

Layered 3D: Tomographic Image Synthesis for Attenuation-based Light Field and High Dynamic Range Displays

Gordon Wetzstein¹

Douglas Lanman²

Wolfgang Heidrich¹

Ramesh Raskar²

¹University of British Columbia

²MIT Media Lab



Figure 1: Inexpensive, glasses-free light field display using volumetric attenuators. (Left) A stack of spatial light modulators (e.g., printed masks) recreates a target light field (here for a car) when illuminated by a backlight. (Right) The target light field is shown in the upper left, together with the optimal five-layer decomposition, obtained with iterative tomographic reconstruction. (Middle) Oblique projections for a viewer standing to the top left (magenta) and bottom right (cyan). Corresponding views of the target light field and five-layer prototype are shown on the left and right, respectively. Such attenuation-based 3D displays allow accurate, high-resolution depiction of motion parallax, occlusion, translucency, and specularity, being exhibited by the trunk, the fender, the window, and the roof of the car, respectively.

Abstract

We develop tomographic techniques for image synthesis on displays composed of compact volumes of light-attenuating material. Such volumetric attenuators recreate a 4D light field or high-contrast 2D image when illuminated by a uniform backlight. Since arbitrary oblique views may be inconsistent with any single attenuator, iterative tomographic reconstruction minimizes the difference between the emitted and target light fields, subject to physical constraints on attenuation. As multi-layer generalizations of conventional parallax barriers, such displays are shown, both by theory and experiment, to exceed the performance of existing dual-layer architectures. For 3D display, spatial resolution, depth of field, and brightness are increased, compared to parallax barriers. For a plane at a fixed depth, our optimization also allows optimal construction of high dynamic range displays, confirming existing heuristics and providing the first extension to multiple, disjoint layers. We conclude by demonstrating the benefits and limitations of attenuation-based light field displays using an inexpensive fabrication method: separating multiple printed transparencies with acrylic sheets.

Keywords: computational displays, light fields, autostereoscopic 3D displays, high dynamic range displays, tomography

Links: [DL](#) [PDF](#) [WEB](#) [VIDEO](#)

1 Introduction

3D displays are designed to replicate as many perceptual depth cues as possible. As surveyed by Lipton [1982], these cues can be classified by those that require one eye (monocular) or both eyes (binocular). Artists have long exploited monocular cues, including perspective, shading, and occlusion, to obtain the illusion of depth with 2D media. Excluding motion parallax and accommodation, existing 2D displays provide the full set of monocular cues. As a result, 3D displays are designed to provide the lacking binocular cues of disparity and convergence, along with these missing monocular cues.

Current 3D displays preserve disparity, but require special eyewear (e.g., LCD shutters, polarizers, or color filters). In contrast, automultiscopic displays replicate disparity and motion parallax without encumbering the viewer. As categorized by Favalora [2005], such glasses-free displays include parallax barriers [Ives 1903; Kanolt 1918] and integral imaging [Lippmann 1908], volumetric displays [Blundell and Schwartz 1999], and holograms [Slinger et al. 2005]. Holograms present all depth cues, but are expensive and primarily restricted to static scenes viewed under controlled illumination [Klug et al. 2001]. Research is addressing these issues [Blanche et al. 2010], yet parallax barriers and volumetric displays remain practical alternatives, utilizing well-established, low-cost fabrication. Furthermore, volumetric displays can replicate similar depth cues with flicker-free refresh rates [Favalora 2005].

This paper considers automultiscopic displays comprised of compact volumes of light-attenuating material, which we dub “Layered 3D” displays. Differing from volumetric displays with light-emitting layers, overlaid attenuation patterns allow objects to appear beyond the display enclosure and for the depiction of motion parallax, occlusion, and specularity. While our theoretical contributions apply equally well to dynamic displays, such as stacks of liquid crystal display (LCD) panels, our prototype uses static printing to demonstrate the principles of tomographic image synthesis. Specifically, we produce multi-layer attenuators using 2D printed transparencies, separated by acrylic sheets (see Figures 1 and 2).

1.1 Contributions

We establish the theoretical and practical limitations of volumetric, attenuation-based displays and present techniques to optimize their performance. Specific contributions are as follows:

- We demonstrate the practical benefits of multi-layered attenuators for light field display, compared to dual-layer devices. We also establish theoretical limitations of all such displays.
- We show how tomographic principles allow light fields to be optimally displayed by volumetric or layered attenuators.
- We demonstrate that multi-layer displays achieve higher resolution, extended depth of field, and increased brightness, as compared to conventional parallax barriers.
- We show our tomographic method also encompasses 2D HDR display. We demonstrate the first HDR display using multiple, disjoint attenuators. For the case of dual-layer architectures, we confirm existing heuristic algorithms [Seetzen et al. 2004].

1.2 Overview of Benefits and Limitations

We summarize our interpretation of the relative benefits and limitations of our approach in Table 1. Unlike many volumetric displays, we exploit multiplicative light absorption across multiple layers, rather than additive light emission. Such spatially-varying attenuation is inexpensively fabricated, without moving parts, using either *static* 2D or 3D printing, or by layering *dynamic* spatial light modulators (e.g., LCD, LCoS, or DMD). Modulation of light allows objects to appear beyond the display and for depiction of occlusion and specularity. We present an optimal decomposition of light fields into two or more static layers, improving upon the method of Lanman et al. [2010]. Compared to parallax barriers, our multi-layer generalization enhances resolution, increases depth of field, and improves dynamic range. Finally, we give the first non-heuristic construction of HDR displays with two or more disjoint layers.

Volumetric attenuators share the limitations of other multi-layer displays, particularly increased cost and complexity compared to monolithic or dual-layer designs. As described by Bell et al. [2008; 2010], multi-layer LCDs exhibit decreased brightness, moiré, and color crosstalk, with additional layers exacerbating problems. Similar limitations are expected with other spatial light modulators. If fabricated with 2D/3D printing, fidelity is restricted by limited-contrast media and by scattering, misalignment, and interreflections. Similar to other automultiscopic displays, including parallax barriers and integral imaging, our design exhibits a finite depth of field. Most significantly, benefits of volumetric attenuators are best realized by simultaneously increasing the number of layers and the display thickness. Finally, our constrained linear least-squares optimization requires solving a large, but sparse, system of equations; currently, solutions cannot be found in real-time, limiting applications to static images or to pre-rendered content.

2 Related Work

Blundell and Schwartz [1999] define a volumetric display as permitting “the generation, absorption, or scattering of visible radiation from a set of localized and specified regions within a physical volume”. As such, our proposed attenuation-based display falls within this broad category; yet, unlike most volumetric displays, layers interact in a multiplicative rather than additive manner [Kooi and Toet 2003]. Through modulation, spatio-angular frequencies are created that correspond with surfaces lying outside the display enclosure, but with decreased brightness. Here we compare our approach to existing additive and multiplicative volumetric displays.



Figure 2: *Prototype multi-layer display. (Left) We fabricate a multi-layer display by separating transparencies with acrylic sheets and back-illuminating with a light box (e.g., an LCD panel). (Right) Printed transparencies and acrylic layers.*

	Volumetric	Integral Imaging	Parallax Barriers	Multi-Layer
2D Resolution	high	low	low	high
3D Resolution	high	moderate	moderate	high
Brightness	high	high	low	moderate-high
Contrast	moderate	moderate	moderate	high
Complexity	high	low	low	moderate
Flip Animations	no	yes	yes	low resolution*

Table 1: *Benefits of multi-layer displays. We enable sharper, brighter images than existing automultiscopic displays, but require additional layers and possibly thicker enclosures. *Flip animations, with uncorrelated multi-view imagery, are handled by enforcing parallax-barrier-like downsampling, described in Section 7.1.*

Additive Volumetric Displays. Many volumetric displays exploit high-speed projection synchronized with mechanically-rotated screens. Such swept volume displays were proposed as early as 1912 [Favalora 2005] and have been continuously improved [Cossairt et al. 2007]. While requiring similar mechanical motion, Jones et al. [2007] instead achieve light field display, preserving accurate perspective and occlusion cues, by introducing an anisotropic diffusing screen and user tracking. Related designs include the Seelinder by Yendo et al. [2005], exploiting a spinning cylindrical parallax barrier and LED arrays, and the work of Maeda et al. [2003], utilizing a spinning LCD panel with a directional privacy filter. Several designs have eliminated moving parts using electronic diffusers [Sullivan 2003], projector arrays [Agocs et al. 2006], and beam-splitters [Akeley et al. 2004]. Whereas others consider projection onto transparent substrates, including water drops [Barnum et al. 2010], passive optical scatterers [Nayar and Anand 2007], and dust particles [Perlin and Han 2006].

Display Fabrication. Our prototype is inspired in part by recent work on fabricating spatially-varying reflectance [Matusik et al. 2009], subsurface scattering [Dong et al. 2010; Hašan et al. 2010], and binary opacity [Mitra and Pauly 2009]. We explore the benefits of similarly constructing continuously-varying volumetric attenuation. While this could be achieved directly with 3D printers [Sagi 2009], we propose a low-cost alternative: separating printed masks with acrylic sheets (see Figure 2). By generalizing parallax barriers, such volumetric attenuators improve spatio-angular resolution for 3D display and enhance contrast for high dynamic range (HDR) imaging. Such attenuation-based displays are shown to be inexpensive alternatives for passive 3D signage, where large depth of field, high dynamic range, and glasses-free viewing are desired.

Multiplicative Volumetric Displays. We are not the first to propose multi-layer, attenuation-based displays. To our knowledge, Loukianitsa and Putilin [2002; 2006] were the first to consider 3D displays with three or more attenuating layers. Their optimization differs from ours, instead employing neural networks. Mitra and Pauly [2009] introduce the related problem of shadow art creation, wherein a single volume, with binary-valued opacity, is optimized to project a target set of shadows. Our system addresses similar

issues, including resolving inconsistent views, instead using least-squares optimization appropriate for light fields, rather than the as-rigid-as-possible deformations required for visual hull reconstruction from shadows. In a closely-related work, Gotoda [2010] introduces a similar optimization scheme. In comparison, we support continuously-varying volumetric attenuators with a tomographic formulation, document our constrained least-squares optimization for 3D display and HDR imaging, build a prototype, and establish the upper bound on depth of field for multi-layer displays. Lanman et al. [2010] introduce content-adaptive parallax barriers, optimizing dual-layer displays with temporally-varying attenuation found with non-negative matrix factorization. In comparison, our iterative tomographic reconstruction provides the optimal rank-1 decomposition, in the two-layer case, guaranteed to converge to the global, rather than a local, minimum. In contrast to simple layered displays, similar to earlier multiplane cameras [Disney 1940], we depict continuously-varying depth both within and outside the display.

3 Tomographic Image Synthesis

This section describes how volumetric attenuators are optimally constructed to emit a target light field using tomographic principles. In this paper we present our analysis in flatland, with a straightforward extension to 3D volumes and 4D light fields. First, we consider the forward problem: modeling the light field emitted by a backlit volumetric attenuator. The logarithm of the emitted light field is shown to equal to the negative Radon transform of the attenuation map. Second, we consider the inverse problem: synthesizing an attenuation map to approximate a target light field. We find the optimal solution, in the least-squares sense, using a series expansion method based on iterative tomographic reconstruction principles. Third, we describe how to apply these principles to the generation of images for *layered* attenuation displays.

3.1 Modeling Volumetric Attenuation

In flatland, a 2D volumetric attenuator is modeled by a continuously-varying attenuation map $\mu(x, y)$, such that the intensity I of a transmitted light ray \mathcal{C} is given by the Beer-Lambert law

$$I = I_0 e^{-\int_{\mathcal{C}} \mu(r) dr}, \quad (1)$$

where I_0 is the incident intensity [Hecht 2001]. Additional scattering and reflection losses are assumed to be negligible. For convenience, the logarithm of the normalized intensity \bar{I} is defined as

$$\bar{I} = \ln\left(\frac{I}{I_0}\right) = -\int_{\mathcal{C}} \mu(r) dr. \quad (2)$$

This section considers a volumetric attenuator composed of a single slab, of width w and height h , such that $\mu(x, y)$ can be non-zero only within the interval $|x| < w/2$ and $|y| < h/2$. A relative two-plane light field parameterization $l(u, a)$ is adopted [Chai et al. 2000]. As shown in Figure 3, the u -axis is coincident with the x -axis, which bisects the slab horizontally. The orientation of ray (u, a) is defined by the slope $a = s - u = d_r \tan(\theta)$, where d_r is the distance of the s -axis from the u -axis.

In conventional parallel beam tomography [Kak and Slaney 2001], the Radon transform $p(u, a)$ encodes all possible line integrals through the attenuation map, along each ray (u, a) , such that

$$p(u, a) = \int_{-\frac{h}{2}}^{\frac{h}{2}} \int_{-\frac{w}{2}}^{\frac{w}{2}} \mu(x, y) \delta(d_r(x - u) - ay) dx dy, \quad (3)$$

where $\delta(\xi)$ denotes the Dirac delta function. Substituting into Equation 2 gives the following expression for the light field $l(u, a)$

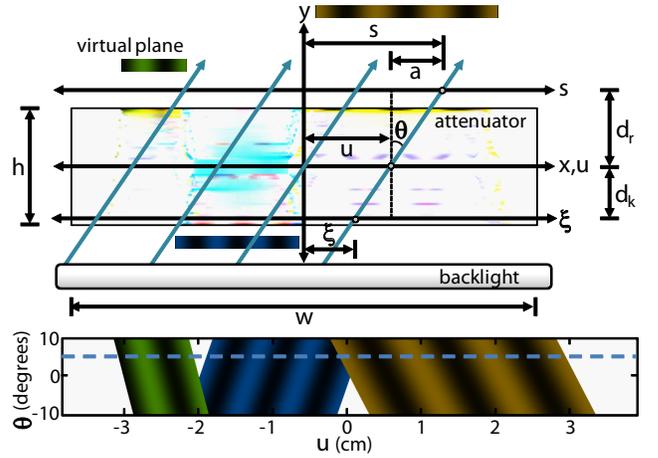


Figure 3: Tomographic analysis of attenuation-based displays. (Top) A volumetric attenuator $\mu(x, y)$ is optimized, using Equation 11, to emit a light field approximating the virtual planes. (Bottom) The target light field $l(u, a)$, with the dashed line denoting the oblique projection with rays passing through the attenuator above.

emitted when a volumetric attenuator is illuminated by a backlight producing the incident light field $l_0(u, a)$:

$$\bar{l}(u, a) = \ln\left(\frac{l(u, a)}{l_0(u, a)}\right) = -p(u, a). \quad (4)$$

In practice, backlights produce uniform illumination so $l_0(u, a) = l_{\max}$ and the light field is normalized so $l(u, a) \in (0, l_{\max}]$. To summarize, tomographic analysis reveals a simple forward model: the logarithm of the emitted light field is equivalent to the negative Radon transform of the attenuation map. For a fixed linear angle $a = a_0$, a 1D slice of the Radon transform $p(u, a_0)$ corresponds to an oblique projection of the attenuation map and, correspondingly, to an emitted oblique view $\bar{l}(u, a_0)$, as shown in Figure 3.

3.2 Synthesizing Light Fields

With parallel beam tomography, an estimate of the attenuation map $\tilde{\mu}(x, y)$ is recovered from the projections $p(u, a)$ using the inverse Radon transform, conventionally implemented using the filtered backprojection algorithm [Kak and Slaney 2001]. A direct application of this algorithm yields an estimate for a volumetric attenuator capable of emitting the target light field $\bar{l}(u, a)$:

$$\tilde{\mu}(x, y) = -\int_{-\infty}^{\infty} \bar{l}(x - (a/d_r)y, a) da. \quad (5)$$

Geometrically, the spatially-filtered oblique views $\bar{l}(u, a)$ are propagated through the attenuation volume, along the rays in Equation 3. However, a high-pass filter $\hat{h}(f_u)$ must be applied first to obtain a sharp estimate of $\mu(x, y)$, where the hat symbol denotes the 1D Fourier transform [Bracewell and Riddle 1967]. This filter is implemented in the ray domain or frequency domain, as follows.

$$\bar{l}(u, a) = \bar{l}(u, a) \otimes h(u) = \int_{-\infty}^{\infty} \hat{h}(f_u) \hat{\bar{l}}(f_u, a) e^{2\pi j f_u u} df_u \quad (6)$$

For this application, the oblique views are known *a priori*, so an ideal ramp filter $\hat{h}(f_u) = |f_u|$ should be used.

Equations 5 and 6 fail to provide a practical method for constructing volumetric attenuators. First, high-pass filtering often leads to

negative attenuations, prohibiting fabrication. Similar to the binary attenuation volumes considered by Mitra and Pauly [2009], the target set of oblique views may be inconsistent with any single attenuation map. While filtered backprojection can accommodate minor inconsistencies due to measurement artifacts, gross inconsistencies introduce negative attenuations. Second, filtered backprojection requires projections vary over a full hemisphere (i.e., $-\infty < a < \infty$). In practice, oblique views may be known only over a limited set of angles (e.g., if the light field is rendered or captured with cameras along a limited baseline). In this case, alternative methods are required to produce a sharp estimate of the attenuation map.

Iterative reconstruction algorithms present a flexible alternative to traditional transform methods [Herman 1995]. While exhibiting greater computational complexity, such methods better account for inconsistent projections over limited angles. We consider a particular *series expansion method*, for which attenuation is modeled by a linear combination of N_b non-negative basis functions $\bar{\phi}_k(x, y)$:

$$\mu(x, y) = \sum_{k=1}^{N_b} \alpha_k \bar{\phi}_k(x, y). \quad (7)$$

The basis can be conventional voxels used in volume rendering [Drebin et al. 1988], or general functions with compact or extended support. Our choice of $\bar{\phi}_k(x, y)$ is discussed in Section 6.1.

Substituting Equation 7 into Equations 3 and 4 gives the following solution to the forward rendering problem.

$$\bar{l}(u, a) = - \sum_{k=1}^{N_b} \alpha_k \int_{-\frac{h}{2}}^{\frac{h}{2}} \int_{-\frac{w}{2}}^{\frac{w}{2}} \bar{\phi}_k(x, y) \delta(d_r(x-u) - ay) dx dy \quad (8)$$

This expression leads to a linear system of equations, when considering a discrete light field \bar{l}_{ij} , such that

$$\bar{l}_{ij} = - \sum_{k=1}^{N_b} \alpha_k P_{ij}^{(k)}, \quad (9)$$

where (i, j) are the discrete indices corresponding to the continuous coordinates (u, a) . The projection matrix $P_{ij}^{(k)}$ is given by

$$P_{ij}^{(k)} = \int_{-\frac{h}{2}}^{\frac{h}{2}} \int_{-\frac{w}{2}}^{\frac{w}{2}} \bar{\phi}_k(x, y) \delta(d_r(x - (i\Delta u)) - (j\Delta a)y) dx dy, \quad (10)$$

corresponding to line integrals through every basis function k along each ray (i, j) . This system is expressed in matrix-vector form as $\mathbf{P}\alpha = -\bar{\mathbf{l}} + \bar{\mathbf{e}}$, where $\bar{\mathbf{e}}$ is the approximation error. As surveyed by Herman [1995], a wide variety of iterative reconstruction algorithms exist to solve this system, primarily differing in computational complexity and in constraints placed on the error $\bar{\mathbf{e}}$. However, by the Weber-Fechner law, the human visual system responds approximately linearly to logarithmic changes in illumination [Reinhard et al. 2010]. As a result, we cast attenuation map synthesis as the following non-negative linear least-squares problem.

$$\arg \min_{\alpha} \|\bar{\mathbf{l}} + \mathbf{P}\alpha\|^2, \text{ for } \alpha \geq 0 \quad (11)$$

Although requiring iterative optimization, this formulation as a convex optimization problem yields an optimal attenuation map, in the least-squares sense, that emits a target light field with consistent views. We observe that this problem is efficiently solved using optimization methods described in detail in Section 6.1.

3.3 Layered Attenuation-based Displays

So far, we have considered an attenuating volume with either a continuous spatially-varying absorption coefficient, or an expansion into a discrete set of basis functions that are uniformly distributed over the display volume. While such volumes could be fabricated with recent rapid prototyping hardware, other manufacturing processes, such as stacks of LCD panels or semi-transparent slides, are better represented as a finite number of discrete attenuation layers.

Our analysis naturally extends to such multi-layered attenuators. Rather than directly constructing attenuation maps, each mask controls spatially-varying transmittance in a single plane. Following Figure 3, ray (u, a) is modulated by N_l layers such that

$$l(u, a) = l_0(u, a) \prod_{k=1}^{N_l} t_k(u + (d_k/d_r)a), \quad (12)$$

where $t_k(\xi)$ is the transmittance of mask k (separated by a distance d_k). Taking the logarithm gives the forward model

$$\bar{l}(u, a) = \sum_{k=1}^{N_l} \ln t_k(u + (d_k/d_r)a) = - \sum_{k=1}^{N_l} a_k(u + (d_k/d_r)a), \quad (13)$$

where $a_k(\xi) = -\ln t_k(\xi)$ is the absorbance. Similar to Equation 9, the linear system $\bar{l}_{ij} = - \sum_{k=1}^{N_l} a_k P_{ij}^{(k)}$ is obtained for a



Figure 4: Multi-layer 3D display. The “dice” scene is rendered for vantage points to the right and left of the display, shown at the top and bottom of the left column, respectively. Corresponding views of the five-layer prototype (see Figure 2) are compared to the right. Inset figures denote the position of layers relative to the scene. Unlike conventional additive volumetric displays, both transparency and occlusion are accurately depicted, allowing for faithful representation of motion parallax for opaque objects. (Bottom) Clockwise from the upper left are two, three, and five layer decompositions, ordered from the front to the back of the display. All layers are uniformly-spaced and span the same total display thickness.

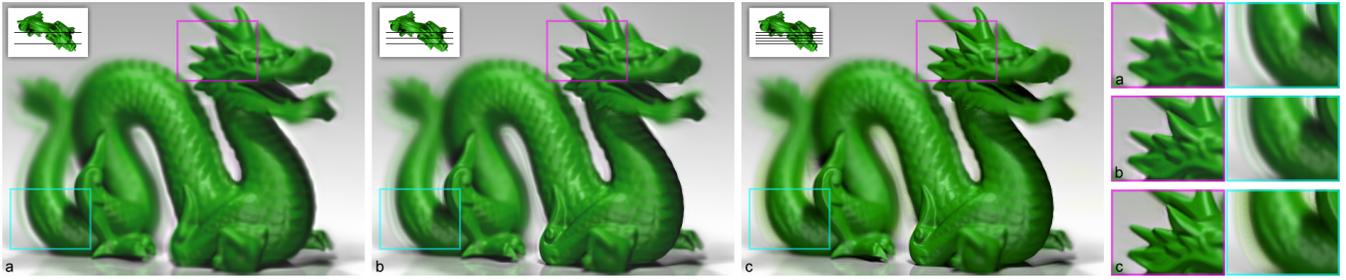


Figure 5: Influence of multi-layer depth of field on reconstruction artifacts. Shown from left to right are reconstructions of the “dragon” scene, using two, three, and five layers, seen when viewing directly in front. Magnified regions are compared on the right. The magenta region, located on the head and inside the display, is rendered with increasing resolution as additional layers are incorporated. The cyan region, located on the tail and behind the display, exhibits noticeable halo artifacts, similar to the dice in Figure 4. As described in Section 4.2, the upper bound on depth of field for multi-layer displays indicates spatial resolution is inversely proportion to scene depth. Since high spatial frequencies are required to depict the tail edge, artifacts result without proper prefiltering. In this example, an approximation of the prefilter in Equation 17 is applied assuming a five-layer decomposition, allowing artifacts to persist in the two-layer and three-layer decompositions.

discrete set of rays (i, j) . For multi-layered attenuators, the form of the projection matrix $P_{ij}^{(k)}$ is modified, now encoding the intersection of every ray with each mask. Thus, a similar optimization solves the inverse problem of constructing an optimal multi-layered attenuator. Practically, however, layers have a finite contrast (i.e., maximum transmissivity and opacity) and Equation 11 is solved as a *constrained* least-squares problem. As an additional benefit, our optimization encompasses additive decompositions, simply by interpreting \bar{l}_{ij} as the light field, rather than its logarithm, and a_k as the negative of the emittance. Section 6 describes additional issues that must be addressed for accurate fabrication (e.g., handling limited contrast and color gamut, achieving accurate mechanical alignment, and mitigating scattering and reflection).

4 Application to 3D Display

This section assesses multi-layer light field decompositions for automultiscopic 3D display. First, we document the qualitative performance of our tomographic algorithm, providing intuition into its behavior and how design parameters influence reconstruction accuracy. Second, we establish the quantitative upper bound on depth of field for all multi-layer displays, informing system design and motivating display prefiltering [Zwicker et al. 2006]. Third, through experimental studies, we develop rules for optimizing design parameters, including the number of layers and display thickness, to minimize artifacts. We conclude by evaluating scenes with varying degrees of disparity, occlusion, translucency, and specularly.

4.1 Assessing Performance for 3D Display

Consider the “dice” scene in Figure 4. Each die is approximately 1.5 cm on a side, with the scene extending 5 cm in depth. In this example and for all others in this paper and using the prototype, we assume a 5.7 cm \times 7.6 cm display with a thickness of 1.25 cm, with evenly-spaced layers. All layers have a resolution of 171 dots per inch (i.e., 149 μ m pixels). The target scene is rendered as a light field with 7×7 oblique projections, spanning a field of view of ± 5 degrees from the display surface normal. Following Figure 3, the light field is parameterized with the u -axis bisecting the middle of the display and the s -axis coincident with the front layer. The scene is transformed so the red die is enclosed within the display, with other dice extending beyond the surface.

We obtain multi-layer decompositions using the tomographic algorithm in Section 3.3. Figure 4 shows masks for two, three, and five layers. First, we observe that 3D objects can be displayed both

inside and outside the enclosure. This illustrates the primary benefit of multiplicative displays over conventional additive volumetric displays: through modulation, spatio-angular frequencies are created corresponding to objects outside the display. Modulation also allows occlusion to be accurately depicted. Second, objects inside or near the display are rendered at full-resolution and with the same brightness as the target light field, representing the primary benefits compared to conventional automultiscopic displays (e.g., parallax barriers and integral imaging). Third, although intentionally rendered with a finite depth of field, halos appear around objects outside the enclosure, with additional layers mitigating these errors.

As shown in Figure 4, optimized masks exhibit predictable structure. Although not produced using filtered backprojection, the qualitative performance of iterative reconstruction can best be anticipated using this simple procedure; if applied, each view would first be sharpened, using Equation 6, and then smeared through the layers to assign absorbance by Equation 5. For a point on a virtual object within the display, a sharp image forms on the closest layer, since smeared views align there, and defocused images form on layers above or below. This is observed for the red die in Figure 4 and the wheel of the car in Figure 1, both appearing sectioned over the layers spanning their physical extent. For objects inside, our decomposition functions as the multiplicative equivalent to depth filtering in additive multi-layer displays [Akeley et al. 2004; Suyama et al. 2004]. However, iterative reconstruction enforces physical constraints on attenuation, resolves inconsistencies between views in a least-squares sense, and constructs attenuation patterns to illustrate objects beyond the display (e.g., the cyan and yellow dice).

4.2 Characterizing Depth of Field

The *depth of field* of an automultiscopic display characterizes the maximum spatial frequency that can be depicted, without aliasing, in a plane parallel to the display at a given distance. As described by Zwicker et al. [2006], depth of field is determined by the spectral properties of a display. For parallax barriers and integral imaging considered in flatland, discrete sampling of rays (u, a) produces a spectrum $\hat{l}(f_u, f_a)$ limited to a rectangle. Following Chai et al. [2000], the spectrum of a Lambertian surface, located a distance d_o from the middle of the display, is the line $f_a = (d_o/d_r)f_u$. Thus, the spatial cutoff frequency is found by intersecting this line with the spectral bandwidth. For parallax barriers and integral imaging, Zwicker et al. [2006] use this construction

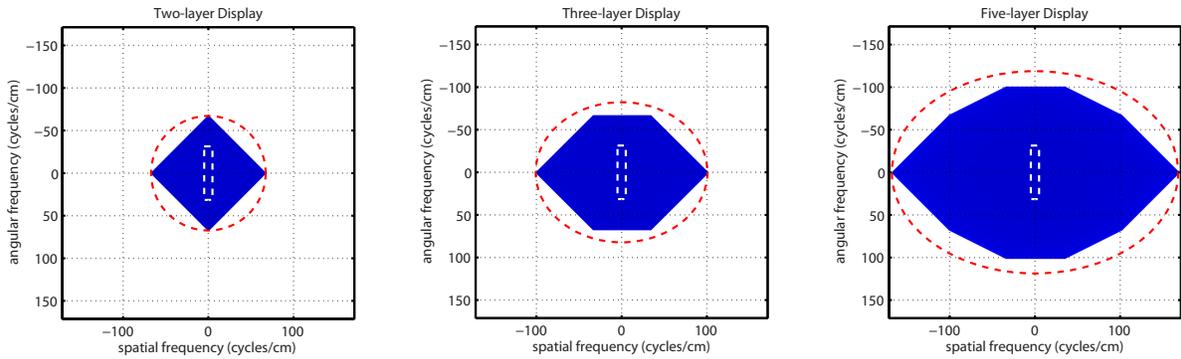


Figure 6: Spectral support for multi-layer displays. The spectral support (shaded blue) is illustrated for two-layer (left), three-layer (middle), and five-layer (right) displays, evaluated using the geometric construction given by Equation 15. Note that the shaded area indicates the achievable region of non-zero spectral support, with a detailed derivation in Appendix E of the supplementary material. The system parameters correspond with the prototype in Section 4.1, with the variable $d_r = h/2$. The ellipse corresponding to the upper bound on achievable spatio-angular frequencies is denoted by a dashed red line. Note that the spectral support of a multi-layer display exceeds the bandwidth of a conventional, two-layer automultiscopic display with a similar physical extent, shown as a dashed white line and given by Equation 14.

to show the spatial frequency f_ξ in a plane at d_o must satisfy

$$|f_\xi| \leq \begin{cases} \frac{f_0}{N_a}, & \text{for } |d_o| + (h/2) \leq N_a h \\ \left(\frac{h}{(h/2) + |d_o|} \right) f_0, & \text{otherwise} \end{cases}, \quad (14)$$

where N_a is the number of views, h is the display thickness, and $f_0 = 1/2p$ is the cutoff frequency for layers with pixels of width p .

As shown in Figure 5 and observed by Gotoda [2010], multi-layer displays exhibit finite depth of field, which, to date, has not been quantitatively described. We observe the upper bound on depth of field is found by similarly considering the maximum spectral bandwidth achievable with multiple layers. The Fourier transform of Equation 12 expresses the spectrum of any multi-layer display as

$$\hat{l}(f_u, f_a) = \bigotimes_{k=1}^{N_l} \hat{t}_k(f_u) \delta(f_a - (d_k/d_r)f_u), \quad (15)$$

where \bigotimes denotes repeated convolution (see Figure 6). Here, the backlight is uniform, such that $l_0(u, a) = 1$, and the light field is normalized such that $l(u, a) \in (0, 1]$. The upper bound on depth of field is found by intersecting the line $f_a = (d_o/d_r)f_u$ with the boundary of maximum spectral support given by Equation 15, using the fact that each mask spectrum $\hat{t}_k(f_\xi)$ has an extent of $\pm f_0$.

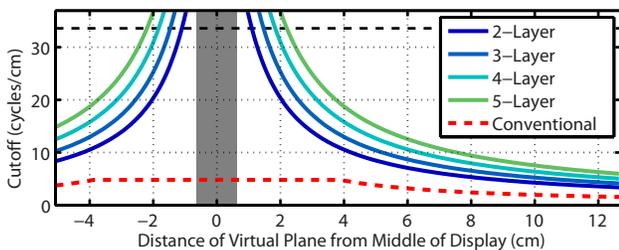


Figure 7: Upper bound on multi-layer depth of field. The spatial cutoff frequency is shown for conventional parallax barriers and integral imaging, using Equation 14, and for N_l -layer displays, using Equation 17. Parameters correspond with the prototype. Spatial resolution exceeds conventional architectures, particularly near or within the display enclosure (shaded region). The maximum spatial resolution of a single mask is denoted by a horizontal dashed line, indicating full-resolution display is possible within the enclosure.

For two layers, this construction yields the upper bound

$$|f_\xi| \leq \left(\frac{h}{(h/2) + |d_o|} \right) f_0. \quad (16)$$

Comparing Equations 14 and 16 indicates that parallax barriers and integral imaging, both of which employ fixed spatio-angular trade-offs, achieve the optimal resolution for objects located far from the display, but reduce resolution for objects close to the display by up to a factor of N_a . As shown in this section, our decompositions more fully realize the achievable spatio-angular bandwidth, obtaining higher-resolution images nearby.

The upper bound for multiple layers is assessed by similar methods, with the geometric construction providing the exact upper bound. However, we observe that repeated convolution of N_l mask spectra $\hat{t}_k(f_u, f_a)$, each with extent $|f_u| \leq f_0$ and constrained to the line $f_a = (d_k/d_r)f_u$, converges to a bivariate, zero-mean Gaussian distribution by the central limit theorem [Chaudhury et al. 2010]. As derived in detail in Appendix E of the supplementary material, the covariance matrix of this distribution is equal to the sum of the covariance matrices for each mask. Thus, contours of the light field spectrum $\hat{l}(f_u, f_a)$ will be ellipses. As before, intersecting the line $f_u = (d_o/d_r)f_a$ with the ellipse bounding the spectral bandwidth gives an approximate expression for the upper bound, as follows.

$$|f_\xi| \leq N_l f_0 \sqrt{\frac{(N_l + 1)h^2}{(N_l + 1)h^2 + 12(N_l - 1)d_o^2}} \quad (17)$$

Figure 7 compares the upper bound for multi-layer and conventional displays. As before, this upper bound indicates the potential to increase the resolution for objects close to the display; yet, even in the upper bound, multi-layer displays exhibit a finite depth of field similar to existing automultiscopic displays. For distant objects, resolution remains inversely proportional to object depth.

4.3 Optimizing Display Performance

To construct a practical multi-layer display, such as the prototype in Figure 2, one must select two key design parameters: the total thickness h and the number of layers N_l , where layers are assumed to be uniformly distributed such that $d_k \in [-h/2, h/2]$. The upper bound on the depth of field informs selection of h and N_l , yet, with the proposed optimization algorithm, further experimental assessment is required for clear design rules—since the upper bound may

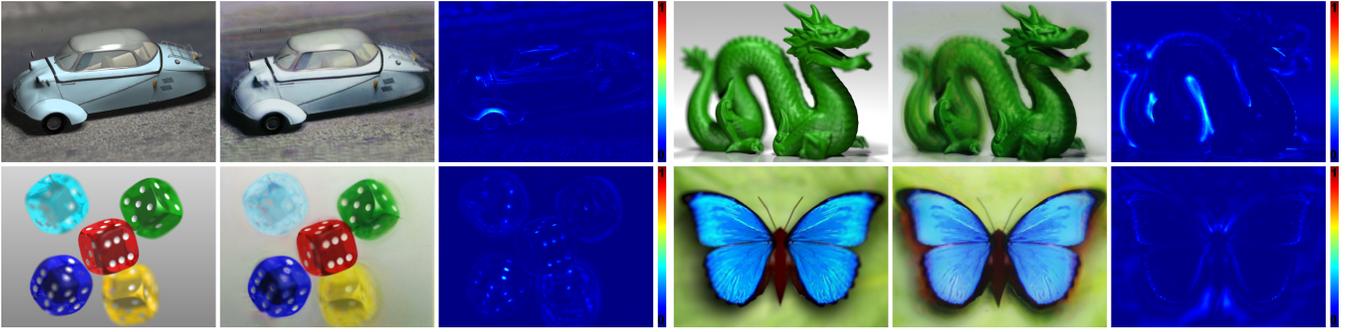


Figure 8: Multi-layer display performance. For each scene from left to right, a direct view of the target light field is compared to a prototype photograph and the absolute error between simulated and target views. Due to proper prefiltering, artifacts are evenly distributed in depth, occurring mostly near high-contrast edges (e.g., dragon silhouette and car fender). Note the specular highlight on the car roof is preserved, with reduced contrast, together with translucency of the windows. Additional results are included in the supplementary material and video.

not be achievable for all scenes. As shown in Figure 8, optimization increases spatial resolution compared to conventional displays, but also introduces artifacts. In the remainder of this section, we analyze how display parameters are selected to minimize artifacts.

For a conventional parallax barrier, with pixel width p , field of view α , and N_a views, the separation h_b between the layers is

$$h_b = \frac{N_a p}{2 \tan(\alpha/2)}, \quad (18)$$

with $h_b = 0.6$ cm for the display and light field parameters listed in Section 4.1. As shown by the red line in Figure 7, we expect parallax barriers to create lower-resolution images than multi-layer decompositions. This is confirmed in Figure 13. However, the upper bound does not indicate whether a high peak signal-to-noise ratio (PSNR) is obtained for a given display configuration.

As shown in Figure 9, a database of light fields facilitates display optimization. Views are rendered with a manually-selected depth of field, approximating combined light field anti-aliasing and display prefilers (see Section 6.1). Several observations can be made regarding general design principles. First, PSNR is maximized by enclosing the scene within the display (for a sufficiently large number of layers). Thus, multi-layer displays can be operated in a mode akin to additive volumetric displays, wherein high resolution is achieved for contained objects. However, particularly for mobile applications, displays must be thinner than the depicted volume. Second, addressing this, we find, for a fixed display thickness (e.g.,

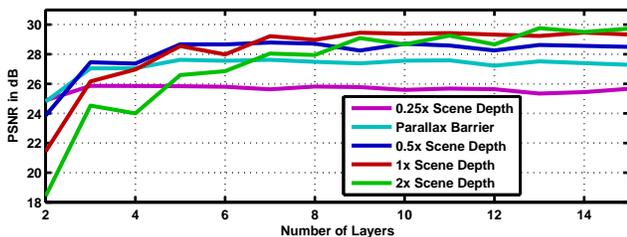


Figure 9: Minimizing artifacts with optimal display designs. To determine rules for optimizing the thickness h and number of layers N_l , our decomposition was applied to four scenes: “butterfly”, “dice”, “dragon”, and “car”. As shown above, the average PSNR is plotted as a function of N_l and h . For a fixed thickness, a finite number of masks closely approximates the maximum-achievable PSNR, obviating fabrication of continuously-varying attenuation.

that of a conventional parallax barrier), addition of layers increases PSNR. However, artifacts persist even with a large number of layers. Thus, our prototype closely approximates the performance of volumetric attenuators, despite relatively few layers. Third, for a fixed number of layers N_l , there is an optimal display thickness determined by the desired depth range. In summary, tomographic image synthesis obtains high-PSNR reconstructions with small numbers of static layers in compact enclosures. Through such simulations, the optimal design parameters h and N_l can be determined depending on form factor or system complexity, respectively, subject to image fidelity and depth range requirements.

5 Application to HDR Display

Our tomographic analysis shows how displays composed of two or more layers can present 4D light fields. Multi-layer LCD panels are currently beginning to enter the consumer market, for example in the form of the Nintendo 3DS parallax barrier display [Jacobs et al. 2003]. Once such displays are available, they can not only be used for 3D display, but also for increasing the dynamic range of 2D images [Seetzen et al. 2004; Reinhard et al. 2010].

With non-negligible separations between attenuators, multi-layer HDR decomposition becomes a 3D display problem, since all viewpoints must produce an accurate rendition of the 2D image within the target field of view. Our constrained tomographic solver inherently accounts for the limited contrast of each layer, thereby allowing simultaneous optimization of dynamic range and accurate multi-view imagery. In a 2D HDR display mode, the target light field encodes a single plane (e.g., coincident with the front layer), with a texture given by the desired HDR image. Figure 10 shows the result from a parallax-free 2D HDR display prototype. The optimized layers in the lower row account for the non-zero black level of the printing process and are scaled appropriately before printing. Note our tomographic algorithm naturally handles decomposition into more than two disjoint layers; consult the video and supplementary material for comparisons with differing numbers of layers.

In the prototype, the target light field has been optimized for a grid of 7×7 viewpoints, all showing the same HDR image. As described in Section 6, each layer is printed on a low-contrast inkjet transparency. For all experiments in this section, we follow a standard procedure for HDR image display, wherein the dynamic range is expanded for the luminance channel only [Seetzen et al. 2004; Reinhard et al. 2010], with chrominance assigned to the front layer after optimization. Although color contrast is limited in this fashion, the human visual system is more sensitive to luminance. In practice,



Figure 10: Multi-layer, parallax-free HDR image display. The target HDR image is shown, with tone mapping, on the upper left. Photographs of a multi-layer HDR display prototype, incorporating one, two, and three layers (each with a contrast of 3.3:1), are shown along the upper row; note how image contrast is improved by applying our tomographic image synthesis method together with multiple, physically-disjoint attenuation layers. The absolute errors between the simulated reconstructions and the target image are shown on the bottom left.

such luminance-chrominance decompositions reduce system complexity, since monochromatic displays can be used for underlying layers, mitigating color crosstalk and moiré.

The optimized layers, as evaluated using Equation 11 and shown in Figure 10, suggest a trend for optimal constructions: the front layer is a sharpened target image and underlying layers appear blurred. This configuration is effective since a sharpened front layer preserves high spatial frequencies in a view-independent manner, with blurred underlying layers similarly enhancing dynamic range despite changes in viewpoint. The degree of blur is determined by the parallax, as defined by the field of view and layer spacing. Note that Seetzen et al. [2004] originally motivated blurring the back layer as a means to tolerate misalignment between layers; we observe that parallax due to spaced layers produces similar alignment issues. The solution in both cases is a low resolution or blurred back layer. Hence, our results are consistent with existing methods for dual-modulation HDR image display, pairing high-resolution LCD front panels with low-resolution LED backlights. As described by Seetzen et al. [2004], such displays use a “heuristic” image synthesis method motivated by both physiological and hardware constraints: the target HDR image is blurred to the resolution of the rear layer, while the front panel displays a compensated sharpened image.

Figure 11 compares this heuristic approach to our tomographic optimization, concluding that qualitatively-similar patterns are produced for dual-layer decompositions—despite significant differences between LED distributions and our assumed high-resolution backlight. This indicates existing HDR displays are near-optimal, in the least-squares sense, and further advocates for low-resolution, spatially-programmable backlighting.

Additional layers increase dynamic range, yet not equally well for all spatial frequencies. Optimal performance occurs with no layer separation, but may not be practically achievable. As separation increases, depicting high spatial frequencies at high contrast becomes difficult over the full field of view. This is a fundamental limitation of multi-layered HDR displays, as well as existing dual-modulation architectures. In Figure 12, we characterize this effect by considering the light field of a plane, coincident with the front layer, containing maximum-contrast sinusoidal textures. This figure charts the modulation transfer function (MTF) [Hecht 2001]: the achieved Michelson contrast, divided by maximum possible contrast, as a function of target spatial frequency. Contrast is averaged over all light field views. We conclude that building multi-layer displays capable of both 3D and 2D HDR modes involves a careful interplay

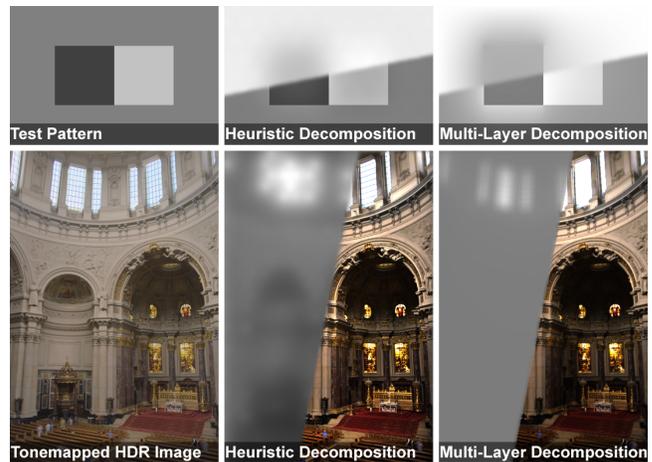


Figure 11: Heuristic vs. tomographic HDR image synthesis. We compare our tomographic method, assuming a high-resolution backlight, to heuristic algorithms [Seetzen et al. 2004] designed to operate with a low-resolution backlight. The backlight and front panel images are divided into upper-left and lower-right halves in each example, respectively. Note our optimization produces similar patterns to prior heuristics. This confirms existing architectures achieve near-optimal results for HDR display, while providing a generalization to multiple layers.

between design constraints; large depth of field for 3D applications necessitates larger gaps, limiting the field of view and maximum spatial frequency for which 2D HDR display is achieved.

6 Implementation

This section describes our attenuation-based display prototypes, developing software for tomographic image synthesis and hardware for multi-layer displays comprising printed transparencies.

6.1 Software

Light fields are rendered with POV-Ray as 7×7 oblique projections within a field of view of ± 5 degrees. A depth-dependent 4D anti-aliasing filter is applied [Levoy and Hanrahan 1996]; this is practi-

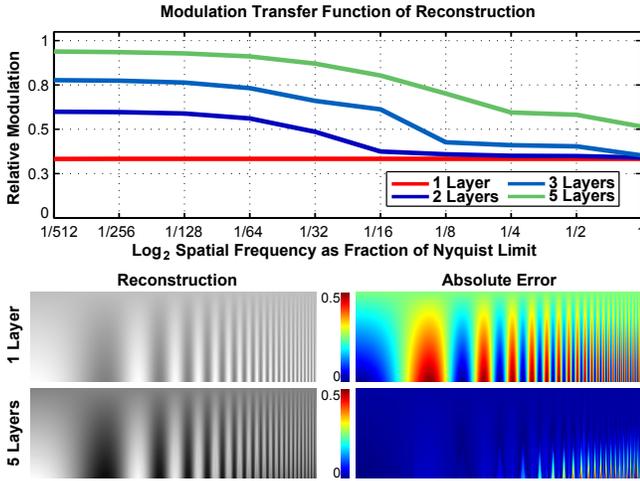


Figure 12: Contrast-resolution tradeoff in multi-layer HDR. (Top) MTF plots illustrating the reconstructed contrast, for sinusoids rendered with maximum contrast, as achieved using multiple disjoint layers with a fixed display thickness. (Bottom, Left) Reconstructed direct views of a light field for a single plane, coincident with the front layer, with a texture that increases in spatial frequency, from left to right, and contrast, from bottom to top. (Bottom, Right) Absolute error maps indicate multiple layers increase dynamic range, but high spatial frequencies with high contrast are difficult to display with disjoint, multi-layer architectures.

cally achieved by rendering each view with a limited depth of field (i.e., a finite aperture camera). As discussed in Section 4.2, volumetric and multi-layer displays also exhibit a limited depth of field, leading to a circle of confusion linearly proportional to distance from the display. As a result, a combined anti-aliasing and display prefilter is approximated by a single selection of the camera aperture in POV-Ray, with the number of angular views determined by the user, based on the desired depth of field.

Following standard practice in tomographic reconstruction [Herman 1995], we solve Equation 11 using a series expansion into a set of normalized linear basis functions, rather than with a discrete voxel representation. This yields smoother reconstructions and mitigates artifacts occurring due to discrete sampling of rays within the projection matrix. We note, however, that linear reconstruction filters applied to the logarithm of the light field do not correspond to the same filters applied to the light field directly.

We use a sparse, constrained, large-scale trust region method [Coleman and Li 1996] to solve Equation 11. Limited layer contrast is incorporated as a constraint. For each light field, with $384 \times 512 \times 7 \times 7$ samples, the three color channels are evaluated independently. While neglecting color crosstalk between layers, such crosstalk could be incorporated into the optimization at the expense of greater memory requirements and longer computation time. A set of N_l masks, each with 384×512 pixels, are considered. For the light fields studied in this work, the solver typically converges in 8-14 iterations. On average, optimization takes 12 minutes, including projection matrix computation, for five layers using a 2.4 GHz 64-bit Intel Core 2 workstation with 8 GB of RAM, although significant acceleration could be achieved with a GPU-based solver.

6.2 Hardware

Our display prototypes (see Figure 2) consist of five layers separated by clear acrylic sheets, each 0.3175 cm thick. Layers are

printed on transparencies at 300 dots per inch using an Epson Stylus 2200 inkjet printer with six color primaries. In practice, inter-reflections between layers do not lead to visual artifacts. Similarly, reflections on the display surface are minimized by dimming ambient lighting (an anti-reflection coating could be applied if needed). Moiré is not observed with our layered fabrication, likely due to ink naturally blending neighboring pixels. We ignore scattering and diffraction due to the ink pigments, although both likely place an upper limit on the achievable spatio-angular resolution. As shown in Figures 8 and 10, the resulting prototype demonstrates accurate 3D and 2D HDR display using tomographic image synthesis.

7 Discussion

7.1 Benefits and Limitations

Summarized earlier in Section 1.2 and Table 1, capabilities of multi-layer displays are comparable to volumetric and automultiscopic displays, particularly parallax barriers and integral imaging. While volumetric displays faithfully reproduce perceptual depth cues, most cannot represent objects beyond the display; many require moving assemblies and often cannot depict opaque scenes. Yet, the inexpensive alternatives of parallax barriers and integral imaging cannot reproduce high-resolution images, even in the display plane (see Figure 13). Our multi-layer displays present a unique set of capabilities within this space. Foremost, depth of field can exceed that of conventional automultiscopic displays, allowing representation of objects floating beyond the enclosure. Within or near the display, full-resolution depiction is routinely achieved (see Figure 8). Multi-layer displays approach the brightness of integral imaging, with additional layers further enhancing contrast. Finally, similar to volumetric displays, accommodation is preserved for objects within the display, with sufficiently-dense layering.

These benefits come with increased mechanical and computational complexity, as well as the introduction of reconstruction artifacts. For static displays, this complexity has proven manageable, involving additional alignment and separation of printed transparencies. As found in Section 4.3, performance is enhanced by simultaneously increasing the number of layers and the thickness of the display. Although the necessary “thick printing” processes for static displays may prove feasible for larger scales (e.g., by adapting general 3D printing to our multi-planar application), embodiments are currently limited by the capabilities of our iterative reconstruction. With our implementation, which stores the sparse projection matrix directly, the display dimensions, image resolution, and number of layers lie at the upper extent afforded by system memory. However, matrix-free optimization is routinely applied to resolve similar issues in computed tomography [Herman 1995]; this approach, however, may further increase computation time. To reduce computation time to interactive rates, several possibilities are available, including precomputed inverses of the tomography matrix, with subsequent enforcement of the physical constraints, as well as GPU implementations of our tomographic decomposition.

While we believe future dynamic displays using LCD stacks or similar technologies will be feasible, the current multi-layer hardware appears best-suited for static 3D signage, being moderately more expensive and complex than parallax barriers, yet providing significantly enhanced resolution and brightness. Two market applications are seen within this scope, differing by the relation of inter-laced views. First, for 3D display, multi-view imagery is correlated, for which Figure 8 demonstrates our method is well-suited. Second, primarily for advertising, multi-view imagery can be uncorrelated. Known colloquially as *flip animations*, different pictures are presented depending on viewpoint. Tomographic image synthesis does not produce compelling decompositions with such sequences (see



Figure 13: Benefits of multi-layer automultiscopic displays. Simulated views of the “dragon” are shown, from left to right, using integral imaging, parallax barriers, and our multi-layer approach. Display parameters are selected to allow 7×7 views, leading to a similar reduction in resolution for conventional methods. Compared to parallax barriers, our approach is 49 times brighter and exceeds the resolution of both methods by a similar factor.



Figure 14: Flip animations with automultiscopic displays. The “numbers” light field consists of a 3×3 mosaic of views, each depicting a different Arabic numeral. Central views with parallax barriers, our direct method, and the constrained form of our method are shown from left to right. Note that we assume a multi-layer display with three layers; each mask is assumed to have a contrast ratio of 4:1. While a direct application of Equation 11 leads to significant crosstalk, constraining the solver with similar downsampling yields an accurate flip animation, with enhanced contrast.

(Figure 14). In contrast, parallax barriers and integral imaging depict such sequences, although at reduced resolution. Yet, the addition of an equivalent spatial downsampling constraint to our solver yields similar results with enhanced contrast (Figure 14, right). As with other automultiscopic displays, limited depth of field is best exploited by transforming scene content to predominantly occupy the region near the display. Finally, a wide field of view is desirable. Our solver considers a finite field of view, outside of which artifacts may occur, although graceful degradation to a 2D image is observed. In contrast, conventional automultiscopic displays create periodic repetition of the central viewing zone, albeit with the cyclic appearance of pseudoscopic imagery.

7.2 Future Work

Any commercial implementation must address prototype limitations. Manual layer alignment is slow and error-prone; specialized multi-planar printing may be developed that, while similar to conventional rapid prototyping, can deliver higher throughput at lower cost. For example, existing 3D printers support 2D printing on fabricated, albeit opaque, 3D surfaces [Z Corporation 2010]. Similarly, enhanced optimization methods are needed to allow larger, higher-resolution displays with greater numbers of layers. Such efforts will be particularly important to transfer this approach to dynamic multi-view video supporting real-time interaction.

Our multi-layer generalization of parallax barriers opens the door to similar modifications of existing display technologies. As documented in Figure 9, addition of layers alone cannot eliminate artifacts. Similar to Lanman et al. [2010], it may be possible to exploit temporal modulation to obtain more accurate light field

reconstructions. While the upper bound on depth of field indicates a potentially significant gain in spatio-angular resolution, factorization methods must first be developed for such dynamically-modulated stacked displays. In contrast, combinations of additive and multiplicative layers may yield similar gains, while also enhancing brightness. Such displays are efficiently modeled with the emission-absorption volume rendering equation [Sabella 1988]. Finally, our formulation facilitates the development of non-planar, volumetric displays with arbitrary curved surfaces.

8 Conclusion

With the resurgence of 3D cinema, development is focused on the missing cues of disparity, convergence, accommodation, and motion parallax. Given stereoscopic displays now achieve compelling results with cumbersome glasses, efforts fixate on high-fidelity, glasses-free alternatives. Yet, within the cyclic resurgence of 3D, technologies for low-cost 3D signage receive limited attention. We propose an inexpensive multi-layer generalization of parallax barriers, demonstrating brighter, higher-resolution images are possible by augmenting this century-old technology. While the prototype achieves low-cost, high-contrast 3D display of static scenes, tomographic image synthesis can be applied to multi-view video with layered compositions of dynamic spatial light modulators. In this manner we hope to inspire others to consider the unrealized potential of dynamic, yet compact, volumetric displays with a full complement of elements for light emission, absorption, and scattering.

Acknowledgements

We thank the reviewers for their insightful feedback and recognize the support of members of the Imager and Camera Culture groups at the University of British Columbia and the MIT Media Lab, respectively. We also thank Dolby and Samsung Electronics for their sponsorships and David Roodnick for helping to build the prototypes. Gordon Wetzstein was supported by a UBC Four Year Fellowship. Wolfgang Heidrich was supported under the Dolby Research Chair in Computer Science at UBC. Ramesh Raskar was supported by an Alfred P. Sloan Research Fellowship.

References

- AGOCS, ET AL. 2006. A large scale interactive holographic display. In *IEEE Virtual Reality*, 311–312.
- 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.
- BARNUM, P. C., NARASIMHAN, S. G., AND KANADE, T. 2010. A multi-layered display with water drops. *ACM Trans. Graph.* 29, 76:1–76:7.
- BELL, G. P., CRAIG, R., PAXTON, R., WONG, G., AND GALBRAITH, D. 2008. Beyond flat panels: Multi-layered displays with real depth. *SID Digest* 39, 1, 352–355.
- BELL, G. P., ENGEL, G. D., SEARLE, M. J., AND EVANICKY, D., 2010. Method to control point spread function of an image. U.S. Patent 7,742,239.
- BLANCHE, P.-A., ET AL. 2010. Holographic 3-d telepresence using large-area photorefractive polymer. *Nature* 468, 80–83.
- BLUNDELL, B., AND SCHWARTZ, A. 1999. *Volumetric Three-Dimensional Display Systems*. Wiley-IEEE Press.

- BRACEWELL, R. N., AND RIDDLE, A. C. 1967. Inversion of fan-beam scans in radio astronomy. *Astrophysical Journal* 150, 427–434.
- CHAI, J.-X., TONG, X., CHAN, S.-C., AND SHUM, H.-Y. 2000. Plenoptic sampling. In *ACM SIGGRAPH*, 307–318.
- CHAUDHURY, K. N., MUÑOZ-BARRUTIA, A., AND UNSER, M. 2010. Fast space-variant elliptical filtering using box splines. *IEEE Trans. Image* 19, 9, 2290–2306.
- COLEMAN, T., AND LI, Y. 1996. A reflective newton method for minimizing a quadratic function subject to bounds on some of the variables. *SIAM Journal on Optimization* 6, 4, 1040–1058.
- COSSAIRT, O. S., NAPOLI, J., HILL, S. L., DORVAL, R. K., AND FAVALORA, G. E. 2007. Occlusion-capable multiview volumetric three-dimensional display. *Applied Optics* 46, 8, 1244–1250.
- DISNEY, W. E., 1940. Art of animation. U.S. Patent 2,201,689.
- DONG, Y., WANG, J., PELLACINI, F., TONG, X., AND GUO, B. 2010. Fabricating spatially-varying subsurface scattering. *ACM Trans. Graph.* 29, 62:1–62:10.
- DREBIN, R. A., CARPENTER, L., AND HANRAHAN, P. 1988. Volume rendering. *ACM SIGGRAPH* 22, 65–74.
- FAVALORA, G. E. 2005. Volumetric 3D displays and application infrastructure. *IEEE Computer* 38, 37–44.
- GOTODA, H. 2010. A multilayer liquid crystal display for autostereoscopic 3D viewing. In *SPIE-IS&T Stereoscopic Displays and Applications XXI*, vol. 7524, 1–8.
- HAŠAN, M., FUCHS, M., MATUSIK, W., PFISTER, H., AND RUSINKIEWICZ, S. 2010. Physical reproduction of materials with specified subsurface scattering. *ACM Trans. Graph.* 29, 61:1–61:10.
- HECHT, E. 2001. *Optics*. Addison Wesley.
- HERMAN, G. T. 1995. Image reconstruction from projections. *Real-Time Imaging* 1, 1, 3–18.
- IVES, F. E., 1903. Parallax stereogram and process of making same. U.S. Patent 725,567.
- JACOBS, A., ET AL. 2003. 2D/3D switchable displays. *Sharp Technical Journal*, 4, 1–5.
- JONES, A., MCDOWALL, I., YAMADA, H., BOLAS, M., AND DEBEVEC, P. 2007. Rendering for an interactive 360° light field display. *ACM Trans. Graph.* 26, 40:1–40:10.
- KAK, A. C., AND SLANEY, M. 2001. *Principles of Computerized Tomographic Imaging*. Society for Industrial Mathematics.
- KANOLT, C. W., 1918. Photographic method and apparatus. U.S. Patent 1,260,682.
- KLUG, M., HOLZBACH, M., AND FERDMAN, A., 2001. Method and apparatus for recording one-step, full-color, full-parallax, holographic stereograms. U.S. Patent 6,330,088.
- KOOI, F. L., AND TOET, A. 2003. Additive and subtractive transparent depth displays. In *SPIE Enhanced and Synthetic Vision*, vol. 5081, 58–65.
- LANMAN, D., HIRSCH, M., KIM, Y., AND RASKAR, R. 2010. Content-adaptive parallax barriers: optimizing dual-layer 3D displays using low-rank light field factorization. *ACM Trans. Graph.* 29, 163:1–163:10.
- LEVOY, M., AND HANRAHAN, P. 1996. Light Field Rendering. In *ACM SIGGRAPH*, 31–42.
- LIPPMANN, G. 1908. Épreuves réversibles donnant la sensation du relief. *Journal of Physics* 7, 4, 821–825.
- LIPTON, L. 1982. *Foundations of the Stereoscopic Cinema*. Van Nostrand Reinhold.
- LOUKIANITSA, A., AND PUTILIN, A. N. 2002. Stereodisplay with neural network image processing. In *SPIE-IT&T Stereoscopic Displays and Virtual Reality Systems IX*, vol. 4660, 207–211.
- MAEDA, H., HIROSE, K., YAMASHITA, J., HIROTA, K., AND HIROSE, M. 2003. All-around display for video avatar in real world. In *IEEE/ACM ISMAR*, 288–289.
- MATUSIK, W., AJDIN, B., GU, J., LAWRENCE, J., LENSCH, H. P. A., PELLACINI, F., AND RUSINKIEWICZ, S. 2009. Printing spatially-varying reflectance. *ACM Trans. Graph.* 28, 128:1–128:9.
- MITRA, N. J., AND PAULY, M. 2009. Shadow art. *ACM Trans. Graph.* 28, 156:1–156:7.
- NAYAR, S., AND ANAND, V. 2007. 3D display using passive optical scatterers. *IEEE Computer Magazine* 40, 7, 54–63.
- PERLIN, K., AND HAN, J. Y., 2006. Volumetric display with dust as the participating medium. U.S. Patent 6,997,558.
- PUTILIN, A. N., AND LOUKIANITSA, A., 2006. Visualization of three dimensional images and multi aspect imaging. U.S. Patent 6,985,290.
- REINHARD, E., WARD, G., DEBEVEC, P., PATTANAİK, S., HEIDRICH, W., AND MYŠKOWSKI, K. 2010. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*. Morgan Kaufmann.
- SABELLA, P. 1988. A rendering algorithm for visualizing 3D scalar fields. *ACM SIGGRAPH* 22, 51–58.
- SAGI, O., 2009. PolyJet Matrix Technology: A new direction in 3D printing. <http://www.objet.com>, June.
- SEETZEN, H., HEIDRICH, W., STUERZLINGER, W., WARD, G., WHITEHEAD, L., TRENTACOSTE, M., GHOSH, A., AND VOROZCOVS, A. 2004. High dynamic range display systems. *ACM Trans. Graph.* 23, 3, 760–768.
- SLINGER, C., CAMERON, C., AND STANLEY, M. 2005. Computer-generated holography as a generic display technology. *Computer* 38, 8, 46–53.
- SULLIVAN, A. 2003. A solid-state multi-planar volumetric display. In *SID Digest*, vol. 32, 207–211.
- SUYAMA, S., OHTSUKA, S., TAKADA, H., UEHIRA, K., AND SAKAI, S. 2004. Apparent 3-D image perceived from luminance-modulated two 2-D images displayed at different depths. *Vision Research* 44, 8, 785–793.
- YENDO, T., KAWAKAMI, N., AND TACHI, S. 2005. Seelinder: the cylindrical lightfield display. In *ACM SIGGRAPH Emerging Technologies*.
- Z CORPORATION, 2010. ZPrinter 650. <http://www.zcorp.com>, January.
- ZWICKER, M., MATUSIK, W., DURAND, F., AND PFISTER, H. 2006. Antialiasing for automultiscopic 3D displays. In *Eurographics Symposium on Rendering*.