Near-Eye Light Field Displays

Douglas Lanman David Luebke
NVIDIA

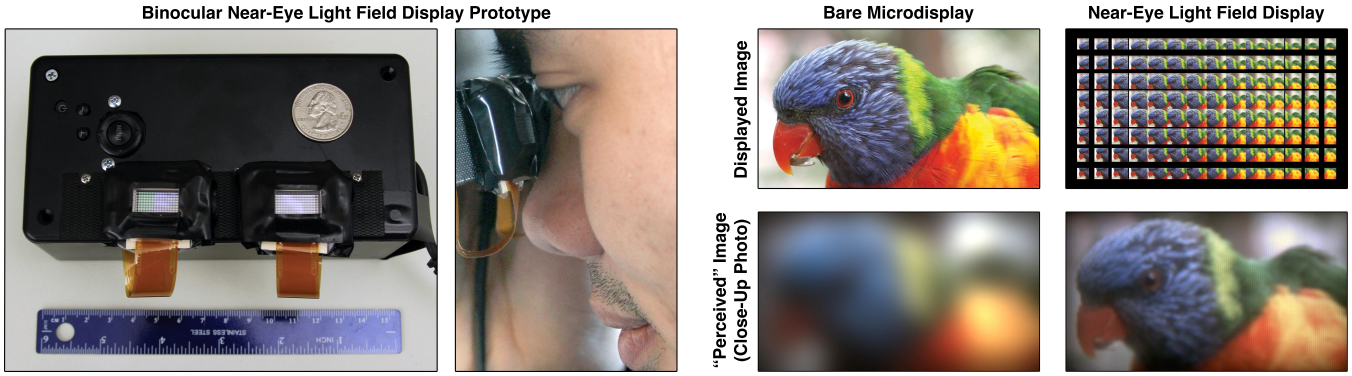

Figure 1: *Enabling thin, lightweight near-eye displays using light field displays. (Left) Our binocular near-eye display prototype is shown, comprising a pair of OLED panels covered with microlens arrays. This design enables a thin head-mounted display, since the black box containing driver electronics could be waist-mounted with longer ribbon cables. (Right) Due to the limited range of human accommodation, a severely defocused image is perceived when a bare microdisplay is held close to the eye (here simulated as a close-up photograph of an OLED). Conventional near-eye displays require bulky magnifying optics to facilitate accommodation. We propose near-eye light field displays as thin, lightweight alternatives, achieving comfortable viewing by synthesizing a light field corresponding to a virtual scene located within the accommodation range (here implemented by viewing a microdisplay, depicting interlaced perspectives, through a microlens array).*

Abstract

We propose near-eye light field displays capable of depicting sharp images by synthesizing light fields corresponding to virtual scenes located within a viewer's natural accommodation range. Our primary contribution is to evaluate the capability of microlens arrays to achieve practical near-eye light field displays. While sharing similarities with existing integral imaging displays and microlens-based light field cameras, we optimize performance in the context of near-eye viewing. As with light field cameras, near-eye light field displays support continuous accommodation of the eye throughout a finite depth of field; as a result, binocular configurations provide a means to address the accommodation-convergence conflict occurring with existing stereoscopic near-eye displays, producing quantified artifacts in approximated retinal defocus blur. We construct a binocular prototype and a real-time, GPU-accelerated stereoscopic light field renderer. Through simulations and experiments, we motivate near-eye light field displays as thin, lightweight alternatives to conventional near-eye displays.

1 Introduction

Near-eye displays project images directly into a viewer's eye, encompassing both head-mounted displays (HMDs) and electronic viewfinders. Such displays confront a fundamental problem: the unaided human eye cannot accommodate (focus) on objects placed in close proximity (see Figure 1). As reviewed by Rolland and Hua [2005], a multitude of optical solutions have been proposed since Sutherland [1968] introduced the first graphics-driven HMD. The majority of such designs emulate the behavior of a simple magnifier: synthesizing an enlarged image of a miniaturized display, appearing to be located within the viewer's natural accommodation range. To be of practical utility, a near-eye display should provide high-resolution, wide-field-of-view imagery with compact, comfortable magnifying optics. However, current magnifier designs typically require multiple optical elements to minimize aberrations, leading to bulky eyewear with limited fields of view that have, to date, prohibited widespread consumer adoption.

Conventional displays emit light isotropically. In contrast, a light field display supports the control of individual rays of light, modulating radiance as a function of position and direction across its surface. We consider a simple near-eye architecture: placing a light field display directly in front of a user's eye (or a pair of such displays for binocular viewing). As shown in Figure 1, sharp imagery is depicted by synthesizing a light field for a virtual display (or a general 3D scene) within the viewer's unaided accommodation range. As characterized in this work, near-eye light field displays provide a means to achieve thin, lightweight HMDs with wide fields of view and to address accommodation-convergence conflict in binocular configurations; however, these benefits come at a high cost: spatial resolution is significantly reduced with microlens-based designs, although with commensurate gains in depth of field and in accurate rendering of retinal defocus blur. It is our goal to assess the viability of this trade-off, in light of fundamental design limitations, anticipated reductions in microdisplay pixel pitches, and the emerging demand for comfortable, yet immersive, HMDs.

1.1 Contributions

Our primary technical contributions include the following:

- We motivate near-eye light field displays capable of synthesizing sharp images of virtual 3D objects located within a viewer's natural accommodation range, despite being in close proximity to the eye. We demonstrate such displays are ideally suited for thin-form-factor, wide-field-of-view HMDs.
- We demonstrate the capability of microlens arrays to create near-eye light field displays, optimizing the trade-offs between design parameters, including: form factor, spatial resolution, field of view, depth of field, and retinal blur.
- We construct a prototype binocular near-eye light field display, containing a pair of OLED displays covered with microlens arrays. We also implement a real-time stereoscopic light field renderer using GPU-accelerated ray tracing.

2 Related Work

2.1 Head-Mounted Displays

Commercial interest in HMDs has increased with the announcement of Google Glass[1] and the Oculus Rift[2]. The former is restricted to a narrow field of view located at the periphery of a viewer's visual field, constituting the first commercial variant of the "digital eye glasses" pioneered by Mann [2012]. Similar to Olson et al. [2011], the latter leverages smartphone technology to provide an immersive, affordable HMD, although with cumbersome headgear.

We are not the first to propose incorporating magnifier arrays within HMDs. Shaoulov et al. [2004] describe the magnification properties for stacks of two dissimilar microlens arrays. In a closely related work, Massof et al. [2003] achieve a 150×100 degree binocular field of view using a spherical array of microdisplays viewed through a curved, multifaceted lens array—a design now commercialized by Sensics, Inc. We emphasize that tiled HMDs are encompassed by our analysis; however, we introduce their use as light field displays, characterizing their ability to address accommodation-convergence conflict and correct for user's optical aberrations when applied in denser, thinner arrays to a single microdisplay.

2.2 Microlens Array Imaging

Imaging devices routinely incorporate microlens arrays. Anderson [1979] creates a unit magnification relay system using a pair of identical microlens arrays. Hutley et al. [1994] introduce *moiré magnification*: forming a magnified, periodic image when viewing an array of identical objects through a microlens array with a similar period. We recognize near-eye light field displays exploit a related phenomenon; however, the underlying imagery is aperiodic and forms magnified views corresponding to general 3D scenes, rather than periodic planar objects.

Near-eye light field displays share many limitations and benefits with light field cameras. Our microlens-based design is based on the integral imaging methods introduced by Lippmann [1908]. By capturing a photograph through a microlens array, an interlaced set of elemental images are recorded, corresponding to a series of off-axis perspective projections. As characterized by Jin et al. [2004], integral imaging mirrors many of the properties of our near-eye display: extending depth of field and field of view, at the cost of decreased spatial resolution. Similar trade-offs are characterized by Ng et al. [2005] for their hand-held light field camera, now commercialized by Lytro, Inc. We emphasize that a key benefit of light field cameras is to allow post-capture *refocusing*; equivalently, near-eye light fields support *focusing* of the eye, with the retina performing a similar integration of optically-aligned, but differing, views.

2.3 3D Displays

Most existing stereoscopic displays suffer from *accommodation-convergence conflict*: presenting accurate binocular disparity (supporting convergence on any point), but only allowing the viewer to accommodate on the display surface. As a result, content is restricted to a "zone of comfort" close to the surface, mitigating accommodation-convergence conflict. However, as characterized by Held et al. [2012], estimation of depth from binocular disparity is accurate near the plane of fixation, whereas retinal blur is more precise elsewhere (i.e., estimating depth from defocus). For near-eye applications, our microlens-based display depicts approximately-correct retinal blur via "super multiview" imagery [Takaki 2006], in which disparity is depicted *across a single pupil*.

[1] http://g.co/projectglass [2] http://www.oculusvr.com

Near-eye light field displays are not the only means to address accommodation-convergence conflict. Akeley et al. [2004] introduce a multifocal display, rendering images across three separate semi-transparent planes. Rolland et al. [2000] describe a related HMD architecture, incorporating a multilayer display similar to the DepthCube [Sullivan 2003]. Rather than requiring multiple physical displays, Love et al. [2009] synthesize a virtual multifocal display using a fast switchable lens synchronized with a single display. While such architectures provide near-correct accommodation cues, their construction currently prohibits thin-form-factor HMDs.

Several works have proposed placing integral imaging displays close to the eye. Pamplona et al. [2010] apply this configuration to estimate refractive errors. Pamplona et al. [2012] extend this work to correct optical aberrations. We recognize similar methods can be applied with our prototype, mitigating the need for corrective eyewear. In comparison to these closely-related works, we emphasize that our efforts differ in scope; we target general-purpose 3D display, rather than estimation and correction of refractive errors. Furthermore, we are the first to optimize the optical design and underlying rendering algorithms to enable thin, lightweight HMDs.

3 Implementation and Assessment

3.1 Hardware Implementation

Two near-eye light field displays were implemented: a static film-based prototype and a dynamic OLED-based prototype. As shown in Figure 2 and provided with the supplementary physical materials, a light valve technology (LVT) film recorder was used to develop 3.75×3.75 cm color transparencies at 120 pixels per mm (ppmm). The transparencies were backlit to emulate high-resolution microdisplays. The magnifier array comprised a Fresnel Technologies #630 rectangular plano-convex microlens sheet, with lens focal length $f = 3.3$ mm and lens width $w_l = 1.0$ mm. The microlenses were oriented with the planar surface facing the viewer. The separation between the microlens array and the transparency was manually adjusted to form a virtual image at a distance $d_o \approx 1.0$ m (as assessed by a focused camera). For an eye relief $d_e = 2.5$ cm, first-order optical analysis provides the following estimates of design parameters: spatial resolution $N_p = 534\times534$ pixels, field of view $\alpha = 67\times67$ degrees, and eye box width $w_e = 7.6\times7.6$ mm (i.e., the maximum lateral displacement relative to the display center). The depth of field extends over the interval $23.2 \leq d'_o < \infty$ cm.

As shown in Figure 1 and in the supplementary video, a binocular OLED-based prototype was constructed using components from a Sony HMZ-T1 personal media viewer. The case and magnifying eyepieces were removed, exposing a pair of Sony ECX322A microdisplays connected by ribbon cables to a driver board and a push button controller. Each 15.36×8.64 mm microdisplay has 1280×720 24-bit color pixels (i.e., 83.3 ppmm). Replicating the LVT-based prototype, Fresnel Technologies #630 microlenses were cut and affixed to each microdisplay (see Figure 3). The electronics were mounted in a plastic enclosure, with the modified microdisplays mounted to Velcro strips to adjust the interpupillary distance.

We estimate the following design parameters for each modified eyepiece: spatial resolution $N_p = 146\times78$ pixels, field of view $\alpha = 29.2\times16.0$ degrees, and eye box width $w_e = 7.6\times7.6$ mm. The depth of field extends over $30.6 \leq d'_o < \infty$ cm. Considering the electronics enclosure as a separate component (which could be waist-mounted with longer cables), the prototype allows an HMD with a form factor corresponding to a pair of modified eyepieces, each measuring $42\times31\times10$ mm and with a 0.7 gram microlens array. In comparison, an unmodified HMZ-T1 eyepiece is $43\times31\times37$ mm with a 57.7 gram lens.

3.2 Software Implementation

The software implementation addresses two challenges: real-time, stereoscopic light field rendering and efficient calibration and correction of mechanical alignment errors and optical aberrations. The LVT-based and OLED-based prototypes contain magnifier arrays with 35×35 and 14×8 lenses, respectively. A direct extension of rasterization would require rendering one projection of the 3D scene for each lens, although only for pixels spanning the corresponding elemental image. As an alternative, we modified the NVIDIA OptiX GPU-accelerated ray tracing engine [Parker et al. 2010] to support quad buffering in OpenGL—providing the HDMI 1.4a frame-packed 3D format required by the HMZ-T1. As shown in the supplementary video, frame rates for the sample scenes varied from 15–70 Hz using a 3.2 GHz Intel Core i7 workstation with 8 GB of RAM and an NVIDIA Quadro K5000 graphics card.

During assembly, horizontal and vertical stripes are displayed on the OLEDs. The microlens array is rotated such that the stripes appear aligned to the microdisplay pixel grid. In practice, this procedure achieves accurate rotational alignment, but the lateral displacement must be corrected by translating the rendered images. Similarly, the manufactured focal length and lenslet width may differ from specifications; both of which are manually tuned using test images. Similar to Pamplona et al. [2012], spherical aberrations can be corrected, independently for each eye, by scaling the depth of the rendered scene. Viewers with minimal astigmatism report that, after calibration, the prototype can be viewed comfortably without eyeglasses. These alignment and correction tasks reduce to defining the design parameters in a configuration file. A simple calibration routine, presenting a set of test images including a Snellen chart, allows the user to interactively adjust these parameters.

3.3 Assessment

The performance of the prototypes is illustrated by close-up photographs in Figures 1 and 2–4 and in the supplementary video. All imagery was captured using a 1600×1200 Point Grey Flea3 digital camera with a Fujinon 2.8–8 mm varifocal lens supporting a minimum f-number of 1.2 (selected to approximate the human eye).

A central benefit of near-eye light fields displays is to support approximate retinal defocus blur, consistent with convergence of the eyes. A stereoscopic pair is shown in Figure 3, demonstrating the perception of a user holding the device close to his face, as in Figure 1. The accuracy of retinal imagery is compared in Figure 4. In this example, a scene containing two dominant object planes is viewed through a prototype eyepiece. We note that, both in simulation and experiment, retinal focus and defocus cues are achieved.

4 Benefits and Limitations

The primary benefits of near-eye light field displays are reduced thickness and weight (e.g., as achieved by substituting a compact microlens array for compound magnifying optics). Such displays also approximate retinal defocus blur and, correspondingly, offer a means to address accommodation-convergence conflict with binocular configurations. As demonstrated in the supplementary video and physical materials, our design can also correct for the viewer's optical aberrations (i.e., their eyeglasses prescription), utilizing the approach of Pamplona et al. [2012] in a near-eye configuration.

Reduction in spatial resolution is the primary impediment to near-eye light field displays. However, multiple technology trends are driving development of high-resolution microdisplays. Commercially, such displays are integral to electronic viewfinders and consumer HMDs (e.g., personal media players). Furthermore, government investment is advancing the state of the art for military HMDs (e.g., night vision systems). Yet, in the near term, microdisplay resolutions are limited and conventional designs have greater commercial potential, despite the inability to address accommodation-convergence conflict and demand for thinner form factors.

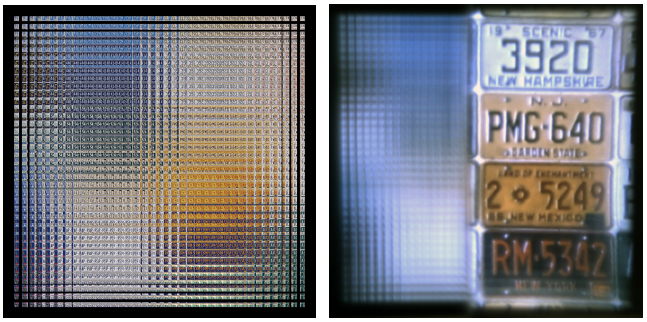

Figure 2: *The LVT-based prototype. (Left) The interlaced set of elemental images required to depict a set of license plates separated by* 1.0 *m from the viewer. (Zoom in to see the detailed microstructure.) (Right) A microlens array is placed over the right half of the developed LVT film. A sharp image is perceived, when viewed through the microlenses, whereas the bare LVT appears defocused.*

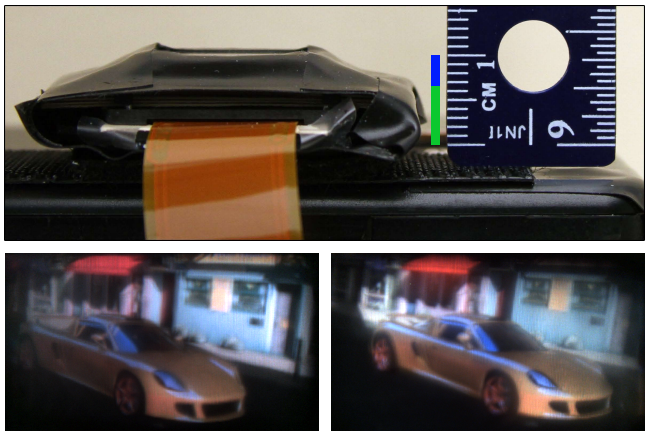

Figure 3: *The OLED-based prototype. (Top) A modified eyepiece is shown in profile. The ruler illustrates a thickness of* 1 *cm. The microlens array (length shaded blue) is thinner than the OLED enclosure (length shaded green). (Bottom) A stereoscopic pair of photographs of the binocular prototype. The left-eye and right-eye images are interchanged to facilitate cross-fused stereo viewing.*

Pixel pitches cannot reduce below the diffraction limit imposed by individual lenses. Following Goodman [2004], far-field diffraction by a square lens, of width w_l and with microdisplay separation d_l, limits the pixel pitch p to $\lambda d_l / w_l$. For our microlens array, the pixel pitch could be reduced to 2.5 micrometers (400 ppmm)—allowing the OLED prototype resolution to increase by a factor of five to 711×383 pixels. Further gains could be realized with wider lenses, at the cost of decreased retinal blur accuracy, and by increasing the dimensions of the microdisplay, also yielding a wider field of view.

Microlens-based designs share the limitations of other integral imaging displays, including the creation of periodic eye boxes. Collimators and other viewing-angle-limiting materials could be introduced to eliminate repetition. Structured illumination could be applied, similar to Levoy et al. [2009], to more accurately calibrate the mapping from microdisplay pixels to emitted rays. In this circumstance, our ray general tracing engine has an additional benefit over direct rasterization: allowing arbitrary pixel-to-ray mappings.

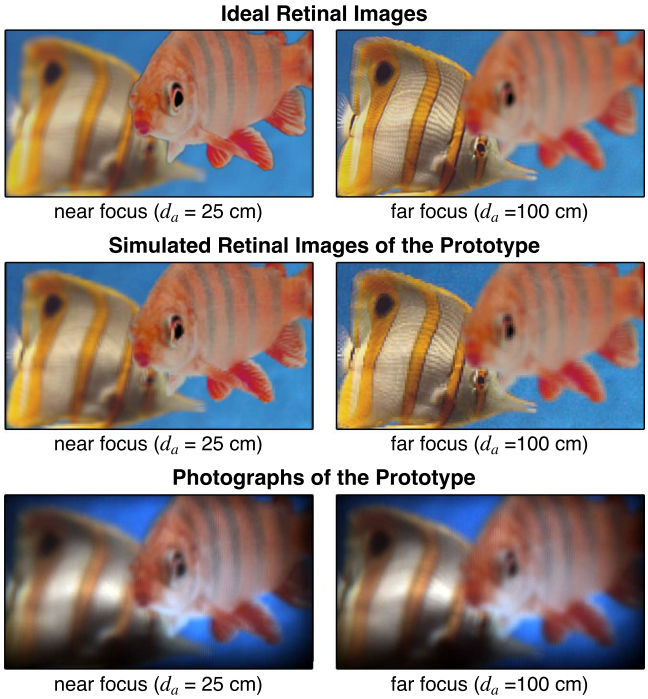

Figure 4: *Approximating retinal blur. The fish on the left and right are located distances of* 25 *cm and* 100 *cm away, respectively. (Top) Simulation of ideal retinal images for a* 4 *mm pupil diameter. (Middle) A ray tracing model, implemented in Matlab, simulates the retinal images for the OLED-based prototype in Section 3. (Bottom) Photographs of the prototype are in close agreement. Note that retinal defocus blur is approximated by averaging overlapping views from lenses spanning the circle of confusion.*

5 Conclusion

Near-eye displays are poised to enter the consumer market. However, emerging devices present one of two restrictive solutions: either narrow-field-of-view displays, located in the periphery of a viewer's visual field, or bulky designs held in place with tight straps. These compromises are necessary to achieve lightweight, eyeglasses-like form factors, with the former, or to obtain wide-field-of-view, immersive experiences, with the latter. Commercial near-eye displays have not yet met these demands with thin, lightweight designs. The microlens-based display architecture we propose is aimed at achieving wide-field-of-view, immersive experiences with compact, comfortable eyewear—offering a new path to practical head-mounted displays that trades spatial resolution for significant improvements in field of view, weight, and form factor.

References

AKELEY, K., WATT, S. J., GIRSHICK, A. R., AND BANKS, M. S. 2004. A stereo display prototype with multiple focal distances. *ACM Trans. Graph. 23*, 804–813.

ANDERSON, R. H. 1979. Close-up imaging of documents and displays with lens arrays. *Applied Optics 18*, 4, 477–484.

ARMITAGE, D., UNDERWOOD, I., AND WU, S. 2006. *Introduction to Microdisplays*. Wiley.

BURDEA, G. C., AND COIFFET, P. 2003. *Virtual Reality Technology*. Wiley-IEEE Press.

GOODMAN, J. 2004. *Introduction to Fourier Optics*. Roberts and Company Publishers.

HELD, R. T., COOPER, E. A., AND BANKS, M. S. 2012. Blur and disparity are complementary cues to depth. *Current Biology 22*, 5, 426–431.

HUTLEY, M. C., HUNT, R., STEVENS, R. F., AND SAVANDER, P. 1994. The moiré magnifier. *Pure and Applied Optics: Journal of the European Optical Society Part A 3*, 2, 133–142.

JIN, F., JANG, J.-S., AND JAVIDI, B. 2004. Effects of device resolution on three-dimensional integral imaging. *Optics Letters 29*, 12, 1345–1347.

LEVOY, M., ZHANG, Z., AND MCDOWALL, I. 2009. Recording and controlling the 4D light field in a microscope using microlens arrays. *Journal of Microscopy 235*, 2, 144–162.

LIPPMANN, G. 1908. Épreuves réversibles donnant la sensation du relief. *Journal of Physics 7*, 4, 821–825.

LOVE, G. D., HOFFMAN, D. M., HANDS, P. J., GAO, J., KIRBY, A. K., AND BANKS, M. S. 2009. High-speed switchable lens enables the development of a volumetric stereoscopic display. *Optics Express 17*, 18, 15716–15725.

MANN, S. 2012. Through the glass, lightly. *IEEE Technology and Society Magazine 31*, 3, 10–14.

MASSOF, R. W., BROWN, L. G., SHAPIRO, M. D., BARNETT, G. D., BAKER, F. H., AND KUROSAWA, F. 2003. Full-field high-resolution binocular HMD. *SID Digest 34*, 1, 1145–1147.

NG, R., LEVOY, M., BRÉDIF, M., DUVAL, G., HOROWITZ, M., AND HANRAHAN, P. 2005. Light field photography with a hand-held plenoptic camera. Tech. Rep. CTSR 2005-02, Stanford.

OLSON, J., KRUM, D., SUMA, E., AND BOLAS, M. 2011. A design for a smartphone-based head mounted display. In *IEEE Virtual Reality*, 233–234.

PAMPLONA, V. F., MOHAN, A., OLIVEIRA, M. M., AND RASKAR, R. 2010. NETRA: interactive display for estimating refractive errors and focal range. *ACM Trans. Graph. 29*.

PAMPLONA, V. F., OLIVEIRA, M. M., ALIAGA, D. G., AND RASKAR, R. 2012. Tailored displays to compensate for visual aberrations. *ACM Trans. Graph. 31*, 4, 81:1–81:12.

PARKER, S. G., BIGLER, J., DIETRICH, A., FRIEDRICH, H., HOBEROCK, J., LUEBKE, D., MCALLISTER, D., MCGUIRE, M., MORLEY, K., ROBISON, A., AND STICH, M. 2010. OptiX: a general purpose ray tracing engine. *ACM Trans. Graph. 29*, 4, 66:1–66:13.

ROLLAND, J., AND HUA, H. 2005. Head-mounted display systems. *Encyclopedia of Optical Engineering*, 1–13.

ROLLAND, J. P., KRUEGER, M. W., AND GOON, A. 2000. Multifocal planes head-mounted displays. *Applied Optics 39*, 19.

SHAOULOV, V., MARTINS, R., AND ROLLAND, J. P. 2004. Compact microlenslet-array-based magnifier. *Optics Letters 29*, 7.

SULLIVAN, A. 2003. A solid-state multi-planar volumetric display. In *SID Digest*, vol. 32, 207–211.

SUTHERLAND, I. E. 1968. A head-mounted three dimensional display. In *AFIPS Fall Joint Computer Conference*, 757–764.

TAKAKI, Y. 2006. High-density directional display for generating natural three-dimensional images. *Proceedings of the IEEE 94*.