

Manufacturing Application-Driven Foveated Near-Eye Displays

Kaan Akşit, Praneeth Chakravarthula, Kishore Rathinavel, Youngmo Jeong, Rachel Albert, Henry Fuchs, and David Luebke

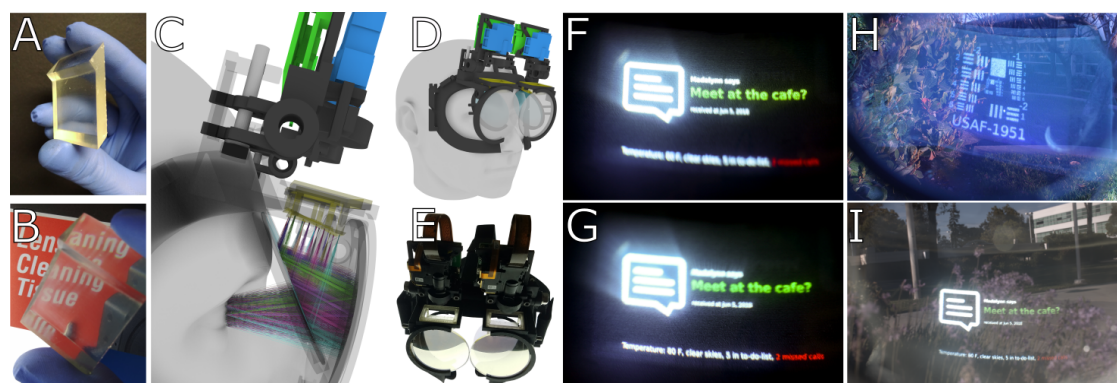


Fig. 1. Manufacturing freeform optical components is a slow, expensive, and labor-intensive task that restricts computational optical designers to only existing components which may be sub-optimal, and limits their ability to iterate and explore new designs. (A) we present a rapid manufacturing technique involving 3D printing, optical bonding, and vacuum forming as a prototyping tool for testing and validating ideas before creating a final product. (B) A sample optical waveguide built using our rapid manufacturing technique shows that the properties of refraction and total internal reflection match our theoretical expectations with negligible deviations. (C) We choose the optically challenging scenario of near-eye displays to evaluate both our manufacturing technique, and our computational methodology for calculating interchangeable freeform surfaces for a given near-eye display optical layout. (D-E) We design and manufacture a completely untethered near-eye display prototype. (F-G) We show that our proposal can provide optically correct depth cues at various scenarios. In this particular case, bottom portion of the field of view is at a closer optical distance with respect to rest. (H-I) Our prototype is also bright enough to be used in direct sunlight conditions, with sufficient resolution for rendering text.

Abstract—Traditional optical manufacturing poses a great challenge to near-eye display designers due to large lead times in the order of multiple weeks, limiting the abilities of optical designers to iterate fast and explore beyond conventional designs. We present a complete near-eye display manufacturing pipeline with a day lead time using commodity hardware. Our novel manufacturing pipeline consists of several innovations including a rapid production technique to improve surface of a 3D printed component to optical quality suitable for near-eye display application, a computational design methodology using machine learning and ray tracing to create freeform static projection screen surfaces for near-eye displays that can represent arbitrary focal surfaces, and a custom projection lens design that distributes pixels non-uniformly for a foveated near-eye display hardware design candidate. We have demonstrated untethered augmented reality near-eye display prototypes to assess success of our technique, and show that a ski-goggles form factor, a large monocular field of view ($30^\circ \times 55^\circ$), and a resolution of 12 cycles per degree can be achieved.

Index Terms—Near-eye displays, See-through Displays, Application Adaptive Displays, Computational Displays, Augmented Reality Displays, 3D printed optical components, Waveguides, projection displays

1 INTRODUCTION

Augmented reality (AR) near-eye displays (NEDs) promise to be the next breakthrough mobile platform, providing a gateway to countless AR applications that will improve our day-to-day lives [80]. To fulfill the promise of immersive and natural-looking scenes, as described by Kress and Sterner [40], AR NED designers need to solve difficult optical design challenges including providing sufficient resolution levels, eyebox, and field of view (FoV). A major impediment to achieving natural images and a key cause of discomfort is the vergence-accommodation conflict (VAC) [27], which is caused by a mismatch between the binocular disparity of a stereoscopic image and the optical

focus cues provided by the display. Mainstream strategies [33] for tackling these challenges involve dynamic display mechanisms that can generate accurate visuals in all possible optical depths, which greatly increases the complexity of the NED design problem. Other obstacles to widespread adoption of AR NEDs include providing price-wise accessibility, requiring a reasonable amount of computation and power, and providing a thin and a light-weight form factor suitable for daily use. All of these problems are still waiting to be addressed and even small steps towards a possible solution requires large lead time, and a massive effort in place.

In this work, we explore 3D printing as an option for rapid production of optical components. To target solutions for near-eye display designs, inexpensive interchangeable optical components can therefore be made into arbitrary shapes using commodity 3D printers. This opens a path towards building much simpler optical AR NED designs driven by the needs of specific AR applications. Our work lead to the following contributions:

- Kaan Akşit, Youngmo Jeong, Rachel Albert, David Luebke are with Nvidia E-mail: kaksit@nvidia.com.
- Kishore Rathinavel, Praneeth Chakravarthula, Henry Fuchs are with UNC Chapel Hill E-mail: fuchs@cs.unc.edu.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxxx

- We introduce a rapid manufacturing technique based on 3D printing, optical bonding, and vacuum forming. We show that complex optical components such as diffusers, lenses, and optical waveguides can be built for development purposes in a day using our

	Focus mechanism	See-through	FoV	resolution	eyebow	form factor	compute overhead	gaze-tracking
Pinlight displays [51]	light fields	yes	wide	low	small	thin	high	no
Freeform optics [34]	light fields	yes	narrow	moderate	moderate	moderate	high	no
HOE [37]	light fields	yes	moderate	low	large	moderate	high	yes
HOE [50]	holographic	yes	wide	moderate	small	N/A	high	yes
Focus tunable light engine [48]	varifocal	yes	narrow	moderate	small	bulky	moderate	yes
Multi-focal plane display [32]	varifocal	yes	narrow	moderate	moderate	bulky	high	yes
Membrane [24]	varifocal	yes	wide	low	large	bulky	low	yes
Varifocal HOE [1]	varifocal	yes	wide	moderate	large	moderate	low	yes
Multi-Focal Display [44]	multi-focal	yes	narrow	low	large	thin	high	no
Focal Surface Display [52]	focal Surface	no	narrow	moderate	narrow	moderate	high	no
This work	static focal surface	yes	wide	moderate	large	moderate	low	no

Table 1. A comparison of see-through accommodation supporting near-eye displays, including the virtual reality near-eye display implementation of [52]. This table is modeled after those in Dunn et al. [24], Akşit et al. [1] and Matsuda et al. [52]. Our prototype demonstrates a unique combination of a good form factor, resolution, FoV, and eyebow. Note that, in our chart, a moderate FoV is define as 40 - 60°, moderate resolution is defined as 10 - 20 cpd, and a moderate eyebow is defined as 5 - 10 mm. Moderate values are adapted from [15,52].

sanding-free technique without requiring 6 – 8 weeks of lead times, large investments, or intensive labor,

- We provide a computational optical design methodology to create freeform projection screens for various applications, accelerated by training a machine learning model. We show that the exit pupils in our designs can either be gaze-adaptive or fixed,
- To understand the demands of various AR applications, we derive depth characteristics for real-life scenes based on egocentric RGB-D datasets that are publicly available and from our own additional data collection, we identify suitable applications accordingly,
- We show a completely untethered application-adaptive AR NED prototype with interchangeable optics, ski goggle form factor, wide FoV, high brightness, and small computational and power demands,
- We design and manufacture a custom projection lens. We steer the projection lens in front of a projector. We show that a foveated display design candidate can be built by requiring only one display mechanism .

Overall, our proposal provides ease of replication, customization, significant cost reduction in production, practicality, and freedom of design. Our proposal avoids the use of computation-intensive elements commonly found in other designs elsewhere in the literature, while tackling a challenging optical scenario of near-eye display designs.

2 RELATED WORK

Our proposal promises a design and manufacturing methodology for building an accommodation supporting foveated AR NED using 3D printed optical components. 3D printing is a process of manufacturing 3D physical objects through various additive manufacturing (AM) techniques. We refer curious readers to Vaezi et al. [79] for a review on various AM techniques, and Bickel et al. [9] for a review on the state of the art on stylized fabrication. In this section, we provide a comprehensive review on 3D printed optics used for displays, along with a review on recent accommodation supporting NEDs, and curved displays.

2.1 Printed optics for displays

The work of Willis et al. [88] investigated the feasibility of 3D printed optical components for sensing, illumination, and displays, and also explored the possibility of using a per-pixel optical waveguide to relay light from an emissive display to a light out-coupling planar surface, achieving dot-shaped pixels 1.2 mm in size. The per-pixel optical waveguides discussed in the work of Willis et al. [88] refer to an array of optical waveguides, each dedicated to delivering information carried by a single light beam to a single location to form one particular pixel on a screen surface. The work of Brockmeyer et al. [11] extended the approach of 3D printing per-pixel optical waveguides to curved surfaces, with their prototype providing pixels at a millimeter scale

over a spherical display surface. Recent work by Pereira et al. [65] proposed an algorithm that maximizes the optical efficiency of light going through 3D printed per-pixel optical waveguides, showing that more complex display surfaces with millimeter scale pixels can be 3D printed. Unlike conventional fixed pitched lenslet arrays, the work of Tompkin et al. [78] proposed custom 3D printed content-adaptive lenslet arrays for glasses-free 3D display applications, reporting pixels of size 0.4×2.5 mm. Papas et al. [61] built a multi-view display using custom 3D printed lenslet arrays and unstructured static printed patterns. Other recent fabrication methods have focused on the surfaces of 3D printed objects, including reproducing specific characteristics of reflectance [53, 85] or refraction [62]. However, existing solutions for display applications of printed optics do not produce sufficiently small pixel sizes (micron scale) to support a practical NED design. Complementary to printing optics for a NED design, the work of Mueller et al. [57] describes a rapid prototyping of functional NED housing by combining Lego bricks and 3D printed parts.

Our manufacturing technique moves away from per-pixel waveguides. In our design, a single optical waveguide projects micron scale pixels onto a complex diffusive display surface rather than millimeter scaled ones found in the literature. Our process uses only a single 3D printing medium, whereas previous 3D printed waveguides required one medium for the core and one medium for cladding. To the best of our knowledge, our work is the first to target NED applications using 3D printed optical components.

2.2 Accommodation supporting near-eye displays

Accommodation supporting NEDs [33] address the problem of VAC by matching the binocular disparity of virtual objects with correct optical focal cues for various depth planes.

Multi-Plane Displays Early on, Akeley et al. [2] demonstrated the benefits of a fixed-viewpoint volumetric desktop display using flat multi-planes, which allowed them to generate near-correct focus cues without tracking the eye location or gaze. Recently such displays have been revisited with improved scene decomposition and gaze-contingent varifocal multi-plane capabilities [54, 58]. However, such displays have large power and computational demands, and complex hardware that would be difficult to miniaturize. These constraints limit their usefulness to perceptual experiments identifying the needs of future NED designs. The work of Hu et al. [32] demonstrated a time multiplexed multi-plane display in the form of a wearable AR NED with a narrow field of view ($30^\circ \times 40^\circ$). Lee et al. [44] proposed a compact multi-plane AR NED composed of a waveguide and a holographic lens, which demonstrated a FoV of $38^\circ \times 19^\circ$. Zhang et al. [93] proposed a stack of switchable geometric phase lenses to create a multi-plane additive light field VR NED, providing approximate focus cues over an 80° FoV. Both the works of Lee et al. [45] and Hu et al. [32] demonstrated time multiplexed multi-plane AR NEDs with FoVs of 30° to 40° , respectively.

Light Field Displays Light field NEDs promise nearly correct optical accommodative cues, but this comes at the cost of significant

resolution loss. Lanman and Luebke [43] introduced a VR Near-Eye Light Field Display (NELD) that uses microlenses as the relay optics, showing a prototype with a FoV of $29.2^\circ \times 16.0^\circ$, leading to a resolution of 2–3 cpd. Huang et al. [35] developed VR NELDs further, demonstrating a prototype with a diagonal binocular FoV of 110° , leading to a resolution of (3–4 cpd). The work of Akşit et al. [3] created a VR NED using a pinhole mask in front of an AMOLED display, and demonstrated full-color images with a diagonal binocular FoV of 83° with 2–3 cpd. By using a see-through sparse backlight mechanism, the work of Maimone et al. [51] introduced a single-color AR NED prototype with a diagonal FoV of 110° , and a resolution of 2–3 cpd.

Varifocal Displays Another solution for solving VAC is a varifocal display, which dynamically changes the optical properties of the display. Although varifocal displays offer large computation benefits, they require precise gaze tracking. Liu et al. [48] used a tunable lens system combined with a spherical mirror, demonstrating 28° of diagonal FoV with 10–14 cpd, which switches depth from 8 D to infinity ($\sim 0.1D$) within 74 ms. Another study by Konrad et al. [66] also took advantage of an electrically tunable lens system, and demonstrated 36° diagonal FoV. The solution of Konrad et al. allowed depth switching from 10 D to infinity ($\sim 0.1D$) within 15 ms, and provided 5–6 cpd. Dunn et al. [24] provided a monocular FoV of $> 60^\circ$ and a fast varifocal mechanism of 300 ms that switches depth from 5 D to infinity ($\sim 0.1D$). Most recently, the work of Akşit et al. [1] proposed using Holographic Optical Elements (HOEs) as a part of an AR varifocal NED system, offering a FoV of 60° with 18 cpd, however the varifocal mechanism is not fast enough (410 ms) when switching from 5 D to infinity ($\sim 0.1D$).

Static and Dynamic Holographic NEDs Holography promises an accurate representation of four-dimensional (4D) light fields, however the limitations of such displays include a small eyebox, large computational demand, large calibration times, and the design trade-off between limited resolution or a bulky form factor. Static holograms encoded into HOEs have been used in various NED types as optical combiners [37, 44, 50] or projection surfaces [1], although the static holograms in these displays do not provide 4D light fields. On the other hand, dynamic holographic VR NEDs can be achieved using phase-only Spatial Light Modulators (SLMs) which can encode holograms [50, 52, 72], enabling a glasses-like form factor, and a wide FoV (80°).

Inspired by the recent static and dynamic holographic NEDs [1, 24, 50, 52, 89], we propose a ski-goggles form factor AR NED design with wide FoV and application-dependent accommodation support. In Table 1, we also provide a comparison of the characteristics of other state-of-the-art NED prototypes found in the literature.

2.3 Curved and Freeform Screens

Our proposal promises a diffusive freeform projection screen surface that can mimic the depth properties of a target depth curvature. Therefore, we look into the literature to identify the state-of-the-art incurved or freeform screens. Researchers have explored desktop-sized static curved displays [8, 10, 31, 41, 84] and very large format, immersive, static curved displays [7, 28, 39, 77]. These curved displays are typically cylindrical or spherical in their surface profile. The work of Brockmeyer et al. [11] demonstrated a static desktop-sized display which achieves a freeform surface profile, inspiring our work. Researchers have also shown manually configurable flexible displays that use Organic LEDs [90], thin electroluminescent films [59], and electronic-ink [25]. Recently, a dynamically shaped changing display was demonstrated by Leithinger et al. [47]. For a more exhaustive survey on non-planar displays, we refer interested readers to the following papers: [4, 46, 67]. To the best of our knowledge, we are the first to claim usage of 3D printing to create freeform projection screen surfaces for AR NED designs.

Relevant to our work, previous body of work on caustic based lens design [21, 70, 92] demonstrated freeform lens calculation process using a single collimated light sources as an input, high precision expensive machinery, and among with large optical paths. Our design process differs as it takes advantage from machine learning, uses multiple

spherical light sources as inputs (a display with many pixels), and relies on our fast-pace cost-effective manufacturing process.

2.4 Foveated Displays

In our proposal, we evaluate 3D printed optical components that can project foveated images, and we also explore the steering of projected images on a diffusive projection screen for foveation purposes. Therefore, we review relevant optical hardware in the foveated display literature. The earliest example of a gaze-contingent visual stimulus was presented by Reder in 1973 [68], paving the way for further research into foveated imagery. Later on, the first proposal for foveated display hardware appeared in the works of Baldwin et al. [5] as a variable resolution transparency magnified by large concave mirrors. A year later, the work of Spooner et al. [73] showed another style of desktop size foveated display hardware, which combines two different displays to provide high resolution images at the fovea, and low resolution images in the periphery. To our knowledge, the work of Shenker et al. [71] is the first to realize the concept of combining two different displays in a NED configuration, in the form of a steerable foveal inset with 20 cpd resolution created using fiber-optics and pancake type optical relays. Later, the work of Howlett et al. [30] followed the path of combining two different displays in a NED configuration to build a complete telepresence system with cameras. Rolland et al. [69] combined two displays using a beam-splitter in a NED setting. In their design, a high-resolution inset with 24 cpd resolution is relayed to the fovea of the eye using microlenses with a FoV of $13.30^\circ \times 10.05^\circ$, while a lower resolution display at 6 cpd spans across a FoV of $50^\circ \times 39^\circ$ through a magnifier lens. The work of Godin et al. [26] describes a dual projector layout in order to realize a stereoscopic desktop-sized display with a fixed foveal region. Most recently, Lee et al. [44] proposed a compact AR NED comprised of a waveguide and a holographic lens that combines two displays. Their design has a FoV of $38^\circ \times 19^\circ$ and eliminates the needs of a gaze tracking hardware. We refer curious readers the following set of papers for detailed perceptual and computational benefits of foveation in computer graphics: [38, 63, 64].

Rather than combining two different displays into one to provide foveated imagery, our proposal experiments with steering a single display per eye, where each display is equipped with a custom foveated lens design and linear mechanical stages. Our mechanism for steering is based on mechanical actuators, resulting in increased power demands and slower steering due to the extra bulk. On the other hand, our proposal promises less computational demand, fewer design trade-offs with respect to other display hardware found in the NED literature, and also promises to avoid the perceptual issues that arise from the integrity and alignment of dual display per eye hardware.

3 APPLICATION-DRIVEN NEAR-EYE DISPLAYS

We seek to target the requirements of the human visual system (HVS). The HVS has a binocular field of view of $\sim 190^\circ$ [83]. Maximal human visual acuity under ideal conditions is ~ 65 cpd at the fovea [76, 86, 87], although it quickly drops to about 2.5 cpd at 35° eccentricity in the periphery [75]. Moreover, human depth acuity is similarly at its peak at the fovea and drops linearly with greater eccentricity [19].

Towards this goal, we have landed on an optical design for AR applications as shown in Figure 2. Our base optical design consists of a specialized light engine, a freeform projection surface producing a target depth plane that is determined by the desired AR application, and an off-axis semi-reflective curved beam combiner. We explain our optical layout by following the light path, starting from the projector. We also discuss our analysis of the depth levels in real-life scenes that guided our design, and explore an extension of our work towards custom projection optics for foveation purposes.

3.1 Projection light engines

The freeform projection surfaces in our proposal are illuminated using projection displays, referred to as light engines. A simplified equivalent of a light engine is an emissive pixelated display coupled with projection optics. Projection optics generate a field curvature [81], causing each pixel of the emissive display to be imaged with sharpest focus at a




Fig. 2. The optical layout for our Augmented Reality Near-Eye Display with freeform projection surfaces. We computationally optimize the shape of the freeform projection surface for a given target depth curvature. Images on the freeform projection surface are relayed to a viewer’s eye using an off-axis curved beam combiner.

specific distance, thereby forming a 3D surface. The field curvature of a complex optical system with $m + 1$ surfaces can be described as

$$\frac{1}{R_p} = \sum_{i=0}^m \frac{n_{i+1} - n_i}{r_i n_{i+1} n_i}, \quad (1)$$

where n_i represents the refractive index and r_i represents the radius of curvature for the i^{th} surface. For a simple thin lens, the field curvature can be calculated as

$$R_p = n f_{\text{thin}}, \quad (2)$$

Where f_{thin} represents the focal length of a thin lens, and n represents the refractive index of a thin lens material. However, the field curvature of a light engine can also be designed in various ways: to achieve a flat surface, telecentric lenses [82], separated thin lenses, field-flattener lenses, or meniscus lenses can be used [81]. The point spread function (PSF) describes the response of an imaging system to a point source. Assuming a circular lens or an effective circular aperture with coherent illumination, the PSF of a light engine is theoretically limited in spot size, which can be described using the diffraction-limited Rayleigh’s resolution criterion,

$$w_0 = \frac{2.44 \lambda_i f}{D_L}, \quad (3)$$

where f is the effective focal length or throw distance, D_L is the effective aperture size, and λ_i represents the wavelength of the light. A pixel on an emissive display has an active area that emits light, called pixel spread, which is typically one of the limiting factors determining the maximum resolution of a NED design. Our design may produce variable pixel spread over the freeform projection surface due to discrepancies between the surface height and the field curvature of the projection optics. Light engines that are based on focus-sweeping lenses [36], or scanning lasers [29] provide an improved constant spot size over a wider range; however, such solutions typically come with trade offs such as loss of contrast for sweeping lenses or limited minimum spot sizes due to maximum scanning mirror sizes. Our optical layout proposal is compatible with those alternative solutions as well.

3.2 Optical Path folding

As our design relies on freeform projection surfaces, placing the light engines directly behind the curved projection surface would lead to both an increased eye relief and a larger form factor, although in some cases this may not even be physically possible due to the anthropometry of the average human. Therefore, one can use free-space beam splitters, mirrors, or optical waveguides to fold the optical projection path. Both the waveguide and beam-splitter approaches are applicable to our optical layout, and the physical arrangement of both designs are depicted in Figure 3.




Fig. 3. There are two options for optical path folding in our proposal. Sketches show (Left) an optical waveguide design merged with our freeform projection surfaces, (right) traditional birdbath optics with beam-splitters.

For the waveguide approach, the light rays emitted from the light engine arrive at a prismatic interface S_0 which has a tilt of Θ_w , as annotated in Figure 3. Rays at the interface of S_0 refract and propagate inside the medium. At the interface S_1 , rays arriving with an angle of incidence larger than the critical angle tend to reflect as dictated by Fresnel coefficients, and causing a phenomena of Total Internal Reflection (TIR). Reflected rays from S_1 arrive at the final projection freeform surface annotated with S_2 in Figure 3. Curious readers may find a review on applications of waveguides in NED designs in the work of Cameron et al. [16].

3.3 Beam Combiners

Our design use standard spherical beam combiners in conjunction with more complex freeform projection surfaces to avoid complex beam combiner designs or multi-element optics directly in front of the display. In our assessments, we observed that a larger radius of curvature for a spherical beam combiner leads to larger eye relief, and therefore also a larger projection volume A , as depicted in Figure 2. To avoid strict tolerances in manufacturing we recommend curvatures in the range of 40 – 80 mm, which maximizes the volume without falling below a thickness of 3 mm or exceeding a thicknesses of 10 mm. Our proposal is applicable for other forms of off-axis curved beam combiner surfaces [14].

3.4 Freeform Projection Surfaces

Using the ray tracing model in Figure 2, we investigate on the required shape of the projection screen surface using a fixed shape off-axis beam combiner. We simulate rays starting from sample points at various target depth curvatures spanning different optical depths, $P_N(x, y, z)$, $P_M(x, y, z)$, $P_F(x, y, z)$, reaching to a planar exit pupil with a size of

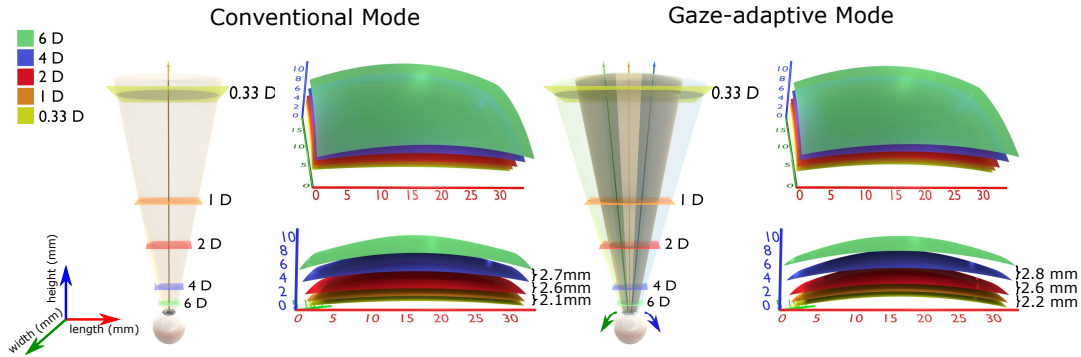


Fig. 4. Calculated freeform projection surfaces for various depth ranges in conventional and gaze-adaptive mode. These surfaces are designed for our wearable prototype using a spherical beam combiner with a 76.4 mm radius of curvature. In the conventional mode we optimize these surfaces for a fixed nominal gaze condition, whereas in the gaze-adaptive mode we optimize points on each projection surface by rotating the eye to points of interest in the visual field of view.

20 mm × 20 mm. A sample ray R_0 shows this path in Figure 2. From the exit pupil, we trace target rays backwards to the off-axis beam combiner, which has a tilt of θ_{BC} with respect to the optical axis as shown in Figure 2. The tilt is chosen such that the reflection of the target rays are on or above the typical height of the eyebrows [91]. An example reflected ray is annotated as R_1 in Figure 2.

To determine the location of the freeform projection surface we find the 3D point locations, where the reflected target rays are closest to each other, forming the smallest possible spot size for a given set of rays. Given the the starting sample points $P_N(x, y, z)$, $P_M(x, y, z)$, $P_F(x, y, z)$, the point candidates on the freeform projection surface are therefore $P'_N(x, y, z)$, $P'_M(x, y, z)$, $P'_F(x, y, z)$. We restrict the eye relief to a range of 35 – 55 mm to accommodate assembly using either a direct optical path or an optical path folding method.

The optics of a NED typically contain aberrations that cause angular deviations of ray bundles as a function of changes in gaze and eye position. Recently, Mercier et al. [54] pointed out that the alignment of focus planes can be affected by gaze changes due to the optics of a multiplane NED. On the other hand, the most recent literature on 3D foveated desktop displays with accommodation support [38, 74] suggests that the fovea is the area most sensitive to depth differences. In light of these recent findings, we expand the concept of “pupil-swim” to our optical design methodology by approximating the freeform projection surface in a gaze adaptive way. Rather than optimizing the entire surface for a single gaze condition at once, Figure 2 shows how we keep the rotation center P_e fixed in space and rerun our ray tracing model with the target depth point determined by the gaze condition. Note that the location of P_e can be personalized depending on the interpupillary distance of the subject [23] or the way the NED fits on the subject’s head. For these simulations, we accommodate the worst-case scenario by using an exit pupil of 8 mm × 8 mm, which is the largest possible aperture of the human eye under scotopic (dark-adapted) conditions [22]. In Figure 4, we simulate different required surfaces using both classical and gaze adaptive optical design methodologies to represent several rectangular planar virtual images at various depth levels. Therefore, it is also possible to generate freeform surfaces based on an eye location of an user.

To accelerate calculations of freeform projection surfaces, we identify an alternative regression-based model. Using a small number of point candidates on a freeform projection surface $P'_N(x, y, z)$, $P'_M(x, y, z)$, $P'_F(x, y, z)$ and their corresponding target depth points $P_N(x, y, z)$, $P_M(x, y, z)$, $P_F(x, y, z)$, we trained a Gaussian Process Regression (GPR) model that follows the general form

$$\begin{bmatrix} P'_x \\ P'_y \\ P'_z \end{bmatrix} = k^T C^{-1} \begin{bmatrix} P_x \\ P_y \\ P_z \end{bmatrix}^T, \quad (4)$$

where P'_x , P'_y , and P'_z represent an estimated projection surface point

for a given target depth point represented with p_x , p_y and p_z , C^{-1} corresponds to the inverse of the covariance matrix for a provided training set ($P'_N(x, y, z)$, $P'_M(x, y, z)$, $P'_F(x, y, z)$, $P_N(x, y, z)$, $P_M(x, y, z)$, $P_F(x, y, z)$), k^T represents the transpose of a vector that contains the similarity measures between depth points used in training $P_N(x, y, z)$, $P_M(x, y, z)$, $P_F(x, y, z)$ and a given target depth point represented with p_x , p_y and p_z . The similarity measure that we used in our training is a generalized Radial Basis Function (RBF) kernel of the following form:

$$\kappa(a, b) = e^{-\frac{\|a-b\|}{2\sigma}}, \quad (5)$$

in which σ represents a free parameter. We tune σ until we obtain a good estimate for our test data. Note that similarity measures vector calculated for each estimation, k^T , and the inverse of the covariance matrix for a given training dataset, C^{-1} are both based on the introduced RBF kernel, $\kappa(a, b)$. Using this regression model, we generate estimates for projection surfaces represented in the form of dense point clouds rather than running a ray tracing-based optimization for each target point in the space, which also saves a designer from a hassle of going into porting code to GPU as most of the modern machine learning libraries are already GPU accelerated.

3.5 Target Depth Fields

To determine the optimal freeform projection surface for various augmented reality contexts, we explore depth characteristics of real-life scenes. This kind of analysis can illuminate the accommodative demands required in various augmented reality use cases. We therefore analyze several scenarios such as “Walking indoors”, “Social interaction”, “Working at a desk”, and “Office workspace”.

We compute the average depth maps for a number of scenes as shown in Fig 5. For the first three columns, we use egocentric RGB-D datasets provided by FIIT Egocentric RGB-D dataset [60], NUS3D-Saliency Dataset [42], and EgoDexter [56]. Additionally, we capture our own egocentric RGB-D dataset for the office workspace case. Our dataset was collected using an Xtion Pro Carmine RGB-D sensor with the OpenNI and OpenCV APIs.

In our calculations, we take into account the depth sensitivity of the HVS, such that the peak depth sensitivity is at the fovea, with a linear angular fall-off rate into the periphery [19]. The average depth map is therefore based on the following formulation:

$$c(i, j) = p - s \times d(i, j), \quad (6)$$

where the visual depth sensitivity function c is defined in terms of peak depth sensitivity at the fovea p , the angular linear fall-off in depth-sensitivity s , and the angular distance $d(i, j)$ from the fovea. We define a weighting function w as:

$$w = g * c, \quad (7)$$

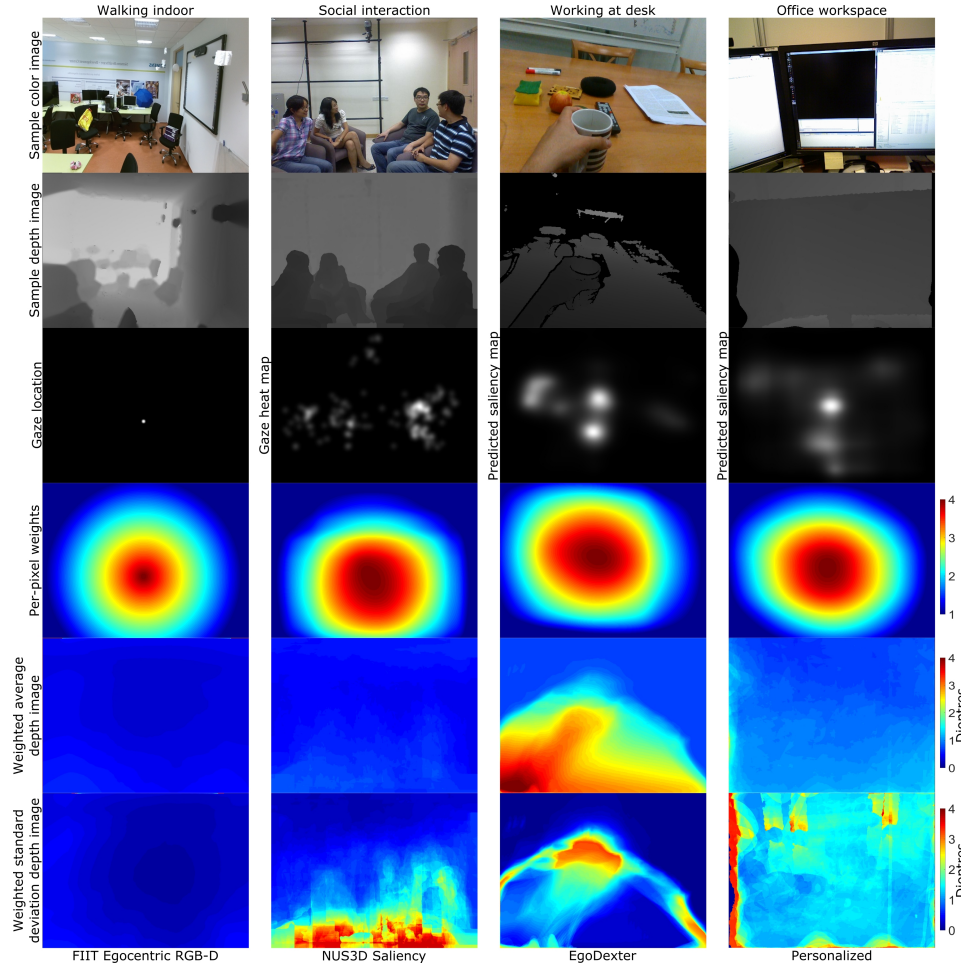


Fig. 5. The process of generating average depth maps from a dataset of depth images. Labels at the bottom of each column indicate the dataset used. Labels at the top of each column indicate the application or situation in which the dataset was collected. Datasets used in the first three columns are FIIT Egocentric RGB-D dataset [60], NUS3D-Saliency Dataset [42], and EgoDexter [56].

where c represents visual depth sensitivity, and g represents gaze locations or gaze heat maps. If gaze information is not available in a given dataset, we rely on saliency prediction maps generated using code provided by [20]. Finally, we arrive at weighted average for a given dataset as:

$$\bar{p} = \frac{\sum w_i \cdot p_i}{\sum w_i}, \quad (8)$$

and we calculate a weighted standard deviation for a given dataset as

$$\sigma = \sqrt{\frac{\sum w_i \cdot (p_i - \bar{p})^2}{\sum w_i}}. \quad (9)$$

The Depth of Field (DoF) of the HVS provides a depth volume inside which the scene is acceptably sharp. The size of this volume can be 0.4 D to 1.2 D for eye aperture diameters in the range of 8 mm to 2 mm, respectively [17, 18]. The average standard deviations for the datasets in Figure 5 suggest that the activity of walking indoors has a variation that falls within this DoF range. The social interaction activity poses a challenging scenario for our approach in the bottom portion of FoV, but a large portion of the scene could still be static in depth. The other two activities, working at a desk and the office workspace, include more close distance interactions and therefore may require two or three static surfaces to be optically combined in order to provide the required accommodation support.




Fig. 6. A foveated lens can distribute pixels of a display in a non-uniform fashion by intentionally introducing pinclusion distortions to the generated final images.

3.6 Foveated projection optics

We also explore projection optics as one possible extension of our system. Foveated displays present a higher resolution image at the location of the user's gaze and progressively less detail with increasing eccentricity, reducing the total number of pixels required. We use projection optics to distribute pixels in a non-uniform fashion over a projected image. In a typical light engine, the pixel locations are distributed evenly in the focus plane and can be expressed as a normalized distribution from -1 to 1 . To create a foveated light engine, we apply an exponent to the evenly spaced distribution (i.e. $\{-1 \dots 0 \dots 1\}^2$), obtaining the best possible PSF at the center $(0, 0)$ with fewer pixels in the periphery, as depicted in Figure 6, very similar to Pinclusion distortions. Our intended design is a single add-on lens in front of a projection light

engine, we restrict our search to solutions with two surfaces as part of a single lens to maintain simplicity and practicality in our final design, while a multi-surface (> 2) system has the potential to provide ease in design process.




Fig. 7. A series of photographs showing optical components built using our manufacturing technique based on 3D printing, optical bonding (OP), and vacuum forming (VF). Photographs of sample parts produced by wet-sanding are provided. With respect to wet-sanding (WS) approach, note that our technique promises least amount of labor work for prototyping 3D printed optics.

4 MANUFACTURING AND IMPLEMENTATION

We provide the details of our manufacturing technique, and demonstrate our NED prototypes accordingly. Additional photographs and videos of our development are also available in our supplementary material.

4.1 Printing optical components

A sub-branch of AM techniques, Stereolithography (SL), has recently garnered interest. SL relies on shining a narrow UV light beam on a vat of liquid photopolymer, followed by curing the polymer to build each layer of the 3D structure in micron resolution (typically 0.25 – 100 μm). There are optically clear liquid photopolymers compatible with SL 3D printers (VeroClear, Formlabs Clear resin), which has similar optical characteristics with Poly(methyl methacrylate) (PMMA). 3D printing optical components with such optically clear liquid photopolymers opens up interesting possibilities, and provides a cheap and rapid manufacturing alternative to optical designers, as opposed to several weeks of a single design iteration. We push the boundaries of 3D printing based manufacturing with optical designs based on refraction, and TIR principles. Using a Formlabs 2 3D printer with a clear resin (Formlabs FLGCL02), We investigate manufacturing lenses, prisms, and waveguides. Moving forward, we will explain the evaluation of our technique by starting from known technique of wet sanding, and later we will show how we save time, and effort by avoiding polishing or sanding with the help of vacuum formers.




Fig. 8. Two photographs showing a view of a mobile phone display magnified by (Left) an off-the-shelf lens (Thorlabs LA1401), and (right) a clone of the same lens manufactured using our technique, to provide a visual comparison of optical quality.

Wet sanding For the wet sanding approach, we print a desired target shape. It can be any arbitrary shape. In our, first example we print a cube, for the rest purposes, cut the corners, and hand-polish the cube with sandpaper starting from coarse to fine grades. We pour water on the sandpaper from time to time to avoid scratches on the cube surface. Our hand-polishing is done by attaching the sandpaper to a rotating drill, reducing the amount of the time needed to polish a surface (~ 5 hours). At each stage, we observe the surface using a microscope to ensure smooth and even polishing. This process continues until we obtain an optically transparent cube as shown in Figure 7.

We use the wet-sanding technique to replicate an off-the-shelf plano-convex lens (Thorlabs LA1401) by printing the lens shape and polishing it. We provide a comparison between an actual lens made out of optical glass (BK7) and a replicated copy as shown in Figure 8. The imaging characteristics of the replica lens are a good approximation, although there is some yellowing in color and minor haze which limits the smallest resolvable point. We believe major component of haze arrives from layered structure of printed pieces. Final product, the clone lens can also be seen in Figure 7 labeled “Lens A”.

We also test whether optical components that use total internal reflection can be made with sanding. We design a waveguide similar to the design sketched in Figure 3 as a part of a NED prototype. In this particular design, one of the surfaces of the waveguide has to be diffusive, and this diffusive surface is designed to replicate a planar image at a 1D distance for a NED display design based on Figure 2 with a tilt of $\theta_{BC} = 25^\circ$ and a 50 mm radius of curvature spherical beam combiner. We first polish the diffusive surface of the designed waveguide to slightly improve the surface roughness and to remove large scratches. Using a diffusive tape on the polished surface, we are able to build a clear diffusive surface that has sufficiently large diffusion cones. The resultant waveguide can be observed as “Prism A” in Figure 7. The light travelling inside Prism A has a long optical path, therefore the haze caused by bulk scatter produces poor optical characteristics rendering it unsuitable for our purposes. The major drawback of this technique is the the extensive time and labor needed to create optically clear freeform surfaces.

Optical bonding To decrease the haze caused by bulk scattering and the length of the sanding process, we explore optical bonding of our 3D printed parts to optical glass. We revisit the design of “Prism A” in Figure 7, and decouple it into two pieces, a right angle prism and a diffusive surface. The right angle prism in our design is an off-the-shelf component (Thorlabs PS908). We print the diffusive surface part and skip the polishing step, and instead bond it to the right angle prism using an off-the-shelf optical adhesive (Norland 68). We apply the diffusive tape after polishing the diffusive surface to remove large scratches. The resulting piece can be observed as “Prism B” in Figure 7. The haze and light efficiency of the optical part is subjectively improved. In our supplementary material, we provide more media on the “Prism B”-based NED prototype, and we also provide an example of bonding




Fig. 9. Two photographs showing our Near-Eye Displays that use interchangeable freeform projection surfaces with birdbath optics. The photograph on the left shows our untethered wearable Near-Eye Display prototype, and the photograph on the right shows our bench-top prototype with projection lens steering capabilities.

two lenses, and bonding a printed lens with an existing lens where different zones in the combined lens provide different magnifications. Overall, optical bonding reduces sanding time, labor, and resulting haze, but it necessitates segmentation of the optical pieces, which limits the final design of the target geometric shape due to the dependence on existing optical components.

Vacuum forming To provide more design freedom for the final target surface shape, move away from the intensive labor required in wet sanding, and also achieve a more automated pipeline, we explore an alternative technique using vacuum forming. We first print the target shape and apply an excessive amount of optical adhesive on the surface to be smoothed. Using a vacuum former (Formmech 508DT), we bend a clear acrylic plastic piece around printed target shape that has been covered with optical adhesive. During the vacuum forming process, we visually check to ensure there are no air gaps in the optical adhesive layer. We use an ultra-violet light source (Formlabs Form Cure) to cure the optical adhesive for ~ 5 h. Once the part is fully cured, we remove the 3D printed part from the vacuum-formed clear acrylic mold. We provide documentation related to each of these steps in our supplementary material. Sample parts that were manufactured using vacuum forming are shown as “Diffuser” and “Lens B” in Figure 7. “Lens B” is a custom-designed lens which can be placed 15 mm in front of a projector to focus at a short throw distance of 50 mm with a projection size of 20×30 mm. Unlike off-the-shelf lenses, “Lens B” can project a large enough image to cover the entire surface of the diffuser in our NED prototypes. Overall, we found that our final vacuum forming pipeline saves time and labor, providing a reasonable gateway towards rapid development of novel optical designs. Although the technique described here cannot provide a production-quality optical piece yet, it opens up the possibility for a computational optics researcher to investigate and iterate on a given lens design in a day, rather than a months-long time frame, while decreasing the iteration cost greatly.

4.2 Prototypes

Our prototypes largely use 3D printed optical and mechanical components, augmented with a handful of off-the-shelf components. The optical paths of the prototypes described in this section use the optical layout described in Figure 2, and use an optical path folding with beam-splitters as sketched in Figure 3. The projection light engine we are using is a commercially available pico-projector that combines RGB LEDs to generate a white light source. Combined time-multiplexed white light in our light engines is modulated using LCoS devices (OVP921, 1280×720 pixels, 60 Hz from ImagineOptix). The beam-splitters used in our prototypes are off-the-shelf 50 : 50 economy beamsplitters (Thorlabs EBP2), and tilted 22.5° with respect to the optical axis. The beam combiners are Zeonex-based custom-made spherical concave mirrors (Diverse Optics) with a clear aperture of 60 mm, reflectivity of 80%, and a transmission of 20%. We build two prototypes to explore effectiveness of our approach: an untethered wear-

able prototype without foveation and a benchtop prototype to explore possibilities with foveated lenses.

Wearable prototype Aside from the optical components we have already described, the wearable prototype shown in Figure 9 also uses a 3D printed housing to accommodate the optical and electronic drive circuits. This particular prototype is equipped with a freeform projection surface that can cover 30° vertical FoV and 55° horizontal FoV while using off-the-shelf lenses (Thorlabs LA1304 and LC1439). Our wearable prototype is connected to a small compute model located at the back of the head via the head strap, as shown in Figure 9. The small computer in our system is made of two Raspberry Pi Zero W units, a USB battery (Anker PowerCore 5000mAh), and a 3D printed housing. We are able to obtain usage times of up to 3 hours with this battery, thereby providing a standalone untethered device. We use OpenGL-ES to render images on our small computer. We provide CAD models of our design along with sample rendering code that runs on our mini computers. Our core intention in this prototype is to show that it is feasible to realize a nearly complete NED product with this approach.

Bench-top prototype To evaluate the optical qualities of our manufacturing technique, we also build a benchtop prototype as shown in Figure 9. The benchtop prototype serves as a base to explore the idea of steering off-the-shelf projection lenses, which can augment a foveal inset in a foveated display or a higher pixel density projector and the idea of a steerable foveated projection lens. Our benchtop prototype is equipped with a freeform projection surface that can cover a 30° vertical FoV and a 55° horizontal FoV. We have designed and printed an x-y stage that moves within a 1 cm by 1 cm area to move a projection lens in front of our light engine. The projection lenses that we used are an off-the-shelf lens (Thorlabs LA1304) and a foveated projection lens design candidate used in front of an off-the-shelf lens (Thorlabs 1304). Our x-y stage is equipped with two off-the-shelf linear actuators (Actuonix PQ12-30-6R). In this setup a projection lens can be moved within 400 – 600 msec in the temporal direction, allowing relocation of the center of projection with speeds of 40 degrees per second for the FoV of our prototype. We control the linear stages using a joystick connected to an Arduino UNO board. Our supplementary documentation also includes CAD models and Arduino code for the benchtop prototype along with our optical simulation design to identify a foveated lens candidate.

5 EVALUATION

5.1 Optical characteristics

Using our bench-top prototype, shown in Figure 9, we evaluate the optical characteristics of elements produced with our manufacturing technique. These evaluation photographs are captured with a Canon EOS Rebel T6i camera body and a Canon 24 – 70 mm 1 : 2.8 lens, and were taken in a dark room using exposure times of 50 ms and an f-number of $f/9$. To mimic the behavior of a human eye, we place the camera at an eye relief of 40 mm.




Fig. 10. Three plots showing the resolution characteristics of our benchtop prototype targeting depth levels of 0.3 D, 1 D, and 2 D using a dedicated freeform projection surface, an off-the-shelf lens (Thorlabs LA1304) for our projector, and a two-axis lens moving stage to steer the projected images to different locations along a freeform projection surface. The three colored lines correspond to central, nasal periphery, and temporal periphery visual fields of view.

Eyebow To evaluate the properties of the eyebow, we capture photographs through our bench-top prototype as presented in Figure 11. The evaluation setup consists of a 3D printed diffuser which has been designed to display a rectangular virtual image at a depth of 1 D with a primary viewpoint position 5 mm below the optical axis defined in Figure 2. Average pixel brightness values in a region of interest that is manually selected around the logo show that the brightest and the sharpest viewpoint inside the eyebow is decentered 5 – 10 mm downward, indicating that our optical design and the photographs captured for the eyebow measurement are well-aligned. Additionally, Figure 11 shows photographs from different locations within the eyebow, and provides an understanding on eyebow dimensions. In our experimentation, relative to the target viewpoint, we find that the width of the eyebow is 15 mm from left to right and the height is 20 mm from top to bottom. Subjectively we also observe that the eyebow of our bench-top prototype is large enough to accommodate gaze changes within the provided FoV. If necessary, the geometric distortions visible in the figure could be corrected on-demand in software using an eye tracker as our optical model is well aligned with the actual prototypes.

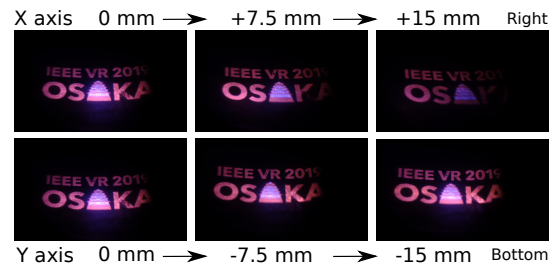


Fig. 11. Evaluating the eyebow properties of our benchtop prototype with a series of photographs taken from different locations inside the eyebow. The two rows show the view as seen through the corresponding viewpoint location in the benchtop prototype. The top row demonstrates the eyebow properties along the X axis relative to the center point and the bottom row shows eyebow properties along Y axis relative to center.

Field of view Our prototypes are designed with a FoV of $55^\circ \times 30^\circ$. In our benchtop prototype, we experiment with two different projection lenses to focus the projector for a short throw distance: an off-the-shelf lens (Thorlabs LB1378-ML), and a custom-designed 3D printed projection lens. Using the off-the-shelf lens the projector is only able to show the entire image with a FoV of $25^\circ \times 16^\circ$, while our custom-designed 3D printed projection lens can cover the entire FoV at once using the same projector.

Resolution We investigate the resolution characteristics of our benchtop prototype using various 3D printed diffusers that have been designed to generate rectangular virtual images at various depth levels (0.33 D, 1 D, 2 D). To provide the greatest pixel density across each 3D printed diffuser surface, we use an off-the-shelf lens (Thorlabs LA1304), which covers a smaller FoV than entire supported FoV. To achieve a high-resolution inset, we steer the off-the-shelf lens in front of the projector. In all our experiments, we adjust the focus of the projector

according to the given 3D printed diffuser surface. Our evaluation follows an industry standard, ISO 12233 slanted-edge Modulation Transfer Function (MTF) method [13]. All the captured photographs are pre-processed by a low-pass filter to decrease high-frequency noise from the background and the sensor. Our raw data, along with the MTF measurements, are provided as in Figure 10. We report the best resolution as 11 cpd in the central portion of the FoV, across all 3D printed diffusers at each of the different depth levels.

Brightness Consumer level NEDs are typically designed to operate in an in-door environment, and may not provide good images under direct sunlight conditions. Our prototype can work under strong sunlight conditions as demonstrated in Figure 1.

Accommodation support We investigate accommodation support in our bench-top prototype. To show the accommodation characteristics of our freeform projection surfaces, we target a depth curvature that progressively changes in optical depth from farther (0.3 D) to closer (5 D), from bottom to top. We compute a freeform surface using our methodology and demonstrate the accommodation characteristics of our benchtop prototype in Figure 1. We also show that complex depth targets can also be calculated using our regression model, and translated into a freeform surface as in Figure 13. Designers have to consider the discontinuities of a depth-target as it can decrease usable area on a freeform projection surface as such cases lead to sharp edges in depth.

Foveation We explore on foveation in hardware using our bench-top prototype. We design and manufacture a new lens design that can be placed in front of an existing off-the-shelf lens (Thorlabs LA1304). Together with off-the-shelf lens, our foveated lens leads to a lens assembly projects intentionally distorted larger images as with respect to a stand-a-lone off-the-shelf lens (Thorlabs LA1304) scenario as depicted in Figure 6. Our final lens assembly covers entire FoV provided by our freeform diffusers ($55^\circ \times 30^\circ$). We move the lens assembly using linear actuators to change the highest resolution portion of the image, leading to results provided in Figure 12. We believe this provides a design candidate for a foveation based near-eye display, which greatly simplifies foveated near-eye display hardware.

5.2 Limitations and Future work

Manufacturing Our proposed optical manufacturing technique promises a simple and an effective solution to a major challenge in the daily lives of optical designers. Optical quality of the pieces manufactured using our technique is shown to be a good match for manufacturing our NED design. The results regarding to the resolution of our NED design (12 cpd) shows resolution levels that are beyond a typical consumer level NED (5 – 10 cpd). However, it should be noted that this is partially due to a well-considered choice in optical design. For example, in our foveated projection lens design, a light engine uses a large surface to project a beam representing a single pixel, therefore aberration related issues caused by imperfection of an optical piece are much less pronounced. Some other optical layouts for different applications or NED designs may have more strict requirements. A good example of such optical designs are the ones that require much shorter focal lengths with much shorter aperture sizes (microscopes).




Fig. 12. Three photographs (A-B-C) showing change in resolution characteristics of our bench-top prototype as our in-house designed and manufactured foveated lens is shifted together with an off-the-shelf lens (LA1304). Entire assembly containing our foveated lens and a LA1304 is shifted 5 mm along one direction. Each photograph also contains a zoomed version of a rectangular region highlighted with red, green, and blue colors. Note that optical distortions introduce intentionally distributed pixels more densely at portions highlighted with colored dashed circles.




Fig. 13. Evaluating a freeform projection surface with a complex scene information. Photograph on the left shows diffusive side of a freeform projection surface manufactured in-house. Photograph on the right shows view through bench-top prototype, while the manufactured freeform projection surface is illuminated with white light.

As our manufacturing technique is a prototyping tool at this stage, vacuum forming manually doesn't guarantee very accurate manufacturing of optical pieces. As highlighted in Figure 8, the optical clarity of the 3D printed pieces is not as good as glass-based equivalents, this is due to two reasons: the material characteristics, and the layer-by-layer nature of additive manufacturing. Materials used in 3D printers also cause a yellowing effect due to its material properties. Perhaps following a careful study of error analysis caused by vacuum forming, and computationally designing optical pieces by taking errors into account that are originated from vacuum forming may provide improved results in achieving a surface as in a target design.

Conventional plastic-based lenses used in product-level VR NEDs are typically manufactured using an injection molding technique. Inspired by injection molding, we have also most recently experimented with 3D printed negative pieces processed using our technique, in which we fill the enclosed separable negative pieces with optical adhesive to manufacture a target optical design. As final piece is only made out of a cured optical adhesive, harvested final piece from a well cured mold promises to avoid yellowing effect of a 3D printer resin, and bulk scatter caused by layer-by-layer printing. We believe building upon our findings, and targeting solutions in the future that can merge to a dedicated automation system for making of 3D printed parts would lead to a hassle free technique as the current pipeline still requires a little bit of expertise and manual labor.

Optical characteristics In our prototypes, we use beam-splitters to fold the path between the beam combiner and the freeform projection surface. Some of the disadvantages of using beam-splitters include a limit on the maximum FoV, ghost images caused by reflections from both surfaces of the beam-splitter, and the lack of a clear path toward a thin glasses-like form factor without sacrificing FoV. HOEs promise a thin form factor representation of optical elements using a thin film, which fixes many of the problems related to using beam-splitters. Following the steps of Akşit et al. [1], we investigate building an AR NED prototype which mimics the properties of a focal surface by physically reshaping an HOE using vacuum forming and 3D printing. Although we show that this can work, we observe that the HOE deteriorates over

time, making this technique impractical. Our findings on curved HOEs are shared in the supplementary material. We find that the HOEs can be physically [55] or virtually [12] curved using different methods, and it is also possible to combine freeform beam combiners and freeform projections surfaces [6], all of which we will explore in greater detail in the near future.

Application-driven projection surfaces Our exploration of the RGB-D datasets shown in Figure 5 suggest that target use cases involving close depth ranges (i.e. working at a desk, or the office workspace) may pose a challenging scenario for our approach. The work of [52] addresses this problem at the cost of a large computation and power demand. To provide a solution for these challenging scenarios, we will consider upgrading our proposal with multiple (2 – 3) interchangeable freeform surfaces with a single projector in a time multiplexed fashion following the work of Liu et al. [49]. Perhaps another simple approach to introduce dynamic behaviour to a static projection surface is through introducing an actuation mechanism to support changing focus as in the work of Akşit et al. [1].

Foveation Our investigation on a design candidate for a foveated NED that greatly reduces the complexity of an optical NED leads to a promising outcome. However, further perceptual studies on requirements in speed of actuation, rendering seamlessly without any artifacts that may cause due to actuation, and an understanding on design trade space of foveated projection lenses are open questions at this point, which we will investigate in a greater detail at a near future as a continuation of our investigation. Like every other foveated display hardware and software, a gaze tracking mechanism is required for our design, while using a foveated lens or actively changing projection location on a freeform diffuser.

6 CONCLUSION

From printed artificial organs to industrial-grade replaceable parts, commodity 3D printing offers new ways to manufacture and design customized components. We believe that 3D printing can play an increasingly important role in designing augmented reality near-eye displays. Customizing optical designs allows researchers to prototype previously non-existent unconventional optics, and explore the needs of near-eye display designs at a much more rapid pace.

With this goal in mind, we introduce a manufacturing methodology using commodity tools. Our methodology can be thought of as a cost-effective process for discovering the best optical design for a given system or application. We examine real-life scenes to inform the optical needs of augmented reality devices. Based on these assessments, we identify a computational approach for designing various freeform optical pieces to target different use cases. We extended our designs to investigate on optical designs that can support foveation in hardware. All of these computational tools help us to prototype an application-driven untethered foveated display. Some hurdles remain in our practical implementation, which we believe can be addressed with further research. We hope our work will inspire researchers to

customize optical elements for near-eye displays using our optics manufacturing techniques, and help them to formulate their designs in a simpler way.

REFERENCES

- [1] K. Akşit, W. Lopes, J. Kim, P. Shirley, and D. Luebke. Near-eye varifocal augmented reality display using see-through screens. *ACM Trans. Graph. (SIGGRAPH)*, (6), 2017.
- [2] K. Akeley, S. J. Watt, A. R. Girshick, and M. S. Banks. A stereo display prototype with multiple focal distances. In *ACM transactions on graphics (TOG)*, volume 23, pages 804–813. ACM, 2004.
- [3] K. Akşit, J. Kautz, and D. Luebke. Slim near-eye display using pinhole aperture arrays. *Applied optics*, 54(11):3422–3427, 2015.
- [4] J. Alexander, A. Roudaut, J. Steimle, K. Hornbaek, M. Bruns Alonso, S. Follmer, and T. Merritt. *Grand Challenges in Shape-Changing Interface Research*. ACM, 12 2017.
- [5] D. Baldwin. Area of interest: Instantaneous field of view vision model. Image Generation/Display Conference, 1981.
- [6] A. Bauer, E. M. Schiesser, and J. P. Rolland. Starting geometry creation and design method for freeform optics. *Nature communications*, 9(1):1756, 2018.
- [7] H. Benko and A. D. Wilson. Multi-point interactions with immersive omnidirectional visualizations in a dome. In *ACM International Conference on Interactive Tabletops and Surfaces, ITS '10*, pages 19–28, New York, NY, USA, 2010. ACM.
- [8] H. Benko, A. D. Wilson, and R. Balakrishnan. Sphere: Multi-touch interactions on a spherical display. In *Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology, UIST '08*, pages 77–86, New York, NY, USA, 2008. ACM.
- [9] B. Bickel, P. Cignoni, L. Malomo, and N. Pietroni. State of the art on stylized fabrication. In *Computer Graphics Forum*. Wiley Online Library, 2018.
- [10] J. Bolton, K. Kim, and R. Vertegaal. Snowglobe: A spherical fish-tank vr display. In *CHI '11 Extended Abstracts on Human Factors in Computing Systems, CHI EA '11*, pages 1159–1164, New York, NY, USA, 2011. ACM.
- [11] E. Brockmeyer, I. Poupyrev, and S. Hudson. Papillon: designing curved display surfaces with printed optics. In *Proceedings of the 26th annual ACM symposium on User interface software and technology*, pages 457–462. ACM, 2013.
- [12] F.-K. Bruder, H. Bang, T. Fäcke, R. Hagen, D. Hönel, E. Orselli, C. Rewitz, T. Rölle, D. Vukicevic, and G. Walze. Precision holographic optical elements in bayfol hx photopolymer. In *Practical Holography XXX: Materials and Applications*, volume 9771, page 977103. International Society for Optics and Photonics, 2016.
- [13] P. D. Burns. Slanted-edge mtf for digital camera and scanner analysis. In *Is and Ts Pics Conference*, pages 135–138. SOCIETY FOR IMAGING SCIENCE & TECHNOLOGY, 2000.
- [14] O. Cakmakci, B. Moore, H. Foroosh, and J. P. Rolland. Optimal local shape description for rotationally non-symmetric optical surface design and analysis. *Optics express*, 16(3):1583–1589, 2008.
- [15] O. Cakmakci and J. Rolland. Head-worn displays: a review. *Journal of display technology*, 2(3):199–216, 2006.
- [16] A. A. Cameron. Optical waveguide technology and its application in head mounted displays. In *Proc. SPIE*, volume 8383, page 83830E, 2012.
- [17] F. W. Campbell. The depth of field of the human eye. *Journal of Modern Optics*, 4(4):157–164, 1957.
- [18] W. Charman and H. Whitefoot. Pupil diameter and the depth-of-field of the human eye as measured by laser speckle. *Journal of Modern Optics*, 24(12):1211–1216, 1977.
- [19] K. J. Ciuffreda, B. Wang, and B. Vasudevan. Conceptual model of human blur perception. *Vision Research*, 47(9):1245 – 1252, 2007.
- [20] M. Cornia, L. Baraldi, G. Serra, and R. Cucchiara. Predicting Human Eye Fixations via an LSTM-based Saliency Attentive Model. *ArXiv e-prints*, Nov. 2016.
- [21] G. Damberg and W. Heidrich. Efficient freeform lens optimization for computational caustic displays. *Optics express*, 23(8):10224–10232, 2015.
- [22] S. De Groot and J. Gebhard. Pupil size as determined by adapting luminance. *JOSA*, 42(7):492–495, 1952.
- [23] N. A. Dodgson. Variation and extrema of human interpupillary distance. In *Stereoscopic Displays and Virtual Reality Systems XI*, volume 5291, pages 36–47. International Society for Optics and Photonics, 2004.
- [24] D. Dunn, C. Tippets, K. Torell, P. Kellnhofer, K. Akşit, P. Didyk, K. Myszkowski, D. Luebke, and H. Fuchs. Wide field of view varifocal near-eye display using see-through deformable membrane mirrors. *IEEE Transactions on Visualization and Computer Graphics*, 23(4):1322–1331, 2017.
- [25] G. H. Gelinck, H. E. A. Huitema, M. Mil, E. Veenendaal, P. J. G. Lieshout, F. Touwslager, S. F. Patry, S. Sohn, T. Whitesides, and M. D. McCreary. A rollable, organic electrophoretic qvga display with fieldshielded pixel architecture. *Journal of the Society for Information Display*, 14(2):113–118, 2006.
- [26] G. Godin, P. Massicotte, and L. Borgeat. High-resolution insets in projector-based stereoscopic displays: principles and techniques. In *Electronic Imaging 2006*, pages 60550F–60550F. International Society for Optics and Photonics, 2006.
- [27] D. M. Hoffman, A. R. Girshick, K. Akeley, and M. S. Banks. Vergence–accommodation conflicts hinder visual performance and cause visual fatigue. *Journal of vision*, 8(3):33–33, 2008.
- [28] T. Höllerer, J. Kuchera-Morin, and X. Amatriain. The allosphere: A large-scale immersive surround-view instrument. In *Proceedings of the 2007 Workshop on Emerging Displays Technologies: Images and Beyond: The Future of Displays and Interacton*, EDT '07, New York, NY, USA, 2007. ACM.
- [29] S. T. Holmstrom, U. Baran, and H. Urey. Mem laser scanners: a review. *Journal of Microelectromechanical Systems*, 23(2):259–275, 2014.
- [30] E. M. Howlett. High-resolution insets in wide-angle head-mounted stereoscopic displays. In *Stereoscopic Displays and Applications III*, volume 1669, pages 193–204. International Society for Optics and Photonics, 1992.
- [31] C.-H. Hsu, Y.-L. Wu, W.-H. Cheng, Y.-J. Chen, and K.-L. Hua. Holotube: a low-cost portable 360-degree interactive autostereoscopic display. *Multimedia Tools and Applications*, 76(7):9099–9132, Apr 2017.
- [32] X. Hu and H. Hua. High-resolution optical see-through multi-focal-plane head-mounted display using freeform optics. *Optics express*, 22(11):13896–13903, 2014.
- [33] H. Hua. Enabling focus cues in head-mounted displays. *Proceedings of the IEEE*, 105(5):805–824, 2017.
- [34] H. Hua and B. Javidi. A 3D integral imaging optical see-through head-mounted display. *Optics express*, 22(11), 2014.
- [35] F.-C. Huang, D. Luebke, and G. Wetzstein. The light field stereoscope. *ACM SIGGRAPH Emerging Technologies*, page 24, 2015.
- [36] D. Iwai, S. Mihara, and K. Sato. Extended depth-of-field projector by fast focal sweep projection. *IEEE transactions on visualization and computer graphics*, 21(4):462–470, 2015.
- [37] C. Jang, K. Bang, S. Moon, J. Kim, S. Lee, and B. Lee. Retinal 3d: Augmented reality near-eye display via pupil-tracked light field projection on retina. *ACM Trans. Graph.*, 36(6):190:1–190:13, Nov. 2017.
- [38] J. Kim, Q. Sun, F.-C. Huang, L.-Y. Wei, D. Luebke, and A. Kaufman. Perceptual studies for foveated light field displays. *arXiv preprint arXiv:1708.06034*, 2017.
- [39] A. Kolb, M. Lambers, S. Todt, N. Cuntz, and C. Rezk-Salama. Immersive rear projection on curved screens. In *2009 IEEE Virtual Reality Conference*, pages 285–286, March 2009.
- [40] B. Kress and T. Stamer. A review of head-mounted displays (hmd) technologies and applications for consumer electronics. In *SPIE Defense, Security, and Sensing*. International Society for Optics and Photonics, 2013.
- [41] P. Kumar and P. Maes. Penetra3d: A penetrable, interactive, 360-degree viewable display. In *3D User Interfaces (3DUI), 2015 IEEE Symposium on*, pages 169–170. IEEE, 2015.
- [42] C. Lang, T. V. Nguyen, H. Katti, K. Yadati, M. Kankanhalli, and S. Yan. Depth matters: Influence of depth cues on visual saliency. In *Computer vision–ECCV 2012*, pages 101–115. Springer, 2012.
- [43] D. Lanman and D. Luebke. Near-eye light field displays. *ACM Transactions on Graphics (TOG)*, 32(6):220, 2013.
- [44] S. Lee, J. Cho, B. Lee, Y. Jo, C. Jang, D. Kim, and B. Lee. Foveated retinal optimization for see-through near-eye multi-layer displays (invited paper). *IEEE Access*, PP(99):1–1, 2017.
- [45] S. Lee, Y. Jo, D. Yoo, J. Cho, D. Lee, and B. Lee. Tomoreal: Tomographic displays. *arXiv preprint arXiv:1804.04619*, 2018.
- [46] S. M. Lee, J. H. Kwon, S. Kwon, and K. C. Choi. A review of flexible oleds toward highly durable unusual displays. *IEEE Transactions on Electron Devices*, 64(5):1922–1931, May 2017.
- [47] D. Leithinger, S. Follmer, A. Olwal, and H. Ishii. Shape displays: Spatial

- interaction with dynamic physical form. *IEEE Computer Graphics and Applications*, 35(5):5–11, Sept 2015.
- [48] S. Liu, D. Cheng, and H. Hua. An optical see-through head mounted display with addressable focal planes. In *Mixed and Augmented Reality, 2008. ISMAR 2008. 7th IEEE/ACM International Symposium on*, pages 33–42. IEEE, 2008.
- [49] S. Liu, Y. Li, P. Zhou, X. Li, N. Rong, S. Huang, W. Lu, and Y. Su. A multi-plane optical see-through head mounted display design for augmented reality applications. *Journal of the Society for Information Display*, 24(4):246–251, 2016.
- [50] A. Maimone, A. Georgiou, and J. S. Kollin. Holographic near-eye displays for virtual and augmented reality. *ACM Transactions on Graphics (TOG)*, 36(4):85, 2017.
- [51] A. Maimone, D. Lanman, K. Rathinavel, K. Keller, D. Luebke, and H. Fuchs. Pinlight displays: wide field of view augmented reality eyeglasses using defocused point light sources. In *ACM SIGGRAPH 2014 Emerging Technologies Booth 203*. ACM, 2014.
- [52] N. Matsuda, A. Fix, and D. Lanman. Focal surface displays. *ACM Transactions on Graphics (TOG)*, 36(4):86, 2017.
- [53] W. Matusik, B. Ajdin, J. Gu, J. Lawrence, H. Lensch, F. Pellacini, and S. Rusinkiewicz. Printing spatially-varying reflectance. In *ACM Transactions on Graphics (TOG)*, volume 28, page 128. ACM, 2009.
- [54] O. Mercier, Y. Sulai, K. Mackenzie, M. Zannoli, J. Hillis, D. Nowrouzezahrai, and D. Lanman. Fast gaze-contingent optimal decompositions for multifocal displays. *ACM Transactions on Graphics (TOG)*, 36(6):237, 2017.
- [55] J. Moore and L. F. Holland. Systems, devices, and methods for splitter optics in wearable heads-up displays, May 1 2018. US Patent 9,958,682.
- [56] F. Mueller, D. Mehta, O. Sotnychenko, S. Sridhar, D. Casas, and C. Theobalt. Real-time hand tracking under occlusion from an egocentric rgb-d sensor. In *Proceedings of International Conference on Computer Vision (ICCV)*, October 2017.
- [57] S. Mueller, T. Mohr, K. Guenther, J. Frohnhofen, and P. Baudisch. fabrication: fast 3d printing of functional objects by integrating construction kit building blocks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 3827–3834. ACM, 2014.
- [58] R. Narain, R. A. Albert, A. Bulbul, G. J. Ward, M. S. Banks, and J. F. O'Brien. Optimal presentation of imagery with focus cues on multi-plane displays. *ACM Transactions on Graphics*, 34(4):59:1–12, Aug. 2015. To be presented at SIGGRAPH 2015, Los Angeles.
- [59] S. Olberding, M. Wessely, and J. Steimle. Printscreen: Fabricating highly customizable thin-film touch-displays. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology*, UIST '14, pages 281–290, New York, NY, USA, 2014. ACM.
- [60] V. Olesova, W. Benesova, and P. Polatsek. Visual attention in egocentric field-of-view using rgb-d data. In *Ninth International Conference on Machine Vision*, pages 103410T–103410T. International Society for Optics and Photonics, 2017.
- [61] M. Papas, T. Houit, D. Nowrouzezahrai, M. H. Gross, and W. Jarosz. The magic lens: refractive steganography. *ACM Trans. Graph.*, 31(6):186–1, 2012.
- [62] M. Papas, W. Jarosz, W. Jakob, S. Rusinkiewicz, W. Matusik, and T. Weyrich. Goal-based caustics. In *Computer Graphics Forum*, volume 30, pages 503–511. Wiley Online Library, 2011.
- [63] D. J. Parkhurst and E. Niebur. Variable-resolution displays: A theoretical, practical, and behavioral evaluation. *Human Factors*, 44(4):611, 2002.
- [64] 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):179, 2016.
- [65] T. Pereira, S. Rusinkiewicz, and W. Matusik. Computational light routing: 3d printed optical fibers for sensing and display. *ACM Transactions on Graphics (TOG)*, 33(3):24, 2014.
- [66] G. W. R. Konrad, E.A Cooper. Novel optical configurations for virtual reality: Evaluating user preference and performance with focus-tunable and monovision near-eye displays. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI16)*, 2016.
- [67] M. K. Rasmussen, E. W. Pedersen, M. G. Petersen, and K. Hornbæk. Shape-changing interfaces: A review of the design space and open research questions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '12, pages 735–744, New York, NY, USA, 2012. ACM.
- [68] S. M. Reder. On-line monitoring of eye-position signals in contingent and noncontingent paradigms. *Behavior Research Methods & Instrumentation*, 5(2):218–228, 1973.
- [69] J. P. Rolland, A. Yoshida, L. D. Davis, and J. H. Reif. High-resolution inset head-mounted display. *Applied optics*, 37(19):4183–4193, 1998.
- [70] Y. Schwartzburg, R. Testuz, A. Tagliasacchi, and M. Pauly. High-contrast computational caustic design. *ACM Transactions on Graphics (TOG)*, 33(4):74, 2014.
- [71] M. Shenker. Optical design criteria for binocular helmet-mounted displays. In *Display System Optics*, volume 778, pages 70–79. International Society for Optics and Photonics, 1987.
- [72] L. Shi, F.-C. Huang, W. Lopes, W. Matusik, and D. Luebke. Near-eye light field holographic rendering with spherical waves for wide field of view interactive 3d computer graphics. *ACM Trans. Graph.*, 36(6):236:1–236:17, Nov. 2017.
- [73] A. M. Spooner. The trend towards area of interest in visual simulation technology. Technical report, NAVAL TRAINING EQUIPMENT CENTER ORLANDO FL, 1982.
- [74] 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):192, 2017.
- [75] L. N. Thibos, F. E. Cheney, and D. J. Walsh. Retinal limits to the detection and resolution of gratings. *J. Opt. Soc. Am. A*, 4(8):1524–1529, Aug 1987.
- [76] L. N. Thibos, D. L. Still, and A. Bradley. Characterization of spatial aliasing and contrast sensitivity in peripheral vision. *Vision research*, 36(2):249–258, 1996.
- [77] Y. Tokuda, M. A. Norasikin, S. Subramanian, and D. Martinez Plasencia. Mistform: Adaptive shape changing fog screens. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, pages 4383–4395, New York, NY, USA, 2017. ACM.
- [78] J. Tompkin, S. Heinzle, J. Kautz, and W. Matusik. Content-adaptive lenticular prints. *ACM Transactions on Graphics (TOG)*, 32(4):133, 2013.
- [79] M. Vaezi, H. Seitz, and S. Yang. A review on 3d micro-additive manufacturing technologies. *The International Journal of Advanced Manufacturing Technology*, 67(5-8):1721–1754, 2013.
- [80] D. Van Krevelen and R. Poelman. A survey of augmented reality technologies, applications and limitations. *International journal of virtual reality*, 9(2):1, 2010.
- [81] Y. Wang. Advanced theory of field curvature. 2014.
- [82] M. Watanabe and S. K. Nayar. Telecentric optics for focus analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(12):1360–1365, 1997.
- [83] P. Webb. Bioastronautics data book. 1964.
- [84] M. Weiss, S. Voelker, C. Sutter, and J. Borchers. Benddesk: Dragging across the curve. In *ACM International Conference on Interactive Tabletops and Surfaces*, ITS '10, pages 1–10, New York, NY, USA, 2010. ACM.
- [85] T. Weyrich, P. Peers, W. Matusik, and S. Rusinkiewicz. Fabricating microgeometry for custom surface reflectance. *ACM Transactions on Graphics (TOG)*, 28(3):32, 2009.
- [86] D. R. Williams. Aliasing in human foveal vision. *Vision research*, 25(2):195–205, 1985.
- [87] D. R. Williams. Visibility of interference fringes near the resolution limit. *J. Opt. Soc. Am. A*, 2(7):1087–1093, Jul 1985.
- [88] K. Willis, E. Brockmeyer, S. Hudson, and I. Poupayev. Printed optics: 3d printing of embedded optical elements for interactive devices. In *Proceedings of the 25th annual ACM symposium on User interface software and technology*, pages 589–598. ACM, 2012.
- [89] W. Wu, P. Llull, I. Tosic, N. Bedard, K. Berkner, and N. Balram. Content-adaptive focus configuration for near-eye multi-focal displays. In *Multi-media and Expo (ICME), 2016 IEEE International Conference on*, pages 1–6. IEEE, 2016.
- [90] J. S. Yoo, S. H. Jung, Y. C. Kim, S. C. Byun, J. M. Kim, N. B. Choi, S. Y. Yoon, C. D. Kim, Y. K. Hwang, and I. J. Chung. Highly flexible amoled display with integrated gate driver using amorphous silicon tft on ultrathin metal foil. *Journal of Display Technology*, 6(11):565–570, Nov 2010.
- [91] J. W. Young. Head and face anthropometry of adult us civilians. Technical report, FEDERAL AVIATION ADMINISTRATION WASHINGTON DC OFFICE OF AVIATION MEDICINE, 1993.
- [92] Y. Yue, K. Iwasaki, B.-Y. Chen, Y. Dobashi, and T. Nishita. Poisson-based continuous surface generation for goal-based caustics. *ACM Transactions on Graphics (TOG)*, 33(3):31, 2014.
- [93] T. Zhan, Y.-H. Lee, and S.-T. Wu. High-resolution additive light field near-eye display by switchable pancharatnam–berry phase lenses. *Optics Express*, 26(4):4863–4872, 2018.