

MIT Open Access Articles

Content-adaptive parallax barriers: optimizing dual-layer 3D displays using low-rank light field factorization

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Douglas Lanman, Matthew Hirsch, Yunhee Kim, and Ramesh Raskar. 2010. Content-adaptive parallax barriers: optimizing dual-layer 3D displays using low-rank light field factorization. In ACM SIGGRAPH Asia 2010 papers (SIGGRAPH ASIA '10). ACM, New York, NY, USA, , Article 163 , 10 pages. DOI=10.1145/1866158.1866164 <http://doi.acm.org/10.1145/1866158.1866164>

As Published: <http://dx.doi.org/10.1145/1866158.1866164>

Publisher: Association for Computing Machinery

Persistent URL: <http://hdl.handle.net/1721.1/67076>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike 3.0



Content-Adaptive Parallax Barriers: Optimizing Dual-Layer 3D Displays using Low-Rank Light Field Factorization

Douglas Lanman Matthew Hirsch Yunhee Kim Ramesh Raskar
MIT Media Lab - Camera Culture Group

<http://cameraculture.media.mit.edu/contentadaptive>




Figure 1: 3D display with content-adaptive parallax barriers. We show that light field display using dual-stacked LCDs can be cast as a matrix approximation problem, leading to a new set of content-adaptive parallax barriers. (Left, Top) A 4D light field, represented as a 2D array of oblique projections. (Left, Bottom) A dual-stacked LCD displays the light field using content-adaptive parallax barriers, confirming both vertical and horizontal parallax. (Middle and Right) A pair of content-adaptive parallax barriers, drawn from a rank-9 decomposition of the reshaped 4D light field matrix. Compared to conventional parallax barriers, with heuristically-determined arrays of slits or pinholes, content adaptation allows increased display brightness and refresh rate while preserving the fidelity of projected images.

Abstract

We optimize automultiscopic displays built by stacking a pair of modified LCD panels. To date, such dual-stacked LCDs have used heuristic parallax barriers for view-dependent imagery: the front LCD shows a fixed array of slits or pinholes, independent of the multi-view content. While prior works adapt the spacing between slits or pinholes, depending on viewer position, we show both layers can also be adapted to the multi-view content, increasing brightness and refresh rate. Unlike conventional barriers, both masks are allowed to exhibit non-binary opacities. It is shown that any 4D light field emitted by a dual-stacked LCD is the tensor product of two 2D masks. Thus, any pair of 1D masks only achieves a rank-1 approximation of a 2D light field. Temporal multiplexing of masks is shown to achieve higher-rank approximations. Non-negative matrix factorization (NMF) minimizes the weighted Euclidean distance between a target light field and that emitted by the display. Simulations and experiments characterize the resulting *content-adaptive parallax barriers* for low-rank light field approximation.

Keywords: Automultiscopic 3D Displays, Parallax Barriers, Light Fields, Weighted Non-negative Matrix and Tensor Factorization

1 Introduction

Thin displays that present the viewer with the illusion of depth have become a driving force in the consumer electronics and entertainment industries, offering a differentiating feature in a market where the utility of increasing 2D resolution has brought diminishing returns. In such systems, binocular depth cues are achieved by presenting different images to each of the viewer's eyes. An attractive class of 3D displays, known as automultiscopic or multi-view autostereoscopic displays, present view-dependent imagery without requiring special eyewear. These displays provide binocular and motion parallax cues. Conventional designs use either optically-attenuating parallax barriers or refracting lens arrays [Konrad and Halle 2007]. In both cases an underlying 2D display (e.g., an LCD) is covered with a second optical element. For barrier designs, a second LCD is placed slightly in front of the first, whereas lens arrays are directly affixed to the underlying display. In most commercial systems either parallax barriers or lenticular sheets are used for horizontal-only parallax. However, pinhole arrays or integral lens sheets can achieve simultaneous vertical and horizontal parallax. While automultiscopic displays realized with lenses are brighter than barrier designs, lens-based designs reduce spatial resolution.

Recent work has focused on increasing the resolution and optical transmission of automultiscopic displays. All 3D displays share three components: display elements, viewers, and multi-view content. Prior works consider the benefits of prefiltering multi-view content for a particular display device [Zwicker et al. 2007]. Others consider adapting display elements (e.g., the spacing within parallax barriers) depending on viewer position [Perlin et al. 2000; Peterka et al. 2008]. We propose *content-adaptive parallax barriers* for which display elements are optimized for the multi-view content. The resulting generalized parallax barriers significantly differ from existing heuristics (e.g., grids of slits or pinholes) and allow




Figure 2: *Conventional versus content-adaptive parallax barriers. (Left) In a conventional parallax barrier display the front panel contains a uniform grid of slits or pinholes. The viewer sees each pixel on the rear panel through this grid, selecting a subset of visible pixels depending on viewer location. A uniform backlight, located behind the rear layer, enables the rear layer to act as a conventional 2D display. (Right) Rather than heuristic barriers, we consider dual-stacked LCDs as general spatial light modulators that act in concert to recreate a target light field by attenuating rays emitted by the backlight. Unlike conventional barriers, both masks can exhibit non-binary opacities.*

increased optical transmission and refresh rate, while preserving the fidelity of displayed images. Furthermore, displays employing these barriers can similarly benefit from prefiltered multi-view content and can adapt to the viewer position, allowing simultaneous optimization of all three significant system components.

1.1 Contributions

We present techniques to optimize the performance of automultiscopic dual-stacked LCDs. Specific contributions are as follows.

- We show that any 4D light field emitted by a fixed pair of masks can be modeled by the tensor product of two 2D masks. Thus, a 1D mask pair only achieves a rank-1 approximation of a 2D light field, modeled by the outer product of the masks.
- We demonstrate higher-rank approximations using temporal multiplexing of mask pairs. A prototype dual-stacked LCD is used to analyze the performance, although off-the-shelf LCDs are not designed for this purpose.
- We show that time-multiplexed light field display using dual-stacked LCDs can be cast as a matrix approximation problem. We propose a new class of content-adaptive parallax barriers, comprising a dynamic set of 2D mask pairs optimized for a given 4D light field. These masks allow compression of the light field into a reduced set of mask pairs, increasing the effective display refresh rate and reducing perceptible flicker.
- To obtain physically-realizable masks, optimized for a given viewer position, we apply non-negative matrix factorization (NMF) to minimize the weighted Euclidean distance between a target light field and that projected by a dual-stacked LCD.

1.2 Benefits and Limitations

As with any automultiscopic display, users do not require special eyewear. Full-resolution 2D display can be restored by setting one LCD to be fully-transparent; while recent lenticular designs also revert to full-resolution 2D modes [Woodgate and Harrold 2003], conventional refractive architectures must trade spatial for angular resolution. Content adaptation increases spatial resolution through temporal multiplexing of mask pairs, exploiting the trend of increasing LCD refresh rates. Unlike time-shifted parallax barriers [Kim et al. 2007], which similarly enhance spatial resolution, the proposed masks adapt to the target light field to further increase display brightness. Arbitrary viewing regions are supported, outside of which the reconstruction is unconstrained, similar to existing methods in which parallax barriers are optimized for one or more tracked viewers [Perlin et al. 2000; Peterka et al. 2008].

The system shares the limitations of other dual-stacked LCDs, including decreased brightness due to transmission through a second liquid crystal layer, moiré, color-channel crosstalk, and discrete viewing zones [Dodgson 2009]. We rely on the human eye to integrate multiple images. Thus, a fast display is required for high-rank approximations. The approach gracefully degrades to lower-rank on displays with reduced refresh rates. Conventional parallax barriers exhibit a beneficial property not shared by our design: periodic view replication; while physically-incorrect, repeated views extend the viewing zone beyond the central field of view. There exist many alternatives for the per-frame optimization we employ. We propose a specific NMF-based scheme, with a weighted multiplicative update. However, our implementation is not fully optimized, requires an initial estimate, and may converge to a sub-optimal solution with decreased PSNR compared to time-shifted parallax barriers.

2 Related Work

2.1 Automultiscopic 3D Display

Automultiscopic displays date back over a century to the work of Lippmann [1908]. More recently, Isono et al. [1993] introduced dual-stacked LCDs to achieve programmable parallax barriers. Commercial implementations [Jacobs et al. 2003] use similar designs for switchable 2D/3D modes (i.e., by disabling or enabling one layer). The resolution of barrier-based displays was characterized by Hoshino et al. [1998]. In a closely-related work, Kim et al. [2007] enhance the spatial resolution of dual-stacked LCDs using time-shifted parallax barriers. Perlin et al. [2000] and Peterka et al. [2008] implement dynamic parallax barriers to optimize the image quality for one or more tracked viewers. In comparison to these works, in which the display is adapted to the viewer, we consider the benefits of simultaneously adapting the display to the content.

2.2 Light Field Capture, Display, and Compression

The two-plane light field parameterization [Levoy and Hanrahan 1996] has been applied to automultiscopic displays by Zwicker et al. [2007] for prefiltering of multi-view content to prevent aliasing. Jones et al. [2007] demonstrate a light field display with a full 360 degree field of view. Light field cameras have mirrored the development of automultiscopic displays, including the use of parallax barriers [Ives 1903] and lenslet arrays [Ng et al. 2005]. In closely-related works [Veeraraghavan et al. 2007; Lanman et al. 2008; Hirsch et al. 2009], heterodyne cameras were designed by placing generalized tiled-broadband patterns close to the sensor; unlike the barriers introduced in this work, these patterns do not adapt to




Figure 3: A thin, dual-layer display (e.g., a dual-stacked LCD) allows depth perception without special eyewear. Multi-view content is rendered or photographed and represented as a 4D light field. Content-adaptive parallax barriers are obtained by applying non-negative matrix factorization to the input light field, increasing display brightness and refresh rate compared to conventional barriers. These mask pairs are displayed using the dual-layer display, emitting a low-rank approximation of the input light field and enabling depth perception.

the incident light field. Our work touches on light field compression [Levoy and Hanrahan 1996; Zwicker et al. 2007]. Here we evaluate low-rank approximations of a 4D light field tensor defined by a series of image pairs. Similar low-rank approximations have been proposed for reflectance field capture [Garg et al. 2006]. In Fourier optics, Ozaktas et al. [2002] analyze coherent illumination using linear algebra. Levin et al. [2009] also exploit the dimensionality gap of 4D light fields. Our general treatment of dual-layer light field modulators is similar to that of Levoy et al. [2009].

3 Content-Adaptive Parallax Barriers

In this section dual-stacked LCDs are analyzed as general spatial light modulators that act in concert to recreate a light field by attenuating rays emitted by the backlight. It is shown that any fixed pair of masks only creates a rank-1 approximation of a light field. Higher-rank approximations are achieved with time multiplexing. We optimize 3D display with dual-stacked LCDs using a matrix approximation framework. This leads to content-adaptive parallax barriers allowing brighter displays with increased refresh rates.

3.1 Light Field Analysis of Parallax Barriers

A general parallax barrier display, containing two mask layers and a backlight, can be analyzed as a light field display device. The following analysis adopts an *absolute two-plane parameterization* of the 4D light field. As shown in Figure 2, an emitted ray is parameterized by the coordinates of its intersection with each mask layer. Thus, the ray (u, v, s, t) intersects the rear mask at the point (u, v) and the front mask at the point (s, t) , with both mask coordinate systems having an origin in the top-left corner.

In a practical automultiscopic display one is primarily concerned with the projection of optical rays within a narrow cone perpendicular to the display surface (see Figure 2), since most viewers will be located directly in front of the device. The distinct images viewable within this region are referred to as the “central views” projected by the display. As a result, a *relative two-plane parameterization* proves more convenient to define a target light field; in this parameterization, an emitted ray is defined by the coordinates (u, v, a, b) , where (u, v) remains the point of intersection with the rear plane and (a, b) denotes the relative offset of the second point of intersection such that $(a, b) = (s - u, t - v)$. As shown in Figure 1, a 2D slice of the 4D light field, for a fixed value of (a, b) , corresponds to a skewed orthographic view (formally an oblique projection).

A general pair of 2D optical attenuation functions, $f(u, v)$ and $g(s, t)$, is defined with the absolute parameterization. These functions correspond to the rear and front masks, respectively. The emitted 4D light field $L(u, v, s, t)$ is given by the product

$$L(u, v, s, t) = f(u, v)g(s, t), \quad (1)$$

assuming illumination by a uniform backlight. In practice, the masks and the emitted light field are discrete functions. The discrete pixel indices are denoted as (i, j, k, l) , corresponding to the continuous coordinates (u, v, s, t) , such that the discretized light field is $\mathbf{L}[i, j, k, l]$ and the sampled masks are $\mathbf{f}[i, j]$ and $\mathbf{g}[k, l]$. When considering only a 2D slice of the 4D light field, the resulting 2D light field matrix $\mathbf{L}[i, k]$ is given by the outer product

$$\mathbf{L}[i, k] = \mathbf{f}[i] \otimes \mathbf{g}[k] = \mathbf{f}[i]\mathbf{g}^T[k], \quad (2)$$

with the masks represented as column vectors $\mathbf{f}[i]$ and $\mathbf{g}[k]$. Note that Equation 1 can be compactly expressed as an outer product only by adopting an absolute two-plane parameterization. For 4D light fields, Equation 2 can be generalized so the light field is given by the tensor product of the 2D masks as follows.

$$\mathbf{L}[i, j, k, l] = \mathbf{f}[i, j] \otimes \mathbf{g}[k, l] \quad (3)$$

3.2 Rank Constraints for Parallax Barriers

From Equation 2 it is clear that a fixed pair of 1D masks can only produce a rank-1 approximation of any given 2D light field matrix. Similarly, a fixed pair of 2D masks also produces a rank-1 approximation of the discrete 4D light field tensor via Equation 3. To our knowledge, this restriction has not been previously described for parallax barrier displays and provides an important insight into their inherent limitations. Figure 4 and Section 5 evaluate the rank




Figure 4: The rank of the bunny light field [Stanford Computer Graphics Laboratory 2008] is assessed. (Left) The central image captured by translating a camera within a 17×17 grid. (Middle) A 2D slice, along the dashed red line, of the 4D light field. (Right) The rank is assessed by the singular value decomposition of the 2D slice. The reconstruction error, measured in terms of the peak signal-to-noise ratio (PSNR), is plotted as a function of the number of singular values included in a given low-rank approximation. In this example, the numerical matrix rank is equal to 17 (i.e., the number of views contained in the light field slice). However, reconstruction with at least three singular values leads to a PSNR greater than 30 dB (generally accepted for lossy compression).

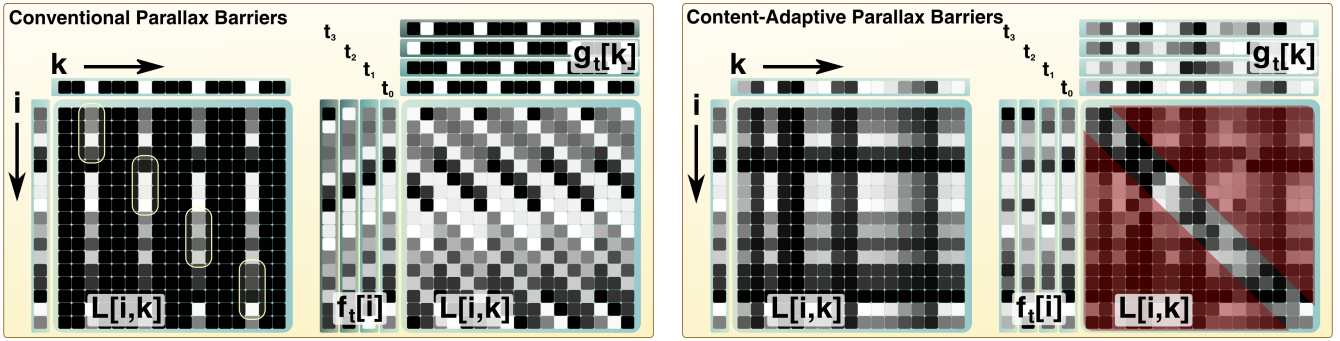


Figure 5: Rank constraints for parallax barriers. (Left) Conventional parallax barriers, following Equation 2, approximate the light field matrix (center) as the outer product of mask vectors (above and to the left). The resulting rank-1 approximation accurately reproduces the circled elements (corresponding to the central views in Figure 2). Note that most columns are not reconstructed, reducing display resolution and brightness. Periodic replicas of the central views are created outside the circled regions. (Middle Left) Time-shifted parallax barriers achieve higher-rank reconstructions by integrating a series of rank-1 approximations, each created by a single translated mask pair. (Middle Right) Content-adaptive parallax barriers increase display brightness by allowing both masks to exhibit non-binary opacities. Here a rank-1 approximation is demonstrated using a single mask pair. (Right) Rank- T approximations are achieved using temporal multiplexing of T content-adaptive parallax barriers via Equations 6 and 9. In practice, the light field will be full rank without enforcing periodic replication (as created by conventional parallax barriers). As a result, we do not constrain rays (shown in red) outside the central view in Equation 9.

of several synthetic and captured light fields; except for the special case when all objects appear in the plane of the display, the rank is typically greater than one. Thus, dual-stacked LCDs employing fixed mask pairs produce rank-deficient approximations; however, perceptually-acceptable approximations can be obtained using conventional parallax barriers, at the cost of decreasing the achievable spatial resolution and image brightness.

As shown in Figure 5, a conventional parallax barrier display employs a heuristic front mask given by

$$\mathbf{g}_{pb}[k, l] = \begin{cases} 1 & \text{if } k \bmod N_h = 0 \text{ and } l \bmod N_v = 0, \\ 0 & \text{otherwise,} \end{cases} \quad (4)$$

where N_h and N_v are the number of skewed orthographic views along the horizontal and vertical display axes, respectively. Thus, the front mask is either a uniform grid of slits or pinholes. Under this definition, the rear mask $\mathbf{f}[i, j]$ is defined such that Equation 3 is satisfied for every ray passing through a non-zero outer mask pixel; thus, the rear mask is given by

$$\mathbf{f}_{pb}[i, j] = \mathbf{L}[i, j, N_h \lfloor i/N_h \rfloor, N_v \lfloor j/N_v \rfloor], \quad (5)$$

when the resolutions of the front and rear masks are equal. Note that, for regions outside the central field of view, periodic replicas of the skewed orthographic views will be projected. These replicas result from viewing neighboring regions of the rear mask through the parallax barrier [Hoshino et al. 1998]. While not correctly capturing the true parallax resulting from steep viewing angles, periodic replication remains a beneficial property of conventional parallax barriers, allowing viewers to see perceptually-acceptable imagery outside the central viewing zone.

In theory, conventional parallax barrier displays achieve perfect reconstruction for any light field ray passing through a non-zero front mask pixel (within the central viewing region). However, as shown in Figure 5, no rays are projected for dark pixels on the front plane. The reconstructed light field will have significant reconstruction errors, when measured using the Euclidean distance between corresponding elements of the target light field. In practice, however, a viewer is separated by a distance that is significantly larger than the slit or pinhole spacing. Spatial low-pass filtering, as performed by the human eye, minimizes perceptual artifacts introduced by parallax barriers (i.e., blending the region between neighboring parallax

barrier gaps). As a result, the occluded regions between slits or pinholes are not perceptually significant; however, these occluded regions significantly reduce the display brightness.

3.3 Time Multiplexing for Higher-Rank Approximation

Despite their practical utility, parallax barriers remain undesirable due to severe attenuation through a slit or pinhole array, as well as reduced spatial resolution of the output light field. Recently, time-shifted parallax barriers have been proposed to eliminate spatial resolution loss [Kim et al. 2007]. In such schemes, a stacked pair of high-speed LCDs is used to sequentially display a series of translated barriers $\mathbf{g}_{pb}[k, l]$ and corresponding underlying masks $\mathbf{f}_{pb}[i, j]$. If the complete mask set is displayed at a rate above the flicker fusion threshold, no image degradation will be perceived.

We generalize the concept of temporal multiplexing for parallax barriers by considering all possible mask pairs rather than the restricted class defined by Equations 4 and 5. Any sequence of T 1D mask pairs creates (at most) a rank- T decomposition of a 2D light field matrix such that

$$\mathbf{L}[i, k] = \sum_{t=1}^T \mathbf{f}_t[i] \otimes \mathbf{g}_t[k] = \sum_{t=1}^T \mathbf{f}_t[i] \mathbf{g}_t^T[k], \quad (6)$$

where $\mathbf{f}_t[i]$ and $\mathbf{g}_t[k]$ denote the rear and front masks for frame t , respectively. Time-multiplexed light field display using dual-stacked LCDs can be cast as a matrix (or more generally a tensor) approximation problem. Specifically, the light field matrix must be decomposed as the matrix product

$$\mathbf{L} = \mathbf{F}\mathbf{G}, \quad (7)$$

where \mathbf{F} and \mathbf{G} are $N_i \times T$ and $T \times N_k$ matrices, respectively. Column t of \mathbf{F} and row t of \mathbf{G} are the masks displayed on the rear and front LCD panels during frame t , respectively. Further observe that a similar expression as Equation 6 can be used to approximate 4D light fields as the summation of multiple tensor products of 2D mask pairs as follows.

$$\mathbf{L}[i, j, k, l] = \sum_{t=1}^T \mathbf{f}_t[i, j] \otimes \mathbf{g}_t[k, l] \quad (8)$$

3.4 Finding Non-negative Mask Sequences

Each mask pair $\{\mathbf{f}_t[i, j], \mathbf{g}_t[k, l]\}$ must be non-negative, since it is illuminated by an incoherent light source (i.e., the rear LCD backlight). We seek a *content-adaptive* light field factorization $\tilde{\mathbf{L}} = \mathbf{F}\mathbf{G}$ that minimizes the weighted Euclidean distance to the target light field \mathbf{L} , under the necessary non-negativity constraints, such that

$$\arg \min_{\mathbf{F}, \mathbf{G}} \frac{1}{2} \|\mathbf{L} - \mathbf{F}\mathbf{G}\|_{\mathbf{W}}^2, \text{ for } \mathbf{F}, \mathbf{G} \geq 0, \quad (9)$$

where the reconstruction error is given by

$$\frac{1}{2} \|\mathbf{L} - \mathbf{F}\mathbf{G}\|_{\mathbf{W}}^2 = \sum_{ijkl} [\mathbf{W} \circ (\mathbf{L} - \mathbf{F}\mathbf{G}) \circ (\mathbf{L} - \mathbf{F}\mathbf{G})]_{ijkl}. \quad (10)$$

Here \circ denotes the Hadamard product for element-wise multiplication of matrices. Unlike conventional barriers, the field of view can be adapted to one or more viewers by specifying elements of the weight matrix \mathbf{W} (i.e., the Euclidean norm will be minimized where \mathbf{W} is large). The weight matrix plays a crucial role, ensuring a low-rank approximation can obtain high reconstruction accuracy by artificially reducing the rank of the target light field. General 4D light fields are handled by reordering as 2D matrices, with 2D masks reordered as vectors, allowing a similar matrix approximation scheme to be applied.

Equation 9 can be solved using non-negative matrix factorization. Prior numerical methods include the multiplicative update rule [Lee and Seung 1999]. We use the weighted update introduced by Blondel et al. [2008]. Initial estimates $\{\mathbf{F}, \mathbf{G}\}$ are refined as follows.

$$\mathbf{F} \leftarrow \mathbf{F} \circ \frac{[(\mathbf{W} \circ \mathbf{L})\mathbf{G}^T]}{[(\mathbf{W} \circ (\mathbf{F}\mathbf{G}))\mathbf{G}^T]} \quad \mathbf{G} \leftarrow \mathbf{G} \circ \frac{[\mathbf{F}^T(\mathbf{W} \circ \mathbf{L})]}{[\mathbf{F}^T(\mathbf{W} \circ (\mathbf{F}\mathbf{G}))]} \quad (11)$$

Typical mask pairs produced by the optimization procedure are shown in Figures 1, 6, and 14. Note that, if Equation 9 was not constrained to weighted, non-negative factorizations, singular value decomposition (SVD) could be applied; however, Srebro and Jaakkola [2003] have shown that solving a weighted SVD problem also requires an iterative algorithm with multiple local minima. In our implementation of Equation 11, the masks are initialized with random values uniformly distributed on $[0, 1]$; alternative strategies, including seeding with conventional parallax barriers, did not yield reconstructions with reduced errors or increased transmission. After each iteration the mask elements are truncated to the range $[0, 1]$. In conclusion, we propose the resulting non-negative, content-adaptive parallax barriers as a generalization of traditional parallax barrier displays, in which images displayed on both LCD layers are jointly optimized, independently for each target automultiscopic video frame.




Figure 6: Intuition behind local parallax barriers. (Left) From left to right and top to bottom: oblique projections of a step edge seen as a viewer moves in similar directions. A rank-9 decomposition produces a set of mask pairs. (Middle) A rear-panel mask. (Right) A front-panel mask. Note that optimization appears to produce a local parallax barrier, rotated to align with the step edge.

4 The Structure of Content-Adaptive Barriers

Content-adaptive parallax barriers exhibit predictable structure. Consider the masks shown in Figures 1 and 14, as well as those in the supplementary materials: flowing, fringe-like patterns are consistently observed. We interpret that content-adaptive parallax barriers are locally-similar to conventional parallax barriers, but rotated to align to nearby edges in the light field. Intuitively, parallax is only perceived as a viewer moves perpendicular to an edge, thus a rotated *local parallax barrier* (i.e., an array of slits) is sufficient to project the correct 4D light field in such local regions. This is similar to the ‘‘aperture problem’’, wherein a windowed, translated grating appears to move perpendicular to the stripe orientation.

Qualitatively, the front-panel masks exhibit flowing, slit-like barriers aligned perpendicular to the *angular gradient* of the 4D light field (see Figure 13), defined using a relative parameterization as

$$\nabla_{ab} L(u, v, a, b) = \left(\frac{\partial L}{\partial a}, \frac{\partial L}{\partial b} \right). \quad (12)$$

The rear-panel masks exhibit rotated spatially-multiplexed images similar to conventional parallax barriers. In Figure 6 we consider a region centered on a depth discontinuity. Locally, the scene is modeled by two fronto-parallel planes (i.e., a step edge). A 4D light field, containing 3×3 oblique projections, is rendered so the disparity between projections is 10 pixels. The front-panel masks contain perturbed lines that run parallel to the edge (i.e., perpendicular to the angular gradient). Their average spacing equals the angular resolution (3 pixels) and they span a region equal to the product of the disparity and the number of views minus one (± 10 pixels from the edge). The masks exhibit random noise away from the edge, approximating a scene without parallax. Following Lee and Seung [1999], Equation 11 converges to a local stationary point, but not necessarily the global minimum; as a result, the observed local parallax barriers possess some randomization due to convergence to a local minima. Additional examples are shown in Figure 14.

Although we can predict mask structure, we cannot provide an analytic solution. This remains a promising future direction. However, the local parallax barrier interpretation gives intuition into the benefits and limitations of the method. Unlike 2D pinhole arrays, adaptation creates 1D slits that transmit more light. Consider $N_h \times N_v$ views of a sphere. With pinholes, each front mask is a grid of $N_h \times N_v$ tiles with one transparent pixel. We create local barriers following the angular gradient (e.g., the sphere boundary). Near discontinuities each $N_h \times N_v$ block of the front mask contains slits with an average of no less than $\min(N_h, N_v)$ transparent pixels. Thus, the average achievable brightness increase is $\min(N_h, N_v)$. We conclude that one significant benefit of content-adaptive parallax barriers is to allow simultaneous horizontal and vertical parallax, while preserving the brightness of conventional parallax barriers (i.e., arrays of slits) that support horizontal-only parallax.

5 Implementation and Analysis

5.1 Implementation

This section discusses the details of constructing a dual-stacked LCD using modified panels, validates its performance using conventional and content-adaptive parallax barriers, and assesses the performance compared to prior automultiscopic displays.

5.1.1 Hardware

As shown in Figure 8, a dual-stacked LCD was constructed using a pair of 1680×1050 Viewsonic FuHzion VX2265wm 120 Hz LCD

panels. The panels have a pixel pitch of $282\ \mu\text{m}$ and are separated by 1.5 cm. However, as described in Section 6, masks are displayed at half the native resolution. Thus, for a typical light field with an angular resolution $N_h \times N_v$ of 5×3 views, the prototype supports an $11^\circ \times 7^\circ$ field of view; a viewer sees correct imagery when moving within a frustum with similar apex angles.

The rear layer is an unmodified panel, whereas the front layer is a spatial light modulator (SLM) fashioned by removing the backlight from a second panel. The front polarizing diffuser and rear polarizing film are removed. The front polarizing diffuser is replaced with a transparent polarizer, restoring the spatial light modulation capability of the panel. Without such modifications, the polarizers in the front panel completely attenuate light polarized by the rear panel. Eliminating the redundant rear polarizer of the front panel increases light transmission. The LCD panels are driven separately via DVI links from a dual-head NVIDIA Quadro FX 570 display adapter, automatically synchronizing the display refreshes.

5.1.2 Software

Light fields are rendered with POV-Ray [Persistence of Vision Pty. Ltd. 2004] and masks are represented by a series of texture pairs. Each color channel is factorized independently. The displays are driven at 120 Hz with a custom OpenGL application. Gamma compression is applied to ensure mask intensity varies linearly with the encoded value; a gamma value of $\gamma = 2.2$ was measured for our LCDs. Mask optimization uses a multi-threaded C++ implementation written with the POSIX Pthreads API; a single-threaded version is provided with the supplementary code. An Intel Xeon 8-core 3.2 GHz processor with 8 GB of RAM is used for optimization and display. For a typical light field with 5×3 views, each with a resolution of 840×525 pixels, the optimization takes approximately 10 seconds per iteration. In practice, at least 50 iterations are required for the PSNR to exceed 30 dB, leading to an average run-time of eight minutes per frame (see Figure 7). As observed by Zwicker et al. [2007], the target light field should be prefiltered to prevent aliasing. Such prefiltering was not applied in our implementation, causing additional artifacts in the right-hand column of Figure 9.




Figure 7: Approximation error as a function of NMF iteration. The average PSNR of the reconstruction is plotted for a rank-9 decomposition of the light fields shown in Figures 1 and 13.

5.2 Analysis

As with any 3D display, a viewer is concerned with resolution (both spatial and angular), brightness, refresh rate, and reconstruction error. Experiments and simulations assess the performance of content-adaptive parallax barriers, as compared to time-shifted parallax barriers [Kim et al. 2007]. Two primary benefits result from content-adaptive parallax barriers: increased display brightness and increased display refresh rate; we analyze each in turn.

5.2.1 Increasing Display Brightness

Following Section 4, content-adaptive parallax barriers appear to exhibit local parallax barrier structure. Using this interpretation we




Figure 8: Prototype automultiscopic display using dual-stacked LCDs. (Left) Side view of the prototype. From right to left: (a) rear LCD with backlight, (b) spacer, (c) front LCD, and (d) replacement polarizing sheet. (Right) Central view of a synthetic scene rendered with content-adaptive parallax barriers. Video results using this prototype are included in the supplementary material.

previously predicted an average brightness increase by a maximum factor of $\min(N_h, N_v)$. The supplementary code was used to render a diverse set of light fields containing varying degrees of disparity, contrast, and geometric complexity. Select light fields are shown in Figures 1 and 13, with additional examples included in the supplementary material and video. As shown in Figure 10, the peak signal-to-noise ratio (PSNR) of the reconstruction is measured as a function of the attempted increase in brightness (i.e., the target light field is multiplied by the desired gain).

Figure 10 demonstrates that content-adaptive parallax barriers can increase display brightness, in comparison to time-shifted parallax barriers. These examples use $T = N_h N_v$ time-multiplexed mask pairs (i.e., identical to the number of masks required with time-shifted parallax barriers). As predicted, when brightness is enhanced by the theoretically-predicted factor of $\min(N_h, N_v)$ (i.e., $3 \times$ brighter in these examples), the PSNR of the reconstruction remains above 30 dB; for greater increases in brightness, artifacts become readily apparent. We observe that the PSNR is finite (i.e., artifacts are present), even when no increase in brightness is attempted. This indicates a limitation of the current optimization procedure. As described in Section 4, Equation 11 is not guaranteed to converge to the global minimum. Thus, artifacts persist even for the step edge in Figure 6. Furthermore, no local parallax barrier (i.e., an array of slits) can represent regions with both horizontal and vertical parallax. For such regions, increasing the brightness by any factor will lead to artifacts under the local parallax barrier interpretation (e.g., the checkerboard corners in Figure 9). However, since a PSNR greater than 30 dB is generally accepted for lossy compression, content adaptation achieves significant increases while presenting images that retain the fidelity of the target light field.

5.2.2 Increasing Display Refresh Rate with Compression

Content adaptation can also increase the effective refresh rate of the automultiscopic display. Consider the prototype system, supporting a native 120 Hz refresh rate. In this case, only five masks can be time-multiplexed before the effective refresh rate drops below 24 Hz and flicker becomes readily apparent. Thus, supporting simultaneous horizontal and vertical parallax becomes challenging.

Fortunately, content-adaptive parallax barriers allow the light field to be compressed using a set of $T < N_h N_v$ mask pairs. Theoretically, rank-1 light fields occur in a single case: when a textured

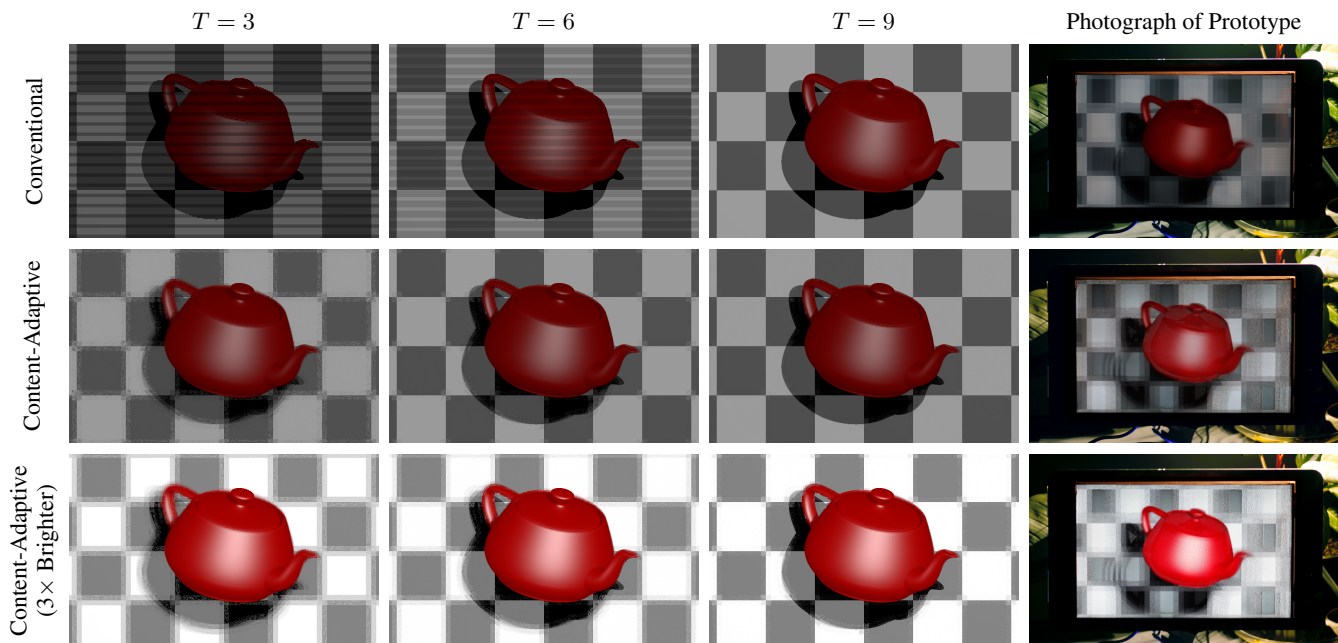


Figure 9: Increasing display brightness and refresh rate. Content-adaptive barriers are compared to time-shifted barriers [Kim et al. 2007], with the exposure normalized so the relative image brightness is consistent with observation. Following Section 5, light field compression is achieved for $T < N_h N_v$ mask pairs. Reconstructions with three, six, and nine time-multiplexed mask pairs are shown in the first three columns from the left, respectively. Experimental photographs (fourth column) are compared to predicted images (third column). All images correspond to the central oblique view for the light field in Figure 1. While content-adaptive barriers produce some high-frequency artifacts, even with nine mask pairs, they can compress the light field with higher PSNR than conventional barriers (see Figure 11). As shown along the bottom row, adaptation also allows the brightness to be increased with minimal degradation in image fidelity (see Figure 10).

plane is displayed in the plane of the rear LCD panel (i.e., for a light field without any parallax). Experimentally, rank grows (to the number of views $N_h N_v$) as the plane is translated away from the rear LCD. For example, consider the 2D slice of a captured light field shown in Figure 4. As described by Chai et al. [2000], the separation of a plane from the rear LCD determines the skew of the 2D light field slice. Thus, distant objects require higher-rank approximations. However, in this example, 17 views were reconstructed with a PSNR greater than 30 dB using three mask pairs. Scenes with limited parallax and depth variation require fewer masks.

Figure 11 illustrates compression trends typical with content-adaptive parallax barriers. As before, artifacts are present even when $T = N_h N_v$ mask pairs are used; however, in this case the PSNR exceeds 45 dB. Examples of the predicted and experimentally-measured artifacts are shown in Figure 9. We conclude that, as with increasing brightness, content adaptation reveals a novel trade-off between automultiscopic display brightness, refresh rate, and reconstruction error. Additional results, including high-resolution stills, masks, and video sequences are included in the supplementary material and video.

6 Benefits and Limitations

All dual-stacked LCDs contend with a similar set of challenges, including moiré and color-channel crosstalk. The proposed method also confronts the further challenges of display flicker and the current limitations of NMF algorithms. We discuss solutions for each of these issues in the remainder of this section.

Moiré: Viewing one LCD through another causes visible fringes (moiré) to appear. Commercial dual-stacked LCDs have eliminated moiré by increasing the blur introduced by the front diffuser on the

rear LCD [Bell et al. 2008]. Ideally, the diffuser should only blur neighboring color subpixels—minimizing moiré while preserving spatial resolution. We use this solution in our prototype; a thin paper vellum sheet is placed against the rear LCD. Experimentally, the diffuser eliminates moiré, however the image resolution is reduced from 1680×1050 to 840×525 ; a custom diffuser, with a properly-selected point spread function, would prevent this reduction.

Color-channel Crosstalk: Each LCD color filter transmits a range of wavelengths. The relative transmission, as a function of wavelength, is known as the color filter transmission spectrum. The transmission spectra exhibit some overlap in commercial panels. Since the panels we use are not optimized for dual-stacked configurations, the overlapping transmission spectra cause visible color-channel crosstalk (see the supplementary material for experimental measurements). This crosstalk is ignored in our optimization; while allowing independent decompositions for each channel, this simplification results in visual artifacts. In a commercial implementation, the transmission spectra could be designed with minimal overlap. However, to minimize crosstalk for grayscale regions, we initialize each color channel with the same random set of values. As shown in Figure 1, the deterministic optimization algorithm leads to grayscale masks that minimize crosstalk in these regions.

Flicker: Humans perceive an intermittent light source as steady when it varies between 16–60 Hz, depending on illumination conditions. For dim stimuli in darkened rooms, 16 Hz is a commonly-accepted lower bound. Our prototype can multiplex up to eight mask pairs at 15 Hz. Multiplexing five mask pairs achieves a 24 Hz refresh, equivalent to cinematic projection. However, 240 Hz LCDs are commercially available and allow doubling the decomposition rank without altering the refresh rate. Thus, our method will benefit from the trend of LCDs with increased refresh rates. However,




Figure 10: Approximation error as a function of gain in brightness. The average PSNR of the reconstruction is plotted for a rank-9 decomposition of the light fields shown in Figures 1 and 13. For time-shifted parallax barriers, transmission can be increased either by enlarging slits/pinholes or by brightening the rear LCD. The latter is considered here, however simulations of the former also confirm time-shifted parallax barriers cannot achieve a PSNR greater than 15 dB when increasing brightness by a factor greater than two.

as observed by Woods and Sehic [2009], such high-speed panels may require further optimization for autostereoscopic applications, rather than their current focus on reducing motion blur.

Non-negative Matrix Factorization: As described by Lee and Seung [1999], multiplicative update rules (e.g., Equation 11) are easy to code, but are not as efficient as other algorithms [Chu et al. 2004]. Our algorithm scales linearly with the number of light field elements and the decomposition rank. The prototype requires at least 50 iterations to converge, resulting in an average run-time of around eight minutes per frame, preventing interactive content. However, as shown in the supplementary video, masks can be precomputed to allow dynamic content. Figure 1, 6, and 14 show our optimization produces high-frequency patterns, even in uniform regions. Regularized NMF algorithms [Zhang et al. 2008] remain a promising direction of future work; preliminary results with smoothed masks are included in the supplementary material and Figure 12.

7 Discussion and Future Work

The analysis of dual-stacked LCDs, as rank-constrained light field displays, points the way along a new direction—one in which the display elements themselves are independently optimized for the target light field. While we show one technique for obtaining content-adaptive parallax barriers that optimize optical transmission and effective refresh rate, it is our hope that future work will reveal a wider range of optimization techniques and classes of adaptive masks. The weight matrix may be used to achieve other effects; weights could be selected to support multiple viewers or a wider field of view. This is a timely development, as the power and availability of computation has made real-time optimization a reality in many fields. In addition, content-adaptive parallax barriers will benefit from the trend of increasing LCD refresh rates. Refined cost functions, such as those that incorporate human perceptual effects, may provide superior results.

Our optimization is reminiscent of that used with computer generated holograms [Slinger et al. 2005], as well as band moiré images [Hersch and Chosson 2004]. It may be possible to obtain analytical interpretations of our results, possibly through a frequency-domain analysis; the structure of the masks we obtain appear to mimic local parallax barriers and suggest a broadband nature, reminiscent of the masks used in heterodyne light field cameras [Veeraraghavan et al. 2007; Lanman et al. 2008; Hirsch et al. 2009].

Any commercial implementation must address the current prototype limitations, including: moiré, color-channel crosstalk, and flicker. Our theory only applies to dual-layer displays, yet extensions allowing more layers may reveal additional benefits, including increased fidelity and reduced requirements on the number of mask




Figure 11: Approximation error as a function of decomposition rank. The average PSNR of the reconstruction is plotted for a rank- T decomposition of the light fields in Figures 1 and 13. For 3×3 views, a theoretical PSNR of infinity is achieved with 9 time-shifted conventional parallax barriers. In comparison, content-adaptive barriers achieve higher PSNR than conventional barriers when fewer frames are used. Experimental and predicted images with varying degrees of compression are shown in Figure 9.

pairs. Generalizing to arbitrary numbers of spatial light modulators, volumetric occluders could be designed to modify a uniform backlight to reproduce a light field; such occluders may function as the display equivalent of the volumetric occluders used for light field capture with reference structure tomography [Brady et al. 2004].

8 Conclusion

Parallax barriers have persisted as a display technology, nearly unchanged, since their introduction over a century ago [Ives 1903]. A first-principles analysis of emerging high-speed, dual-stacked LCDs is timely and necessary to move beyond earlier heuristics. With content-adaptive parallax barriers we hope to inspire others to reinvestigate barrier displays. Using a light field analysis, dual-stacked LCDs are shown to produce low-rank approximations, even with the benefit of time multiplexing. While conventional parallax barriers provide perceptually-acceptable imagery, content-adaptive parallax barriers reveal a new design space. In this paper, light field display is cast as a matrix approximation problem, exploiting the inherent compressibility of light fields, not just for transmission, but also for display applications. Such content-adaptive displays, in which individual optical elements are optimized in real-time, offer new opportunities for automultiscopic design.




Figure 12: Regularized NMF for smooth masks. (Top) The spheres light field (see Figure 8) is decomposed via Equation 11. The masks contain high-frequency patterns, even in uniform regions without parallax. (Bottom) Spatial smoothness is achieved by convolving the masks with a Gaussian filter after every 10 iterations. The filter standard deviation is reduced over time, allowing high-frequencies to appear only in later iterations. Note the close agreement with the local barrier interpretation in Section 4.




Figure 13: Predicting the structure of content-adaptive parallax barriers. (Left) The blocks light field containing a stack of three colored glass blocks. From left to right and top to bottom: oblique projections seen as a viewer moves from left to right and top to bottom. (Right) Streamlines of the angular gradient of the light field, evaluated following Section 4, visualized using line integral convolution [Cabral and Leedom 1993]. Note that streamline direction predicts the orientation of the local parallax barriers appearing in the right-hand side of Figure 14. Consider the windowed region within the blue rectangle (rendered in the same position for all figures on this page). As shown on the left, the light field primarily exhibits horizontal parallax within this window. Thus, the streamlines run vertically on the right; similarly, the corresponding region on the right-hand side of Figure 14 exhibits vertically-oriented slits. As described in Section 4, the resulting local parallax barrier is sufficient to project this windowed region of the light field. Note similar correspondences within the red and green windows.




Figure 14: A content-adaptive parallax barrier mask pair. A rank-9 decomposition of the blocks light field, shown in Figure 13, was evaluated using Equation 11 in Section 3. A single mask pair is shown, with the rear and front masks to the left and right, respectively. To enhance the visibility of the emergent local parallax barriers, only the luminance channel of the light field is processed.

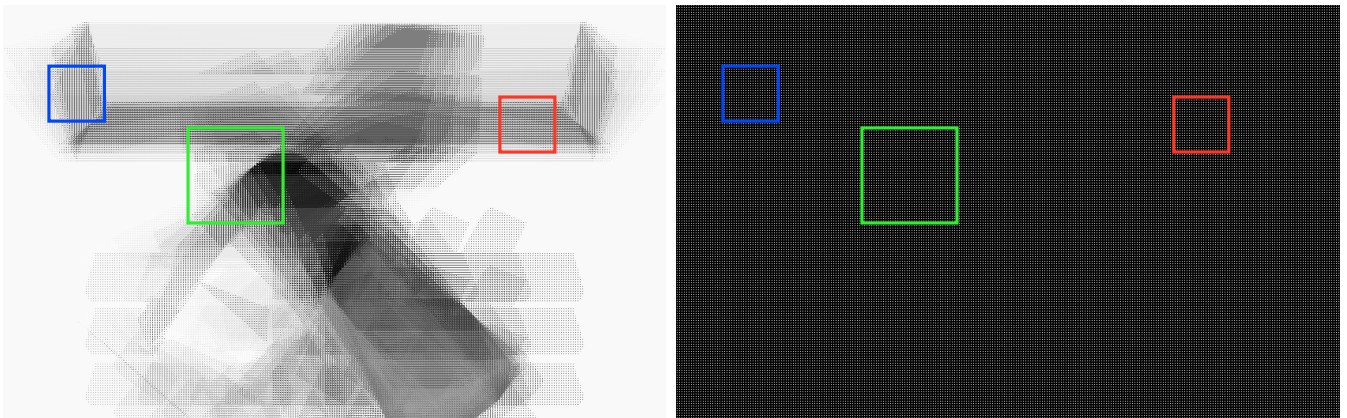


Figure 15: A conventional parallax barrier mask pair. A set of nine time-shifted conventional parallax barriers [Kim et al. 2007] were evaluated using Equations 4 and 5 in Section 3. A single mask pair is shown, with the rear and front masks to the left and right, respectively. For comparison with Figure 14, only the luminance channel of the light field is processed.

Acknowledgements

We thank the reviewers for insightful feedback, the Camera Culture and Information Ecology groups for their support, and Samsung Electronics for its sponsorship. Szymon Jakubczak contributed to an earlier version of this work presented at SIGGRAPH 2010. Thomas Baran and Gabriel Taubin contributed useful discussions. Douglas Lanman is supported by NSF Grant CCF-0729126, Yunhee Kim by NRF of Korea Grant 2009-352-D00232, and Ramesh Raskar by an Alfred P. Sloan Research Fellowship.

References

- 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.
- BLONDEL, V. D., HO, N.-D., AND VAN DOOREN, P. 2008. Weighted nonnegative matrix factorization and face feature extraction. *Image and Vision Computing*.
- BRADY, D. J., PITSIANIS, N. P., AND SUN, X. 2004. Reference structure tomography. *J. Opt. Soc. Am. A* 21, 7, 1140–1147.
- CABRAL, B., AND LEEDOM, L. C. 1993. Imaging vector fields using line integral convolution. In *ACM SIGGRAPH*, 263–270.
- CHAI, J.-X., TONG, X., CHAN, S.-C., AND SHUM, H.-Y. 2000. Plenoptic sampling. In *ACM SIGGRAPH*, 307–318.
- CHU, M., DIELE, F., PLEMMONS, R., AND RAGNI, S. 2004. Optimality, computation, and interpretation of nonnegative matrix factorizations. *SIAM Journal on Matrix Analysis*.
- DODGSON, N. A. 2009. Analysis of the viewing zone of multi-view autostereoscopic displays. In *SPIE Stereoscopic Displays and Applications XIII*, 254–265.
- GARG, G., TALVALA, E.-V., LEVOY, M., AND LENSCH, H. P. A. 2006. Symmetric photography: Exploiting data-sparseness in reflectance fields. In *EGSR*, 251–262.
- HERSCH, R. D., AND CHOSSON, S. 2004. Band moiré images. *ACM Trans. Graph.* 23, 3, 239–247.
- HIRSCH, M., LANMAN, D., HOLTZMAN, H., AND RASKAR, R. 2009. BiDi screen: A thin, depth-sensing LCD for 3D interaction using light fields. *ACM Trans. Graph.* 28, 5.
- HOSHINO, H., OKANO, F., ISONO, H., AND YUYAMA, I. 1998. Analysis of resolution limitation of integral photography. *J. Opt. Soc. Am. A* 15, 8, 2059–2065.
- ISONO, H., YASUDA, M., AND SASAZAWA, H. 1993. Autostereoscopic 3-D display using LCD-generated parallax barrier. *Electronics and Communications in Japan* 76, 7, 77–84.
- IVES, F. E., 1903. Parallax stereogram and process of making same. United States Patent 725,567.
- JACOBS, A., MATHER, J., WINLOW, R., MONTGOMERY, D., JONES, G., WILLIS, M., TILLIN, M., HILL, L., KHAZOVA, M., STEVENSON, H., AND BOURHILL, G. 2003. 2D/3D switchable displays. *Sharp Technical Journal*, 4, 15–18.
- JONES, A., MCDOWALL, I., YAMADA, H., BOLAS, M., AND DEBEVEC, P. 2007. Rendering for an interactive 360° light field display. *ACM Trans. Graph.* 26, 3.
- KIM, Y., KIM, J., KANG, J.-M., JUNG, J.-H., CHOI, H., AND LEE, B. 2007. Point light source integral imaging with improved resolution and viewing angle by the use of electrically movable pinhole array. *Optics Express* 15, 26, 18253–18267.
- KONRAD, J., AND HALLE, M. 2007. 3-D displays and signal processing. *IEEE Signal Processing Magazine* 24, 6, 97–111.
- LANMAN, D., RASKAR, R., AGRAWAL, A., AND TAUBIN, G. 2008. Shield fields: Modeling and capturing 3D occluders. *ACM Trans. Graph.* 27, 5.
- LEE, D. D., AND SEUNG, H. S. 1999. Learning the parts of objects by non-negative matrix factorization. *Nature* 401, 788–791.
- LEVIN, A., HASINOFF, S. W., GREEN, P., DURAND, F., AND FREEMAN, W. T. 2009. 4D frequency analysis of computational cameras for depth of field extension. In *ACM SIGGRAPH*, 1–14.
- LEVOY, M., AND HANRAHAN, P. 1996. Light field rendering. In *ACM SIGGRAPH*, 31–42.
- 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. Epreuves reversibles donnant la sensation du relief. *Journal of Physics* 7, 4, 821–825.
- 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., Stanford University.
- OZAKTAS, H. M., YÜKSEL, S., AND KUTAY, M. A. 2002. Linear algebraic theory of partial coherence: Discrete fields and measures of partial coherence. *J. Opt. Soc. Am. A* 19, 8, 1563–1571.
- PERLIN, K., PAXIA, S., AND KOLLIN, J. S. 2000. An autostereoscopic display. In *ACM SIGGRAPH*, 319–326.
- PERSISTENCE OF VISION PTY. LTD., 2004. Persistence of vision raytracer (version 3.6). <http://www.povray.org>.
- PETERKA, T., KOOIMA, R. L., SANDIN, D. J., JOHNSON, A., LEIGH, J., AND DEFANTI, T. A. 2008. Advances in the dyanalax solid-state dynamic parallax barrier autostereoscopic visualization display system. *IEEE TVCG* 14, 3, 487–499.
- SLINGER, C., CAMERON, C., AND STANLEY, M. 2005. Computer-generated holography as a generic display technology. *IEEE Computer* 38, 8, 46–53.
- SREBRO, N., AND JAAKKOLA, T. 2003. Weighted low-rank approximations. In *ICML*, 720–727.
- STANFORD COMPUTER GRAPHICS LABORATORY, 2008. The Stanford light field archive. <http://lightfield.stanford.edu>.
- VEERARAGHAVAN, A., RASKAR, R., AGRAWAL, A., MOHAN, A., AND TUMBLIN, J. 2007. Dappled photography: Mask enhanced cameras for heterodyned light fields and coded aperture refocusing. *ACM Trans. Graph.* 26, 3, 69.
- WOODGATE, G. J., AND HARROLD, J. 2003. High efficiency reconfigurable 2D/3D autostereoscopic display. In *SID DIGEST*.
- WOODS, A. J., AND SEHIC, A. 2009. The compatibility of LCD TVs with time-sequential stereoscopic 3D visualization. In *SPIE Stereoscopic Displays and Applications XX*.
- ZHANG, T., FANG, B., LIU, W., TANG, Y. Y., HE, G., AND WEN, J. 2008. Total variation norm-based nonnegative matrix factorization for identifying discriminant representation of image patterns. *Neurocomputing* 71, 10–12, 1824–1831.
- ZWICKER, M., VETRO, A., YEA, S., MATUSIK, W., PFISTER, H., AND DURAND, F. 2007. Resampling, antialiasing, and compression in multiview 3-D displays. *IEEE Signal Processing Magazine* 24, 6, 88–96.