

Supplementary Document

Large-Area Fabrication-Aware Computational Diffractive Optics

KAIXUAN WEI*, King Abdullah University of Science and Technology, Saudi Arabia

HECTOR A. JIMENEZ-ROMERO*, King Abdullah University of Science and Technology, Saudi Arabia

HADI AMATA*, King Abdullah University of Science and Technology, Saudi Arabia

JIPENG SUN, Princeton University, United States of America

QIANG FU, King Abdullah University of Science and Technology, Saudi Arabia

FELIX HEIDE, Princeton University, United States of America

WOLFGANG HEIDRICH, King Abdullah University of Science and Technology, Saudi Arabia

ACM Reference Format:

Kaixuan Wei, Hector A. Jimenez-Romero, Hadi Amata, Jipeng Sun, Qiang Fu, Felix Heide, and Wolfgang Heidrich. 2025. Supplementary Document Large-Area Fabrication-Aware Computational Diffractive Optics. *ACM Trans. Graph.* 44, 6 (December 2025), 15 pages. <https://doi.org/10.1145/3763358>

CONTENTS

Contents	1
A Additional Fabrication Details	2
A.1 Direct-write Grayscale Lithography	2
A.2 Nanoimprint Lithography	2
B Additional Implementation Details	3
B.1 Neural Lithography Model Architecture	3
B.2 Hologram Reconstruction Loss	3
B.3 Broadband Imaging PSF Loss	4
B.4 Additional RAW Image Processing Details	5
B.5 Additional ISP Post-Processing Details	6
C Additional Results and Analysis	7
C.1 More Results on Lithography Model Calibration and Evaluation	7
C.2 Additional Experimental Results on Broadband Imaging	8
D Additional Experimental Prototype Details	14
D.1 Broadband Imaging Setup	14
References	15

*indicates joint first authorship.

© 2025 Association for Computing Machinery.

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Graphics*, <https://doi.org/10.1145/3763358>.

A ADDITIONAL FABRICATION DETAILS

A.1 Direct-write Grayscale Lithography

In this work, we employ direct-write grayscale lithography using the Heidelberg Instruments DWL 66+ mask writer to fabricate intricate microstructured patterns with high precision. This advanced lithographic technique leverages the system's ability to finely modulate the laser exposure across 1023 discrete grayscale levels. Such fine resolution in exposure control allows for the accurate modulation of the energy dose imparted to the photoresist, which is crucial for defining complex topographies.

We utilize AZ® 4562 photoresist, one of the most reliable and versatile positive-tone resists available, widely recognized for its excellent response to variations in exposure dose. Its high sensitivity and broad dynamic range make it exceptionally well-suited for grayscale lithography, enabling the formation of smooth, continuous three-dimensional relief structures directly within the resist layer. These 3D profiles can be generated in a single exposure step, eliminating the need for physical photomasks or additional processing stages such as reactive ion etching or resist reflow. The process begins with spin-coating a layer of AZ® 4562, a positive-tone photoresist, onto a soda-lime glass substrate. The coated sample is then soft-baked on a hot plate at 120 °C for 3 minutes. Following a 3-hour rest period to ensure photoresist stability, the design is written using the DWL 66+ mask writer. Development is carried out using AZ® 726 MIF developer for 25 seconds, carefully calibrated to achieve the desired depth profile.

A.2 Nanoimprint Lithography

To permanently transfer the patterned photoresist structures into a robust optical material, we employ room-temperature nanoimprint lithography (NIL) using the Obducat Eitre 3 system. This process does not require thermal assistance; instead, it relies solely on the application of uniform pressure to achieve conformal contact between the mold and the imprint material. The initial grayscale structures, fabricated in AZ® 4562 photoresist on soda-lime glass substrates, serve as high-resolution master molds for the replication process.

The micro- and nanoscale diffractive optical element (DOE) patterns are transferred into OrmoComp, a UV-curable organic-inorganic hybrid polymer known for its excellent optical clarity, low volumetric shrinkage, and strong mechanical stability. These properties make OrmoComp particularly well-suited for faithfully replicating fine surface relief structures at ambient conditions.

During imprinting, a thin layer of uncured OrmoComp is dispensed onto a target Soda-lime substrate, and the master mold is pressed into the polymer under controlled pressure. Once conformal contact is achieved, the resist is exposed to UV light through the transparent substrate, initiating photopolymerization and solidifying the replicated structure. The mold is then released, yielding a high-fidelity replica of the original 3D relief with minimal deformation or feature loss. This pressure-only NIL approach simplifies processing, reduces thermal stress, and enables cost-effective fabrication of high-performance micro-optical components.

B ADDITIONAL IMPLEMENTATION DETAILS

B.1 Neural Lithography Model Architecture

We construct a small yet effective fully-convolutional neural network as the neural lithography model to learn the mapping from design layouts (on a $1\ \mu\text{m}$ -spacing grid) to their corresponding high-resolution AFM measurements, amounting to a super-resolution task. For designs with coarser features (such as $2\ \mu\text{m}$), we first perform nearest-upsampling from the original grid to the $1\ \mu\text{m}$ -spacing grid, and then apply the neural network to get final high-resolution predictions. The detailed architecture is specified in Table S1.

The neural lithography model consists of a convolutional neural network (CNN) with N residual blocks (set to 2 by default), each containing a convolutional layer followed by a ReLU activation function. These blocks extract hierarchical feature representations from the input design layout. To generate a high-resolution output, the network employs a projection layer followed by a PixelShuffle operation for upsampling. Skip connections, including a global skip connection that adds the input to the projection output, are used to stabilize gradient flow and mitigate the vanishing gradient problem. This is particularly important for fabrication-aware optimization, where gradients must propagate effectively back to the input design layout to enable precise adjustments.

For very large-scale designs, we adapt the default architecture by reducing the number of blocks N and the intermediate channel dimension F_i to optimally balance prediction precision with computational efficiency. The final upsampling to the target grid is achieved with a defined ratio r .

Table S1: Architecture of the Neural Lithography Model. The model is a convolutional neural network (CNN) with N residual blocks and skip connections, designed for feature extraction and upsampling. Here, r denotes the upsample ratio. By default, we set $F_0 = F_i = 100$ and $N = 2$. For very large-scale designs (e.g., ultra-high-definition 2-D holograms, see Fig. 11), fewer blocks and a lower channel dimension are used to balance precision and computational efficiency.

Component	Operation	Output Shape	Details
Input	–	$(B, H, W, 1)$	Batch, Height, Width, Channels
Initial Feature Extraction	Conv + ReLU	(B, H, W, F_0)	Kernel: 3×3 , Features: F_0
Residual Blocks (N blocks)	Conv + ReLU Add	(B, H, W, F_i)	Kernel: 3×3 Residual connection
Projection	Conv	(B, H, W, r^2)	Kernel: 1×1
Global Skip Connection	Add	(B, H, W, r^2)	Adds the initial input
Upsampling	PixelShuffle	$(B, H \cdot r, W \cdot r, 1)$	Rearranges channels to spatial dimensions
Output	–	$(B, H \cdot r, W \cdot r, 1)$	Final high-resolution output

B.2 Hologram Reconstruction Loss

To optimize a DOE for generating a 2-D hologram, we minimize a scale-invariant mean square error loss with an additional energy regularization. Let

$$I \in \mathbb{R}^{H \times W}$$

denote the predicted hologram intensity with spatial resolution $H \times W$. Let

$$T \in \mathbb{R}^{H \times W}$$

be the target hologram pattern, and let $\epsilon > 0$ be a small constant for numerical stability. The total loss is

$$\mathcal{L}_p = \mathcal{L}_{\text{si-mse}} + \beta \mathcal{L}_{\text{energy}},$$

where $\beta > 0$ is a weighting factor for the energy regularization term.

1. *Scale-Invariant MSE Loss.* The scale-invariant mean square error (si-MSE) loss minimizes the difference between the predicted hologram intensity I and the target T , allowing for an optimal scalar scaling factor α . The optimal scale is computed as

$$\alpha = \frac{\sum_{h=1}^H \sum_{w=1}^W I_{h,w} T_{h,w}}{\sum_{h=1}^H \sum_{w=1}^W I_{h,w}^2 + \epsilon}.$$

The si-MSE loss is then defined as

$$\mathcal{L}_{\text{si-mse}} = \frac{\sum_{h=1}^H \sum_{w=1}^W (\alpha I_{h,w} - T_{h,w})^2}{\sum_{h=1}^H \sum_{w=1}^W T_{h,w}^2 + \epsilon}.$$

This loss normalizes the MSE by the target's energy to ensure scale invariance.

2. *Energy Loss.* To prevent trivial solutions and encourage the total intensity to approach unity, we define the energy loss as

$$\bar{I} = \frac{1}{HW} \sum_{h=1}^H \sum_{w=1}^W I_{h,w}, \quad \mathcal{L}_{\text{energy}} = 1 - \bar{I} + \epsilon.$$

This term penalizes deviations of the average intensity \bar{I} from 1, with ϵ ensuring numerical stability.

B.3 Broadband Imaging PSF Loss

To optimize a DOE for broadband imaging, we combine several complementary loss terms. Let

$$I \in \mathbb{R}^{C \times H \times W}$$

denote the predicted PSF intensities for C spectral channels and spatial resolution $H \times W$. Let

$$T \in \mathbb{R}^{C \times H \times W}$$

be the target PSF pattern, and let $\epsilon > 0$ be a small constant for numerical stability. We write the image center pixel as (h_0, w_0) . The total loss is

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{psf}} + \mathcal{L}_{\text{focus}} + \mathcal{L}_{\text{cst}} + \mathcal{L}_{\text{energy}}$$

1. *PSF Matching Loss.* This term encourages the predicted PSF to align with the target distribution T . We take a log-barrier over the spatial inner product to penalize mismatches:

$$\mathcal{L}_{\text{psf}} = -\frac{1}{C} \sum_{c=1}^C \ln \left(\sum_{h=1}^H \sum_{w=1}^W I_{c,h,w} T_{c,h,w} + \epsilon \right).$$

By summing over all pixels (h, w) and channels c , this loss drives high correlation between predicted and target PSF intensities.

2. *Focus (Center-Energy) Loss.* To ensure that most energy is concentrated at the PSF center, we maximize the mean central pixel intensity:

$$\mathcal{L}_{\text{focus}} = -\ln \left(\frac{1}{C} \sum_{c=1}^C I_{c,h_0,w_0} + \epsilon \right).$$

A larger central value reduces this term, encouraging a sharp focal peak.

3. *Consistency Loss.* Real PSFs should exhibit both uniform peak energy (energy consistency) and similar spatial structure across channels (channel consistency).

Energy consistency measures per-channel fluctuations of the central energy around its mean:

$$\mu_{\text{cen}} = \frac{1}{C} \sum_{c=1}^C I_{c,h_0,w_0}, \quad \mathcal{L}_{\text{energy_cst}} = \frac{\frac{1}{C} \sum_{c=1}^C |I_{c,h_0,w_0} - \mu_{\text{cen}}|}{\mu_{\text{cen}} + \epsilon}.$$

Channel consistency enforces that each spectral channel's PSF shape does not deviate excessively from its own spectral channel mean:

$$\bar{I}_{h,w} = \frac{1}{C} \sum_{c=1}^C I_{c,h,w}, \quad \mathcal{L}_{\text{channel_cst}} = \frac{\sum_{h=1}^H \sum_{w=1}^W |\bar{I}_{h,w} - I_{c,h,w}|}{\sum_{h=1}^H \sum_{w=1}^W \bar{I}_{h,w} + \epsilon}.$$

The combined consistency loss is

$$\mathcal{L}_{\text{cst}} = \mathcal{L}_{\text{energy_cst}} + \mathcal{L}_{\text{channel_cst}}.$$

4. *Total-Energy Loss.* We also maximize the overall PSF energy to prevent trivial zero solutions:

$$\bar{I} = \frac{1}{CHW} \sum_{c,h,w} I_{c,h,w}, \quad \mathcal{L}_{\text{energy}} = -\ln(\bar{I} + \epsilon).$$

B.4 Additional RAW Image Processing Details

For single-DOE broadband imaging, we capture RAW images and perform the image reconstruction (deconvolution) in the linear-RGB space, thus entailing a RAW-to-linear-RGB preprocessing. As in [Brooks et al. 2019], our RAW-to-linear-RGB conversion comprises four stages: channel packing with level normalization, white-balance gain computation, bilinear demosaicing, and vignetting correction. Let

$$Y \in \{0, \dots, 2^{14} - 1\}^{H \times W}$$

be the raw 14-bit Bayer mosaic, with per-channel black levels $\mathbf{b} = [b_R, b_{G_1}, b_B, b_{G_2}]^T$ and fixed white level $W = 2^{14}$. Denote the 2×2 Bayer offsets by (Δ_c^r, Δ_c^c) for $c \in \{R, G_1, B, G_2\}$.

1. *Channel Packing & Level Normalization.* Define

$$P_c(x, y) = \frac{Y(2x + \Delta_c^r, 2y + \Delta_c^c) - b_c}{W - b_c}, \quad c \in \{R, G_1, B, G_2\},$$

with subsequent clamping to $[0, 1]$. This yields $P \in [0, 1]^{4 \times \frac{H}{2} \times \frac{W}{2}}$.

2. *White-Balance Gain Computation & Application.* Let the raw white-balance vector be $\mathbf{w} = [w_R, w_G, w_B, w_G]^T$, normalized to $\mathbf{g} = \mathbf{w}/w_G = [g_R, 1, g_B, 1]^T$. We form $\widehat{P}_c(x, y) = g_c P_c(x, y)$ for each channel c .

3. *Bilinear Demosaicing.* Denote by $\mathcal{U}\{\cdot\}$ the bilinear upsampler from $\frac{H}{2} \times \frac{W}{2}$ to $H \times W$. Then

$$R(i, j) = \mathcal{U}\{\widehat{P}_R\}(i, j), \quad B(i, j) = \mathcal{U}\{\widehat{P}_B\}(i, j),$$

and

$$G(i, j) = \frac{1}{2} \left(\mathcal{U}\{\widehat{P}_{G_1}\}(i, j) + \mathcal{U}\{\widehat{P}_{G_2}\}(i, j) \right).$$

Stacking gives the linear RGB tensor

$$X(c, i, j) \in [0, 1], \quad c \in \{R, G, B\}.$$

4. *Vignetting Correction.* To correct for lens vignetting, we capture a uniform white reference RAW Y_{ref} and process it identically to obtain $X_{\text{ref}}(c, i, j)$. We then compute a single-channel shading mask

$$M(i, j) = \frac{1}{3} \sum_{c \in \{R, G, B\}} X_{\text{ref}}(c, i, j),$$

and divide the target image by this mask:

$$X_{\text{corr}}(c, i, j) = \frac{X(c, i, j)}{M(i, j) + \varepsilon},$$

where ε is a small constant to avoid division by zero. Finally, we clamp X_{corr} to $[0, 1]$ to produce the vignetting-corrected, linear RGB image $\widehat{X}(c, i, j)$, which is then passed on to our inverse-filtering step.

B.5 Additional ISP Post-Processing Details

After the inverse filtering, we post-process the resulting restored linear-RGB images into sRGB images with simple image signal processing (ISP). Given a white-balanced, linear RGB image

$$X = X(c, i, j) \in [0, 1]^{3 \times H \times W},$$

we convert it into display-ready sRGB in three stages: color correction, tone mapping, gamma compression.

1. *Color Correction.* We apply a camera-to-RGB color correction matrix $\mathbf{C} \in \mathbb{R}^{3 \times 3}$ to each pixel:

$$\widetilde{X}(i, j) = \mathbf{C} X(i, j),$$

where $X(i, j) \in \mathbb{R}^3$ is the column vector of linear RGB at spatial location (i, j) .

2. *Global Tone Mapping.* We apply a smoothstep-like curve element-wise to compress highlights:

$$T(u) = \frac{1 - \sin(3 \arcsin(0.5 - u))}{2}, \quad u \in [0, 1].$$

Thus the tone-mapped image is

$$Y(i, j) = T(\widetilde{X}(i, j)) \in [0, 1]^3.$$

3. *Gamma Compression.* To approximate the sRGB transfer, we raise each channel to the power $1/\gamma$:

$$\widehat{X}(i, j) = \max(Y(i, j), \varepsilon)^{1/\gamma},$$

with $\gamma = 2.2$ (typical) and small $\varepsilon = 10^{-8}$ for numerical stability. The result \widehat{X} is the sRGB image fed to display or perceptual evaluation.

C ADDITIONAL RESULTS AND ANALYSIS

C.1 More Results on Lithography Model Calibration and Evaluation

We provide several extra results related to the lithography model calibration and evaluation, which include

- (1) The design layout pattern for contrast curve calibration (Fig. S1), which extends Fig. 4 in the main paper.
- (2) More evaluation results on lithography model forward predictability (Fig. S2), which extends Fig. 5 in the main paper.
- (3) Comparison of DOE design approaches (Table S2), which extends Fig. 3 in the main paper. The simulation computational holographic display results suggest that, of the fabrication-aware models, our approach outperforms Xu et al. [2025]. The performance of a DOE designed with the Xu’s model drops significantly (from 31.2 dB to 19.6 dB) when evaluated with our more accurate neural lithography model instead of its own. This indicates that the simplified physical model of Xu et al. lacks the fidelity to fully capture the complex fabrication process, limiting its effectiveness as a design surrogate. In contrast, our model achieves superior performance by providing a more precise representation of the true fabrication kernel both in simulation and real-world experiments (Fig. S3).
- (4) Quantitative comparison of neural lithography models with different upsampling factors (Table S3). While all models achieve similar PSNR when evaluated at their native resolution, their performance diverges significantly when evaluated at a fixed, high-resolution (10×) standard. The model trained with only 2× upsampling suffers a severe performance drop (from 21.8 to 17.4 dB). Performance consistently improves with higher training upsampling ratios. The 4× model shows acceptable performance (20.1 dB), while the 8× model performs nearly identically to the 10× model (20.8 dB vs. 20.9 dB), demonstrating that 8× upsampling is sufficient to accurately capture fine-grained fabrication effects with negligible model mismatch.

Table S2: Comparison of DOE Design Approaches (extension of Fig.3). We report quantitative evaluations (PSNR in db) of different diffractive optical element design strategies for creating 2-D hologram (*c.f.*, the task setting of Fig.3) : conventional (at 2 μm -spacing grid), conventional with nearest upsampling, physical lithography model of Xu et al. [2025] and our neural lithography model (all at 250-nm-spacing grid). Diagonal entries show performance when training and evaluation methods match; off-diagonals reveal model-mismatch assessment.

Eval \ Train	Conventional	Nearest Upsample	Xu et al. [2025]	Our Approach
Conventional	34.9	14.9	16.0	12.0
Nearest Upsample	16.8	39.2	20.7	14.1
Xu et al.	13.5	15.0	31.2	17.9
Our Approach	12.4	12.9	19.6	21.9

Table S3: Comparison of Neural Lithography Models with Different Upsampling Ratios. We report 2-D hologram reconstruction quality (PSNR in dB) for models trained with different upsampling factors. The first row shows performance when models are evaluated using their native (same) upsampling ratio. The second row shows performance when all models are evaluated using a fixed 10× upsampling ratio.

Eval \ Train	2× Upsample	4× Upsample	8× Upsample	10× Upsample
Same	21.8	21.8	21.9	20.9
10× Upsample	17.4	20.1	20.8	20.9

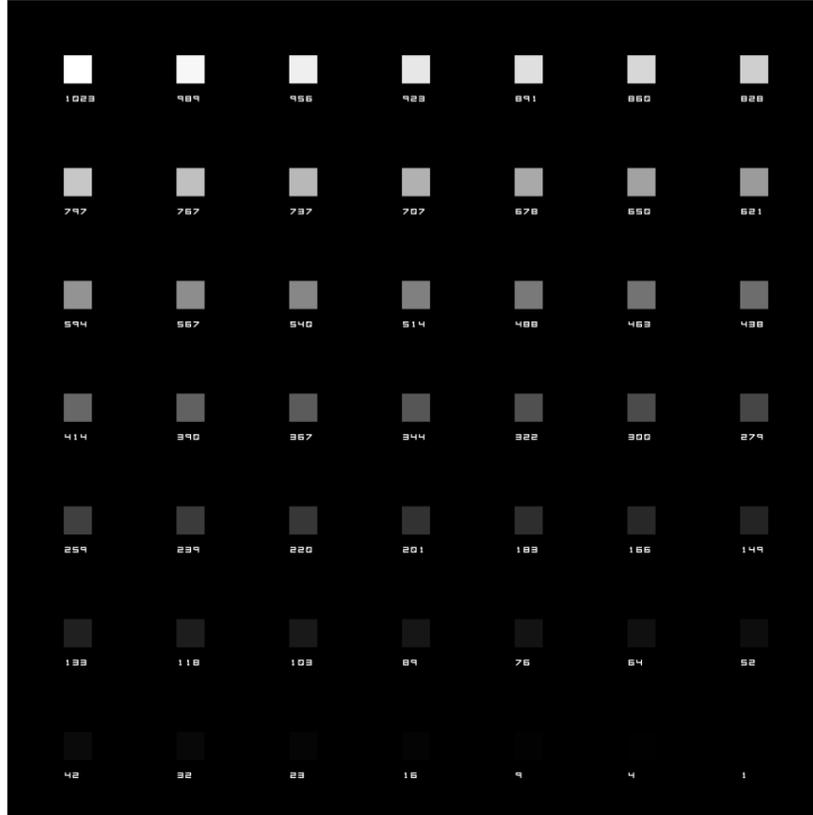


Fig. S1: Contrast Curve Calibration (extension of Fig.4). A 7×7 test pattern of uniform patches with incrementally varying gray values (1 to 1023, specified below each patch) is used to measure and plot a contrast curve, showing the relationship between gray values and developed resist depths.

C.2 Additional Experimental Results on Broadband Imaging

We provide additional comparison on raw measurements and inverse-filtered reconstructions to demonstrate the effectiveness of our fabrication-aware approach.

Figure S4 presents the raw measurements for the reference target (Column 1), a Fresnel DOE with analytical quadratic phase at $\lambda_0 = 550$ nm (Column 2), the conventional nearest-upsampling DOE (Column 3), and our fabrication-aware neural-upsampling DOE optimized via the broadband PSF loss (Column 4). The analytical quadratic phase is defined as:

$$\phi_{\text{quad}}(x, y) = \frac{\pi}{\lambda_0 f} (x^2 + y^2), \quad \lambda_0 = 550 \text{ nm},$$

Though the Fresnel DOE in Column 2 appears sharper at first glance, it exhibits pronounced chromatic aberration that varies with the scene's spectral distribution (compare Rows 1 and 2). This inconsistency forces reconstruction algorithms to rely heavily on cross-channel priors [Heide et al. 2013], undermining robustness. In contrast, DOEs optimized with our broadband PSF loss are engineered to produce spectrally invariant point-spread

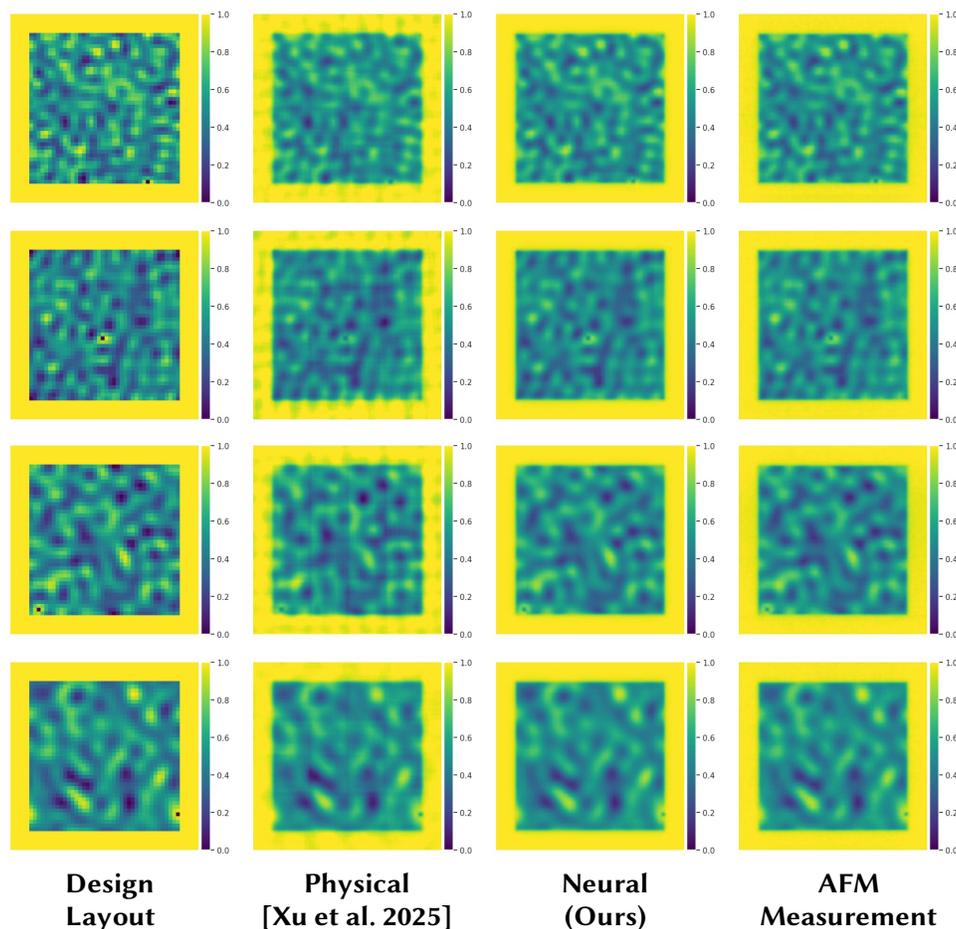


Fig. S2: Lithography Model Forward Predictability (extension of Fig.5). We show design patterns from the constructed evaluation set, their corresponding AFM measurements, and the predictions from two lithography models: a physical model [Xu et al. 2025] and our proposed neural lithography model.

functions, effectively eliminating strong color fringing across the visible band. Crucially, real-world measurements (Column 4 of Fig. S4) demonstrate that the fabrication-aware neural upsampling design delivers higher spatial resolution and significantly reduced chromatic aberration compared to the conventional nearest-neighbor upsampling DOE (Column 3). This performance gap between simulated design and fabricated device highlights the importance of incorporating fabrication and broadband effects into the optimization.

Next, we present additional inverse-filtering results (see Sec. 5.4). The new captured scenes further confirm that our fabrication-aware DOE produces measured PSFs nearly identical to those predicted in simulation. We highlight the scene in Row 2 of Fig. S5 to demonstrate how residual chromatic aberration—stemming from imperfect fabrication modeling—induces pronounced ringing artifacts in the downstream non-blind reconstruction.

Finally, we experimentally measured the PSFs of both the conventional nearest-neighbor upsampling DOE and the proposed fabrication-aware DOE. To validate the accuracy of our lithography model, we performed inverse

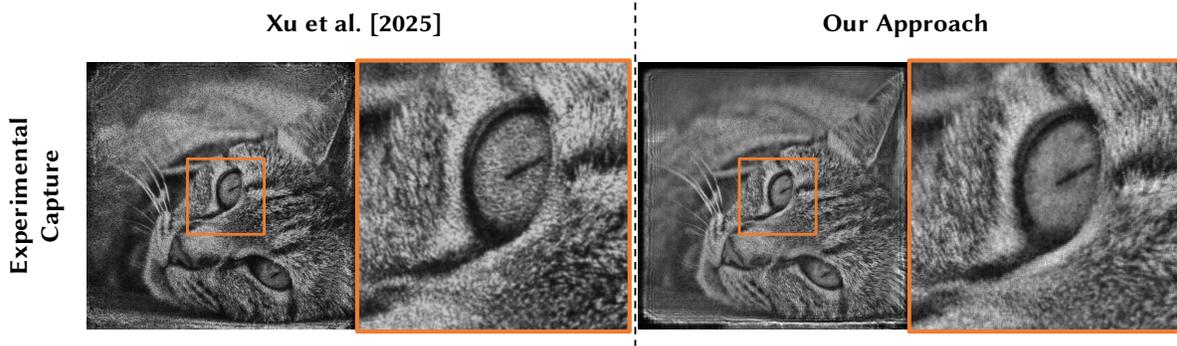


Fig. S3: Experimental Comparisons on Computational Holographic Display. We present captured 2-D holograms from both Xu et al.'s method [Xu et al. 2025] and our approach (both optimized with $8\times$ upsampling on a 250-nm-spacing grid). Compared to the linear physical model used by Xu et al., our method produces cleaner textures with fewer speckle artifacts.



Fig. S4: Raw Measurement Comparisons of broadband imaging DOEs. The classic Fresnel DOE yields sharper images but exhibits significant scene-dependent chromatic aberration (e.g., purple artifacts), challenging to correct algorithmically.

filtering using both the experimentally measured PSFs and their simulated counterparts. If the deconvolution results obtained from the same measurement closely match when using simulated and experimental PSFs, this indicates that the predicted PSF reliably approximates the measured one. The results in Fig. S6 confirm this alignment, demonstrating the effectiveness of the proposed fabrication-aware model. This is consistent with the modulation transfer functions corresponding to both the simulated and experimental PSFs in Fig. S7.



Fig. S5: Experimental Validation of Broadband Imaging Performance. Inverse deconvolution results from additional captured scenes verify close agreement between measured and simulated point spread functions using our fabrication-aware DOE design. Representative results (Row 2) reveal residual chromatic artifacts arising from fabrication errors, manifesting as ringing artifacts post-deconvolution as for the conventional design.

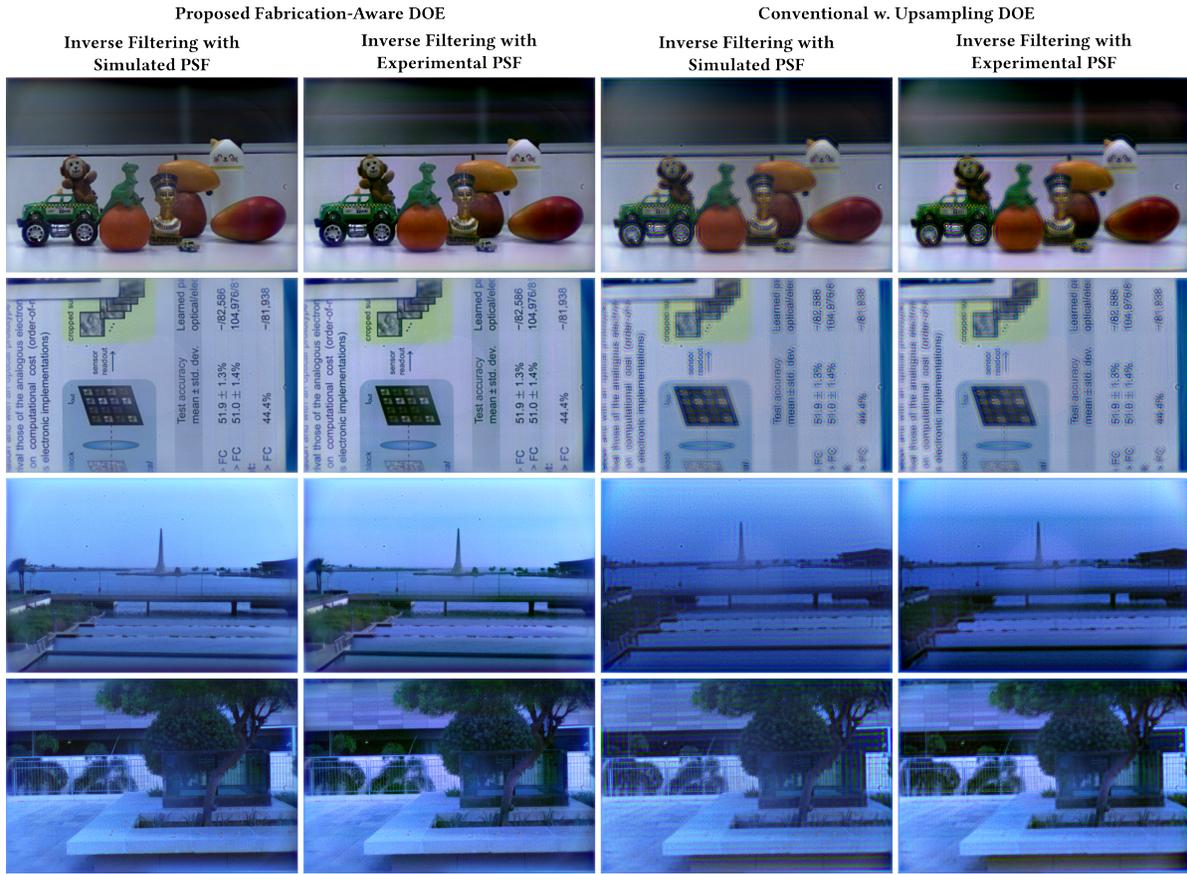


Fig. S6: Comparison of Inverse Filtering with Simulated and Experimental PSFs. Deconvolution results obtained using both measured and simulated PSFs for the conventional nearest-neighbor upsampling DOE and the proposed fabrication-aware DOE. The close agreement between results validates that the simulated PSFs reliably approximate the *de facto* ones, demonstrating the accuracy of the proposed lithography model.

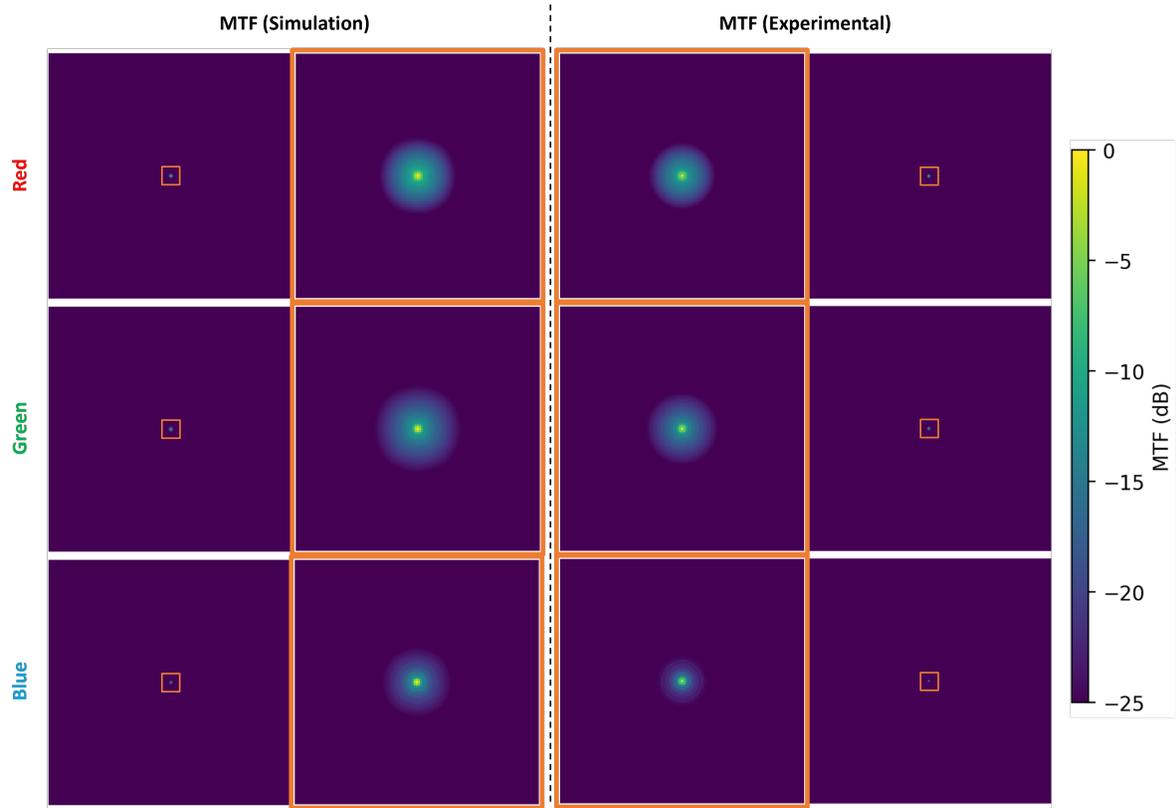


Fig. S7: Simulated vs. Experimental MTF. A comparison of the modulation transfer function (log scale) from the fabrication-aware simulation and the experimental measurement based on the captured PSF.



Fig. S8: Experimental Assemblies for assessing the performance of diffractive optical elements in broadband imaging applications.

D ADDITIONAL EXPERIMENTAL PROTOTYPE DETAILS

D.1 Broadband Imaging Setup

Fig. S8 depicts the assemblies adapted for broadband imaging evaluation, designed to capture visually similar scenes. To mount the fabricated DOEs, we used telescopic lens tubes (Thorlabs SM1NR1), which enabled precise positioning at the correct focal distance. The DOEs were attached to custom acrylic holders, 2.8 mm thick, cut using a laser cutter and designed with sufficient clearance to accommodate the square-shaped elements. For reference imaging, we used an achromatic lens with a focal length of $f = 80$ mm, in combination with an iris aperture set to 2 mm. The capturing device was a full-frame Canon camera (EOS 5D Mark IV), and the data were recorded in CR2 RAW format.

REFERENCES

- Tim Brooks, Ben Mildenhall, Tianfan Xue, Jiawen Chen, Dillon Sharlet, and Jonathan T Barron. 2019. Unprocessing images for learned raw denoising. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 11036–11045.
- Felix Heide, Mushfiqur Rouf, Matthias B Hullin, Bjorn Labitzke, Wolfgang Heidrich, and Andreas Kolb. 2013. High-quality computational imaging through simple lenses. *ACM Transactions on Graphics (ToG)* 32, 5 (2013), 1–14.
- Yunpeng Xu, Zihan Zang, Haoqiang Wang, Yanjun Han, Hongtao Li, Yi Luo, Lai Wang, Changzheng Sun, Bing Xiong, Zhibiao Hao, Jian Wang, and Lin Gan. 2025. Fabrication-integrated design for diffractive optical elements. *Optica* 12, 2 (Feb. 2025), 228. <https://doi.org/10.1364/OPTICA.539824>