Holographic display étendue expansion with a binary π-metasurface

CONGLI WANG[1], ZHIHAO ZHOU[2], ETHAN TSENG[1], VICTOR CHU[1], WEI TING CHEN[3], JOHANNES E. FRÖCH[2,4], ARKA MAJUMDAR[2,4], AND FELIX HEIDE[1,*]

[1] *Department of Computer Science, Princeton University, Princeton, New Jersey 08544, USA*
[2] *Department of Electrical and Computer Engineering, University of Washington, Seattle, Washington 98195, USA*
[3] *SNOChip Inc., Princeton, New Jersey 08540, USA*
[4] *Department of Physics, University of Washington, Seattle, Washington 98195, USA*
[*] *fheide@princeton.edu*

Published in Optics Letters, Vol. 51, No. 19 (2026)

Holographic displays trade field of view against eyebox size because a spatial light modulator provides fixed étendue. We pair a commercial 4-bit piston-mode phase light modulator (PLM) with a pixel-registered binary metasurface designed for an achromatic π phase step. Grouping the PLM into $s \times s$ superpixels while retaining native mask sampling increases the modeled angular-support area by s^2 with one passive element. Scalar simulations and three-color experimental captures at 450, 520 and 660 nm show matched operation from one fixed mirror bias. A fixed-depth diffractive element leaks 7.0% and 11.9% of the power into an un-expanded ghost at 450 and 660 nm, which the metasurface's full-wave-simulated zero order holds to 0.3% and 2.7%.

https://doi.org/10.1364/OL.613712

The field of view and the eyebox of a holographic near-eye display are coupled through the space–bandwidth product of the spatial light modulator (SLM) that forms the phase pattern [1, 2]. A modulator with $N_x \times N_y$ pixels and pitches p_x, p_y has angular spans $\Delta\theta_x \sim \lambda/p_x$ and $\Delta\theta_y \sim \lambda/p_y$, but only $N_x N_y$ controllable spatial modes, and that mode count fixes the display's étendue, the product of field of view and eyebox area. Increasing angular support at a fixed mode count therefore shrinks the eyebox over which the image can be delivered, producing the familiar field-of-view versus eyebox trade-off.

A productive route around this limit is to follow the modulator with a static, high-resolution phase element (an "étendue expander") that scatters light into a wider angular support while a co-designed modulator pattern re-focuses that expanded bandwidth into the target image [3–5]. Recent neural étendue expanders learn a grayscale height map jointly with a corpus of holograms and reach ultra-wide fields [5]. These systems occupy a high-dimensional design point, with a continuously varying mask learned end-to-end through a fixed differentiable image-formation model. Metasurfaces have likewise entered holographic displays, as co-designed waveguide couplers [6], polarization-multiplexing SLM companions [7], and synthetic-aperture waveguide holography [8].

Here we take a deliberately minimal approach. Our expander is a dispersion-engineered binary metasurface consisting of SiN nanopillars [9–11]: a two-level mask that is designed to impose the same π phase step at all three primary wavelengths, in contrast to a diffractive optical element (DOE) whose phase scales as $2\pi(n-1)h/\lambda$. It is fabricated once, carries no free parameters at run time, and is driven by a commercial 4-bit piston-mode phase light modulator (PLM). The modulator phase is found by plain first-order optimization of a differentiable far-field model rather than a learned camera- or hardware-in-the-loop model [12], and the expansion is achieved in one fixed passive element that needs no per-wavelength reconfiguration. We describe the system [Fig. 1], evaluate its three-color design in simulation against a fixed-depth diffractive baseline [Fig. 2], and compare the scalar model with experimental captures at 450, 520 and 660 nm [Fig. 3].

Figure 1(a) shows the experimental layout. A fiber-coupled RGB laser at 450, 520 and 660 nm, collimated at the PLM, passes through a beamsplitter to the PLM; the modulated return is folded by the beamsplitter into a $4f$ relay of two achromatic doublets, with a Fourier-plane DC stop and iris that block the PLM zero order, that images the PLM onto the metasurface at unit magnification for mirror-to-pixel registration. A Fourier-transform lens then forms the far-field hologram on a monochrome camera. Component part numbers, the camera-plane sampling calibration (a measured $0.66\times$ scale applied whenever simulation and capture are compared) and the measured modulator look-up tables are given in Supplement 1.

The phase pattern is generated by a Texas Instruments PLM (device XDLPM670E1FYZ), a piston-mode micro-electro-mechanical device with 10.8 μm phase cells and 4-bit phase, operated at one fixed mirror-bias voltage for all three wavelengths with the chromatic calibration held in three measured phase look-up tables (Supplement 1).

The passive étendue expander is a silicon-nitride metasurface fabricated by SNOChip Inc., patterned on a 0.5 mm fused-silica substrate, with an active area of 13.8 mm $\times$ 8.6 mm and a supercell pitch of 10.8 μm that registers one-to-one with the PLM mirrors [Fig. 1]. Each supercell consists of either 191-nm- or 336-nm-diameter SiN nanopillars, 640 nm tall, on a 386 nm center-to-center lattice. The two pillar populations are dispersion-engineered so that their transmissions maintain a π phase difference across the three primaries, to within 13° on average over 400–700 nm in full-wave simulation [Fig. 1(c), Supplement 1], with a simulated transmission of 82, 79 and 93% at 450, 520 and

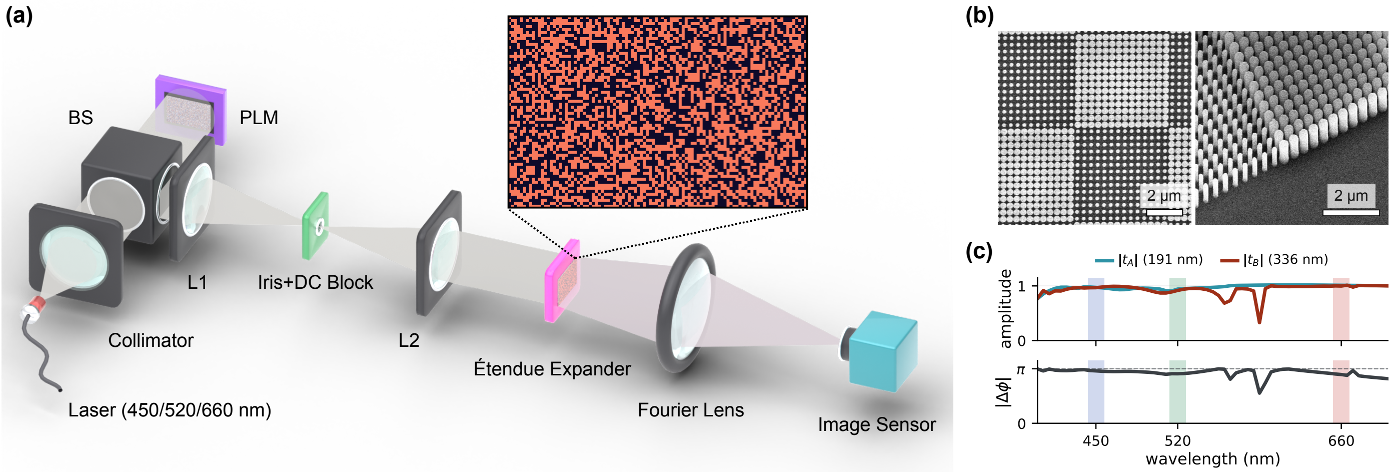

Fig. 1. System overview. (a) A fiber-coupled laser illuminates the 4-bit piston-mode PLM through a beamsplitter (BS); a 1:1 relay (lenses L_1, L_2, with a Fourier-plane iris and zero-order DC block) images the PLM onto the pixel-registered binary metasurface, and a Fourier lens forms the far field on the camera. The widening envelope indicates the expanded angular support. (b) Scanning electron micrographs (top and oblique views) of the fabricated nanopillars: 191 and 336 nm diameters on a 386 nm lattice. (c) Simulated amplitude and phase difference of the two meta-atom states; the phase difference stays within 13° of π on average across 400–700 nm, apart from the 560–595 nm resonance (Supplement 1).

660 nm. Figure 1(b) presents scanning electron micrographs of a representative region of the fabricated sample.

Holograms are synthesized by direct optimization of the modulator phase. Let $\phi_{\mathrm{P}}^{(k)}$ be the PLM phase of the kth of the K multiplexed solutions, on a coarse $(R/s) \times (C/s)$ superpixel grid, where the integer étendue expansion factor s groups $s \times s$ mirrors into one superpixel, and let $\phi_{\mathrm{E}}(\mathbf{x}) = \pi m(\mathbf{x})$, $m \in \{0, 1\}$, be the fixed metasurface phase on the native $R \times C$ grid. The operator U_s expands each PLM sample into a constant $s \times s$ mirror block; the metasurface mask is never resized. Assuming coherent, unit-amplitude illumination and a single Fraunhofer propagation, a modulator phase ϕ produces the far-field intensity

$$\mathcal{I}[\phi](\mathbf{u}) = \left|\mathcal{F}\left\{\exp\left(i\left[\,U_s\phi(\mathbf{x}) + \phi_{\mathrm{E}}(\mathbf{x})\,\right]\right)\right\}\right|^2, \tag{1}$$

where $\mathbf{x}$ and $\mathbf{u}$ are transverse coordinates in the PLM (and pixel-registered metasurface) plane and in the image plane, respectively, and $\mathcal{F}$ is the discrete Fourier transform. Each displayed pattern is then the solution of

$$\phi_{\mathrm{P}}^{(k)} \in \arg\min_{\phi} \left\| \widehat{\mathcal{I}}[\phi] - T \right\|_2^2, \qquad \widehat{\mathcal{I}}[\phi] = \frac{\mathcal{I}[\phi]}{\|\mathcal{I}[\phi]\|_1}, \tag{2}$$

with T the power-normalized target. We solve Eq. (2) by direct first-order descent with Adam [13] (learning rate 0.1, 2000 iterations), initialized from a different random seed for each $k = 1, \ldots, K$, which is what makes the K solutions distinct. Figure 2 multiplexes $K = 8$ such solutions per color, so the camera-integrated reconstruction is speckle-suppressed. The metasurface arm optimizes against the fabricated design's simulated complex transmissions $t_{A,B}(\lambda)$ in place of ϕ_{E} in Eq. (1), whereas the DOE simulation re-propagates the π-designed solutions through the chromatic DOE: a hologram co-designed with the real element against one deployed without recalibration. The unscattered fraction $|c_0(\lambda)|^2$ in Fig. 2 is analytic, set by the DOE's fixed-depth material dispersion with the resin indices of Ref. [5], whereas the metasurface curve is the full-wave simulated zero order of the fabricated meta-atom pair, which a scalar thin-element estimate built from the two states' simulated transmissions reproduces (Supplement 1). Figure 2 uses the continuous optimized phase, modeling an ideal phase-only SLM so that the DOE-versus-metasurface difference is not confounded by phase modulator quantization.

The $s \times s$ grouping gives the PLM an effective pitch sp, so its unassisted angular support scales as $\lambda/(sp)$ along each axis. The registered mask retains pitch p and therefore supplies spatial frequencies out to the native scale λ/p (2.4°, 2.8° and 3.5° full width at 450, 520 and 660 nm). Relative to the grouped PLM alone, the modeled support expands by s per axis, or s^2 in area, while the number of controllable PLM phase values decreases by the same areal factor. A future system would pair a natively driven modulator with an s-times finer mask; the pixel-registered grouping emulates that configuration with one element, at the same controllable fraction $1/s^2$. We report this modeled areal expansion generically as s^2.

Coherent illumination and a broadly scattering expander leave a single reconstruction strongly speckled. The cause is the mode deficit: a phase-only modulator with a fixed mode count cannot fill the enlarged support with one pattern, so some output modes stay uncontrolled for any expander, and the simulated DOE and metasurface share the residual. The mask's randomness sets where that energy lands, as a diffuse halo rather than discrete orders. This is the space–bandwidth trade-off quantified for PLM-based holography by Chu et al. [14], and it is reducible: co-designing the expander with the holograms recovers much of the loss at fixed modulator resolution [5], which the fixed binary mask here deliberately forgoes. We suppress it instead by temporal multiplexing, loading K solutions of Eq. (2) from different random seeds across the subframes the drive emits within one video frame, eight per color when all three are driven together and all 24 when the colors are captured sequentially; the camera integrates them and the measured three-color contrast σ/μ at $s = 2$ falls from 0.65 to 0.44 and 0.41 at $K = 1$, 8 and 24 [15, 16] (Supplement 1). Earlier demonstrations remove the residual before display, by low-pass filtering the captures to the controllable bandwidth [3] or by averaging 20 patterns [5]; neither quotes a speckle contrast, so we report ours, noting

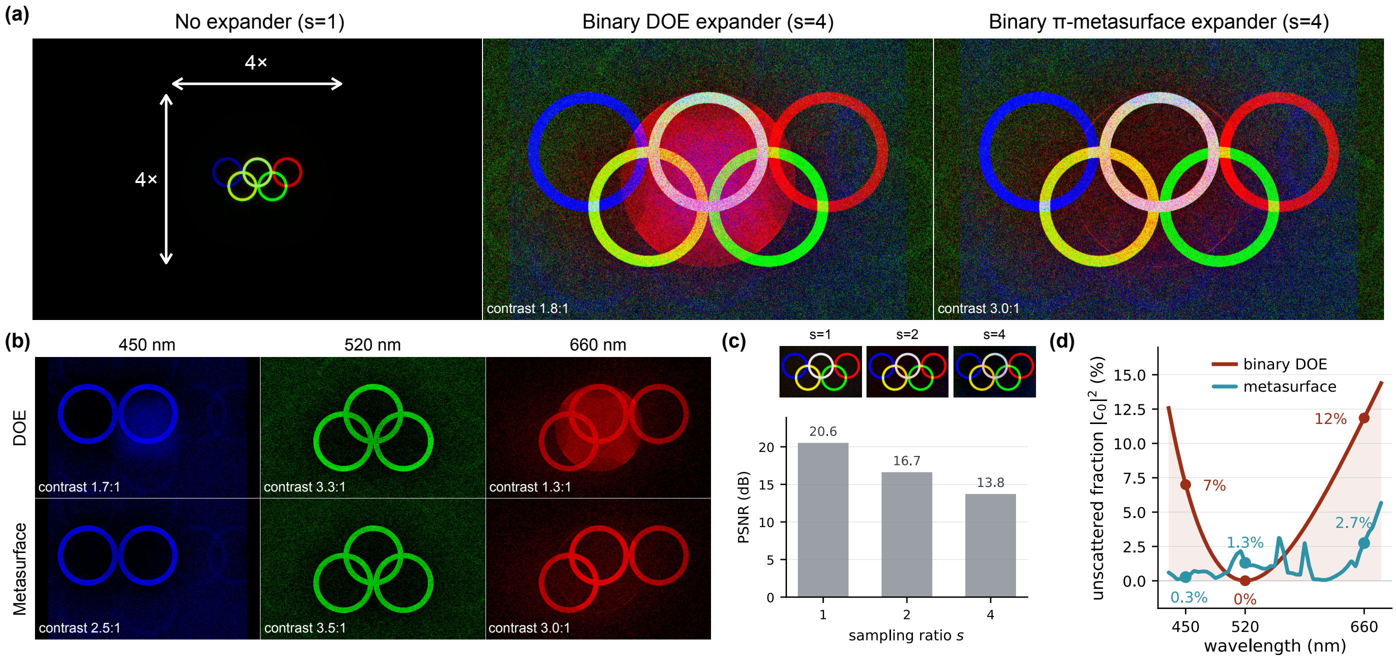

Fig. 2. Simulated étendue expansion. (a) Full-color (450/520/660 nm) reconstructions of a color-ring target with no expander ($s = 1$), a binary DOE designed for π at 520 nm ($s = 4$), and the fabricated metasurface ($s = 4$, holograms optimized against its simulated response), on an ideal continuous phase-only SLM with $K = 8$ multiplexed solutions per color. The target grows with the expanded support, 4× per axis (16× in area) over the un-expanded reference; both passive elements address the enlarged field, and each leaks its unscattered fraction into the un-expanded central zone, dominant for the DOE and faintly visible at 660 nm for the metasurface, in proportion to the zero orders of (d). Contrast labels are signal-to-background ratios inside that zone (Supplement 1). (b) Per-wavelength channels of the two expanders' $s = 4$ reconstructions with their un-expanded ghosts. (c) Increasing the étendue expansion factor s at a fixed target extent enlarges the modeled areal support by s^2 at a cost in peak signal-to-noise ratio (PSNR). (d) Unscattered fraction $|c_0(\lambda)|^2$: 7.0%/0.0%/11.9% at 450/520/660 nm for the DOE against the metasurface's full-wave simulated zero order of 0.3%/1.3%/2.7% and a measured 9.1%/7.5%/3.5%, referenced through the simulated transmission (Supplement 1); the DOE is exact only at its design wavelength.

that the subframe depth is fixed by the drive budget rather than chosen.

Figure 2(a) compares simulated full-color far fields at $s = 4$, a 16× areal expansion; the $s = 2$ configuration of the experimental captures is retained in Supplement 1. Without the expander the reconstruction is confined to the un-expanded s=1 support. A conventional binary DOE, the same mask layout as a fixed-depth surface-relief element with a π step at 520 nm, expands the support but fails in color: because its phase step is wavelength-dependent [1], the mask leaves an unscattered fraction $|c_0(\lambda)|^2$ that passes as an un-expanded ghost copy of the hologram [Fig. 2(b)]. Increasing the étendue expansion factor raises the modeled areal support by s^2: at $s = 1$, 2 and 4 the PSNR falls from 20.6 to 16.7 and 13.8 dB [Fig. 2(c)]. That unscattered fraction is 7.0% at 450 nm and 11.9% at 660 nm for the DOE, against 0.27%, 1.31% and 2.75% for the meta-atom pair in full-wave simulation [Fig. 2(d)]; the DOE is exact at its 520 nm design point. The measured fraction is 9.1%, 7.5% and 3.5%, above the simulation; what it supports is the achromatic bound rather than absolute suppression, the worst channel staying below that element's 11.9% and far below the 51.2% of a 660 nm design. More DOE alternatives are quantified in Supplement 1.

Figure 3 shows three-color experimental reconstructions beside the scalar model's renderings of the same displayed patterns, all captured at $s = 2$, so the expansion itself is exercised experimentally rather than in simulation alone. The measured far fields are qualitatively consistent with the modeled structure, supporting use of the scalar-Fraunhofer model, and the 450, 520 and 660 nm channels reconstruct each target at matched positions from the single fixed mirror bias, demonstrating three-color operation through the one passive element. The captures also probe the achromatic π step directly: every hologram is optimized against the nominal π mask at each wavelength, so a chromatic phase step would surface in the captures alone, as an un-expanded ghost raising the background inside the content region. Instead the measured background contrast stays at a common fraction of the π-model prediction at all three primaries, 0.25 in the median, with only a mild red deficit consistent with the documented stray light, whereas a 520 nm fixed-depth element in the same simulated metric shows a wavelength-selective penalty reaching 1.7× at 660 nm (Supplement 1). The rows differ where the physical display path enters: a broad background haze and detail softened by residual speckle leave the measured composites 2 to 5 dB below simulation in color PSNR on most subjects, the gap vanishing on the smoothest natural subjects (per-subject numbers in Supplement 1). The background is set by the mode deficit and by registration and stray light, together with the element's unscattered order, which the defocus redistributes (Supplement 1). Each capture averages the $K = 24$ separately optimized patterns the sequential single-laser drive allows; the expansion itself requires no multiplexing, since every subframe carries the full support. See Supplement 1 for the measured performance at other subframe counts K.

In summary, a commercial 4-bit piston-mode PLM and a

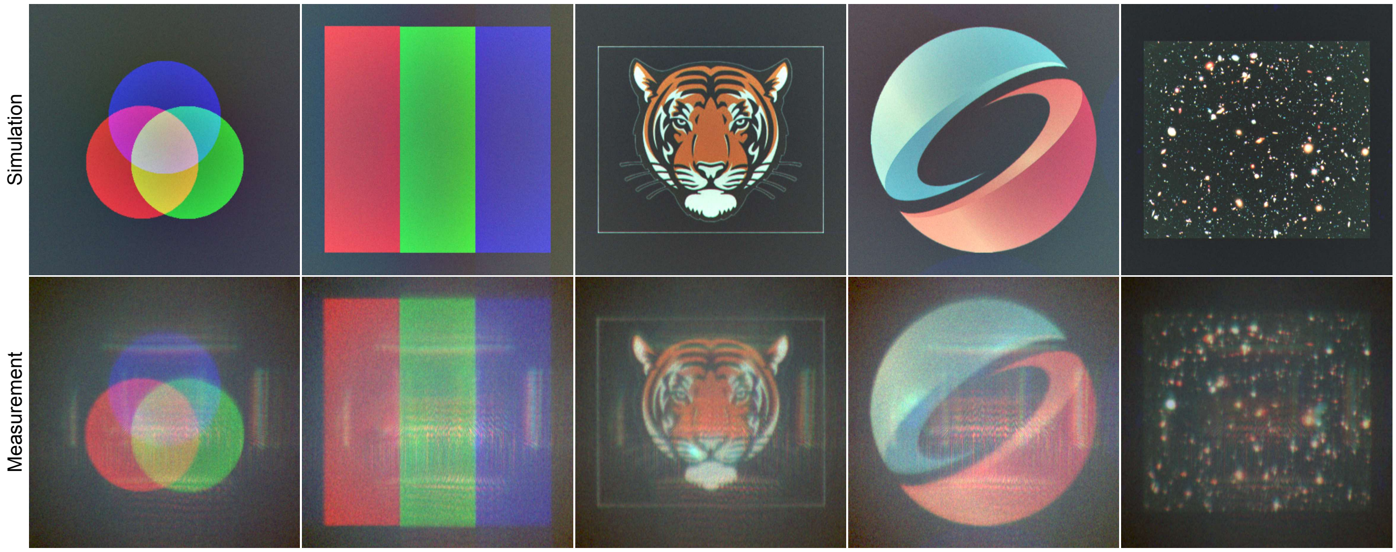

Fig. 3. Three-color measurement against the scalar model at the expanded support. Scalar-model renderings of the displayed patterns (top) and measured full-color reconstructions (bottom) of five targets, captured at 450, 520 and 660 nm at $s = 2$, so every capture operates on the expanded support. Each capture averages $K = 24$ separately optimized patterns per color within one exposure to suppress speckle; the simulation row propagates the same patterns through the scalar model, so the rows differ only by the physical display path. The holograms pre-compensate the 16.8 mm sensor defocus, spreading the residual zero order into a low-level background. Each channel is scaled to its optimization target by one mean-match gain, a digital channel balance applied identically to both rows, and each subject's rows share one display scale and magnification (Supplement 1).

pixel-registered binary π-metasurface give modeled areal support expansion of s^2 from a single passive element, with the modulator phase found by direct differentiable optimization and no camera-in-the-loop learned model. Three-color experimental captures at 450, 520 and 660 nm reconstruct the targets on the expanded $s = 2$ support from one fixed mirror bias and agree qualitatively with the scalar model, where a fixed-depth diffractive equivalent would leak an un-expanded ghost at the outer primaries. The appeal is that étendue expansion moves off the modulator and into a fabricated element, so reaching larger s or a wider band becomes a mask-design question rather than a per-wavelength calibration problem. The present system is a proof of principle of this passive achromatic mechanism rather than a display prototype. The quality at large s reflects the modulator's fixed mode budget. A neural proxy [12] trained on experimental measurements to absorb the residual un-modeled physics is then a natural route toward achromatic étendue-expanded displays limited neither by support nor by coherent noise, and a joint-design metasurface expander could be co-optimized with the holograms to recover some of the lost reconstruction quality at fixed modulator resolution [5].

Funding. Defense Advanced Research Projects Agency (ENVision program); U.S. National Science Foundation (NNCI-1542101, NNCI-2025489); Gordon and Betty Moore Foundation; NSF CAREER Award (2047359); Packard Foundation Fellowship; Sloan Research Fellowship; Sony Young Faculty Award; Project X Innovation Award; Amazon Science Research Award; Bosch Research Award.

Acknowledgment. The authors thank Yiling Li for help with the experiments. Part of this work was conducted at a National Nanotechnology Coordinated Infrastructure (NNCI) site at the University of Washington.

Disclosures. E.T. is Founder and CTO of Cephia AI. W.T.C. is CTO of SNOChip Inc. A.M. is a co-founder of Tunoptix. F.H. is a co-founder of Algolux (now Torc Robotics), Head of AI at Torc Robotics, and a co-founder of Cephia AI. The remaining authors declare no conflicts of interest.

Data availability. Data underlying the results presented in this paper are not publicly available at this time but may be obtained from the authors upon reasonable request.

Supplemental document. See Supplement 1 for supporting content.

REFERENCES

1. J. W. Goodman, *Introduction to Fourier Optics* (W. H. Freeman, 2017), 4th ed.
2. C. Chang, K. Bang, G. Wetzstein, *et al.*, Optica **7**, 1563 (2020).
3. G. Kuo, L. Waller, R. Ng, and A. Maimone, ACM Trans. Graph. **39**, 66 (2020).
4. J. Park, K. Lee, and Y. Park, Nat. Commun. **10**, 1304 (2019).
5. E. Tseng, G. Kuo, S.-H. Baek, *et al.*, Nat. Commun. **15**, 2907 (2024).
6. M. Gopakumar, G.-Y. Lee, S. Choi, *et al.*, Nature **629**, 791 (2024).
7. S.-W. Nam, Y. Kim, D. Kim, and Y. Jeong, ACM Trans. Graph. **42**, 202 (2023).
8. S. Choi, C. Jang, D. Lanman, and G. Wetzstein, Nat. Photonics **19**, 854 (2025).
9. W. T. Chen, A. Y. Zhu, and F. Capasso, Nat. Rev. Mater. **5**, 604 (2020).
10. W. T. Chen, A. Y. Zhu, V. Sanjeev, *et al.*, Nat. Nanotechnol. **13**, 220 (2018).
11. S. Wang, P. C. Wu, V.-C. Su, *et al.*, Nat. Nanotechnol. **13**, 227 (2018).
12. Y. Peng, S. Choi, N. Padmanaban, and G. Wetzstein, ACM Trans. Graph. **39**, 185 (2020).
13. D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *International Conference on Learning Representations (ICLR)*, (2015). arXiv:1412.6980.
14. V. Chu, O. Pueyo-Ciutad, E. Tseng, *et al.*, ACM Trans. Graph. **44**, 219 (2025).
15. L. Golan and S. Shoham, Opt. Express **17**, 1330 (2009).
16. V. Bianco, P. Memmolo, M. Leo, *et al.*, Light Sci. Appl. **7**, 48 (2018).