



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

Yicheng Zhan¹(✉) , Xiangjun Gao² , Long Quan² , and Kaan Aksit¹

¹ University College London, London, UK
yicheng_zhan2001@outlook.com

² Hong Kong University of Science and Technology (HKUST), Hong Kong, China

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 3D content presentations [29]. 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)

Y. Zhan and X. Gao—Equal contribution.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/978-3-032-37574-2_31.

is to design compact and efficient representations that preserve high-frequency details of holograms while being scalable [52, 59].




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]).

Conventional image representation methods, such as Implicit Neural Representation (INR) [41, 45, 53], optimize a continuous implicit function to represent an image. However, implicit functions generally favor continuous low-frequency data, making it difficult to capture the high-frequency details typical of holograms. In parallel, autoencoder-based approaches [6, 42, 54, 58], typically pretrained 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 [64, 69, 70] 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 [11, 13, 66]. 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 [18], 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 [39]), 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 [66] 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 [23] 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 frameworks. For readers less familiar with the underlying principles of CGH, we present the essential concepts in Sect. 2.2. Our code is available at [65].

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 [6, 42, 58] and INRs [32, 45, 46, 53] 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 [19, 28, 35, 60], several works have extended Gaussian primitives from neural rendering to natural image encoding [64, 69, 70], 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 (Fig. 2).

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 natural images that capture only smooth, low-frequency intensity variations, a hologram simultaneously




Fig. 2. (Simulated) Smooth hologram shows high quality at the pupil center but degrade severely with pupil shifts, whereas random hologram remains visible (Source Image: [56]).

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 [23] and Random POH with directly optimized phase [47]. For more details of different hologram formats, please refer to Sect. 3.3.

Light Propagation. The core of CGH is free-space propagation based on scalar diffraction theory [20]. 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 [36, 71]. 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 [10, 25, 26, 30, 48], or trains neural networks to directly predict hologram pixels [15, 40, 50, 51, 67], both of which yield large solution spaces and hinder scalability. More recently, Gaussian primitives have been introduced to bridge computer graphics and holography. Here we clarify the positioning of our work relative to these efforts. *3D novel-view-synthesis (NVS) CGH* methods operate on multi-view inputs: Gaussian Wave Splatting (GWS) [13] and Random-phase GWS [11] leverage pretrained 2DGS scenes [24] for geometry-aware modeling of interference and diffraction, while complex-valued holographic radiance fields [66] demonstrate that 3D Gaussians can directly represent volumetric holographic scenes.




Fig. 3. Hologram reconstruction via free-space light propagation (Source Image: [27]).

In contrast, *natural-image Gaussian* representations [64, 69, 70] contain no wave optics and underperform on holograms, whereas *per-pixel CGH* [8, 15, 26] is dense and incurs a large solution space. A concurrent effort further applies 2D Gaussians to compress double-phase holograms [17]. Distinct from all of these, our method is a compact complex-valued representation designed specifically for single-view hologram optimization and rendering, 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. 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: [27].)

3.1 Complex-Valued 2D Gaussian Primitives

Figure 4 shows the training pipeline of our method. Building on the development of complex-valued holographic radiance fields [66], we extend the representation to define a complex-valued 2D Gaussian primitive for single-hologram optimization. Whereas [66] 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 Sect. 4). We apply activation functions to enforce valid parameter ranges (See Suppl Sect. 2.4 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 Sect. 2.4 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 [18] 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 $\Delta x \Delta f_x = \frac{1}{2}$ [18]. 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** [18], $\psi(x) = \exp(-\beta^2(x - x_0)^2) \cdot \exp(j(2\pi f_0 x + \varphi))$. 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. 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. 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 samples. As shown in Fig. 5, 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. 5. 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.

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 Sect. 2.2. Specifically, we adopt the BLASM [36] 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 $\mathcal{L}_{MSE} = \frac{1}{L} \sum_{l=1}^L \|\hat{I}_l - I_l\|^2$, where $\hat{I}_l$ denotes the ground-truth multi-plane focal stack images generated by Kavaklı et al. [26]. To further improve the defocus region’s image quality, we utilize the reconstruction loss $\mathcal{L}_{recon}$ by Kavaklı et al. [25], computed as

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

where M_l is the binary mask for depth plane Π_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 [49], the commercial holographic displays are predominantly **Liquid Crystal based** phase-only displays that do not support displaying complex holograms directly [5]. 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) [23] 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 $\mathbf{H}_{\text{rand}}$ 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

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

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 [66] 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 [28] and adapted to 2D screen-space. For the details of development and derivation, please refer to Suppl Sect. 2.

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 Sect. 3.

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 [61] with empirically-selected, parameter-specific learning rates: positions 10^{-2} (cosine-annealed to 10^{-3} [34]), 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 Sect. 6). Depth maps are provided by Depth Anything v2 [62] and MiDaS [43]. 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, 50, 51, 67]. 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 Sect. 1.

5 Evaluation

In this section, we conduct a comprehensive evaluation of our method against baselines. We use the first 50 images from DIV2K [57] 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 [68], averaged over across planes, together with parameter counts, peak VRAM usage, and training time. For results of training progression, varying depth-plane, and varying propagation distances, please refer to Suppl Sect. 6, 7, and 8.




Fig. 6. 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. (Color figure online)

5.1 Runtime and Memory Performance

Figure 6 presents the performance results of our kernel; the development detail is provided in Sect. 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* [6]	11.6	0.09	0.79	2.7 G	2.5 M	-
MLP [45]	7.5	0.04	0.85	9.9 G	1.0 M	6.9
SIREN [53]	7.6	0.05	0.84	13.1 G	1.0 M	7.8
Image-GS [64]	17.2	0.29	0.70	1.3 G	2.4 M	1.6
Instant-GI [70]	23.5	0.56	0.56	3.4 G	2.8 M	0.9
GI [69]	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




Fig. 7. Qualitative comparison of simulated reconstructions at near and far focal planes. Our method uses a 5:1 parameter ratio, while existing Gaussian-based approaches [64,69,70] use equal primitive counts for the two real components of the complex field. *GI* denotes GaussianImage [69]; *Instant-GI* uses network-predicted initialization with variable primitive counts (Source Image: [4,31]).

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 [45] and SIREN [53] yield PSNR below 8 dB and LPIPS around 0.85, while pretrained TAESD [6] 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 [64] achieving 17.2 dB and GI [69] and Instant-GI [70] 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. 7, 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.

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. 8, we show

Table 2. Quantitative comparison across our method, learned CGH and optimization methods. (* Naive Opt refers to Random POH baseline used in [9,30,40,48]). 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 [15]	28.3	0.92	0.31	7.1 G	3.9 M	90	31
TensorV2 [51]	27.1	0.94	0.29	8.7 G	0.1 M	80	48
GWS [13]	28.2	0.76	0.41	2.5 G	1.1 M	5.1	6840
U-Net [67]	27.2	0.91	0.35	6.3 G	2.2 M	100	18
Multi-color [26]	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 [8]	25.3	0.47	0.48	3.5 G	2.0 M	2.8	-
Multi-color [26]	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	-

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 [15], Tensor V2 [51], and U-Net [67]), 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, we extend its development to support single-plane POH rendering. Compared with GWS [13], 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 [26] remaining below 21 dB. Wirtinger Holography [8] 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. Comparison of simulated reconstructions at near and far focus using different optimization and learned CGH methods. *Multi-color* refers to [26]; *U-Net* [44] refers to typical learned CGH networks widely used in [2, 12, 15, 22, 33, 40, 67] (Source Image: [37]).

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

Parameter Reduction Ratio	Params	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Render (ms)
Dense Per-pixel	4.0 M	32.3	0.893	0.29	4.02
2:1	2.0 M	31.9	0.891	0.30	2.58
3:1	1.3 M	31.5	0.885	0.31	2.33
5:1	0.8 M	30.7	0.863	0.33	2.13
7:1	0.6 M	30.3	0.856	0.34	1.90
10:1	0.4 M	29.4	0.835	0.37	1.72

5.4 Ablation Study

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. Figure 9 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.

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. Figure 10 reports quantitative results in simulation. Our conversion procedure yields higher-quality reconstructions and, as a side effect,




Fig. 9. 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: [55]).

Table 4. Ablation study of POH conversion losses and Gaussian definition choices, evaluated on the test image [3] and [21].

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

appears to reduce the noise artifacts commonly seen in Random POH methods [16, 26, 30]. The zoomed insets in Fig. 10 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 Sect. 9.

Experiment. Figure 11 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. 10. 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, 21, 55]).




Fig. 11. Experimental comparison of reconstructions at near and far focal planes between our method and independently optimized Random POH (Source Image: [3,37]).

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 [63] pipeline. Our method achieves higher reconstruction quality (+6.3 dB) with 33% fewer parameters (12 vs. 9×2 per pair). As shown in Fig. 12, 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.




Fig. 12. Simulated reconstructions between complex-valued 2D Gaussians and naive paired real-valued Gaussians. (Source: [21])

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. 10 and Fig. 11, our conversion procedure appears to suppress noise artifacts 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 [7, 14] target on complementary objectives rather than serving as direct competitors*. In Suppl Sect. 5, 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 non-linear 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 [9, 10, 14, 38].

7 Conclusion

In this paper, we extend the complex-valued 3D Gaussian formulation introduced for holographic radiance fields [66] 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 [18] attaining the minimum space–frequency uncertainty, supported by a differentiable rasterizer and GPU-optimized light propagation. Unlike end-to-end POH methods, our representation serves as a general intermediate convertible to *both* smooth and random POH, decoupling the representation from hardware-specific encoding; using it as structural guidance to reduce the search space of the hologram optimization. Experiments show that our method enables a 5:1 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. Alphab.fr: Chinese new year 2016 in london (2016). [Openverse](#)
4. anotherlunch.com: Soft pretzel, carrots, grapes, cheddar cheese - easylunchboxes & tomica bento (2011). [Openverse](#)
5. Blanche, P.A.: Holography, and the future of 3d display. *Light: Advanced Manufacturing* **2**(4), 446–459 (2021). <https://doi.org/10.37188/lam.2021.028>, <https://www.light-am.com/en/article/doi/10.37188/lam.2021.028>
6. Bohan, O.B.: Tiny autoencoder for stable diffusion (2023). Accessed 22 May 2024
7. Chakravarthula, P., et al.: Pupil-aware holography. *ACM Trans. Graph.* **41**(6) (2022). <https://doi.org/10.1145/3550454.3555508>
8. Chakravarthula, P., Peng, Y., Kollin, J., Fuchs, H., Heide, F.: Wirtinger holography for near-eye displays. *ACM Trans. Graph.* **38**(6) (2019). <https://doi.org/10.1145/3355089.3356539>,
9. 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) (2020). <https://doi.org/10.1145/3414685.3417846>
10. 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 2024. Association for Computing Machinery, New York, NY, USA (2024). <https://doi.org/10.1145/3680528.3687600>
11. 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)
12. Chen, K., Wen, A., Zhang, Y., Chakravarthula, P., Sun, Q.: View synthesis for 3d computer-generated holograms using deep neural fields. *Opt. Express* **33**(9), 19399–19408 (2025)
13. Choi, S., Chao, B., Yang, J., Gopakumar, M., Wetzstein, G.: Gaussian wave splatting for computer-generated holography. *ACM Trans. Graph.* **44**(4) (2025). <https://doi.org/10.1145/3731163>
14. 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)
15. 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) (2021). <https://doi.org/10.1145/3478513.3480542>
16. Chu, V., et al.: Artifact-resilient real-time holography. *ACM Trans. Graph.* **44**(6) (2025). <https://doi.org/10.1145/3763361>
17. 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>
18. Gabor, D.: Theory of communication. part 1: the analysis of information. *J. Inst. Electr. Eng.-part III: Radio Commun. Eng.* **93**(26), 429–441 (1946)
19. Gao, X., et al.: Mani-gs: Gaussian splatting manipulation with triangular mesh. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 21392–21402 (2025)
20. Goodman, J.W.: *Introduction to Fourier Optics*. Roberts and Company Publishers (2005)

21. 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 (2020). [Openverse](#)
22. Hossein Eybposh, M., Cairra, N.W., Atisa, M., Chakravarthula, P., Pégard, N.C.: Deepcgh: 3d computer-generated holography using deep learning. *Opt. Express* **28**(18), 26636–26650 (2020)
23. Hsueh, C.K., Sawchuk, A.A.: Computer-generated double-phase holograms. *Appl. Opt.* **17**(24), 3874–3883 (1978)
24. 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)
25. Kavaklı, 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)
26. Kavaklı, K., Shi, L., Urey, H., Matusik, W., Akşit, K.: Multi-color holograms improve brightness in holographic displays. In: *ACM SIGGRAPH ASIA 2023 Conference Proceedings*, ACM, Sydney, NSW, Australia (2023). <https://doi.org/10.1145/3610548.3618135>
27. kennejima: Guy fieris vegas kitchen & bar (2014). [Openverse](#)
28. 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)
29. 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)
30. Kuo, G., Schiffers, F., Lanman, D., Cossairt, O., Matsuda, N.: Multisource holography. *ACM Trans. Graph. (Tog)* **42**(6), 1–14 (2023)
31. kuujinbo: Manette, wa graffiti wall (2012). [Openverse](#)
32. 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)
33. 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)
34. 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>
35. 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). <https://doi.org/10.1145/3680528.3687694>
36. Matsushima, K., Shimobaba, T.: Band-limited angular spectrum method for numerical simulation of free-space propagation in far and near fields. *Opt. Express* **17**(22), 19662–19673 (2009)
37. Mildenhall, B., et al.: Local light field fusion: practical view synthesis with prescriptive sampling guidelines. *ACM Trans. Graph. (TOG)* (2019). <https://bmild.github.io/llff/>
38. Nam, S.W., Kim, Y., Kim, D., Jeong, Y.: Depolarized holography with polarization-multiplexing metasurface. *ACM Trans. Graph. (TOG)* **42**(6), 1–16 (2023)

39. Paszke, A., et al.: Pytorch: a portable deep learning library. In: Advances in Neural Information Processing Systems 32, pp. 8024–8035. Curran Associates, Inc. (2019). <http://papers.neurips.cc>
40. Peng, Y., Choi, S., Padmanaban, N., Wetzstein, G.: Neural holography with camera-in-the-loop training. ACM Trans. Graph. (TOG) **39**(6), 1–14 (2020)
41. 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 2025. Association for Computing Machinery, New York, NY, USA (2025). <https://doi.org/10.1145/3721250.3742993>
42. 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)
43. Ranftl, R., Lasinger, K., Hafner, D., Schindler, K., Koltun, V.: Towards robust monocular depth estimation: mixing datasets for zero-shot cross-dataset transfer. IEEE Trans. Pattern Anal. Mach. Intell. **44**(3) (2022)
44. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
45. Rumelhart, D.E., Hinton, G.E., Williams, R.J.: Learning representations by back-propagating errors. Nature **323**(6088), 533–536 (1986)
46. Saragam, 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)
47. Schiffers, F., et al.: Stochastic light field holography. In: IEEE International Conference on Computational Photography (ICCP) (2023)
48. 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) (2025). <https://doi.org/10.1145/3732935>
49. 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) (2017). <https://doi.org/10.1145/3130800.3130832>
50. 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)
51. Shi, L., Li, B., Matusik, W.: End-to-end learning of 3d phase-only holograms for holographic display. Light: Sci. Appl. **11**(1), 247 (2022)
52. Shi, L., Webb, R., Xiao, L., Kim, C., Jang, C.: Neural compression for hologram images and videos. Opt. Lett. **47**(22), 6013 (2022)
53. Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G.: Implicit neural representations with periodic activation functions. Adv. Neural. Inf. Process. Syst. **33**, 7462–7473 (2020)
54. Song, Y., Dhariwal, P., Chen, M., Sutskever, I.: Consistency models. In: ICML (2023)
55. Sutherland, B.: Wasabi rainbow sushi set (2010). [Openverse](#)
56. over 10 million views Thanks !!, S.W.: A pair of ecuadorian amazon red-lored parrots (2009). [Openverse](#)

57. Timofte, R., Agustsson, E., Van Gool, L., Yang, M.H., Zhang, L., Lim, B.: Ntire 2017 challenge on single image super-resolution: methods and results In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops, (2017)
58. Wang, Y., et al.: 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)
59. Wang, Y., Chakravarthula, P., Sun, Q., Chen, B.: Joint neural phase retrieval and compression for energy-and computation-efficient holography on the edge. *ACM Trans. Graph.* **41**(4) (2022)
60. Wu, G., et al.: 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)
61. Xie, X., Zhou, P., Li, H., Lin, Z., Yan, S.: Adan: adaptive nesterov momentum algorithm for faster optimizing deep models. *IEEE Trans. Pattern Anal. Mach. Intell.* (2024)
62. Yang, L., et al.: Depth anything v2. [arXiv:2406.09414](https://arxiv.org/abs/2406.09414) (2024)
63. Ye, V., et al.: gsplat: an open-source library for gaussian splatting. *J. Mach. Learn. Res.* **26**(34), 1–17 (2025)
64. 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) (2025)
65. Zhan, Y., Gao, X., Quan, L., Akşit, K.: Complex-valued 2d gaussian representation for computer-generated holography (2026). <https://github.com/complight/Complex-Valued-2D-Gaussian-Representation>. source code
66. Zhan, Y., Shin, D.H., Baek, S.H., Akşit, K.: Complex-valued holographic radiance fields. *ACM Trans. Graph.* **45**(3) (2026). <https://doi.org/10.1145/3804450>
67. Zhan, Y., Sun, Q., Shi, L., Matusik, W., Akşit, K.: Configurable holography: towards display and scene adaptation. arXiv preprint [arXiv:2405.01558](https://arxiv.org/abs/2405.01558) (2024)
68. 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
69. Zhang, X., et al.: Gaussianimage: 1000 fps image representation and compression by 2d gaussian splatting. In: European Conference on Computer Vision, pp. 327–345. Springer, Cham (2024)
70. Zhang, Y., et al.: Image-gs: content-adaptive image representation via 2d gaussians. In: Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers, pp. 1–11 (2025)
71. 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 2024, Association for Computing Machinery, New York, NY, USA (2024). <https://doi.org/10.1145/3681758.3697989>