Exponentially-wide étendue displays using a tilting cascade

Sagi Monin¹ Aswin C. Sankaranarayanan² Anat Levin¹

Abstract—A fundamental limitation of spatial light modulation (SLM) devices is that their étendue, defined as the product of the display's size and angular range, is bounded by the number of pixel units. Current SLMs are woefully inadequate in meeting the spatial and angular ranges of many real applications. In particular using these SLMs to generate realistic computer-controlled holographic displays would require scaling the number of display units by a few orders of magnitude. In this work, we suggest that rather than excessively increasing the pixel count, étendue can be expanded by augmenting the display units with tilting capabilities. Furthermore, we show that tiltable displays can be realized using a cascade of binary tilt layers, each capable of tilting the light towards one of two orientations. With proper design, the étendue-expansion factor scales *exponentially* with the number of display layers; hence, a very small number of such layers can effectively realize wide expansion factors. We implement a proof of concept display and demonstrate its applicability for displaying multi-view or holographic content with increased size and angular field-of-view.

Index Terms—Spatial light modulation, Holographic displays, Etendue

1 INTRODUCTION

Spatial Light Modulators (SLMs) are one of the most promising ideas at the front of technology and science. Their applications range from 3D, augmented reality and holographic displays, to wavefront shaping techniques that remove aberrations caused either by atmospheric turbulence or by tissue layers. A typical SLM consists of an array of controllable/programmable units, each of which is capable of modulating the *phase* of incident light, rather than its intensity. This allows them to modulate a complex-valued incident wavefront in a programmable way.

As much as this idea is promising and offers enormous flexibility in the shaping of light, existing SLMs are extremely limited. A straightforward Nyquist analysis shows that the maximal angle at which SLMs can spread light is inversely bounded by their pixel pitch. Therefore, bending light at large angles invariably requires a pixel pitch that is significantly smaller than what is available by current SLM technology. While it is possible to demagnify pixels using lenses, this unavoidably reduces the image size as well. Further, it can be shown that any magnification or demagnification with lenses preserves the so-called *étendue* of the display, which is defined as the product of its spatial size and its maximal tilt angle. The étendue is fundamentally constrained only by the pixel count of the display. In the context of near-eye VR/AR displays the display area and maximal tilt angle limit the display eye-box and field of view (FoV). Supporting a human field of view of 120° and an eye-box of a centimeter, would require an array size of about $10^5 \times 10^5$ pixels [1]; unfortunately, this is three orders of magnitude above the capacity of current technology. Even

 sagim@campus.technion.ac.il, saswin@andrew.cmu.edu, anat.levin@ee.technion.ac.il if the hardware technology will ever meet the gap, the data buffers for such enormous arrays and the computational content creation challenges are far from being trivial. At the same time, reducing pixel pitch to support large tilt angles would result in resolution far beyond the limit of the human retina. Hence, expanding étendue by the direct approach of increasing the number of pixels may be excessive.

A recent class of techniques [2], [3], [4], [5], [6] expand the étendue of a display by inserting a fixed phase mask in front of the programmable SLM; where the pitch of the phase mask is significantly *smaller* than that of the SLM. The higher resolution mask allows the display to spread light over a wider angular range. However, as the mask is fixed it cannot adapt to image content and the amount of control we have on the wavefront is limited.

Here, we take a different approach to étendue expansion. Rather than excessively increasing the pixel count, we suggest augmenting the display with a simple capability— *-using programmable tiltable display units*. Classical SLMs display a roughly constant phase over the area of each pixel, providing a *piecewise constant* approximation to the desired phase surface. If each display pixel could tilt the wavefront, we can generate *piecewise linear* phase surfaces. The étendue of the display is then governed by the maximal tilt angle of these units, as larger tilt angles can allow us to spread light toward a wider range of directions. Moreover, we can decouple the strong dependence between the pixel pitch of the SLM and the maximal tilt angle it can produce.

Some prototypes for tiltable micro-mirror units are available on the market, e.g. [7]. A review of state-of-the-art technology is available by [8] and discussed in Sec. 2.3. However, current devices still do not offer a sufficiently large number of units with a wide non-binary range of tilt angles.

In this work, we implement a piecewise linear phase modulation by using existing technology components. Specifically, we combine a standard SLM implementing

^{• &}lt;sup>1</sup> Department of Electrical Engineering, Technion.

 ² Department of ECE, Carnegie Mellon University.

piecewise constant phase modulation, along with a cascade of binary-tilt units, optically aligned with each other using a relay system. Each layer in the cascade uses an liquid crystal display (LCD) array as a switch to select between vertical and horizontal polarization states, independently at each of its pixels.

Using polarization-sensitive gratings, we induce a local tilt to the wavefront by $+\theta$ or $-\theta$, depending on the polarization of light. By combining Nyquist analysis with simple binary coding algorithms, we show that the number of required panels scales only *logarithmically* with the desired angular expansion. Put another way, by combining *K* binary tilt layers we can expand the angular range of the display by an exponentially increasing factor of 2^{K} .

We build a proof of concept display with three layers and demonstrate its applicability in the design of multi-view and holographic displays, with increased size and angular fieldof-view (FoV).

Beyond our specific implementation, we hope that this paper motivates the development of tiltable micro-mirror technology, as an alternative route for improving SLM technology and expanding their étendue.

2 RELATED WORK

2.1 SLM applications

SLMs have many applications in multiple fields of science. While most of these applications would benefit from larger étendue, many of them can make good usage of piecewise linear modulation with large tilt angles, and do not really require a larger number of pixel units.

One of the popular applications of SLMs is holographic displays. The SLM is used to display a wavefront that will generate different images based on different positions and focus conditions of the viewer's eye. Holography is one of the leading approaches for the design of near-eye displays [9], due to the compact form factor it offers. In a near-eye display, the SLM plane is placed right in front of the viewer's eye, and as the viewer focuses much further away, he will see the Fourier/Fresnel propagation of the SLM wavefront to another plane. Assuming for simplicity that the viewer's eye focuses at infinity, he will see the Fourier transform of the phase pattern on the SLM plane [10], [11], [12], [13] and we broadly refer to this class of algorithms as Fourier domain displays.

In an alternative operation mode of holograms [1], [14], [15], [16], [17], [18] the viewer focuses directly on the display or within a small distance from it. The SLM is used to display hogel units generating a light field, or multiview content such that a viewer observes different images from different viewpoints. Below we broadly refer to these algorithms as image domain displays.

In multi-focal displays [19], three-dimensional content is shown to a user by placing objects on a discrete number of different focal planes, which are optically placed at different depths from the viewer, using e.g. a focus tunable lens. The main shortcoming of this idea is that the different focal planes are made transparent. As a result, next to depth discontinuous one does not get occlusion cues, and rather, depending on the viewing direction, content from different depth layers may overlap. Recently Chang *et al.* [20] suggested a simple way to fix this, using a small pixel dependent tilt of the light cone in each image position. This tilt is implemented by co-locating a phase based SLM with a standard intensity display. As with previous SLM based systems, the maximal tilt angle the SLM could realize largely limits the light cone supported by this display. We note that while the display requires larger tilt angles, the slopes involved vary rather smoothly and agree well with the piecewise linear phase model of this paper.

Focal Surface Displays [21] use an SLM inside a near-eye multi-focal display to adaptively vary the focal power over the image area.

The light arriving at the SLM plane is not necessarily generated by some back illumination of the display, but can utilize any form of energy received from the environment as a light sensitive display [14], [22]. These displays aim to mimic the BRDF and other light sensitive variations of materials in the real world. In this work, the phase masks design draws inspiration from geometric optics micro-facet [23], [24] surface models. As a result, it leads to piecewise linear phase surfaces rather than the highly varying phase patterns produced by general purpose phase retrieval algorithms. As such, they form another good application case for our advocated piecewise linear phase display.

SLMs were also used in [25], [26] to redistribute energy on a display plane and generate high dynamic range displays. While a standard display simply blocks light at pixels that correspond to dark regions in the image, this display uses SLM light steering abilities to redirect light from dark parts of the image plane toward bright parts, allowing for a better contrast between bright and dark parts and better usage of the display energy power.

Last but not least, one of the most important applications of SLMs is the ability to reshape a wavefront to correct for optical aberrations. This idea is applied in almost any field of science from astronomy to microscopy. In astronomy, the light arriving from far planets is aberrated due to the index of refraction variations on its way through the atmosphere. Assuming this aberration can be measured, an SLM is used to bend light rays before they are collected by the detector, to undo the atmospheric aberration and achieve a sharper image [27], [28]. At the other end of the scale SLMs have also been successfully used in microscopy to undo aberrations caused by index of refraction variations in tissue and other biological subjects [29].

2.2 Overcoming étendue limits

Previous approaches dealing with the limited étendue of holographic near-eye displays use a tracking system to track the position of the viewer pupil inside the eye box and steer the hologram towards it [30], [31], [32], [33]. A Similar approach is also used in non holographic near eye displays such as in [34] where they use pupil tracking and steer the image to achieve a foveated display. Our approach, in contrast, seeks to expand the étendue of the actual display.

Alternatively, static phase masks whose resolution is higher than the SLM have been attached to the display [2], [3], [4], [5], [6]. SLM content is optimized subject to the higher scattering angles of the phase mask, to produce wider content. While this can largely extend étendue, the fixed phase mask does not allow full control over the output. Also, accurate calibration and alignment of the mask propose non-trivial challenges. Our display extends over these masks as it offers programmable capabilities.

In [35] a static diffractive optical element is placed in front of the eye that has the structure of a periodic lens array. As such it creates multiple copies of the eye-box allowing the hologram to be viewed from multiple eye positions. However, the copies replicate the same hologram and do not allow the display of viewpoint dependent effects.

2.3 Tiltable micro-mirror technology

Binary tiltiable digital micro-mirror arrays form a mature technology, widely used in projectors, but binary states do not support the continuous tilts required for étendueexpanded displays. There have been attempts to combine the micro-mirrors with a fast laser that will allow using intermediate tilt angles [36], but such synchronization is not trivial.

There are many instances of research-grade prototypes that have shown the feasibility of programmable tilt micromirrors [37], [38] supporting continuous tilts. Most of these mount a micro-mirror on top of micro-actuators, which produce tilts and displacements using electrostatic, thermal or piezo effects. An alternate approach is to fabricate micromirror arrays using liquid metal pivots; here, a micro-mirror is bonded to a liquid metal droplet whose shape is changed via electrowetting [39]. This allows for the micro-mirror to be subjected to both analog tilts and z-displacements. However, most prototypes mentioned in prior work have a limited number of micro-mirror units, which does not meet holography requirements. The best commercially-available device we have tracked is offered by the Fraunhofer Inst. [7], providing analog 2D tilts with an array of 512×320 micro-mirrors.

To motivate the applicability of tiltable micro-mirror displays, in this paper we implement a proof-of-concept tiltable display using a cascade of commercially available SLMs. Our construction approximates the desired tiltable-array, as it can achieve a wide range of tilts but only supports a quantized set of angles in the desired range.

3 ÉTENDUE AND ITS EXPANSION

Notation. Below we use the notation $\vec{x} = (x, y)$ to denote a 2D position, and $\vec{\theta}$ to denote the first two coordinates of a direction (unit norm 3D) vector. For example a tilt in angle θ along the *x* dimension corresponds to a direction vector $(\sin(\theta), 0, \cos(\theta))$. For the small angles we target in this paper the paraxial approximation implies $\sin(\theta) \approx \theta$, $\cos(\theta) \approx 1$, and this tilt would be represented using the vector $\vec{\theta} = (\theta, 0)$.

3.1 SLM and étendue limits

We start with some background on SLM operation and explain the étendue limits. We refer the reader to [1] for a detailed derivation of these principles.

SLMs are arrays of pixels, such that the phase of each unit can be computer controlled, allowing the display of programmable phase masks. However, as the phase is



Fig. 1: Étendue expansion using tilts: (a) A normal SLM can tilt an incoming plane wave by angles in the range $[-0.5\Omega_{\theta}, 0.5\Omega_{\theta}]$. (b) If the incoming wave is tilted at angle θ it can be tilted by the SLM to the range $[\theta - 0.5\Omega_{\theta}, \theta + 0.5\Omega_{\theta}]$. (c) If the incoming wave is a combination of slopes in the range $[-0.5\Omega_S, 0.5\Omega_S]$, after passing through the SLM we can generate angles in the range $[-0.5\Omega_S - 0.5\Omega_{\theta}, 0.5\Omega_S + 0.5\Omega_{\theta}]$. In the illustration green and blue mark different slopes.

roughly constant within each pixel unit, below we model the resulting phase mask as a *piecewise-constant* function $\phi_c(\lfloor \vec{x} \rfloor)$, with $\lfloor \vec{x} \rfloor$ denoting the rounding of point \vec{x} to the nearest grid unit. We denote the SLM pixel pitch with Δ_x . A display with $\mathcal{N} \times \mathcal{N}$ units has a width and height of

$$\Omega_x = \Delta_x \mathcal{N}.\tag{1}$$

To send light toward direction $\vec{\theta}$ we need to display on the SLM a sinusoid of the form $\phi_c(\vec{x}) = \frac{2\pi}{\lambda} \vec{\theta} \cdot \vec{x}$. An SLM can only display a piecewise constant approximation of this sinusoid. Thus Nyquist limit implies that realizable tilts are limited to the range $\vec{\theta} \in [-\frac{1}{2}\Omega_{\theta}, \frac{1}{2}\Omega_{\theta}] \times [-\frac{1}{2}\Omega_{\theta}, \frac{1}{2}\Omega_{\theta}]$ with

$$\Omega_{\theta} = \frac{\lambda}{\Delta_x}.$$
 (2)

The maximal tilt angle is inversely limited by the SLM pitch.

Below we use \mathcal{E} to denote the *étendue* of the display, quantifying the product of its spatial and angular ranges. Using Eqs. (1) and (2), we see that étendue is essentially a constant controlled only by the number of pixel units in the display (and the wavelength of light):

$$\boldsymbol{\mathcal{E}}^2 = \Omega_x^2 \cdot \Omega_\theta^2 = (\lambda \mathcal{N})^2. \tag{3}$$

One can use a relay system to generate a smaller copy of the SLM plane. As pixels are de-magnified the tilt angle increases. However, this unavoidably also scales down the image area. It is easy to show that the étendue of the display is maintained fixed when it is viewed through any lens system.

In the context of near-eye AR and VR displays, the spatial area of the display sets a limit on the eye-box, namely the range of positions where the viewer pupile can be positioned, and the maximal tilt angle limits the display field of view (FoV). The étendue of existing displays is extremely limited, and forms one of the largest barriers on the practical utility of digital holographic displays. Some calculations show that for the high end SLMs available on

the market today (pitch $\Delta = 3.74 \mu m$ and 4160×2464 pixels), spreading light over a wide angular range thus supporting a desired FoV of 120° , leads to a tiny image area (or eye-box) of 1mm. Supporting a display area of 1cm leads to a tiny FoV of 12° .

This state of affairs is unfortunate, as supporting both a wide tilt and a wide display area requires devices with a pixel count of three orders of magnitude higher than current technology. Even if technology will meet the gap, the data buffers and the content creation algorithms will pose nontrivial challenges. At the same time, such a large number of pixel units would result in an angular resolution which is way above what the human retina can resolve. Here we suggest that rather than excessively increasing the number of display units, étendue can be expanded using devices with different functionality.

3.2 Étendue with piecewise-linear phase modulation

This paper suggests to implement étendue expansion using piecewise-linear phase modulations, of the form

$$\phi(\vec{x}) = \phi_l(\vec{x}) + \phi_c(\vec{x}) = \vec{s}(\lfloor \vec{x} \rfloor) \cdot \vec{x} + \phi_c(\lfloor \vec{x} \rfloor), \quad (4)$$

where $\phi_c(\lfloor \vec{x} \rfloor)$ is piecewise constant as before, and $\phi_l(\vec{x}) = s(\lfloor \vec{x} \rfloor) \cdot \vec{x}$ is a *piecewise linear* function, where $\vec{s}(\lfloor \vec{x} \rfloor)$ is a 2D tilt vector $\vec{s} = (s_x, s_y)$ denoting the horizontal and vertical slope of pixel $\lfloor \vec{x} \rfloor$. Such a display effectively tilts the wavefront at each SLM pixel.

If the tilting slopes take the range $|s_x| \leq \frac{1}{2}\Omega_S, |s_y| \leq \frac{1}{2}\Omega_S$, and the piecewise constant phase can implement tilts in the range $[-\frac{1}{2}\Omega_{\theta}, \frac{1}{2}\Omega_{\theta}]$, their combination can spread light over the extended angular range:

$$\Omega_{\theta} = \Omega_{\theta} + \Omega_S, \tag{5}$$

see Fig. 1 for an illustration.

For the same display area this increases the étendue by a factor $q_{\mathcal{E}} \times q_{\mathcal{E}}$ with

$$q_{\mathcal{E}} = \frac{\mathcal{E}}{\mathcal{E}} = \frac{\Omega_{\theta}}{\Omega_{\theta}} = \frac{\Omega_S}{\Omega_{\theta}} + 1.$$
 (6)

4 BINARY TILT CASCADE

As mentioned in Sec. 2.3 tiltable-micro mirror technology is improving, but existing devices have a limited number of micro-mirror units, which is not yet sufficient for a rich holographic display. In this paper we offer to approximate a tiltable micro-mirror display by concatenating multiple layers that can tilt the light toward a binary set of directions. We start by analyzing binary tilts and characterizing the space of tilts they can support. Hardware realization is discussed in Sec. 5.

We define a binary tilt layer as one that provides the ability to programmably choose between two tilts at each pixel. Such a device introduces a phase function of the form

$$\phi_l(\vec{x}) = b(\lfloor \vec{x} \rfloor) \ \vec{s} \cdot \vec{x},\tag{7}$$

where $b(\lfloor \vec{x} \rfloor) \in \{-1, 1\}$ is a binary variable that allows for the selection of a ramp with slope $+\vec{s}$ or $-\vec{s}$. While these layers can take only a binary state, they can implement a linear ramp over the pixel area rather than a constant phase as the SLM.



Fig. 2: Tilting cascade schematic: The display is constructed as a cascade of panels, where the first one is a standard piecewise constant SLM with tilt range $[-\Omega_{\theta}/2, \Omega_{\theta}/2]$, and successive panels tilt a pixel unit toward 2 binary orientations. Panels are imaged on successive ones via a 4f system, and the total number of different tilt angles expressed through the cascade is *exponential* in the number of panels.

Setup. For simplicity, we start by considering tilts along the x dimension alone. That is, our tilt vectors will have the form $\vec{s} = (s, 0)$, where s is a scalar.

Our setup, illustrated in Fig. 2 is basically a cascade of displays. Each layer is imaged onto the next one via a pair of lenses forming a 4f system, all panels are transmissive and have the same pitch Δ_x . One of them is a standard SLM implementing a piecewise constant phase function. The rest are binary tilt layers. Each binary tilt panel realizes a different binary set of slopes, and we denote the slope of the *k*th panel as $\pm s^k$.

As successive panel layers are imaged on each other, the resulting phase is the sum of the phase at all layers of the display:

$$e^{\frac{2\pi i}{\lambda}\sum_{k=0}^{K}\phi^{k}(\vec{\boldsymbol{x}})} = e^{\frac{2\pi i}{\lambda}\left(\phi_{0}(\lfloor \vec{\boldsymbol{x}} \rfloor) + \sum_{k=1}^{K}b^{k}(\lfloor \vec{\boldsymbol{x}} \rfloor)\vec{\boldsymbol{s}}^{k}\cdot\vec{\boldsymbol{x}}\right)}.$$
 (8)

Although the order of the layers can be permuted, the notation of Eq. (8) assumes the first layer is implemented with a piecewise constant phase $\phi^0(\vec{x}) = \phi_c^0(\lfloor \vec{x} \rfloor)$, and successive layers implemented with binary tilts of slopes $\pm s^k$.

Optimal slope basis. Layers $1, \ldots, K$ of Eq. (8) provide a range of linear ramps, but clearly with a finite number of layers only a finite number of slopes can be expressed. However, with a proper choice of the basis tilts, given K layers we can express an exponentially large set of 2^{K} different slopes, and this is equivalent to converting integer numbers into a sequence of bits in a binary basis. Specifically, we select the slopes of the different layers with exponentially increasing tilt 'powers'.

Assume the first piecewise constant layer can span slopes in the range $[-\Omega_{\theta}/2, \Omega_{\theta}/2]$. Now, let us add a binary tilt layer of slope $s^1 = \pm \Omega_{\theta}/2$. Selecting $s^1 = \Omega_{\theta}/2$, and allowing for the first layer to set slopes in the range $[-\Omega_{\theta}/2, \Omega_{\theta}/2]$, we can generate any slope between $[0, \Omega_{\theta}]$; similarly, by selecting $s^1 = -\Omega_{\theta}/2$, we can generate the slopes $[-\Omega_{\theta}, 0]$. Thus the combination of a $s^1 = \pm \Omega_{\theta}/2$ binary tilt and a piecewise constant layer can express any tilt in the range $[-\Omega_{\theta}, \Omega_{\theta}]$. In the next binary layer we select $s^2 = \pm \Omega_{\theta}$, thereby expanding the range of slopes we can generate to $[-2\Omega_{\theta}, 2\Omega_{\theta}]$. Generalizing this idea, we set $s^k = 2^{k-2}\Omega_{\theta}$, which allows us to generate any slope in the range $[-2^{K-1}\Omega_{\theta}, 2^{K-1}\Omega_{\theta}]$. We note that increasing the slope powers exponentially is critical for spanning the widest range. If all *K* slopes had been set to be equal, their linear combinations would only span a range of 2K different values rather than 2^{K} different values.

How many panels are required? Suppose we want to express all slopes $|s| \leq \hat{\Omega}_{\theta}$ using displays of pitch Δ_x . As the native angular range of a piecewise constant panel of this pitch is $\Omega_{\theta} = \lambda/\Delta_x$, we need $\log_2(\hat{\Omega}_{\theta}/\Omega_{\theta}) = \log_2(\hat{\Omega}_{\theta}\Delta_x/\lambda)$ panels to expand the slope in one axis (e.g. horizontal). For example if $\Delta_x = 4\mu m, \lambda = 0.5\mu m$ and the native angular range of a piecewise constant display is $\Omega_{\theta} = 1/16$, to expand it to an angular range of $\hat{\Omega}_{\theta} = 0.5$ we need $\log(8) = 3$ layers in each axis.

2D tilting. While for the simplicity of the analysis the previous discussion considered 1D tilts, extending the idea to two dimensional tilts is straightforward, and we use *K* layers of horizontal tilts and another *K* layers of vertical tilts. E.g. to express all 2D slopes in the range $|s| \leq \hat{\Omega}_{\theta}$, in both horizontal and vertical axes, we need $2 \log_2(\hat{\Omega}_{\theta} \Delta_x / \lambda)$ panels.

Linear vs. exponential expansion. To emphasize the advantage of the binary tilt cascade, it worth considering other strategies for expanding étendue by combining K standard piecewise constant SLMs. For example we could place a few displays panels one next to each other, increasing the spatial display area. Alternatively, we could concatenate Kstandard SLMs via 4f relay and introduce sub-pixel shifts between them to achieve a higher pitch. However, the expansion in étendue obtained by such constructions is only *linear* in the number of panels. In contrast by using a binary tilt with different angles s^k in each layer, we expand the tilt range by an exponential factor.

5 HARDWARE CONSTRUCTION

One way to implement binary tilt layers in hardware would be using binary digital micro-mirror devices that are commonly found in projectors. However, these only support a fixed tilt angle of $\pm 12^{\circ}$ and we cannot customize them for our needs.

Here we implement binary tilting using Liquid Crystal Display (LCD) panels. LCDs are one of the common devices used to realize SLMs. As we review in App. A.1, they have been used to modulate either intensity or phase, by adjusting the arrangement and orientation of polarizes before/after the LCD panels. However, here we exploit their birefringence properties in a different way, and use them to rotate the polarization of light. Given independent control of polarization in each pixel, polarization sensitive grating is used to split the light into two different tilt angles.

5.1 Implementing binary tilts using polarizationsensitive optics

To implement binary tilting we use polarization sensitive grating (Edmund optics 16-590). These surfaces tilt an incident light by θ or $-\theta$ degrees, depending on its polarization. By combining it with the programmable polarization created by the LCD, we can tilt different pixels in different



Fig. 3: Tilting rays by Fourier shifting: We use two polarizationsensitive gratings to shift the rays in the Fourier-plane, which introduces a tilt after the second lens. The colors, yellow and orange, show the two different polarization states and the paths they take after the grating layers. The shaded colors show the virtual path from which each polarized state appears to originate. The spacing between the gratings controls the tilts the rays experience. The LCD at the input image plane selects a polarization state, and this polarization leads to an independent tilt of light at each pixel.

directions. To maintain a display with a thin form factor, ideally we want to attach all LCD panels and grating layers to each other. However, this would require grating layers with custom tilt angles. To work with off-the-shelf components, we used polarization gratings with a fixed angle of $\pm 5^{\circ}$.

We used a 4f relay to image one panel on the next one and place two grating in the Fourier plane of the relay system. Our construction is based on the observation that to tilt rays, we need to shift them in the Fourier plane. As illustrated in Fig. 3, a set of rays emerging from a pixel generates a collimated beam in the Fourier plane. By independently controlling the LCD polarization in that pixel we can adjust the polarization of that beam. The first grating layer tilts this beam left or right depending on its polarization and the second undoes the tilt, bringing it back to its original direction. As the original direction is preserved the rays will focus back into the same pixel at distance f after the 2nd lens of the relay. However, the two grating layers generate a tilt. While passing the gap between the grating layers, all rays shifted by a small amount left or right, depending on their original polarization. As a result of this shift, they appear as if they have originated from a different point on the Fourier plane. Recall that rays emerging from each point on the Fourier plane have an equivalent orientation after the 2nd lens of the 4f system. Hence the shift of the rays passed between the grating layers is translated into tilt as they refract via the 2nd lens of the 4f system. Varying the distance between the grating layers controls the tilt angle applied in this cascade layer.

5.2 Prototype

A schematic of our setup is available in Fig. 4, and a component list is provided in App. A.3. Our prototype implements a $\times 4$ étendue expansion along the horizontal dimension alone. Tilting in the vertical axis would require rotating the grating layers by 90°. The prototype uses a sequence of three LCD panels. Two of them were combined with a polarization sensitive grating to implement binary tilt ($\times 4$ horizontal expansion), and the last one was used in the standard operation mode of LCDs, as a simple piecewise



Fig. 4: System: A collimated laser passes through two angle tilt layers and then illuminates a final SLM used as either a conventional piecewise-constant phase modulator. L1-L6 - lenses, P - polarizer, A - aperture, QWP - quarter-wave plate retarder, G - grating, and BS - beam-splitter. Note that in the setup image we only show the Fourier plane camera. The colors mark the different parts of the system.

constant phase modulator. Note that unlike the simplified description in Sec. 4, in our implementation the piecewise constant modulation layer is the last rather than the first in the cascade.

We chose *transmissive* LCD panels (Holoeye SLM LC 2012) over common reflective LCDs, first due to their affordable price and second, since a transmissive cascade is simpler to set and align compared to a reflective one. These transmissive panels are of lower quality compared to advance reflective technology, consisting of only 1024×768 units with a $36\mu m$ pitch. Due to increasing peripheral aberrations, we used only the central area of size 768×768 pixels. The transmissive panels do not fully obey the simplified Jones matrix model and phase modulation is coupled with intensity modulation. Also, with a 532nm illumination they support a phase range of $[0, 1.8\pi]$ rather than $[0, 2\pi]$.

The LCD swaps between orthogonal linear polarization states while the grating separates according to orthogonal circular polarization. To this end, we used quarter waveplate retarders between the LCDs and grating.

After the 2nd binary layer we use a polarizer to merge the two polarization states, by projecting both polarization states into a single linear polarization. This linear polarization is used to illuminate the piecewise constant phase SLM layer.

At the cascade output we used two cameras, one focused on the last panel imaging its content, and one imaging the Fourier transform of the phase SLM output. These two cameras serve us for demonstrating the operation of the display in different modes.

Calibration. For good results we had to carefully calibrate various imperfections in our system. Even after calibration, various unmodeled optical imperfections keep limiting the quality of our results. Our calibration follows standard ideas presented in [10], [40], [41], [42], [43]. The first step is calibrating the relationship between the voltage applied in a pixel and the resulting phase and amplitude. The second issue has to do with the fact that while the 4f relay is supposed to image one SLM plane on the next one, real optics does not fully follow the idealized model. We account

for it using a combination of phase distortion and geometric distortion, as detailed in App. A.2.

In particular, as discussed in App. A.2, optical aberrations in our prototype cause different tilt angles to accumulate slightly different shifts through the optical path, and pixels implementing different tilts overlay each other in the final SLM plane, see Fig. 11. As a result, trying to tilt every pixel in our prototype independently in a different direction leads to severe overlaps at the final plane, causing various artifacts. Therefore, in our experimental demonstration below we have only used phase masks ϕ_l that assign neighboring pixels similar slopes.

6 RESULTS

We demonstrate the operation of our SLM system in two display modes which we broadly term Fourier and image domain display.

6.1 Fourier display

In near-eye displays, the SLM is usually placed close to the viewer's eye. The viewer focuses on a plane much beyond the SLM plane, hence what the viewer will see is a hologram corresponding to the propagation of the wavefront presented on the display to some far plane. To simplify the discussion we assume the viewer focuses his eye at infinity and hence sees the Fourier transformation of the presented wavefront.

To generate content for a Fourier hologram we seek a phase mask defined as in Eq. (4) as $\phi(x) = \phi_c(\vec{x}) + \phi_l(\vec{x})$. The Fourier transform of this phase mask should correspond to a desired target image. As in [2] we minimize the following cost:

$$\underset{\phi_c}{\operatorname{argmin}} \left\| g * \left| \mathcal{F} \left(e^{i(\phi_c(\vec{x}) + \phi_l(\vec{x}))} \right) \right|^2 - g * I \right\|^2, \qquad (9)$$

where g is a low pass filter encoding the fact that the angular resolution of the hologram is often higher than what the human retina can resolve. As explained in [44] this allows



Fig. 5: Display capture: We used our display with piecewise linear slope patches of 16×16 and 128×128 pixels. While we cannot implement fully random masks [2] we also compared against it in simulation. Big patches allow for better contrast in simulation, and the random phase masks lead to the lowest contrast. In practice the contrast of small patches is even lower than the simulation prediction due to the challenge in accurately calibrating phase modulators. For the snake image we present a cross-section (red, blue, and magenta lines), showing the contrast difference of the original simulation, and the even larger contrast difference in the real capture. Yellow box on the target image marks the original range of the display. Next to each image we visualize s_x , the slope (tilt) assignment in different display pixels, see definition in Eq. (4). Mountain source image by r-z (CC BY 2.0).

finding better solutions to the above optimization problem, since by reducing the resolution one can match the number of entries in the target to the number of degrees of freedom (the number of programmable pixels) in the display.

While we can jointly optimize both the piecewise constant and the piecewise linear parts, optimizing the slopes while accounting for all system aberrations is challenging and subject to many local minima. In most examples below we set the piecewise linear part ϕ_l using application specific rules, and optimize only the piecewise constant part ϕ_c . We solve Eq. (9) with L-BFGS solver [45], with the same Butterworth low-pass filter used in [2]. Average unoptimized code run time is 90 seconds.

We start by testing our display under generic slope assignments, where we divide the SLM into patches of size $q \times q$ and randomly assign each patch one of the discrete sets of slopes supported by our prototype. We optimize the piecewise constant phase component subject to these slopes.

In [44] the effect of the patch size q is analyzed in detail, and they explain that small patches result in low contrast holograms, while the contrast can be significantly improved when using big patches of consistent slopes. This is due to the fact that very small patches act as a highly random filter, spreading significant speckle noise throughout the spectrum. The low-pass filter in Eq. (9), filters the speckles, reducing contrast in the resulting images.

Our display can help validate this analysis in a physical system. In Fig. 5 we compared results as a function of the patch size. We evaluate both a synthetic simulation of the image formation following the model of Eq. (9), as well as real capture from the setup. In all cases, the lowest reconstruction error was obtained with big patches. While we did not have a physical random mask as in [2], we have simulated its expected performance. This is conceptually similar to patch size q = 1, and indeed the high randomness of this mask led to lower contrast. For very sparse images such as the snake image, big patches lead to significant improvement in contrast. While all results in this paper are normalized to have maximal intensity 1, for the snake example we also show a cross-section along one column in the image, comparing the contrast of the small and big patches before normalization. This shows that the big patches achieve better light efficiency (higher maximal values at the desired boundary point) as well as reduce background clutter.

When capturing real images of the display the difference between big patches and small ones is even more drastic than in simulation. The reason is that despite calibrating and modeling the aberrations of our display, there are still deviations between the simulated light propagation and the one the display performs in practice. Bigger patches are less sensitive to these imperfections. We note that good calibration of the optics is one of the bigger challenges of previous étendue expansion approaches that used a fixed phase mask [2], [5], [6], leading to a big discrepancy between simulation and practice. The usage of big patches largely alleviates calibration challenges.

As studied in [44], in practice one should use patches of intermediate sizes, since very big patches do not generate content over the full eye-box. As our slopes are programmable, we can also use time multiplexing to permute the part of the display implementing different tilt angles.

To eliminate speckle artifacts, the display capture results in Figs. 5–7 average 10 different holograms optimized to display the same target. In App. A.6 we compare this with a capture of a single hologram. Also, as SLM modulation is imperfect there is usually residual diffraction in the DC direction. To eliminate it we add a vertical linear ramp on the displayed phase and tilt the displayed image away from the DC direction.

Content-dependent dynamic range. Another advantage of our programmable slope assignment is that more light can be sent to certain regions of the display, depending on the desired content. For that, rather than assigning the slope of each $q \times q$ patch uniformly, we assign them according to the histogram of energy in different regions of the target image. For example, if the target image has 90% of its energy on the left part and 10% of its energy on the right part, 90% of the patches would display a left tilt while only 10% of them would display a right tilt. Fig. 6 and App. Figs. 14 and 15 illustrate examples where the energy of the target is not uniform. Our display can achieve higher contrast and reduced artifacts, when compared with fixed slope masks that spread an equal amount of light in all directions.

Focal stacks. Fig. 7 demonstrates the display of a 3D target, with 2 objects in two different depth planes. As explained in App. A.4, the hologram is optimized to generate the desired focal stack image at each focus depth, and the image demonstrates capturing the same hologram while varying the camera focus between these planes.

6.2 Image domain display

In near-eye display implementations most of the étendue "budget" is used to display wide 2D images without supporting a very wide eye-box. That is, the display does not show view-dependent images when the eye position shifts. An alternative mode of usage for SLMs is to display multi-view content, or a light field. In particular, multi stereoscopic displays are often located far enough from the viewer's eye so that the viewer focuses at the display plane rather than on its Fourier transform. In these applications, the display should send different images towards different viewing directions. A lot of the research efforts were devoted to the display of such multi-view content, at the price of reducing the spatial image resolution.

Below we demonstrate the applicability of our device for displaying such multi-view content, using time multiplexing and spatial multiplexing.

Time multiplexing: The simplest way to display view dependent content is to time multiplex the desired output. The

first two layers are used to tilt the display output toward one of 4 viewpoints, where the entire area of these layers displays a single slope. In each time frame, the entire area of the last LCD panel is used to display a desired image into one direction, and we loop through the different viewing directions. In this mode we use the last LCD panel as a standard intensity display rather than a phase modulator, which we achieve by adding a linear polarizer at the output. Fig. 8 visualizes the output of this mode, showing that there is very minimal leaking of content directed to one viewpoint toward the other viewpoints.

This approach resembles attempts to use timemultiplexing for multi-view displays by rotating the entire display plane, in synchrony with a projector displaying images of different viewpoints [46], [47].

Spatial multiplexing: An alternative approach for displaying multi-view content is to multiplex the SLM area trading the multi-view content for spatial image resolution.

To this end, we divide the angular range of the display $\widehat{\Omega}_{\theta}$ into K bins and aim to display K different images $I_1, \ldots I_k$ toward K different viewing directions. In practice we display only K-1 independent images and use the last one to collect residual energy $I_K(\vec{x}) = 1 - \sum_{k=1}^{K-1} I_k(\vec{x})$. We denote $e(\vec{x}) = [I_1(\vec{x}), \ldots I_K(\vec{x})]$ the probability vector defining the desired angular distribution of energy emitted from display pixel \vec{x} .

The display content is usually generated using the hogel principle [1], [14], [15], [16], [17], [18]. This assumes the minimal unit that a viewer can resolve is of size $q \times q$ SLM pixels. Consider the $q \times q$ bin around pixel \vec{x}_o . We want to use it to display a hogel, namely diffraction pattern that will propagate a certain amount of light to each of the K viewing directions, as defined by the target vector $e(\hat{\mathbf{x}}_{o})$. While multiple strategies to design this content were proposed, we follow the scheme of [14], who design hogel content based on geometric optics micro-facet principles [23], [24]. By geometric optics principles, to send light to the kth direction, the hogel should simply display a ramp $\phi(\vec{x}) = \vec{\theta} \cdot \vec{x}$ at the corresponding angle $\vec{\theta}^k$. The strategy of [14] *adapts* the slope of the phase surface to the desired $e(\hat{\mathbf{x}}_o)$ distribution in each pixel. Thus, portion $e(\hat{\mathbf{x}})_k$ of the hogel area should have a slope at the k'th direction, that is, be a micro-facet tilted at θ^k . See Fig. 9 for an illustration of the phase mask presented on the display. This scheme leads to piecewise linear phase surfaces, forming a good match with the piecewise linear modulation capabilities of our display. Once we assign the slopes of the piecewise linear layers we can fine tune the phase using wave-optics based optimization, but in our experiments this only resulted in minor improvements.

Fig. 10 demonstrates 7 views of a hogel based display. Overall our display supports an angular range $\hat{\Omega}_{\theta} = 4\Omega_{\theta}$ and the 7 views plus an 8th residual image (not shown) are spaced over this range at intervals of $\frac{1}{2}\Omega_{\theta}$. Our piecewise linear component supports only 4 tilts. Each of the 4 cones is split to display two directional images by displaying sinusoids of two different frequencies on the piecewise constant layer. We used groups of q = 16 pixels in each hogel unit. As our SLM is only 768 pixels wide, the images we could display are of very limited resolution. The image size can be largely increased using SLMs with a larger pixel count such



Fig. 6: Content dependent slope assignment: We demonstrate the advantage of optimizing the amount of light directed to different regions of the target vs uniform illumination. In both images we have a bright foreground (car) in front of a dark background (trees). If we don't optimize the illumination the car is much darker. Additional effects are that the background trees disappear when the car passes them and we can see more artifacts from replicas. Red boxes mark the position of these artifacts and the desired target. Next to each image we visualize s_x , the slope (tilt) assignment in different display pixels. See App. Figs. 14 and 15 for additional results.



Fig. 7: Focal stack holography: We demonstrate our display's ability to generate focal-stack content, by showing targets at two different distances. The images illustrate capturing the same hologram while focusing the camera on both front and back planes. Next to each image we visualize s_x , the slope (tilt) assignment in different display pixels.

as the PLUTO or GAEA SLMs.

7 LIMITATIONS

Lack of per-pixel control. As discussed in Sec. 5.2, our prototype results in large aberrations when neighboring pixels are tilted independently. Hence, in practice we only use it to tilt larger patches of pixels by the same angle. This limits the space of piecewise linear functions it can realize. **Quantized tilts.** The cascade proposed in this paper can only tilt light by a discrete subset of angles. Therefore to display a continuous tilt angle $\theta \in [-0.5\hat{\Omega}_{\theta}, 0.5\hat{\Omega}_{\theta}]$ we

realize the rounded angle $\lfloor \theta / \Omega_{\theta} \rfloor$ using the binary layers and the residual $\theta - \lfloor \theta / \Omega_{\theta} \rfloor$ using the standard piecewise constant SLM layer. This quantization leads to some replica artifacts in the resulting images which could be removed if the display would support a continuous range of tilts.

2D expansion. While our current prototype only expands étendue along the horizontal axis, implementing a vertical expension is a simple matter of rotating some of the polarization sensitive grating by 90 degrees. This does come at the cost of doubling the number of cascade layers.

Miniaturization. Another complicating factor is the number of SLM layers needed; while our prototype does realize



Fig. 8: Temporal multiplexing results: By sequentially changing the tilts on the first two layers and using the last panel for intensity modulation we can display different images in different viewing directions.



Fig. 9: Image formation in spatial multiplexing mode: The display is divided to hogels of size $q \times q = 16 \times 16$ pixels, and the area of each hogel is divided according to the desired energy distribution in this spatial position. e.g. if 25% of the hogel energy should be sent to direction 1, then 25% of the hogel area display slope $\vec{\theta}^1$. For example we demonstrate four hogels out of the entire display. On the right side we present the histogram of energies for the 8 different views (vector $e(\vec{x})$), and how each hogel area is divided displaying different slopes, hence sending different intensities at different directions. The numbers correspond to the different generated images, and image 8 acts as a residual image.



Fig. 10: Hogel results: We use our device to display 7 independent views, using spatial multiplexing. The rightmost image visualizes s_x , the slope (tilt) assignment in different display pixels, corresponding to the different generated images.

an exponential expansion of angular ranges with a small number of SLM layers, optically aligning these layers is often a challenge and further, the use of 4f relays adds to the size and bulk of the device. For some applications such as table-top displays, the form factor of the system may be less of an issue, but for near-eye displays there is an interest in thin lightweight devices. There are, however, multiple approaches for resolving this. We note that the reason we used 4f relays and placed gratings in their Fourier planes is that we wanted to work with off the shelf gratings. By getting customized polarization sensitive gratings whose tilt angles match the exact tilt we want to apply in each layer, we can attach these gratings directly to the SLM plane, and avoid the 4f relay. Thus we can concatenate different SLM panels one after the other as in [48].

Alternatively, piecewise-linear phase modulation would be naturally implemented with tiltable micro-mirror arrays. Such prototype devices have been demonstrated, and we hope our paper can motivate their development. In particular, the Fraunhofer Institution [7] offers a device providing analog 2D tilts with an array of 512×320 micro-mirrors. While this device would not allow independent tilt of each SLM unit, it can be optically relayed with a standard SLM, such that each micro-mirror unit is aligned with a subgroup of 4×4 or 8×8 pixels, tilting them together. As analyzed in [44], piecewise linear phase modulations that tilt groups of pixels simultaneously in the same direction can provide good étendue-expansion capabilities. Some calculation shows that given the pitch and maximal tilt range of the device in [7] one can obtain a 4×4 étendue-expansion over a PLUTO SLM and 2×2 étendue-expansion over a GAEA SLM.

Color display. Our current prototype only displays monochromatic images. As in previous work [10], [11], [13] one can display color images by using field sequential mode, temporally multiplexing lasers at three different wavelengths. Each wavelength would involve a separate calibration process.

Content optimization. As another limitation we acknowledge that all content optimization in this paper uses simple gradient descent strategies. Smarter content design strategies e.g. exploiting neural networks has been proposed in the literature [11], [13] and could probably be incorporated to further improve our results.

8 CONCLUSIONS

This paper explores how to increase the étendue of an SLM display. Rather than excessively increasing the pixel count, we suggest the usage of tiltable display units. We implement a prototype that realizes a discrete number of tilts with a cascade of SLMs, where one layer acts as a standard piecewise constant SLM and successive layers implement binary tilts. We show that with proper construction, these binary tilts can lead to wide expansion angles, growing exponentially with the number of layers.

SLMs have a wide range of applications beyond the ones demonstrated in this paper. Piecewise linear phase modulators which expand the tilting range of these SLM can offer significant advances for many other tasks, such as multifocal displays [20], light sensitive displays [14], [22], Focal Surface displays [21], and high dynamic range displays [25], [26]. While these applications require large tilt angles, many of them involve piecewise smooth slopes as supported by our display, rather than highly varying phases. Hence, we hope that this paper motivates the development of novel SLMs that go beyond piecewise constant phase offset at each pixel.

ACKNOWLEDGMENTS

We thank Marina Alterman, Rick Chang and Wei-Yu Chen for their valuable feedback on the manuscript. This research was supported by H2020 European Research Council (635537); Israel Science Foundation (1947/20); United States-Israel Binational Science Foundation (2008123/2019758); US National Science Foundation (2008464).

REFERENCES

- [1] L. Shi, F.-C. Huang, W. Lopes, W. Matusik, and D. Luebke, "Neareye light field holographic rendering with spherical waves for wide field of view interactive 3d computer graphics," ACM Trans. Graph., vol. 36, no. 6, nov 2017.
- G. Kuo, L. Waller, R. Ng, and A. Maimone, "High resolution [2] étendue expansion for holographic displays," ACM Trans. Graph., vol. 39, no. 4, jul 2020.
- [3] S.-H. Baek, E. Tseng, A. Maimone, N. Matsuda, G. Kuo, Q. Fu, W. Heidrich, D. Lanman, and F. Heide, "Neural Étendue expander for ultra-wide-angle high-fidelity holographic display," 09 2021.
- E. Buckley, A. Cable, N. Lawrence, and T. Wilkinson, "Viewing [4] angle enhancement for two- and three-dimensional holographic displays with random superresolution phase masks," Appl. Opt., vol. 45, no. 28, pp. 7334-7341, Oct 2006.
- H. Yu, K. Lee, J. Park, and Y. Park, "Ultrahigh-definition dynamic [5] 3d holographic display by active control of volume speckle fields," Nature Photonics, vol. 11, 01 2017.
- J. Park, K. Lee, and Y. Park, "Ultrathin wide-angle large-area [6] digital 3d holographic display using a non-periodic photon sieve," Nature Communications, vol. 10, 03 2019.
- Franhuofer, "Fraunhofer imps," 2021. [Online]. Available: {https: [7] //www.ipms.fraunhofer.de/en/Components-and-Systems/ Components-and-Systems-Actuators/Optical-Actuators/ spatial-light-modulators.html}
- Y. Song, R. M. Panas, and J. B. Hopkins, "A review of micromirror [8] arrays," Precision Engineering, vol. 51, pp. 729 - 761, 2018.
- G. Koulieris, K. Akşit, M. Stengel, R. Mantiuk, K. Mania, and C. Richardt, "Near-eye display and tracking technologies for virtual and augmented reality," *Computer Graphics Forum*, vol. 38, [9] no. 2, pp. 493-519, Jun. 2019.
- [10] A. Maimone, A. Georgiou, and J. S. Kollin, "Holographic neareye displays for virtual and augmented reality," vol. 36, no. 4, pp. 85:1-85:16, Jul. 2017.
- [11] S. Choi, M. Gopakumar, Y. Peng, J. Kim, and G. Wetzstein, "Neural 3d holography: Learning accurate wave propagation models for 3d holographic virtual and augmented reality displays," ACM Trans. Graph. (SIGGRAPH Asia), 2021.
- [12] J. Zhang, N. Pegard, Z. Jingshan, H. Adesnik, and L. Waller, "3d computer-generated holography by non-convex optimization," Optica, vol. 4, p. 1306, 10 2017
- [13] L. Shi, B. Li, C. Kim, P. Kellnhofer, and W. Matusik, "Towards realtime photorealistic 3D holography with deep neural networks," Nature, vol. 591, no. 7849, pp. 234–239, 2021. [14] D. Glasner, T. Zickler, and A. Levin, "A reflectance display," ACM
- Trans. Graph., vol. 33, no. 4, pp. 61:1-61:12, Jul. 2014.
- [15] N. Padmanaban, Y. Peng, and G. Wetzstein, "Holographic Near-Eye Displays Based on Overlap-Add Stereograms," ACM Trans. Graph. (SIGGRAPH Asia), no. 6, 2019.
- [16] T. Yatagai, "Stereoscopic approach to 3-d display using computergenerated holograms," Appl. Opt., vol. 15, no. 11, pp. 2722-2729, Nov 1976.
- [17] M. Yamaguchi, H. Hoshino, T. Honda, and N. Ohyama, "Phaseadded stereogram: calculation of hologram using computer graphics technique," *Proceedings of SPIE*, vol. 1914, pp. 25–31, 09 1993. [18] H. Kang, T. Yamaguchi, and H. Yoshikawa, "Accurate phase-
- added stereogram to improve the coherent stereogram," Applied optics, vol. 47, pp. D44-54, 08 2008.

- [19] K. Akeley, "Achieving near-correct focus cues using multiple image planes," Ph.D. dissertation, Stanford University, 2004.
- [20] J.-H. Chang, A. Levin, B. Kumar, and A. Sankaranarayanan, "Towards occlusion-aware multifocal displays," ACM Transactions on Graphics, vol. 39, 07 2020.
- [21] N. Matsuda, A. Fix, and D. Lanman, "Focal surface displays," ACM Transactions on Graphics (TOG), vol. 36, no. 4, pp. 86:1-86:14, Jul. 2017. [Online]. Available: http://doi.acm.org/10.1145/ 3072959.3073590
- [22] A. Levin, H. Maron, and M. Yarom, "Passive light and viewpoint sensitive display of 3d content," in ICCP, 2016.
- [23] M. Ashikhmin, S. Premoze, and P. S. Shirley, "A microfacet-based BRDF generator," in ACM SIGGRAPH, Jul. 2000.
- [24] T. Weyrich, P. Peers, W. Matusik, and S. Rusinkiewicz, "Fabricating microgeometry for custom surface reflectance," SIGGRAPH, 2009.
- [25] G. Damberg, J. Gregson, and W. Heidrich, "High brightness hdr projection using dynamic freeform lensing," ACM Transactions on Graphics (TOG), vol. 35, no. 3, pp. 24:1–24:11, May 2016. [Online]. Available: http://doi.acm.org/10.1145/2857051
- [26] R. Hoskinson, B. Stoeber, W. Heidrich, and S. Fel, "Light reallocation for high contrast projectionusing an analog micromirror array," SIGGRAPH ACM Transactions on Graphics, 2010.
- [27] R. Tyson, Principles Of Adaptive Optics. CRC Press, 2015.
- [28] R. K. Tyson, "Adaptive optics and ground-to-space laser communications," Appl. Opt., vol. 35, no. 19, pp. 3640-3646, Jul 1996.
- [29] I. M. Vellekoop and A. P. Mosk, "Focusing coherent light through opaque strongly scattering media," Opt. Lett., 2007.
- [30] R. Haeussler, S. Reichelt, N. Leister, E. Zschau, R. Missbach, and A. Schwerdtner, "Large real-time holographic displays: from prototypes to a consumer product," Proc. SPIE, vol. 7237, 02 2009.
- [31] C. Jang, K. Bang, G. Li, and B. Lee, "Holographic near-eye display with expanded eye-box," ACM Trans. Graph., vol. 37, no. 6, dec 2018.
- [32] M. Kim, S. Lim, G. Choi, Y. Kim, H. Kim, and J. Hahn, "Expanded exit-pupil holographic head-mounted display with high-speed digital micromirror device," ETRI Journal, vol. 40, 05 2018.
- [33] M.-H. Choi, Y.-G. Ju, and J.-H. Park, "Holographic near-eye display with continuously expanded eyebox using two-dimensional replication and angular spectrum wrapping," Optics Express, vol. 28, pp. 533–547, 01 2020.
- [34] J. Kim et al., "Foveated ar: Dynamically-foveated augmented reality display," ACM Trans. Graph., vol. 38, no. 4, jul 2019.
- X. Xia, Y. Guan, A. State, P. Chakravarthula, T.-J. Cham, and [35] H. Fuchs, "Towards eyeglass-style holographic near-eye displays with statically," in 2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR), 2020, pp. 312-319.
- [36] B. Hellman and Y. Takashima, "Angular and spatial light modulation by single digital micromirror device for multi-image output and nearly-doubled etendue," Opt. Express, vol. 27, no. 15, pp. 21 477-21 496, Jul 2019.
- [37] Y. Song, R. M. Panas, and J. B. Hopkins, "A review of micromirror arrays," Precision Engineering, vol. 51, pp. 729–761, 2018.
- [38] P. J. Gilgunn, "Soi-cmos-mems electrothermal micromirror arrays," Ph.D. dissertation, Carnegie Mellon University, 2010.
- [39] H. Zeng, Z. Wan, and A. D. Feinerman, "Tilting micromirror with a liquid-metal pivot," Journal of microelectromechanical systems, vol. 15, no. 6, pp. 1568–1575, 2006. [40] Y. Peng, S. Choi, N. Padmanaban, and G. Wetzstein, "Neu-
- ral Holography with Camera-in-the-loop Training," ACM Trans. Graph. (SIGGRAPH Asia), 2020.
- [41] S. Choi, J. Kim, Y. Peng, and G. Wetzstein, "Optimizing Image Quality for Holographic Near-eye Displays with Michelson Holography," OSA Optica, 2021.
- [42] P. Chakravarthula, E. Tseng, T. Srivastava, H. Fuchs, and F. Heide, "Learned hardware-in-the-loop phase retrieval for holographic near-eye displays," ACM Trans. Graph. (SIGGRAPH Asia), 2020.
- [43] K. Kavakli, H. Urey, and K. Akşit, "Learned holographic light transport: invited," Appl. Opt., vol. 61, no. 5, pp. B50-B55, Feb 2022
- [44] A. C. S. Sagi Monin and A. Levin, "Analyzing phase masks for wide etendue holographic displays," ICCP, 2022
- [45] D. C. Liu and J. Nocedal, "On the limited memory bfgs method for large scale optimization," Mathematical Programming, vol. 45, no. 1, pp. 503–528, Aug 1989.
- [46] O. S. Cossairt, J. Napoli, S. L. Hill, R. K. Dorval, and G. E. Favalora, "Occlusion-capable multiview volumetric threedimensional display," Applied Optics, vol. 46, no. 8, pp. 1244-1250,

Mar 2007. [Online]. Available: http://ao.osa.org/abstract.cfm? URI=ao-46-8-1244

- [47] A. Jones, I. McDowall, H. Yamada, M. Bolas, and P. Debevec, "Rendering for an interactive 360° light field display," ACM Transactions on Graphics (TOG), vol. 26, no. 3, Jul. 2007. [Online]. Available: http://doi.acm.org/10.1145/1276377.1276427
- [48] G. Wetzstein, D. Lanman, M. Hirsch, and R. Raskar, "Tensor displays: compressive light field synthesis using multilayer displays with directional backlighting," ACM Trans. Graph., vol. 31, no. 4, p. 80, 2012.
- [49] J. W. Goodman, Introduction to Fourier Optics. McGraw-Hill Book Company, 1968.
- [50] A. Hermerschmidt, S. Osten, S. Krüger, and T. Blümel, "Wave front generation using a phase-only modulating liquid-crystal-based micro-display with hdtv resolution," in *International Congress on Optics and Optoelectronics*. International Society for Optics and Photonics, 2007, pp. 65840E–65840E.
- [51] G. Lazarev, A. Hermerschmidt, S. Krüger, and S. Osten, "Lcos spatial light modulators: Trends and applications," Optical Imaging and Metrology: Advanced Technologies], W. Osten and N. Reingand editors, Wiley-VCH Verlag, 2012.
- [52] A. C. Sankaranarayanan, V. Saragadam, V. Rengarajan, R. Tadano, T. Zhuang, H. Oyaizu, and J. Murayama, "Programmable Spectral Filter Arrays for Hyperspectral Imaging," sep 2021. [Online]. Available: http://arxiv.org/abs/2109.14450



Sagi Monin Sagi Monin received his B.Sc. and M.Sc. (Summa cum Laude) degrees in Electrical Engineering from the Technion. He is currently a Ph.D. student at department of Electrical Engineering, Technion, Israel, supervised by Prof. Anat Levin.



Aswin C. Sankaranarayanan Aswin C. Sankaranarayanan is a professor in the ECE department at CMU, where he is the PI of the Image Science Lab. His research interests are broadly in compressive sensing, computational photography, signal processing and machine vision. His doctoral research was in the University of Maryland where his dissertation won the distinguished dissertation award from the ECE department in 2009. Aswin is the recipient of the best paper awards at CVPR

2019 and ICCP 2021, the CIT Dean's Early Career Fellowship, the Spira Teaching award, the NSF CAREER award, the Eta Kappa Nu (CMU Chapter) Excellence in Teaching award, and the Herschel Rich Invention award from Rice University



Anat Levin Anat Levin is an Associate Prof. at the department of Electrical Engineering, Technion, Israel, doing research in the field of computational imaging. She received her Ph.D. from the Hebrew University at 2006. During the years 2007- 2009 she was a postdoc at MIT CSAIL, and during 2009-2016 she was an Assistant and Associate Prof. at the department of Computer Science and Applied Math, the Weizmann Inst. of Science.

Exponentially-wide étendue displays using a tilting cascade

Supplementary appendix

APPENDIX A

A.1 Liquid crystal displays overview

The effect of a liquid crystal (LC) cell on a fully polarized, monochromatic light wave can be described using the formalism of *Jones matrices* [49].

We describe a fully polarized incident light using a *Jones vector*, expressing the projection of the polarization vector on some predefined x, y basis:

$$\vec{U} = \begin{pmatrix} U_x \\ U_y \end{pmatrix},\tag{10}$$

where U_x, U_y are complex scalars. The action of a polarizing element such as a LC cell is described using a 2×2 Jones matrix *L*. The resulting state of polarization $\vec{U'}$ is given by

$$\vec{U}' = L\vec{U}.\tag{11}$$

As explained in [50], [51] the Jones matrix which describes the action of the LC cell e.g. PLUTO SLM (Holoeye Inc.) is given by

$$L = \begin{pmatrix} e^{i\beta} & 0\\ 0 & 1 \end{pmatrix}, \tag{12}$$

where the birefringence β can be computer controlled by applying different voltage in each pixel. If the incident light is linearly polarized parallel, expressed by a Jones vector of the form $U = (U_x, 0)^T$, the multiplication with *L* can change its phase.

The same principle can be used for amplitude modulation, if the illumination is polarized at 45° , and the outgoing light is passing via a polarizer with the orthogonal orientation.

In our case we also use illumination oriented at 45° relative to the birefringence axis, described by a Jones vector of the form $U = (U_0, U_0)^T$. We note that by programming the pixel to display $\beta = 0$ we maintain the original polarization, and by displaying $\beta = \pi$ we rotate to the orthogonal polarization state.

A.2 System calibration

Our systems requires careful calibration of several parameters.

Calibration of SLM for polarization rotating: For each angle tilt layer we use an SLM to either rotate the polarization of the light in the pixel or maintain the polarization. For the two SLMs in our system which are used to binary tilt the light, we calibrate the relationship between the voltage applied to each pixel and the output polarization. Once calibrated, for each SLM we use two values, one value for rotating the polarization in each pixel and another maintaining the polarization of the light in the same direction.

Calibration of SLM for phase modulation: For phase modulation calibration we use the suggested procedure



Fig. 11: Distortion between different linear slopes: Our system suffers from distortion which is introduced by different optical elements. We use the angle tilt layer to generate a checkerboard pattern, resulting in a pattern at one view-point and its negative (the complementary checkerboard pattern) in a different viewpoint. The different colors match the different view-points. In an ideal system the green and red squares should be complementary, but aberrations cause them to shift and overlay. The amount of distortion shift varies over the image and is stronger in the periphery of the SLM.

described by the manufacturer of self-reference interference calibration. We place a mask with two holes in front of the SLM and illuminate the mask with a collimated polarized light. The light after the mask illuminates the SLM at two different halves and is then collected by an objective lens and projected onto a camera creating an interference pattern. By keeping one half of the SLM with a constant voltage, and changing the voltage of the second half we measure the phase shift in the interference pattern and calculate the phase shift. Since our SLM has coupling between the amplitude and phase, we also measure the intensity when changing the voltage. We use both amplitude and phase measurements to create a look up table between voltage and phase-amplitude values.

Mapping between SLM:

In an idealized model the 4f system we have between two consecutive SLMs should perfectly image one on top the other. In practice as we use simple lenses significant distortion is present and we could not map one SLM to the other using a simple homography transformation. To this end we followed the registration procedure of [52]. We first find a mapping between each SLM to our camera and later find mapping between the different SLMs. We start by describing our procedure for one SLM. We start by displaying N rows on the SLM and capturing them with the camera. From this image of rows we get a set of N rows from the SLM labeled as $\{r_i\}$ which are mapped to coordinates $\{x_i^r, y_i^r\}$ of the camera. We calculate a transformation between the coordinates and the rows by solving:

$$\min_{C} \|f_r(C, x_i^r, y_i^r) - r_i\|^2,$$
(13)

where $f_r()$ is a third order polynomial function, and C are the corresponding coefficients. We use RANSAC to solve Eq. (13). We similarly repeat this process for different columns of the SLM resulting in columns $\{c_i\}$, coordinates $\{x_i^c, y_i^c\}$ and a polynomial fit $f_c()$. We fit horizontal and vertical polynomials for each SLM. Next we find an inverse mapping between the camera to the SLM. This is done by generating multiple random points on the SLM and using the obtained polynomial fits to calculate corresponding points on the camera. With these points we again solve Eq. (13) obtaining the inverse transformations from the camera to the SLM. This process is then repeated for each SLM individually, resulting in six polynomial fits from the SLMs to the camera and six polynomial fits from the camera to the different SLMs. To calculate the mapping between the SLMs we again generate multiple random points and calculate their mappings from the first SLM to the camera and finally to the target SLM. Using the corresponding points we again solve Eq. (13) and result with the mapping between the different SLMs.

In addition, we observe that light with different tilts pass through different paths and suffer from slightly different distortions. In Fig. 11 the difference distortions are illustrated. The first SLM is used to create a checkerboard pattern tilting the white squares to one direction and the black squares to a different direction. The two images are then captured by a camera with an aperture in the Fourier plane in front of the camera, blocking one of the two tilts. In Fig. 11 we display both images using different colors to show different angle tilts. If the different tilts suffer the same distortion we expect the to images to complete without any overlap. However, in the resulting image we see significant overlap between the two images especially in the peripheral of the SLMs. To account for the distortion of every tilt individually, we repeat the above process for each tilt direction independently.

Phase aberration calibration: In addition to geometric distortions that effect the mapping between the different SLMs, our system suffers from phase aberrations. To calibrate for these aberrations we follow the works of [10], [31]. We attempt to map the optical aberrations with a small set of Zernike polynomials:

$$u_{Zernike}(\vec{x}) = \sum_{k} c_k Z_k(\vec{x}), \qquad (14)$$

where Z_k and c_k are the k^{th} Zernike polynomial and coefficient, respectively. We limit the search to five low order Zernike basis polynomial corresponding to: focus, vertical



Fig. 12: Optical aberrations: Result of calibrated aberrated wavefront illuminating the phase-modulation SLM.

and oblique astigmatism, and horizontal and vertical comas. To compute the coefficients we use our SLM to generate a binary pattern, and iterate over different values of coefficients, looking to maximize the intensity inside the binary pattern. As in [10], we notice that the aberrations vary over the SLM, with significantly stronger aberrations in the peripheral of the SLM. To overcome this we spatially divide the SLM into 4×4 sub-regions and find Zernike coefficients for each one. The final calibrated phase is presented in Fig. 12. In contrary to geometric distortion which are significantly effected by the linear phase added by tilts, the aberrations are not as sensitive to the linear phase and we used the same phase correction with all tilt directions.

A.3 Component list

Our system is constructed of multiple optical elements. As a light source we used a 532nm laser [CPS532 - Thorlabs], with a collimating lens (AC508-200-A-ML Thorlabs) and a polarizer (LPVISC100 Thorlabs). For the angle tilt layers, in each layer we used: one SLM (LC-2012 Holoeye) to rotate the polarization of each pixel, two lenses (AC508-200-A-ML Thorlabs) for the 4f system, two polarization gratings (16-590 Edmund optics), and as the gratings operate with circular polarization we required two quarter-wave plate retards (WPQ20ME-532 Thorlabs and 88-257 Edmund optics). For phase modulation we used an additional SLM (LC-2012 Holoeve) with two polarizers before and after the SLM. To capture the holography results we used a camera attached with a Cannon lens (Prosilica GT3400 with 180mm Cannon lens). To capture temporal and spatial multiplexing results we used Grasshopper3 camera (Teledyne FLIR) with a 50mm Nikon camera lens.

A.4 Focal stack optimization

To optimize SLM content to generate a focal stack target we follow the Frensel propagation models used in previous work. The image I_z that an SLM produces at focal depth z(which is now different than the Fourier plane of Fig. 1) can be computed using the Frensel propagation operator:

$$\mathcal{F}\left(h_z(\vec{\boldsymbol{x}})e^{i\phi(\vec{\boldsymbol{x}})}\right),\tag{15}$$

with

$$h_z(\vec{x}) = e^{\frac{2\pi i z}{\lambda}} e^{i\pi\lambda z \|\vec{x}\|^2}.$$
 (16)

With this goal in mind, we search for SLM content minimizing the error with a target focal stack

$$\underset{\phi}{\operatorname{argmin}} \sum_{z} \left\| g * \left| \mathcal{F} \left(h_{z}(\vec{x}) \cdot e^{i(\phi_{c}(\vec{x}) + \phi_{l}(\vec{x}))} \right) \right|^{2} - g * I_{z} \right\|^{2}.$$
(17)

Fig. 7 in the main paper visualizes the outcome of this optimization.

A.5 Additional results

We present additional capture results. For results shown in the main text, we add comparisons between simulations and real capture. In Fig. 13 we compare the captured results with simulated ones. In Figs. 14 and 15 we show additional results of content dependent dynamic range, where we can see that the resulting artifacts can also be seen in simulations.

A.6 Compare average vs. single capture

In Fig. 16 we show a comparison between a single capture and averaging 10 captures of different holograms generating the same target. As different solutions lead to somewhat different speckles, by temporal multiplexing a few solutions we can reduce speckle artifacts.



Fig. 13: Display capture: We used our display with piecewise linear slope patches of 16×16 and 128×128 pixels. Big patches allow for better contrast in simulation. In practice the contrast of small patches is even lower than the simulation prediction due to the challenge in accurately calibrating phase modulators.



Fig. 14: Content dependent slope assignment: We demonstrate the advantage of optimizing the amount of light directed to each part of the image vs using equal distribution of the light. If we don't optimize the amount of light the images suffer from more artifacts and the resulting image is less bright for the same exposure. Candle source image by Tahmid Munaz (CC BY-SA 2.0)



Fig. 15: Content dependent slope assignment: Since the light is not distributed correctly we can see that parts of the fish are darker in comparison to the target image, and that image suffers from non-uniform intensity in regions it should be uniform (marked in red boxes). Fish image by Dave Huth (CC BY-NC 2.0)



Fig. 16: We show a comparison between averaging 10 images vs a single capture for all experimental results (including appendix material). Mountain source image by r-z (CC BY 2.0); Candle source image by Tahmid Munaz (CC BY-SA 2.0); Fish image by Dave Huth (CC BY-NC 2.0)