

A Reflectance Display

Daniel Glasner*
Harvard University

Todd Zickler
Harvard University

Anat Levin
Weizmann Institute of Science

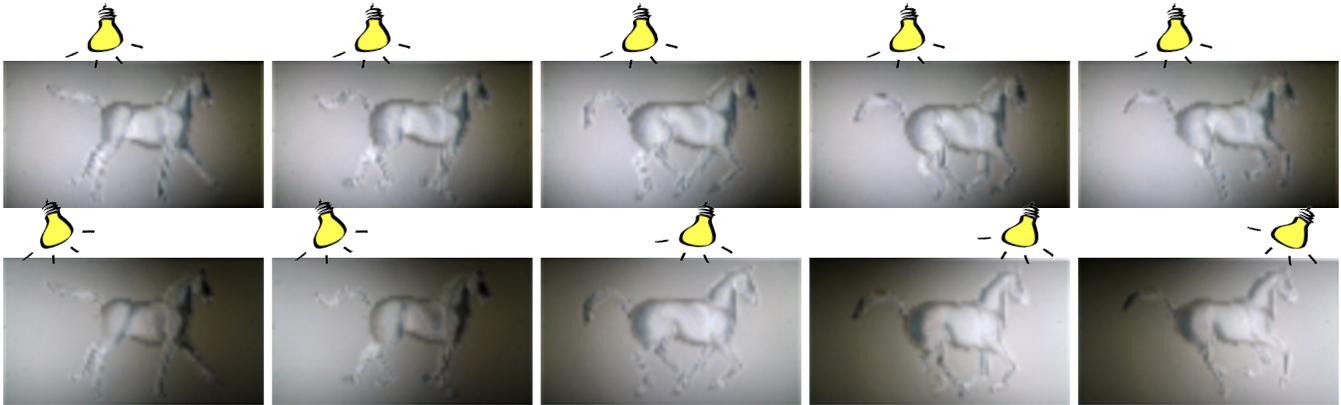


Figure 1: We introduce a reflectance display: a dynamic digital array of dots, each of which can independently display a custom, time-varying reflectance function. The display passively reacts to illumination and viewpoint changes in real-time, without any illumination-recording sensors, head tracking, or on-the-fly rendering. In this example the time-varying reflectance functions create a “reflectance video” that gives the illusion of a dynamic 3D model being physically-shaded by the room’s ambient lighting. The top row shows a time-sequence of photographs of the dynamic display from a stationary viewpoint under fixed ambient lighting, and the bottom row shows how the display reacts to changes in ambient lighting by passively inducing the appropriate 3D shading effects.

Abstract

We present a *reflectance display*: a dynamic digital display capable of showing images and videos with spatially-varying, user-defined reflectance functions. Our display is *passive*: it operates by phase-modulation of reflected light. As such, it does not rely on any illumination recording sensors, nor does it require expensive on-the-fly rendering. It reacts to lighting changes instantaneously and consumes only a minimal amount of energy. Our work builds on the wave optics approach to BRDF fabrication of Levin et al. [2013]. We replace their expensive one-time hardware fabrication with a programmable liquid crystal spatial light modulator, retaining high resolution of approximately 160 dpi. Our approach enables the display of a much wider family of angular reflectances, and it allows the display of dynamic content with time varying reflectance properties—“reflectance videos”. To facilitate these new capabilities we develop novel reflectance design algorithms with improved resolution tradeoffs. We demonstrate the utility of our display with a diverse set of experiments including display of custom reflectance images and videos, interactive reflectance editing, display of 3D content reproducing lighting and depth variation, and simultaneous display of two independent channels on one screen.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Viewing and display algorithms;

Keywords: Digital displays, BRDF design, wave optics.

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

1 Introduction

Display technology has advanced significantly in recent years, producing higher definition, richer color, and even display of 3D content. However, the overwhelming majority of current displays are insensitive to the illumination in the observer’s environment. This imposes a significant barrier to achieving an immersive experience because scenes can only be displayed under the same illumination conditions in which they were captured. If the illumination in the observer’s environment changes during playback, there are no corresponding effects on the shading, highlight positions, or cast shadows that are witnessed on the display.

To address this aspect of display technology we introduce a computer-controlled spatially-varying reflectance display: a computer-controlled array of dots that can each independently produce a user-defined reflectance function. Instead of specifying a gray level or RGB value at every dot, the control signal to the display specifies a set of per-dot angular reflectance functions, and these indicate the amount of light each dot will reflect for different configurations of the observer’s viewpoint and the lighting in their environment. A single control signal produces a “reflectance image” that passively allows the observer to witness different things as

*Part of this work was done while the author was at The Weizmann Institute of Science.

they move their head, move the display, or alter their environment's lighting; and the display can be updated at up to 60 Hz to produce dynamic "reflectance videos" like the one illustrated in Fig. 1.

Most previous approaches to light-sensitive displays [Nayar et al. 2004; Cossairt et al. 2008; Hirsch et al. 2009; Hirsch et al. 2012; Hirsch et al. 2013; Horisaki and Tanida 2013] require an active feedback loop, with illumination sensors that measure the ambient light, and dedicated hardware that quickly re-renders or looks-up images of the scene in the corresponding lighting configuration. In contrast, our display is *passive*. It does not have its own source of radiant power; it operates purely by controlling how it reflects the light it receives from the environment. This has two main advantages. First, it draws only negligible power to maintain its state, making it more energy efficient and mobile. Second, it reacts to illumination changes instantly, at the speed of light.

Our display builds on the wave optics approach to BRDF fabrication introduced by Levin et al. [2013]. That approach uses a lithographic process to permanently etch patterns of micron-sized features onto silicon wafers, thereby modulating the phase of incoming light waves to control reflections under different view and lighting configurations. This allows fabricating spatially-varying reflectance functions at spatial resolutions up to 220 reflectance dots per inch (dpi), with each reflectance dot exhibiting two dimensions of angular variation, parameterized by the 2D half-vector. In this paper, we create spatially-varying reflectance functions with similar resolution that are *programmable*. Our approach uses a spatial light modulator (SLM) consisting of an array of liquid crystal cells. The refractive index of each cell can be independently controlled by adjusting its voltage; and by programming the refractive indices across the array, we modulate the phase of incoming light waves to produce *dynamic* spatially-varying reflectance.

Due to fabrication constraints, Levin et al. [2013] were restricted to using patterns of planar facets at only two distinct height levels, and hence achieved only binary modulations of phase. This imposed several restrictions on the achievable 2D reflectance: i) all reflectances are symmetric about the half-vector; ii) all reflectances exhibit a specular spike; and iii) tradeoffs between spatial and angular resolutions are sub-optimal. All of these restrictions are overcome by our approach, because SLMs enable a dense set of phase modulations via fine-scale refractive index control in each liquid crystal cell. To exploit this, we introduce a new surface design algorithm that replaces the stochastic sampling of Levin et al. [2013] with a direct optimization. This facilitates a wider variety of reflectances, including non-symmetric ones, and achieves better tradeoffs between angular and spatial resolutions. Our ability to induce 2D reflectances that are *non-symmetric* functions of the half-vector is particularly useful, since it allows displaying realistic view and lighting effects, such as depth parallax, shading variations, and the motion of glossy highlights and cast shadows.

We demonstrate an initial prototype with 54×96 dots, and a dot pitch of $160 \mu\text{m}$. Each dot of the display is comprised of multiple SLM cells, and can be independently programmed to induce custom 2D reflectance functions with a refresh rate of 60 Hz. The cell pitch imposes a constraint on the angular range of view and light directions the display can support. In our prototype, this angular range is limited to 5° . The display's angular range and its spatial resolution will increase proportionally with future advances in SLM technology that allow more cells and smaller cell pitch.

1.1 Related Work

[Hullin et al. 2013] provide an excellent survey of computational fabrication and display of material appearance. Below we discuss technologies that are most related to our work.

1.1.1 Active Displays:

One of the first efforts to display relightable objects was by Nayar et al. [2004], who proposed a device that measures 2D illumination conditions using a wide-angle camera embedded in the display frame, and renders and displays the corresponding 2D image in real time. Koike et al.'s BRDF display [2008] extends this by adding a layer of lenslets to the LCD panel, obtaining an autostereoscopic display. Other active approaches that sense the 4D illumination field include [Cossairt et al. 2008; Hirsch et al. 2009; Hirsch et al. 2012; Hirsch et al. 2013; Horisaki and Tanida 2013]. The main disadvantage of active systems is the expense of computation, which requires dedicated hardware and significant amounts of power.

1.1.2 Passive Displays

In contrast to the active setups described above, passive displays do not have their own source of radiant power. They operate by reflecting the light they receive from the environment. Passive setups can be further classified into those that do not have any electronically-controlled parts in the optic pathway, such as fabricated surfaces with custom BRDFs or the device of [Fuchs et al. 2008]; and those that electronically adjust the displayed content, such as our display and the E-ink displays mentioned below.

E-ink: The best-known working principle used for passive digital displays is electrophoretic reflective technology, branded as E-ink [Comiskey et al. 1998]. It is based on microcapsules containing black and white pigments that are oriented by applying an external electric field. This technology is very attractive for its modest power consumption and its readability in bright environments. Color extensions of this technology utilize transparent color filter arrays.

BRDF fabrication is another passive approach, aimed at generating surfaces with user-defined appearance properties [Weyrich et al. 2009; Finckh et al. 2010; Papas et al. 2011; Kiser et al. 2012; Dong et al. 2010; Hašan et al. 2010; Matusik et al. 2009; Patow and Pueyo 2005; Patow et al. 2007; Weyrich et al. 2007; Malzbender et al. 2012; Lan et al. 2013]. A number of these works [Weyrich et al. 2009; Papas et al. 2011] are based on micro-facet theory [Torrance and Sparrow 1967] and similar geometric optics extensions, modeling the surface micro-structure as a collection of small facets, usually mirrors. The facets' local orientations can be adjusted to control the overall spread of reflected light. While this approach produces impressive results, the geometric optics model is valid only for relatively large facets, limiting the spatial resolution of BRDF dots. Levin et al. [2013] proposed a BRDF fabrication approach that takes into account fine-scale wave optics effects. This allows producing spatially-varying reflectances at resolutions high enough to match that of a human observer. Their approach, reviewed in Sec. 3.1, serves as the basis for our display.

Programmable BRDFs: Pioneering attempts to create dynamic, computer-controlled BRDFs were made by [Hullin et al. 2011; Ochiai et al. 2012; Ochiai et al. 2013]. These techniques are limited to a single BRDF without spatial variation, akin to a BRDF display with a single dot. Hullin et al. [2011] create a programmable BRDF by placing two actuators on the surface of water and exciting waves to achieve varying degrees of anisotropic roughness. Ochiai et al. [2012; 2013] projected images through a bubble of colloidal material. The bubble is modulated by ultrasonic vibrations, allowing control of the transparency and surface state.

1.1.3 Holography

Holography is based on recording the fine interference pattern between a coherent reference plane wave and the waves scattered from a 3D object [Redman 1968; DeBitetto 1969; Klug et al. 1993;

Benton and Bove 2007; Yaroslavsky 2004]. When the recorded fringe pattern is viewed under proper illumination, the observer perceives the object’s 3D parallax. While basic holograms rely on coherent monochromatic illumination, extensions like rainbow holograms [Benton and Bove 2007] can be viewed under white light while sacrificing vertical parallax. Volume holograms [Kogelnik 1969; Goodman 1968] allow for more general white light operation but their design and construction is significantly harder. Computer generated holograms induce the experience of virtual 3D objects that need not exist in the real world [Dallas 1980; Tricoles 1987; Lucente 1993; Ahrenberg et al. 2006; Ahrenberg et al. 2008; Yaroslavsky 2004], and programmable holographic displays based on these ideas have been implemented using SLM technology [Benton 1991; St-Hilaire et al. 1992; Kim et al. 2005; Smalley et al. 2013]. Generating such holograms often requires a time-consuming computation, which starts with a sufficiently-fine point-cloud representation of the object and calculates the fringe field using Fresnel or Fourier transforms. The *hogel* concept introduced by [Lucente 1994] and other works such as [Ziegler et al. 2007] offer more efficient and practical methods for converting a light field into a hologram. These methods are designed for monochromatic illumination but nevertheless can be used to display color holograms with active displays by multiplexing three monochromatic holograms and illuminating with red, green, and blue lasers. In contrast, our reflectance display is a passive one designed to work under ambient, full-spectrum (white) illumination.

2 Spatial Light Modulation Technology

A light wave impinging on a surface accumulates spatially-varying changes as it interacts with micro-scale surface structures. By designing the micro-scale structure we can control the shape of the scattered wave. To implement our dynamic reflectance display we require an SLM—a computer controlled array of microscopic cells that allow us to modulate the reflected light wave. In this section we briefly survey existing SLM technology. SLM devices can modulate either the intensity or the phase at each cell. As explained in Appendix B, intensity modulation only allows generating symmetric reflectance functions, and to implement a general family of reflectances we are interested in phase modulation. As explained below, two existing technologies for phase modulation operate by 1) manipulating the cell’s height; 2) manipulating the cell’s refractive index.

Implementing phase modulation: As light interacts with the surface of phase-based SLMs, it accumulates small, spatially-varying, wavelength-dependent phase changes. We represent this accumulated effect by a function called the phase modulation $\phi_\lambda(x)$, which controls the shape of the scattered wave and is at the heart of reflectance design.

In [Levin et al. 2013] the incident lightwave is reflected from a silicon wafer with micron-sized height variations $z(x) = z_0 + \Delta z(x)$ as illustrated in Fig. 2(a). The light reflected from surface point $(x, z(x))$ travels an additional distance of $2\Delta z(x)$ compared to light reflected at (x, z_0) , and thus it accumulates a phase change (or phase retardation) of $\phi_\lambda(x) = 4\pi\Delta z(x)/\lambda$.

In contrast, a refractive index SLM consists of micron-sized transparent cells over a reflective silicon mirror. The cells all have the same height, but their refractive indices are electronically-controlled as shown in Fig. 2(b). The incoming light travels through the refractive layer and is reflected by the silicon mirror. Since the cells have different refractive indices, the light accumulates a spatially-varying phase change of $\phi_\lambda(x) = 4\pi dn_\lambda(x)/\lambda$, where $n_\lambda(x)$ denotes the spatially-varying, wavelength-dependent refractive index at x , and d is the thickness of the layer.

Existing technology: Microscopic height variations can be achieved with piston-based digital micromirror devices (DMD).

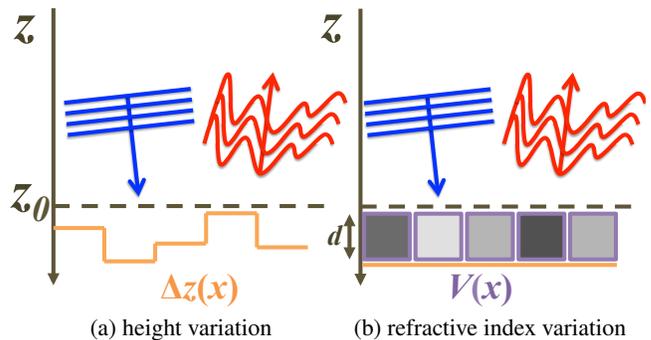


Figure 2: Implementing phase modulation. Two methods for implementing phase modulation are shown. An incoming plane wave with wavelength λ is shown in blue; the wave reflected from the surface in red. In (a), phase is modulated by a surface made of small steps with varying heights (5 steps are shown). A step at $(x, z_0 + \Delta z(x))$ causes a phase modulation of $\phi_\lambda(x) = 4\pi\Delta z(x)/\lambda$. On the right, a similar effect is achieved using liquid crystal cells over a reflective silicon layer. The LC cells are transparent cells of thickness d , whose refractive index $n_\lambda(x)$ can be adjusted by applied voltage $V(x)$. The resulting modulation is given by $\phi_\lambda(x) = 4\pi dn_\lambda(x)/\lambda$.

The number of mirrors and pitch of current piston-based devices is not sufficient for our purposes.

Liquid-crystal SLMs achieve manipulation of the second kind with an array of microscopic cells whose refractive indices can be dynamically adjusted by applying voltage. Multiple types of these SLMs are available, and we use *nematic* liquid crystals. The liquid crystal molecules have an elongated structure and exhibit *birefringence*, meaning that they exhibit different refractive indices for light polarized in different directions. By applying voltage to the cell we can rotate the molecules and change the birefringence properties. This can be exploited in multiple ways. If the light passes through a linear polarizer which transfers only the extraordinary polarization component, effectively we adjust the refractive index of the cell, and the LCD acts as a *phase* modulator, which is our main interest here. When two crossed polarizers are used, *amplitude* modulation is achieved, a technology commonly used in computer monitors. Lanman et al. [2011] use a liquid crystal array as a polarization *rotator*, an effect which can be achieved by combining it with a phase plate.

To achieve phase modulation, we use the liquid crystal on silicon “PLUTO” SLM by Holoeye Inc. It consists of 1920×1080 refractive liquid crystal cells above a reflective silicon background layer. The cell pitch is $8 \mu\text{m}$. A general drawback of our approach is that we must add a linear polarizer in front of the SLM, because cell modulation applies only for polarized light. This induces a 50% decrease in reflective light efficiency, but since the polarizer can be physically attached as an integral part of the display, we need not make any assumptions about the polarization of the ambient illumination.

3 Controlling Reflectance

In this section we review how the reflectance from a surface can be controlled using microscopic structures. The main result is that angular reflectance is the Fourier transform of the microscopic spatial phase pattern. Using the Fourier relation we analyze the tradeoffs between the achievable spatial and angular resolution and range.

To simplify notation we restrict our attention to vectors that lie in a plane. As shown in Fig. 2, we use coordinates x and z respectively

symbol	meaning	units
\mathbf{l}	lighting direction	unit vector
\mathbf{v}	viewing direction	unit vector
\mathbf{h}	half-vector	
λ	wavelength	mm
ϕ	phase	
V	voltage	volts
x	spatial argument	mm
ω	angular argument	radians
\mathbf{a}	intermediate modulation function	
\mathfrak{A}	Fourier transform of \mathbf{a}	
Δ_d	spatial pitch	mm
Δ_a	angular pitch	radians
Ω_d	spatial range	mm
Ω_a	angular range	radians
Δ_c	coherence length	mm
Δ_o	SLM pitch	mm
q	number of cells in a dot	

Table 1: Notation and units.

to index positions tangential and orthogonal to a reference surface plane. Table 1 provides a list of other symbols.

3.1 Reflectance Under Extended Sources

To understand the relation between the phase modulation and reflectance, we review the main results from [Levin et al. 2013]. Our analysis begins with the notion of a display *dot* Δ_d , which is a reflectance unit of the display. When the display is viewed at a distance such that a display area of Δ_d is mapped to the area smaller than a single camera pixel, or is below the diffraction blur limit of the observer’s eye, the appropriate reflectance effects will be witnessed. While the reflectance is a spatially-varying function, in what follows we analyze the reflectance at a single dot, and let x parameterize the position within it.

We denote the observer’s viewing direction by unit vector $\mathbf{v} = (\mathbf{v}_x, \mathbf{v}_z)^T$, and we initially assume a coherent source in one direction, denoted by unit vector $\mathbf{l} = (\mathbf{l}_x, \mathbf{l}_z)^T$. We denote by $\mathbf{h} = (\mathbf{h}_x, \mathbf{h}_z)^T = (\mathbf{l} + \mathbf{v})/2$ the *half-vector*, which is not unit-length (cf. [Rusinkiewicz 1998]). The reflectance functions derived below are a function of the half-vector alone rather than a general function of both lighting and viewpoint.

Let us define an intermediate function

$$\mathbf{a}_\lambda(x) = e^{i\phi_\lambda(x)}, \quad (1)$$

where $\phi_\lambda(x)$ is the spatially-varying, wavelength-dependent phase modulation induced by small patterns of height or refractive index. The Fourier transform of \mathbf{a} , evaluated over the dot width Δ_d , is

$$\mathfrak{A}_\lambda(\omega_x) \propto \int_{x \in \Delta_d} e^{-i2\pi\omega_x x} \mathbf{a}_\lambda(x) dx. \quad (2)$$

Following a standard derivation (see for example Appendix A of [Levin et al. 2013] and references within) the spectral reflectance under coherent light is calculated as the intensity (squared amplitude, or *power spectrum*) of the Fourier transform, evaluated at the wavelength-adjusted half-angle direction:

$$R_c(\mathbf{h}, \lambda) = \left| \mathfrak{A}_\lambda \left(\frac{2\mathbf{h}_x}{\lambda} \right) \right|^2. \quad (3)$$

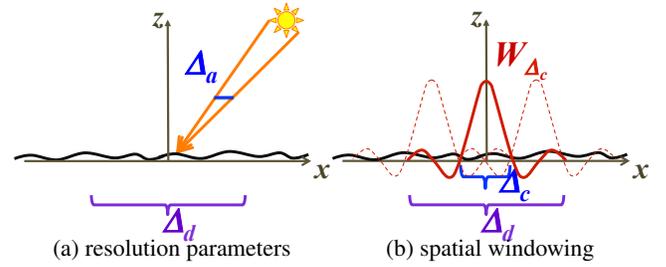


Figure 3: Resolution parameters in incoherent illumination. (a) *Angular resolution:* illumination from an area light source covering a subtended angle Δ_a . *Spatial resolution:* The dot size Δ_d . (b) *Incoherent reflectance* is formed by applying coherent reflectance rules (Fourier transform) over coherent windows of size $\Delta_c = \lambda\Delta_a^{-1}$. The reflectance from all coherent areas inside a dot area Δ_d is averaged to form the incoherent reflectance.

Natural illumination: Spatially-coherent laser sources produce plane waves in a single direction, which can be thought of as emitted by infinitesimally narrow point sources. In contrast, natural sources have finite area and thus produce illumination over a range of directions. As shown in Fig. 3, we denote by Δ_a the solid angle that a (far-field) light source, centered at direction \mathbf{l} , subtends relative to the surface.

The analysis of [Levin et al. 2013] shows that the spectral reflectance from a spatially-incoherent source is the averaged coherent *intensity* over all illumination directions within its solid angle. Equivalently, R_{ic} is the rect-filtered power spectrum:

$$R_{ic}(\mathbf{h}, \lambda) = R_c(\mathbf{h}, \lambda) \otimes \Pi_{\Delta_a/\lambda} \quad (4)$$

where $\Pi_{\Delta_a/\lambda}$ denotes a rectangular filter of width Δ_a/λ .

Finally, using the convolution theorem, Levin et al. show that the total reflectance from dot Δ_d under spatially-incoherent illumination is the average of the power spectra computed over all windows of size $\Delta_c \triangleq \lambda\Delta_a^{-1}$ within the dot:

$$R_{ic}(\mathbf{h}, \lambda) \propto \int_{x_0 \in \Delta_d} \left| \int \mathbf{a}_\lambda(x - x_0) W_{\Delta_c}(x) e^{-\frac{2\pi i}{\lambda}(2\mathbf{h}_x x)} dx \right|^2 dx_0. \quad (5)$$

Here, W_{Δ_c} is a sinc window of width Δ_c as shown in Fig. 3(b). The quantity $\Delta_c = \lambda\Delta_a^{-1}$ is called the *coherence length* of the illumination. It is the small region over which the incoherent light at the surface still exhibits significant cross-correlation (and therefore significant interference). The coherence length is inversely proportional to the solid angle of the source.

Based on these derivations, the reflectance design problem is as follows. We are given the intended solid angle of the source Δ_a and the desired spectral reflectance $R_{ic}(\mathbf{h}, \lambda)$, and we are tasked with choosing a dot pitch $\Delta_d \geq \Delta_c$ and solving for a modulation function $\phi_\lambda(x)$ by inverting Eqs. (4) or (5). As explained below, choosing the dot pitch is important for controlling resolution and design flexibility.

3.2 Spatial and Angular Tradeoffs

Before proceeding to reflectance design it is important to understand what can be achieved in terms of spatial and angular resolution and range. Here, we explain the relevant parameters, their relation to the SLM characteristics, and the design tradeoffs.

Our reflectance display will have a certain spatial extent Ω_d determined by the size of the SLM. Likewise, it will control reflection over a certain range of view and light configurations satisfying

$-\Omega_a \leq \mathbf{l}_x + \mathbf{v}_x \leq \Omega_a$ for some *angular range* Ω_a . The SLM cell pitch, denoted by Δ_o , is the physical quantity that controls the angular range Ω_a through the Fourier relationship of Eq. (3), in the following way. The Nyquist-Shannon sampling theorem guarantees that by sampling at Δ_o -intervals one can reconstruct a band-limited signal with frequencies in the range $[-\Omega_w, \Omega_w]$ where $\Omega_w = 0.5\Delta_o^{-1}$. Since the reflectance at \mathbf{v} , \mathbf{l} is given by the Fourier transform evaluated at frequency $w = (\mathbf{l}_x + \mathbf{v}_x)/\lambda$, the reflectance at wavelength λ can be viewed (“reconstructed”) in the angular range $-\Omega_a^\lambda \leq \mathbf{l}_x + \mathbf{v}_x \leq \Omega_a^\lambda$, where

$$\Omega_a^\lambda = \frac{\lambda}{2\Delta_o}. \quad (6)$$

Beyond this range, there are low energy side-lobe replicas caused by the piecewise-constant structure of the SLM’s phase modulation, as we analyse in Appendix A.

Within these ranges Ω_d and Ω_a , the spatial and angular resolutions are determined by the sizes (and thus number) of spatial and angular “dots” in which we have independent control of the reflectance values. As introduced in Sec. 3.1 we denote the *spatial dot pitch* by Δ_d , and the *angular pitch*, by Δ_a . It is helpful to interpret the angular pitch as a lower bound on the solid angles of the sources that will illuminate the display. If the sources are narrow enough for the solid angles to drop below this bound, there will be insufficient angular smoothing in Eq. (4), and our design will not guarantee artifact-free appearance. Sources larger than this bound do not pose a problem, however, since they will simply induce an angular integral of the reflectance, analogous to rendering with BRDFs and extended sources.

Trade-offs and multiplexing: We have design flexibility when choosing dot pitch Δ_d and angular pitch Δ_a , but these values are closely linked. In particular, the number of SLM cells that are spanned by one display dot, $q = \Delta_d/\Delta_o$, induces a lower bound on the angular pitch the device can support. By Eq. (3) the reflectance from a dot is a Fourier transform of q controllable SLM cells and thus within the angular range $[-\Omega_a^\lambda, \Omega_a^\lambda]$ we can independently control q angular dots of size

$$\tilde{\Delta}_a \triangleq 2\Omega_a^\lambda/q = \lambda/(q\Delta_o) = \lambda/\Delta_d. \quad (7)$$

We call this the *Fourier angular pitch*. In practice, it is useful to target a dot pitch Δ_d whose induced Fourier angular pitch $\tilde{\Delta}_a = \lambda\Delta_d^{-1}$ is lower than the desired angular pitch, so that $\Delta_a > \tilde{\Delta}_a$, because it facilitates spatial multiplexing and provides more flexibility in reflectance design. As described in Sec. 3.1, the angular pitch defines the coherence length ($\Delta_c = \lambda\Delta_a^{-1}$), and when the dot pitch is larger than the coherence length $\Delta_d > \Delta_c$, the reflectance is the average of the power spectra of multiple Δ_c -sized lengths inside the dot. This provides design flexibility because instead of finding a single phase modulation $\phi_\lambda(x)$ over length $\Delta_c = \Delta_d$ to match a desired reflectance, we can find multiple phase modulations over lengths $\Delta_c < \Delta_d$ such that their *averaged* power spectrum produces the desired reflectance.

Sampling vs. optimization design: The work of Levin et al. [2013] is an extreme example of multiplexing. Due to fabrication constraints, their etching process was limited to binary height values and hence to binary phase modulations $\phi_\lambda(x) \in \{\phi_1, \phi_2\}$. In addition to restricting reflectance to be symmetric and contain a specular spike as reviewed in Appendix B, this presents the computational challenge of inverting Eqs. (4) and (1) over binary variables. Their solution is to use a dot pitch that is much larger than the coherence length, allowing one to generate binary Δ_c -sized phase modulations pseudo-randomly using a suitable stochastic process, and still have them converge, by averaging, to the desired reflectance.

In contrast, our SLM can instantiate a dense range of distinct phase modulations, so we are not restricted to binary values. Therefore, in Sec. 4 we choose to *optimize* for an efficient phase modulation, rather than use a pseudo-random one. This approach allows constructing a wider range of reflectances. It also provides a better tradeoff between angular and spatial resolutions, because by explicitly optimizing, we can achieve the desired reflectance with fewer Δ_c -sized areas within a dot. Thus, for a fixed angular resolution Δ_a , which in turn determines the coherence area Δ_c , we can work with smaller dots. Alternatively, given a fixed dot pitch Δ_d , we can achieve finer angular resolution Δ_a .

Specifications of our prototype: The SLM we use has spatial range $\Omega_d = 15.36\text{mm} \times 8.64\text{mm}$, with 1920×1080 cells and a cell pitch of $\Delta_o = 8\ \mu\text{m}$. To generate a single spatial dot of reflectance, we use $q = 20$ cells, resulting in a dot pitch of $\Delta_d = 160 \times 160\ \mu\text{m}^2$ and a spatial resolution of 54×96 reflectance dots.

The Fourier angular pitch induced by the $160\ \mu\text{m}$ dot pitch is very small: $\tilde{\Delta}_a \approx 0.17^\circ$ at $\lambda = 0.5\ \mu\text{m}$. We allow for threefold spatial multiplexing by setting the angular pitch to $\Delta_a = 0.5^\circ$. This is a reasonable lower bound on the extent of typical light sources. For example, it is the angle subtended by a 5cm source placed less than 6m from the display. (As a point of reference, Levin et al. [2013] targeted a larger angular pitch of $\Delta_a = 2^\circ$.)

Increasing range and resolution: The spatial range of the display scales linearly with the physical extent of the SLM device. It can also be extended by tiling multiple SLMs. Our prototype’s relatively large $8\ \mu\text{m}$ cell pitch limits the angular range Ω_a (via Eq. (6)) to approximately 5° for red light and 3° for violet. Since smaller cell pitch translates directly to larger angular range, this limitation will soften substantially as hardware improves. Devices with $6\ \mu\text{m}$ cell pitch are already commercially available (LETO, Holoeye Inc.), and devices with $2\ \mu\text{m}$ cell pitch are in development [IMEC 2011]. Complementary to these developments, one can also increase angular range by viewing the device through a lens. Basic lens rules imply that the product of the spatial magnification and angular magnification is constant, regardless of lens power and distance. Adding a lens to a reflectance display of size Ω_d and angular range $[-\Omega_a, \Omega_a]$ will produce an (inverted) display of size $M\Omega_d$ and angular range $[\frac{\Omega_a}{M}, \frac{-\Omega_a}{M}]$. Thus, adjusting the distance between the display and the lens one can vary M and increase the angular range. However, this solution comes at a cost of reducing the spatial range Ω_d , so is especially useful when tiling multiple SLMs spatially.

4 Reflectance Design

In this section we discuss the problem of reflectance design: given a target reflectance function we wish to find a phase modulation that generates this reflectance. We start by discussing design constraints imposed by the SLM characteristics. Next, we formulate an optimization problem that accounts for the complete Fourier optics model over all wavelengths simultaneously. As the resulting optimization problem is highly non-convex, we introduce an initialization strategy based on geometric optics principles.

4.1 Design Constraints

SLM calibration: For each illumination wavelength, one can calibrate the SLM to obtain a mapping between the voltage assignment $V(x, y)$ and the resulting index of refraction of each cell, or more directly, to the phase modulation of the lightwave, $\phi_\lambda(V)$. If a voltage assignment V produces a refractive index $\eta_\lambda(V)$ and the cell thickness is d , we achieve a modulation

$$\phi_\lambda(V) = \frac{4\pi d n_\lambda(V)}{\lambda}. \quad (8)$$

In practice the refractive index varies between wavelengths. Fixing a reference wavelength λ_o , we denote refractive index variation as

$$\tilde{\alpha}_\lambda = \frac{n_\lambda(V)}{n_{\lambda_o}(V)}. \quad (9)$$

The voltage-to-phase mappings at different wavelengths are linearly related:

$$\phi_\lambda(V) = \alpha_\lambda \phi_{\lambda_o}(V), \quad \text{with: } \alpha_\lambda = \frac{\lambda_o}{\lambda} \tilde{\alpha}_\lambda. \quad (10)$$

This simple linear scaling allows formulating our optimization over the modulation $\phi = \phi_{\lambda_o}$ instead of over the voltage assignment V .

We calibrate the voltage-to-phase mapping $\phi_\lambda(V)$ using an interferometric method suggested by Holoeye [Hermerschmidt et al. 2007]. Setting $\lambda_o = 633\text{nm}$, we find the refractive index ratios $\tilde{\alpha}_\lambda$ over the visible spectrum to be within the range [0.996, 1.284].

Modulation range: The thickness of the LC cells in the SLM implies an upper bound on the maximal phase modulation it can apply. The SLM we used allows for a phase modulation of up to 6π . That is, for $\lambda_o = 633\text{ nm}$, $\phi_{\lambda_o}(V)$ can take values in the range $[0, 6\pi]$. Since the modulation function in Eq. (1) depends only on ϕ modulo 2π , this means that for a given wavelength there are multiple voltage assignments that produce the same phase modulation. Since we are optimizing to achieve a target reflectance at multiple wavelengths simultaneously this redundancy offers us more flexibility and facilitates optimization.

4.2 Optimization

Given a target angular reflectance function $R_T(\mathbf{h})$, we wish to find a phase modulation $\phi(x, y) = \phi_{\lambda_o}(V(x, y))$ using a collection of $q \times q$ SLM cells (in our implementation $q = 20$) that produces the desired angular reflectance at many wavelengths over the visible spectrum. This optimization problem is reminiscent of that of computer generated holograms (CGH). Our problem differs from the CGH setting, however, in that we are looking for a phase modulation that produces the target reflectance at multiple wavelengths in the visible range *simultaneously*. This is a more complex optimization problem that cannot be solved by algorithms typically used for CGH, such as Gerchberg-Saxton [1972].

To simplify notation we present the one-dimensional case in which the modulation and reflectance are functions of x . We minimize

$$L(\phi) = \sum_{\lambda \in \Lambda, \mathbf{h}_x \in \mathcal{H}} (|R_T(\mathbf{h}_x) - R(\mathbf{h}_x, \lambda; \phi)|^2) + \eta \|\nabla(\phi)\|_1 \quad (11)$$

s.t. $0 \leq \phi \leq 6\pi$ (12)

where Λ denotes a discrete sample of wavelengths in the visible range, and $\mathcal{H} \subseteq \Omega_a$ is a discrete set of angles. $R(\mathbf{h}_x, \lambda; \phi)$ is the reflectance generated by the SLM with a phase modulation ϕ , and $\|\nabla(\phi)\|_1$ is a regularization term that promotes smoothness in the spatial variation of the modulation ϕ , weighted by a scalar η .

$R(\mathbf{h}_x, \lambda; \phi)$ can be computed using Eq. (4) with

$$\mathbf{a}_\lambda(x) = \exp\{i\alpha_\lambda \phi(x)\}. \quad (13)$$

Note that using the model of Eq. (10), we have $\mathbf{a}_\lambda(x) = \mathbf{a}_{\lambda_o}(x)^{\alpha_\lambda}$, and $R(\mathbf{h}_x, \lambda; \phi)$ is simply the blurred power spectrum of $\mathbf{a}_{\lambda_o}(x)^{\alpha_\lambda}$.

Optimization strategy: Since joint optimization over all 1920×1080 SLM cells is impractical, we optimize the phase independently at each dot, assuming the phase of neighboring dots is repeated periodically. This simplification allows us to efficiently evaluate Eq. (5) using the fast Fourier transform. The resulting opti-

mization is exact in uniform regions, but can suffer from artifacts in nonuniform ones if large variations are present between the reflectances of neighboring dots. We use Matlab's nonlinear optimization to minimize Eq. (11).

As mentioned in Sec. 3.2, the angular resolutions of our reflectance functions is limited by the angular pitch Δ_a . To allow for a feasible solution it is important to ensure that the target $R_T(\mathbf{h}_x)$ has no energy at higher angular frequencies. Thus, when required, we low-pass filter the target angular functions accordingly.

The nonlinear optimization can benefit greatly from a good initialization. Here we suggest an initialization strategy based on geometric optics principles.

Geometric optics initialization: A micro-facet model such as Torrance-Sparrow [1967] models the surface as a collection of small, randomly-oriented facets, with each facet behaving as an ideal mirror. The distribution of facet normals within a dot determines the intensity of light reflected in different directions. Following this idea, we seek a surface $S = (x, z(x))$, for which the distribution of normals matches the target reflectance $R_T(\mathbf{h}_x)$. As discussed in Sec. 2 a surface with height $z(x)$ applies a phase modulation of $\phi_\lambda(x) = 2z(x)\frac{2\pi}{\lambda}$ to a reflected monochromatic wave. Thus, given a surface S whose gradient distribution matches R_T we initialize our wave optics optimization with a phase modulation:

$$\phi^{init}(x) = 2z(x)\frac{2\pi}{\lambda_o}. \quad (14)$$

To compute $z(x)$ for $x \in [0, \Delta_d]$, we treat R_T as a probability density and compute a surface whose gradients are distributed according to this density. Specifically, we compute

$$z(x) = \Delta_d \int_0^{\frac{x}{\Delta_d}} G(y) dy, \quad (15)$$

where $G(y): [0, 1] \mapsto [-\Omega_a, \Omega_a]$ is a 'gradient density' function whose values are distributed according to R_T . To obtain function G we define the cumulative reflectance distribution

$$\tilde{R}_T(\mathbf{h}_x) = \int_{-\Omega_a}^{\mathbf{h}_x} R_T(\omega) d\omega \quad (16)$$

and G as its generalized inverse

$$G(y) = \inf_{\mathbf{h}_x} \left\{ \tilde{R}_T(\mathbf{h}_x) \geq y \right\}. \quad (17)$$

In Fig. 4(a) we show an example of a target R_T and its cumulative distribution $\tilde{R}_T(\mathbf{h}_x)$. We also compare the normalized histogram of G to R_T .

If the resulting initialization does not conform with our optimization constraints in Eq. (12), we use $\phi_{\lambda_o}(x) \bmod 6\pi$ and let the nonlinear optimization proceed from there.

The geometric optics initialization is usually quite close to the target reflectance. The differences are mainly due to wave interference and chromatic aberration. The former can be understood as follows. From a geometric optics viewpoint any surface whose gradients are distributed according to R_T should produce the desired reflectance. However, at small scales, interference is prominent and the reflectance can deviate from the geometric prediction. We observe that interference is stronger when the surface slope (G in Eq. (15)) has larger variation, so one of the advantage of the initialization procedure above is that it generates smooth slopes. This is demonstrated in Fig. 4, where the surface in Fig. 4(c) is generated by integrating a random permutation of the slopes of Fig. 4(b). Since both surfaces have equivalent gradient distributions, from a geometric

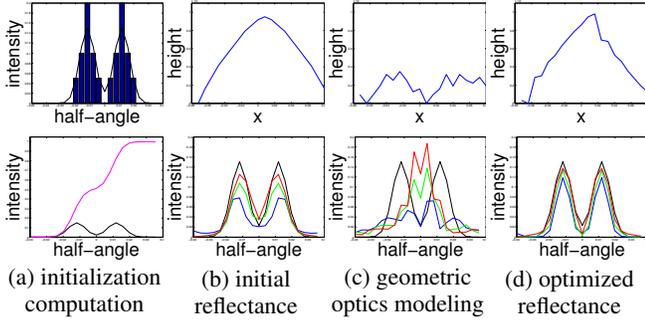


Figure 4: 1D optimization example for a target reflectance R_T , plotted black in the top panel of (a) and in the bottom panels of (a)-(d). Function G from Eq. (17) is computed by inverting the cumulative distribution $\int_{-\Omega_a}^{h_x} R_T(\omega) d\omega$, plotted pink in (a). Top panel of (a) verifies that the normalized histogram of G matches the target R_T . Top panel of (b) shows the height surface $z(x)$ from Eq. (15), and the bottom shows the resulting reflectance $R(\mathbf{h}, \lambda)$ at $\lambda = \{450, 550, 650\}$ using blue, green, and red plots respectively. Panel (c) demonstrates the benefit of smoothness in the geometric-optics initialization. The surface in the top panel has the same gradient distribution as that in (b) and should produce the same reflectance according to geometric optics; but bottom panels of (b) and (c) reveal that reflectances predicted by wave optics are very different. Panel (d) shows the surface and reflectance generated by our non-linear optimization, initialized from surface (b). The optimization reduces chromatic artifacts.

optics viewpoint they should produce the same reflectances. Yet, when the full wave optics model is taken into account, they produce very different reflectances. The reflectance of the smoother surface, Fig. 4(b), is much closer to the target reflectance R_T .

The quality of the geometric optics initialization is also affected by chromatic aberration. Since the refractive index is not perfectly constant across wavelengths (the ratio $\tilde{\alpha}_\lambda$ in Eq. (9) can be as large as 1.2), the conversion from surface to phase in Eq. (14) is imperfect for wavelengths $\lambda \neq \lambda_o$. This results in chromatic artifacts, but these are usually not significant and can be corrected by the non-linear optimization of Eq. (11), as illustrated in Fig. 4(d). The non-linear optimization can be executed from a random initialization as well, but the geometric optics initialization generally leads to better solutions. We observe that the angular shapes of the reflectances produced with random initializations are often reasonable, but that they suffer from more substantial chromatic artifacts than those produced with geometric optics initializations.

Initializing two-dimensional targets: In general, the target reflectance is function of two angular dimensions, \mathbf{h}_x and \mathbf{h}_y . As we will see, in some applications it is separable, or $R_T(\mathbf{h}_x, \mathbf{h}_y) = P(\mathbf{h}_x)Q(\mathbf{h}_y)$, and we can compute a two-dimensional geometric-optics initialization using an outer product of two one-dimensional height functions. For target reflectances that are not separable, we have explored a variety of strategies. For the rotated Gaussian lobe of Fig. 9 we initialize by rotating the surface corresponding to a separable, axis-aligned lobe. For the simultaneous views in Fig. 13, we initialize by subdividing each dot into a spatial mosaic of two separable initializations. Designing initialization algorithms for other reflectances is an interesting direction for future work.

5 Experiments and Applications

We demonstrate various applications of our display, including: displaying custom reflectance images and reflectance videos; interac-

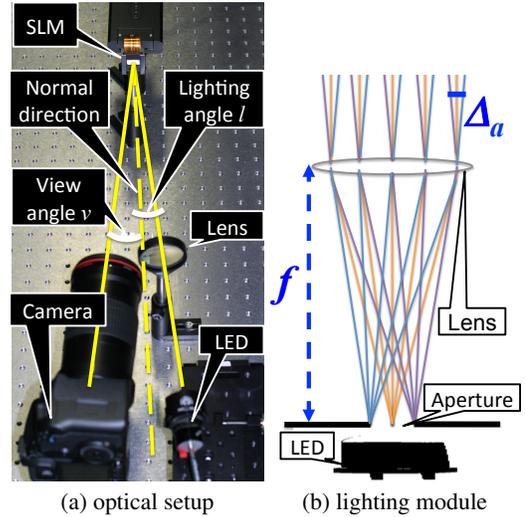


Figure 5: Experimental setup. Our optical setup is shown in (a). The lighting module (seen in the lower right part of (a)) illustrated in (b) consists of an apertured LED and a lens. Every point on the aperture is a spherical point source (three such points are shown in the illustration). The light from the aperture point sources is collimated by the lens, generating plane waves spanning an angular range of size Δ_a . The LED is mounted on an $x - y$ plane translational stage allowing a shift of the illumination direction. Following the optical path in (a), the light from the lens is modulated and reflected by the SLM and then recorded with a camera.

tive reflectance editing; displaying light-sensitive or multi-view 3D scenes; and simultaneously displaying two independent video channels in different view directions.

Our experimental setup is illustrated in Fig. 5. To illuminate the display with a controlled source we place a white LED at the focal distance of a lens. Rays emitted from a single point on the LED are formed into a parallel beam by the lens, and we change the light direction at the SLM by shifting the LED within the focal plane. The LED source is a finite area source rather than an infinitesimally small point source. By placing an aperture between the LED and the lens, we control the active area and thus the illumination solid angle Δ_a , since the light incident on the SLM is a collection of plane waves emitted by all points inside the aperture area. Since the SLM operates on linearly polarized light, a polarizer must be used. Such a polarizer can be attached to the SLM as an integral part of the display. Finally, the modulated light is reflected from the SLM and captured by a camera.

Figure conventions: We visualize reflectance functions as 2D images parameterized by the \mathbf{h}_x (horizontal) and \mathbf{h}_y (vertical) components of the half-vector. When displaying photographs of the SLM, we visualize the corresponding light and view directions with a small polar plot that shows the orthogonal projection of unit-magnitude direction-vectors onto the xy -plane. Lighting direction is indicated by an orange dot and viewing direction by a blue one.

SLM Appearance: Fig. 6 shows a close-up view of the SLM, captured with a macro lens at 1:1 magnification. The image, reminiscent of sensor data captured by a plenoptic camera [Ng et al. 2005], helps us understand the display’s operation. The top two rows in Fig. 6 show how the anisotropic reflectance of Fig. 7 is achieved. Focusing on a single dot (two bottom rows of the figure) we see that when the dot is illuminated at angle l , the portion of the dot that has a gradient in that direction lights up. As we shift the

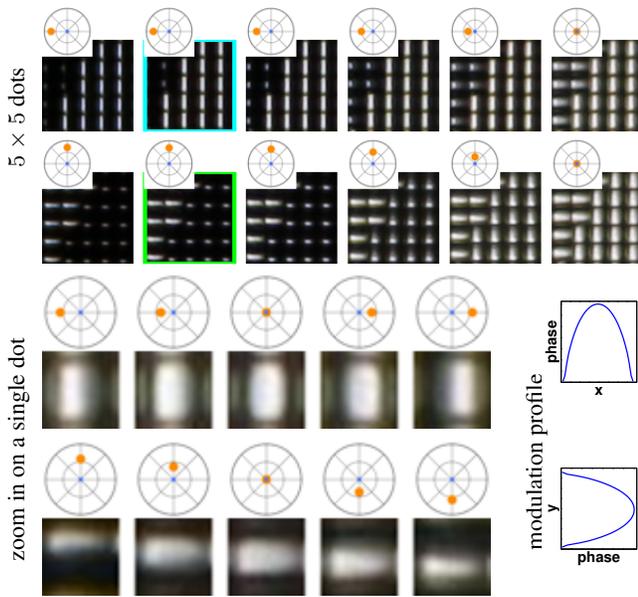


Figure 6: SLM appearance, close up. These images show a close up view of the SLM displaying two anisotropic reflectances at perpendicular orientations, the frames indicated with a green and a cyan boundary correspond to regions from Fig. 7. The images were captured with a macro lens at 1:1 magnification. The different frames in the top two rows show the reflection from one patch of 5×5 dots (each dot is represented by one pixel at the resolution of the images in Fig. 7), under varying illumination directions and a fixed viewpoint, indicated by the corresponding polar plots. The bottom two rows zoom in on the reflectance from a single dot, a profile of the modulation encoded in this dot is shown on the right. Under illumination from direction 1 the dot reflects light from the area in which the gradient of the modulation is 1. When illuminating at angle 1, the dot area with the corresponding slopes lights up. The lit area shifts as the illumination direction varies.

illumination direction, the bright portion of the dot shifts with it.

5.1 Applications

We now present a number of example applications which demonstrate the utility and capabilities of our display.

Reflectance images: Our display can show images with custom, user-defined reflectance functions. For example, Fig. 7 shows a reflectance image with two anisotropic reflectance functions. These appear bright from one light direction and dark from another. Switching between horizontal and vertical illuminations, we observe the negative image, and the optical illusion is inverted. Fig. 8 demonstrates an unusual “anti-mirror” reflectance, which reflects light in oblique directions but not in the mirror direction. Note that our display does not exhibit the specular spike artifact that was seen in the anti-mirror reflectance created photolithographically by Levin et al. [2013]. As explained in Appendix B, this is due to the SLM’s dense range of phase modulations, unlike the binary modulations of the photolithographic approach.

Dynamic reflectance videos: The Holoeye PLUTO SLM offers a refresh rate of 60 Hz, and since we don’t need to perform any sensing or on-the-fly calculations as in active systems, we can easily display reflectance videos, or dynamic sequences of reflectance images. Sample clips are shown in the accompanying video.

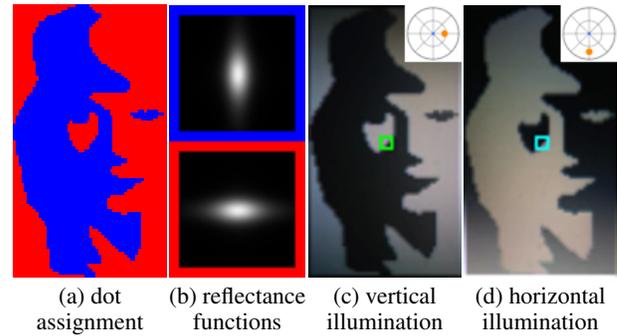


Figure 7: Anisotropic reflectance example. Each dot is assigned one of two anisotropic reflectance functions at opposite orientations. In (c) and (d) we show two images of the SLM taken from the same viewpoint, under different illuminations: vertical and horizontal, respectively. As we change the illumination direction, we observe the negative image and the optical illusion is inverted. The regions indicated by green and cyan squares in (c) and (d) are shown in high resolution in Fig. 6.

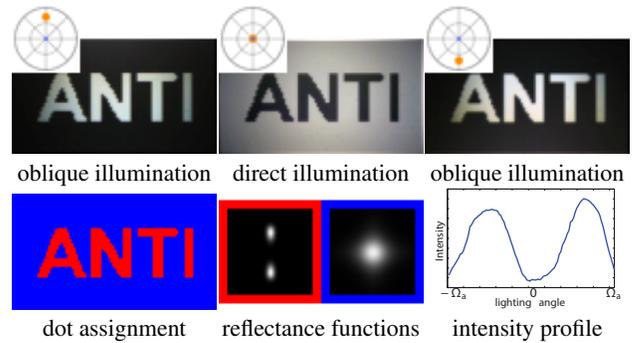


Figure 8: Anti mirror reflectance. The displayed pattern shown on the bottom left consists of a standard isotropic background and an anti-mirror reflectance at the text regions, this anti-mirror reflectance reflects light in oblique illumination directions but not in the mirror direction. The top row presents three images of the SLM taken from different illumination directions. On the bottom right we show a plot of the intensity reflected from the letters as a function of illumination angle. Note that unlike [Levin et al. 2013], we don’t suffer from the specular spike artifact.

Interactive reflectance editing: We can use our display for interactive reflectance editing. Fig. 9 demonstrates the concept, where a user chooses a region of a virtual scene and manipulates its target reflectance interactively. By precomputing the phase modulations for a family of reflectances—parametrized here by lobe width, lobe eccentricity, and lobe orientation—the display is updated to show the scene with the manipulated reflectance in real time. The supplementary video illustrates an interaction session.

Visual artifacts and their causes: In Figs. 7–9 and the following ones, we witness non-idealities caused by a variety of physical phenomena. We see spatial non-uniformity in the lighting of the SLM because the output of our light source is not perfectly uniform and our camera is not orthographic, causing different points on the camera’s sensor to observe the SLM from slightly different angles. Also, since the SLM’s refractive index is not constant across wavelength, the observations contain slight chromatic variations and are not perfectly gray. Finally, since the reflectances are designed and optimized at each dot independently, under the implicit assumption

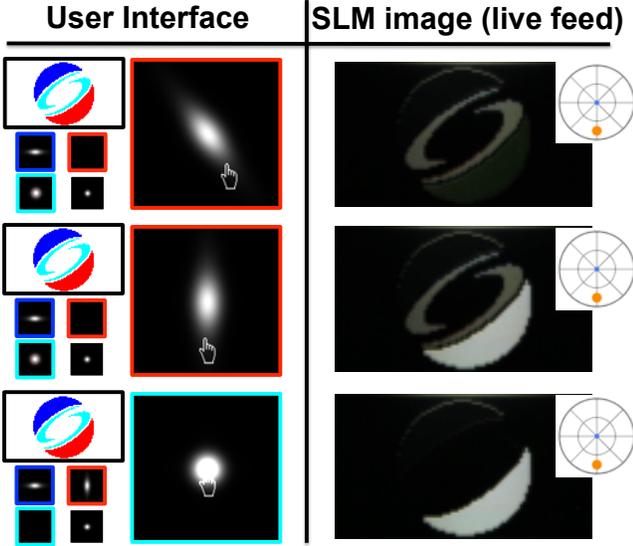


Figure 9: Interactive reflectance editing. Using a computer interface (left) the user manipulates the target reflectance of a virtual, SIGGRAPH-logo scene, modifying the orientation of the reflectance lobe for the red region, and the width of the lobe for the cyan region. The SLM display (right) updates in real time, allowing the user to interact with, and inspect, the scene while moving the ambient illumination. A live interaction session is illustrated in the supplementary video.

that each dot’s pattern is repeated periodically at its neighboring dots, there can be slight inconsistencies at the boundaries between regions that have different reflectances. An example of this is the bright arc at the top of the logo in the right of Fig. 9.

5.1.1 Light-sensitive and View-sensitive 3D scenes

We use our display to create the illusion of a 3D scene by presenting reflectance images that are optimized to emulate 3D shading variations or motion parallax. For a static 3D scene we can display a single reflectance image that, when observed from a fixed viewpoint as the light moves, induces appropriate 3D shading effects such as cast shadows (Fig. 10). Alternatively, we can display a different reflectance image that produces motion parallax when observed from different viewpoints under fixed lighting (Fig. 12). We can achieve the same effects for dynamic 3D scenes, simply by displaying reflectance videos instead of static reflectance images (Fig. 1). These results¹ can be better appreciated in the supplementary video.

First consider illumination variation. Given a fixed view \mathbf{v}_o and a dense set of q illumination directions $\mathbf{l}_1, \dots, \mathbf{l}_q$, we render a corresponding set of target images $\{I_{\mathbf{v}_o, \mathbf{l}_i}\}_{i=1}^q$ of the 3D scene at a spatial resolution that matches our reflectance display (e.g., bottom row of Fig. 10). We define a target reflectance $R_{T,p}$ for dot p of our display that encodes the light-parameterized intensity profile at corresponding pixel p in the target image set:

$$R_{T,p} \left(\frac{\mathbf{v}_o + \mathbf{l}_i}{2} \right) = I_{\mathbf{v}_o, \mathbf{l}_i}(p). \quad (18)$$

Alternatively, we can fix the illumination \mathbf{l}_o and define a target reflectance that encodes the appearance of a 3D scene from a dense

¹The animated horse model is courtesy of Robert Sumner and Jovan Popović from the Computer Graphics Group at MIT. The Ariadne Bust is courtesy of Lincoln 3D Scans.

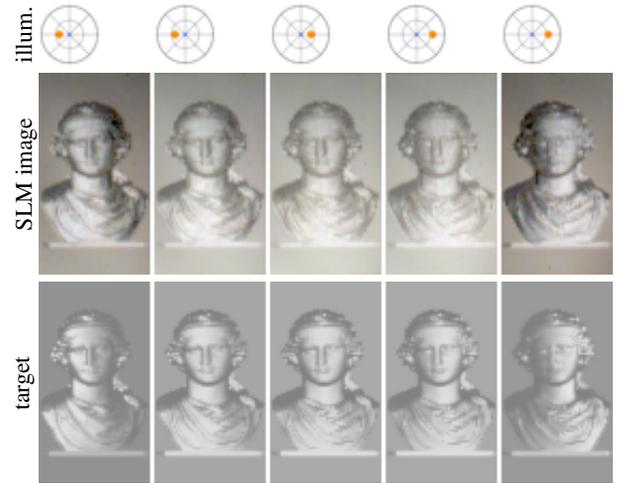


Figure 10: Display of lighting variation. A single reflectance image displayed on our SLM creates the appearance of a 3D model under varying illumination. Renderings of a triangle mesh at multiple illumination directions (as indicated by the polar plots) from a fixed viewpoint, were used to generate a target reflectance at each SLM dot. Images of the SLM from the corresponding illumination directions accurately reproduce highlights and cast shadows (see video).

set of views $\mathbf{v}_1, \dots, \mathbf{v}_q$ (e.g., bottom row of Fig. 12). Since our reflectances depend only on the half-angle, Eq. (18) fully determines the reflectance functions. This means that we do not have additional degrees of freedom to generate reflectances that simultaneously encode separate viewpoint and lighting variations.

When we define 2D target reflectances we must ensure they sum to one. This normalization constraint can be achieved in a variety of ways, but to simplify the design process for our experiments, we simply target view or lighting changes that are constrained to lie in the horizontal plane and use the reflectance values at vertically-offset directions as free variables during normalization. A one-dimensional sequence of target values for a dot $\{I_{\mathbf{v}_o, \mathbf{l}_i}(p)\}_{i=1}^q$ constrains only the x -component of the reflectance $R_T^1(\mathbf{h}_x)$, and we define a full 2D target reflectance as an outer product

$$R_T(\mathbf{h}_x, \mathbf{h}_y) = R_T^1(\mathbf{h}_x) \otimes R_T^1(\mathbf{h}_y), \quad (19)$$

where $R_T^1 \in \mathbb{R}^{1 \times q}$ is a one-dimensional rect filter with width chosen so that $\sum_{\mathbf{h}_x, \mathbf{h}_y} R_T(\mathbf{h}_x, \mathbf{h}_y) = 1$. Given this target, we find the phase modulation using the optimization scheme of Sec. 4.

Note that the ability to create non-symmetric reflectances is critical to producing view and lighting effects like those in Figs. 10 and 12. Such non-symmetric reflectances cannot be achieved with the binary designs of Levin et al. [2013], as explained in Appendix B. Examples of non-symmetric reflectances are shown in Fig. 11.

Our display’s ability to produce view-sensitive parallax effects is very similar to classical holograms, but unlike classical holograms our display does not rely on monochromatic illumination. Similar view-dependent effects can be generated with active light field displays such as [Lanman et al. 2011; Wetzstein et al. 2012]. A non-programmable display of lighting variation similar to our result was produced by Mann [1995], who multiplexed holograms to encode the appearance of an object under multiple lighting directions.

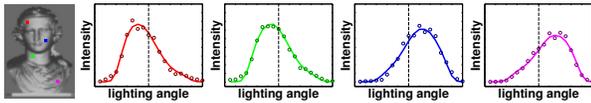


Figure 11: Non-symmetric Reflectance. We show example reflectances from the experiment of Fig. 10. The target reflectance is a solid line and the result of our optimization is shown as circles. Note that these are all non-symmetric reflectances, highlighting the benefits of an SLM producing continuous phase modulation.

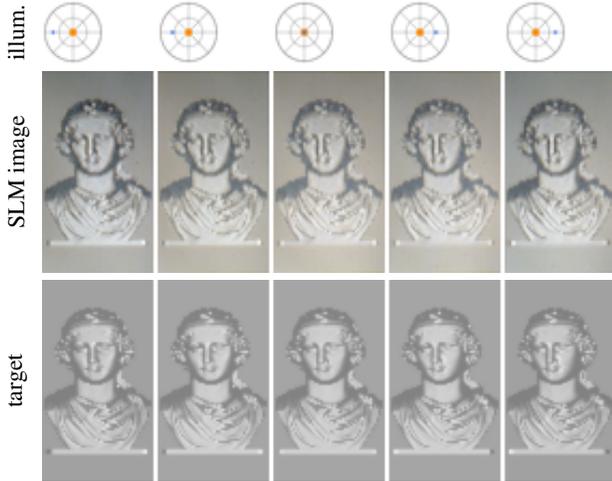


Figure 12: Display of viewpoint variation. A single reflectance image displayed on our SLM creates the appearance of a 3D model from varying viewpoints. Renderings of a triangle mesh at multiple viewing directions (as indicated by the polar plots) under a fixed illumination, were used to generate a target reflectance at each SLM dot. Images of the SLM from the corresponding view directions accurately reproduce depth parallax (see video).

5.1.2 Simultaneous View: Two Sequences on One Screen

Our display allows two people to view two different videos simultaneously on the same screen. We achieve this by playing a single reflectance video that encodes two independent channels. A viewer looking at the display under some fixed illumination \mathbf{l}_o from angle \mathbf{v}_1 sees one sequence while a viewer at angle \mathbf{v}_2 sees another.

In the example of Fig. 13, we simultaneously display two binary-valued videos by multiplexing in time and space copies of four different reflectances. These reflectances are shown in the figure, and they account for the four possible combinations of simultaneous binary values in the two viewing directions. The full sequence can be viewed in the supplementary video.

In this example, the views are symmetric across the x -axis, $\mathbf{v}_1 = (-\theta, 0, 1)$, $\mathbf{v}_2 = (\theta, 0, 1)$, and the reflectances selectively direct light to small angular cones around \mathbf{v}_1 and \mathbf{v}_2 . The reflectances must integrate to unity, so when two dark values are simultaneously required, the residual energy is directed toward extreme vertical angles that are outside of the observers' fields-of-view.

6 Discussion

This work introduces a programmable, spatially-varying, dynamic reflectance display. The display does not rely on illumination-recording sensors, nor does it require expensive on-the-fly computation. It reacts to lighting changes instantaneously and con-

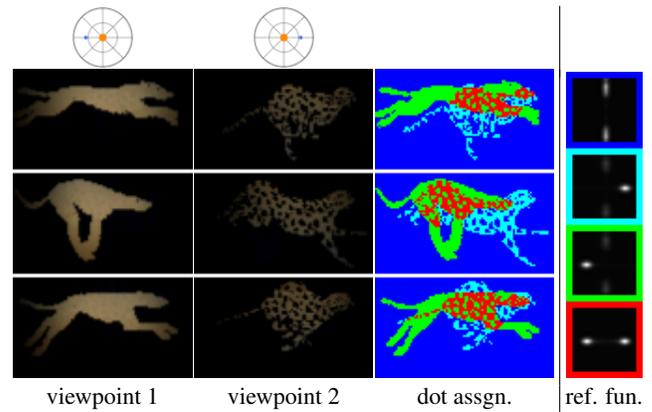


Figure 13: Simultaneous view. We use our display to simultaneously play two independent channels on one screen. We show three frames, as observed from two different viewpoints at the same time. The dot assignments and the reflectances are shown on the right.

sumes only a minimal amount of energy. Our display builds on the wave-optics approach to lithographic BRDF fabrication of Levin et al. [2013]. It maintains the benefits of that approach, such as a high spatial resolution of approximately 160dpi. In addition, it introduces new capabilities, such as: i) display of more general reflectances, including non-symmetric ones; and ii) display of dynamic reflectance videos. To exploit these capabilities, we introduce a new optimization algorithm, which has the added benefit of improving the trade-off between spatial and angular resolutions.

The display suffers from a number of limitations. We do not have the ability to display albedo variations and color. The angular resolution of the display is quite limited and as a result it cannot display real wide angle diffuse reflectances. Finally our prototype is capable of showing fairly small images with only 54×96 dots. However, as discussed in Sec. 3.2 advances in SLM technology will alleviate the angular resolution problem as well as facilitate display of larger images.

Acknowledgements

Funding was provided by BSF, ERC, Intel ICRI-CI, and NSF awards IIS-1012454 and IIS-1212928.

References

- AHRENBURG, L., BENZIE, P., MAGNOR, M., AND WATSON, J. 2006. Computer generated holography using parallel commodity graphics hardware. *Opt. Express* 14, 17 (Aug), 7636–7641.
- AHRENBURG, L., BENZIE, P., MAGNOR, M., AND WATSON, J. 2008. Computer generated holograms from three dimensional meshes using an analytic light transport model. *Appl. Opt.* 47, 10 (Apr), 1567–1574.
- BENTON, S. A., AND BOVE, V. M. 2007. *Holographic Imaging*. Wiley-Interscience.
- BENTON, S. 1991. Experiments in holographic video imaging. In *SPIE*, vol. 08, 247–267.
- COMISKEY, B., ALBERT, J. D., YOSHIKAWA, H., AND JACOBSON, J. 1998. An electrophoretic ink for all-printed reflective electronic displays. *Nature* 394, 6690, 253–255.
- COSSAIRT, O., NAYAR, S. K., AND RAMAMOORTHY, R. 2008. Light Field Transfer: Global Illumination Between Real and Synthetic Objects. *ACM SIGGRAPH* (Aug).

- DALLAS, W. J. 1980. Computer-generated holograms. *The Computer in Optical Research of Topics in Applied Physics* 41, 291–366.
- DEBITETTO, D. J. 1969. Holographic panoramic stereograms synthesized from white light recordings. *Appl. Opt.* 8, 1740–1741.
- DONG, Y., WANG, J., PELLACINI, F., TONG, X., AND GUO, B. 2010. Fabricating spatially-varying subsurface scattering. *ACM Trans. Graph.* 29, 4 (July), 62:1–62:10.
- FINCKH, M., DAMMERTZ, H., AND LENSCH, H. P. A. 2010. Geometry construction from caustic images. In *ECCV*, 464–477.
- FUCHS, M., RASKAR, R., SEIDEL, H.-P., AND LENSCH, H. P. A. 2008. Towards passive 6D reflectance field displays. *ACM Trans. Graph.* 27, 3.
- GERCHBERG, R. W., AND SAXTON, W. O. 1972. A practical algorithm for the determination of the phase from image and diffraction plane pictures. *Optik* 35, 237.
- GOODMAN, J. W. 1968. *Introduction to Fourier Optics*. McGraw-Hill Book Company.
- HAAŠAN, M., FUCHS, M., MATUSIK, W., PFISTER, H., AND RUSINKIEWICZ, S. 2010. Physical reproduction of materials with specified subsurface scattering. *ACM SIGGRAPH* 29, 3.
- HERMERSCHMIDT, A., OSTEN, S., KRÜGER, S., AND BLÜMEL, T. 2007. Wave front generation using a phase-only modulating liquid-crystal-based micro-display with hdtv resolution. In *International Congress on Optics and Optoelectronics*, International Society for Optics and Photonics, 65840E–65840E.
- 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.
- HIRSCH, M., IZADI, S., HOLTZMAN, H., AND RASKAR, R. 2012. 8D display: a relightable glasses-free 3d display. In *Proceedings of the 2012 ACM international conference on Interactive tabletops and surfaces*, ACM, 319–322.
- HIRSCH, M., IZADI, S., HOLTZMAN, H., AND RASKAR, R. 2013. 8D: interacting with a relightable glasses-free 3D display. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, 2209–2212.
- HORISAKI, R., AND TANIDA, J. 2013. Reflectance field display. *Opt. Express* 21, 9 (May), 11181–11186.
- HULLIN, M. B., LENSCH, H. P. A., RASKAR, R., SEIDEL, H.-P., AND IHRKE, I. 2011. Dynamic display of BRDFs. In *EUROGRAPHICS*, 475–483.
- HULLIN, M. B., IHRKE, I., HEIDRICH, W., WEYRICH, T., DAMBERG, G., AND FUCHS, M. 2013. Computational fabrication and display of material appearance. In *Eurographics State-of-the-Art Reports (STAR)*.
- IMEC, 2011. IMEC scientific report 2011. <http://www.imec.be/ScientificReport/SR2011/1414043.html>.
- KIM, S.-C., MOON, J.-W., LEE, D.-H., SON, K.-C., AND KIM, E.-S., 2005. Holographic full-color 3D display system using color-LCoS spatial light modulator.
- KISER, T., EIGENSATZ, M., NGUYEN, M. M., BOMPAS, P., AND PAULY, M. 2012. Architectural caustics controlling light with geometry. In *Advances in Architectural Geometry*.
- KLUG, M. A., HALLE, M. W., LUCENTE, M. E., AND PLESNIAK, W. J. 1993. Compact prototype one-step ultragram printer. *Proc. SPIE* 1914, 15–24.
- KOGELNIK, H. 1969. Coupled-wave theory for thick hologram gratings. *Bell System Technical Journal* 48, 2909.
- KOIKE, T., AND NAEMURA, T. 2008. BRDF display: interactive view dependent texture display using integral photography. In *Proceedings of the 2008 workshop on Immersive projection technologies/Emerging display technologies*.
- LAN, Y., DONG, Y., PELLACINI, F., AND TONG, X. 2013. Bi-scale appearance fabrication. *ACM Trans. Graph.* 32, 4.
- LANMAN, D., WETZSTEIN, G., HIRSCH, M., HEIDRICH, W., AND RASKAR, R. 2011. Polarization fields: dynamic light field display using multi-layer LCDs. *ACM Trans. Graph.* 30, 6, 186.
- LEVIN, A., GLASNER, D., XIONG, Y., DURAND, F., FREEMAN, B., MATUSIK, W., AND ZICKLER, T. 2013. Fabricating BRDFs at high spatial resolution using wave optics. *ACM SIGGRAPH*.
- LUCENTE, M. E. 1993. Interactive computation of holograms using a look-up table. *J. of Electronic Imaging* 2, 1, 28–34.
- LUCENTE, M. 1994. *Diffraction-specific Fringe Computation for Electro-holography*. PhD thesis. AAI0575566.
- MALZBENDER, T., SAMADANI, R., SCHER, S., CRUME, A., DUNN, D., AND DAVIS, J. 2012. Printing reflectance functions. *ACM Trans. Graph.* 31, 3.
- MANN, S. 1995. Recording Lightspace so Shadows and Highlights Vary with Varying Viewing Illumination. 2538–2540.
- MATUSIK, W., AJDIN, B., GU, J., LAWRENCE, J., LENSCH, H. P., PELLACINI, F., AND RUSINKIEWICZ, S. 2009. Printing spatially-varying reflectance. *ACM SIGGRAPH Asia* 28, 5.
- NAYAR, S., BELHUMEUR, P., AND BOULT, T. 2004. Lighting Sensitive Display. *ACM Trans. on Graphics* 23, 4, 963–979.
- NG, R., LEVOY, M., BREDIF, M., DUVAL, G., HOROWITZ, M., AND HANRAHAN, P. 2005. Light field photography with a handheld plenoptic camera. *Stanford U. Tech Rep CSTR 2005-02*.
- OCHIAI, Y., OYAMA, A., AND TOYOSHIMA, K. 2012. A colloidal display: membrane screen that combines transparency, BRDF and 3D volume. In *ACM SIGGRAPH Emerging Technologies*.
- OCHIAI, Y., OYAMA, A., HOSHI, T., AND REKIMOTO, J. 2013. Reflective, deformable, colloidal display: a waterfall-based colloidal membrane using focused ultrasonic waves. In *SIGGRAPH Posters*, 49.
- PAPAS, M., JAROSZ, W., JAKOB, W., RUSINKIEWICZ, S., MATUSIK, W., AND WEYRICH, T. 2011. Goal-based caustics. *Computer Graphics Forum (Proc. Eurographics)* 30, 2 (Apr.).
- PATOW, G., AND PUEYO, X. 2005. A survey of inverse surface design from light transport behavior specification. *Comput. Graph. Forum* 24, 4, 773–789.
- PATOW, G., PUEYO, X., AND VINACUA, A. 2007. User-guided inverse reflector design. *Comput. Graph.* 31, 3 (June), 501–515.
- REDMAN, J. 1968. The three-dimensional reconstruction of people and outdoor scenes using holographic multiplexing. *Proceedings of SPIE Seminar-in Depth on Holography* 15, 117122.
- RUSINKIEWICZ, S. 1998. A new change of variables for efficient BRDF representation. In *EGSR*.
- SMALLEY, D. E., SMITHWICK, Q. Y. J., BOVE, V. M., BARABAS, J., AND JOLLY, S. 2013. Anisotropic leaky-mode modulator for holographic video displays. *Nature*, 7454, 313317.
- ST-HILAIRE, P., BENTON, S. A., LUCENTE, M. E., AND HUBEL, P. M., 1992. Color images with the MIT holographic video display.

- TORRANCE, K. E., AND SPARROW, E. M. 1967. Theory for off-specular reflection from roughened surfaces. *JOSA* 57, 9.
- TRICOLES, G. 1987. Computer generated holograms: an historical review. *Appl. Opt.* 26, 20 (Oct), 4351–4357.
- WETZSTEIN, G., LANMAN, D., HIRSCH, M., AND RASKAR, R. 2012. Tensor displays: compressive light field synthesis using multilayer displays with directional backlighting. *ACM Trans. Graph.* 31, 4, 80.
- WEYRICH, T., DENG, J., BARNES, C., RUSINKIEWICZ, S., AND FINKELSTEIN, A. 2007. Digital Bas-Relief from 3D Scenes. *ACM SIGGRAPH* 26, 3 (Aug.).
- WEYRICH, T., PEERS, P., MATUSIK, W., AND RUSINKIEWICZ, S. 2009. Fabricating microgeometry for custom surface reflectance. *ACM SIGGRAPH* 28, 3 (Aug.).
- YAROSLAVSKY, L. 2004. *Digital Holography and Digital Image Processing*. Kluwer Academic Publishers.
- ZIEGLER, R., BUCHELI, S., AHRENBURG, L., MAGNOR, M., AND GROSS, M. 2007. A bidirectional light field-hologram transform. In *Computer Graphics Forum*, vol. 26, 435–446.

A Piecewise Constant Surface Effects

As the SLM consists of units of size $\Delta_o = 8\mu m$, the phase modulation it applies is roughly piecewise constant over units of this size. Our goal here is to derive the properties of piecewise constant phase modulations. We note that the piecewise model is only an approximation since in practice there is cross talk between adjacent cells, nevertheless, the analysis helps understand our display.

Let $R(\mathbf{h}_x, \lambda; \phi)$ denote the reflectance at angle \mathbf{h}_x and wavelength λ under the phase modulation ϕ . The reflectance in the display range $-\Omega_a^\lambda \leq 2\mathbf{h}_x \leq \Omega_a^\lambda$ is determined by ϕ as explained in Sec. 3.1. We show that outside this range one observes replicas with reduced contrast. Let $\tilde{\mathbf{h}}_x$ denote the remainder mapping of $2\mathbf{h}_x$ into the range $[-\Omega_a^\lambda, \Omega_a^\lambda]$, $2\tilde{\mathbf{h}}_x \equiv 2\mathbf{h}_x \pmod{\Omega_a^\lambda}$, then

$$R(\mathbf{h}_x, \lambda; \phi) = \text{sinc}\left(\frac{\mathbf{h}_x}{\Omega_a^\lambda}\right) \cdot R(\tilde{\mathbf{h}}_x, \lambda; \phi). \quad (20)$$

To see this, consider a surface represented by discrete impulses at spacing Δ_o , the Fourier transform would consist of replicas of the transform in the $[-\Omega_a^\lambda, \Omega_a^\lambda]$ range, and $R(\mathbf{h}_x, \lambda; \phi) = R(\tilde{\mathbf{h}}_x, \lambda; \phi)$, as illustrated in Fig. 14. The piecewise constant surface modulation is the impulse train convolved with a rect of width Δ_o . Therefore, by the convolution theorem its Fourier transform is the Fourier transform of the impulse train multiplied by the Fourier transform of that rect which is a sinc of width Ω_a^λ as in Eq. (20).

To summarize, for a piecewise constant surface, outside the controlled angular range of $[-\Omega_a^\lambda, \Omega_a^\lambda]$ we see replicas of the main reflectance. This effect is reminiscent of parallax barrier multiscopic displays. The intensity of these replicas is reduced by a sinc function. In practice, cross talk between SLM dots further smooths the phase modulation, reducing the contrast in side lobes.

B Binary Phase Limitations

Due to fabrication constraints Levin et al. [2013] restricted their height surface (and phase modulation $\phi_\lambda(x)$) to be binary. We show that reflectances produced with binary phase modulation as well as by amplitude modulation SLMs suffer from two main restrictions: i) reflectances must be symmetric and ii) include a specular spike.

To see this note that with amplitude modulation SLMs the function $\alpha_\lambda(x)$ of Eq. (1) takes real (and positive) values. The Fourier

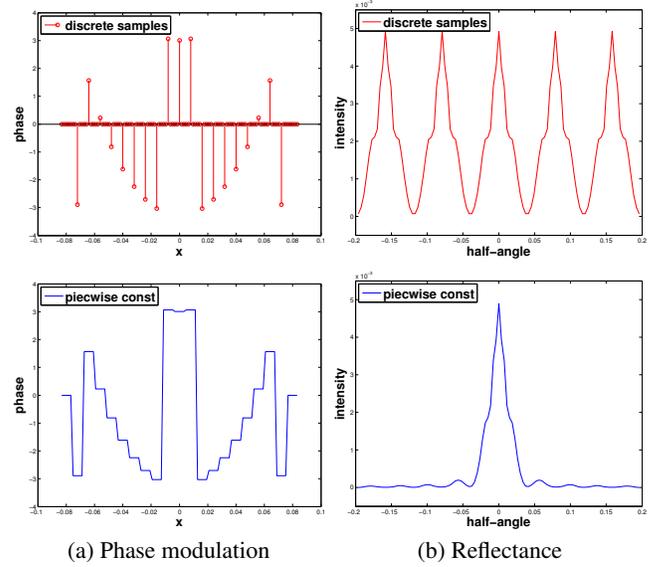


Figure 14: Piecewise constant phase. A piecewise constant phase modulation is shown in the bottom panel of (a) its reflectance can be understood by thinking of it as the convolution of a discretely sampled phase modulation, as shown on the top, with a rect function. The Fourier transform of the impulse train is periodic as shown in the top panel of (b). By the convolution theorem the Fourier transform of the piecewise constant modulation is the Fourier transform of the impulse train multiplied by the Fourier transform of a rect which is a sinc, the resulting reflectance is shown in the bottom panel of (b).

transform of a real signal is known to be symmetric, and as the reflectance $|\mathfrak{A}_\lambda|^2$ is derived in Eqs. (2) and (3) as the Fourier power spectrum of $\alpha_\lambda(x)$, it is limited to symmetric functions as well.

We show that binary phase modulation is equivalent to real valued $\alpha_\lambda(x)$. Consider a binary valued $\alpha_\lambda(x) \in \{\zeta_1, \zeta_2\}$. There exist two complex scalars $\alpha, \beta \in \mathbb{C}$ s.t. $\tilde{\alpha}_\lambda = \alpha\alpha_\lambda + \beta$ and $\tilde{\alpha}_\lambda(x) \in \{-1, 1\}$. Let us denote by $\mathfrak{A}_\lambda, \tilde{\mathfrak{A}}_\lambda$ the Fourier transforms of $\alpha_\lambda, \tilde{\alpha}_\lambda$ respectively. Since $\tilde{\alpha}_\lambda$ is real valued, its Fourier transform $\tilde{\mathfrak{A}}_\lambda$ is symmetric, namely $\tilde{\mathfrak{A}}_\lambda(\omega) = \tilde{\mathfrak{A}}_\lambda(-\omega)^*$. We show that this implies that $|\mathfrak{A}_\lambda|^2$ must be symmetric as well.

Let $c(x) = \beta$ be a constant function in the primal domain. It's Fourier transform $\mathcal{C}(\omega)$ is an impulse (delta) at $\omega = 0$. Since $\mathfrak{A}_\lambda(\omega) = 1/\alpha(\tilde{\mathfrak{A}}_\lambda(\omega) + \mathcal{C}(\omega))$, we get that $|\mathfrak{A}_\lambda|^2$ must be symmetric as well. Additionally, unless $\beta = 0$, \mathfrak{A}_λ has a non zero DC term, and thus the reflectance has a specular spike at the mirror direction. By choosing the two heights appropriately it is possible to have $\beta = 0$ and get rid of the specular spike at one wavelength, but this can't be done for all wavelengths simultaneously. In contrast, our SLM allows us to apply a dense range of distinct phase modulations. This enables us to design non-symmetric reflectances without a specular spike as we demonstrate in Sec. 5.1.