Passive Light and Viewpoint Sensitive Display of 3D Content

Anat Levin,   Haggai Maron,   Michal Yarom
Weizmann Institute of Science

Abstract

We present a 3D light-sensitive display. The display is capable of presenting simple opaque 3D surfaces without self occlusions, while reproducing both viewpoint-sensitive depth parallax and illumination-sensitive variations such as shadows and highlights. Our display is passive in the sense that it does not rely on illumination sensors and on-the-fly rendering of the image content. Rather, it consists of optical elements that produce light transport paths approximating those present in the real scene.

Our display uses two layers of Spatial Light Modulators (SLMs) whose micron-sized elements allow us to digitally simulate thin optical surfaces with flexible shapes. We derive a simple content creation algorithm utilizing geometric optics tools to design optical surfaces that can mimic the ray transfer of target virtual 3D scenes. We demonstrate a possible implementation of a small prototype, and present a number of simple virtual 3D scenes.

1. Introduction

Despite significant recent advances in multiscopic display technology [2, 29, 50, 12], the majority of these displays are still limited in an important aspect: they can only display a scene captured under a fixed-illumination condition, and they do not naturally react to illumination in the viewer's environment. However, changing the environment lighting can vary appearance substantially by shifting highlights and shadows. These effects are extremely important for shape and material perception, and thus for realism.

This paper makes a step towards a *3D light-sensitive display*, capable of presenting a viewpoint-sensitive depth content as well as spatially varying material reflectance properties that will accurately react to the interaction between environment lighting and scene geometry.

A popular approach taken in [39, 34, 8, 22, 20, 21, 23] is to build an *active* light-sensitive display in which a camera is used to capture the illumination in the viewing environment, and the displayed scene is re-rendered to match the sensed illumination. In contrast, this paper proposes a *passive* display device. It does not have its own source of radiant power and does not rely on illumination sensors. It operates using an optical setup controlling how environment lighting is reflected. The passive approach has two main advantages. First, since computationally expensive rendering engines are not required, 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.

Some passive devices have been proposed in recent years, however they all consider restricted subsets of light-sensitive effects. For example, some displays consider only 2D surfaces. Other displays consider 3D targets, but assume that the viewer position is fixed. An alternative model considers only distant light sources, which produce spatially-invariant illumination on the display area. The family of previous models is summarized in Table 1 and reviewed in Sec. 1.1.1.

In this paper we attempt to display 3D objects while addressing general viewpoint and lighting, which can vary both spatially and angularly. To simplify the problem, we restrict our attention to the display of smooth and opaque surfaces without self occlusions. As we explain below, while the general space of all incoming and outgoing directions accounts for an 8D lightfield space, the type of scenes considered in this paper allows us to reduce the problem into a 5D non-linear sub-manifold.

Our display follows [16], and uses Spatial Light Modulation (SLM) technology. SLMs are arrays of micron-sized units offering dynamic computer control over surface micro-structures, an effect that can be used to control the reflected light. However, while the single SLM layer used in [16] is inherently limited to 4D subspaces of the 8D lightfield, our design uses a cascade of two such SLMs, an arrangement which can significantly expand the set of light transport models we can support.

Due to their micron-sized units, SLM appearance should be analyzed using wave optics models. However, designing the SLM content using wave optics directly (e.g. Ye *et al*. [53]), results in cost functions involving a combinatorial number of local minima that are hard to optimize and analyze. To address this difficulty, we suggest design algorithms that are based on geometric optics tools, providing a quick approximation to the wave model. These algorithms are easier to optimize, and improve our understanding of the problem.

Our first contribution is to derive a simple closed-form solution for the display assignment in the case of thin bas-relief surfaces. Given a target virtual surface, this solution uses the first display layer as a lens, which focuses on the second layer the rays that would otherwise have focused on the target surface. The second layer can then directly copy surface details from a scaled version of the target surface.

As a second contribution, we suggest an approximation strategy extending the bas-relief case to more general surfaces with varying depth, whose optimization involves simple quadratic programming. We demonstrate the strategy and its limitation through simulations.

Our third contribution is the implementation of a small programable display prototype. It can present simple smooth opaque 3D surfaces, while reproducing both viewpoint-parallax and illumination-dependent appearance effects such as varying highlights and attached shadows. The initial prototype has a high spatial resolution. However, due to the limitation of existing SLMs, it has a small size and a very limited angular range. However, SLM technology is rapidly advancing and may soon relax these limitations as any improvement in SLM resolution directly translates into a wider angular range.

1.1. Related work

An excellent survey of the state-of-the-art in computational fabrication and display of light-sensitive appearance was recently published by Hullin *et al.* [25]. Below we discuss some technologies that are most related to our work.

1.1.1 Passive displays

Passive displays are desirable because they react instantaneously to environment lighting, and because of their modest power consumption. We consider here both hardware fabricated objects, and real digital displays.

In its most general form, a light-sensitive display should allow a mapping $L(x_{in}, \alpha_{in}, x_{out}, \alpha_{out})$ from any incident ray (x_{in}, α_{in}) to any outgoing ray (x_{out}, α_{out}). Since each spatial coordinate x_{in}, x_{out} as well as each angular coordinate $\alpha_{in}, \alpha_{out}$ can vary in 2D, the variability of all these aspects accounts for an 8D display. Previous passive light-sensitive displays simplify the task by considering various lower-dimensional subsets of the problem. We discuss the leading strategies below, and summarize them in Table 1.

2D display: The first family of light-sensitive surfaces were designed to generate goal-based caustics [13, 40, 32, 45, 54, 42]. These are refractive surfaces whose structure is designed to generate special 2D caustic images under proper illumination. An alternative line of research studied

Dim.	Model	Reference
2D	$L(\cdot, \alpha_{in}^o, x_{out}, \cdot)$	Caustics
	$L(x_{in}^o, x_{out}^o, \alpha_{in} + \alpha_{out})$	Dot BRDF
4D	$L(x_{in}, x_{out} = x_{in}, \alpha_{in} + \alpha_{out})$	Planar BRDF display
	$L(\cdot, \cdot, x_{out}, \alpha_{out})$	3D display
	$L(x_{in}, \alpha_{in}, x_{out}^o, \alpha_{out}^o)$	[37, 16]
6D	$L(x_{in}, \alpha_{in}, x_{out} = x_{in}, \alpha_{out})$	[14, 53]
	$L(\cdot, \alpha_{in}, x_{out}, \alpha_{out})$	?
8D	$L(x_{in}, \alpha_{in}, x_{out}, \alpha_{out})$	?
5D	$L(x_{in}, \alpha_{in}, x_{out}, \alpha_{out})$	This work
	Non-lin rest. on scene content	

Table 1: *A summary of strategies for reducing the dimensionality of passive light-sensitive displays. An entry marked by " · " indicates that this dimension is assumed uniform or that the output is invariant to it. An entry with superscript ^{o} implies that its value is fixed and known when content is designed. The "?" mark indicates that to our best knowledge, this configuration was not yet demonstrated.*

BRDF display [51, 26], but was restricted to demonstrating a single BRDF dot rather than spatially-varying BRDF. While in its most general form a BRDF function of the form $R(\alpha_{in}, \alpha_{out})$ has 4D of freedom (2D lighting × 2D viewpoint), most BRDF fabrication-and-display strategies consider BRDFs that are mostly a function of the half-vector $\alpha_{in} + \alpha_{out}$ rather than an independent function of both $\alpha_{in}, \alpha_{out}$, restricting its dimensionality to 2D.

4D displays: Another class of special-appearance light-sensitive surfaces [11, 18, 38, 41, 37, 35, 3] are those targeting 4D subsets of the lightfield. For example, [37] suggested a 4D light-sensitive display $L(x_{in}, \alpha_{in})$ that assumes a fixed viewpoint and adjusts the appearance of every point on the 2D display to any 2D illumination direction. Alternatively, BRDF fabrication and displays [35, 36, 16] present a 2D image, such that in each spatial pixel one can control the BRDF that varies as a function of the half-vector, effectively controlling a 4D function of the form $L(x, \alpha_{in} + \alpha_{out})$.

6D displays: A more expressive type of light-sensitive displays consider 6D lightfields of the form $L(x, \alpha_{in}, \alpha_{out})$, without restricting them to vary only as a function of the half vector $\alpha_{in} + \alpha_{out}$. Note that this model does not trivially support the display of 3D scenes since if the scene contains depth variations, then not all points are located on a single reference plane. Thus a ray hitting a reference plane at point x_{in} may continue to propagate until it hits a surface point at a different depth, reflects and emerges out hitting the reference plane at a different point $x_{out} \neq x_{in}$. While in theory one can use 6D lightfields to display 3D scenes under distant, spatially-uniform illumination at the price of some energy loss, we are not aware of any previous attempt

to implement this idea. In contrast, our proposed display supports general illumination, which can vary spatially and angularly. This freedom is achieved at the price of simplifying assumptions on the scene content.

Fuchs *et al.* [14] fabricated a 6D light-sensitive lightfield, using an elaborate arrangement of lenslets. This flattens the 6D light transport tensor onto a plane, where it is modulated by an attenuating transparent layer. Roughly speaking, one set of lenslets separates the incoming illumination directions into distinct spatial positions, where they are modulated by a second set of lenslets and an attenuation layer to produce the desired reflected field. This approach suffers from two main limitations. First, due to the high dimension of the problem, the flattened lightfield tensor suffers from inherent resolution limits, and in practice the prototype was limited to a quantized set of $7 \times 7 \times 6 \times 5$ directions. A second issue is that the design assumes that the illumination source arrives from the back of the display while the viewer is located at the other side. Generalizing this arrangement to a more natural setup in which lighting and viewer are located on the same side of the display is non trivial, since the light would have to travel back through the lenslet layers.

In [53] a surface with 6D variation was fabricated by cascading multiple transparencies, each of them printed with micron-sized features. Using wave optics tools, they predicted the propagation of light through the layers, and defined an optimization problem to find a micro pattern whose diffracting wave matches a target appearance. While this approach considers light transport using accurate wave models, the resulting optimization problem is highly non-convex, thus hard to optimize and analyze.

1.1.2 Holography

Holography is based on recording the fine interference pattern between the wave scattered from a 3D object and a mutually-coherent reference plane wave [43, 10, 4, 52]. When the recorded fringe pattern is viewed under proper illumination, the observer perceives the object's 3D parallax. Computer-generated holograms induce the experience of virtual 3D objects that need not exist in the real world [9, 49, 1, 52], and programmable holographic displays based on these ideas have been implemented using SLM technology [47, 31, 46].

Volume holography: Most existing holograms fall under the *thin* holograms category, in the sense that the fringe pattern is recorded on a relatively thin plane, as achieved by a single planar SLM. These holograms are essentially equivalent to 4D lightfields, since they can only record the scattering wave under a single illumination direction. When illuminated by a plane wave whose angle is different from the recording angle, a rotated version of the scene is reconstructed. In contrast, in *volume* holograms [33, 17, 6], the
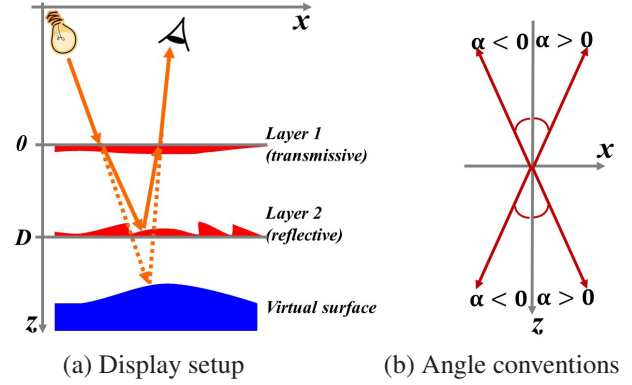

Figure 1: *Setup: (a) Our display consists of two layers at fixed depths on which we can realize programable thin optical surfaces. These surfaces are designed such that they direct rays in the same way as a target 3D virtual surface. Dashed rays mark the expected path of rays to and from the virtual surface, while solid ones denote the path via the display layers. While the path inside the display is different than the virtual path, the ray via the display and the ray from the virtual surface coincide upon exiting the display. (b) The sign convention for angles in our system.*

fringe patterns are recorded over a volume whose thickness is a few hundred wavelengths. The result is that the object is reconstructed only if illuminated by the exact recording angle and wavelength. Expanding this idea, one can store on the hologram multiple scenes from multiple illumination directions and achieve illumination-sensitive image reconstruction. Stratified volume holograms [28, 7, 24] attempt to approximate the volume principle using discrete layers. Volume holograms are usually manufactured by exposing a photosensitive glass. Recently, computational optimization approaches that can find micro-scale patterns to print on a stratified hologram have been proposed [5, 30, 15, 53]. Volumetric holograms are very relevant to our 8D display goal. However, while the design and analysis of volume holograms usually involves complex electromagnetic equations, we propose here simpler geometric optics analysis, which is directly tailored to our goals, and suggest that two planar layers already provide some of the more useful volumetric control flexibilities.

2. Our Model

Setup: Figure 1(a) illustrates the display geometry, consisting of two planar layers placed at a distance D from each other, on which we conceptually position relatively thin optical elements whose shape will be denoted here as $z_1^d(x), z_2^d(x)$. The front layer is transmissive (e.g., a glass) and the back layer reflective (a mirror). Every incoming ray is refracted through the front layer, crosses the distance to the back layer, is reflected by the back layer, propagates back, and is refracted through the front layer toward the

viewer. Our goal is to find a shape assignment for the layers, so as to mimic the appearance of a target virtual 3D surface. That is, the display layers should be designed such that any ray hitting the front layer of the display will emerge from it at the exact same position and orientation as it would have emerged had it traveled straight to the virtual surface and been reflected from it.

Our assumptions: We will design our display using the following simplifying assumptions on the scene and setup. (i) The display is illuminated and viewed only within a bounded angular range $\alpha_{in}, \alpha_{out} \in [-\Omega_M, \Omega_M]$. (ii) Our 3D scene contains opaque surfaces without rapid depth variations. (iii) A ray at angle $\alpha \in [-\Omega_M, \Omega_M]$ does not intersect the surface at more than one point. That is, no self occlusion are present, and hence cast shadows do not appear. (iv) Each surface point can have an independent BRDF. As is the case with most physical BRDFs when no subscattering is present, the BRDFs are only a function of the half vector between the incoming and outgoing rays at the intersection point. That is, the BRDF is a function of $\alpha_{in} + \alpha_{out}$, rather than an independent function of both $(\alpha_{in}, \alpha_{out})$. (v) Since our angles of interest are small we use the paraxial approximation to ray optics. In the supplementary file we show empirically that this approximation is reasonable for incident and viewing angles up to about 20^o.

Dimensionality: As explained in Sec. 1.1.1, general light-sensitive displays that can support arbitrary transformations between any incoming ray to any outgoing ray essentially lie in 8D. Under the assumptions mentioned above, the light-transport functions of our scenes lie in a non-linear 5D submanifold: for each 2D point on the display plane we have one degree of freedom in selecting the surface height and two degrees of freedom in selecting a BRDF.

Yet, this 5D display offers new capabilities that were not demonstrated by previous 6D displays [14, 53] discussed in Sec. 1.1.1, which mostly demonstrated spatially and angularly varying appearance effects on a 2D display plane. While this was not yet demonstrated in previous work, a 6D display can mimic a light-sensitive appearance of a 3D scene only if the illumination source is assumed to be distant enough so that its output is spatially uniform over the display area. In contrast, our display supports general illumination that can vary spatially and angularly.

SLMs and their implications: As will be discussed in Sec. 5, in practice the display layers in our setup (Fig. 1) are realized digitally using two reflective Spatial Light Modulators (SLMs). Since the layers consist of very small elements, to accurately analyze how light propagates through such a display, we must account for the wave nature of light. Unfortunately, a direct design using this waves model quickly leads to a hard non-linear optimization. To make the problem tractable, we analyze the surfaces using geometric optics tools. This choice is motivated by the empirical observation of Glasner *et al.* [16] that the geometric optics model provides reasonable approximation when the surfaces are sufficiently smooth, and can provide a good initialization for a more accurate wave optics optimization.

While we defer the description of the SLM operation to Sec. 5, it is worth mentioning here two main limitations of this implementation: (i) the optical surfaces that we can implement are rather thin, and (ii) the maximal surface slope they can express is bounded $|\partial z_k^d(x)/\partial x| \leq S_M$. As will be explained in Sec. 5, this constraint also implies that the range of angles by which the display can be illuminated and viewed is limited.

Notations: For ease of notation, we will perform most of our derivation in flatland. As mentioned above, we denote the optical surfaces programmed on the display layers as $z_k^d(x)$ where $k = 1, 2$ is the layer index, and the virtual target surface height function as $z^v(x)$. For optical derivation we often rely on the surface slope which we denote here as:

$$s_k^d(x) = \frac{\partial z_k^d(x)}{\partial x}, \quad s^v(x) = \frac{\partial z^v(x)}{\partial x}. \tag{1}$$

We assume w.l.o.g. that the optical axis of the system coincides with the z axis, and that the front layer of the display is positioned at the $z = 0$ plane, see setup and coordinate system in Fig. 1. We denote a light ray by a 3×1 vector $\mathbf{r} = (x, \alpha, 1)^T$ where x is the spatial distance from the optical axis at the ray origin, and α denotes its angle with respect to the optical axis. When an input ray $\mathbf{r}_{in}$ propagates through a scene it is transferred to an output ray $\mathbf{r}_{out} = A(\mathbf{r}_{in})$, where the transformation A is determined by the scene geometry. In this paper we will denote by A_v, A_d the transformations through the target virtual scene and through the display layers correspondingly, where the input and output rays $\mathbf{r}_{in}, \mathbf{r}_{out}$ are recorded at their intersection with the $z = 0$ plane. We will also be interested in the forward and backward transformations $A_v^f, A_d^f, A_v^b, A_d^b$, where A_v^f, A_d^f map an input ray $\mathbf{r}_{in}$ from the $z = 0$ plane to its intersection with the virtual surface or the back layer of the display. Similarly A_v^b, A_d^b refer to the backward transformation from the back surface to the $z = 0$ plane.

Display goals: Our goal is to find two thin optical surfaces $z_1(x), z_2(x)$ to be presented on the display layers, so as to mimic the ray transfer of the target virtual surface. For uniform surfaces, this reduces to the request that $A_d = A_v$, i.e. every incoming ray $\mathbf{r}_{in}$ hitting the front layer should exit the display in the same position and orientation as the ray $\mathbf{r}_{out}$ that would have emerged if the incoming ray had actually been reflected by the target virtual surface.
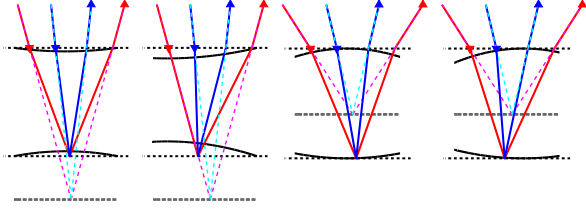
Figure 2: *Optically-equivalent pairs. For each virtual mirror (dashed gray line), we find a pair of display surfaces (solid black curves) at fixed depths, producing an equivalent ray transformation. Dashed rays mark reflection at the target virtual surface and solid rays refraction and reflection via the display layers. Despite the different paths inside the display, the output position and direction of the display rays coincide with the rays reflecting at the virtual surface. Moreover, two rays that meet on the target surface also meet on the back layer of the display. The one-to-one mapping between points on the virtual surface and points on the back layer of the display allows us to copy the texture of the target surface. For each of two virtual surfaces, we present two types of solutions, illustrating two different values of t_1^d in Eq. (5).*

If the target surface contains a spatially-varying texture of intensity or BRDF, ensuring that $A_d = A_v$ is insufficient. For instance, if the surface contains spatially-varying albedo, we do not only have to ensure that rays emerge at the right direction, but also that their intensity is properly attenuated. In order to consistently attenuate all rays that reach a single point on the display, these rays should also meet at a single point on the virtual surface. That is, we need a unique mapping from points on the target surface to points on the back layer of the display. This situation is illustrated in Figure 2. In the next sections we show that this is equivalent to making the front layer focus on the back layer the rays that would otherwise have focused on the target virtual surface, i.e. the front layer should act as a lens imaging the surface $z^v(x)$ on the back display layer.

3. Bas-Relief Surfaces

We start the derivation of the display assignment by considering a specific case for which we can derive a simple closed-form solution. This case will also help us build intuition about the display operation. We consider here bas-relief surfaces, or essentially "normal maps". These are relatively planar surfaces, which we will describe as

$$z^v(x) = \zeta^v(x) + c^v, \qquad (2)$$

where c^v is the plane depth and ζ^v are small-bounded depth variations obeying $|\zeta^v(x)| \leq \epsilon$. Despite the modest depth variations, $\zeta^v(x)$ can have highly-varying gradient orientations, leading to a non-uniform appearance. As an example, consider the thin surface in Figure 3. We note that the plane depth c^v can be arbitrary and may not coincide with the

depth of any of the display layers.

We start by assuming that the surface to be represented is a pure mirror. In Sec. 3.2, we show that BRDF variations can be added as a small modification to the same formula, by considering micro-features on the target surfaces.

3.1. Bas-relief: use top layer as a focusing lens

For thin bas-relief surfaces, we have a simple closed-form solution for the display layers. We use the front layer as a lens that images the target plane on the back layer of the display. That is, the front layer focuses on the back layer the rays that would otherwise have focused on the target virtual surface. The back layer can then present a copy of the target surface, scaled to match the lens magnification. To this copy of the target, we need to add a curved mirror that corrects the angle at which rays are focused.

Claim 1. *Given a thin bas-relief surface $z^v(x) = \zeta^v(x) + c^v$ at depth c^v, we can reproduce its ray transfer using two display layers at distance D apart. The front layer is a parabolic lens*

$$z_1^d(x) = -\frac{1}{2}p_1^d \cdot x^2 - t_1^d \cdot x + c_1^d. \qquad (3)$$

and the back layer a parabolic mirror plus a scaled copy of the target

$$
\begin{aligned}
z_2^d(x) &= -\frac{1}{2}p_2^d \cdot x^2 - t_2^d \cdot x + z^v\left(\frac{c^v}{D}x + c^v t_1^d\right) - c^v \\
&= -\frac{1}{2}p_2^d \cdot x^2 - t_2^d \cdot x + \zeta^v\left(\frac{c^v}{D}x + c^v t_1^d\right), \quad (4)
\end{aligned}
$$

where t_1^d is a free parameter and

$$p_1^d = \frac{c^v - D}{Dc^v}, \quad p_2^d = \frac{D - c^v}{D^2}, \quad t_2^d = -\frac{c^v}{D}t_1^d = \frac{p_2^d}{p_1^d}t_1^d. \qquad (5)$$

The proof is provided in supplementary Appendix A.

Figure 2 illustrates the task of the front layer on a simple planar scene. This layer focuses on the back layer any pair of rays that would otherwise meet on the virtual plane. Similarly, Figure 3 illustrates the display layers for a bas-relief surface. The first layer is a parabolic lens and the second one a quadratic mirror plus a scaled copy of the target. Note that Eq. (5) implies that if the virtual layer lies behind the 2nd display layer, i.e. $c^v > D$, we get $p_1^d > 0$. That is, the front layer acts as a positive lens focusing rays. The situation is inverted if the virtual layer lies before the back layer $c^v < D$. Also, when the front layer acts as a positive lens, it shrinks the distances between spatial points, hence the copy of the bunny is scaled down. When one uses a negative lens at the front layer, the bunny bas-relief is expanded.
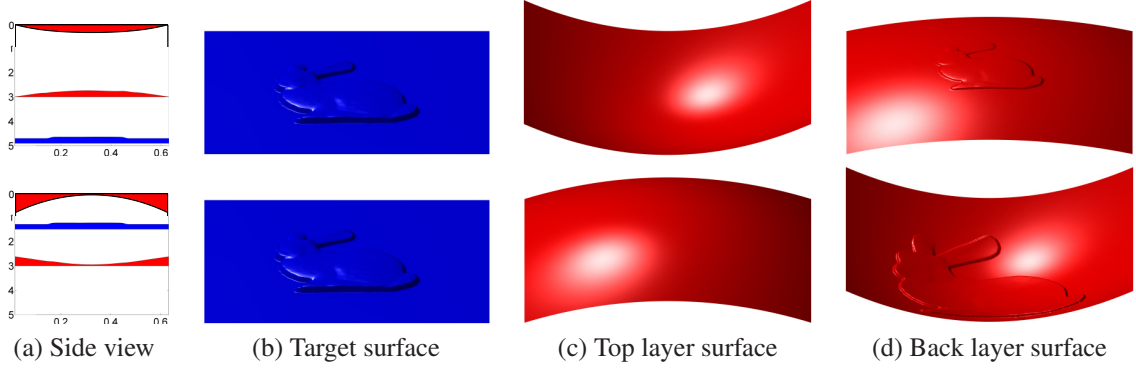
(a) Side view   (b) Target surface   (c) Top layer surface   (d) Back layer surface

Figure 3: *A target bas-relief scene containing a thin bunny surface, where the target depth is below the second display layer (first row), or above it (second row). As suggested by Claim 1, for a surface below the 2nd layer the first layer is a positive lens and the second one is a negative mirror with a copy of the bunny. The situation is inverted for a surface above the 2nd layer. Since a positive/negative lens on the top surface shrinks/expands spatial distances, the copy of the bunny on the back layer is smaller/wider.*

3.2. BRDF

xx The above derivation considered the display of a virtual 3D mirror surface. To simulate surfaces with richer BRDFs, which reflect on a wider range of directions, we should add surface roughness. For that we follow the ideas of the micro-facet theory [48], explaining BRDF formation by suggesting that if we zoom in closely on a seemingly smooth surface we will see rough micro-features whose size is lower than the resolution observable by the viewer. A collimated set of rays entering the area of pixel x at angle α_{in} will emerge over a range of angles $\{\alpha_{out}\}$, since each ray hits a different surface normal (see Fig. 4(a)). The exact angular spread depends on the distribution of normal orientations over the dot area. This strategy was used by [44, 16], and we follow it here. To add BRDFs to our target surface, we select a unit size Δ_d that is below the resolvable viewing resolution. We add to the smooth 3D shape small bumps of width Δ_d. The distribution of normals on the bumpy surface determines the BRDF.

As a simple example, consider a BRDF that is uniform over an angular range β around the direction of reflection and reflects zero energy outside this range. Denoting the local surface slope by s we can express this reflectance as:

$$R(\alpha_{in}, \alpha_{out}) = \begin{cases} \frac{1}{2\beta} & |\alpha_{out} - (\alpha_{in} + 2s)| \leq \beta \\ 0 & |\alpha_{out} - (\alpha_{in} + 2s)| > \beta \end{cases} \quad (6)$$

To mimic this BRDF we can add to the surface a bump defined by a parabola:

$$\zeta(x - x_o) = \frac{\beta}{\Delta_d}(x - x_o)^2, \quad (7)$$

whose slope over the range $[-\Delta_d/2, \Delta_d/2]$ covers exactly the angular range $[-\beta, \beta]$. Note that since we work here in the paraxial regime, the angle of reflection is approximated as $\alpha_{out} = 2s + \alpha_{in}$. Also we can ignore foreshortening as
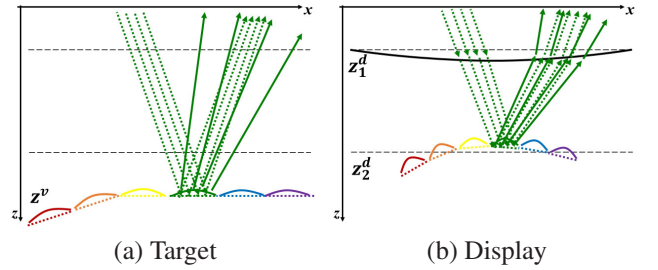

(a) Target   (b) Display

Figure 4: *(a) To add BRDF to our target surfaces, we follow the micro-facets model, adding to the seemingly smooth surface small micro surfaces whose size is below the resolvable spatial resolution. The angular variation of the micro-surface causes incoming rays to spread over a wider angular range. In the figure, dashed rays represent rays as directed by a smooth mirror surface, and solid rays represent ray casting when the micro-bump features are added. (b) To generate the same angular spread with our display, we add a scaled version of the target micro-bumps to the 2nd layer of our display.*
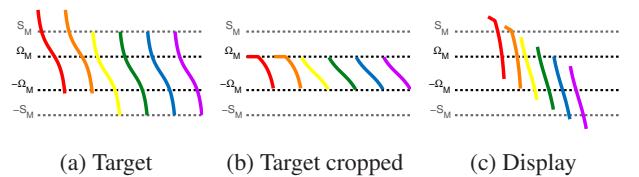

(a) Target   (b) Target cropped   (c) Display

Figure 5: *Slopes of the surfaces in Fig. 4. (a) Slopes of the original target in Fig. 4(a). (b) Cropped slopes: the slopes required to generate the actual BRDF can be large. However, given a bounded lighting and viewing range $[-\Omega_M, \Omega_M]$, only slopes within a fixed range are reflecting light, thus the bump slopes can be cropped without affecting appearance. (c) Slopes on the display's back layer (Fig. 4(b)), equal to the sum of the target slope plus the slope of the parabolic surface on the 2nd display layer. Since the slope of a parabola is linear, we effectively add a linear ramp to the target slopes. The increased slopes are problematic when the display slope bound S_M is not sufficiently wide.*

for our small angles we can assume $cos(\alpha) \approx 1$.

In our display we follow a similar strategy and add small bumps to the back layer of the display, as in Fig. 4(b). Once BRDF is added to the target, adjusting it to the display is straightforward. Using the 2nd layer construction formula in Claim 1, we can copy to the back layer a scaled version of the target surface bumps $\zeta^v(x)$, as in Eq. (4).

With this micro-facet model the BRDFs at each spatial point only have 2D of freedom, as reflection at each facet point is a function of the half angle $\alpha_{in} + \alpha_{out}$ and not an independent function of both viewing and illumination angles.

Since we illuminate the surface only from angles $|\alpha_{in}| < \Omega_M$, any surface point on the target whose slope is larger than the angular range, i.e., $|z'(x)| > \Omega_M$, can only reflect rays outside our angular range. Thus if the BRDFs spread light over a range wider than Ω_M, we can crop the parts of the bump area that do not contribute to our angular range, leading to better light efficiency. Also this can possibly make the surface realizable with a smaller slope bound S_M. This cropping is illustrated in Figure 5(b).

3.3. Depth range vs. slope range

As mentioned above, the surfaces we can present on the display layers should have a bounded slope range $|s_k^d(x)| < S_M$. It is worth noting that this limits the range of scene depths we can present, simply because the layer power in Claim 1 increases with depth, resulting in higher slopes. As illustrated in Fig. 5(c), the slopes on the second layer are a sum of the depth-dependent power p_2^d (Eq. (5)), and the slopes of the BRDF bumps, hence wider BRDFs limit the supported depth even further.

For general surfaces these tradeoffs do not follow simple rules, but to gain intuition we consider here the case of planar surfaces with spatially uniform BRDF. In this simple case we can derive the limit on depth range explicitly.

Claim 2. *Consider a planar surface of width W, whose BRDF is spatially uniform, and suppose the maximal slope of the BRDF bumps ζ^v is $S_{BRDF} = \max_x |\zeta^{v\prime}(x)|$. If the maximal slopes of the surfaces z_k^d we can present on the SLM layers is bounded by S_M, then the depth c^v of the target planes we can express on the display should satisfy:*

$$\frac{1}{2}p_2^d W = \frac{D - c^v}{2D^2}W \leq S_M - c^v/D \cdot S_{BRDF} \quad (8)$$

$$\frac{1}{2}p_1^d W = \frac{D - c^v}{2Dc^v}W \leq S_M \quad (9)$$

Proof. Consider a target planar mirror surface with width W at depth c^v, defined by $z^v(x) = c^v$, $\zeta^v(x) = 0$. The closed-form solution from Claim 1 states that the surfaces we need to place on the display layers are parabolas with curvatures $p_1^d = \frac{c^v - D}{Dc^v}$, $p_2^d = \frac{D - c^v}{D^2}$. Over a spatial area

of $[-0.5W, 0.5W]$, these parabolas have slope ranges

$$\left[-\frac{1}{2}p_k^d W, \frac{1}{2}p_k^d W\right] \quad (10)$$

Thus, the planes we can present using SLMs whose slope range is bounded S_M, should have depths c^v such that $\frac{1}{2}p_k^d W \leq S_M$.

Suppose now that our target surface is a plane with non-mirror BRDF. Let ζ^v be the bumps generating this BRDF, and denote the maximal slope of this bump by $S_{BRDF} = \max_x |\zeta^{v\prime}(x)|$. Adding the bumps to the back layer as in Eq. (4), implies that at every local point we add a function with slopes which can reach $c^v/D \cdot S_{BRDF}$, thus the maximal c^v we can support is even more limited, and should satisfy $\frac{1}{2}p_2^d W \leq S_M - c^v/D \cdot S_{BRDF}$. $\square$

To demonstrate the tradeoff of Claim 2, consider Figure 6, demonstrating simple lightfields in flatland at a viewing range $\Omega_M = 0.2$ (equivalent to 12^o). We placed the second layer at depth $D = 2cm$. The target scene here was a simple plane, with some texture to distinguish between neighboring points. We considered targets at depths $z^v = 1cm$ and $z^v = 1.5cm$ and width $W = 1cm$. When the slope bound $S_M = 0.4$ was large enough, we could found good values for the display layers in both cases, and the resulting lightfields are illustrated in Fig. 6(b) (note that the first surface, farther from the 2nd layer, has a stronger lightfield slope). When the bound was reduced to $S_M = 0.3$, the surface farther from the 2nd layer $z^v = 1$ is no longer feasible and the optimization strategy described in the next section produced a lightfield with a lower slope (Fig. 6(c),top), meaning that the solution essentially simulates a plane whose depth is closer to D. Note that the scenes in Fig. 6(b,c) assumed a lambertian BRDF, and the BRDF bumps in this case span a slope range of $S_{BRDF} = \Omega_M = 0.2$, even before the depth-dependent layer curvature is added. Another way to make the plane at depth $z^v = 1$ feasible is to use a narrower depth range; e.g., Fig. 6(d) illustrates the solutions with $S_M = 0.3, S_{BRDF} = 0.04$. The disadvantage of a narrow glossy BRDF is that the scene often appears dark, especially from large viewing/illumination angles (see dark rows in Fig. 6(d)), and at surface areas with steep normals. The narrow BRDF is especially important if we attempt to use a slope range lower than the angular range of the display $S_M < \Omega_M$, in which case to find a feasible solution one must use a rather narrow BRDF so that $S_{BRDF} < \Omega_M$.

As we will see in the results section, these tradeoffs limit the depth variation and the parallax we can obtain with our prototype display, given the limited slope range S_M that the SLMs can support. The dark-regions phenomenon will also be quite apparent in the prototype implementation.
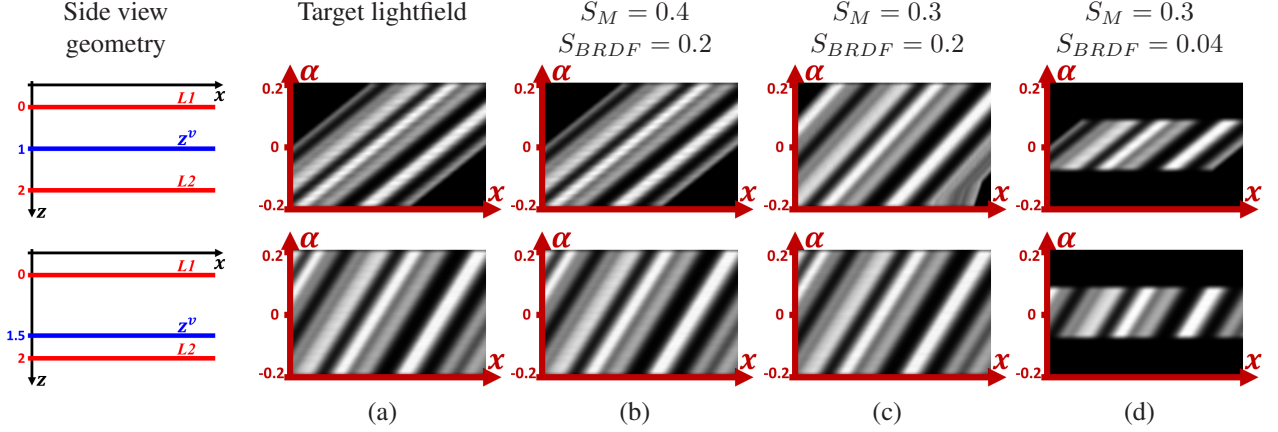
Figure 6: *Visualizing the display slope range constraint, on a flat target at two depths: $z^v = 1$ and $z^v = 1.5$. (a) Flatland lightfield of the target scene as a function of output angular and spatial coordinates (x_{out}, α_{out}), under uniform collimated illumination at $\alpha_{in} = 0$. (b) Lightfield of the display in case S_M is large enough. (c) When slope bound S_M is reduced there is no feasible solution for the first scene, where the surface is farther from the back layer of the display. As a result the optimization resulted in a 1st layer whose power matches a closer plane, as evident by the fact that the lightfield slope in (c,top) does not match the target slope in (a,top). The second target at $z^v = 1.5$ already requires lower slopes so it can be displayed correctly even with the tighter S_M bound. (d) One way to allow a feasible solution for a plane at depth $z^v = 1$ without increasing S_M is to use a more glossy BRDF, with a narrower angular spread. As a result, however, the scene appears dark from some angles within our viewing range.*

4. Global solution

The derivation in the previous section considered thin bas-relief surfaces. In this section, we show how to extend it to more general surfaces that can contain larger depth variations. Although the solutions we can find in the general case are not exact, in practice for many scenes we can find visually plausible approximations. Our strategy will follow the intuition of the planar case: (i) We attempt to design a front layer such that rays focusing at a single point on the target, will focus at a single point on the back layer. (ii) We use the mapping between the target and the back layer in order to copy surface details to the back layer.

In our implementation we use a discretized representation of the surfaces. We denote the set of grid points separated by units Δ_g by $\{x^j\}$, at which we aim to solve for a surface $z_k^d(x^j)$, for $k = 1, 2$.

4.1. Solving for the front layer

To derive the desired properties of the surface z_1^d we refer to Claim 1, which states that around x^j the desired surface power should be

$$p_1^{d,j} = \frac{z^v(x^j) - D}{D z^v(x^j)}, \qquad (11)$$

where $z^v(x^j)$ is the depth of the target surface at point x^j.

Let us denote by $\mathcal{N}(x^j)$ the area where rays emerging from

the point $(x^j, z^v(x^j))$ intersect the front layer plane at $z = 0$. For every keypoint i such that $x^i \in \mathcal{N}(x^j)$ we want the front layer to have the power $p_1^{d,j}$ given in Eq. (11).

We solve for the values $z_1^d(x^j)$ that best approximate this requirement, by minimizing the cost:

$$\sum_j \sum_{i | x^i \in \mathcal{N}(x^j)} \left(-p_1^{d,j} - \frac{z_1^d(x^{i+1}) + z_1^d(x^{i-1}) - 2z_1^d(x^i)}{\Delta_g^2} \right)^2 \qquad (12)$$

To respect the bounds on the slope of the display layer, we minimize Eq. (12) subject to the constraint that

$$\left| s_1^d(x^j) \right| = \left| \frac{z_1^d(x^{j+1}) - z_1^d(x^j)}{\Delta_g} \right| \leq S_M. \qquad (13)$$

Note that when $\{x^j\}$ represent scene points at different depths, the target powers $p_1^{d,j}$ are different. Since every grid point x^i usually belongs to multiple neighborhoods $\mathcal{N}(x^j)$, we cannot expect a zero-error solution. With the least-square cost of Eq. (12), the typical power we get at every grid point is some averaging of the ones at nearby neighborhoods. As will be illustrated in Figure 9, the implication is that the disparities our display obtains in practice are lower than the target disparities.

In addition to the bounded slope constraints of Eq. (13), the next subsection derives constraints that should ensure a feasible solution for the back layer.

Algorithm 1 Display design.

1: Solve for the top layer as a quadratic program, minimizing:

$$\sum_{j} \sum_{x^i \in \mathcal{N}(x^j)} \left(-p_1^{d,j} - \frac{z_1^d(x^{i+1}) + z_1^d(x^{i-1}) - 2z_1^d(x^i)}{\Delta_g^2} \right)^2 \tag{16}$$

Subject to:

$$\left| \frac{z_1^d(x^{j+1}) - z_1^d(x^j)}{\Delta_g} \right| \leq S_M$$

$$\left| \frac{z^v(x^j) \cdot s^v(x^j)}{D} - \frac{(z_1^d(x^{j+1}) - z_1^d(x^j))}{\Delta_g} \right| \leq S_M$$

2: Use ray tracing to find the mapping f between points on the back layer to points on the virtual surface.
3: Construct a back layer z_2^d with slopes:

$$s_2^d(x) = \frac{z^v(f(x)) \cdot s^v(f(x))}{D} - s_1^d(f(x)) \tag{17}$$

4.2. Setting the back layer

Using the front layer we can define a mapping f from points on the 2nd layer to points on the virtual surface, and the inverse mapping $g = f^{-1}$ from the target surface to the back layer. We can determine the shape of the second layer using arguments similar to Claim 4 in the supplementary appendix. The slope $s_2^d(x)$ of the 2nd layer at point x should be set such that rays hitting this point will cross the 1st layer at the same point as rays hitting the target surface at point $(f(x), z^v(f(x)))$ (the proof in the appendix implies that this can be achieved in a unique way as long as the first layer focuses the target on the back layer). A short calculation leads to the following rule:

$$s_2^d(g(x)) = \frac{z^v(x) \cdot s^v(x)}{D} - s_1^d(x), \tag{14}$$

or

$$s_2^d(x) = \frac{z^v(f(x)) \cdot s^v(f(x))}{D} - s_1^d(f(x)). \tag{15}$$

Note that Eqs. (14) and (15) can also be seen as a direct generalization of Claim 1: the slope of the back layer is a scaled version of the virtual surface one, plus a term which is the inverse of the 1st layer slope and should undo its curving.

Back layer feasibility in the front layer optimization: As mentioned above, a feasible back layer surface should have slopes bounded by S_M. To ensure that the back layer is feasible, we want to account for the constraints on the

back layer when solving for the front layer in Eq. (12). For that, we use the fact that the slope in the back layer is a linear function of the front one, given by Eq. (14). Thus we minimize Eq. (12) subject to the constraint

$$\left| \frac{z^v(x^j) \cdot s^v(x^j)}{D} - s_1^d(x^j) \right| \leq S_M. \tag{18}$$

The display design scheme is summarized in Algorithm 1.

In practice, the resolution of the 1st layer is much lower than that of the 2nd layer, and we use that to speed the optimization by selecting a coarse grid spacing Δ_g. To account for this, we replace Eq. (18) by the constraints corresponding to the minimal and maximal derivative s^v in a neighborhood of size Δ_g around x^j.

5. Implementation

After deriving the optical surfaces that we wish to present on our display layers, we discuss here the practical implementation. While one can fabricate the display layers in hardware, we choose to use computer-controlled Spatial Light Modulators (SLM) technology, with the additional advantage of allowing dynamic content.

5.1. Spatial Light Modulators

We used the liquid crystal on silicon "PLUTO" SLM offered by Holoeye Inc. It consists of 1920×1080 refractive liquid crystal cells with a pitch of $\Delta_o = 8\mu m$ over a reflective silicon background layer. When used with polarized light, the cells allow computer control over the index of refraction $\eta_\lambda(x)$ in each individual cell, and these refractive indices can be adjusted dynamically at a speed of 60Hz.

Since the cells are so small, the reflected energy should be analyzed using wave optics tools. In supplementary Appendix C we briefly review the results of [16], which analyzed SLMs in terms of both wave and geometric optics and derived the conditions under which they are equivalent. We also provide simulation results illustrating the quality of the geometric optics approximation.

5.2. Implementing optical surfaces with SLMs

We can use the SLM to implement an optical surface of interest as derived in Sec. 4. We simply sample the surface at intervals of $\Delta_o = 8\mu m$ and set the SLM phase of a pixel x as $\phi(x) = 4\pi d_o \eta_\lambda(x)/\lambda = 4\pi z(x)/\lambda$. It has been observed by [16] that using this conversion will produce reasonable results provided that the slopes of the optical surfaces are smooth over spatial coherence windows of size Δ_c, whose size is determined by the angular resolution the display aims to support. The exact relation between spatial

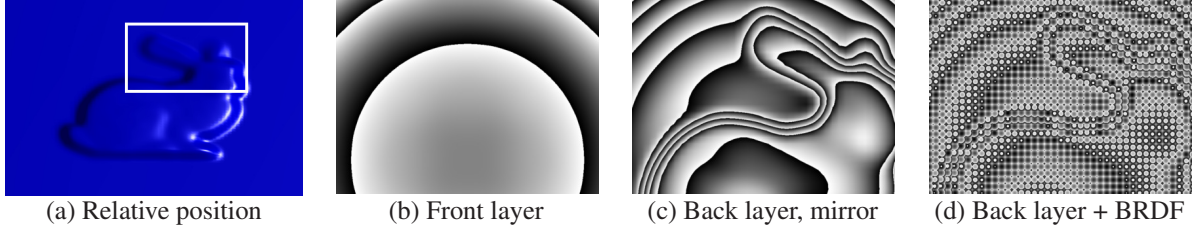| (a) Relative position | (b) Front layer | (c) Back layer, mirror | (d) Back layer + BRDF |

Figure 7: *A crop from the SLM assignment resulting from the surfaces in Figure 3. (a) Relative position. (b) Front layer. (c) Back layer for a mirror reflectance. (d) Back layer with BRDF micro-features. The surfaces in (b,c) are obtained from Figure 3(c,d) modulo the maximal phase retardation of the SLM.*

coherence and angular resolution is detailed in the supplementary appendix. In our design we selected the width Δ_d of the added BRDF bumps to be a few times wider than the spatial coherence Δ_c. For example, assuming the viewing/ illumination angle cannot go below $\Delta_a = 0.5^o - 1^o$, which is the angle subtended by a $5cm$ source placed $3-6m$ from the display, leads to a coherence length of $\Delta_c = 30-60\mu m$. The bumps we added occupied 20×20 or 10×10 SLM units, resulting in a dot area of $\Delta_d = 80-160\mu m$, which is below the minimal resolution observed by a human standing a reasonable distance from the surface. In the supplementary, we also evaluate the accuracy of the geometric optics approximation through simulations. With 1920×1080 SLM cells and $10 \times 10 - 20 \times 20$ cell dots, we are able to display images of 192×108 or 96×54 dots, where a "dot" corresponds to a pixel in the viewer resolution.

The SLM implementation suffers from two main restrictions, on angular range and thickness. As explained in supplementary Appendix C, the first restriction results from the fact that since the SLM has a limited pitch of $\Delta_o = 8\mu m$ pitch, the maximal surface slope s we can present is limited as $|s| \leq S_M^\lambda$:

$$S_M^\lambda = \frac{\lambda}{4\Delta_o}. \qquad (19)$$

For the red illumination at $\lambda = 633nm$, we get $S_M = 0.02 = 1.13^o$. Restricting surface normals to $|s| \leq S_M$, the experiments in Sec. 6.2 are designed to support an angular range of $\Omega_M = 2S_M$. That is, it supports illumination and viewing direction between -2.25^o to 2.25^o, an overall range of 4.5^o. Since a smaller cell pitch translates directly into a larger angular range, this limitation will soften substantially as hardware improves. Devices with $6mm$ cell pitch are already commercially available (LETO, Holoeye), and devices with $2mm$ cell pitch are in development [27].

A second restriction of SLMs is that they have limited thickness. Therefore, rather than presenting $\phi(x) = 4\pi z(x)/\lambda$, we present $\phi(x) = \text{mod}(4\pi z(x)/\lambda, T)$, where T is the maximal phase derived from the thickness of the SLMs. As clipping is a function of wavelength, this results in some chromatic artifacts. Fig. 7 illustrates the SLM assignment $\phi(x)$, corresponding to the display surfaces of Fig. 3.
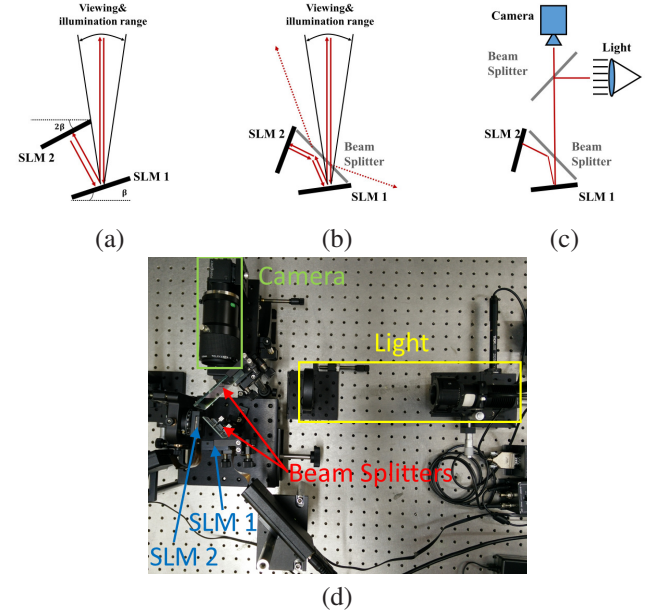

Figure 8: *Experimental setup: while we envision a display with one transmissive and one reflective layer, due to the limited resolution of available transmissive SLMs we used two reflective ones. (a) Positioning the two reflective SLMs with proper tilt between them can simulate the configuration in Figure 1(a). (b) To avoid occlusions between the SLMs without using a large tilt angle, we used a beam splitter. (c) The display should support every illumination and viewing angle inside a target angular range. To allow demonstrating the display from every illumination and viewing direction, without occlusions between the light source and the camera, we combined them through another beam splitter. (d) An image of our physical setup.*

5.3. Experimental setup

To implement the system described above we need one transmissive and one reflective SLM. While liquid crystal cells are transmissive in their nature, advanced SLMs rely on LCoS (Liquid Crystal on a Silicon) background technology, and are effectively reflective. Since the transmissive

SLMs available on the market are of significantly lower resolution, we used two reflective SLMs.

The simplest reflective setup is illustrated in Figure 8(a). One SLM layer is positioned at angle β to the camera image plane, a ray reflects from it toward a 2nd SLM positioned at angle 2β from the camera so that a ray oriented along the optical axis will return to the camera along the same axis if the SLMs are acting as flat mirrors. Ideally, the tilt angle β should be relatively small to minimize SLM aberrations, however occlusions limit the minimal angle at which we can view both mirrors. In practice, we avoid occlusions using an equivalent construction with a beam splitter, as in Fig. 8(b). Finally, we want to allow the camera and illumination to be positioned anywhere within the same angular range $-\Omega_M \leq \alpha_{in}, \alpha_{out} \leq \Omega_M$. To allow us to image the display in the lab under all these configurations, without occlusions between the camera and the source, we recombine them through a second beam splitter resulting with the setup of Fig. 8(c). An image of our setup is available in Fig. 8(d). In this setup the distance between the two layers is $D = 35mm$.

We used a band-limited red LED illumination centered at $\lambda = 633nm$ with a $20nm$ bandwidth. As demonstrated by [16], illuminating smooth SLM patterns of the type we are using here with broadband white illumination usually leads to modest chromatic artifacts, but we did not test this. We used a monochromatic recording sensor, resulting in monochromatic images. The LED source was placed at the focal length of a lens to produce collimated illumination.

The setup was imaged using the Computar TEC-M55 telecentric lens. In practice the level of telecentricity achieved by the lens is limited, implying that the camera projection is not perfectly orthographic. As a result, not all pixels are viewed from exactly the same direction, causing some spatial variations in the viewed images. Another restriction on the angular variation that we could present is that the acceptance angle of the camera is about 1^o.

6. Results

As listed in Sec. 2, the main assumptions inherited in the algorithm are the inability to model self-occlusions and subscattering. At the same time, the limited pitch of existing SLMs limit the slope range S_M, and as explained below that significantly limits the angular range of the display as well as the range of depth content it can support. To distinguish between algorithmic restrictions and hardware restrictions we present here two types of results: (i) Synthetic simulations that allow us to demonstrate our algorithm and study its limitations independently of the restrictions of current SLM technology. (ii) Real images captured with our prototype of two SLMs, subject to their limited slope range S_M.

Figure 9: *The leftmost and rightmost viewpoints of synthetic surfaces superimposed in different color channels. Our display (2nd row) produces weaker disparity compared to the target (top row) due to neighborhood averaging of the 1st layer power.*

6.1. Synthetic results

In our synthetic simulations, we targeted an angular range of $\Omega_M = 0.2$ (corresponding to a viewing angle of $[-12, 12]^o$, which is a reasonable viewing range for most autosteroscopic displays), a viewing dot resolution of $\Delta_d = 0.1mm$ and scenes whose depths vary at the range $[1.5, 2]cm$, this allowed for a nice disparity shift of 20 dots between the farthest and closest scene parts.

In Figures 9–11, we present a few scenes that were designed for our display. In each case we demonstrate geometric-optics ray-tracing rendering of the target compared to the solution of our display design algorithm. Please also see videos in the supplementary file.

To visualize the disparity obtained by our display, we synthesize in Fig. 9 the leftmost and rightmost views in our range as two color channels of the same image, so that a shift from a red feature to a blue one corresponds to the feature depth and is stronger for the front parts. One can notice that while our algorithm obtained some depth-dependent disparity, its output disparity is lower than what is observed at the target. As explained after Eq. (12), this is due to the fact that rays through every grid point on the first layer reach scene points at different depths, and the target powers of these depths are averaged spatially. However, the reduced disparities here are not related to a bound on the slope as illustrated in Fig. 6. In fact, without constraining the optimization explicitly, the slopes used by the display layers for the scenes in Figure 9 are all below $S_M = 0.25$. For an angular range $\Omega_M = 0.2$, the BRDF bumps alone require slopes in the range $[-0.2, 0.2]$ so the geometry did not add much to the slope beyond it.
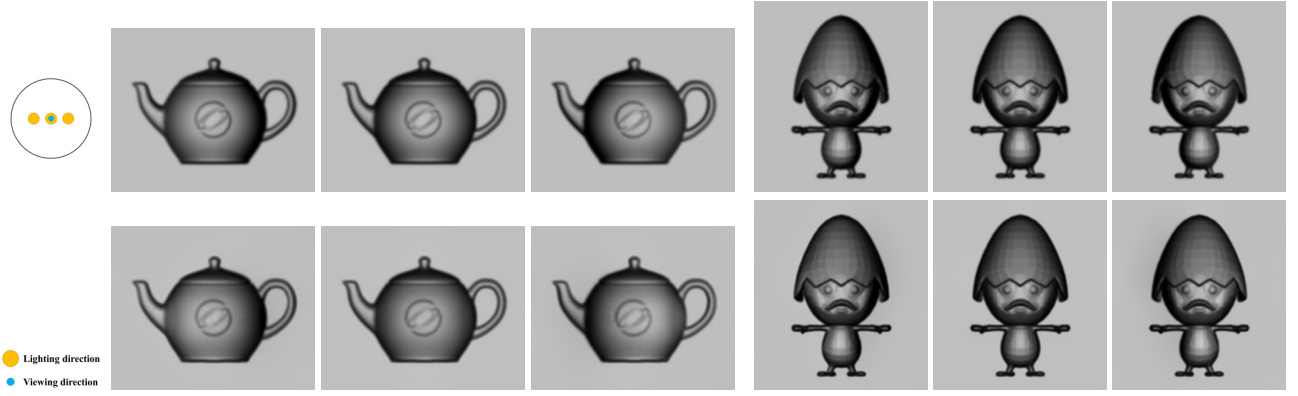
Figure 10: *Synthetic surfaces under collimated, spatially-uniform illumination from varying angles. Our display (2nd row) accurately reproduces attached shadows and highlights in the target (top row).*
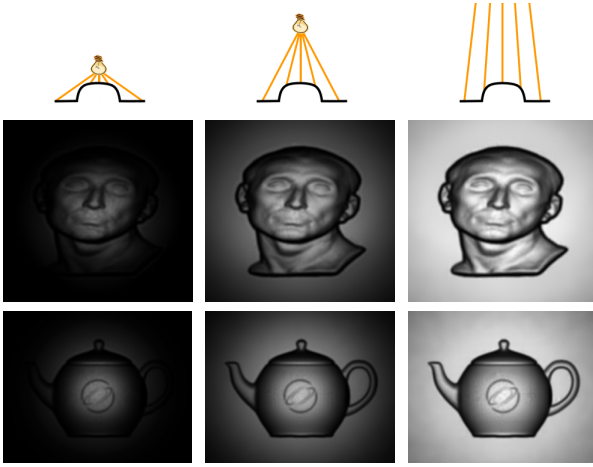

Figure 11: *Spatially varying illumination: a point source at varying heights above the surface, illuminating different surface points at radially varying angles, leading to spatially varying brightness. The top-row icons illustrate the source position relative to the surface. Note that the individual image brightness is normalized independently.*

Figure 10 illustrates scenes from a fixed viewpoint, under collimated (spatially uniform) illumination from varying orientations. See also animation in the supplementary file. While as explained earlier, our scenes do not contain occlusions within the angular range of the display and hence no cast shadows are present, our display does produce accurate reproduction of illumination-sensitive effects in the target, such as highlight and cast shadows.

Our display supports spatially-varying illumination effects. Figure 11 demonstrates this using a point light source at varying distances. As the source distance vary, the angle at which light rays hit surface points varies, affecting the reflected intensity. For example, when the source is very

close, the periphery of the image is illuminated at high grazing angles and thus appears dark. Note that the images have been normalized to have equal energy.

6.2. Prototype results

Unlike the wide angular and depth ranges demonstrated in the synthetic examples of 6.1, existing SLM technology provide limited range and resolution. As mentioned in Sec. 5, inherited by the limited pitch, the major restriction on the SLMs we used is the limited slope range $S_M = 0.02$, which corresponds to 1.14^o degrees. Aiming to display some angular variation, we attempted to stretch the angular range up to $\Omega_M = 2S_M = 0.04$. The limited slope ranges restricted our scene content in two ways. The first restriction comes from the fact that since the target angular range is larger than the slope range of the display $\Omega_M > S_M$, the BRDFs we could add were narrower than the angular range, in the sense that any surface point looks dark through a subset of viewing angles within our range. In particular surface points with steep gradients appear dark from anywhere inside our viewing range. A second restriction, as derived in Claim 2, is that the slope range S_M limits the depth range we can program on our display. This restriction results from the fact that if the distance $|c^v - D|$ from the target surface to the second layer is large, the powers of the display layers increase. In our setup with $D = 35mm$, applying Claim 2 for a planar target whose area is about half of the display width and with a narrow mirror BRDF implies that the depth of the target surface can vary in the range $[30-42]mm$. This depth range, combined with the limited viewing range, translates into depth parallax of maximum 6 Δ_d-sized dots between the farthest and closest scene points (recall that dots are the pixels in the viewable resolution). If we attempt to add a BRDF such that the slope of our bumps ζ^v ranges between $[-0.01, 0.01]$, the depth range reduces to $[32-38]mm$, and the maximal parallax reduces to about 2 dots.
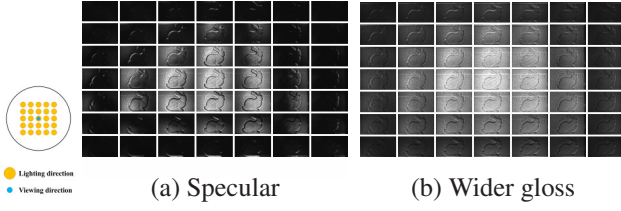
(a) Specular   (b) Wider gloss

Figure 12: *The bas-relief bunny from Figure 7 with a mirror reflectance and wider BRDF.*
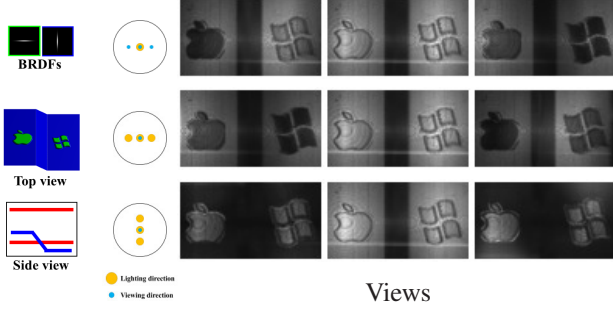

Views

Figure 13: *Left: A target surface of a step edge (two planes at two different depths). The target includes a spatially-varying texture with horizontally and vertically anisotropic-reflecting BRDFs. Right: a few views of the prototype display, from varying viewing and lighting directions that are indicated using their projection on the x-y plane. Since the surface consists of anisotropic BRDFs, the background appears bright when the illumination varies horizontally and dark when it moves vertically. The opposite effect is obtained inside the logos.*

Demonstrating all variations in an 8D lightfield is challenging. Yet we demonstrate images from multiple illumination and viewing directions. We used a collimated source that illuminates the display area uniformly. In all figures we denote the viewing/lighting directions as blue/orange points representing their projection to the $x - y$ plane.

Figure 12 demonstrates the bas-relief bunny surface from Figs. 3 and 7 with both a mirror appearance and a wider BRDF. As this surface is relatively thin no depth parallax is observed, yet varying illumination varies appearance due to the variations in gradient orientations.

Figure 13 demonstrates a step edge whose depth variation is more dominant, with two logos on each side of the step. The logo areas demonstrate spatial variations in BRDF. Recall that our BRDFs vary as a function of the mean of lighting and viewing directions (the half vector) and not as an independent function of both of them, thus we can present the BRDF as a 2D function rather than a 4D one, demonstrating the amount of energy reflected as a function of the x, y components of $\alpha_{in} + \alpha_{out}$. In this example we used two anisotropic BRDFs with opposite orientations, one for the background and one for the inner logo area. As illustrated in Fig. 13, when the display is lit with horizontal lighting
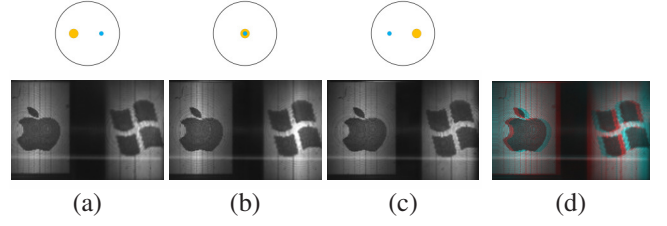

(a)   (b)   (c)   (d)

Figure 14: *Another slicing of the step edge lightfield from Fig. 13. The illumination direction is chosen as a symmetric reflection of the viewing direction, reducing variations in image brightness and emphasizing viewpoint parallax. In (d) we demonstrate the parallax by fusing the leftmost and rightmost images as different color channels.*

the background appears bright and the logo is dark. When we move the illumination to a vertical direction the background becomes dark and the logo turns bright. In Fig. 14 we demonstrate the same scene with a mirror BRDF (allowing for a somewhat larger depth range). In this figure the lighting direction was selected as the reflection of the mirror direction, so that the half angle between the lighting and viewing direction is always $(0,0)$, implying that no brightness change is observed when shifting viewpoints. This makes viewpoint-dependent depth parallax more observable. To demonstrate this parallax, in Fig. 14(d) we fuse the leftmost and rightmost images as two color channels of one image.

Fig. 15 demonstrates the Barbapapa figure with character details on is upper facet emphasized using a different BRDF. The gradient at the stage edges lie outside our angular range and thus appear dark. Varying the illumination direction shrinks and expands the size of the shadow region, providing a sense of 3D. This is mostly an attached shadow effect, which happens since the orientation of surface normals at the border of the step vary rapidly. Given the relatively narrow angular extent of our BRDFs, small changes in illumination angle affect which pixels appear bright.

7. Discussion

This paper presents a 3D light-sensitive display capable of producing the illumination and viewpoint sensitive appearance of simple virtual 3D scenes. Our contributions are: (i) Suggesting a setup based on only two planar SLMs used to display programable thin optical surfaces. (ii) Deriving a simple closed-form solution for the display of thin bas-relief surfaces, which reduces to using the first layer as a lens focusing the target surface on the second layer, and then using the second layer to copy all surface details from the target. (iii) Proposing an approximate optimization framework for the display of more general surfaces with larger depth variations, involving a simple convex optimization task with a
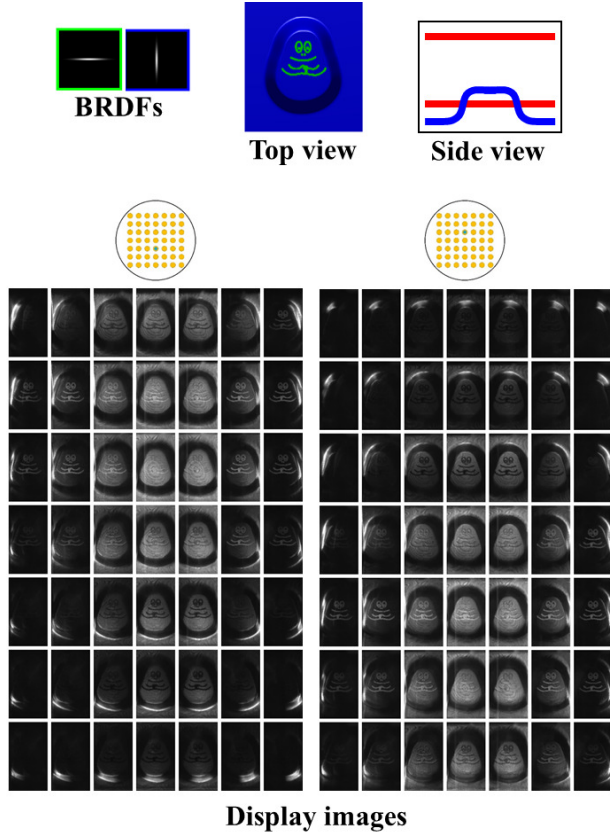
BRDFs　　**Top view**　　**Side view**

Display images

Figure 15: *Our prototype displaying a Barbapapa character, imaged from two viewpoints under a 7×7 grid of lighting directions. The background had an anisotropic BRDF in the vertical direction, hence appears bright when illumination is varied vertically, while the character details have vertically anisotropic BRDF and appears bright as illumination is varied on the horizontal axis.*

global solution. (iv) Building a small prototype and demonstrating the display of a few simple 3D scenes.

Among the limitations of the display: (i) Due to the limited size and pixel pitch of existing SLMs, the spatial size and angular range of our current prototype are quite limited. (ii) The limited angular extent limits the range of viewable surface normals in our target scenes. This also reduces the amount of depth variation we can include in the displayed scenes. (iii) Our algorithm assumes all surfaces are opaque and does not consider interesting light-transport effects such as refraction and subsurface scattering. (iv) The surfaces we support contain no self occlusions within the angular range of interest, and hence cast shadows cannot be demonstrated. (v) The display is restricted to gray scenes and does not support color variations.

Our design algorithms rely on geometric optics models rather than accurate wave optics ones. Our main motivation was to achieve simple optimization algorithms, which will allow us to gain intuition on the structure of the so-

lution and its limitations. We were also motivated by the empirical observations of Glasner *et al.* [16], stating that when the SLM content is sufficiently smooth, the geometric optics prediction provides a reasonable approximation of wave optics effects. Nevertheless, we believe the results can be improved if wave optics is taken into account. We hope that the geometric optics solution can serve as a useful initialization for a non-convex optimization accounting for electromagnetic effects.

While any improvement in SLM technology would directly translate into improvement in our display, another way to relax the limitations is to use a cascade of more than two SLMs. The first advantage of multiple SLMs is that they provide a larger set of adjustable variables and hence more flexibility in the approximation of harder scenes with larger depth variation. A second advantage is that multiple layers can allow increasing the angular range of the display. While the ray bending of each layer is bounded by S_M, their concatenation can achieve larger bending. Finally, if the layers can be placed close to each other, electromagnetic effects such as Bragg diffraction can be used to further enhance the power of the display.

References

[1] L. Ahrenberg, P. Benzie, M. Magnor, and J. Watson. Computer generated holograms from three dimensional meshes using an analytic light transport model. *Appl. Opt.*, 47(10):1567–1574, Apr 2008.

[2] K. Akeley, S. J. Watt, A. R. Girshick, and M. S. Banks. A stereo display prototype with multiple focal distances. *ACM Trans. Graph.*, 23(3):804–813, Aug. 2004.

[3] I. Baran, P. Keller, D. Bradley, S. Coros, W. Jarosz, D. Nowrouzezahrai, and M. Gross. Manufacturing layered attenuators for multiple prescribed shadow images. *Eurographics*, 31(2):603–610, 2012.

[4] S. A. Benton and V. M. Bove. *Holographic Imaging*. Wiley-Interscience, 2007.

[5] S. Borgsmüller, S. Noehte, C. Dietrich, T. Kresse, and R. Männer. Computer-generated stratified diffractive optical elements. *Appl. Opt.*, 42(26):5274–5283, Sep 2003.

[6] D. Brady and D. Psaltis. Control of volume holograms. *JOSA*, 9:11671182, 1992.

[7] D. Chambers, G. Nordin, and S. Kim. Fabrication and analysis of a three-layer stratified volume diffractive optical element high-efficiency grating. *Opt. Express*, 11(1):27–38, Jan 2003.

[8] O. Cossairt, S. K. Nayar, and R. Ramamoorthi. Light Field Transfer: Global Illumination Between Real and Synthetic Objects. *ACM SIGGRAPH*, 2008.

[9] W. J. Dallas. Computer-generated holograms. *The Computer in Opt. Research of Topics in Appl. Physics*, 41:291–366, 1980.

[10] D. J. DeBitetto. Holographic panoramic stereograms synthesized from white light recordings. *Appl. Opt.*, 8(8):1740–1741, Aug 1969.

[11] Y. Dong, J. Wang, F. Pellacini, X. Tong, and B. Guo. Fabricating spatially-varying subsurface scattering. *ACM Trans. Graph.*, 29(4):62:1–62:10, July 2010.

[12] G. E. Favalora. Volumetric 3d displays and application infrastructure. *Computer*, 38(8):37–44, Aug. 2005.

[13] M. Finckh, H. Dammertz, and H. P. A. Lensch. Geometry construction from caustic images. In *ECCV*, 2010.

[14] M. Fuchs, R. Raskar, H. Seidel, and H. Lensch. Towards passive 6D reflectance field displays. *SIGGRAPH*, 2008.

[15] T. Gerke and R. Piestun. Aperiodic volume optic. *Nature Photonics*, 2010.

[16] D. Glasner, T. Zickler, and A. Levin. A reflectance display. *ACM SIGGRAPH*, 2014.

[17] J. W. Goodman. *Introduction to Fourier Optics*. McGraw-Hill Book Company, 1968.

[18] M. Hašan, M. Fuchs, W. Matusik, H. Pfister, and S. Rusinkiewicz. Physical reproduction of materials with specified subsurface scattering. *ACM SIGGRAPH*, 2010.

[19] E. Hecht. *Optics*. Pearson education. Addison-Wesley, 2002.

[20] M. Hirsch, S. Izadi, H. Holtzman, and R. Raskar. 8D display: a relightable glasses-free 3d display. In *ACM ITS*, 2012.

[21] M. Hirsch, S. Izadi, H. Holtzman, and R. Raskar. 8D: interacting with a relightable glasses-free 3D display. In *SIGCHI*, 2013.

[22] M. Hirsch, D. Lanman, H. Holtzman, and R. Raskar. BiDi screen: a thin, depth-sensing LCD for 3D interaction using light fields. *ACM Trans. Graph.*, 28(5), 2009.

[23] R. Horisaki and J. Tanida. Reflectance field display. *Opt. Express*, 21(9):11181–11186, May 2013.

[24] Z. Hu, A. Yan, D. Liu, X. Wang, and L. Liu. Anisotropic diffraction of stratified volume holographic grating in photorefractive media. *Optik - Int. J. for Light and Electron Opt.*, 117(8):393 – 397, 2006.

[25] M. B. Hullin, I. Ihrke, W. Heidrich, T. Weyrich, G. Damberg, and M. Fuchs. Computational fabrication and display of material appearance. In *Eurographics State-of-the-Art Reports (STAR)*, May 2013.

[26] M. B. Hullin, H. P. A. Lensch, R. Raskar, H.-P. Seidel, and I. Ihrke. Dynamic display of BRDFs. In *EUROGRAPHICS*, pages 475–483, 2011.

[27] IMEC. IMEC scientific report 2011. http://www.imec.be/ScientificReport/SR2011/1414043.html, 2011.

[28] R. V. Johnson and A. R. T. Jr. Stratified volume holographic optical elements. *Optics News*, 14(12):30–31, Dec 1988.

[29] A. Jones, I. McDowall, H. Yamada, M. Bolas, and P. Debevec. Rendering for an interactive 360 light field display. *ACM Trans. Graph.*, 26(3), July 2007.

[30] T. Kämpfe, E.-B. Kley, A. Tünnermann, and P. Dannberg. Design and fabrication of stacked, computer generated holograms for multicolor image generation. *Appl. Opt.*, 46(22):5482–5488, Aug 2007.

[31] S.-C. Kim, J.-W. Moon, D.-H. Lee, K.-C. Son, and E.-S. Kim. Holographic full-color 3D display system using color-LCoS spatial light modulator, 2005.

[32] T. Kiser, M. Eigensatz, M. M. Nguyen, P. Bompas, and M. Pauly. Architectural caustics controlling light with geometry. In *Advances in Architectural Geometry*, 2012.

[33] H. Kogelnik. Coupled-wave theory for thick hologram gratings. *Bell System Technical Journal*, 48(2909), 1969.

[34] T. Koike and T. Naemura. BRDF display: interactive view dependent texture display using integral photography. In *EDT-IPT*, pages 6:1–6:4, 2008.

[35] Y. Lan, Y. Dong, F. Pellacini, and X. Tong. Bi-scale appearance fabrication. *ACM Trans. Graph.*, 32(4), 2013.

[36] A. Levin, D. Glasner, Y. Xiong, F. Durand, B. Freeman, W. Matusik, and T. Zickler. Fabricating BRDFs at high spatial resolution using wave optics. *ACM SIGGRAPH*, 2013.

[37] T. Malzbender, R. Samadani, S. Scher, A. Crume, D. Dunn, and J. Davis. Printing reflectance functions. *ACM Trans. Graph.*, 31(3):20:1–20:11, 2012.

[38] W. Matusik, B. Ajdin, J. Gu, J. Lawrence, H. P. Lensch, F. Pellacini, and S. Rusinkiewicz. Printing spatially-varying reflectance. *ACM SIGGRAPH Asia*, 28(5), 2009.

[39] S. Nayar, P. Belhumeur, and T. Boult. Lighting Sensitive Display. *ACM Trans. on Graphics*, 23(4):963–979, Oct 2004.

[40] M. Papas, W. Jarosz, W. Jakob, S. Rusinkiewicz, W. Matusik, and T. Weyrich. Goal-based caustics. *Eurographics*, 2011.

[41] G. Patow and X. Pueyo. A survey of inverse surface design from light transport behavior specification. *Comput. Graph. Forum*, 24(4):773–789, 2005.

[42] G. Patow, X. Pueyo, and A. Vinacua. User-guided inverse reflector design. *Comput. Graph.*, 31(3):501–515, June 2007.

[43] J. Redman. The three-dimensional reconstruction of people and outdoor scenes using holographic multiplexing. *SPIE Seminar-in Depth on Holography*, 15:117122, 1968.

[44] O. Rouiller, B. Bickel, J. Kautz, W. Matusik, and M. Alexa. 3d-printing spatially varying brdfs. *IEEE Computer Graphics and Applications*, 33, 2013.

[45] Y. Schwartzburg, R. Testuz, A. Tagliasacchi, and M. Pauly. High-contrast computational caustic design. *ACM SIGGRAPH*, 2014.

[46] D. E. Smalley, Q. Y. J. Smithwick, V. M. Bove, J. Barabas, and S. Jolly. Anisotropic leaky-mode modulator for holographic video displays. *Nature*, (7454):313317, 2013.

[47] P. St-Hilaire, S. Benton, M. Lucente, and P. Hubel. Color images with the MIT holographic video display, 1992.

[48] K. E. Torrance and E. M. Sparrow. Theory for off-specular reflection from roughened surfaces. *J. Opt. Soc. Am.*, 57(9):1105–1112, Sep 1967.

[49] G. Tricoles. Computer generated holograms: an historical review. *Appl. Opt.*, 26(20):4351–4357, Oct 1987.

[50] G. Wetzstein, D. Lanman, M. Hirsch, and R. Raskar. Tensor displays: compressive light field synthesis using multilayer displays with directional backlighting. *ACM SIGGRAPH*, 2012.

[51] T. Weyrich, P. Peers, W. Matusik, and S. Rusinkiewicz. Fabricating microgeometry for custom surface reflectance. *ACM. SIGGRAPH*, 2009.

[52] L. Yaroslavsky. *Digital Holography and Digital Image Processing*. Kluwer Academic Publishers, 2004.

[53] G. Ye, S. Jolly, V. M. Bove, Q. Dai, R. Raskar, and G. Wetzstein. Toward BxDF Display using Multilayer Diffraction. *SIGGRAPH Asia*, 33(6), 2014.

[54] Y. Yue, K. Iwasaki, B.-Y. Chen, Y. Dobashi, and T. Nishita. Poisson-based continuous surface generation for goal-based caustics. *ACM Trans. Graph.*, 2014.

A. Appendix: Bas-relief focusing rules

In this appendix we prove Claim 1 of the main paper, deriving the closed-form display assignment for bas-relief surfaces.

A.1. Paraxial model

In order to derive the structure of our display, we need to express the ray transfer operators $A_v, A_d, A_v^f, A_d^f, A_v^b, A_d^b$ mentioned above. As is often done in the optics literature [19], we use the paraxial approximation of ray optics. The paraxial approximation of ray optics is valid when all rays are assumed to have a small angle $\alpha \approx \tan(\alpha)$ and a small distance x relative to the optical axis of the system. We denote a light ray by a 3×1 vector $\mathbf{r} = (x, \alpha, 1)^T$. When the entries of the ray are small we can approximate the ray propagation transformation of an optical system as a linear matrix multiplication, encoding a first-order Taylor expansion of the non-linear ray propagation formulas:

$$\mathbf{r}_{out} = A\,\mathbf{r}_{in} \qquad (20)$$

For a system of two layers as in our display, each ray passes through a few simple transformations inside the system: 1) Refraction by the front layer. 2) Free-space transfer from the front layer to the back layer. 3) Reflection by the back layer. 4) Transfer from the back layer to the front layer. 5) Inverse refraction through the front layer. To describe the operation of our system we thus need to express three basic transformations: (a) Free-space transfer, (b) Refraction, and (c) Reflection. We will show how to express each of these individual transformations. To describe the full system, we will have to concatenate these building blocks. Since we are working with a linear approximation, the concatenation reduces to simple matrix multiplication.

To start with, consider the case of free-space transfer: a ray $\mathbf{r} = (x, \alpha, 1)^T$ exiting the front layer at direction α will reach the back layer without changing its direction, while its spatial position is shifted by $D \cdot \alpha$ (where D denotes the distance between the layers). This transformation can be expressed as

$$\mathbf{r}_{out} = T_D\,\mathbf{r}_{in}, \qquad T_D = \begin{pmatrix} 1 & D & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}. \qquad (21)$$

To describe the ray transformation associated with the mirror-layer reflection, note that this transformation changes only the ray angle and not the intersection point. Since the slopes of interest in this paper are small, the mirror normal can be expressed as $\mathbf{n}_2^d(x) \approx (s_2^d(x), -1)^T$. Recall that for a an incoming ray $\mathbf{l}$, the reflected vector $\mathbf{v}$ is defined as $\mathbf{v} = \mathbf{l} - 2 < \mathbf{n}, \mathbf{l} > \mathbf{n}$, and for small angles a

first-order approximation results in

$$\mathbf{v} \approx (\alpha_{in} + 2s_2^d(x), 1)^T, \qquad (22)$$

or in other words, the reflected angle satisfies $\alpha_{out} = \alpha_{in} + 2s_2^d(x)$. In a similar way, it can be shown [19] that refraction by a transmissive layer is described by

$$\alpha_{out} = \alpha_{in} + (\eta - 1)s_1^d(x) \qquad (23)$$

where η is the refractive index of the transmissive material. To simplify notation we omit the factor $(\eta - 1)$ in our derivation.

In this paper, the surfaces we place on the front layer are relatively smooth, and thus we can use a second-order approximation of the form

$$z_1^d(x) \;=\; -p_1^d(x - x^o)^2 - t_1^d(x - x^o) + c_1^d, \quad (24)$$

and a first-order approximation to its slope:

$$s_1^d(x) \;=\; \frac{\partial z_1^d(x)}{\partial x} = -p_1^d(x - x^o) - t_1^d. \qquad (25)$$

Assuming the spatial position x_{in} of an incoming ray $\mathbf{r}_{in}$ is sufficiently close to an optical center x^o, we can express the refraction operation as a linear matrix multiplication:

$$\mathbf{r}_{out} = L_{p_1^d, t_1^d}\,\mathbf{r}_{in}, \qquad L_{p_1^d, t_1^d} = \begin{pmatrix} 1 & 0 & 0 \\ -p_1^d & 1 & p_1^d x^o - t_1^d \\ 0 & 0 & 1 \end{pmatrix}. \quad (26)$$

Note that in the optics literature the p_1^d factors are usually referred to as the power of the optical surface, since they determine by how much rays are bent. The t_1^d components are non-zero if the optical surface is not centered at $x = x^o$, but at $x^o - t_1^d/p_1^d$.

Having defined the basic building blocks of a propagation in an optical system, we can describe the ray transfer of our system by concatenating these transformations. Since we are working with linear approximations, we can simply multiply the individual matrices. For example, the transformation A_d^f that a ray entering the system will pass from the front layer to the last one can be described by concatenating the refraction and transfer transformations:

$$\mathbf{r}_{int} = A_d^f\,\mathbf{r}_{in}, \qquad A_d^f = T_D \cdot L_{p_1^d, t_1^d}, \qquad (27)$$

and similarly,

$$A_d^b = L_{p_1^d, t_1^d} \cdot T_D. \qquad (28)$$

For a virtual surface whose local depth is c^v the forward and backward transformation are straightforward

$$A_v^f = A_v^b = T_{c^v}. \qquad (29)$$

A.2. Deriving the front layer

With the definition of ray transfer operators from Sec. A.1, we are ready to derive the structure of the display layers and prove Claim 1. We start with the front layer, whose goal is to focus on the back layer the rays that would otherwise have met on the target surface. That is, to allow unique mapping between points on the target surface and points on the back display layer. Below we use the notation $A(j, :)$ to refer to the j-th row of the matrix A.

Claim 3. *Given a virtual plane at depth c^v, a lens satisfying*

$$A_v^f(1,:)\mathbf{r}^1 = A_v^f(1,:)\mathbf{r}^2 \Rightarrow A_d^f(1,:)\mathbf{r}^1 = A_d^f(1,:)\mathbf{r}^2 \quad (30)$$

can have an arbitrary center t_1^d, but should have power

$$p_1^d = \frac{c^v - D}{Dc^v}. \quad (31)$$

Proof. We can compute the desired power in Eq. (31) directly from the lens law. However, since it will be helpful later, we will use the ray transfer matrix formulation. Let $\mathbf{r}^o = (x^o, \alpha, 1)^T$ be a ray reaching point x^o on the 2nd layer. This ray originated from the ray $(A_d^f)^{-1}\mathbf{r}^o$ at the $z = 0$ plane, and in the virtual world it will be mapped to $A_m\mathbf{r}^o$ with $A_m = A_v^f(A_d^f)^{-1}$. If we want that all rays $\mathbf{r}^o$ at spatial position x^o on the 2nd layer will be mapped to a single point on the virtual plane, regardless of their angle α, we need to ensure that $A_m(1, 2) = 0$. Namely, since this A_m element determines how much the angle affects the target spatial point, nulling it implies that the target spatial point is independent of α. We use Eqs. (27) and (29) to express A_m, and select a p_1^d value that leads to $A_m(1, 2) = 0$, producing Eq. (31). As it will be useful later, we include the formula for the resulting mapping A_m:

$$A_m = \begin{pmatrix} \frac{c^v}{D} & 0 & c^v t_1^d \\ \frac{c^v - D}{Dc^v} & \frac{D}{c^v} & t_1^d \\ 0 & 0 & 1 \end{pmatrix} \quad (32)$$

□

The mapping A_m in Eq. (32) specifies explicitly the mapping from points on the back display layer to points on the virtual plane:

$$f(x^o) = \frac{c^v}{D}x^o + c^v t_1^d \quad (33)$$

As a sanity check, note that since the lens is imaging points at depth D to points at depth c^v it is expected to have a magnification of c^v/D, as evident by Eq. (33).

A.3. Deriving the back layer

Having derived the shape of the front layer, we need to find an assignment for the back layer that will ensure that the display mimics the ray transfer of the virtual surface and $A_d = A_v$.

Recall that rays hitting the back layer at a spatial point x^o are reflected according to the local surface normal. For every point x^o we should set this normal so as to achieve the desired output. We note from Eq. (32) that rays meeting at a single point on the virtual plane will be mapped by the front layer to a single point on the back layer. However, the arrival angle is not preserved, but passes an affine transformation of the form $\alpha_{new} = \frac{D}{c^v}\alpha_{old} + \frac{c^v - D}{Dc^v}x^o + t_1^d$. That is, the farther the point x^o is from the optical axis, the larger the ray bending. To undo this bending, we should select the surface normals on the back layer accordingly.

We show below that this normal should be a scaled version of the normal at the corresponding point on the virtual surface, plus a term that will undo the curving of the front layer. This leads to the following rule:

Claim 4. *Given a quadratic front layer set by Eq. (31), selecting the back layer as a quadratic curve plus a scaled version of the target surface*

$$z_2^d(x) = -\frac{1}{2}p_2^d(x)^2 - t_2^d x + z^v\left(\frac{c^v}{D}x + c^v t_1^d\right) - c^v \quad (34)$$

with

$$p_2^d = \frac{D - c^v}{D^2}, \quad t_2^d = \frac{c^v}{D}t_1^d \quad (35)$$

will result in outgoing rays that simulate the ray transfer of the target $A_d = A_v$.

Proof. Before diving into the proof it is worth mentioning that while our solution uses a smooth surface in the front layer, the target surface and the back layer can have high-frequency variations. Thus some of the linearization strategies used in Claim 3 and Eq. (26) are not used when designing the back layer.

For every point x^o on the virtual surface or on the 2nd layer we can express reflection as

$$\mathbf{r}^r = M_s\mathbf{r}^o, \quad M_s = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 2s \\ 0 & 0 & 1 \end{pmatrix}, \quad (36)$$

where s is the local surface slope at x^o.

Let $\mathbf{r}^o = (x^o, \alpha, 1)^T$ be a ray hitting point x^o on the back layer. If we want the display to produce the same ray transfer as the virtual surface, we need to ensure that

$$A_v^b M_{s^v(f(x^o))}A_v^f(A_d^f)^{-1}\mathbf{r}^o = A_d^b M_{s_2^d(x^o)}A_d^f(A_d^f)^{-1}\mathbf{r}^o \quad (37)$$

where $s_2^d(x^o)$ denotes the slope of the back layer at x^o, and $s^v(f(x^o))$ the slope of the virtual surface at the matching point $f(x^o)$. A short calculation shows that this holds if we
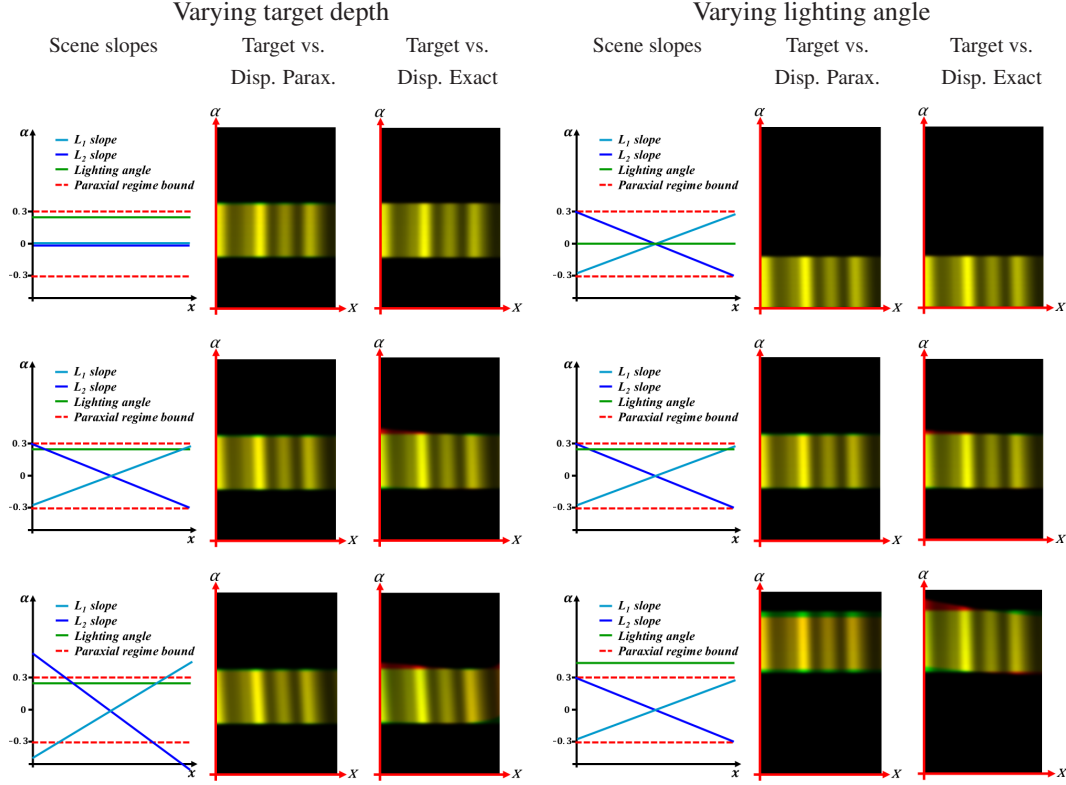
Figure 16: *The effect of the paraxial approximation: We used a planar target scene demonstrated at varying depths (left) and illumination angles (right). In each case we show the slopes of the display, and the lightfield resulting from spatially uniform illumination. To emphasize rendering errors, we overlay the desired lightfield (green channel) with the one produced by the display (red channel) assuming paraxial or exact ray tracing. When the depth of the target plane is close to the 2nd layer, the slopes of the display surfaces is low and the paraxial model agrees with the output of exact ray tracing. As the target depth increases, slopes increase and exact ray tracing differs from the target (see red rim at the edge of the illuminated range). Similarly, errors are observed even if the depth of the target plane is reasonable, but the illumination angle exceeds the range of the paraxial angles.*

select the local slope as

$$s_2^d(x^o) = \frac{c^v}{D} s^v(f(x^o)) - p_2^d x_o + t_2^d, \qquad (38)$$

with the p_2^d, t_2^d values of Eq. (35).

To complete the proof, we need to plug into Eq. (38) the matching point $f(x^o)$ provided in Eq. (33). We note that scaling the x coordinate of z^v by c^v/D automatically scales its slopes. Thus, using Eqs. (33) and (38), we reach the desired result in Eq. (34). $\square$

B. Appendix: Paraxial approximation

An important simplification made in this paper is the analysis of ray reflection/ refraction using the paraxial approximation. This approximation is valid for small angles at which we can use $\sin(\theta) \approx \theta$ and $cos(\theta) = 1$. In Figure 16 we test the validity of this approximation. We compare here

exact ray tracing with ray tracing using the paraxial model, both with our display and on the target surface. We see that the paraxial and the exact models agree for angles up to 20^o degrees. In the synthetic case we target a viewing/ illumination range of $[-12^o, 12^o]$ degrees, which is well within the scope of the paraxial approximation, and a sufficiently reasonable range for a multiscopic display. The SLM experiments in the paper, however, support a much lower range due to the limited SLM pitch.

C. Appendix: Wave optics analysis of SLMs

C.1. Spatial Light Modulators (SLMs)

As mentioned, in this paper we used the liquid crystal on silicon "PLUTO" SLM offered by Holoeye Inc. It consists of 1920×1080 refractive liquid crystal cells with a pitch of $\Delta_o = 8\mu m$ over a reflective silicon background layer. When used with polarized light, the cells allow computer

control over the index of refraction $\eta_\lambda(x)$ in each individual pixel and these refractive indices can be adjusted dynamically at a speed of 60Hz.

Since the cells are so small the reflected energy should be analyzed using wave optics tools. We briefly review here the basic results, and refer the reader to [36, 16] for further derivation. Our main goal in this review is to summarize how one can approximately convert the wave optics derivation to the geometric optics terminology used in this paper.

Wave optics derivation of reflectance: We denote the lighting and viewing directions (both unit vectors) by $\mathbf{l}$, $\mathbf{v}$ respectively, and the half vector by $\mathbf{h} = (\mathbf{h_x}, \mathbf{h_z})^T = (\mathbf{l} + \mathbf{v})/2$ (note that $\mathbf{h}$ here is not normalized). An incident light is described as a complex wave. Due to surface micro-structure, this wave accumulates phase changes as it interacts with the surface, reshaping the scattered wave field. We denote the local phase change at spatial surface position x as $\phi(x)$. For a reflective surface as in [36], these phase changes are achieved as a result of small variations in surface height, so $\phi(x) = \left(\frac{2\pi}{\lambda}\right) 2z(x)$. A liquid crystal display (LCD) can digitally adjust the refractive index $\eta_\lambda(x)$ of each cell, and can thus achieve similar phase variations $\phi(x) = 4\pi d_o \eta_\lambda(x)/\lambda$, where d_o is the thickness of the liquid crystal layer.

We denote the surface modulation by $\mathfrak{a}(x, \lambda)$, as a function of spatial position and wavelength:

$$\mathfrak{a}(x, \lambda) = e^{-\frac{i2\pi}{\lambda}\mathbf{h_z}\phi(x)} \qquad (39)$$

The basic result, which can be found in many standard textbooks [17], states that for coherent monochromatic illumination and for small incident angles, the reflectance is proportional to the Fourier power spectrum of the modulation function $\mathfrak{a}(x)$, denoted here by $\mathfrak{A}(\omega_x)$, at a spatial frequency determined by the half vector:

$$R_c(\mathbf{h}, \lambda) \propto \left| \mathfrak{A}\left(\frac{2\mathbf{h_x}}{\lambda}\right) \right|^2. \qquad (40)$$

To connect this result with the standard geometric optics derivation, consider a slanted flat-mirror surface, which is described by a perfectly linear phase function $\phi(x) = s \cdot x$. The modulation function $\mathfrak{a}(x)$ is then a single-frequency sinusoid, and the resulting Fourier transform is a delta at a single direction, implying that all the energy is reflected toward a single direction. A simple calculation can show that this direction is indeed the direction of mirror reflectance, as we would expect from a mirror.

Natural Illumination: The main difference between coherent laser sources and natural illumination sources is that natural sources have finite area and thus produce illumination over a range of directions. We denote here by Δ_a the solid angle that a (far-field) light source, centered at direction $\mathbf{l}$, subtends relative to the surface.

The spectral reflectance from a spatially-incoherent source is the averaged coherent *intensity* over all illumination directions within its solid angle:

$$R_{ic}(\mathbf{h}, \lambda) = \left| \mathfrak{A}_\lambda\left(\frac{2\mathbf{h_x}}{\lambda}\right) \right|^2 \otimes \Pi_{\Delta_a/\lambda} \qquad (41)$$

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

Using the convolution theorem, it can be shown that the total reflectance from a surface dot with area Δ_d under spatially-incoherent illumination, is the average of the Fourier power spectra computed over all local windows of size $\Delta_c \triangleq \lambda\Delta_a^{-1}$ within the dot:

$$R_{ic}(\mathbf{h}, \lambda) \propto \int_{x_0 \in \Delta_d} \left| \int \mathfrak{a}_\lambda(x - x_0) W_{\Delta_c}(x) e^{-\frac{2\pi i}{\lambda}(2\mathbf{h_x} x)} dx \right|^2 dx_0. \qquad (42)$$

Here, W_{Δ_c} is a sinc window of width Δ_c. 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.

C.2. Implementing optical surfaces with SLMs

We can use the SLM to implement an optical surface of interest as derived in Sec. 4. We simply sample the surface at intervals of $\Delta_o = 8\mu m$ and set the SLM phase of a pixel x as $\phi(x) = 4\pi d_o \eta_\lambda(x)/\lambda = 4\pi z(x)/\lambda$. It has been empirically observed by [16] that using this conversion between geometric surfaces to waves will produce reasonable results provided that the slopes of the optical surfaces are smooth over windows of size Δ_c; that is, the surface is smooth over the small coherent units at which wave interference occurs. To account for this requirement, we selected the width Δ_d of the BRDF bumps we added to be a few times wider than the spatial coherence Δ_c. For example, an illumination angle $\Delta_a = 0.5^o - 1^o$, which is the angle subtended by a $5cm$ source placed $3 - 6m$ from the display, produces a coherence length of $\Delta_c = 30 - 60\mu m$. The bumps we added occupied 20×20 or 10×10 SLM units, resulting in a dot area of $80 - 160\mu m^1$. To reassure the model of [16], we include below a short simulation of the geometric vs. wave optics model predictions.

---

[1]This calculation also emphasizes that standard LCD displays whose pitch is in the order of $100\mu m$, essentially have a pitch larger than the coherence length of a typical natural illumination. If the display phase and intensity are uniform over Δ_c windows, wave interference does not apply and they cannot produce any angularly-varying illumination effects. Thus, wide LCD cells can only attenuate ray intensity but they cannot bend its direction.

The SLM implementation suffers from two main restrictions, on angular range and thickness. The first restriction results from the fact that since the SLM has a limited resolution ($\Delta_o = 8\mu m$ pitch), it cannot display surfaces with higher frequency. For example, consider a slanted surface of the form $z(x) = s \cdot x$. It results in the phase $\phi(x) = 4\pi s x/\lambda$. The SLM should present the sinusoid $e^{i\phi(x)}$, sampled at Δ_o intervals. From the Nyquist theorem, to represent a sinusoid reliably we must sample it with spacing below its frequency, implying that the maximal surface slope s that we can present is limited by the pitch of the SLM, $|s| \leq S_M^\lambda$:

$$S_M^\lambda = \frac{\lambda}{4\Delta_o}. \qquad (43)$$

For the red illumination at $\lambda = 633nm$, this results in $S_M = 0.02 = 1.13^o$. Restricting surface normals to $|s| \leq S_M$ also implies that the display can be viewed only when $|\mathbf{h_x}| < S_M$ or $|\mathbf{l_x} + \mathbf{v_x}| < 2S_M$. In practice, the experiments in Sec. 6.2 are designed to support an angular range of $\Omega_M = 2S_M$. That is, it supports illumination and viewing directions between -2.25^o to 2.25^o, an overall range of 4.5^o. Since a smaller cell pitch translates directly into a larger angular range, this limitation will soften substantially as hardware improves. Devices with $6mm$ cell pitch are already commercially available (LETO, Holoeye), and devices with $2mm$ cell pitch are in development [27].

A second restriction of SLMs is that they have limited thickness. Therefore, rather than presenting $\phi(x) = 4\pi z(x)/\lambda$, we present $\phi(x) = \mod(4\pi z(x)/\lambda, T)$, where T is the maximal phase derived from the thickness of the SLMs. Most SLMs enable a phase range of up to 2π, but we chose an SLM allowing a higher phase retardation of up to 6π at $\lambda = 700nm$ (and a higher range at lower wavelengths). While the chopping to the height function modulo T results in exactly the same phase from the wave optics perspective, the main problem introduced is that the clipping is a function of wavelength, causing some artifacts when the SLM is viewed under natural polychromatic illumination. Figure 7 illustrates the SLM assignment $\phi(x)$, corresponding to the display surfaces of Figure 3.

In our implementation, the SLMs have 1920×1080 liquid crystal cells of pitch $\Delta_o = 8\mu m$. We display surfaces with 192×108 or 96×54 grid dots and in correspondence, use 10×10 or 20×20 cells to sample the BRDF at each dot. Overall this matches to a spatial resolution of $\Delta_d = 80\mu m$ or $\Delta_d = 160\mu m$, which is below the minimal resolution observed by a human standing a reasonable distance from the surface. The layer size is $W \times H = 15.4mm \times 8.6mm$.

Geometric vs. wave optics simulation: In Figure 17 we simulate the accuracy of the geometric optics approximation. We compare the geometric optics and the wave optics predictions of surface reflectance, for a quadratic surface whose BRDF is uniform over an angular range $\beta = 0.02$

(see definition in Eq. (6)). Note that since the basic surface is quadratic, the normal orientations vary linearly with x. Thus every local spatial neighborhood spreads light toward a different set of directions determined by the basic orientation of the surface normal plus the BRDF spread. The basic surface orientation results in a slanted illumination area in the lightfields of Fig. 17.

We computed the lightfield resulting from a wave optics reflectance assuming collimated incoming illumination at angle $\alpha_{in} = 0^o$ and angular resolution of $\Delta_a = 1^o$. A wave optics prediction of the same continuous surface relatively agrees with the geometric prediction (Fig. 17(b)). However, in practice an SLM realizes a quantized approximation to the desired continuous surface, and the surface is also taken modulo the SLM thickness. Different quantizations are evaluated in Fig. 17(c-e). The $\Delta_o = 8\mu m$ used in this paper (Fig. 17(d)) provides a reasonable approximation up to weak chromatic side lobes outside the angular viewing zoom of interest. A coarser quantization (e.g. Fig. 17(e)) completely loses the shape of the reflectance. Note that as discussed above, the applicable quantization resolution is a function of the angular range of interest, and the BRDF spread β in Figure 17 is selected according to Eq. (43) to match the $\Delta_o = 8\mu m$ SLM pitch quantization.
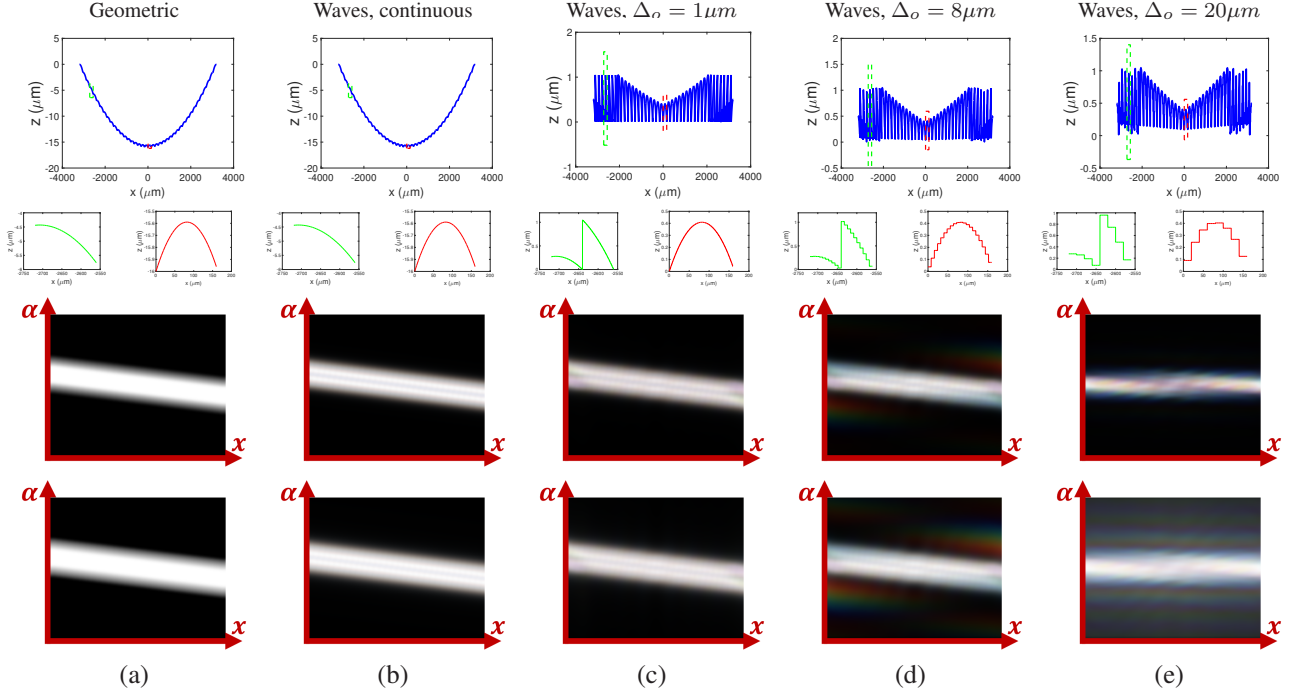
Figure 17: *Reflection with a wave optics model. We compare the geometric optics and the wave optics prediction of surface reflectance for a quadratic surface with uniform BRDF, with angular resolution $\Delta_a = 1^o$. (a) Geometric optics reflectance. (b) The wave optics prediction on the same continuous surface is relatively consistent with the geometric optics one. (c-e) In practice the surface displayed on the SLM is the continuous one: (i) modulo the SLM thickness and (ii) quantized to match SLM pitch. The quantization level affects the quality of the approximation. (c) $\Delta_o = 1\mu m$ provides an indistinguishable approximation of the continuous case. (d) The $\Delta_o = 8\mu m$ pitch used by our SLM provides a reasonable approximation apart from some weak chromatic lobes. (e) A coarser thresholding can completely destroy the shape of the reflectance. Top row: the input surface. 2nd row: zooming on two micro-texture units of the surface. 3rd row: the resulting lightfield in flatland assuming a collimated illumination at angle $\alpha_{in} = 0^o$. 4th row: gamma-amplified version of the lightfield revealing some of the side lobs.*