

Mix-and-Match Holography

YIFAN PENG, The University of British Columbia, King Abdullah University Of Science And Technology

XIONG DUN, King Abdullah University Of Science And Technology

QILIN SUN, King Abdullah University Of Science And Technology

WOLFGANG HEIDRICH, King Abdullah University Of Science And Technology, The University of British Columbia




Fig. 1. Overview of mix-and-match holography. We computationally design pairs of diffractive optical elements that encode multiple holograms or target images under different geometric alignments or DOE pairings (top left: encoding scheme, bottom left: fabricated DOEs). By changing the geometric alignment or pairing different DOEs, the individual encoded holograms can be de-multiplexed. Illumination with a pre-determined lighting condition produces the corresponding target image, that can be generalized in a wide range of scenarios (selected results presented on the right: through offset layer pairing (a), combinational pairing (b), multiview images' encoding (c), and animated images through linear translation (d)). The design process makes use of a combination of iterative phase retrieval methods and complex matrix factorization.

Computational caustics and light steering displays offer a wide range of interesting applications, ranging from art works and architectural installations to energy efficient HDR projection. In this work we expand on this concept by encoding several target images into pairs of front and rear phase-distorting surfaces. Different target holograms can be decoded by mixing and matching different front and rear surfaces under specific geometric alignments. Our approach, which we call mix-and-match holography, is made possible by moving from a refractive caustic image formation process to a diffractive, holographic one. This provides the extra bandwidth that is required to multiplex several images into pairing surfaces.

We derive a detailed image formation model for the setting of holographic projection displays, as well as a multiplexing method based on a combination of phase retrieval methods and complex matrix factorization. We demonstrate several application scenarios in both simulation and physical prototypes.

CCS Concepts: • **Computing methodologies** → **Computer graphics**; • **Hardware** → *Displays and imagers*;

This work is supported by the KAUST baseline funding, as well as the UBC 4YF Doctoral Award.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2017 ACM. 0730-0301/2017/11-ART191 \$15.00
DOI: 10.1145/3130800.3130839

ACM Reference format:

Yifan Peng, Xiong Dun, Qilin Sun, and Wolfgang Heidrich. 2017. Mix-and-Match Holography. *ACM Trans. Graph.* 36, 6, Article 191 (November 2017), 12 pages.

DOI: 10.1145/3130800.3130839

1 INTRODUCTION

Goal-based caustics and computational light steering have for a long time been a promising research topic in computer graphics, with proposed applications in art and architecture (e.g. [Finckh et al. 2010; Kiser et al. 2013; Papas et al. 2011]) as well as energy efficient high dynamic range (HDR) projection [Damberg et al. 2016; Hoskinson et al. 2010]. The general problem can be stated as designing an optical system that arranges micro-structures, material distributions, or other physical parameters of an optical element to produce a target light distribution under a given illumination condition.

Traditionally in graphics, this problem has been tackled using a *refractive* setting, i.e. by optimizing a free-form lens with one flat and one curved or patterned surface, that generates the desired light distribution under collimated white light. The optics community has also extensively explored the *diffractive*, or *holographic* setting, which offers several advantages. First, refractive designs for computational caustics can be quite bulky, while diffractive designs are flat at a macroscopic scale, and can therefore be very thin and light-weight. Second, whereas refractive designs must aim for relatively

smooth surfaces in order to avoid diffraction artifacts, holographic approaches directly incorporate diffraction into the image formation process, and can therefore make use of smaller features, which drastically increases the bandwidth (i.e. the number of degrees of freedom) for the optical design. Finally, we note that, while diffractive optical elements (DOEs) are challenging to prototype, they can be mass-produced at high speed and low cost using methods such as nano imprinting [Ahn and Guo 2009; Chou et al. 1996].

In the research of holographic techniques, angular or wavelength multiplexing have been developed to further increase the bandwidth of holograms so as to encode a large amount of desired data. In our work we propose *mix-and-match holography*, which is a new spatial multiplexing scheme for holography, by which different target holograms are encoded onto pairs of front and rear diffractive optical elements (DOEs, see Figs. 1, 2). These DOEs are designed in such a way that different pairings or geometric alignments represent different target holograms. By mixing and matching different front and rear DOEs and alignments, we can decode the different holograms, and these can in turn be used to produce different target projected images when illuminated under a specific lighting condition. In addition to adding an interesting dynamic component to existing applications of caustic displays, this mix-and-match design enables new applications in security, similar to the work by Naor and Shamir [1994], since the two phase DOEs can be thought of as an encrypted message and a key: both are required to decode the target information, but can be stored separately, unlike, for example, volume holograms.

Our main technical contributions are as follows:

- We re-formulate the goal-driven caustic problem as a holographic reconstruction problem, and derive detailed image formation models for this scenario.
- We exploit the larger design space of diffractive optics by firstly using a modified computer-generated hologram algorithm to retrieve a target phase distribution, and more importantly, using complex matrix factorization to multiplex holograms of several target images onto a pair of diffractive surfaces via mix-and-match holography.
- We synthetically and experimentally validate four application scenarios subject to different mix-and-match schemes, including encoding high-contrast holographic image through offset layers, representing multiple target images through combinatorial pairings, reconstructing animated images through linear translation, and encoding multiview images through a barrier-based display configuration.

2 RELATED WORK

Refractive caustic designs. Throughout the last few years, a large body of computational caustic designs have emerged in the graphics community, that mostly focusing on a refractive setting [Hullin et al. 2012]. We can distinguish between discrete designs [Papas et al. 2012, 2011; Yue et al. 2012] where the surface is broken into smaller elements, and continuous designs (see [Finckh et al. 2010; Kiser et al. 2013] as well as [Schwartzburg et al. 2014; Yue et al. 2014]) that yield a single smooth surface. Discrete designs have the advantage that they allow for discontinuities in the surface, which results in

flatter geometries, and more lightweight, compact objects. On the other hand, the discontinuities also introduce artifacts (including diffraction) that limit the contrast compared to smooth designs. The design process of all the above work is based on ray optics, and thus neglects diffraction.

The concept of multiplexing images from two or more portions has been reported recently. Fuchs et al. [2008] proposed a reflectance field display that combines different optical elements and encoded transparencies to allow for projecting different light fields. Refractive optics in combination with multiple attenuation masks have also been used for generating animations [Baran et al. 2012]. Other solutions include combining structured white light with masked prisms [Hostettler et al. 2015] or combining single source image with a refractive lenslet array [Papas et al. 2012]. This last work in particular can reveal different hidden images when the lens is placed at prescribed orientations on the source image or viewed from different angles. These works on steganography, display, and spectral color production with refractive means serve as an inspiration for our own work with a holographic image formation.

Instead of a purely refractive approach, Damberg et al. [2016; 2015] proposed a design that uses a spatial phase modulator but also a geometric optics image formation model to arrive at a piecewise smooth phase profile with high diffraction efficiency. This work can be interpreted as a hybrid between refractive and diffractive designs. Unfortunately, we found that the piecewise smoothness cannot be maintained while multiplexing multiple images into a single phase pattern.

Diffractive displays. Researchers in graphics have proposed several designs for exploiting diffractive optics in display applications. Recent attempts to fabricate the optical surfaces with desirable BRDFs or BTDFs have been investigated [Glasner et al. 2014; Levin et al. 2013; Ye et al. 2014]. Ye et al. [2014] proposed to use multiple diffractive layers to encode a 4D BRDF or BTDF. Finally, Xue et al. [2014] multiplexed multiple images into a single diffraction grating at different wavelengths of light based on an analytic model.

Computer-generated holography. In the optics community, there has been a considerable amount of work on computer-generated holography (CGH) for decades. Earlier work includes generating binary holograms by computer drawing [Lohmann and Paris 1967] and forming an image by a wavefront reconstruction device called Kinoform [Lesem et al. 1969]. Later development of holography includes the work on computationally designing complex transmittance or reflectance functions for holographic systems (see e.g. [Buckley 2010; Nagahama et al. 2016; Psaltis et al. 1995; Tricoles 1987]). There have been various optimization-based algorithms capable of generating desired phase or amplitude distributions considering variables in both the spatial and the spectral domain. Iterative methods based on Gerchberg-Saxton (GS) search, simulated annealing (SA) algorithm, and direct binary search (DBS) have been applied to the optimization of both monochromatic and broadband DOEs [Fienup 1982; Gerchberg and Saxton 1971; Kim et al. 2012]. Within our work, the hologram optimization is formed on the basis of well-known GS search (see Sec. 4.2).

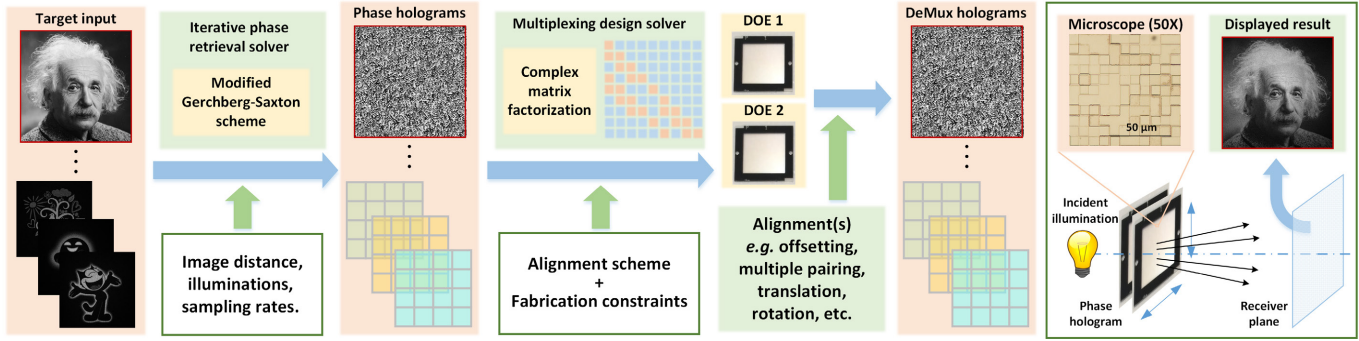


Fig. 2. Processing pipeline. Given a sequence of target inputs and the user defined reconstruction requirements, we first obtain a sequence of high fidelity phase holograms by applying a modified GS-based phase retrieval algorithm. Following that, we multiplex multiple of these phase holograms into a latent complex transmittance function, based on a designated alignment scheme and fabrication constraints. Relying on complex matrix factorization, two factorized matrices (vectors) can then be found, which represent the complex transmittance functions of two diffractive optical elements. Finally, by pairing two phase DOEs under different geometric alignments, we can de-multiplex multiple holograms. The corresponding projected images are formed by illuminating the setup under a specific lighting condition. (Albert Einstein Photograph by Orren Jack Turner)

Multiplexed holograms. Multi-layer or volume holograms can provide further increased bandwidth and thus allow for additional degrees-of-freedom in the design process [Gerke and Piestun 2010]. Using angular multiplexing, Mok [1993] stored many holograms within a single crystal of LiNbO₃ to increase the bandwidth. A survey of this inspiring body of work can be found in [Psaltis and Burr 1998], which also investigates their application potential in data encryption and storage. Our mix-and-match design differs from conventional volume holograms that use reference-recording scheme with either wavelength or angular multiplexing to encode the data. Brand [2011] proposed the specular holography that uses differential geometry to control light steering for several viewing geometries, which is not derived from diffractive optics.

Finally, Shi et al. [2007] and Meng et al. [2007] investigated similar concepts of hiding multiple images using two or more layered phase holograms. However, their setting is quite different from ours in several ways: first, their work results in a static design of a fixed pair of front and rear surface, while we instead use a combination of phase retrieval with complex matrix factorization to enable mix-and-match designs of many pairings. Second, in their scenario, multiple images are represented as essentially a depth-dependent point spread function, i.e. the decoding step involves moving the projection screen. By comparison, our design decodes different holograms and images by changing the relative geometric alignment of the DOEs themselves.

We envision that our mix-and-match holography scheme can be applied in a number of scenarios. For instance, one can leverage diffractive propagation to obtain high quality imaging or encryption visualization with (virtual) alignment for (digital) security applications. Besides, one can obtain actual display functionality with state-of-the-art hardware capabilities, like using fabricated DOEs or commercial spatial light modulators (SLMs).

Compressive displays. Compressive computational displays are of relevance to our work since they partially compensate for the limited data bandwidth of current display hardware (e.g. LCD panels). The fundamental concept behind compressive displays is to encode

high-dimensional data such as light fields into a small number of either multiplicative (e.g. [Heide et al. 2014a; Lanman et al. 2011; Wetzstein et al. 2012]) or additive planes (e.g. [Lee et al. 2016]). Similar methods have been applied to achieve super-resolution effects (e.g. [Berthouzoz and Fattal 2012; Sajadi et al. 2012]). The methodology behind these methods is often based on non-negative matrix or tensor factorization, and related optimization problems. This is also the case for our framework, which is most similar to work by Heide et al. [2014b], but operates in a diffractive setting.

3 OVERVIEW

Our goal is to create a holographic projection, as depicted on the right side of Fig. 2. A pair of a front and rear diffractive optical element placed directly on top of each other forms a hologram (center-right of Fig. 2). When illuminated by either a directional light or a point light (depending on the design), a caustic image is formed at a pre-determined image plane. Moreover, by mixing-and-matching different front and rear DOEs, or by changing their relative geometric alignment, we can extract different holograms, which in turn correspond to different target images. In the following we will distinguish between the (*phase*) *hologram*, which is the diffraction pattern that results in a given image on the image plane, and the (*phase*) *DOEs*, which are the two constituent diffractive optical elements that together can produce one or more holograms, depending on their geometric alignment.

To achieve this goal, we propose mix-and-match holography, with the complete framework illustrated in Fig. 2. This is a novel multiplexing scheme that combines two main technical components. First, we develop a modified Gerchberg-Saxton (GS) algorithm for creating a computer-generated phase hologram (CGH) by computing a target phase function for each input image (Sec. 4.2). Second, we multiplex the retrieved holograms for all images onto a pair of phase-only DOEs, where each hologram is represented by a different geometric alignment of the two phase DOEs. This multiplexing step can be implemented as a complex matrix factorization problem,

subject to enforcing physical constraints (Sec. 4.3). Therefore, we efficiently incorporate phase retrieval methods and modern numerical optimization strategies to significantly increase the design space and overcome the hardware limitations of refractive caustic displays. In the implementation, we describe *four different* application scenarios with different kinds of geometric alignments (Sec. 5).

4 COMPUTATIONAL DESIGN

In this section we discuss in detail the two aspects of the design process: the computer-generated holography and the multiplexing step. First, however, we briefly review the image formation using Fresnel diffractive wave propagation.

4.1 Diffractive Image Formation

We start by analyzing the propagation of light from the hologram plane to the image plane. To this end, we apply scalar diffraction theory [Goodman 2005] under the assumption of paraxial incident light. Using the Fresnel approximation, the light propagation is then given as:

$$\tilde{U}(x, y) = \frac{\exp(jkz_1)}{j\lambda z_1} \int \int_{\Sigma} \tilde{U}(x_1, y_1) \exp\left\{\frac{jk}{2z_1} [(x - x_1)^2 + (y - y_1)^2]\right\} dx_1 dy_1, \quad (1)$$

where (x, y) and (x_1, y_1) refer to positions on the image plane and the hologram plane, respectively. j is the imaginary unit, k the wave number, and z_1 stands for distance from the hologram plane to the image plane. According to Eq. 1, one can calculate the integral over the whole DOE transmittance aperture Σ to compute the complex amplitude distribution $\tilde{U}(x, y)$, including both amplitude and phase information. Please note that Eq. 1 represents essentially a Fourier transformation.

Note that $\tilde{U}(x_1, y_1)$ is the wave *exiting* the DOE plane, i.e. after diffraction in the hologram has already occurred. It is related to the *incident* wave \tilde{U}_0 at the hologram via the complex transmittance function $t(x_1, y_1)$:

$$\tilde{U}(x_1, y_1) = \tilde{U}_0(x_1, y_1) \cdot t(x_1, y_1) \quad (2)$$

with

$$t(x_1, y_1) = A(x_1, y_1) \cdot \exp(j\Phi(x_1, y_1)), \quad (3)$$

where A represents the amplitude and Φ represents the phase of each point on the hologram plane.

In our case, we design a phase hologram, meaning that the amplitude is constant across the plane, and thus the transmittance function is determined exclusively by the phase distribution $\Phi(x_1, y_1)$, which is physically represented by the height profile $h(x_1, y_1)$ of a DOE:

$$\Phi(x_1, y_1) = (n - 1) \frac{h(x_1, y_1)}{\lambda} 2\pi. \quad (4)$$

Putting all equations together, we can now design a height profile h on a substrate of refractive index n , that results in a given complex target amplitude $\tilde{U}(x, y)$ on the image plane. In the following, both the transmittance function and the height profile are pixelated into discrete squares of constant phase retardation and height, respectively. We call these squares *cells* in order to avoid confusion with pixels in the target images.

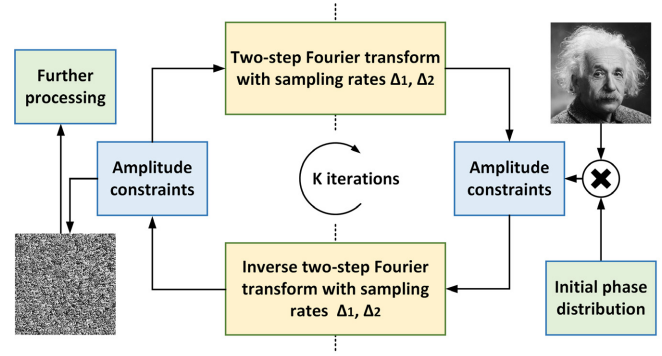


Fig. 3. Flowchart of our modified double-sampling GS algorithm. By relying on a virtual intermediate sampling plane (see Appendix for details), the sampling rates on the image plane and the hologram plane can be efficiently decoupled and tuned subject to design intentions. This results in a two-step FFT light propagation process. The input is a gray-scale target image (intensity distribution), while the output is a phase hologram.

4.2 Hologram Optimization

Our hologram optimization approach (summarized in Fig. 3) follows the general Gerchberg-Saxton (GS) framework and its extended versions [Fienup 1982; Gerchberg and Saxton 1971; Shimobaba et al. 2013]: we start with a target image in the image plane and convert it to a wave representation $\tilde{U}^{(0)}(x_1, y_1)$, where the phase at each point is random, and the amplitude represents the image intensity at each point. This complex waveform is then propagated backwards from the image plane to the hologram plane using Eq. 1 to obtain $\tilde{U}^{(0)}(x, y)$. Since we are designing a phase-only hologram, the amplitude is set to 1 across the hologram plane, yielding an updated waveform $\tilde{U}^{(1)}(x, y)$, which is then forward-propagated to the image plane (again using Eq. 1), to give $\tilde{U}^{(1)}(x_1, y_1)$. In the image plane, an amplitude constraint is applied by enforcing that the amplitude matches the intensity distribution of the target image, before starting the next iteration.

As mentioned before, both the forward and the backward transport are essentially Fourier transformations. However, a single-step discrete Fourier transform (DFT) as used in a standard GS algorithm has two constraints. First, the sampling rate on the image plane is determined by the wavelength and the propagation distance. Second, the sampling pattern on the image plane is restricted by the sampling pattern on the hologram plane, and vice versa. We would like to design holograms that expand the projection, such that a small DOE can produce a large projected image. This requires different sampling rates and regions of interest on the hologram plane versus the image plane.

To overcome this challenge, we use a two-step sampling scheme similar to Okada et al. [2013]. This procedure uses a combination of two FFTs with a virtual intermediate plane, to replace the original single forward FFT process. Refer to the Appendix for detailed derivations. Another modification is to use a spherical incident wave (i.e. a point light) rather than a regular planar wave, similar to the work by Nagahama et al. [2016]. We position the center of this spherical wave at the virtual intermediate plane. This strategy is

effective to enlarge the projection angle, and offers the side benefit of a relatively uniform light distribution that reduces speckle noise.

4.3 Mix-and-match Multiplexing

After the CGH step, we have obtained a target transmittance function t_k for each of the k holograms. We would now like to multiplex the k transmittance functions onto two DOEs under different alignments. We assume that the two DOEs are in direct contact, so that the alignment has two degrees of freedom for translation, and one for rotation, although in this work we make no use of the rotational degree of freedom. We denote the complex transmittance of cell i in the first DOE as a_i , while similarly, b_j denotes the complex transmittance of the cell j in the second DOE. As a result, the synthesized complex transmittance of each paired cell fragment $t_{i,j}$ can be obtained as

$$t_{i,j} = w_{i,j}(a_i \cdot b_j), \quad (5)$$

where $w_{i,j}$ here denotes the physical overlap of cell i and cell j ($i \in [1, n], j \in [1, m]$) in a given geometric configuration.

We note that each individual t_k can be expressed by a set of $t_{i,j}$ subject to a specific geometric alignment of the two phase DOEs. As shown in Fig. 4, the multiplexing procedure can then be expressed as a mapping via a function f_w , defined as:

$$t_{i,j} \leftarrow f_w(t_k). \quad (6)$$

In vector-matrix form, all the input target holograms can be stacked into a single matrix \mathbf{T}

$$\mathbf{T} = \mathbf{W} \circ (\mathbf{A}\mathbf{B}^\dagger), \quad (7)$$

where \mathbf{B}^\dagger is the complex conjugate of \mathbf{B} , and \circ denotes the Hadamard product. We include a sparse weighted matrix \mathbf{W} to select only those cells that are physically overlapped.

Matrix factorization. With this notation, it is easy to see that the computational design of multiplexing is in fact a complex matrix factorization problem of finding two phase distributions that best approximate the target hologram in a least-squares sense. Specifically, we rewrite Eq. 7, by addressing the columns of two complex matrices \mathbf{A} and \mathbf{B} [Heide et al. 2014b; Ho 2008]. Then, the goal becomes to solve the following constrained least-squares optimization problem:

$$\mathbf{A}_{\text{opt}}, \mathbf{B}_{\text{opt}} = \underset{\substack{\mathbf{A} \in \mathbb{C}^{m \times r} \\ |\cdot|=1}, \substack{\mathbf{B} \in \mathbb{C}^{n \times r} \\ |\cdot|=1}}{\text{argmin}} \frac{1}{2} \left\| \mathbf{W} \circ \mathbf{T} - \mathbf{W} \circ \mathbf{A}\mathbf{B}^\dagger \right\|_2^2. \quad (8)$$

Note that r here is the rank of the approximation, with $r = 1$ if both DOEs have static (time-invariant) transmittance functions.

To solve this problem, the following well-known multiplicative update scheme can be used [Ho 2008]:

$$\begin{aligned} \mathbf{A}^{k+1} &\leftarrow \mathbf{A}^k \circ \frac{(\mathbf{W} \circ \mathbf{T})\mathbf{B}^k}{(\mathbf{W} \circ (\mathbf{A}^k \mathbf{B}^{\dagger k}))\mathbf{B}^k}, \\ \mathbf{B}^{k+1} &\leftarrow \mathbf{B}^k \circ \frac{\mathbf{A}^{\dagger k}(\mathbf{W} \circ \mathbf{T})}{\mathbf{A}^{\dagger k}(\mathbf{W} \circ (\mathbf{A}^k \mathbf{B}^{\dagger k}))}. \end{aligned} \quad (9)$$

Note that the division here is the Hadamard (element-wise) matrix division. For these update rules to be well-defined, we need to ensure that the denominators are non-zero by adding a small value ϵ in the algorithm. The above update rule is straightforward and effective,

however, the convergence of this algorithm is slow, which affects computational efficiency.

Alternatively, one can use a gradient descent update scheme [Cichocki and Zdunek 2007; Lin 2007] to solve Eq. 8 efficiently. Assuming function $g(\mathbf{A}^k, \mathbf{B}^k)$ represents the non-negative least squares term in Eq. 8, the update rules now are:

$$\begin{aligned} \mathbf{A}^{k+1} &\leftarrow \mathbf{A}^k - \alpha^k \nabla_{\mathbf{A}} g(\mathbf{A}^k, \mathbf{B}^k), \\ \mathbf{B}^{k+1} &\leftarrow \mathbf{B}^k - \alpha^k \nabla_{\mathbf{B}} g(\mathbf{A}^k, \mathbf{B}^k), \end{aligned} \quad (10)$$

where α_k is a step size that is adapted in every iteration. Specifically, α_k is updated as $\alpha_{k+1} \leftarrow \alpha_k / \beta$ or $\alpha_{k+1} \leftarrow \alpha_k \cdot \beta$, as described in detail in [Lin 2007].

During the complex matrix factorization, it is necessary to enforce the physical constraints of the optical system. This is analogous to enforcing non-negativity in traditional non-negative matrix factorization (NMF) problems. In our case, the physical constraint is that we desire a phase-only modulation in both DOEs, so that the amplitude of the complex transmittance must equal 1 in each cell. We enforce this with the following update rules for the individual components of \mathbf{A} and \mathbf{B} , which we call *conditional projection*:

$$\mathbf{A}_i^{k+1} \leftarrow \begin{cases} \frac{\mathbf{A}_i^{k+1}}{|\mathbf{A}_i^{k+1}|} & |\mathbf{A}_i^{k+1}| > 1, \\ \mathbf{A}_i^{k+1} & \text{else,} \end{cases} \quad (11)$$

where $|\cdot|$ denotes the amplitude of a complex value.

As mentioned, if we assume static DOEs, the matrices \mathbf{A}, \mathbf{B} are of rank 1. Specifically, \mathbf{A} is a $n \times 1$ matrix (aka. a vector), \mathbf{B} is a $m \times 1$ matrix, \mathbf{W} is a $n \times m$ sparse weight matrix, and \mathbf{T} is a $n \times m$ sparse matrix containing the target complex transmittance of each cell. We can rewrite the model in Eq. 8 as:

$$\mathbf{b}_{\text{opt}}, \mathbf{a}_{\text{opt}} = \underset{\substack{\mathbf{a} \in \mathbb{C}^{m \times 1} \\ |\cdot|=1}, \substack{\mathbf{b} \in \mathbb{C}^{n \times 1} \\ |\cdot|=1}}{\text{argmin}} \frac{1}{2} \left\| \mathbf{T} - \mathbf{W} \circ \mathbf{a}\mathbf{b}^\dagger \right\|_{\mathbf{W}}^2. \quad (12)$$

For the rank 1 weighted optimization, we use Newton's updates to enforce a faster convergence. Algorithm 1 shows the specific algorithm we use to solve this rank 1 problem. It consists of an unconstrained, bi-convex rank 1 alternating least-squared weighted matrix factorization, followed by a constrained projection operation after solving each subproblem, as in Eq. 11. In each subproblem, we diagonalize the weighting matrix and compute the outer product, followed by vectorization.

Thus, we have cast the non-convex problem as a sequence of convex optimization problems. The **b-step** and **a-step** in Alg. 1 can be solved using gradient-based or Hessian-based methods, such as fast Newton's updates. Our method shares similarities with Encoded Diffractive Optics [Heide et al. 2016], although we are now dealing with holograms that exhibit more randomness than that of a symmetric imaging lens or a relative smooth phase profile. In addition, we apply a different projection rule (refer to Eq. 11) rather than naively projecting all current optima back to their phase profile. We find that this additional filtering projection leads to a faster and better convergence than state-of-the-art (see supplement for details).

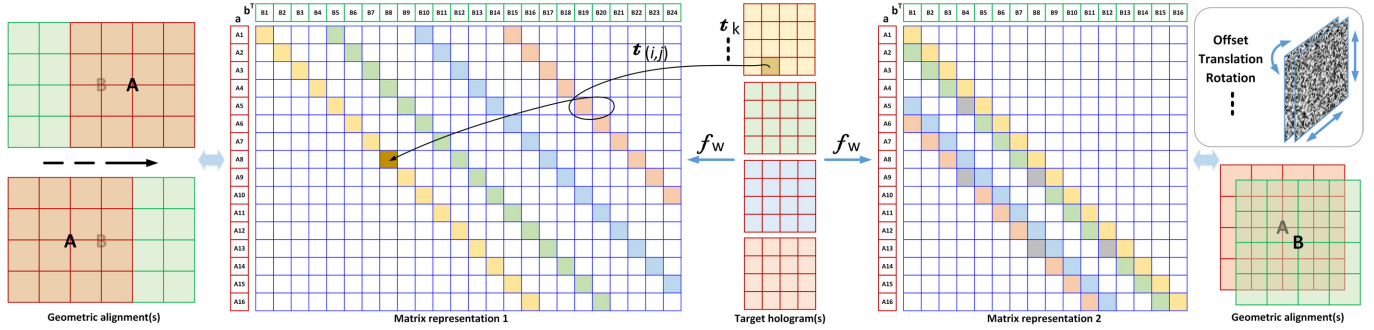


Fig. 4. Illustration of the matrix structure for our multiplexing scheme. Left: matrix formation for translational pairing. Center: a sequence of target holograms for multiplexing. Right: matrix formation for offset pairing. Specially, the matrix can be represented as the outer product of two vectors \mathbf{a} , \mathbf{b} , representing two DOEs, where the physically overlapped cells are selected via different geometric alignments (as shown in different color cell fragments). f_w here indicates an element-wise mapping from one of the target holograms t_k to the target matrix \mathbf{T} .

Algorithm 1 Rank 1 Alternating Least Squares Weighted Complex Matrix Factorization

- 1: $k = 0$, $\mathbf{a}_{\text{opt}}^0 = \mathbf{a}_{\text{init}}$, $\mathbf{b}_{\text{opt}}^0 = \mathbf{b}_{\text{init}}$
 - 2: **repeat**
 - 3: $\mathbf{b}_{\text{opt}}^{k+1} := \underset{\mathbf{b}}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{T} - \mathbf{a}\mathbf{b}^\dagger\|_{\mathbf{W}}^2$ \triangleright **b-step**
 - 4: $\mathbf{b}_{\text{opt}}^{k+1} := \operatorname{proj}_{(|\mathbf{b}_{\text{opt}}^{k+1}| > 1)}(\mathbf{b}_{\text{opt}}^{k+1})$ \triangleright conditional projection
 - 5: $\mathbf{a}_{\text{opt}}^{k+1} := \underset{\mathbf{a}}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{T} - \mathbf{a}\mathbf{b}^\dagger\|_{\mathbf{W}}^2$ \triangleright **a-step**
 - 6: $\mathbf{a}_{\text{opt}}^{k+1} := \operatorname{proj}_{(|\mathbf{a}_{\text{opt}}^{k+1}| > 1)}(\mathbf{a}_{\text{opt}}^{k+1})$ \triangleright conditional projection
 - 7: $k := k + 1$
 - 8: **until** Optimality achieved
-

5 APPLICATIONS AND RESULTS

This section describes four specific applications of mix-and-match holography — high-contrast holographic projection via offset pairing, versatile image reconstructions via multiple pairings, animated image reconstruction via translational pairings, and barrier-based multi-view projection display.

For all our applications we built custom phase DOEs on 4" fused silica wafers using photolithography and reactive ion etching (RIE). For our current proof-of-concept stage, four iterations of lithography and RIE were repeated to produce 16 discrete height levels in each DOE [Peng et al. 2016]. The principal design wavelength is chosen as $\lambda = 550\text{nm}$, and the corresponding total etching depth is 1195nm. The feature size on the DOE is set to $8\mu\text{m}$. This is a convenient size for prototype development, but we note that photolithography can be pushed to around $1\mu\text{m}$ and sub-wavelength feature sizes are possible with e-beam lithography. In both cases, the higher resolution would provide additional degrees of freedom for multiplexing a larger number of holograms. The close-up of Fig. 1 shows a $20\times$ microscopic image of the central areas of our fabricated DOEs. Refer to the supplementary document for additional details on the fabrication.

The pairing of DOEs requires element-wise alignment (i.e. tilt and de-center). Accordingly, alignment marks are added to the individual

DOEs. During experiments, the DOEs are mounted on a 3 degree-of-freedom alignment system for accurate control of position and rotation (see left-most close-up in Fig. 6). For the illumination we use a Thorlabs 532nm laser kit, a $5\times$ beam expander, and a custom lens system to realize different lighting conditions (collimated or spot light). Except for the DOEs, all the auxiliary components are off-the-shelf products.

5.1 Application Scenarios

In Sec. 4, we discussed the general principle of hologram multiplexing in a mix-and match sense. In the following we describe 4 specific scenarios that are enabled by different choices of the layer matrices \mathbf{A} and \mathbf{B} , as well as the weight matrix \mathbf{W} .

Offset layered high-contrast projection. The first application is a “super-resolution” example with just one target image. The feature size in the hologram has a direct impact on contrast and image quality of the target image, but depending on the fabrication process it can be difficult or expensive to further reduce the feature size. This example shows that the factorization from Section 4.3 can be used to approximate a high resolution ($4\mu\text{m}$ feature size) hologram with two lower resolution ($8\mu\text{m}$ feature size) DOEs as long as the global alignment can be controlled with $4\mu\text{m}$ accuracy.

Specifically, in this application, the alignment of the two DOEs is designed for an offset of half the feature size (Fig. 4). Accordingly, the two layers \mathbf{A}, \mathbf{B} are expressed as column vectors $\mathbf{a} \in \mathbb{C}^{n \times 1}_{|\cdot|=1}$ and $\mathbf{b} \in \mathbb{C}^{m \times 1}_{|\cdot|=1}$, respectively. We can efficiently solve this rank 1 optimization problem. Simulation results for this scenario are shown in Fig. 5. From the PSNRs and the close-ups (red and blue rectangles), we see that the offsetting configuration reconstructs an image with higher contrast and less speckle noise than that of a single layer. We also note that the maximum diffraction angle is dependent on the feature size. Decreasing the feature size allows for enlarged diffraction angles, and with our factorization we can achieve this effect without the need to use more precise fabrication techniques. Particularly, we note that this design feature of incorporating super-resolution effects into holograms can not be achieved by state-of-the-art CGH methods.

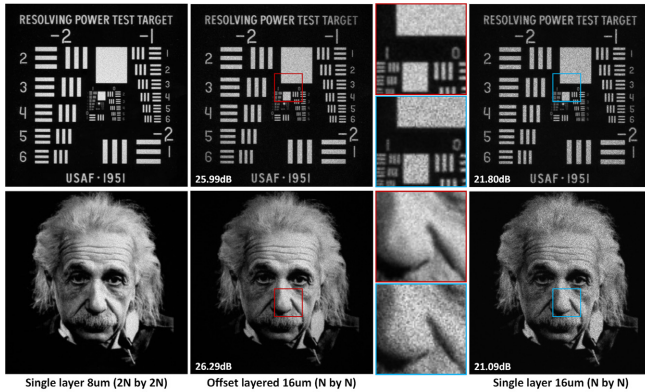


Fig. 5. Synthetic results of the sub-cell offset design. The images are reconstructed from a single DOE with a cell size of $8\ \mu\text{m}$ (left column), two DOEs with a cell size of $16\ \mu\text{m}$ that are offset by half a cell (middle), and a single DOE with $16\ \mu\text{m}$ cells (right). The PSNR are evaluated. Notice that the images are resampled and integrated over the physical size of $200\ \mu\text{m}$ per pixel subject to the angular resolution limit of the human eye at the observation distance of 50cm . The same setting applies to the followings. (Albert Einstein Photograph by Philippe Halsman)

A real result of projecting an image through an offset pairing of DOEs is shown in the center left close-up of Fig. 1 and in Fig. 6. Illuminated by the designated green incident light, we obtain a high-contrast image at the designated distance (center green close-up), whose quality is visually matched to the synthetic result. Note that the observed speckle is inherent in the use of highly coherent laser illumination, and can be reduced with a number of strategies that are orthogonal to our approach [Damberg et al. 2016]. As one would expect, using another wavelength of light that differs from the design wavelength by $>100\ \text{nm}$ results in a blurred image (right red close-up). An additional example is presented in the supplement.

Multiple discrete pairings. A first example of creating multiple images via different pairings of DOEs, is to have each DOE composed of several discrete blocks such that each combination of blocks from the two DOEs represents a hologram for a different target image, as shown in Fig. 7. For instance, if we have 9 target images, they can be assigned to a matrix T containing 3×3 blocks. We create three blocks in the column direction to form a matrix A and three blocks in the row direction to form a matrix B . Specially, we put all the elements on the diagonal of each individual block. Now we have sparse matrices $A \in \mathbb{C}_{|\cdot|=1}^{3m \times m}$, and $B \in \mathbb{C}_{|\cdot|=1}^{m \times 3m}$. Either Eq. 9 or Eq. 10 in Section 4.3 can be used to solve this problem.

In this scenario we only need to fabricate $2n$ DOEs in two groups to represent n^2 images. This application mode also has applications in privacy or security [Naor and Shamir 1994], that differ from prior holographic encryption examples that are discussed in Sec. 2. We can interpret one of the DOEs as a set of encoded messages, while the other is a set of keys. Different keys can unlock different messages. For instance to identify multiple banking accounts under one username, the user holds only one piece (either digitally or in real plate) while the bank holds multiple paired pieces, or vice versa.

Both synthetic and experimental results of this scenario are shown in Figs. 7 for images and 8 for message encoding.

For Fig. 8, we show specifically that, if image quality is not of paramount importance, the feature size of the holograms can be reduced fairly significantly, which makes the mix-and-match patterns more robust to misalignment. The feature size in this example is $50\ \mu\text{m}$, which allows the phase DOEs to be mounted into individual frames (whose procedure is a one-time work and can be done practically) which can be snapped together easily with sufficient precision and accuracy to decode the multiplexed image. A hand-on mix-and-match demo is presented in the supplementary video. We emphasize that this is a stress test for the image quality, which suffers from the lower diffraction efficiency in DOEs with larger features. It is certainly possible to building mechanical frames for alignment accuracy much better than $50\ \mu\text{m}$.

Animated projection under translation. By extending the multiplexing from a discrete set of static pairings to a translational alignment, one can generate animations. The concept is to shift a top layer DOE along one direction to overlap the different regions of a bottom layer DOE (see the left close-up of Fig. 9). At each designated alignment, the overlapped region of two layers generates the hologram for a specific frame of the animation. In our example, the bottom layer is only twice the width of the top layer, but we multiplex 26 frames into this configuration. In this scenario, the two layers A, B are expressed as column vectors $\mathbf{a} \in \mathbb{C}_{|\cdot|=1}^{n \times 1}$ and $\mathbf{b} \in \mathbb{C}_{|\cdot|=1}^{2n \times 1}$, respectively. The experimental results of projecting 26 key frames of animated images through translation pairing of DOEs are shown in Fig. 9.

Barrier-based multiview projection. Theoretically any subregion of a hologram is able to reconstruct partial frequencies of the full target image. This unique property makes it possible to synthesize a hologram that contains a partial design for multiple perspectives. By incorporating coded illumination via a binary barrier or a lenticular lens sheet, as in conventional auto-stereoscopic displays, one can multiplex holograms of multiview images into a single pair of DOEs. A possible configuration is illustrated on the left of Fig. 10. In this setting, the illumination from one direction only reaches certain regions of cells to reconstruct the corresponding perspective image. Note that this application can be combined with sub-cell alignment to recover some of the bandwidth lost by multiplexing multiple views into the same DOE pair.

The synthetic and experimental results of multiview projection are shown on the right of Fig. 10. One can see the change in occlusion between the fish (front) and the coral (back). Readers may also observe a residual image of the mask in the center of the measured image, which does not appear in the synthetic results. This is mainly due to a loss of diffraction efficiency caused by fabrication errors as well as variations in the illumination wavelength.

6 ANALYSIS

The presented prototype architecture and the resulting images that can be achieved demonstrate the feasibility of the concept. In the following we discuss some of the issues that our initial prototypes suffer from.



Fig. 6. Synthetic and experimental results of projecting an image through an offset pair of DOEs, with the diagram of set-up indicated (left). We present synthetic results (center left), experimental results illuminated with green laser light (center right), and experimental results illuminated with red laser light (right-most). The size of DOEs are 1.4cm by 1.4cm individually, and the projection distance is set 1m. The final projected image size is around 12cm by 12cm. (Albert Einstein Photograph by Orren Jack Turner)




Fig. 7. Synthetic (left) and experimental (right) results for all 9 possible combinations of 3 individual patterns on the front DOE and 3 individual patterns on the back DOE. Each individual image is reconstructed by aligning one pattern from group A and one pattern from group B. We use the same incident light source and the same projection distance as in Fig. 6. The final reconstructed image size is around 3.7cm by 3.7cm. (Multiple source images from online resource)

Image artifacts. The observed artifacts include image blur, noisy energy distribution, and image contrast loss, which can be traced back to the limitations of current manufacturing and assembling process. We quantitatively evaluate the perceptual quality of a reconstructed image using the offset pairing design. The variables added include: random Gaussian white noise to approximate fabrication error, and relative in-plane shift and rotation to approximate misalignment of two physical layers. Thus, based on the ground-truth image reconstructed with well-defined conditions (center left subfigure in Fig. 6), we present structural similarity (SSIM) assessment in Tab. 1. Refer to the supplement for full resolution results subject to different testing conditions.

Table 1. SSIM comparisons of synthetic images reconstructed from the offset pairing subject to different experimental conditions. The columns indicate the results under the conditions— 1 weak noise, 2 strong noise, 3 small rotation, 4 large rotation, 5 one directional shift, 6 dual directional shift, 7 weak noise + small rotation, 8 weak noise + large rotation, 9 strong noise + small rotation, 10 strong noise + large rotation, respectively.

Condition	1	2	3	4	5
SSIM	0.994	0.975	0.712	0.525	0.476
Condition	6	7	8	9	10
SSIM	0.420	0.708	0.523	0.697	0.520

We find that relative shift and rotation account for the main factors that degrade the image quality. A direct consequence is the




Fig. 8. Experimental setup (left) and decoded (right) results for all 4 possible combinations. For this demonstration, we have released the feature size to an extreme value of $50\mu\text{m}$ and aligned the DOEs onto mechanical frames, allowing for fast hand-on manipulation. The labels F and R indicate the front and rear DOEs, respectively. By snapping together two frames from each group, the messages is decoded. We use the same incident light source and the same projection distance as in Fig. 7. The final reconstructed image size is around 1.6cm by 1.6cm.

existence of ghost (crosstalk) images due to the overlap of diffraction orders. With careful alignment this artifact can be drastically reduced. Where precise pixel level alignment is not practical, an effective solution is to pad extra zeros outside the target image in the phase map retrieval (see Sec. 4). This trick directly enlarges the physical distance of two neighboring diffraction orders such that the overlap issue can be mitigated within the field. In our experiments, misalignment of less than one cell ($< 8\mu\text{m}$) will not result in significant image degradation. Although tedious, this alignment accuracy is quite feasible in modern manufacturing and engineering methods.

With respect to those multiplexing designs incorporating relative shifts, the two DOEs are no longer in direct contact, but with a very small air gap to avoid scratching the DOEs. In our case, the influence from a possible air gap can be neglected without causing noticeable artifacts. We have quantitatively analyzed this issue in the supplementary material.

Hardware limitations. The accuracy of DOEs fabricated by our photolithography process is highly dependent on the alignment accuracy between the used lithography masks and the stability of etching speed. Empirically, we have designed several pairs of alignment fiducials (from $20\mu\text{m}$ to $1\mu\text{m}$) to assist the alignment, and have used a gradual etching method to enforce a stable etching speed. In practice some of DOEs are well-fabricated (Fig. 6 and Fig. 7) while others (Fig. 10) may suffer from over-etching errors. Like other DOEs with discrete levels of surface relief, our fabrication process cannot guarantee perfect single order diffraction. This higher order diffraction artifact is still difficult to eliminate.

Our prototype is also limited by the micro-meter level accuracy of the optical mounts we built. With additional engineering efforts (as demonstrated for the hand-on pairing scenario) or mass-production, most of these alignment issues could be addressed. We also note that, for mass production, very inexpensive and fast methods exist

for reproducing the same high-resolution phase DOE over and over again. For example, nano-imprinting [Ahn and Guo 2009] can stamp sub-wavelength features directly into polymer film at the scale of rolls of material. Therefore, the aforementioned issues primarily affect prototyping and research, but not potential commercial applications.

Computational performance. Although convergence to a global optimum cannot be guaranteed for either the modified GS algorithm or the multiplexing method, both problems are in practice well behaved. Empirically, the modified GS algorithm converges in around 60 iterations for a zero-padded hologram with a resolution of $7,000 \times 7,000$ pixels. This takes about 15 minutes in our Matlab implementation on a PC with Intel Xeon i7 2.70GHz CPU. The multiplexing step takes another 2 minutes on the same system.

Within our work, we consider two types of efficient factorization algorithms. We analyze their convergence and the run time performance. Some results are shown in Fig. 11. We make two observations – first, Newton’s update generally leads to a faster convergence than the well-known multiplicative update; second, enforcing a conditional projection (recall Sec. 4) further outperforms conventional naive projection for both methods. Please refer to the supplement for detailed quantitative evaluations, including the factorization of a Fresnel lens profile that is derived from analytic calculation.

Notice that we solve a constrained matrix factorization problem over the domain of complex numbers. The physical constraints (i.e. enforcing projection) are somewhat akin to a traditional non-negative matrix factorization (NMF). As in NMF, the resulting non-convexity mirrors more or less the situation with NMF algorithms as well. That said, there may be other formulations, which may yield even better results without projection. We leave this open but interesting question to future work.

Spectral performance. While designing at 550nm, we use easily available 532nm and 650nm laser light sources in Fig. 6. This wavelength deviation may result in diffraction efficiency loss. In Fig. 6, we have presented a photographed result of directly illuminating a green light designed DOE pair with red illumination. Noticeably, the image is degraded because of the wavelength dependency of diffraction. We also present another example showing the result illuminated by a white spot light source with 550nm band-pass filter (half width 10nm), shown in Fig. 12 (center). Although the flux is not as strong as coherence laser illumination, we observe a decoded image with less speckle noise (aka. blurred by the integration over 10nm wavelength). Readers may notice the “grid” effect, which is a ghosting effect due to misalignment, as discussed previously and in the supplement. As expected, the full spectrum white example (right) suffers from severe blur because of color dispersion.

To obtain a color image, one could encode different holograms for different color primaries into the same DOE pair, similar to the parallax barrier setup (see the supplement for an initial result). Nevertheless, it is challenging to obtain full-color display in high definition when trading-off light efficiency, image size, spatial resolution, and system complexity. Multiple modulation pairings, similar to multi-modulator projectors, may be an alternative. We leave this topic for future work.




Fig. 9. The experimental results of projecting an animation through translation DOE pairings. Left: diagram of experimental set up. Top right: synthetic results of selected frames. Bottom right: experimental results of selected frames. The projection distance is set to 1m as in Fig. 6, while here we use a planar incident illumination to simplify the setup. The labels in the bottom-left of each sub-figure indicate the frame number. Refer to the supplementary document for the complete frame sequence.




Fig. 10. Left: system diagram. Top right: synthetic results. Bottom right: experimental results for projecting multiview images via a combination of a parallax barrier and directional incident illumination. The sizes of the DOEs are 1.4cm by 3.2cm individually, where we multiplex 5 perspective holograms. In our experiments, $d=1.2\text{cm}$ and $\Delta=2.24\text{mm}$ (light field image credits at [Wetzstein et al. 2012]).

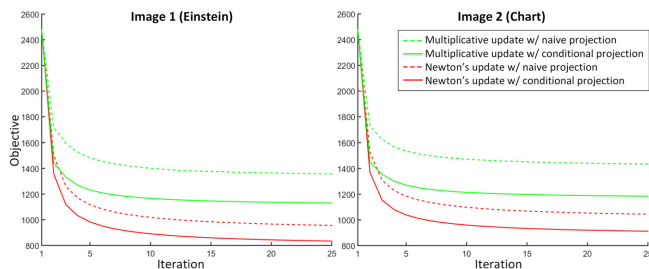


Fig. 11. Convergence for two factorization methods (rank 1 factorization on a 1,751 by 1,751 pixel target image) with and without projection (Eq. 11). The two images tested are same as in Fig. 5.

More than two layers. Another extension that we leave for future work is to increase the number of layers from two to three or more. This would change the matrix factorization problem in Eq. 7 to a tensor factorization problem, similar in spirit to the work by Wetzstein et al. [2012] on tensor displays.




Fig. 12. Results of an offset pairing example illuminated by a coherence green light source (left), broadband green light source (center) and a full spectrum white light source (right). (Lena source image from Wikipedia)

7 CONCLUSION

We have demonstrated that it is viable to encode the holograms for many target images into a pair of diffractive optical elements, where different alignments of the DOEs generate different target images. This is achieved by exploring the large bandwidth of holographic approaches. Computationally, we combine iterative phase retrieval methods and complex matrix factorization for multiplexing holograms onto a pair of diffractive optical elements, subject to

a designated alignment scheme. While the method exhibits some artifacts, which we analyze in detail, the overall image quality is very good for many applications.

In the future, our solution can serve as the basis for further generalizations to arbitrary numbers (and thicknesses) of physical or air layers, as well as more general types of geometric alignment. The unique benefits of multiplexed holograms can potentially inspire a wide range of application scenarios where the refractive caustic designs or single layer diffractive designs may fail.

Applications in security market can be very promising, like public-private encryption visualization of identity purposes. We note that for this application, the randomness of holograms and the effort needed for precise alignment in fact strength the security and privacy. Our solution can be a trigger for many people to start research on a variety of novel encryption modes. The display industry is another field where our solution may pave the way for exploiting diffraction in a variety of augmented reality (AR) designs. In addition, embedding holographic images can make architecture or industrial products more appealing. We believe that animations and other dynamic effects enabled by multiplexed holography enable new and enriched user experiences in many of these fields.

REFERENCES

- Se Hyun Ahn and L Jay Guo. 2009. Large-area roll-to-roll and roll-to-plate nanoimprint lithography: a step toward high-throughput application of continuous nanoimprinting. *ACS Nano* 3, 8 (2009), 2304–2310.
- Ilya Baran, Philipp Keller, Derek Bradley, Stelian Coros, Wojciech Jarosz, Derek Nowrouzezahrai, and Markus Gross. 2012. Manufacturing layered attenuators for multiple prescribed shadow images. In *Computer Graphics Forum*, Vol. 31. Wiley Online Library, 603–610.
- Floraine Berthouzou and Raanan Fattal. 2012. Resolution enhancement by vibrating displays. *ACM Trans. Graph* 31, 2 (2012), 15.
- Matthew Brand. 2011. Specular holography. *Applied Optics* 50, 25 (2011), 5042–5046.
- Edward Buckley. 2010. Holographic projector using one lens. *Optics Letters* 35, 20 (2010), 3399–3401.
- Stephen Y Chou, Peter R Krauss, and Preston J Renstrom. 1996. Nanoimprint lithography. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures Processing, Measurement, and Phenomena* 14, 6 (1996), 4129–4133.
- Andrzej Cichocki and Rafal Zdunek. 2007. Multilayer nonnegative matrix factorization using projected gradient approaches. *Int. J. of Neural Systems* 17, 06 (2007), 431–446.
- Gerwin Damberg, James Gregson, and Wolfgang Heidrich. 2016. High Brightness HDR Projection Using Dynamic Freeform Lensing. *ACM Trans. Graph* 35, 3 (2016), 24.
- Gerwin Damberg and Wolfgang Heidrich. 2015. Efficient freeform lens optimization for computational caustic displays. *Optics Express* 23, 8 (2015), 10224–10232.
- James R Fienup. 1982. Phase retrieval algorithms: a comparison. *Applied optics* 21, 15 (1982), 2758–2769.
- Manuel Finckh, Holger Dammert, and Hendrik PA Lensch. 2010. Geometry construction from caustic images. In *Proc. ECCV*. 464–477.
- Martin Fuchs, Ramesh Raskar, Hans-Peter Seidel, and Hendrik Lensch. 2008. Towards passive 6D reflectance field displays. *ACM Transactions on Graphics (TOG)* 27, 3 (2008), 58.
- Rw W Gerchberg and WO Saxton. 1971. Phase determination from image and diffraction plane pictures in electron-microscope. *Optik* 34, 3 (1971), 275–+.
- Tim D Gerke and Rafael Piestun. 2010. Aperiodic volume optics. *Nature Photonics* 4, 3 (2010), 188–193.
- Daniel Glasner, Todd Zickler, and Anat Levin. 2014. A reflectance display. *ACM Trans. Graph* 33, 4 (2014), 61.
- Joseph W Goodman. 2005. *Introduction to Fourier optics*. Roberts and Company Publishers.
- Felix Heide, Qiang Fu, Yifan Peng, and Wolfgang Heidrich. 2016. Encoded diffractive optics for full-spectrum computational imaging. *Scientific Reports* 6 (2016).
- Felix Heide, James Gregson, Gordon Wetzstein, Ramesh Raskar, and Wolfgang Heidrich. 2014a. Compressive multi-mode superresolution display. *Optics Express* 22, 12 (2014), 14981–14992.
- Felix Heide, Douglas Lanman, Dikpal Reddy, Jan Kautz, Kari Pulli, and David Luebke. 2014b. Cascaded displays: spatiotemporal superresolution using offset pixel layers. *ACM Trans. Graph* 33, 4 (2014), 60.
- Ngoc-Diep Ho. 2008. *Nonnegative matrix factorization algorithms and applications*. Ph.D. Dissertation. ÉCOLE POLYTECHNIQUE.
- Reynald Hoskinson, Boris Stoeber, Wolfgang Heidrich, and Sidney Fels. 2010. Light reallocation for high contrast projection using an analog micromirror array. *ACM Trans. Graph* 29, 6 (2010), 165.
- Rafael Hostettler, Ralf Habel, Markus Gross, and Wojciech Jarosz. 2015. Dispersion-based Color Projection using Masked Prisms. In *Computer Graphics Forum*, Vol. 34. Wiley Online Library, 329–338.
- Matthias B Hullin, Ivo Ihrke, Wolfgang Heidrich, Tim Weyrich, Gerwin Damberg, and Martin Fuchs. 2012. Computational fabrication and display of material appearance. In *Eurographics 2013-State of the Art Reports*. 137–153.
- Ganghun Kim, José A Domínguez-Caballero, and Rajesh Menon. 2012. Design and analysis of multi-wavelength diffractive optics. *Optics Express* 20, 3 (2012), 2814–2823.
- Thomas Kiser, Michael Eigensatz, Minh Man Nguyen, Philippe Bompas, and Mark Pauly. 2013. *Architectural Caustics—Controlling Light with Geometry*. Springer.
- Douglas Lanman, Gordon Wetzstein, Matthew Hirsch, Wolfgang Heidrich, and Ramesh Raskar. 2011. Polarization fields: dynamic light field display using multi-layer LCDs. *ACM Trans. Graph* 30, 6 (2011), 186.
- Seungjae Lee, Changwon Jang, Seokil Moon, Jaebum Cho, and Byoungcho Lee. 2016. Additive light field displays: realization of augmented reality with holographic optical elements. *ACM Trans. Graph* 35, 4 (2016), 60.
- LB Lesem, PM Hirsch, and JA Jordan. 1969. The kinoform: a new wavefront reconstruction device. *IBM Journal of Research and Development* 13, 2 (1969), 150–155.
- Anat Levin, Daniel Glasner, Ying Xiong, Frédo Durand, William Freeman, Wojciech Matusik, and Todd Zickler. 2013. Fabricating BRDFs at high spatial resolution using wave optics. *ACM Trans. Graph* 32, 4 (2013), 144.
- Chuan-bi Lin. 2007. Projected gradient methods for nonnegative matrix factorization. *Neural Computation* 19, 10 (2007), 2756–2779.
- Adolph W Lohmann and DP Paris. 1967. Binary Fraunhofer holograms, generated by computer. *Applied Optics* 6, 10 (1967), 1739–1748.
- XF Meng, LZ Cai, YR Wang, XL Yang, XF Xu, GY Dong, XX Shen, H Zhang, and XC Cheng. 2007. Hierarchical image encryption based on cascaded iterative phase retrieval algorithm in the Fresnel domain. *Journal of Optics A: Pure and Applied Optics* 9, 11 (2007), 1070.
- Fai H Mok. 1993. Angle-multiplexed storage of 5000 holograms in lithium niobate. *Optics Letters* 18, 11 (1993), 915–917.
- Yuki Nagahama, Tomoyoshi Shimobaba, Tetsuya Kawashima, Takashi Kakue, and Tomoyoshi Ito. 2016. Holographic multi-projection using the random phase-free method. *Applied Optics* 55, 5 (2016), 1118–1123.
- Moni Naor and Adi Shamir. 1994. Visual cryptography. In *Workshop on the Theory and Application of Cryptographic Techniques*. Springer, 1–12.
- Naohisa Okada, Tomoyoshi Shimobaba, Yasuyuki Ichihashi, Ryutaro Oi, Kenji Yamamoto, Minoru Oikawa, Takashi Kakue, Nobuyuki Masuda, and Tomoyoshi Ito. 2013. Band-limited double-step Fresnel diffraction and its application to computer-generated holograms. *Optics Express* 21, 7 (2013), 9192–9197.
- Marios Pappas, Thomas Houit, Derek Nowrouzezahrai, Markus Gross, and Wojciech Jarosz. 2012. The magic lens: refractive steganography. *ACM Trans. Graph* 31, 6 (2012), 186.
- Marios Pappas, Wojciech Jarosz, Wenzel Jakob, Szymon Rusinkiewicz, Wojciech Matusik, and Tim Weyrich. 2011. Goal-based Caustics. In *Computer Graphics Forum*, Vol. 30. 503–511.
- Yifan Peng, Qiang Fu, Felix Heide, and Wolfgang Heidrich. 2016. The diffractive achromat full spectrum computational imaging with diffractive optics. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 31.
- Demetri Psaltis and Geoffrey W Burr. 1998. Holographic data storage. *Computer* 31, 2 (1998), 52–60.
- Demetri Psaltis, Kevin Curtis, Michael Levene, Allen Pu, and George Barbastathis. 1995. Holographic storage using shift multiplexing. *Optics letters* 20, 7 (1995), 782–784.
- Weidong Qu, Huarong Gu, Hao Zhang, and Qiaofeng Tan. 2015. Image magnification in lensless holographic projection using double-sampling Fresnel diffraction. *Applied Optics* 54, 34 (2015), 10018–10021.
- Behzad Sajadi, M Gopi, and Aditi Majumder. 2012. Edge-guided resolution enhancement in projectors via optical pixel sharing. *ACM Trans. Graph* 31, 4 (2012), 79.
- Yuliy Schwartzburg, Romain Testuz, Andrea Tagliasacchi, and Mark Pauly. 2014. High-contrast computational caustic design. *ACM Trans. Graph* 33, 4 (2014), 74.
- Yishi Shi, Guohai Situ, and Jinguan Zhang. 2007. Multiple-image hiding in the Fresnel domain. *Optics letters* 32, 13 (2007), 1914–1916.
- Tomoyoshi Shimobaba, Michal Makowski, Takashi Kakue, Minoru Oikawa, Naohisa Okada, Yutaka Endo, Ryuji Hirayama, and Tomoyoshi Ito. 2013. Lensless zoomable holographic projection using scaled Fresnel diffraction. *Optics Express* 21, 21 (2013), 25285–25290.
- G Tricoles. 1987. Computer generated holograms: an historical review. *Applied Optics* 26, 20 (1987), 4351–4360.
- Gordon Wetzstein, Douglas Lanman, Matthew Hirsch, and Ramesh Raskar. 2012. Tensor displays: compressive light field synthesis using multilayer displays with directional

- backlighting. *ACM Trans. Graph.* 31, 4 (2012), 80.
- Gaolei Xue, Juan Liu, Xin Li, Jia Jia, Zhao Zhang, Bin Hu, and Yongtian Wang. 2014. Multiplexing encoding method for full-color dynamic 3D holographic display. *Optics Express* 22, 15 (2014), 18473–18482.
- Genzhi Ye, Sundeep Jolly, V Michael Bove Jr, Qionghai Dai, Ramesh Raskar, and Gordon Wetzstein. 2014. Toward BxDF display using multilayer diffraction. *ACM Trans. Graph.* 33, 6 (2014), 191.
- Yonghao Yue, Kei Iwasaki, Bing-Yu Chen, Yoshinori Dobashi, and Tomoyuki Nishita. 2012. Pixel art with refracted light by rearrangeable sticks. *Computer Graphics Forum* 31, 2pt3 (2012), 575–582.
- Y. Yue, K. Iwasaki, B.-Y. Chen, Y. Dobashi, and T. Nishita. 2014. Poisson-based Continuous Surface Generation for Goal-based Caustics. *ACM Trans. Graph.* (2014).

APPENDIX

Two-step FFT process in modified GS-based phase retrieval. As mentioned in Sec. 4, we apply a double-sampling strategy for optimizing the hologram, which contains a two-step FFT process to eliminate of the sampling limitation imposed by transport with a single discrete Fourier transform (DFT). As shown in Fig. 13, this procedure uses a combination of two FFTs with a virtual intermediate plane, to replace the original single forward FFT process [Okada et al. 2013; Qu et al. 2015]. Specifically, during each step, applying single-step DFT gives two sampling rates:

$$\Delta x_v = \frac{\lambda r}{N \Delta x_0}, \Delta x = \frac{\lambda(r+z)}{N \Delta x_v}, \quad (13)$$

where Δx_v , Δx_0 , and Δx are the sampling rates on virtual intermediate plane, hologram plane, and image plane, respectively. r is the distance from sphere wave light source to hologram plane, while z is the propagated distance from hologram plane to image plane. N is the original sampling number on hologram plane in DFT process.

Thus, we can derive the relationship between Δx and Δx_0 , as follows:

$$\Delta x = \frac{(z+r)\Delta x_0}{r}. \quad (14)$$

Now we successfully derive the sampling configuration without involving λ and N , which gives us flexibility to tune different sampling rates and regions of interest on hologram plane versus image plane. This is particularly beneficial when designing holograms that expand the projection, such that a small DOE can recreate a large projected image. The calculation time for the modified double-sampling diffraction is $O(N^2 \log_2 N)$, where a hologram has the size of N by N pixels.




Fig. 13. Configuration of diffractive wave propagation under a double-sampling scheme, where the red dashed arrows indicate the first (backward) propagation, and the blue solid arrows indicate the second (forward) propagation. By relying on a virtual intermediate sampling plane, the sampling rate Δx on image plane and the sampling rate Δx_0 on hologram plane can be decoupled and efficiently tuned subject to design intentions.