

Large Étendue 3D Holographic Display with Content-adaptive Dynamic Fourier Modulation

BRIAN CHAO, Stanford University, USA
MANU GOPAKUMAR, Stanford University, USA
SUYEON CHOI, Stanford University, USA
JONGHYUN KIM, NVIDIA, USA
LIANG SHI, Massachusetts Institute of Technology, USA
GORDON WETZSTEIN, Stanford University, USA

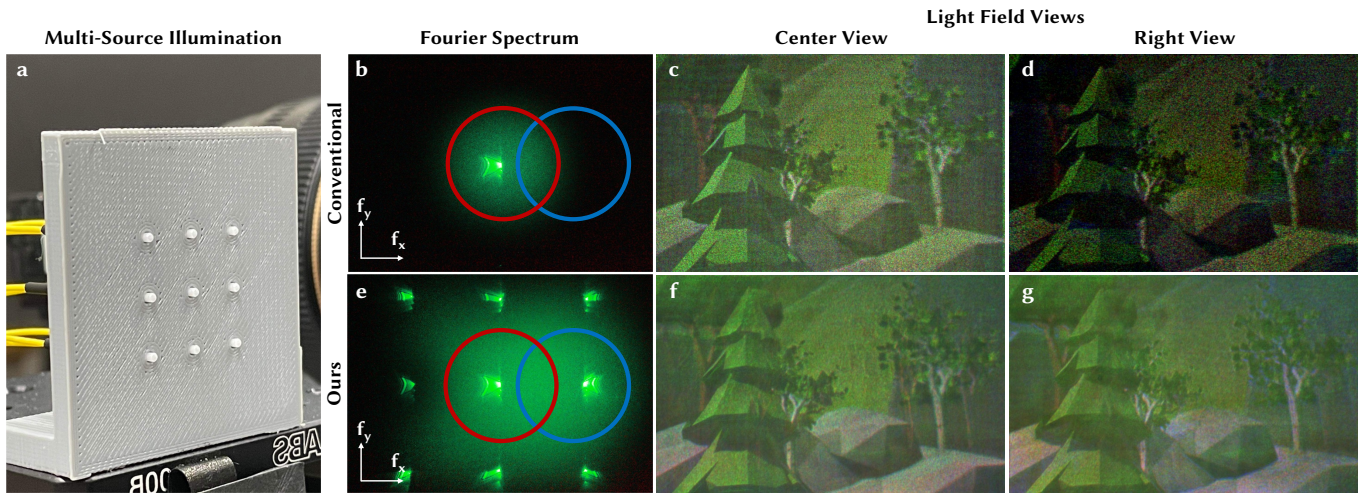


Fig. 1. Conventional holographic displays use a single laser source that provides a limited étendue, here visualized by a recorded spectrum that only covers a small area of the Fourier plane (b). Single-source holograms therefore only support a limited eyebox size, which means that an image can be observed only when the user’s pupil (b, red) is well aligned with the eyebox (c). The image quickly degrades and fades into black as the pupil (b, blue) shifts even a small amount (d). Using multi-source illumination (a), our holographic display creates a significantly expanded coverage of addressable spatial frequencies (e) which, combined with our content-adaptive Fourier modulation strategy, achieves a large étendue with better image quality across an expanded eyebox (f,g).

Emerging holographic display technology offers unique capabilities for next-generation virtual reality systems. Current holographic near-eye displays, however, only support a small étendue, which results in a direct tradeoff between achievable field of view and eyebox size. Étendue expansion has recently been explored, but existing approaches are either fundamentally limited in the image quality that can be achieved or they require extremely

high-speed spatial light modulators. We describe a new étendue expansion approach that combines multiple coherent sources with content-adaptive amplitude modulation of the hologram spectrum in the Fourier plane. To generate time-multiplexed phase and amplitude patterns for our spatial light modulators, we devise a pupil-aware gradient-descent-based computer-generated holography algorithm that is supervised by a large-baseline target light field. Compared with relevant baseline approaches, ours demonstrates significant improvements in image quality and étendue in simulation and with an experimental holographic display prototype.

Authors’ addresses: Brian Chao, Stanford University, Stanford, CA, 94305, USA, brianhc@stanford.edu; Manu Gopakumar, Stanford University, Stanford, CA, 94305, USA, manugopa@stanford.edu; Suyeon Choi, Stanford University, Stanford, CA, 94305, USA, suyeon@stanford.edu; Jonghyun Kim, NVIDIA, Santa Clara, CA, 95051, USA, jonghyunk@nvidia.com; Liang Shi, Massachusetts Institute of Technology, Cambridge, MA, 02139, USA, liangs@mit.edu; Gordon Wetzstein, Stanford University, Stanford, CA, 94305, USA, gordon.wetzstein@stanford.edu.

CCS Concepts: • **Hardware** → **Emerging technologies**; • **Computing methodologies** → **Graphics systems and interfaces**.

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 the author(s) 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](https://permissions.acm.org).

Additional Key Words and Phrases: computational displays, holography, virtual reality

SA Conference Papers ’24, December 3–6, 2024, Tokyo, Japan
© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-1131-2/24/12
<https://doi.org/10.1145/3680528.3687600>

ACM Reference Format:

Brian Chao, Manu Gopakumar, Suyeon Choi, Jonghyun Kim, Liang Shi, and Gordon Wetzstein. 2024. Large Étendue 3D Holographic Display with Content-adaptive Dynamic Fourier Modulation. In *SIGGRAPH Asia 2024 Conference Papers (SA Conference Papers ’24)*, December 3–6, 2024, Tokyo, Japan. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3680528.3687600>

1 INTRODUCTION

Holographic near-eye displays offer unique benefits to virtual and augmented reality (VR/AR) applications. For example, holographic displays can present perceptually realistic 3D images with natural parallax to the user in lightweight device form factors [Gopakumar et al. 2024; Jang et al. 2024; Kim et al. 2022; Maimone et al. 2017]. Yet, the étendue of holographic displays is fundamentally limited by the pixel count of the underlying spatial light modulators (SLMs), preventing current holographic near-eye displays from achieving a large field of view and eyebox simultaneously. This limitation is a fundamental barrier towards making this a practical display technology.

Increasing the pixel count of an SLM seems like the natural solution. However, developing large-area phase-only SLMs with pixel pitches matching the small feature sizes (i.e., tens of nanometers) of analog holographic films [Benton and Bove Jr 2008] is simply not feasible with today’s hardware solutions. To overcome this problem, étendue expansion techniques have been described in the literature, including those based on static, high-resolution masks [Buckley et al. 2006; Kuo et al. 2020; Monin et al. 2022a; Park et al. 2019; Tseng et al. 2024; Yu et al. 2017], pupil replication [Kress and Chatterjee 2020], steered or multi-source illumination [Jang et al. 2018; Jo et al. 2022; Lee et al. 2020, 2022; Monin et al. 2022b], and making use of higher-diffraction orders and pupil optimization [Schiffers et al. 2023; Shi et al. 2024]. However, each of these approaches has its limitations, as mask-based systems do not have sufficient degrees of freedom to achieve a high image quality, pupil replication approaches cannot create natural 3D effects and parallax over the eyebox, and steered sources are hindered by the requirement for high-speed SLMs as well as high diffraction orders (HDOs) that fundamentally limit the image quality. As a result, none of these solutions is able to achieve high-quality 3D holography with a large étendue.

Our work is motivated by the hypothesis that a holographic display requires sufficient degrees of freedom to achieve a large field of view and eyebox simultaneously. In the absence of an extremely high-resolution SLM, this is only achievable using steered or multi-source illumination. We thus build on the latter approach but address its major shortcomings, HDOs and symmetric illumination copies, by introducing a dynamic, programmable amplitude modulation mechanism in the Fourier plane, after the SLM. This unique optical setup allows us to extend steered / multi-source configurations such that they modulate the frequency spectrum of the display image in a content-adaptive manner. For this purpose, we leverage a stochastic optimization approach that factors a target light field into a set of time-multiplexed phase SLM and corresponding Fourier amplitude masks that are displayed in rapid succession while being integrated by the user’s eye.

Using the proposed system, we demonstrate improved 3D image quality over a large étendue, surpassing the performances of existing approaches in both simulation and experiment. Specifically, our contributions include

- A novel optical holographic display configuration that combines a time-multiplexed phase SLM near the image plane and a dynamic amplitude SLM that controls the frequency spectrum.

- A computer-generated holography framework that uses stochastic optimization to factor a target light field into a set of phase–amplitude image pairs.
- Demonstration of improved 3D image quality among high-étendue holographic displays.

Our method should be clearly distinguished from Multisource Holography, a system recently proposed by Kuo et al. [2023] for speckle reduction that also leverages a multi-source laser array. In Multisource Holography, the multi-source laser and two *phase-only* SLMs placed in close proximity are used to remove speckles, but the system étendue remains limited since the spacing between each source is relatively small. In our system, we place the laser sources much farther apart to create a high-étendue backlight for the phase-only SLM to greatly increase the eyebox size and place an *amplitude* display at the Fourier plane.

2 RELATED WORK

Holographic Near-eye Displays. Holographic displays are a promising technology for virtual and augmented reality applications due to their unique capability to display true 3D content and significant progress has been made recently [Chang et al. 2020; Javidi et al. 2021; Pi et al. 2022]. In particular, the advancement in computer graphics, machine learning, and computing infrastructures have enabled real-time hologram rendering based on neural networks [Peng et al. 2020; Shi et al. 2021], significantly improved image quality with end-to-end optimization [Chakravarthula et al. 2020; Choi et al. 2022; Peng et al. 2020], higher light efficiency and brightness with simultaneous control of multiple wavelengths and energy-efficiency loss function [Chao et al. 2023; Kavaklı et al. 2023; Markley et al. 2023], and thin form factors in eyeglasses-like design [Gopakumar et al. 2024; Jang et al. 2024; Kim et al. 2022; Maimone et al. 2017]. Despite offering these unique capabilities, current holographic displays fail to provide a comfortable immersive experience as they cannot simultaneously provide a wide field of view (FoV) and a sufficiently large eyebox (i.e., the region in which a user’s eye perceives the displayed content).

In a given display system, the product of the FoV and the eyebox is a constant, referred to as the *étendue*. For a holographic display, the étendue is directly proportional to the number of pixels in the SLM. A 1080p SLM, for instance, can either support a wide field of view (e.g., 80 degrees) with an eyebox smaller than 1 mm or vice versa. However, increasing the SLM resolution to the point where both large field of view and eyebox can be achieved simultaneously faces significant challenges in manufacturing, cost, and addressing speed, accuracy, and bandwidth. Instead, efforts have been made to increase the étendue of holographic displays without increasing SLM resolution. The approaches fall under two categories: (i) increasing étendue *after the SLM* and (ii) increasing étendue *before the SLM*.

Post-SLM Étendue Expansion. The most representative method in this category is mask-based étendue expansion [Buckley et al. 2006; Kuo et al. 2020; Monin et al. 2022a; Park et al. 2019; Tseng et al. 2024; Yu et al. 2017], where a static mask at a resolution higher than the SLM is placed after the SLM to increase the diffraction angles of the SLM-modulated wavefront, thus increasing the étendue. However, such systems suffer from difficulties in alignment and

from reduced image quality and low contrast because their effective degrees of freedom [Starikov 1982] are insufficient to synthesize a high-quality large-étendue wavefront. Pupil replication [Bigler et al. 2018, 2019; Draper et al. 2019; Draper and Blanche 2021, 2022; Jang et al. 2024; Kress and Chatterjee 2020; Park and Kim 2018] is another popular approach. It is implemented either by putting a pupil-replicating waveguide after the SLM to replicate pupil locations at its out-coupler or using a pupil-replicating holographic optical element (HOE) as the eyepiece, effectively expanding the eyebox of the system. However, pupil-replicating displays cannot display 3D content or natural parallax across the expanded eyebox since the content within the eyebox are merely copies of the same wavefront. Higher-diffraction orders combined with pupil optimization can also be leveraged to slightly expand the eyebox in the single-source case [Schiffers et al. 2023; Shi et al. 2024]. Finally, a regular eyepiece can be replaced by a lens array to partition an unexpanded eyebox into an array of smaller chunks that cover an expanded area [Chae et al. 2023; Wang et al. 2023; Xia et al. 2020]. However, this comes at an explicit cost of image quality and brightness nonuniformity, especially when observed with a small pupil.

Pre-SLM Étendue Expansion. Methods in this category modify the laser illumination to expand étendue either through a multi-source configuration or beam-steering. Jang et al. [2018] used a micro-electromechanical-system (MEMS) mirror to temporally change the laser illumination and steer the resulting pupils over a larger eye box. Lee et al. [2020] implemented the same principle by arranging individual laser diodes into a 2D array and sequentially turning each one on to create temporal directional illumination. Monin et al. [2022b] implemented *per-pixel* beam steering of the phase SLM by using transmissive LCD panels and polarization gratings and demonstrated that the étendue expansion amount scales exponentially with the number of LCD layers. To permanently expand the étendue, Jo et al. [2022] activated all illumination sources simultaneously. They introduced a random mask at the Fourier plane to break the correlation among copies in the spectrum formed by directional illuminations. This effectively eliminates duplicate images within the expanded eyebox. However, they did not demonstrate view-dependent effects across the expanded eyebox and the random mask is not content adaptive, resulting in reduced 3D realism and low image quality. Instead of using multiple laser diodes that are incoherent with each other, Lee et al. [2022] implemented a mutually coherent multi-laser source using a lens array. The mutually coherent sources can interfere constructively and destructively with each other, granting the hologram optimization process more degrees of freedom. However, their system requires eye tracking and a new hologram needs to be optimized for each dynamic pupil location, making the system challenging for real-time applications.

Fourier Modulation. Holographic display systems often require Fourier plane filtering to remove HDOs created by the pixelated structure of the phase SLM [Maimone et al. 2017; Peng et al. 2020; Shi et al. 2021, 2022]. However, it is not straightforward to apply this to a multi-source or beam steering setting since the directional illuminations create shifted copies of the wavefront from normally




Fig. 2. **System Architecture.** Multiple mutually incoherent sources illuminate a fast phase-only SLM, creating a high-étendue backlight. An additional amplitude display is placed at the Fourier plane to remove ghost image artifacts created by HDOs and the multisource illumination. The phase and amplitude patterns are optimized jointly for a target light field in a content adaptive manner. We illustrate our design following the style of the schematic in [Kuo et al. 2023] for easier comparison.

incident illumination and the associated HDOs in the Fourier domain. When using beam-steering, the filter position needs to be dynamically adjusted to block the HDOs of the shifted wavefront. To achieve this, Lee et al. [2020] placed a programmable polarization shutter at the Fourier plane and synchronized the laser sources to filter out the HDOs. However, when using multi-source illumination for étendue expansion, HDOs associated with one illumination intermingle with the shifted wavefront of another illumination at the Fourier domain, making it impossible to separate. It is therefore crucial to model HDOs to precisely characterize how they contribute to different angular views. Explicit modeling of HDOs has been demonstrated for 2D images [Gopakumar et al. 2021] and 3D focal stacks [Kim et al. 2022; Shi et al. 2024], but not for 4D light fields under multi-source illumination.

Inspired by Jo et al. [2022], we employ a multisource laser illumination and additionally place a dynamic amplitude SLM at the Fourier plane to enable content adaptive modulation and time multiplexing. We jointly optimize the patterns for the phase and amplitude SLMs to reproduce a 4D light field rendered over an expanded eyebox. We also explicitly model the HDOs and demonstrate notable improvement in image quality and contrast. Collectively, this new hardware and software co-design enables a dynamic view-dependent holographic display with a large eyebox.

3 METHOD

In this section, we first review the conventional single-source holographic image formation model before introducing the multi-source image formation model of our system.

3.1 Single-Source Holographic Image Formation Model

For on-axis Fresnel holography, a collimated beam from a laser source illuminates an SLM with a normally incident, coherent field u_{src} . The SLM imparts a spatially varying phase delay ϕ to the field which propagates a distance z along the optical axis. The wavefront at this plane can be mathematically described using the angular spectrum method (ASM) [Goodman 2005] as a function of the phase

pattern and distance from the SLM:

$$f(\phi, z) = \mathcal{F}^{-1} \left\{ \mathcal{F} \left\{ e^{i\phi(x,y)} u_{\text{src}}(x, y) \right\} \cdot \mathcal{H}(f_x, f_y; z) \right\}$$

$$\mathcal{H}(f_x, f_y; z) = \begin{cases} e^{i\frac{2\pi}{\lambda} z \sqrt{1 - (\lambda f_x)^2 - (\lambda f_y)^2}} & \text{if } \sqrt{f_x^2 + f_y^2} < \frac{1}{\lambda}. \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

Here, λ is the wavelength of light, x, y are the spatial coordinates on the SLM, f_x, f_y are the frequency coordinates, and \mathcal{H} is the transfer function of the ASM. The operator f models free-space propagation between the parallel SLM and target planes separated by a distance z . For notational convenience, we omit the dependence of the fields on x and y . The intensity generated by a holographic display at a distance z in front of the SLM is therefore $|f(\phi, z)|^2$. If a high-speed SLM is available, a time-multiplexed variant of the image formation is $\sum_{t=1}^T |f(\phi^{(t)}, z)|^2 / T$, where T phase SLM patterns $\phi^{(t)}$, $t = 1, \dots, T$ are rapidly displayed in sequence, and the resulting intensities are averaged by the users' eye [Choi et al. 2022].

3.2 Multi-Source Holographic Image Formation Model with Fourier Modulation

To extend the single-source image formation model to our system, we modify the formulation to incorporate off-axis collimated illumination traveling in direction $\mathbf{k} = (k_x, k_y, k_z)$ and a programmable amplitude mask \mathcal{P} at the Fourier plane of the holographic display system. This results in the model

$$f^{(j)}(\phi, \mathcal{P}, z) = \mathcal{F}^{-1} \left\{ \mathcal{U}_{\text{slm}}^{(j)}(f_x, f_y; \phi) \cdot \mathcal{H}(f_x, f_y; z) \cdot \mathcal{P}(f_x, f_y) \right\}$$

$$\mathcal{U}_{\text{slm}}^{(j)}(f_x, f_y; \phi) = \mathcal{F} \left\{ e^{i\phi(x,y)} u_{\text{src}}^{(j)}(x, y) e^{i\mathbf{k}^{(j)} \cdot \mathbf{x}} \right\} \quad (2)$$

where j is the index of the source, and $u_{\text{src}}^{(j)}(x, y)$ is the complex-valued field modeling any deviations in amplitude and phase of source $j = 1, \dots, J$ from a perfect plane wave $e^{i\mathbf{k}^{(j)} \cdot \mathbf{x}}$, $\mathbf{x} = (x, y, z)$. Moreover, $u_{\text{src}}^{(j)}(x, y)$ can optionally also include per-source, time-dependent modulation, such as switching individual lasers on and off. In our setup, we do not consider this case and assume that all sources are turned on at all times. Please refer to the supplemental material for more discussions about the generalized configuration with amplitude-controllable laser sources $u_{\text{src}}^{(j)}(x, y)$.

3.3 Stochastic Optimization of Light Field Holograms

To reconstruct a light field, we use gradient descent to optimize a set of time-multiplexed phase patterns $\phi^{(t)}$ and corresponding Fourier masks $\mathcal{P}^{(t)}$ by minimizing the following objective:

$$\underset{\{\phi^{(t)}, \mathcal{P}^{(t)}\}}{\text{minimize}} \left\| s \sqrt{\frac{1}{T} \sum_{t=1}^T \sum_{j=1}^J |\text{H2LF}(f^{(j)}(\phi^{(t)}, \mathcal{P}^{(t)}, z))|^2 - \text{If}_{\text{target}}} \right\| \quad (3)$$

where s is a scale factor, $\text{If}_{\text{target}}$ is the amplitude of the target light field, and H2LF is a hologram-to-light field transformation, such as the Short-Time Fourier Transform (STFT) [Padmanaban et al. 2019; Zhang and Levoy 2009].

The memory consumption of the above optimization problem is huge due to time multiplexing, multiple sources, and the explicit modeling of HDOs. Therefore, it is impractical to realize H2LF using the STFT since it reconstructs a whole light field and the memory consumption would explode for a dense $\text{If}_{\text{target}}$. To solve this problem, we devise a *stochastic* version of Eq. 3 that allows us to optimize a single light-field view rather than a full light field in each iteration of the optimization routine. For this purpose, we randomly chose a view p of the target light $\text{If}_{\text{target}}^{(p)}$ in each iteration and run a gradient descent step of Eq. 3. A binary pupil mask $\mathcal{M}^{(p)}$ in the Fourier plane in the hologram-to-light field transform is applied to reconstruct one specific view as

$$\text{H2LF}^{(p)} \left(f^{(j)}(\phi^{(t)}, \mathcal{P}^{(t)}, z) \right) = f^{(j)}(\phi^{(t)}, \mathcal{P}^{(t)} \cdot \mathcal{M}^{(p)}, z),$$

$$\mathcal{M}^{(p)}(f_x, f_y) = \begin{cases} 1, & \text{if } (f_x - c_{x,p})^2 + (f_y - c_{y,p})^2 \leq r_p^2, \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

where $\mathcal{M}^{(p)}$ is a binary pupil mask in the Fourier plane, r_p is the radius of the pupil and $c_{x,p}, c_{y,p}$ are the spatial coordinates of the center of the pupil. This procedure is similar to the pupil-supervision techniques described in [Chakravarthula et al. 2022; Schiffrers et al. 2023; Shi et al. 2024]. Please refer to the supplemental material for more details on our stochastic light field optimization procedure.

3.4 Implementation Details

Since we are using a highly-quantized 4-bit phase SLM, the quantization of pixel values needs to be taken into consideration. Such quantization constraints can be enforced using techniques described in prior work [Choi et al. 2022]. Higher diffraction orders (HDOs) are modeled using the wave propagation model described in [Gopakumar et al. 2021]. We use PyTorch to implement all our algorithms and run optimization.

In all our experiments, the radius r_p of the pupils are set to be 2 mm, resulting in a 4 mm diameter pupil. 81 pupils are equally spaced in the Fourier plane (eyebow plane), where each pupil corresponds to a single view in a 9×9 light field. The illumination directions of the multisource laser are set such that they match the diffraction angle of the $\pm 1^{\text{st}}$ higher diffraction orders of the blue wavelength. This allows the blue spectrum copies to be perfectly tiled in the Fourier plane, while removing the gaps between the red and green wavelength spectrum copies. Please see the supplemental materials for detailed discussion on the choosing the appropriate illumination angles.

4 ANALYSIS

4.1 Optical System Analysis

The *étendue* G of a display is defined as the product of the display area with the solid angle of emitted light:

$$G = 4A \sin^2 \theta, \quad (5)$$

where A is the display area and 2θ is the solid angle of the emission cone of each display pixel. Étendue is conserved through reflections, refractions, and free space propagation in an optical system. When illuminated with a normal incidence light of wavelength λ , the




Fig. 3. **Tradeoff between 2D field of view (FoV) and eyebox size.** Each white line represents the fixed étendue of a holographic display system illuminated by a grid of $\alpha \times \alpha$ sources with different α values. We show the étendue of the systems in log-scaled color maps. As the number of sources increase, the étendue of the system also increases.

diffraction angle θ_{SLM} of an SLM with pixel pitch p can be expressed as $\theta_{\text{SLM}} = \pm \sin^{-1} \frac{\lambda}{2p}$.

For an SLM of physical size $L_x \times L_y$, its étendue G_{SLM} can be expressed as:

$$G_{\text{SLM}} = 4L_x L_y \sin^2 \theta_{\text{SLM}} = \lambda^2 N_x N_y \quad (6)$$

where $N_x \times N_y$ is the pixel resolution of the SLM. This means that the étendue of a holographic display is directly proportional to the number of pixels of the SLM.

In a Fresnel holography display system, the 1D field-of-view (FoV) and eyebox size w can be expressed as follows:

$$\text{FoV} = 2 \tan^{-1} \left(\frac{L}{2g} \right) = 2 \tan^{-1} \left(\frac{Np}{2g} \right), \quad w = \frac{g\lambda}{p}, \quad (7)$$

where L, N are the size of the SLM and the number of SLM pixels in the x or y axis, respectively, and g is the eyepiece focal length. Under paraxial assumptions ($\theta \approx \sin \theta \approx \tan \theta$), we see that the product of the 2D FoV and eyebox of a holographic display system is exactly the étendue of the system:

$$\text{FoV}_x \cdot \text{FoV}_y \cdot w^2 = 2 \tan^{-1} \left(\frac{N_x p}{2g} \right) \cdot 2 \tan^{-1} \left(\frac{N_y p}{2g} \right) \cdot \frac{g^2 \lambda^2}{p^2} \approx \lambda^2 N_x N_y = G_{\text{SLM}} \quad (8)$$

This implies that there is an inherent tradeoff between the FoV and the eyebox of a holographic display.

When the SLM is illuminated with a grid of $\alpha \times \alpha$ off-axis, directional illuminations, the system eyebox is expanded due to shifted copies of the original spectrum. Specifically, if the directional illumination is selected such that the illumination direction matches the higher-order diffraction angles, the system 1D eyebox is exactly expanded by α while the FoV remains the same, resulting in an expanded 1D eyebox size of $w = \frac{\alpha f \lambda}{p}$. Therefore, the 2D étendue of the system is expanded by a factor of α^2 .

We show how the FoV and eyebox size relates to the required number of sources in Fig. 3. We assume an SLM pixel pitch of $10.8 \mu\text{m}$ and resolution of 1000×1000 and laser wavelength of 632.8 nm . The FoV and eyebox size move along each white line in opposite directions as we vary the eyepiece focal length g while the system étendue remains fixed. As we increase the number of sources, the

étendue of the system also increases, as the white lines move further towards to upper-right of the plot.

4.2 Baseline Configurations

We next discuss a number of holographic display system configurations that serve as baselines to our proposed design shown in Fig. 2. Illustrations of these baselines are shown in Fig. 4.

I. Single Source with Fourier Filter. The conventional holographic display setup with a single laser source and a Fourier filter to block HDOs, including [Choi et al. 2022; Maimone et al. 2017; Peng et al. 2020; Shi et al. 2021]. Such systems suffer from small étendue and non-uniform brightness across the eyebox.

II. Single Source with Phase Mask. A high-resolution phase mask is placed in front of the SLM to increase the diffraction angle of the SLM and therefore increase the étendue of the system. The phase masks can be random [Buckley et al. 2006; Kuo et al. 2020; Park et al. 2019; Yu et al. 2017] or optimized [Monin et al. 2022a; Tseng et al. 2024]. These approaches have been shown to expand the étendue at the cost of decreased image quality and contrast.

III. Multiple Sources. Multiple mutually incoherent lasers illuminate the SLM from different angles *simultaneously*. Due to the absence of a Fourier filter, the frequency spectrum contains multiple shifted, potentially overlapping copies of the same hologram. These constraints limit this system's capability to perfectly reconstruct a light field.

IV. Multiple Sources with Fixed Random Fourier Mask. Multiple lasers illuminate the SLM from different angles *simultaneously* while a fixed random mask is placed at the Fourier plane to break the correlation between the image copies, as demonstrated by Jo et al. [2022]. Time multiplexing and content-adaptive filtering are not feasible since the random masks are custom-printed and fixed.

V. Multiple Sources with Dynamic Fourier Filter (ours). Multiple lasers illuminate the SLM from different angles *simultaneously* while an amplitude SLM is placed at the Fourier plane. The amplitude SLM can be dynamically refreshed and is synchronized with the phase SLM, allowing for time-multiplexed and content-adaptive Fourier modulation. We additionally compare with a generalized configuration V^* where the amplitude of the laser sources are controllable rather than fixed. More discussions on this generalized configuration can be found in the supplement.

VI. Steered Illumination. Multiple individually controllable sources or a single, swept source illuminate the SLM from different angles without Fourier filtering, including [An et al. 2020; Chang et al. 2019; Jang et al. 2018; Reichelt et al. 2010; Xia et al. 2023]. Reduced image quality due to HDOs remains an issue due to the lack of filtering. Furthermore, such methods only apply to very high-speed SLMs, because each source is sequentially turned on or steered in sequence.

VII. Steered Illumination with Shifting Fourier Filter. Multiple individually controllable sources or a single, swept source illuminate the SLM from different angles while a synchronizable, dynamic filter is placed at the Fourier plane to filter out the HDOs. One example is




Fig. 4. **Illustration of baseline display configurations.** Single-source configurations (I, II) trade image quality off for étendue expansion factor (II). Multi-source approaches that use all sources simultaneously (III, IV) benefit from a high-étendue “backlight” but operate within limited effective degrees of freedom, which also makes it challenging to achieve a high image quality. Steered illumination approaches (VI, VII) sequentially illuminate the system from different directions and require high-speed SLMs. Without a Fourier mask, the image quality achieved by these systems is also limited by high diffraction orders (VI). We illustrate our design following the style of the schematic in [Kuo et al. 2023] for easier comparison.

the steered illumination system described in [Lee et al. 2020]. This method is still sequential in nature, as each source is turned on one at a time, and requires a very high-speed SLM.

4.3 Assessment

Table 1 and Figure 6 shows the light field reconstruction performance of different baseline configurations in simulation. For configurations with multiple sources (III, IV, V, V*, VI, and VII), we consider a 3×3 grid of sources. We simulate a 800×1280 phase SLM for single-SLM configurations (I, II, III, VI, VII) and an additional 20×20 Fourier display for configurations IV, V, and V*. We run our optimization algorithm on all configurations to reconstruct a 9×9 light field. Additionally, we optimize the single source configuration to reconstruct a smaller, 3×3 light field in Table 1. Time multiplexing is not used for configurations I, II, III, and IV. Please refer to the supplemental material for additional discussions on the optimization parameters, degrees of freedom of the system, and ablation studies on the resolution of the Fourier display.

The naive single-source configuration (I) is able to reconstruct a small 3×3 light field, but fails to reconstruct a larger-baseline 9×9 light field and cannot support uniform brightness across the expanded eyepiece (Fig. 1, b–d). Mask-based étendue expansion techniques (II) reconstruct low-contrast and speckly images. By using multiple sources (III), the eyepiece is expanded but the light field reconstruction quality is poor due to copies created by multiple sources and HDOs. Introducing a fixed random mask at the Fourier plane (IV) improves image quality, although the improvement is limited due to the lack of time multiplexing and content-adaptive Fourier mask optimization. Steered illumination options (VI, VII) achieve decent image quality, however both configurations reconstruct speckly images due to HDOs and can only be implemented using high-speed SLMs due to the large number of required time-multiplexed frames.

Our methods (V, V*) achieves the best image reconstruction quality when using one frame and is better than the steered illumination baseline (VII) while using fewer frames (6 vs. 9). This is achieved through time multiplexing and our novel content-adaptive Fourier modulation optimization framework. Our method successfully removes the copies created by multiple sources and HDOs, reconstructing clean and speckless light field views. More importantly, a minimal increase in degrees of freedom in the Fourier plane (a low-resolution 20×20 Fourier display) is sufficient to achieve good image quality, and we perform extensive experiments to validate this claim in the supplemental material. Finally, although our generalized configuration V* with amplitude-controllable sources achieves the best quantitative image quality, the improvement is marginal (<0.5 dB in terms of PSNR) and suffers from a much higher system complexity. Hence, we opted for configuration V for our hardware implementation.

5 EXPERIMENTAL RESULTS

Hardware Implementation. We implement the proposed 3D holographic display design and evaluate our algorithms on the system. The hardware setup and the optical path are shown in Fig. 5. We implement our multi-source laser by cascading multiple 1:4 fiber splitters (Thorlabs TWQ560HA) and arranging 9 customized fiber tip outputs into a 3×3 array, which is then held together using a custom-printed 3D mount. The spacing between each source is 8.17 mm and a 200 mm lens is used to collimate the multi-source laser. Each collimated source field is, therefore, incident on the phase SLM with a 2.34° incident angle. We use a TI DLP6750Q1EVM phase SLM for phase modulation and a 1080p SiliconMicroDisplay liquid crystal on silicon (LCoS) display for Fourier amplitude modulation. A 75 mm Fourier transform lens is used to image the spectrum of the phase-modulated wavefront onto the amplitude SLM. Our final design has a diagonal field-of-view (FoV) of 7.78 degrees and an eyepiece size of 8.53 mm \times 8.53 mm. A FLIR Grasshopper 2.3 MP

Configuration	Image Quality (PSNR/SSIM)
I. Single Source with Fourier Filter [Choi et al. 2022; Peng et al. 2020; Shi et al. 2021]	46.23 / 0.97 (2D image) 22.43 / 0.46 (3×3 light field) 16.60 / 0.23 (9×9 light field)
II. Single Source with Phase Mask [Kuo et al. 2020]	14.61 / 0.15
III. Multiple Sources	13.35 / 0.24
IV. Multiple Sources with Fixed Random Fourier Mask [Jo et al. 2022]	18.63 / 0.26
V. Multiple Sources with Dynamic Fourier Filter (ours)	21.05 / 0.39 (1 frame) 24.65 / 0.66 (6 frames) 25.17 / 0.71 (9 frames)
V*. Multiple Sources with Dynamic Fourier Filter with Laser Amplitude Control	21.43 / 0.42 (1 frame) 25.02 / 0.69 (6 frames) 25.66 / 0.74 (9 frames)
VI. Steered Illumination	20.02 / 0.43 (9 frames)
VII. Steered Illumination with Shifting Fourier Filter [Lee et al. 2020]	23.95 / 0.51 (9 frames)

Table 1. **Baseline comparisons in simulation.** Light field reconstruction performance of different baseline configurations in terms of PSNR/SSIM. We run our optimization algorithm on all configurations and six test scenes to reconstruct 9×9 light fields and report the average reconstruction performance. We run additional optimizations on the single source configuration to reconstruct a 2D image and a smaller 3×3 light field. Our method achieves the best reconstruction quality for the full light field while using fewer frames. Per-scene PSNR/SSIM can be found in the supplemental materials.




Fig. 5. **Photograph of our multi-source holographic display prototype.** The propagation path is illustrated in red and components are labeled.

color USB3 vision sensor paired with a Canon EF 50mm f/1.4 USM camera lens is used to capture all experimental results. Please refer to the supplemental material for additional details on the degrees of freedom of our system and the relevant optimization parameters.

Experimental Capture Details. To capture light field views, we place pupil masks at the Fourier plane to mimic the movement of the user’s eyes, which is a technique used in prior works [Schiffers et al. 2023; Shi et al. 2024]. We implement this with a Thorlabs SM1D12 adjustable iris on a translation stage at the Fourier plane. To capture focal stacks, we center the pupil at the Fourier plane and adjust the camera focus to capture images at different depths.

Assessment. Experimentally captured results are shown in Figs. 1, 7, Table 2, and in the supplemental material. The PSNR and SSIM

Configuration	Image Quality (PSNR/SSIM)
III. Multiple Sources	12.72 / 0.19
IV. Multiple Sources with Fixed Random Fourier Mask [Jo et al. 2022]	12.46 / 0.17
V. Multiple Sources with Dynamic Fourier Filter (ours)	13.83 / 0.23 (1 frame) 14.43 / 0.40 (6 frames)

Table 2. **Experimentally captured baseline comparisons.** Experimentally captured light field reconstruction performance of different baseline configurations in terms of PSNR/SSIM, averaged across all six test scenes. Our method achieves the best experimental image quality. Per-scene PSNR/SSIM can be found in the supplemental materials.

values are averaged across all captured light field views. We observe the same trends as predicted by our simulations both quantitatively and qualitatively: the single-source configuration only supports a limited eyebox and suffers from severe brightness falloff at peripheral viewpoints; 3D multi-source holography without a Fourier filter cannot achieve a high image quality due to the copies created the multiple sources; a static random mask placed in the Fourier plane only provides limited degrees of freedom and suffers from low contrast; our approach without time multiplexing (i.e., 1 frame) improves the quality over the random mask as it optimizes the amplitude mask pattern in a content-adaptive manner; our method with 6-frame time multiplexing achieves the highest image quality with the largest amount of empirically observed parallax.

6 DISCUSSION

In summary, we present a novel hardware system for étendue expansion and an algorithmic framework for 4D light-field-supervised computer-generated holography. The hardware system includes a multi-source laser array to create a large-étendue coherent backlight for the phase SLM and an amplitude SLM for dynamic Fourier-amplitude modulation. The algorithmic framework includes the joint optimization of time-multiplexed amplitude SLM and phase SLM patterns and a memory-efficient, stochastic light field supervision procedure to create 4D light field holograms. We compare our method with a number of étendue expansion baselines and verify in simulation and experimentally that our system achieves the highest-quality light field reconstruction results for large étendue settings.

Limitations and Future Work. We demonstrate our results on a benchtop display setup but further efforts are required to miniaturize this system. Currently, our multisource laser array is implemented using bulky fiber splitters and could be miniaturized using nanophotonic phased arrays [Sun et al. 2013]. Folding the propagation distance of holograms using optical waveguides could further remove the need of beam splitters and subsequently shrink the form factor, as demonstrated in [Jang et al. 2024; Lin et al. 2018, 2020; Yeom et al. 2021]. We illustrate potential compact designs in the supplemental material. The frame rate of our system is limited by our amplitude display (240 Hz native frame rate). This translates to a ~ 13.33 Hz frame rate when operating in color-sequential mode with 6-frames time-multiplexing. The frame rate can be improved

by using more advanced LCoS displays with frame rate > 720 Hz [Lazarev et al. 2017]. Real-time synthesis of light field holograms are necessary for practical holographic displays, but is not currently not supported by our system. Extending recent neural network-based hologram synthesis methods [Shi et al. 2021, 2022; Yang et al. 2022] to work for 4D light field holograms would be an interesting future direction. Finally, we did not attempt to calibrate a neural network-parameterized wave propagation model of our prototype display system, which has been demonstrated to significantly improve experimentally captured holographic image quality for other types of optical configurations [Choi et al. 2022; Peng et al. 2020].

Conclusion. The novel hardware design and algorithmic framework presented in this work improves the étendue of holographic displays and allows for light field holograms synthesis with improved image quality. These help make holographic displays a more practical technology for augmented and virtual reality applications.

ACKNOWLEDGMENTS

We thank Grace Kuo for helpful advice regarding the implementation of the multisource laser setup. Brian Chao is supported by the Stanford Graduate Fellowship and the NSF GRFP. Manu Gopakumar is supported by the Stanford Graduate Fellowship. Suyeon Choi is supported by the Meta Research PhD Fellowship.

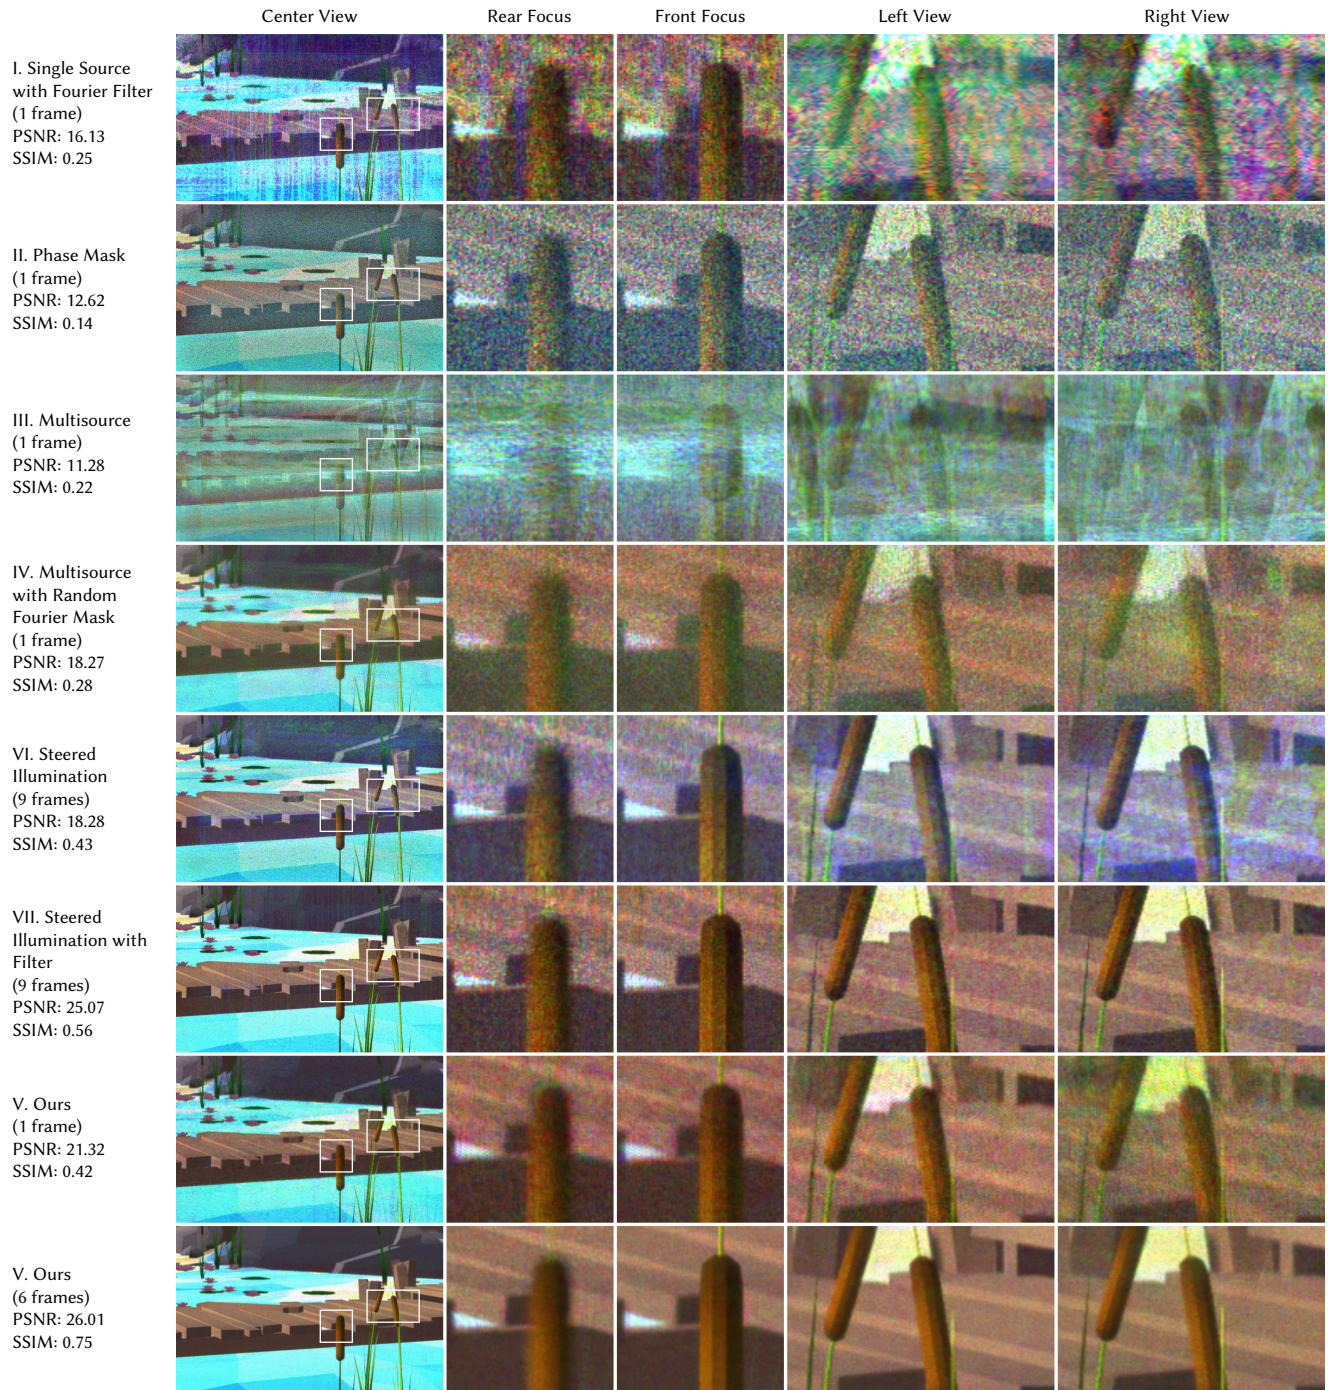


Fig. 6. **Comparison of hardware configurations using simulated reconstruction.** Here, we compare different étendue-expanded holographic display configurations, including the conventional setup with light-field supervision (1st row), phase masks [Kuo et al. 2020] (2nd row), multisource (3rd row), multisource with random Fourier mask [Jo et al. 2022] (4th row), steering (5th row), steering with filter (6th row), and ours with 1-frame and 6-frame time multiplexing (multisource with content-adaptive dynamic Fourier modulation, 7–8th rows). For each configuration, we present the central view (1st column) and insets at two focal slices (rear and front) and two different viewpoints (left and right) in the next four columns. Quantitative evaluations are included as PSNR (dB)/SSIM on the left. Note that all methods use the same, large-baseline target light field for supervision, which degrades the quality of the single-source configuration (I) because it simply does not support such a large eyebox.




Fig. 7. **Focals stack and parallax comparison using experimentally captured results.** We compare the conventional setup and various multi-source holographic display configurations, including the single source setup with filter (1st column), the multisource setup without a filter (2nd column), with a random Fourier mask [Jo et al. 2022] (3rd column), and our configuration with 1-frame (4th column) and 6-frame time multiplexing (5th column). We show the full image from the central viewpoint in the top row, and the insets in the following rows are captured from different viewpoints. We see that the single-source configuration suffers from extreme brightness falloff at peripheral views. Our method achieves the best image quality among all multisource settings.

REFERENCES

- Jungkwon An, Kanghee Won, Young Kim, Jong-Young Hong, Hojung Kim, Yongkyu Kim, Hoon Song, Chilsung Choi, Yunhee Kim, Juwon Seo, Alexander Morozov, Hyunsik Park, Sunghoon Hong, Sungwoo Hwang, Kichul Kim, and Hong-Seok Lee. 2020. Slim-panel holographic video display. *Nature Communications* 11, 1 (10 Nov 2020), 5568. <https://doi.org/10.1038/s41467-020-19298-4>
- Stephen A Benton and V Michael Bove Jr. 2008. *Holographic imaging*. John Wiley & Sons.
- Colton M. Bigler, Pierre-Alexandre Blanche, and Kalluri Sarma. 2018. Holographic waveguide heads-up display for longitudinal image magnification and pupil expansion. *Appl. Opt.* 57, 9 (Mar 2018), 2007–2013. <https://doi.org/10.1364/AO.57.002007>
- Colton M. Bigler, Micah S. Mann, and Pierre-Alexandre Blanche. 2019. Holographic waveguide HUD with in-line pupil expansion and 2D FOV expansion. *Appl. Opt.* 58, 34 (Dec 2019), G326–G331. <https://doi.org/10.1364/AO.58.00G326>
- Edward Buckley, Adrian Cable, Nic Lawrence, and Tim Wilkinson. 2006. Viewing angle enhancement for two- and three-dimensional holographic displays with random superresolution phase masks. *Appl. Opt.* 45, 28 (Oct 2006), 7334–7341. <https://doi.org/10.1364/AO.45.007334>
- Minseok Chae, Kiseung Bang, Dongheon Yoo, and Yoonchan Jeong. 2023. Étendue Expansion in Holographic Near Eye Displays through Sparse Eye-box Generation Using Lens Array Eyepiece. *ACM Trans. Graph.* 42, 4 (July 2023), 1–13.
- Praneeth Chakravarthula, Seung-Hwan Baek, Florian Schiffers, Ethan Tseng, Grace Kuo, Andrew Maimone, Nathan Matsuda, Oliver Cossairt, Douglas Lanman, and Felix Heide. 2022. Pupil-Aware Holography. *ACM Trans. Graph.* 41, 6, Article 212 (nov 2022), 15 pages. <https://doi.org/10.1145/3550454.3555508>
- Praneeth Chakravarthula, Ethan Tseng, Tarun Srivastava, Henry Fuchs, and Felix Heide. 2020. Learned hardware-in-the-loop phase retrieval for holographic near-eye displays. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–18.
- C Chang, K Bang, G Wetzstein, B Lee, and L Gao. 2020. Toward the next-generation VR/AR optics: a review of holographic near-eye displays from a human-centric perspective. *Optica* (2020).
- Chenliang Chang, Wei Cui, Jongchan Park, and Liang Gao. 2019. Computational holographic Maxwellian near-eye display with an expanded eyebox. *Scientific Reports* 9, 1 (10 Dec 2019), 18749. <https://doi.org/10.1038/s41598-019-55346-w>
- Brian Chao, Manu Gopakumar, Suyeon Choi, and Gordon Wetzstein. 2023. High-brightness holographic projection. *Opt. Lett.* 48, 15 (Aug. 2023), 4041–4044.
- Suyeon Choi, Manu Gopakumar, Yifan Peng, Jonghyun Kim, Matthew O’Toole, and Gordon Wetzstein. 2022. Time-Multiplexed Neural Holography: A Flexible Framework for Holographic Near-Eye Displays with Fast Heavily-Quantized Spatial Light Modulators. In *ACM SIGGRAPH 2022 Conference Proceedings* (Vancouver, BC, Canada) (SIGGRAPH ’22). Association for Computing Machinery, New York, NY, USA, Article 32, 9 pages. <https://doi.org/10.1145/3528233.3530734>
- Craig T. Draper, Colton M. Bigler, Micah S. Mann, Kalluri Sarma, and Pierre-Alexandre Blanche. 2019. Holographic waveguide head-up display with 2-D pupil expansion and longitudinal image magnification. *Appl. Opt.* 58, 5 (Feb 2019), A251–A257. <https://doi.org/10.1364/AO.58.00A251>
- Craig T. Draper and Pierre-Alexandre Blanche. 2021. Examining aberrations due to depth of field in holographic pupil replication waveguide systems. *Appl. Opt.* 60, 6 (Feb 2021), 1653–1659. <https://doi.org/10.1364/AO.417756>
- Craig T. Draper and Pierre-Alexandre Blanche. 2022. Holographic curved waveguide combiner for HUD/AR with 1-D pupil expansion. *Opt. Express* 30, 2 (Jan 2022), 2503–2516. <https://doi.org/10.1364/OE.445091>
- Joseph W Goodman. 2005. Introduction to Fourier optics. *Introduction to Fourier optics, 3rd ed.*, by JW Goodman. Englewood, CO: Roberts & Co. Publishers, 2005 1 (2005).
- Manu Gopakumar, Jonghyun Kim, Suyeon Choi, Yifan Peng, and Gordon Wetzstein. 2021. Unfiltered holography: optimizing high diffraction orders without optical filtering for compact holographic displays. *Opt. Lett.* 46, 23 (Dec 2021), 5822–5825. <https://doi.org/10.1364/OL.442851>
- Manu Gopakumar, Gun-Yeal Lee, Suyeon Choi, Brian Chao, Yifan Peng, Jonghyun Kim, and Gordon Wetzstein. 2024. Full-colour 3D holographic augmented-reality displays with metasurface waveguides. *Nature* (08 May 2024). <https://doi.org/10.1038/s41586-024-07386-0>
- Changwon Jang, Kiseung Bang, Minseok Chae, Byoungcho Lee, and Douglas Lanman. 2024. Waveguide holography for 3D augmented reality glasses. *Nature Communications* 15, 1 (02 Jan 2024), 66. <https://doi.org/10.1038/s41467-023-44032-1>
- Changwon Jang, Kiseung Bang, Gang Li, and Byoungcho Lee. 2018. Holographic Near-Eye Display with Expanded Eye-Box. *ACM Trans. Graph.* 37, 6, Article 195 (dec 2018), 14 pages. <https://doi.org/10.1145/3272127.3275069>
- Bahram Javidi, Artur Carnicer, Arun Anand, George Barbastathis, Wen Chen, Pietro Ferraro, J W Goodman, Ryoichi Horisaki, Kedar Khare, Malgorzata Kujawinska, Rainer A Leitgeb, Pierre Marquet, Takanori Nomura, Aydogan Ozcan, Yongkeun Park, Giancarlo Pedrini, Pascal Picart, Joseph Rosen, Genaro Saavedra, Natan T Shaked, Adrian Stern, Enrique Tajahuerce, Lei Tian, Gordon Wetzstein, and Masahiro Yamaguchi. 2021. Roadmap on digital holography [Invited]. *Opt. Express*, OE 29, 22 (Oct. 2021), 35078–35118.
- Youngjin Jo, Dongheon Yoo, Dukho Lee, Minkwan Kim, and Byoungcho Lee. 2022. Multi-illumination 3D holographic display using a binary mask. *Opt. Lett.* 47, 10 (May 2022), 2482–2485. <https://doi.org/10.1364/OL.455348>
- Koray Kavaklı, Liang Shi, Hakan Urey, Wojciech Matusik, and Kaan Akşit. 2023. Multi-color Holograms Improve Brightness in Holographic Displays. In *SIGGRAPH Asia 2023 Conference Papers*. Association for Computing Machinery, 1–11.
- Jonghyun Kim, Manu Gopakumar, Suyeon Choi, Yifan Peng, Ward Lopes, and Gordon Wetzstein. 2022. Holographic Glasses for Virtual Reality. In *ACM SIGGRAPH 2022 Conference Proceedings* (Vancouver, BC, Canada) (SIGGRAPH ’22, Article 33). Association for Computing Machinery, New York, NY, USA, 1–9.
- Bernard C Kress and Ishan Chatterjee. 2020. Waveguide combiners for mixed reality headsets: a nanophotonics design perspective. *Nanophotonics* 10, 1 (2020), 41–74.
- Grace Kuo, Florian Schiffers, Douglas Lanman, Oliver Cossairt, and Nathan Matsuda. 2023. Multisource Holography. *ACM Trans. Graph.* 42, 6, Article 203 (dec 2023), 14 pages. <https://doi.org/10.1145/3618380>
- Grace Kuo, Laura Waller, Ren Ng, and Andrew Maimone. 2020. High Resolution Étendue Expansion for Holographic Displays. *ACM Trans. Graph.* 39, 4, Article 66 (aug 2020), 14 pages. <https://doi.org/10.1145/3386569.3392414>
- Grigory Lazarev, Stefanie Bonifer, Philip Engel, Daniel Höhne, and Gunther Notni. 2017. High-resolution LCOS microdisplay with sub-kHz frame rate for high performance, high precision 3D sensor. In *Digital Optical Technologies 2017*, Bernard C. Kress and Peter Schellkens (Eds.), Vol. 10335. International Society for Optics and Photonics, SPIE, 103351B. <https://doi.org/10.1117/12.2272367>
- Byoungcho Lee, Dongheon Yoo, Jinsoo Jeong, Seungjae Lee, Dukho Lee, and Byoungcho Lee. 2020. Wide-angle speckleless DMD holographic display using structured illumination with temporal multiplexing. *Opt. Lett.* 45, 8 (Apr 2020), 2148–2151. <https://doi.org/10.1364/OL.390552>
- Dukho Lee, Kiseung Bang, Seung-Woo Nam, Byoungcho Lee, Dongyeon Kim, and Byoungcho Lee. 2022. Expanding energy envelope in holographic display via mutually coherent multi-directional illumination. *Scientific Reports* 12, 1 (22 Apr 2022), 6649. <https://doi.org/10.1038/s41598-022-10355-0>
- Wen-Kai Lin, Osamu Matoba, Bor-Shyh Lin, and Wei-Chia Su. 2018. Astigmatism and deformation correction for a holographic head-mounted display with a wedge-shaped holographic waveguide. *Appl. Opt.* 57, 25 (Sep 2018), 7094–7101. <https://doi.org/10.1364/AO.57.007094>
- Wen-Kai Lin, Osamu Matoba, Bor-Shyh Lin, and Wei-Chia Su. 2020. Astigmatism correction and quality optimization of computer-generated holograms for holographic waveguide displays. *Opt. Express* 28, 4 (Feb 2020), 5519–5527. <https://doi.org/10.1364/OE.381193>
- Andrew Maimone, Andreas Georgiou, and Joel S. Kollin. 2017. Holographic Near-Eye Displays for Virtual and Augmented Reality. *ACM Trans. Graph.* 36, 4, Article 85 (jul 2017), 16 pages. <https://doi.org/10.1145/3072959.3073624>
- Eric Markley, Nathan Matsuda, Florian Schiffers, Oliver Cossairt, and Grace Kuo. 2023. Simultaneous Color Computer Generated Holography. In *SIGGRAPH Asia 2023 Conference Papers* (<conf-loc>, <city>Sydney</city>, <state>NSW</state>, <country>Australia</country>, </conf-loc>) (SA ’23, Article 22). Association for Computing Machinery, New York, NY, USA, 1–11.
- Sagi Monin, Aswin C. Sankaranarayanan, and Anat Levin. 2022a. Analyzing phase masks for wide étendue holographic displays. In *2022 IEEE International Conference on Computational Photography (ICCP)*. 1–12. <https://doi.org/10.1109/ICCP54855.2022.9887757>
- Sagi Monin, Aswin C. Sankaranarayanan, and Anat Levin. 2022b. Exponentially-wide étendue displays using a tilting cascade. In *2022 IEEE International Conference on Computational Photography (ICCP)*. 1–12. <https://doi.org/10.1109/ICCP54855.2022.9887737>
- N. Padmanaban, Y. Peng, and G. Wetzstein. 2019. Holographic Near-Eye Displays Based on Overlap-Add Stereograms. *ACM Trans. Graph. (SIGGRAPH Asia)* 6 (2019). Issue 38.
- Jongchan Park, KyeoReh Lee, and YongKeun Park. 2019. Ultrathin wide-angle large-area digital 3D holographic display using a non-periodic photon sieve. *Nature Communications* 10, 1 (21 Mar 2019), 1304. <https://doi.org/10.1038/s41467-019-09126-9>
- Jae-Hyeung Park and Seong-Bok Kim. 2018. Optical see-through holographic near-eye display with eyebox steering and depth of field control. *Opt. Express* 26, 21 (Oct. 2018), 27076–27088.
- Y. Peng, S. Choi, N. Padmanaban, and G. Wetzstein. 2020. Neural Holography with Camera-in-the-loop Training. *ACM Trans. Graph. (SIGGRAPH Asia)* (2020).
- Dapu Pi, Juan Liu, and Yongtian Wang. 2022. Review of computer-generated hologram algorithms for color dynamic holographic three-dimensional display. *Light Sci Appl* 11, 1 (July 2022), 231.
- Stephan Reichelt, Ralf Haussler, Norbert Leister, Gerald Futterer, Hagen Stolle, and Armin Schwertner. 2010. Holographic 3-D Displays - Electro-holography within the Grasp of Commercialization. In *Advances in Lasers and Electro Optics*, Nelson Costa and Adolfo Cartaxo (Eds.), IntechOpen, Rijeka, Chapter 29. <https://doi.org/10.5772/8650>

- Florian Schiffers, Praneeth Chakravarthula, Nathan Matsuda, Grace Kuo, Ethan Tseng, Douglas Lanman, Felix Heide, and Oliver Cossairt. 2023. Stochastic Light Field Holography. *IEEE International Conference on Computational Photography (ICCP)* (2023).
- Liang Shi, Beichen Li, Changil Kim, Petr Kellnhofer, and Wojciech Matusik. 2021. Towards real-time photorealistic 3D holography with deep neural networks. *Nature* 591, 7849 (01 Mar 2021), 234–239. <https://doi.org/10.1038/s41586-020-03152-0>
- Liang Shi, Beichen Li, and Wojciech Matusik. 2022. End-to-end learning of 3D phase-only holograms for holographic display. *Light Sci Appl* 11, 1 (Aug. 2022), 247.
- Liang Shi, DongHun Ryu, and Wojciech Matusik. 2024. Ergonomic-Centric Holography: Optimizing Realism, Immersion, and Comfort for Holographic Display. *Laser & Photonics Reviews* 18, 4 (2024), 2300651. <https://doi.org/10.1002/lpor.202300651>
- A Starikov. 1982. Effective number of degrees of freedom of partially coherent sources. *JOSA* 72, 11 (1982), 1538–1544.
- Jie Sun, Erman Timurdogan, Ami Yaacobi, Ehsan Shah Hosseini, and Michael R Watts. 2013. Large-scale nanophotonic phased array. *Nature* 493, 7431 (2013), 195–199.
- Ethan Tseng, Grace Kuo, Seung-Hwan Baek, Nathan Matsuda, Andrew Maimone, Florian Schiffers, Praneeth Chakravarthula, Qiang Fu, Wolfgang Heidrich, Douglas Lanman, and Felix Heide. 2024. Neural étendue expander for ultra-wide-angle high-fidelity holographic display. *Nature Communications* 15, 1 (22 Apr 2024), 2907. <https://doi.org/10.1038/s41467-024-46915-3>
- Zi Wang, Guoqiang Lv, Yujian Pang, Qibin Feng, Anting Wang, and Hai Ming. 2023. Lens array-based holographic 3D display with an expanded field of view and eyebox. *Opt. Lett.* 48, 21 (Nov. 2023), 5559–5562.
- Xinxing Xia, Yunqing Guan, Andrei State, Praneeth Chakravarthula, Tat-Jen Cham, and Henry Fuchs. 2020. Towards Eyeglass-style Holographic Near-eye Displays with Statically. In *2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 312–319.
- Xinxing Xia, Weisen Wang, Frank Guan, Furong Yang, Xinghua Shui, Huadong Zheng, Yingjie Yu, and Yifan Peng. 2023. Exploring angular-steering illumination-based eyebox expansion for holographic displays. *Opt. Express* 31, 19 (Sep 2023), 31563–31573. <https://opg.optica.org/oe/abstract.cfm?URI=oe-31-19-31563>
- Daeho Yang, Wontaek Seo, Hyeonseung Yu, Sun Il Kim, Bongsu Shin, Chang-Kun Lee, Seokil Moon, Jungkwuen An, Jong-Young Hong, Geeyoung Sung, and Hong-Seok Lee. 2022. Diffraction-engineered holography: Beyond the depth representation limit of holographic displays. *Nature Communications* 13, 1 (12 Oct 2022), 6012. <https://doi.org/10.1038/s41467-022-33728-5>
- Jiwoon Yeom, Yeseul Son, and Kwangsoon Choi. 2021. Crosstalk Reduction in Voxels for a See-Through Holographic Waveguide by Using Integral Imaging with Compensated Elemental Images. *Photonics* 8, 6 (2021). <https://doi.org/10.3390/photonics8060217>
- Hyeonseung Yu, KyeoReh Lee, Jongchan Park, and YongKeun Park. 2017. Ultrahigh-definition dynamic 3D holographic display by active control of volume speckle fields. *Nature Photonics* 11, 3 (01 Mar 2017), 186–192. <https://doi.org/10.1038/nphoton.2016.272>
- Zhengyun Zhang and Marc Levoy. 2009. Wigner distributions and how they relate to the light field. In *2009 IEEE International Conference on Computational Photography (ICCP)*. IEEE, 1–10.