distorted points is also rendered by blue dots on the right, showing that they are fairly close to the target points.

While these distortions increase in the periphery, they may not be perceptible because stereoacuity also decreases in the periphery [McKee 1983; Rawlings and Shipley 1969]. For example, stereoacuity at 10° eccentricity was measured to be 80 arcsec [Fendick and Westheimer 1983]. In comparison, the expected amount of disparity distortion due to a lack of appropriate ocular parallax rendering at 10° eccentricity is 73.76 arcsec (fixating at 4 D). This implies that the disparity distortions of conventional stereo rendering in peripheral vision may be too small to have a measurable effect on depth perception, which was indeed verified by the egocentric distance estimation study in the main paper. However, the micro parallax and gaze-contingent occlusion induced by ocular parallax rendering were shown to improve ordinal depth estimation and create a stronger impression of realistic depth.




Fig. 1. Simulated conditions for estimating disparity distortions caused by conventional stereo rendering (right). The amount of expected disparity distortion is expected to be zero for foveal vision and increasing in the periphery of the visual field (left). When converted to perceived depth distortion (center), the expected amount of distortion, even in the periphery, is negligible and perhaps imperceptible.

2 OCULAR PARALLAX LINEARITY IN DIOPTRIC SPACE

The amount of parallax due to eye rotations that is observed near an occlusion boundary is a function of the magnitude of the eye rotation, the distance to the farther surface, and the separation between two surfaces. Bingham [1993] presents a model taking these factors into account and the ocular parallax observed when the distances are considered in metric units is shown in Figure 2 (left). Clearly, ocular parallax is not linear with separation distance in meters. However, the amount of ocular parallax observed on the retina is linear when the distances are considered in diopters, or inverse meters (Figure 2, right). The simulated ocular parallax results in Figure 2 assumed 15° of eye rotation and a distance between the nodal point and center of rotation of 7.69 mm.




Fig. 2. Model of ocular parallax in metric and dioptric spaces. The plots show the estimated amount of ocular parallax between two objects separated in depth. The differently colored lines indicate the ocular parallax when the back object is set to 1, 2, or 3 D for varying distances to the front object. The left plot shows that the ocular parallax is not linear with metric distance, is linear with dioptric distance.

3 EYE TRACKER CHARACTERIZATION

We characterized the Pupil Labs eye tracker accuracy on a subset of the subjects from the experiments in the main paper (4 subjects, ages 22–28, 1 female). Subjects performed the manufacturer-provided 7-point eye tracker calibration followed by an additional single-point eye-tracker re-calibration to mimic the procedure of the main experiments. Subjects then viewed each of the five X-shaped targets in the order shown in Figure 3 (left); each target subtended 1° of visual field. The targets were placed at the center of their visual field and 15° above, below, to the left, and to the right of the center. The subjects were instructed to look to the center of the targets and indicated they were doing so with a keyboard press which recorded 5 seconds of gaze data. Users repeated this procedure for each target.

The results of the eye tracker accuracy characterization can be found in Figure 3 (right). The error averaged across all subjects and field positions was 1.18° , which is worse than the 1.0° error claimed by Pupil Labs. Clearly, the eye tracker error increases with increasing eccentricity, but was never higher than 1.6° (for the field position #3). One subject wore contacts which is known to deteriorate eye tracker performance. Indeed, the eye-tracker accuracy for this subject was 1.415° while the average eye-tracker accuracy for the subjects without contacts was 1.1° .




Fig. 3. Eye tracker characterization. After calibrating the eye tracker, subjects viewed each of the 5 targets on the left and their gaze data was recorded. The average errors across subjects for each field position is shown on the right. Naturally, the lowest error was observed at the center of the visual field, but the error never rose above 1.6° (target position #3) at the measured field positions.

4 PSYCHOMETRIC FUNCTIONS

We present the full set of psychometric functions from our detection and discrimination threshold experiments in Figure 4. The *psignifit* Python package [Schütt et al. 2016; Wichmann and Hill 2001a,b], was used for fitting psychometric functions to the data using Bayesian inference.




Fig. 4. Detection and Discrimination psychometric functions. The left column shows the psychometric functions for the ocular parallax detection experiment, while the right column presents the functions for the ocular parallax discrimination experiment. Each row of each column corresponds to the psychometric functions measured for one subject for 3 different stimulus conditions. The vertical lines intersecting the psychometric fit indicate the threshold corresponding to the 75% correct response rate, and the horizontal line abutting it representing the 95% confidence interval. Lower thresholds correspond to subjects being more sensitive to the ocular parallax effect.

5 ADDITIONAL DETAILS ON EGOCENTRIC DEPTH STUDY

Multiple reports have shown that egocentric depth perception, or the subjective perceived distance to objects around one's self, is shorter in virtual environments than in real ones [Renner et al. 2013]. Factors like vergence and accommodation [Watt et al. 2005], the quality of graphics [Knapp and Loomis 2003], and even the weight and inertia of the HMD itself have been deemed to contribute to this effect. Ocular parallax rendering could reduce the underestimation by minimizing disparity distortions and providing additional depth information through micro parallax and gaze-contingent occlusions. We investigate participants' egocentric depth estimation in their action or personal space (i.e. within arm's reach), where the ocular parallax effect is strongest. Because verbal estimates have been shown to be variable [Napieralski et al. 2011], we rely on a blind reaching task where users view a target and then, with the display turned off, reach to where it was last seen. The blind reaching task is the standard way to evaluate egocentric distance estimation for objects within arm's reach [Altenhoff et al. 2012].

Sixteen volunteers participated in the study (age range 22-32, 5 females), of which one was excluded for not passing a standard Randot stereo vision test, and two others were excluded due to the eye tracker failing to track their pupils.

Stimuli and study setup. To minimize any potential effects of graphics quality on depth judgment [Knapp and Loomis 2003], the viewing environment for the blind reaching task consisted of a photo-realistic dining room scene as seen in Figure 5. A floating pencil served as the reaching target and was surrounded by objects at different depths and a feature rich background, emphasizing the ocular parallax effect.

The HTC Vive Pro with Pupil Labs eye trackers served as the HMD; 6-DOF head pose tracking was enabled. The hand's position was tracked via an HTC Vive tracker mounted on an optical post that was held upright (Figure 5, top right). To facilitate natural reaching, participants held the tracker in their dominant hand and an HTC Controller in their other hand for interaction.

Conditions. Each participant performed the blind reaching task with ocular parallax enabled and disabled. To maintain a common reach metric between participants, reach target distances were set proportionally to each participants' arm reach. The reach targets appeared at 50, 58, 67, 75, 82, and 90% of the participant's maximum arm reach. The participants completed trials under each of the 12 conditions 5 times and the presentation order of the 60 trials was randomized.

Procedure. Each participant performed an ipd, eye tracker, and maximum arm reach calibration procedure at the beginning of the session. To determine their maximum arm reach, participants were instructed to fully extend their arm directly in front of their head. The distance between the eyes' midpoint and the hand defines the maximum arm reach.

A standard blind reaching task was then conducted [Altenhoff et al. 2012]. Participants first observed the viewing environment to familiarize themselves with its size and object depths. They then began the trials. Each participant performed two practice trials followed by 60 recorded distance estimates with each trial consisting of a target presentation phase and a blind reaching phase. In the final presentation phase, the pencil appeared in front of the users and they were instructed to determine its depth by looking at it, but also to gaze to other scene objects to maximize eye rotation and the ocular parallax effect. Participants were required to view the environment for a minimum of 5 seconds before they were given the ability to turn off the display when they were ready to perform the blind reach. With the display off, they were instructed to reach for the pencil with only their arm while their torso remained steady. They progressed onto the next trial by pulling the controller's trigger when satisfied with their reach position. All head and hand pose tracking data, as well as the pencil's position, were recorded during the blind reaching phase.

Analysis. The reach estimate was computed as the average of the last 10 samples of the participant's reach data. Because 6-DOF head tracking was enabled, natural head movement resulted in deviation from the the head to target distances reported above. To account for this, the data was analyzed along the horizontal 2D vector between the eyes' midpoint and the target's position. Both the target's distance and reach distance were computed along this vector as a percentage of the arm's maximum reach.

Results. Enabling ocular parallax rendering had little to no effect on the egocentric depth perception of the participants. Figure 5 shows each participant's reach, as a percentage of their maximum arm reach, to a presented target distance for the ocular parallax enabled and disabled conditions. Linear models were fit to the two sets of reaches, finding that the slopes for the ocular parallax enabled and disabled modes were 1.141 and 1.128, respectively, while the intercepts were 0.167 and 0.165, respectively. A multiple regression analysis did not show a significant difference between the two conditions.




Fig. 5. Egocentric depth perception. This study investigates whether the additional depth cues in ocular parallax aid egocentric depth perception. The viewing environment including the reach target, a pencil, are shown in the top left. Users viewed the target and then, with the display turned off, reached to it as seen in the top right. Participants' reaches in the ocular parallax enabled and disabled modes are found in the bottom plot.

REFERENCES

- Bliss M. Altenhoff, Phillip E. Napieralski, Lindsay O. Long, Jeffrey W. Bertrand, Christopher C. Pagano, Sabarish V. Babu, and Timothy A. Davis. 2012. Effects of Calibration to Visual and Haptic Feedback on Near-field Depth Perception in an Immersive Virtual Environment. In *Proceedings of the ACM Symposium on Applied Perception (SAP '12)*. ACM, New York, NY, USA, 71–78.
- Geoffrey P. Bingham. 1993. Optical flow from eye movement with head immobilized: “Ocular occlusion” beyond the nose. *Vision Research* 33, 5 (1993), 777 – 789.
- M. Fendick and G. Westheimer. 1983. Effects of practice and the separation of test targets on foveal and peripheral stereoacuity. *Vision Research* 23, 2 (1983), 145 – 150.
- Joshua Knapp and Jack Loomis. 2003. *Visual Perception of Egocentric Distance in Real and Virtual Environments*. Vol. 11. 21–46.
- Suzanne P. McKee. 1983. The spatial requirements for fine stereoacuity. *Vision Research* 23, 2 (1983), 191 – 198.
- Phillip E. Napieralski, Bliss M. Altenhoff, Jeffrey W. Bertrand, Lindsay O. Long, Sabarish V. Babu, Christopher C. Pagano, Justin Kern, and Timothy A. Davis. 2011. Near-field Distance Perception in Real and Virtual Environments Using Both Verbal and Action Responses. *ACM Trans. Appl. Percept.* 8, 3, Article 18 (Aug. 2011), 19 pages.
- Samuel C. Rawlings and T. Shipley. 1969. Stereoscopic Acuity and Horizontal Angular Distance from Fixation*. *J. Opt. Soc. Am.* 59, 8 (Aug 1969), 991–993.
- Rebekka S. Renner, Boris M. Velichkovsky, and Jens R. Helmert. 2013. The Perception of Egocentric Distances in Virtual Environments - A Review. *ACM Comput. Surv.* 46, 2 (2013), 23:1–23:40.
- S. J. Watt, K. Akeley, M. O. Ernst, and M. S. Banks. 2005. Focus cues affect perceived depth. *Journal of Vision* 5, 10 (2005), 834–862.