

Complex-Valued 2D Gaussian Representation for Computer-Generated Holography

Yicheng Zhan^{*1}, Xiangjun Gao^{*2}, Long Quan², and Kaan Aksit¹

¹ University College London, London, UK

² Hong Kong University of Science and Technology (HKUST), Hong Kong, China
yicheng_zhan2001@outlook.com

Abstract. Complex-valued Gaussian primitives have recently been explored for representing holographic radiance fields in 3D novel view synthesis. In this work, we extend this line of research to the hologram optimization domain and propose a structured representation based on complex-valued 2D Gaussian primitives. Inspired by Gabor’s theory, we show that our primitive attains the minimum space–frequency uncertainty and reduces the parameter search space by 5:1 compared to per-pixel parameterization. To enable end-to-end training, we develop a differentiable rasterizer for our representation, integrated with a GPU-optimized light propagation kernel in free space. Extensive experiments show that our method reduces VRAM usage by up to 30% and accelerates optimization by 50% over standard autodiff-based implementations, delivers up to 13 dB higher PSNR than prior Gaussian-based methods, and achieves up to 3200× faster rendering while maintaining reconstruction quality on par with existing CGH approaches. For evaluation, we introduce a conversion procedure that adapts our representation to practical hologram formats, including smooth and random phase-only holograms. By reducing the hologram parameter search space, our representation enables a more scalable hologram estimation in the next-generation computer-generated holography systems.

1 Introduction

Holographic displays are a promising technology for realistic Three-Dimensional (3D) content presentations [33]. Unlike natural images, which record only light intensity, holograms capture intensity, interference, and diffraction phenomena. As shown in Fig 1, compared to natural images, holograms exhibit markedly different spatial characteristics. Therefore, a key challenge in Computer-Generated Holography (CGH) is to design compact and efficient representations that preserve high-frequency details of holograms while being scalable [56, 63].

Conventional image representation methods, such as Implicit Neural Representation (INR) [45, 49, 57], optimize a continuous implicit function to represent an image. However, implicit functions generally favor continuous low-frequency

* Equal Contribution.

data, making it difficult to capture the high-frequency details typical of holograms. In parallel, autoencoder-based approaches [7, 46, 58, 62], typically pre-trained on natural images, rely on learned priors that do not readily generalize to hologram structures, hinting a need for hologram specific adaptations. Additionally, emerging Gaussian-based image representations [69, 74, 75] have recently been proposed for natural image modeling. While effective for natural images, these methods directly encode hologram pixels without modeling interference and diffraction phenomena, leaving room for more specialized hologram representations and better reconstruction quality.

More recently, Gaussian primitives have also been adopted in CGH for 3D novel view synthesis [12, 14, 71]. In this work, we extend this line of research to further discover the potential of Gaussians in the hologram optimization domain.

Instead of modeling a 3D scene from multi-view images, we focus on representing a *single* hologram through complex-valued 2D Gaussian primitives for efficient and scalable hologram estimation. Inspired by Gabor’s theory [21], we show that our primitive

attains the minimum space-frequency uncertainty, making it the compact and efficient primitive for hologram representation. To enable end-to-end optimization, as the computational backbone of our method, we develop a differentiable rasterizer for complex-valued 2D Gaussians, integrated with a GPU-optimized light propagation kernel, enabling efficient and scalable hologram optimization and rendering. Our comprehensive evaluation shows that, our method reduces the parameter search space by a **5:1** ratio compared to dense per-pixel parameterization (the conventional hologram optimization setting), decreases VRAM usage by **30%**, and accelerates optimization by **50%** compared to standard autodiff-based developments (PyTorch [43]), while achieving reconstruction fidelity comparable to prior Gaussian-based and CGH methods. Our contribution is summarized as follows:

- We migrate the complex-valued 3D Gaussians from CGH 3D novel view synthesis [71] to **complex-valued 2D Gaussian** primitives. We show the effective usage of this primitive in hologram optimization by reducing the parameter search space by **5:1** while preserving reconstruction fidelity (degrading gracefully up to a **10:1** upper bound).

Additionally, for evaluation purposes, we introduce a conversion procedure that adapts our complex-valued 2D Gaussian representation to practical hologram formats, including Smooth Phase-only Hologram (POH) via double-phase coding [26] and Random POH via structural guided phase optimization. We hope this work demonstrates the broader potential of Gaussian primitives in CGH and encourages further investigation into compact, scalable hologram optimization




Fig. 1: Comparison between a natural image and different hologram formats. Unlike the smoother pixels in natural images, holograms produce dense high-frequency and random spatial variations that are challenging to represent (Source Image: [1]).

frameworks. For readers less familiar with the underlying principles of CGH, we present the essential concepts in Sec. 2.2. Our code is available at [70].

2 Related Work

2.1 Gaussian and Learned Image Formation

Recent advances in natural image representation have explored compact alternatives to dense pixel-wise parameterizations. Autoencoder-based methods [7, 46, 62] and INRs [36, 49, 50, 57] map natural images into latent spaces or continuous functional representations, enabling structured compression but often favoring smooth, low-frequency content. More recently, building on 3D Gaussian Splatting (3DGS) and its variants [22, 32, 39, 64], several works have extended Gaussian primitives from neural rendering to natural image encoding [69, 74, 75], leveraging 2D Gaussians for efficient image rendering and compression. While these Gaussian-based methods show promising results on representing natural images, they do not model interference and diffraction phenomena, and are therefore not directly suited for hologram representation. In this paper, we explore and propose complex-valued 2D Gaussian primitives as a hologram representation that explicitly integrates light propagation during optimization, targeting both parameter search space reduction and high-quality hologram reconstruction.

2.2 Preliminary Concepts: CGH

CGH is a **computational imaging** task that performs wave-based rendering, synthesizing **holograms** to reconstruct 3D scenes for holographic displays, where the display is typically implemented as Spatial Light Modulator (SLM) as the programmable wavefront device. Unlike natu-




Fig. 2: (Simulated) Smooth hologram shows high quality at the pupil center but degrade severely with (SLM) as the programmable pupil shifts, whereas random hologram remains visible (Source Image: [60]).

ral images that capture only smooth, low-frequency intensity variations, a hologram simultaneously encodes light’s intensity, interference, and diffraction. As shown in Fig 1, this wave-optical data format manifests as dense, high-frequency, and random-valued structures that challenge conventional image representations. Holographic data can be represented as a complex hologram: $\mathbf{H} = A \exp(j\varphi)$ with amplitude A and phase φ , requiring specialized modulators; or as a phase-only hologram: $\mathbf{H}_{\text{POH}} = \exp(j\varphi)$, compatible with commercial holographic displays. POH variants include Smooth POH with spatial multiplexed phase [26] and Random POH with directly optimized phase [51]. For more details of different hologram formats, please refer to Sec. 3.3.

Light Propagation. The core of CGH is free-space propagation based on scalar diffraction theory [23]. As shown in Fig 3, free-space propagation over a distance d can be expressed as a 2D convolution of the source field $\mathbf{U}(0)$ with the spatial impulse response h_d ,

$$\mathbf{U}(d) = \mathbf{U}(0) * h_d, \quad (1)$$

where 0 denotes the source plane at zero distance. Equivalently, in the frequency domain

$$\mathbf{U}(d) = \mathcal{F}^{-1}\{H_d(f_x, f_y) \mathcal{F}\{\mathbf{U}(0)\}\}, \quad (2)$$

where $H_d(f_x, f_y) = \mathcal{F}\{h_d\}$ is the transfer function [40, 76]. A common choice of H is the Band-limited Angular Spectrum Method (BLASM)

$$H_d(f_x, f_y) = \begin{cases} \exp\left(j 2\pi d \sqrt{\frac{1}{\lambda^2} - r^2}\right), & r^2 \leq \frac{1}{\lambda^2}, \\ 0, & \text{else,} \end{cases} \quad (3)$$

where $r^2 = f_x^2 + f_y^2$ and λ denotes the working wavelength. In this paper, we further optimized BLASM kernel to model light propagation, being 50% faster and 30% VRAM-efficient than PyTorch development.

2.3 Holographic Representation

Conventional CGH either optimizes per-pixel holograms [11, 29, 30, 34, 52], or trains neural networks to directly predict hologram pixels [17, 44, 54, 55, 72], both of which yield large solution spaces and hinder scalability. More recently, Gaussian primitives have been introduced to bridge computer graphics and holography. Gaussian Wave Splatting (GWS) [14] and Random-phase GWS [12] leverage pretrained 2DGS scenes [27] for geometry-aware modeling of interference and diffraction, while complex-valued holographic radiance fields [71] demonstrate that 3D Gaussians can directly represent volumetric holographic scenes. A concurrent effort further applies 2D Gaussians to compress double-phase holograms [20]. In this work, we extend this discovery to the hologram optimization domain by proposing a complex-valued 2D Gaussian representation that explicitly integrates light propagation during optimization, achieving both parameter search space reduction and competitive reconstruction fidelity.

3 Method

Problem Definition Given a target image $\mathbf{I}_{\text{target}} \in \mathbb{R}^{C \times H \times W}$ and an optional depth map $\mathbf{D} \in \mathbb{R}^{H \times W}$, we aim to synthesize a complex 3D hologram $\mathbf{H} \in \mathbb{C}^{C \times H \times W}$ whose optical propagation reconstructs $\mathbf{I}_{\text{target}}$ with correct focus and defocus. For $C = 3$ (RGB), each channel of \mathbf{H} is complex-valued (real and imaginary parts), equivalently $\mathbf{H} \in \mathbb{R}^{6 \times H \times W}$.




Fig. 3: Hologram reconstruction via free-space light propagation (Source Image: [31]).




Fig. 4: Overview of our pipeline. Complex-valued 2D Gaussians are rasterized into a complex hologram (amplitude and phase), which is propagated to multiple depth planes using optimized light propagation. Reconstructions are compared with RGB+D derived targets at different focal distances, and we report PSNR, SSIM, and LPIPS. Here, *Light Prop* denotes light propagation (Source Image: [31].)

3.1 Complex-Valued 2D Gaussian Primitives

Fig 4 shows the training pipeline of our method. Building on the development of complex-valued holographic radiance fields [71], we extend the representation to define a complex-valued 2D Gaussian primitive for single-hologram optimization. Whereas [71] models volumetric 3D scenes from multi-view images, our formulation operates on the hologram plane directly, targeting efficient and scalable hologram estimation. Each primitive \mathcal{G}_n is parameterized as

$$\mathcal{G}_n = \{\tilde{\mathbf{x}}_n, \tilde{\mathbf{s}}_n, \theta_n, \mathbf{c}_n, \varphi_n, \tilde{\alpha}_n\}, \quad (4)$$

where $\tilde{\mathbf{x}}_n \in \mathbb{R}^2$ denotes the pre-activation 2D position, $\tilde{\mathbf{s}}_n \in \mathbb{R}^2$ the pre-activation scales, $\theta_n \in \mathbb{R}$ the in-plane rotation angle, $\mathbf{c}_n \in \mathbb{R}^C$ the color amplitudes, $\tilde{\alpha}_n \in \mathbb{R}$ the pre-activation opacity, and $\varphi_n \in \mathbb{R}^C$ the per-channel phase. Naively, a complex-valued field can be represented by pairing two real-valued Gaussians for the real part and the imaginary part, respectively. This demands 18 parameters per primitive pair and two separate renderings, one for each real-valued Gaussian, resulting in extra computation. By contrast, our formulation has 12 parameters, achieving a $\frac{1}{3}$ reduction in parameterization while requiring a single rendering. We keep the amplitude factored as $\alpha_n \mathbf{c}_n$ rather than a single value (ablated in Suppl Sec. 11). We apply activation functions to enforce valid parameter ranges (See Suppl Sec. 9.4 for activation formulations). In the following equations, \mathbf{x}_n , \mathbf{s}_n , and α_n denote the activated parameters obtained from their corresponding pre-activation counterparts $\tilde{\mathbf{x}}_n$, $\tilde{\mathbf{s}}_n$, and $\tilde{\alpha}_n$. The spatial distribution of the Gaussian is defined by a 2D covariance matrix $\Sigma_n = \mathbf{R}(\theta_n) \mathbf{S}_n^2 \mathbf{R}(\theta_n)^\top$, where $\mathbf{R}(\theta_n)$ is the rotation matrix and $\mathbf{S}_n = \text{diag}(\mathbf{s}_n)$. The inverse covariance Σ_n^{-1} is computed analytically (see Suppl Sec. 9.4 for the full covariance calculation). The contribution of \mathcal{G}_n at the pixel coordinate \mathbf{p} is

$$g_n(\mathbf{p}) = \exp\left(-\frac{1}{2}(\mathbf{p} - \mathbf{x}_n)^\top \Sigma_n^{-1}(\mathbf{p} - \mathbf{x}_n)\right), \quad (5)$$

and the complex-valued hologram pixel at \mathbf{p} is

$$\mathbf{H}_n(\mathbf{p}) = \alpha_n \mathbf{c}_n g_n(\mathbf{p}) \exp(j\varphi_n). \quad (6)$$

The hologram is formed by the accumulation of all primitives in pixel grid

$$\mathbf{H} = \left\{ \sum_{n=1}^N \mathbf{H}_n(\mathbf{p}) \mid \mathbf{p} \in [1, W] \times [1, H] \right\}. \quad (7)$$

Connection to Gabor’s Theory. Each pixel on the SLM diffracts light into propagation angles. The Fourier transform of the hologram field maps these pixel values to spatial frequencies, defining the reconstructed image. In holography, we are capturing both spatial (where light originates) and frequency information (the angles at which it propagates) simultaneously. Gabor’s uncertainty principle [21] states that the more precisely we know one of these properties, the less precisely we can know the other—a fundamental constraint for all wave-based systems. For hologram representation, we therefore seek functions that minimize the joint uncertainty between space and frequency domains. The Gaussian is the unique function achieving this theoretical minimum bound $\Delta x \Delta f_x = \frac{1}{2}$ [21]. To illustrate, a rectangular aperture produces a sinc-function Fourier transform with high-frequency oscillations, whereas a Gaussian aperture yields a Gaussian Fourier transform with minimal space–frequency spread. Our primitive $\mathbf{H}_n(\mathbf{p}) = \alpha_n \mathbf{c}_n g_n(\mathbf{p}) \exp(j\varphi_n)$ is the two-dimensional analogue of Gabor’s 1D **elementary signal** [21], $\psi(x) = \exp(-\beta^2(x - x_0)^2) \cdot \exp(j(2\pi f_0 x + \varphi))$. Since a 2D Gaussian factorizes along its principal axes, it attains minimum space–frequency uncertainty in both dimensions, meaning fewer primitives are needed to faithfully represent a given hologram (validated in Suppl Sec. 12). Note that high-frequency content in our reconstructed field does not arise from variations within individual primitives, but from the superposition of N primitives at positions \mathbf{x}_n and from the frequency-dependent phase applied by BLASM.

3.2 Hologram Reconstruction and Optimization

Naively, we can supervise \mathbf{H} with \mathbf{H}_{target} , however, this approach is insufficient in practice. To faithfully represent a hologram, it is essential to explicitly incorporate light propagation during optimization. Given a hologram field $\mathbf{U}(0) = \mathbf{H}$, we simulate free-space propagation using the convolutional methods summarized in Sec. 2.2. Specifically, we adopt the BLASM [40] as H_d and compute $\mathbf{U}(d) = \mathcal{F}^{-1}\{H_d^{\text{BLASM}}(f_x, f_y) \mathcal{F}\{\mathbf{U}(0)\}\}$. To capture depth-dependent effects, we reconstruct $\mathbf{U}(0)$ on L uniformly spaced parallel planes $\Pi_{l=1}^L$ (spacing distance Δz , e.g., 2 mm) along the optical axis, centered at propagation distance d_0 (e.g., 5 mm), yielding $d_1 = d_0 - \frac{L-1}{2}\Delta z$, $d_L = d_0 + \frac{L-1}{2}\Delta z$, and the reconstructed intensity at plane l is $I_l = |\mathbf{U}(d_l)|^2$. Naively, we can directly supervise I against the target image \hat{I} per depth plane l using $\mathcal{L}_{MSE} = \frac{1}{L} \sum_{l=1}^L \|I_l - \hat{I}_l\|^2$, where \hat{I}_l denotes the ground-truth multi-plane focal stack images generated by Kavaklı et al. [30]. To further improve the defocus region’s image quality, we utilize the reconstruction loss \mathcal{L}_{recon} by Kavaklı et al. [29], computed as

$$\mathcal{L}_{recon} = \frac{1}{L} \sum_{l=1}^L (\|I_l - \hat{I}_l\|^2 + \|I_l \cdot M_l - \hat{I}_l \cdot M_l\|^2 + \|I_l \cdot \hat{I}_l - \hat{I}_l \cdot \hat{I}_l\|^2), \quad (8)$$

where M_l is the binary mask for depth plane I_l generated from the target image and its quantized depth. We also employ the SSIM loss, the final training loss is $\mathcal{L} = \mathcal{L}_{recon} + \lambda_1 \cdot \mathcal{L}_{SSIM}$, where $\lambda_1 = 0.005$.

3.3 POH Conversion Procedure

Our complex-valued 2D Gaussians provide a compact, efficient hologram representation in the complex domain. Although there are prototypes utilizing full complex holograms [53], the commercial holographic displays are predominantly **Liquid Crystal based** phase-only displays that do not support displaying complex holograms directly [6]. To bridge this gap, we design a simple yet effective conversion procedure that adapts our complex representation as *structural guidance* to different hologram formats, including Smooth POH and Random POH. Rather than optimizing each device-specific format end-to-end, we keep a single complex representation as a general intermediate convertible to *both* formats, decoupling the representation from hardware-specific encoding.

Smooth POH. Our hologram is represented as $\mathbf{H} = \alpha \mathbf{c} g \cdot \exp(j\varphi)$. We employ Double Phase-Amplitude Coding (DPAC) [26] to convert \mathbf{H} into a smooth, phase-only representation by spatially multiplexing amplitude $A = \alpha \mathbf{c} g$ and phase φ via a checkerboard pattern: $\varphi_{\text{DPAC}}(i, j) = A(i, j)$ if $(i + j)$ is even, and $\varphi(i, j)$ if odd. The converted Smooth POH is $\mathbf{H}_{\text{smooth}} = \exp(j\varphi_{\text{DPAC}})$.

Random POH. We leverage our complex representation as structural guidance for the Random POH conversion. The Random POH is parameterized as $\mathbf{H}_{\text{rand}} = \exp(j\varphi_{\text{rand}})$, where $\varphi_{\text{rand}} \in \mathbb{R}^{C \times H \times W}$ denotes learnable randomly-initialized phase values. Both \mathbf{H} and \mathbf{H}_{rand} are propagated through the BLASM in parallel, obtaining reconstructions at the same depth plane l with $I_l = |\mathbf{U}(d_l)|^2$ and $I_{\text{rand}}^{(l)} = |\mathbf{U}_{\text{rand}}(d_l)|^2$. We jointly optimize the following objectives in both intensity and complex field domains

$$\begin{aligned} \mathcal{L}_{\text{extract}} = \sum_{l=1}^L & \left[\mathcal{L}_{\text{recon}}(I_l, \hat{I}_l) + \mathcal{L}_{\text{recon}}(I_{\text{rand}}^{(l)}, \hat{I}_l) \right. \\ & \left. + \lambda_{\text{comp}} \|I_l - I_{\text{rand}}^{(l)}\|^2 + \lambda_{\text{field}} \|\mathbf{U}(d_l) - \mathbf{U}_{\text{rand}}(d_l)\|_{1, \mathbb{C}} \right], \end{aligned} \quad (9)$$

where $\lambda_{\text{comp}} = 0.1$ and $\lambda_{\text{field}} = 0.01$. Here, $\|\cdot\|_{1, \mathbb{C}}$ is the sum of the L1-norms of the real and imaginary components.

3.4 Efficient CUDA Rendering and Propagation

We develop our hologram representation pipeline using CUDA, covering both complex-valued rasterization and light propagation in the Fourier domain.

Complex-Valued 2D Gaussian Rasterizer. We adapt the tile-based rasterizer from complex-valued holographic radiance field [71] to 2D Gaussians for a *single hologram*. Each primitive contributes a complex value instead of a real scalar. The forward pass decomposes Gaussians into real and imag components via trigonometric evaluation. We retain the 16×16 tile structure with duplicate-with-keys assignment and radix sorting, but store only per-pixel opacity, recovering intermediate values during the backward pass by division for constant VRAM overhead. Gradients of amplitude and phase require trigonometric chain rules with negated sine/cosine terms from the complex exponential derivative. Position, covariance, and opacity gradients follow [32] and adapted to 2D screen-space. For the details of development and derivation, please refer to Suppl Sec. 9.

GPU-optimized Light Propagation. Additionally, as part of the rasterizer, we develop a BLASM kernel that processes spatial frequencies in parallel and evaluates the transfer function to simulate light propagation effectively. Valid frequencies are multiplied by the transfer function using accelerated trigonometric operations. The backward pass applies the conjugate transfer function while preserving bandlimiting, and achieves efficiency through coalesced read-only cache access. For the details of development and derivation, please refer to Suppl Sec. 10.

4 Implementation

We initialize N Gaussians by uniformly sampling image-plane positions $\mathbf{x}_n^{\text{raw}} \sim \text{Uniform}([0, W] \times [0, H])$ and transforming to the unconstrained domain via $\tilde{\mathbf{x}}_n^{\text{init}} = \text{atanh}(2\mathbf{x}_n^{\text{raw}}/[W, H] - 1)$. Scales are set to $\tilde{s}_n = \log([1.5, 5.0])$ pixels, colors \mathbf{c}_n sampled from $[0, 1]$, phases φ_n initialized to zero, and opacity pre-activations fixed at $\tilde{\alpha}_n = -0.5$ (yielding $\alpha_n \approx 0.38$ post-sigmoid). We optimize using Adan [66] with empirically-selected, parameter-specific learning rates: positions 10^{-2} (cosine-annealed to 10^{-3} [38]), scales 5×10^{-3} , amplitudes and phases 2.5×10^{-3} , opacities 2.5×10^{-2} , and rotations 10^{-3} . Training runs for 2000 steps on a single NVIDIA RTX 3090 GPU with a 5:1 parameter reduction ratio, with convergence typically occurring around 1000+ steps (training visualizations in Suppl Sec. 14). Depth maps are provided by Depth Anything v2 [67] and MiDaS [47]. We adopt wavelengths of 639nm, 532nm, and 473nm for red, green, and blue channels respectively, with pixel pitch $3.74\mu\text{m}$, propagation distance 3mm, and volume depth 4mm, consistent with prior work [2, 54, 55, 72]. Our experiments use a LASOS MCS4 RGB laser, Jasper JD7714 phase-only SLM (2400×4094 , $3.74 \mu\text{m}$), and a lens-based optical relay with spatial filtering. Reconstructions are recorded by a Point Grey GS3-U3-23S6M-C lensless sensor mounted on a motorized stage. Hardware details are provided in Suppl Sec. 8.

5 Evaluation

In this section, we conduct a comprehensive evaluation of our method against baselines. We use the first 50 images from DIV2K [61] dataset as the test set; for each method, training and evaluation are performed at a resolution of $3 \times 1024 \times 640$ ($L = 2$). We report mean PSNR, SSIM, and LPIPS [73], averaged over both test images and across planes, together with parameter counts, peak VRAM usage, and training time. For results of training method ($L = 3$), comparing CUDA-based development with the PyTorch baseline. Red downward arrows and percentages indicate the VRAM reduction rate.




Fig. 5: Runtime (bar) and VRAM usage (line) across spatial resolutions for our method ($L = 3$), comparing CUDA-based development with the PyTorch baseline. Red downward arrows and percentages indicate the VRAM reduction rate.

5.1 Runtime and Memory Performance

Fig 5 presents the performance results of our kernel; the development detail is provided in Sec. 3.4. Across all resolutions, our CUDA kernel consistently reduces VRAM usage by 29–36% and accelerates runtime by 40–50% compared to the PyTorch developments. Specifically, memory savings are 28.9%, 30.7%, 35.6%, and 31.9% at 768×512 , 1024×1024 , 1920×1080 , and 2560×1440 , respectively. The runtime gains come primarily from the BLASM components: the Forward pass is 47–67% faster and the Backward pass is 56–58% faster, while the rasterization stages remain unchanged. Consequently, the overall step time is reduced by 39.7%, 42.6%, 47.7%, and 50.1%, with corresponding memory reductions of 28.9%, 30.7%, 35.6%, and 31.9% at the four resolutions. With $L = 3$, our method scales up to a resolution of 3200×1800 (5.8M pixels) without Out of Memory (OOM), validating the scalability and efficiency of our method. **Table 1:** Quantitative comparison of our method, Gaussian-based and learned representation methods. Top 2 metrics are highlighted in orange and yellow (* pretrained).

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	VRAM	Params	Time (min)
TAESD* [7]	11.6	0.09	0.79	2.7 G	2.5 M	-
MLP [49]	7.5	0.04	0.85	9.9 G	1.0 M	6.9
SIREN [57]	7.6	0.05	0.84	13.1 G	1.0 M	7.8
Image-GS [69]	17.2	0.29	0.70	1.3 G	2.4 M	1.6
Instant-GI [75]	23.5	0.56	0.56	3.4 G	2.8 M	0.9
GI [74]	22.6	0.49	0.59	1.1 G	2.4 M	0.8
Ours	30.7	0.86	0.33	2.2 G	0.8 M	1.4

5.2 Comparison With Representation Methods

Reconstruction Fidelity. Table. 1 compares our method with Gaussian-based and learned representations for complex hologram modeling. Learned representations show limited effectiveness: MLP [49] and SIREN [57] yield PSNR below

8 dB and LPIPS around 0.85, while pretrained TAESD [7] reaches 11.6 dB, suggesting that INR and autoencoder-based methods are not well-suited for encoding hologram structures. Gaussian-based methods perform substantially better, with Image-GS [69] achieving 17.2 dB and GI [74] and Instant-GI [75] reaching 22.6 and 23.5 dB, reflecting the advantage of explicit representations for high-frequency hologram content. Our method reports 30.7 dB PSNR, 0.86 SSIM, and 0.33 LPIPS. As shown in Fig 6, prior Gaussian-based methods exhibit structural distortions and blurred details, while learned approaches introduce severe artifacts and structural loss; in contrast, our method preserves sharper reconstructions across near- and far-focus planes.

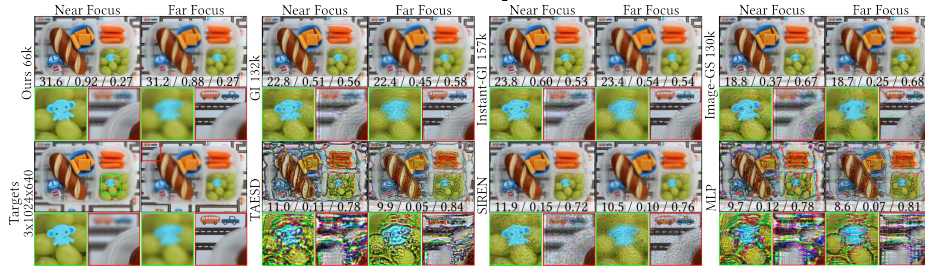


Fig. 6: Qualitative comparison of simulated reconstructions at near and far focal planes. Our method uses a 5:1 parameter ratio, while existing Gaussian-based approaches [69, 74, 75] use equal primitive counts for the two real components of the complex field. *GI* denotes GaussianImage [74]; *Instant-GI* uses network-predicted initialization with variable primitive counts (Source Image: [4, 35]).

Efficiency and Memory Usage. As shown in Table. 1, our method uses 0.8M parameters, fewer than other Gaussian-based approaches (2.4–2.8M). This is due to our compact complex-valued 2D Gaussian definition that avoids the paired real/imaginary parameterization. In terms of memory, our method requires 2.2G VRAM, which is moderate due to the explicit incorporation of light propagation; Gaussian-based methods that omit propagation use less memory but at the cost of fidelity. Training time remains comparable to Gaussian-based methods and is considerably shorter than learned approaches. These results suggest that our representation offers a reasonable trade-off between reconstruction quality, parameter efficiency, and computational cost for hologram modeling.

5.3 Comparison With CGH Methods

Table. 2 shows the quantitative results of our method alongside existing CGH methods for both Smooth and Random POH. Additionally, in Fig 7, we show the qualitative results in an overview and zoomed-in details. Both Random and Smooth POH extracted from our complex-valued 2D Gaussians achieve sharp focus edges and high perceptual quality.

Smooth POH. Our method achieves comparable reconstruction quality to learned CGH methods (NH3D [17], Tensor V2 [55], and U-Net [72]), with PSNR of 29.0 dB, while requiring fewer parameters (0.8 M), lower VRAM (2.4 G), and shorter training time (1.4 min). Since GWS is a novel view synthesis method,

Table 2: Quantitative comparison across our method, learned CGH and optimization methods. (* Naive Opt refers to Random POH baseline used in [10, 34, 44, 52]). Our Smooth POH time covers complex optimization only (POH conversion adds <1 s); Random POH time includes the conversion.

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	VRAM	Params	Time (min)	Render (ms)
Smooth POH							
NH3D [17]	28.3	0.92	0.31	7.1 G	3.9 M	90	31
TensorV2 [55]	27.1	0.94	0.29	8.7 G	0.1 M	80	48
GWS [15]	28.2	0.76	0.41	2.5 G	1.1 M	5.1	6840
U-Net [72]	27.2	0.91	0.35	6.3 G	2.2 M	100	18
Multi-color [30]	27.9	0.74	0.40	3.2 G	4.0 M	5.3	-
Ours	29.0	0.81	0.38	2.4 G	0.8 M	1.4	2.13
Random POH							
Naive Opt*	19.8	0.33	0.60	2.9 G	2.0 M	2.9	-
Wirtinger [9]	25.3	0.47	0.48	3.5 G	2.0 M	2.8	-
Multi-color [30]	20.3	0.35	0.64	3.1 G	4.0 M	3.0	-
Ours	29.4	0.81	0.34	3.4 G	2.0 M	3.8	-

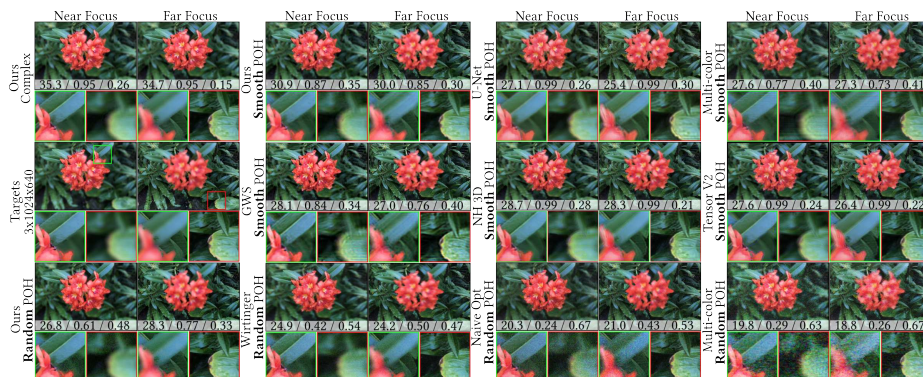


Fig. 7: Comparison of simulated reconstructions at near and far focus using different optimization and learned CGH methods. *Multi-color* refers to [30]; *U-Net* [48] refers to typical learned CGH networks widely used in [2, 13, 17, 25, 37, 44, 72] (Source Image: [41]).

we extend its development to support single-plane POH rendering. Compared with GWS [15], our method achieves higher fidelity and is substantially faster, reducing optimization time from 5.1 to 1.4 minutes (3.6 \times) and render latency from 6840 to 2.13 ms (\sim 3200 \times faster).

Random POH. Existing baselines operate directly on per-pixel phase and yield limited fidelity, with Naive Opt and Multi-color [30] remaining below 21 dB. Wirtinger Holography [9] reaches 25.3 dB PSNR; our method reports 29.4 dB PSNR, 0.81 SSIM, and 0.34 LPIPS, corresponding to gains of +4.1 dB PSNR, +0.34 SSIM, and -0.14 LPIPS over Wirtinger at comparable memory and modest additional training time (3.8 vs. 2.8 min). These results indicate that structural guidance from complex-valued 2D Gaussians and the POH conversion procedure is beneficial for random-valued, per-pixel hologram optimization.




Fig. 8: Simulated reconstructions with varying Gaussian counts. Labels (e.g., *Ours 33k ratio 10:1*) denote the number of complex 2D Gaussians (33k) and the parameter reduction ratio compared to dense per-pixel representation (10:1) (Source Image: [59]).

5.4 Ablation Study

Table 3: Comparison of our method under different parameter reduction ratios.

Parameter Reduction Ratio	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Render (ms)
Dense Per-pixel	32.3	0.893	0.29	4.02
2:1	31.9	0.891	0.30	2.58
5:1	30.7	0.863	0.33	2.13
10:1	29.4	0.835	0.37	1.72

Parameter Space Reduction. Table. 3 evaluates our method under different parameter reduction ratios. We adopt 5:1 as the default used in all main tables, while 10:1 serves as an ablation upper bound rather than a headline claim. Quality degrades gracefully: PSNR remains highly comparable at 2:1 (31.9) and declines only moderately to 30.7 and 29.4 at 5:1 and 10:1, while rendering accelerates accordingly. Fig 8 confirms that reconstructions remain sharp under aggressive reduction. Consistent with Gabor’s theory, our primitive does not surpass per-pixel quality but *retains* it most effectively as parameters decrease, making it an efficient basis for compact hologram representation.

Table 4: Ablation study of POH conversion losses and Gaussian definition choices, evaluated on the test image [3] and [24].

Type	Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Params	Render (ms)
Random POH	w/o guidance	19.1	0.37	0.52	2.0	-
	with guidance (Ours)	30.6	0.88	0.22	2.0	-
Complex	Naive Paired Gaussians	25.5	0.74	0.47	1.2	20.1
	Ours	31.8	0.89	0.31	0.8	2.13

Random POH Conversion Quality. Table. 4 compares our POH conversion to independently optimized Random POH. With structural guidance and a more constrained search space, it improves reconstruction quality by +11.5 dB.

Simulation. Fig 9 reports quantitative results in simulation. Our conversion procedure yields higher-quality reconstructions and, as a side effect, appears to reduce the noise artifacts commonly seen in Random POH methods [19, 30, 34]. The zoomed insets in Fig 9 show reduced noise at both near and far focal planes. We note that this work *focuses on single-view and does not target improved reconstruction fidelity under pupil shift*; the Random POH results are included only to show that our representation can be used to generate Random POH for

commercial phase-only holographic display in practice. Additional experimental results are provided in Suppl Sec. 17.

Experiment. Fig 10 shows experimentally captured results. While the reduction appears less pronounced than in the simulation due to optical imperfections and laser speckle, our method still improves visual quality compared with independently optimized Random POH.




Fig. 9: Simulated reconstructions of our method and independently optimized Random POH at near and far focal planes, using identical training strategies. Insets show corresponding hologram pixels (Source Image: [3, 24, 59]).




Fig. 10: Experimental comparison of reconstructions at near and far focal planes between our method and independently optimized Random POH (Source Image: [3, 41]).

Unified Complex vs. Naive Paired Gaussians. As shown in Table. 4, we compare our unified complex-valued 2D Gaussians against naive paired real-valued 2D Gaussians, which model amplitude and phase using two independent primitives rendered separately with the standard gsplat [68] pipeline. Our method achieves higher reconstruction quality (+6.3 dB) with 33% fewer parameters (12 vs. 9×2 per pair). As shown in Fig 11, the paired formulation introduces implicit spatial misalignment between amplitude and phase Gaussians, which distorts the local complex field and leads to blurred reconstructions. Our unified formulation enforces spatial alignment within each primitive, maintaining sharper details across focal planes. The paired approach is also an order of magnitude slower due to the overhead of two separate rasterization passes. These results show that the unified complex-valued Gaussian, as Gabor’s elementary signal, is a more effective primitive for hologram representation.

6 Discussion

Beyond the core contributions described above, we observe an interesting side effect of our representation in the context of Random POH conversion. As shown in Fig 9 and Fig 10, our conversion procedure appears to suppress noise artifacts




Fig. 11: Simulated reconstructions between complex-valued 2D Gaussians and naive paired real-valued Gaussians. (Source: [24])

in hologram reconstructions. We attribute this behavior to the structured nature of our representation: the complex-valued 2D Gaussians provide a compact, spatially regularized prior to the hologram field, so the subsequent optimization explores a more constrained parameter search space, which yields cleaner reconstructions during Random POH conversion. However, the current conversion procedure remains indirect, relying on a separate Gaussian-guided step. A more elegant approach would be using our representation to directly render Random POH, which we consider as a promising future direction.

Our method focuses on *improving reconstruction quality at the center of the eyebox, whereas existing eyebox-expansion approaches [8, 16] target on complementary objectives rather than serving as direct competitors*. In Suppl Sec. 13, we provide an eyebox-shifting results and spectrum analysis for our random POH. Additionally, the most pronounced performance gains are observed in pure simulation, while experimentally captured results show smaller relative margins. We attribute this gap to several compounding factors: laser speckle introduces multiplicative noise absent in simulation; 8-bit SLM phase quantization and its nonlinear response perturb the encoded field; and optical misalignment, aberrations, and sensor noise further degrade the capture. As these factors affect all methods alike, they compress the relative quality margins rather than change the ranking, which also partly explains the reduced noise-suppression benefit observed on hardware. Therefore, another promising direction for future research is to extend our method to feedforward, real-time video cases and integrate it with multiplexing and camera-in-the-loop techniques to further improve the perceptual quality of the reconstructions [10, 11, 16, 42].

7 Conclusion

In this paper, we extend the complex-valued 3D Gaussian formulation introduced for holographic radiance fields [71] to the hologram optimization domain. We propose a novel representation based on complex-valued 2D Gaussian primitives, each of which acts as Gabor’s elementary signal [21] attaining the minimum space–frequency uncertainty, supported by a differentiable rasterizer and GPU-optimized light propagation. Experiments show that our method enables up to 10:1 parameter reduction, 50% faster optimization, and 30% lower memory consumption, with reconstruction fidelity comparable to existing methods.

Acknowledgements

The authors thank Ye Mao and Dr. Suyeon Choi for the valuable suggestions in the early stage of this work.

References

1. Mango, grapefruit, pomegranate, tropical fruit. Openverse
2. Akşit, K., Itoh, Y.: Holobeam: Paper-thin near-eye displays. In: 2023 IEEE Conference Virtual Reality and 3D User Interfaces (VR). pp. 581–591. IEEE (2023)
3. Alphas: Chinese new year 2016 in london. Openverse (2016)
4. anotherlunch.com: Soft pretzel, carrots, grapes, cheddar cheese - easylunchboxes & toluca bento. Openverse (2011)
5. Appel, M.: Siberian tiger. Openverse (2015)
6. Blanche, P.A.: Holography, and the future of 3d display. *Light: Advanced Manufacturing* **2**(4), 446–459 (2021). <https://doi.org/https://doi.org/10.37188/lam.2021.028>, <https://www.light-am.com/en/article/doi/10.37188/lam.2021.028>
7. Bohan, O.B.: Tiny autoencoder for stable diffusion. Retrieved May **22**, 2024 (2023)
8. Chakravarthula, P., Baek, S.H., Schiffrers, F., Tseng, E., Kuo, G., Maimone, A., Matsuda, N., Cossairt, O., Lanman, D., Heide, F.: Pupil-aware holography. *ACM Trans. Graph.* **41**(6) (Nov 2022). <https://doi.org/10.1145/3550454.3555508>, <https://doi.org/10.1145/3550454.3555508>
9. Chakravarthula, P., Peng, Y., Kollin, J., Fuchs, H., Heide, F.: Wirtinger holography for near-eye displays. *ACM Trans. Graph.* **38**(6) (Nov 2019). <https://doi.org/10.1145/3355089.3356539>, <https://doi.org/10.1145/3355089.3356539>
10. Chakravarthula, P., Tseng, E., Srivastava, T., Fuchs, H., Heide, F.: Learned hardware-in-the-loop phase retrieval for holographic near-eye displays. *ACM Trans. Graph.* **39**(6) (Nov 2020). <https://doi.org/10.1145/3414685.3417846>, <https://doi.org/10.1145/3414685.3417846>
11. Chao, B., Gopakumar, M., Choi, S., Kim, J., Shi, L., Wetzstein, G.: Large Étendue 3d holographic display with content-adaptive dynamic fourier modulation. In: SIGGRAPH Asia 2024 Conference Papers. SA '24, Association for Computing Machinery, New York, NY, USA (2024). <https://doi.org/10.1145/3680528.3687600>, <https://doi.org/10.1145/3680528.3687600>
12. Chao, B., Yang, J., Choi, S., Gopakumar, M., Koiso, R., Wetzstein, G.: Random-phase wave splatting of translucent primitives for computer-generated holography. arXiv preprint arXiv:2508.17480 (2025)
13. Chen, K., Wen, A., Zhang, Y., Chakravarthula, P., Sun, Q.: View synthesis for 3d computer-generated holograms using deep neural fields. *Optics Express* **33**(9), 19399–19408 (2025)
14. Choi, S., Chao, B., Yang, J., Gopakumar, M., Wetzstein, G.: Gaussian wave splatting for computer-generated holography. *ACM Trans. Graph.* **44**(4) (Jul 2025). <https://doi.org/10.1145/3731163>, <https://doi.org/10.1145/3731163>
15. Choi, S., Chao, B., Yang, J., Gopakumar, M., Wetzstein, G.: Gaussian wave splatting for computer-generated holography. arXiv preprint arXiv:2505.06582 (2025)
16. Choi, S., Gopakumar, M., Peng, Y., Kim, J., O’Toole, M., Wetzstein, G.: 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. pp. 1–9 (2022)

17. Choi, S., Gopakumar, M., Peng, Y., Kim, J., Wetzstein, G.: Neural 3d holography: learning accurate wave propagation models for 3d holographic virtual and augmented reality displays. *ACM Trans. Graph.* **40**(6) (Dec 2021). <https://doi.org/10.1145/3478513.3480542>, <https://doi.org/10.1145/3478513.3480542>
18. Chu, V., Pueyo-Ciudad, O., Tseng, E., Schiffers, F., Kuo, G., Matsuda, N., Redo-Sanchez, A., Douglas, L., Cossairt, O., Heide, F.: Artifact-resilient real-time holography. *ACM Trans. Graph.* **44**(6) (Dec 2025). <https://doi.org/10.1145/3763361>, <https://doi.org/10.1145/3763361>
19. Chu, V., Pueyo-Ciudad, O., Tseng, E., Schiffers, F., Kuo, G., Matsuda, N., Redo-Sanchez, A., Lanman, D., Cossairt, O., Heide, F.: Artifact-resilient real-time holography. *ACM Transactions on Graphics* **44**(6) (Dec 2025). <https://doi.org/10.1145/3763361>, publisher Copyright: © 2025 Copyright held by the owner/author(s).
20. Fan, X., Zhan, Y., Mazumdar, A., Akşit, K.: Compressing Double-Phase Holograms using 2D Gaussians. In: Gerrits, T., Teschner, M. (eds.) *Eurographics 2026 - Posters*. The Eurographics Association (2026). <https://doi.org/10.2312/egp.20261010>
21. Gabor, D.: Theory of communication. part 1: The analysis of information. *Journal of the Institution of Electrical Engineers-part III: radio and communication engineering* **93**(26), 429–441 (1946)
22. Gao, X., Li, X., Zhuang, Y., Zhang, Q., Hu, W., Zhang, C., Yao, Y., Shan, Y., Quan, L.: Mani-gs: Gaussian splatting manipulation with triangular mesh. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 21392–21402 (2025)
23. Goodman, J.W.: *Introduction to Fourier optics*. Roberts and Company publishers (2005)
24. HONGSESISHEN: A space station based in a universe inspired by the space opera dune by frank herbert. it circles a planet and a big spaceship of the spice guild is seen in the background. *Openverse* (2020)
25. Hossein Eybposh, M., Caira, N.W., Atisa, M., Chakravarthula, P., Pégard, N.C.: Deepcgh: 3d computer-generated holography using deep learning. *Optics Express* **28**(18), 26636–26650 (2020)
26. Hsueh, C.K., Sawchuk, A.A.: Computer-generated double-phase holograms. *Applied optics* **17**(24), 3874–3883 (1978)
27. Huang, B., Yu, Z., Chen, A., Geiger, A., Gao, S.: 2d gaussian splatting for geometrically accurate radiance fields. In: *ACM SIGGRAPH 2024 conference papers*. pp. 1–11 (2024)
28. Hugo-90: 1950s diamond t. *Openverse* (2012)
29. Kavakh, K., Itoh, Y., Urey, H., Akşit, K.: Realistic defocus blur for multiplane computer-generated holography. In: *2023 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*. pp. 418–426. IEEE (2023)
30. Kavakh, K., Shi, L., Urey, H., Matusik, W., Akşit, K.: Multi-color holograms improve brightness in holographic displays. In: *ACM SIGGRAPH ASIA 2023 Conference Proceedings*. pp. –. ACM, Sydney, NSW, Australia (2023). <https://doi.org/https://doi.org/10.1145/3610548.3618135>
31. kennejima: Guy fieris vegas kitchen & bar. *Openverse* (2014)
32. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.* **42**(4), 139–1 (2023)
33. Kim, D., Nam, S.W., Choi, S., Seo, J.M., Wetzstein, G., Jeong, Y.: Holographic parallax improves 3d perceptual realism. *ACM Transactions on Graphics (TOG)* **43**(4), 1–13 (2024)

34. Kuo, G., Schiffers, F., Lanman, D., Cossairt, O., Matsuda, N.: Multisource holography. *ACM Transactions on Graphics (Tog)* **42**(6), 1–14 (2023)
35. kuujinbo: Manette, wa graffiti wall. Openverse (2012)
36. Lindell, D.B., Van Veen, D., Park, J.J., Wetzstein, G.: Bacon: Band-limited coordinate networks for multiscale scene representation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 16252–16262 (2022)
37. Liu, K., Wu, J., He, Z., Cao, L.: 4k-dmdnet: diffraction model-driven network for 4k computer-generated holography. *Opto-Electronic Advances* **6**(5), 220135–1 (2023)
38. Loshchilov, I., Hutter, F.: SGDR: stochastic gradient descent with warm restarts. In: *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24–26, 2017, Conference Track Proceedings*. OpenReview.net (2017), <https://openreview.net/forum?id=Skq89Scxx>
39. Mallick, S.S., Goel, R., Kerbl, B., Steinberger, M., Carrasco, F.V., De La Torre, F.: Taming 3dgs: High-quality radiance fields with limited resources. In: *SIGGRAPH Asia 2024 Conference Papers*. pp. 1–11 (2024)
40. Matsushima, K., Shimobaba, T.: Band-limited angular spectrum method for numerical simulation of free-space propagation in far and near fields. *Optics express* **17**(22), 19662–19673 (2009)
41. Mildenhall, B., Srinivasan, P.P., Ortiz-Cayon, R., Kalantari, N.K., Ramamoorthi, R., Ng, R., Kar, A.: Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)* (2019), <https://b mild.github.io/llff/>
42. Nam, S.W., Kim, Y., Kim, D., Jeong, Y.: Depolarized holography with polarization-multiplexing metasurface. *ACM Transactions on Graphics (TOG)* **42**(6), 1–16 (2023)
43. Paszke, A., et al.: Pytorch: An imperative style, high-performance deep learning library. In: *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. Curran Associates, Inc. (2019), <http://papers.neurips.cc>
44. Peng, Y., Choi, S., Padmanaban, N., Wetzstein, G.: Neural holography with camera-in-the-loop training. *ACM Transactions on Graphics (TOG)* **39**(6), 1–14 (2020)
45. Peng, Z., Zhan, Y., Spjut, J., Akşit, K.: Assessing learned models for phase-only hologram compression. In: *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Posters. SIGGRAPH Posters '25*, Association for Computing Machinery, New York, NY, USA (2025). <https://doi.org/10.1145/3721250.3742993>, <https://doi.org/10.1145/3721250.3742993>
46. Preechakul, K., Chatthee, N., Wizadwongsa, S., Suwajanakorn, S.: Diffusion autoencoders: Toward a meaningful and decodable representation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10619–10629 (2022)
47. Ranftl, R., Lasinger, K., Hafner, D., Schindler, K., Koltun, V.: Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(3) (2022)
48. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5–9, 2015, proceedings, part III 18*. pp. 234–241. Springer (2015)
49. Rumelhart, D.E., Hinton, G.E., Williams, R.J.: Learning representations by back-propagating errors. *nature* **323**(6088), 533–536 (1986)

50. Saragadam, V., LeJeune, D., Tan, J., Balakrishnan, G., Veeraraghavan, A., Baraniuk, R.G.: Wire: Wavelet implicit neural representations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18507–18516 (2023)
51. Schiffers, F., Chakravarthula, P., Matsuda, N., Kuo, G., Tseng, E., Lanman, D., Heide, F., Cossairt, O.: Stochastic light field holography. IEEE International Conference on Computational Photography (ICCP) (2023)
52. Schiffers, F.A., Kuo, G., Matsuda, N., Lanman, D., Cossairt, O.: Holochrome: Polychromatic illumination for speckle reduction in holographic near-eye displays. ACM Trans. Graph. **44**(3) (May 2025). <https://doi.org/10.1145/3732935>, <https://doi.org/10.1145/3732935>
53. Shi, L., Huang, F.C., Lopes, W., Matusik, W., Luebke, D.: Near-eye light field holographic rendering with spherical waves for wide field of view interactive 3d computer graphics. ACM Trans. Graph. **36**(6) (Nov 2017). <https://doi.org/10.1145/3130800.3130832>, <https://doi.org/10.1145/3130800.3130832>
54. Shi, L., Li, B., Kim, C., Kellnhofer, P., Matusik, W.: Towards real-time photorealistic 3d holography with deep neural networks. Nature **591**(7849), 234–239 (2021)
55. Shi, L., Li, B., Matusik, W.: End-to-end learning of 3d phase-only holograms for holographic display. Light: Science & Applications **11**(1), 247 (2022)
56. Shi, L., Webb, R., Xiao, L., Kim, C., Jang, C.: Neural compression for hologram images and videos. Opt. Lett. **47**(22), 6013 (Nov 2022)
57. Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G.: Implicit neural representations with periodic activation functions. Advances in neural information processing systems **33**, 7462–7473 (2020)
58. Song, Y., Dhariwal, P., Chen, M., Sutskever, I.: Consistency models. ICML (2023)
59. Sutherland, B.: Wasabi rainbow sushi set. Openverse (2010)
60. over 10 million views Thanks !!, S.W.: A pair of ecuadorian amazon red-loreed parrots. Openverse (2009)
61. Timofte, R., Agustsson, E., Van Gool, L., Yang, M.H., Zhang, L., Lim, B., et al.: Ntire 2017 challenge on single image super-resolution: Methods and results. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops (July 2017)
62. Wang, Y., Yang, W., Chen, X., Wang, Y., Guo, L., Chau, L.P., Liu, Z., Qiao, Y., Kot, A.C., Wen, B.: Sinsr: diffusion-based image super-resolution in a single step. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 25796–25805 (2024)
63. Wang, Y., Chakravarthula, P., Sun, Q., Chen, B.: Joint neural phase retrieval and compression for energy-and computation-efficient holography on the edge. ACM Transactions on Graphics **41**(4) (2022)
64. Wu, G., Yi, T., Fang, J., Xie, L., Zhang, X., Wei, W., Liu, W., Tian, Q., Wang, X.: 4d gaussian splatting for real-time dynamic scene rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20310–20320 (2024)
65. www.metaphoricalplatypus.com: Peas and strawberries. Openverse (2013)
66. Xie, X., Zhou, P., Li, H., Lin, Z., Yan, S.: Adan: Adaptive nesterov momentum algorithm for faster optimizing deep models. IEEE Transactions on Pattern Analysis and Machine Intelligence (2024)
67. Yang, L., Kang, B., Huang, Z., Zhao, Z., Xu, X., Feng, J., Zhao, H.: Depth anything v2. arXiv:2406.09414 (2024)

68. Ye, V., Li, R., Kerr, J., Turkulainen, M., Yi, B., Pan, Z., Seiskari, O., Ye, J., Hu, J., Tancik, M., et al.: gsplat: An open-source library for gaussian splatting. *Journal of Machine Learning Research* **26**(34), 1–17 (2025)
69. Zeng, Z., Wang, Y., Yang, C., Guan, T., Ju, L.: Instant gaussianimage: A generalizable and self-adaptive image representation via 2d gaussian splatting. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (October 2025)
70. Zhan, Y., Gao, X., Quan, L., Akşit, K.: Complex-valued 2d gaussian representation for computer-generated holography. https://github.com/complight/Complex-Valued_2D_Gaussian_Representation (2026), source code
71. Zhan, Y., Shin, D.H., Baek, S.H., Akşit, K.: Complex-valued holographic radiance fields. *arXiv preprint arXiv:2506.08350* (2025)
72. Zhan, Y., Sun, Q., Shi, L., Matusik, W., Akşit, K.: Configurable holography: Towards display and scene adaptation. *arXiv preprint arXiv:2405.01558* (2024)
73. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18–22, 2018*. pp. 586–595. Computer Vision Foundation / IEEE Computer Society (2018). <https://doi.org/10.1109/CVPR.2018.00068>, http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_The_Unreasonable_Effectiveness_CVPR_2018_paper.html
74. Zhang, X., Ge, X., Xu, T., He, D., Wang, Y., Qin, H., Lu, G., Geng, J., Zhang, J.: Gaussianimage: 1000 fps image representation and compression by 2d gaussian splatting. In: *European Conference on Computer Vision*. pp. 327–345. Springer (2024)
75. Zhang, Y., Li, B., Kuznetsov, A., Jindal, A., Diolatzis, S., Chen, K., Sochenov, A., Kaplanyan, A., Sun, Q.: Image-gs: Content-adaptive image representation via 2d gaussians. In: *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Papers*. pp. 1–11 (2025)
76. Zheng, C., Zhan, Y., Shi, L., Cakmakci, O., Akşit, K.: Focal surface holographic light transport using learned spatially adaptive convolutions. In: *SIGGRAPH Asia 2024 Technical Communications*. SA '24, Association for Computing Machinery, New York, NY, USA (2024). <https://doi.org/10.1145/3681758.3697989>, <https://doi.org/10.1145/3681758.3697989>

Supplementary Material: Complex-Valued 2D Gaussian Representation for Computer-Generated Holography

8 Hardware Prototype

Figure 12 shows the holographic display prototype used in our experiments. The optical path begins with a laser source (LASOS MCS4) integrating three individual laser lines. Light emitted from a single-mode fibre is collimated by a Thorlabs LA1708-A plano-convex lens with a 200 mm focal length. The linearly polarized, collimated beam is then directed by a beamsplitter (Thorlabs BP245B1) onto a phase-only SLM, the Jasper JD7714 (2400×4094 , $3.74 \mu\text{m}$ pixel pitch). The modulated beam passes through a lens pair (Thorlabs LA1908-A and LB1056-A) with focal lengths of 500 mm and 250 mm, respectively. A pinhole aperture (Thorlabs SM1D12) is placed at the focal plane of the lenses for spatial filtering. Finally, the holographic reconstructions are recorded by a lensless image sensor (Point Grey GS3-U3-23S6M-C, USB 3.0) mounted on an X-stage (Thorlabs PT1/M) with a travel range of 0–25 mm and a positioning precision of 0.01 mm.




Fig. 12: Holographic display prototype (Jasper JD7714) used to evaluate holograms generated by our model.

9 Differentiable 2D Complex-Valued Gaussian Rendering

This section provides detailed mathematical formulations and gradient derivations for our 2D Complex-Valued Gaussian Rasterizer.

9.1 Notation

- $s_x, s_y \in \mathbb{R}^+$ - Activated scaling factors for Gaussian in x and y directions
- $\tilde{s}_x, \tilde{s}_y \in \mathbb{R}$ - Pre-activation scale parameters
- $\theta \in \mathbb{R}$ - Rotation angle of the Gaussian ellipse
- $\mathbf{x} = (x_0, x_1) \in \mathbb{R}^2$ - Activated 2D mean position
- $\tilde{\mathbf{x}} = (\tilde{x}_0, \tilde{x}_1) \in \mathbb{R}^2$ - Pre-activation mean parameters
- $\Sigma \in \mathbb{R}^{2 \times 2}$ - 2D covariance matrix
- $\Sigma^{-1} \in \mathbb{R}^{2 \times 2}$ - Inverse 2D covariance matrix
- Σ_{ij}^{-1} - Elements of inverse covariance where $i, j \in \{0, 1\}$
- $d_x = x - x_0, d_y = y - x_1$ - Distance from pixel to Gaussian center
- $\mathbf{c}_n \in \mathbb{R}^C$ - Color/amplitude values for Gaussian n across C channels
- $\varphi_n \in \mathbb{R}^C$ - Phase values for Gaussian n across C channels
- $\alpha_n \in [0, 1]$ - Activated opacity value for Gaussian n
- $\tilde{\alpha}_n \in \mathbb{R}$ - Pre-activation opacity parameter for Gaussian n
- power - Gaussian exponent term (negative half Mahalanobis distance)
- W, H - Image width and height
- $\epsilon_s = 0.1$ - Scale regularization constant
- $\epsilon_c = 0.1$ - Covariance regularization constant
- $\epsilon_d = 10^{-10}$ - Determinant clamping threshold

9.2 Optimization Algorithm

Our optimization algorithm adapts 2D Gaussian primitives for hologram generation, as summarized in Algorithm 1.

9.3 Tile-Based Rasterizer

Our tile-based rasterizer efficiently computes complex fields across the hologram plane, as detailed in Algorithm 2.

Key features include: (1) parallel tile processing with 16×16 blocks, (2) shared VRAM for batch Gaussian loading, (3) early termination when Gaussian contribution is negligible (e.g., $\alpha_{\text{eff}} < 1/255$), (4) channel-wise complex accumulation, and (5) improved numerical stability via power clamping.

9.4 Forward Pass

Parameter Activation Functions Mean Position Activation (Tanh-based):

$$\mathbf{x} = \left(\frac{\tanh(\tilde{x}_x) + 1}{2} \cdot W, \frac{\tanh(\tilde{x}_y) + 1}{2} \cdot H \right) \quad (10)$$

Algorithm 1 2D Gaussian Hologram Optimization

Require: W, H : hologram resolution
Require: L : number of depth planes

- 1: $M \leftarrow \text{InitPositions}(N)$
- 2: $\mathbf{S}, \mathbf{C}, \alpha \leftarrow \text{InitAttributes}()$
- 3: $\varphi \leftarrow \text{InitPhase}()$
- 4: $\theta \leftarrow \text{InitRotation}()$
- 5: $i \leftarrow 0$
- 6: **while** not converged **do**
- 7: $I_{\text{target}}, D \leftarrow \text{GetTarget}()$
- 8: $U_{\text{complex}} \leftarrow \text{ComplexRasterize2D}(M, \mathbf{S}, \theta, \mathbf{C}, \varphi, \alpha)$
- 9: $P \leftarrow \text{ZeroPad}(U_{\text{complex}})$
- 10: $\{I_{\text{recon}, l}\}_{l=1}^L \leftarrow \text{MultiPlanePropagate}(P)$
- 11: $\mathcal{L} \leftarrow \text{MultiPlaneLoss}(\{I_{\text{recon}, l}\}, I_{\text{target}}, D)$
- 12: $M, \mathbf{S}, \theta, \mathbf{C}, \varphi, \alpha \leftarrow \text{Adan}(\nabla \mathcal{L})$
- 13: $i \leftarrow i + 1$
- 14: **end while**

Scale Activation (Exponential):

$$s_x = \exp(\tilde{s}_x) + \epsilon_s, \quad s_y = \exp(\tilde{s}_y) + \epsilon_s \quad (11)$$

Opacity Activation (Sigmoid):

$$\alpha_n = \sigma(\tilde{\alpha}_n) = \frac{1}{1 + \exp(-\tilde{\alpha}_n)} \quad (12)$$

2D Covariance Matrix Computation The 2D covariance matrix is:

$$\Sigma = R \cdot S^2 \cdot R^T + \epsilon_c \cdot \mathbf{I} \quad (13)$$

where $R = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}$ and $S^2 = \begin{pmatrix} s_x^2 & 0 \\ 0 & s_y^2 \end{pmatrix}$. Expanding:

$$\begin{aligned} \Sigma_{00} &= s_x^2 \cos^2 \theta + s_y^2 \sin^2 \theta + \epsilon_c \\ \Sigma_{01} &= (s_x^2 - s_y^2) \cos \theta \sin \theta \\ \Sigma_{11} &= s_x^2 \sin^2 \theta + s_y^2 \cos^2 \theta + \epsilon_c \end{aligned} \quad (14)$$

2D Covariance Matrix Inversion For 2×2 matrix inversion:

$$\det(\Sigma) = \Sigma_{00}\Sigma_{11} - \Sigma_{01}^2 \quad (15)$$

$$\Sigma^{-1} = \frac{1}{\max(\det(\Sigma), \epsilon_d)} \begin{pmatrix} \Sigma_{11} & -\Sigma_{01} \\ -\Sigma_{01} & \Sigma_{00} \end{pmatrix} \quad (16)$$

Algorithm 2 2D Complex-Valued Tile-Based Rasterization

Require: W, H : hologram dimensions
Require: M, \mathbf{S}, θ : Gaussian positions, scales, rotations
Require: $\mathbf{C}, \alpha, \varphi$: Amplitudes, opacities, phases
Require: C : number of color channels

- 1: **function** COMPLEXRASTERIZE2D($W, H, M, \mathbf{S}, \theta,$
 $\mathbf{C}, \alpha, \varphi$)
- 2: $\Sigma \leftarrow \text{Compute2DCovariance}(\mathbf{S}, \theta)$
- 3: $\Sigma^{-1}, r \leftarrow \text{Invert2DCovariance}(\Sigma)$
- 4: $T \leftarrow \text{CreateTiles}(W, H)$
- 5: $\mathcal{I}, \mathcal{K} \leftarrow \text{DuplicateWithKeys}(M, r, T)$
- 6: $\mathcal{K}_s, \mathcal{I}_s \leftarrow \text{SortByKeys}(\mathcal{I}, \mathcal{K})$
- 7: $\mathcal{R} \leftarrow \text{IdentifyTileRanges}(T, \mathcal{K}_s)$
- 8: $U_{\text{real}}, U_{\text{imag}} \leftarrow \text{InitCanvas}(C, W, H)$
- 9: **for all** Tiles t **in** T **parallel do**
- 10: **for all** Pixels pix **in** t **parallel do**
- 11: $\text{real}_{\text{acc}}[C], \text{imag}_{\text{acc}}[C] \leftarrow 0$
- 12: $\text{range} \leftarrow \text{GetTileRange}(\mathcal{R}, t)$
- 13: **for** g **in** range **do**
- 14: $d_x \leftarrow pix_x - x_{g,0}, d_y \leftarrow pix_y - x_{g,1}$
- 15: $\text{power} \leftarrow -0.5 \cdot (d_x^2 \Sigma_{00}^{-1} +$
 $2d_x d_y \Sigma_{01}^{-1} + d_y^2 \Sigma_{11}^{-1})$
- 16: $G \leftarrow \exp(\max(\text{power}, -50))$
- 17: $\alpha_{\text{eff}} \leftarrow \min(0.99, \alpha_g \cdot G)$
- 18: **if** $\alpha_{\text{eff}} < 1/255$ **then**
- 19: **continue**
- 20: **end if**
- 21: **for** $c \leftarrow 0$ **to** $C - 1$ **do**
- 22: $\text{scale} \leftarrow \mathbf{c}_{g,c} \cdot \alpha_{\text{eff}}$
- 23: $\cos_\varphi, \sin_\varphi \leftarrow \cos(\varphi_{g,c}),$
 $\sin(\varphi_{g,c})$
- 24: $\text{real}_{\text{acc}}[c] += \text{scale} \cdot \cos_\varphi$
- 25: $\text{imag}_{\text{acc}}[c] += \text{scale} \cdot \sin_\varphi$
- 26: **end for**
- 27: **end for**
- 28: **for** $c \leftarrow 0$ **to** $C - 1$ **do**
- 29: $U_{\text{real}}[c, pix] \leftarrow \text{real}_{\text{acc}}[c]$
- 30: $U_{\text{imag}}[c, pix] \leftarrow \text{imag}_{\text{acc}}[c]$
- 31: **end for**
- 32: **end for**
- 33: **end for**
- 34: **return** $U_{\text{real}} + j \cdot U_{\text{imag}}$
- 35: **end function**

Stored as $[\text{inv}_{00}, \text{inv}_{01}, \text{inv}_{11}]$:

$$\begin{aligned}
 \text{inv}_{00} &= \Sigma_{11} / \det_{\text{safe}} \\
 \text{inv}_{01} &= -\Sigma_{01} / \det_{\text{safe}} \\
 \text{inv}_{11} &= \Sigma_{00} / \det_{\text{safe}}
 \end{aligned} \tag{17}$$

Gaussian Evaluation For pixel (x, y) and Gaussian n :

$$\begin{aligned} \text{mahal_dist} &= d_x^2 \cdot \text{inv}_{00} + 2d_x d_y \cdot \text{inv}_{01} \\ &\quad + d_y^2 \cdot \text{inv}_{11} \end{aligned} \quad (18)$$

where $d_x = x - x_x$, $d_y = y - x_y$. With numerical stability:

$$\text{power} = \max(-0.5 \cdot \text{mahal_dist}, -50.0) \quad (19)$$

$$\mathcal{G}_n(x, y) = \exp(\text{power}) \quad (20)$$

Complex Field Rendering For each channel c :

$$\text{real}_c(x, y) = \sum_n \mathbf{c}_{n,c} \cdot \alpha_n \cdot \mathcal{G}_n(x, y) \cdot \cos(\varphi_{n,c}) \quad (21)$$

$$\text{imag}_c(x, y) = \sum_n \mathbf{c}_{n,c} \cdot \alpha_n \cdot \mathcal{G}_n(x, y) \cdot \sin(\varphi_{n,c})$$

$$U_c(x, y) = \text{real}_c(x, y) + j \cdot \text{imag}_c(x, y) \quad (22)$$

9.5 Backward Pass

Gradient Flow Overview For parameter $\tilde{\theta}$:

$$\frac{\partial \mathcal{L}}{\partial \tilde{\theta}} = \frac{\partial \mathcal{L}}{\partial U} \cdot \frac{\partial U}{\partial \theta} \cdot \frac{\partial \theta}{\partial \tilde{\theta}} \quad (23)$$

Real and imaginary components are accumulated separately but remain coupled through shared Gaussian parameters.

Detailed Gradient Derivation

1. Gradient for Color/Amplitude \mathbf{c}_n

For each channel c and Gaussian n :

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \mathbf{c}_{n,c}} &= \sum_{x,y} \alpha_n \cdot \mathcal{G}_n(x, y) \cdot \left(\cos(\varphi_{n,c}) \cdot \frac{\partial \mathcal{L}}{\partial \text{real}_c(x, y)} \right. \\ &\quad \left. + \sin(\varphi_{n,c}) \cdot \frac{\partial \mathcal{L}}{\partial \text{imag}_c(x, y)} \right) \end{aligned} \quad (24)$$

where $\mathcal{G}_n(x, y) = \exp(\text{power})$ with clamping applied.

2. Gradient for Phase φ_n

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \varphi_{n,c}} &= \sum_{x,y} \mathbf{c}_{n,c} \cdot \alpha_n \cdot \mathcal{G}_n(x, y) \cdot \left(-\sin(\varphi_{n,c}) \right. \\ &\quad \left. \cdot \frac{\partial \mathcal{L}}{\partial \text{real}_c(x, y)} + \cos(\varphi_{n,c}) \cdot \frac{\partial \mathcal{L}}{\partial \text{imag}_c(x, y)} \right) \end{aligned} \quad (25)$$

3. Gradient for Pre-activation Opacity $\tilde{\alpha}_n$

First compute gradient w.r.t. activated opacity:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \alpha_n} &= \sum_{c,x,y} \mathbf{c}_{n,c} \cdot \mathcal{G}_n(x,y) \cdot \left(\cos(\varphi_{n,c}) \right. \\ &\quad \left. \cdot \frac{\partial \mathcal{L}}{\partial \text{real}_c(x,y)} + \sin(\varphi_{n,c}) \cdot \frac{\partial \mathcal{L}}{\partial \text{imag}_c(x,y)} \right) \end{aligned} \quad (26)$$

Then apply sigmoid derivative:

$$\frac{\partial \mathcal{L}}{\partial \tilde{\alpha}_n} = \frac{\partial \mathcal{L}}{\partial \alpha_n} \cdot \alpha_n \cdot (1 - \alpha_n) \quad (27)$$

4. Gradient for Pre-activation Mean $\tilde{\mathbf{x}}_n$

The gradient flows through: $\tilde{\mathbf{x}} \rightarrow \mathbf{x} \rightarrow d_x, d_y \rightarrow \text{mahal_dist} \rightarrow \text{power} \rightarrow \mathcal{G}_n$.
For the Mahalanobis distance:

$$\begin{aligned} \frac{\partial \text{mahal_dist}}{\partial x_x} &= -2(d_x \cdot \text{inv}_{00} + d_y \cdot \text{inv}_{01}) \\ \frac{\partial \text{mahal_dist}}{\partial x_y} &= -2(d_x \cdot \text{inv}_{01} + d_y \cdot \text{inv}_{11}) \end{aligned} \quad (28)$$

For the clamped power term:

$$\frac{\partial \text{power}}{\partial \text{mahal_dist}} = \begin{cases} -0.5 & \text{if power} > -50 \\ 0 & \text{otherwise} \end{cases} \quad (29)$$

Tanh activation backward:

$$\begin{aligned} \frac{\partial x_x}{\partial \tilde{x}_x} &= \frac{W}{2} \cdot (1 - \tanh^2(\tilde{x}_x)) \\ \frac{\partial x_y}{\partial \tilde{x}_y} &= \frac{H}{2} \cdot (1 - \tanh^2(\tilde{x}_y)) \end{aligned} \quad (30)$$

5. Gradient for Inverse Covariance Elements

The gradient w.r.t. inverse covariance (stored as 3 elements):

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \text{inv}_{00}} &= \sum_{x,y} \frac{\partial \mathcal{L}}{\partial \mathcal{G}_n} \cdot \mathcal{G}_n \cdot \left(-\frac{1}{2} \right) \cdot d_x^2 \\ \frac{\partial \mathcal{L}}{\partial \text{inv}_{01}} &= \sum_{x,y} \frac{\partial \mathcal{L}}{\partial \mathcal{G}_n} \cdot \mathcal{G}_n \cdot (-1) \cdot d_x \cdot d_y \\ \frac{\partial \mathcal{L}}{\partial \text{inv}_{11}} &= \sum_{x,y} \frac{\partial \mathcal{L}}{\partial \mathcal{G}_n} \cdot \mathcal{G}_n \cdot \left(-\frac{1}{2} \right) \cdot d_y^2 \end{aligned} \quad (31)$$

6. Gradient for Pre-activation Scales \tilde{s}_x, \tilde{s}_y

Through $\tilde{s} \rightarrow s \rightarrow \Sigma \rightarrow \Sigma^{-1} \rightarrow$ power:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial s_x} = & 2s_x \left(\cos^2 \theta \cdot \frac{\partial \mathcal{L}}{\partial \Sigma_{00}} + \cos \theta \sin \theta \right. \\ & \left. \cdot \frac{\partial \mathcal{L}}{\partial \Sigma_{01}} + \sin^2 \theta \cdot \frac{\partial \mathcal{L}}{\partial \Sigma_{11}} \right) \end{aligned} \quad (32)$$

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial s_y} = & 2s_y \left(\sin^2 \theta \cdot \frac{\partial \mathcal{L}}{\partial \Sigma_{00}} - \cos \theta \sin \theta \right. \\ & \left. \cdot \frac{\partial \mathcal{L}}{\partial \Sigma_{01}} + \cos^2 \theta \cdot \frac{\partial \mathcal{L}}{\partial \Sigma_{11}} \right) \end{aligned} \quad (33)$$

Exponential backward:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \tilde{s}_x} &= \frac{\partial \mathcal{L}}{\partial s_x} \cdot \exp(\tilde{s}_x) \\ \frac{\partial \mathcal{L}}{\partial \tilde{s}_y} &= \frac{\partial \mathcal{L}}{\partial s_y} \cdot \exp(\tilde{s}_y) \end{aligned} \quad (34)$$

7. Gradient for Rotation θ

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \theta} = & 2(s_y^2 - s_x^2) \cos \theta \sin \theta \cdot \frac{\partial \mathcal{L}}{\partial \Sigma_{00}} \\ & + (s_x^2 - s_y^2)(\cos^2 \theta - \sin^2 \theta) \cdot \frac{\partial \mathcal{L}}{\partial \Sigma_{01}} \\ & + 2(s_x^2 - s_y^2) \cos \theta \sin \theta \cdot \frac{\partial \mathcal{L}}{\partial \Sigma_{11}} \end{aligned} \quad (35)$$

10 Differentiable Light Propagation

This section provides detailed mathematical formulations and gradient derivations for our BLASM method.

10.1 BLASM Algorithm

Algorithm 3 details the BLASM method used for hologram reconstruction.

10.2 Forward Pass

Spatial Frequency Computation For a hologram of size $N_x \times N_y$ with pixel pitch Δx , the spatial frequencies at index (i_x, i_y) are computed as:

$$\begin{aligned} L_x &= N_x \cdot \Delta x, \quad L_y = N_y \cdot \Delta x \\ f_x(i_x) &= \frac{i_x - N_x/2}{L_x} \\ f_y(i_y) &= \frac{i_y - N_y/2}{L_y} \end{aligned} \quad (36)$$

where the zero-frequency component is centered at $(N_x/2, N_y/2)$ following FFT-shift convention.

Algorithm 3 Bandlimited Angular Spectrum Method

Require: $\tilde{U}(f_x, f_y, 0)$: Fourier-domain input field
Require: λ : wavelength, d : propagation distance
Require: Δx : pixel pitch, N_x, N_y : resolution
Require: a : aperture size (optional)

- 1: **function** PROPAGATEFIELD($\tilde{U}, \lambda, d, \Delta x, N_x, N_y, a$)
- 2: $k \leftarrow 2\pi/\lambda$
- 3: $L_x \leftarrow N_x \cdot \Delta x$
- 4: $L_y \leftarrow N_y \cdot \Delta x$
- 5: $\tilde{U}_{\text{out}} \leftarrow \text{InitEmpty}(N_x, N_y)$
- 6: **for all** (i_x, i_y) **in parallel do**
- 7: $f_x \leftarrow (i_x - N_x/2)/L_x$
- 8: $f_y \leftarrow (i_y - N_y/2)/L_y$
- 9: $f_x^{\text{max}} \leftarrow \frac{1}{\lambda\sqrt{(2d/L_x)^2 + 1}}$
- 10: $f_y^{\text{max}} \leftarrow \frac{1}{\lambda\sqrt{(2d/L_y)^2 + 1}}$
- 11: **if** $|f_x| \geq f_x^{\text{max}}$ **or** $|f_y| \geq f_y^{\text{max}}$ **then**
- 12: $\tilde{U}_{\text{out}}[i_x, i_y] \leftarrow 0$
- 13: **continue**
- 14: **end if**
- 15: $k_z^2 \leftarrow k^2 - (2\pi)^2(f_x^2 + f_y^2)$
- 16: $k_z \leftarrow \begin{cases} \sqrt{k_z^2} & \text{if } k_z^2 > 0 \\ 0 & \text{otherwise} \end{cases}$
- 17: $\cos_\phi, \sin_\phi \leftarrow \cos(k_z d), \sin(k_z d)$
- 18: $H \leftarrow \cos_\phi + j \sin_\phi$
- 19: $\tilde{U}_{\text{out}}[i_x, i_y] \leftarrow \tilde{U}[i_x, i_y] \cdot H$
- 20: **end for**
- 21: **if** $a > 0$ **then**
- 22: **for all** (i_x, i_y) **in parallel do**
- 23: $dx \leftarrow i_x - N_x/2 + 0.5$
- 24: $dy \leftarrow i_y - N_y/2 + 0.5$
- 25: **if** $dx^2 + dy^2 \geq a^2$ **then**
- 26: $\tilde{U}_{\text{out}}[i_x, i_y] \leftarrow 0$
- 27: **end if**
- 28: **end for**
- 29: **end if**
- 30: **return** \tilde{U}_{out}
- 31: **end function**

Bandlimit Computation The maximum spatial frequencies that can propagate without aliasing are computed per-thread:

$$\begin{aligned}
 f_x^{\text{max}} &= \frac{1}{\lambda\sqrt{(2d/L_x)^2 + 1}} \\
 f_y^{\text{max}} &= \frac{1}{\lambda\sqrt{(2d/L_y)^2 + 1}}
 \end{aligned} \tag{37}$$

where λ is the wavelength and d is the propagation distance.

Transfer Function Evaluation For spatial frequency (f_x, f_y) , the wave vector component along propagation direction is:

$$k_z^2 = k^2 - (2\pi)^2(f_x^2 + f_y^2) \quad (38)$$

where $k = 2\pi/\lambda$ is the wave number. The longitudinal wave vector is:

$$k_z = \begin{cases} \sqrt{k_z^2} & \text{if } k_z^2 > 0 \\ 0 & \text{otherwise} \end{cases} \quad (39)$$

The transfer function is:

$$H(f_x, f_y, d) = \begin{cases} e^{jk_z d} & \text{if } |f_x| < f_x^{\max}, |f_y| < f_y^{\max} \\ 0 & \text{otherwise} \end{cases} \quad (40)$$

The complex exponential is evaluated using:

$$e^{jk_z d} = \cos(k_z d) + j \sin(k_z d) \quad (41)$$

computed with hardware-accelerated `sincosf` or `sincos` functions.

Field Propagation The propagated field in Fourier domain is:

$$\tilde{U}(f_x, f_y, d) = \tilde{U}(f_x, f_y, 0) \cdot H(f_x, f_y, d) \quad (42)$$

For complex multiplication with input $\tilde{U}_{\text{in}} = \text{real}_{\text{in}} + j \cdot \text{imag}_{\text{in}}$ and transfer function $H = \cos(k_z d) + j \sin(k_z d)$:

$$\begin{aligned} \text{real}_{\text{out}} &= \text{real}_{\text{in}} \cos(k_z d) - \text{imag}_{\text{in}} \sin(k_z d) \\ \text{imag}_{\text{out}} &= \text{real}_{\text{in}} \sin(k_z d) + \text{imag}_{\text{in}} \cos(k_z d) \end{aligned} \quad (43)$$

Aperture Filtering When aperture size $a > 0$, circular filtering is applied in a separate kernel pass:

$$\tilde{U}_{\text{out}}(i_x, i_y, d) = \begin{cases} \tilde{U}_{\text{out}}(i_x, i_y, d) & \text{if } (i_x - o_x)^2 + (i_y - o_y)^2 < a^2 \\ 0 & \text{otherwise} \end{cases} \quad (44)$$

where $(o_x, o_y) = (N_x/2 - 0.5, N_y/2 - 0.5)$ is the centered offset.

10.3 Backward Pass

The backward pass computes gradients with respect to the input Fourier field $\tilde{U}(f_x, f_y, 0)$ given gradients of the output $\partial\mathcal{L}/\partial\tilde{U}(f_x, f_y, d)$.

Complex Conjugate Transfer Function The gradient flows through the conjugate transfer function:

$$\frac{\partial \mathcal{L}}{\partial \tilde{U}(f_x, f_y, 0)} = \frac{\partial \mathcal{L}}{\partial \tilde{U}(f_x, f_y, d)} \cdot H^*(f_x, f_y, d) \quad (45)$$

where $H^*(f_x, f_y, d) = e^{-jk_z d}$ is the conjugate, equivalent to backward propagation:

$$\begin{aligned} H^*(f_x, f_y, d) &= \cos(-k_z d) + j \sin(-k_z d) \\ &= \cos(k_z d) - j \sin(k_z d) \end{aligned} \quad (46)$$

Gradient Complex Multiplication For input gradients $\frac{\partial \mathcal{L}}{\partial \tilde{U}_{\text{out}}} = \text{grad}_{\text{real}} + j \cdot \text{grad}_{\text{imag}}$:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \text{real}_{\text{in}}} &= \text{grad}_{\text{real}} \cos(-k_z d) - \text{grad}_{\text{imag}} \sin(-k_z d) \\ \frac{\partial \mathcal{L}}{\partial \text{imag}_{\text{in}}} &= \text{grad}_{\text{real}} \sin(-k_z d) + \text{grad}_{\text{imag}} \cos(-k_z d) \end{aligned} \quad (47)$$

Bandlimiting in Backward Pass The same bandlimiting conditions apply:

$$\frac{\partial \mathcal{L}}{\partial \tilde{U}(f_x, f_y, 0)} = \begin{cases} 0 & \text{if } |f_x| \geq f_x^{\text{max}} \text{ or } |f_y| \geq f_y^{\text{max}} \\ \text{computed} & \text{otherwise} \end{cases} \quad (48)$$

Gradients only flow through physically valid propagating modes within the bandlimit.

11 Justification for Amplitude Decomposition

Although the primitive contribution $\alpha_n \mathbf{c}_n g_n(\mathbf{p})$ is a product, the three factors play distinct roles: $g_n(\mathbf{p})$ is a spatial kernel determined by the covariance (a function of \mathbf{p} , not a free parameter); $\mathbf{c}_n \in \mathbb{R}^C$ provides the per-channel RGB amplitude; and α_n is a channel-shared opacity bounded to $[0, 1]$ by a sigmoid for optimization stability, consistent with the standard GaussianImage formulation [74]. Collapsing α_n and \mathbf{c}_n into a single unconstrained amplitude removes this bounded control and reduces the optimizer’s effective degrees of freedom. As shown in Table. 5, the merged variant consistently degrades reconstruction quality, even when matched in parameter count.


12 Gabor’s Theory: A 1D Thought Experiment

To isolate the benefit of Gabor primitives from the full hologram pipeline, we conduct a controlled study on a complex-valued 1D signal. At equal parameter budgets, we reconstruct the signal with Gabor atoms versus uniform pixel

Table 5: Ablation on amplitude decomposition (50 DIV2K images). Merging α and \mathbf{c} into a single amplitude \mathbf{A} degrades quality.

Variant	Params	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Merged ($\mathbf{A} = \alpha \cdot \mathbf{c}$)	1.8 M	29.5	0.84	0.40
Merged + matched params	2.0 M	29.5	0.84	0.40
Ours (α, \mathbf{c}, g)	2.0 M	30.7	0.86	0.33

samples. As shown in Fig. 13, Gabor atoms consistently achieve higher PSNR than pixel samples across budgets. This supports our central claim: Gabor’s theory does not assert that Gaussians *surpass* per-pixel quality at full budget, but that they *retain* quality most effectively as the budget shrinks, owing to their minimum space–frequency uncertainty.

**Fig. 13:** Thought experiment on a complex-valued 1D signal: Gabor atoms versus uniform pixel samples at equal parameter budgets. At matched budgets, Gabor atoms retain higher PSNR, most notably as the budget shrinks.

13 Eyebow Analysis and Single-Viewpoint Reconstruction

Our method targets high-fidelity reconstruction at the center of the eyebow. As shown in Fig. 14, our Random POH exhibits a uniform energy spectrum than Smooth POH, similar to typical random POH [30]: at pixel-level, scene structures stay recognizable. In contrast, the smooth POH from NH3D [17] shows stronger concentration and sharper peaks. A uniform spectrum alone, however, does not guarantee a uniform eyebow, it requires *explicit* supervision across pupil positions [8,18], so our center-viewpoint objective is complementary to with such methods. As shown in Fig. 15, once our representation is optimized under pupil-shift supervision, it reconstructs stably across shifted pupil positions, confirming that it extends naturally to the eyebow-expansion setting.

14 Training Steps Visualization

Figure 16 illustrates the training progression of our method, showing simulated reconstructions near the focal plane along with the corresponding complex holograms and POH visualizations. The image quality becomes stable at around 1000+ steps.

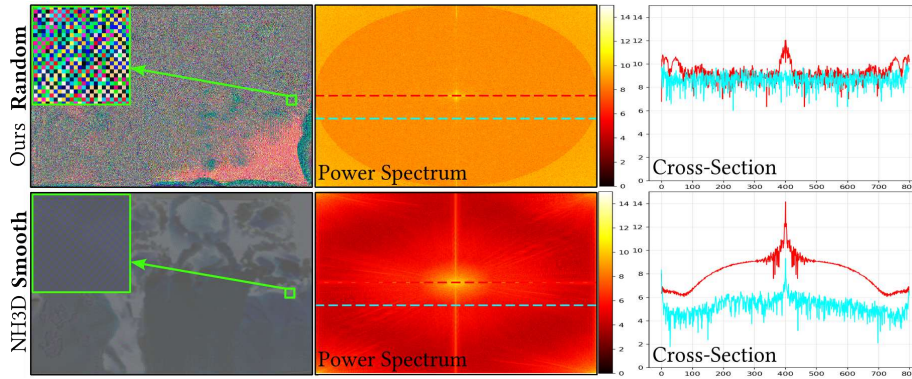


Fig. 14: Shifting eyebox result and spectrum analysis.




Fig. 15: Pupil-shift reconstructions when our representation is optimized under eyebox supervision.

15 Different Depth Planes Visualization

Figure 17 illustrates the simulated reconstructions of our method across different depth planes ($L = 1, 2, 3$), showing consistent preservation of fine structures from near to far focus. The results demonstrate that our representation maintains image fidelity across varying focal depths as reflected by metrics.

16 Different Propagation Distances Visualization

Figure 18 illustrates simulated reconstructions of our method at varying propagation distances, demonstrating consistent preservation of fine structural details from near to far focus. The results indicate that our representation maintains high image fidelity across a wide range of propagation distances and remains robust even at long distances, such as $50mm$, as reflected by the evaluation metrics.




Fig. 16: Comparison of simulated reconstructions at different training stages; for convenience of space, only the near focal plane is presented. The corresponding complex-valued 2D Gaussian hologram and the extracted random and double POH are shown in parallel. Results at 2000 steps are evaluated using PSNR, SSIM, and LPIPS. (Source Image: [5])

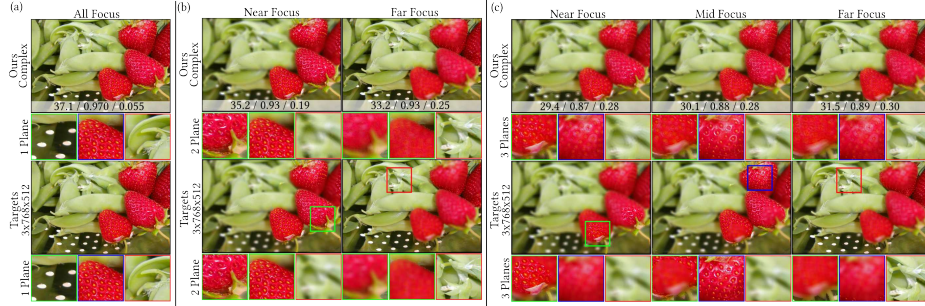


Fig. 17: Comparison of simulated reconstructions of our method for different depth planes. Results are evaluated using PSNR, SSIM, and LPIPS. (Source Image: [65])



Fig. 18: Comparison of simulated reconstructions of our method for different propagation distances, ranging from 2mm to 50mm and the volume depth is 4mm. Results are evaluated using PSNR, SSIM, and LPIPS. (Source Image: [4, 28])

17 Extra Experimentally Captured Results

Fig 19, Fig 20, Fig 21, and Fig 22 present experimentally captured results across five distinct scenes at resolution of $3 \times 2048 \times 1280$. Compared to the independently trained Random POH, our method achieves an effective suppression of noise without relying on additional time-multiplexing [16], wavelength-multiplexing [34, 52], or light-field-based methods [33, 51], which often costs substantial memory and computational overhead for better image quality.

Although the captures obtained using Smooth POH also exhibit good image quality with clear focus and defocus, they suffer from reduced contrast and brightness relative to Random POH. To partially mitigate this degradation, a different set of laser powers was applied during acquisition, which, however, introduces a noticeable shift in the overall color tone compared with the captures from Random POH.

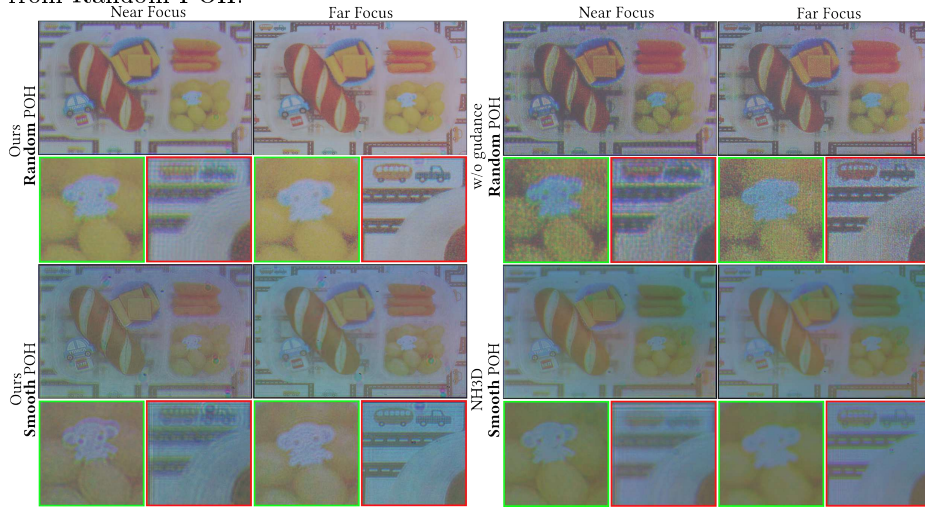


Fig. 19: Comparison of experimentally captured results of our method with Random POH, Smooth POH, NH3D [17], and an independently trained Random POH model. Results are evaluated using PSNR, SSIM, and LPIPS. (Source Image: [4])

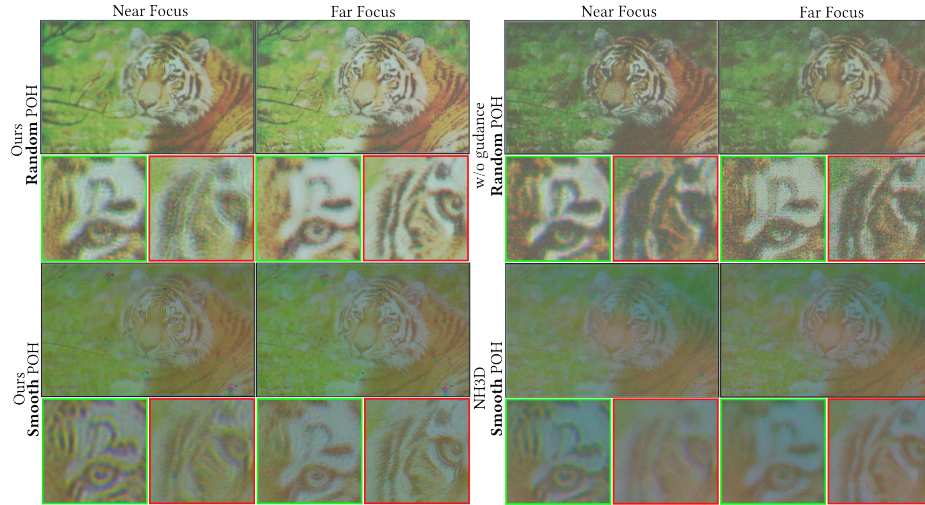


Fig. 20: Comparison of experimentally captured results of our method with Random POH, Smooth POH, NH3D [17], and an independently trained Random POH model. Results are evaluated using PSNR, SSIM, and LPIPS. (Source Image: [5])

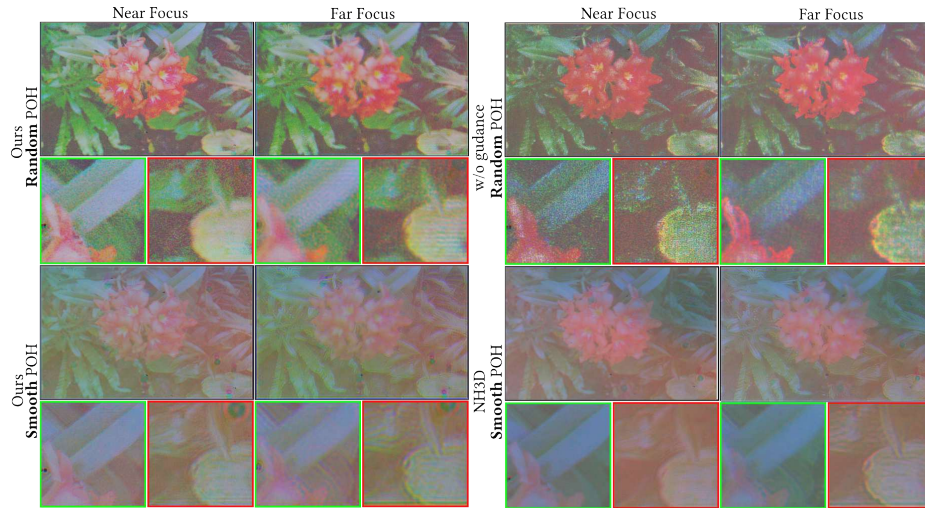


Fig. 21: Comparison of experimentally captured results of our method with Random POH, Smooth POH, NH3D [17], and an independently trained Random POH model. Results are evaluated using PSNR, SSIM, and LPIPS. (Source Image: [41])

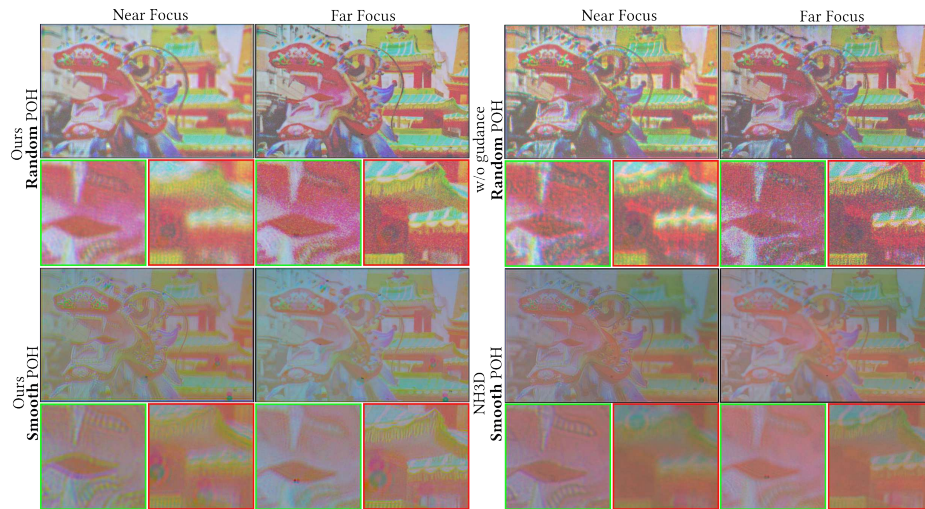


Fig. 22: Comparison of experimentally captured results of our method with Random POH, Smooth POH, NH3D [17], and an independently trained Random POH model. Results are evaluated using PSNR, SSIM, and LPIPS. (Source Image: [3])