

Hologram Representation via Quadratic Phase Gaussian Splatting

HAOLONG WANG, Swansea University, United Kingdom
 YICHENG ZHAN, University College London, United Kingdom
 KAAAN AKŞIT, University College London, United Kingdom
 SIMENG QIU, Swansea University, United Kingdom

We introduce Complex-Valued Quadratic Phase Gaussian (CVQPG), a novel hologram representation method that augments each 2D Gaussian primitive with a quadratic phase profile controlled by a learnable curvature parameter. Against the planar Gaussian baseline, CVQPG improves the average PSNR of holographic reconstructions by 0.19 dB (RGB) and 0.33 dB (grayscale) at equal primitive counts, and by 0.05 dB (RGB) and 0.08 dB (grayscale) at equal parameter counts, where it still leads in all visual quality metrics. Our frequency-domain analysis shows that CVQPG better preserves the mid-to-high frequency band of natural images, indicating that modulating primitive wavefronts is an effective and lightweight enhancement.

CCS Concepts: • **Computing methodologies** → **Computer graphics**; **Rasterization**; **Computer vision representations**.

1 Introduction

Computer-Generated Holography (CGH) plays an essential role in applications such as 3D displays [4, 11, 13]. Given the dense spatial variations in holograms, designing efficient representations that preserve high-frequency details remains a major challenge in CGH. Implicit Neural Representations (INR) [7] favor low-frequency content, and Gaussian image representations [15] fit 2D Gaussians to pixels for compression. A separate line of research tackles hologram novel view synthesis using Gaussians [3, 12]. Closest to us, Zhan et al. [11] place complex-valued Gaussians directly on the 2D hologram plane to model interference and diffraction, enabling efficient and high-fidelity reconstruction.

The work by Zhan et al. assumes a flat phase profile for each primitive; the representational capacity of such Gaussian primitives remains an under-explored question. We further explore this topic by representing full-complex holograms using the complex-valued Gaussian modulated with a quadratic phase wavefront, which explicitly models the field with an additional learnable curvature parameter. Our work makes the following contributions:

- **Quadratic Phase Gaussian with Curvature Parameter.** Our proposed quadratic phase Gaussian primitive adds one curvature parameter to modulate its wavefront. We derive a per-primitive curvature bound from the display’s sampling limit, avoiding high-frequency aliasing.

Our comprehensive experiments with various settings show that CVQPG achieves a generalized quality improvement and outperforms the baseline in parameter effectiveness. We investigate the frequency domain features of CVQPG and ablate its components. Our code is available at [GitHub:gsmark36/complex-valued-quadratic-phase-gaussian](https://github.com/GitHub:gsmark36/complex-valued-quadratic-phase-gaussian).

Authors’ Contact Information: Haolong Wang, Swansea University, United Kingdom, haolong.wang@swansea.ac.uk; Yicheng Zhan, University College London, United Kingdom, yicheng.zhan.21@ucl.ac.uk; Kaan Akşit, University College London, United Kingdom, k.aksit@ucl.ac.uk; Simeng Qiu, Swansea University, United Kingdom, simeng.qiu@swansea.ac.uk.

2 Methods

Problem Definition. The holographic display can reconstruct a 3D scene as multi-plane intensity images by propagating light from a complex hologram $\mathbf{H} \in \mathbb{C}^{N_c \times H \times W}$ with N_c color channels. To synthesize the intensity target for each depth plane, we take an image $\mathbf{I}_{target} \in \mathbb{R}^{N_c \times H \times W}$ and its depth map $\mathbf{D} \in \mathbb{R}^{H \times W}$ as the input for the target synthesis function f_{synth}^m at plane m . The optimization problem is therefore formulated as:

$$\hat{\mathbf{H}} \leftarrow \underset{\mathbf{H}}{\operatorname{argmin}} \sum_{m=1}^M \mathcal{L}(f_{recon}^m(\mathbf{H}), f_{synth}^m(\mathbf{I}_{target}, \mathbf{D})), \quad (1)$$

where M denotes the total number of depth planes, f_{recon}^m is the reconstruction function at plane m , and $\mathcal{L}$ represents the loss function that measures the difference between the reconstructed image and the synthesized target.

Complex-Valued Planar Gaussian. Our method builds upon the complex-valued 2D Gaussian representation [11], which we refer to as the planar Gaussian baseline model. It encodes a hologram as a set of N Gaussian primitives. Each primitive is parameterized as $\mathcal{G} = \{\alpha, \mathbf{c}, \mathbf{x}, \mathbf{s}, \theta, \boldsymbol{\varphi}\}$, where $\alpha \in \mathbb{R}$ denotes the opacity, $\mathbf{c} \in \mathbb{R}^{N_c}$ the per-channel color amplitude, $\mathbf{x} \in \mathbb{R}^2$ the 2D position, $\mathbf{s} \in \mathbb{R}^2$ the scales, $\theta \in \mathbb{R}$ the rotation angle, and $\boldsymbol{\varphi} \in \mathbb{R}^{N_c}$ the per-channel phase. The distribution of the Gaussian primitive is determined by the 2D covariance matrix $\Sigma = \mathbf{R}\mathbf{S}\mathbf{S}^\top\mathbf{R}^\top$, where $\mathbf{R}$ is the rotation matrix of angle θ , and $\mathbf{S} = \operatorname{diag}(\mathbf{s})$ is the scaling matrix of scales $\mathbf{s}$. Therefore, the contribution of a Gaussian primitive at pixel coordinate $\mathbf{p}$ is

$$\mathcal{G}(\mathbf{p}) = \alpha \cdot \mathbf{c} \cdot \exp\left(-\frac{1}{2}(\mathbf{p} - \mathbf{x})^\top \Sigma^{-1}(\mathbf{p} - \mathbf{x})\right) \cdot \exp(j\boldsymbol{\varphi}). \quad (2)$$

The phase profile of each primitive is flat across its effective area, equivalent to a plane wave emitted from a Gaussian ellipse.

Complex-Valued Quadratic Phase Gaussian. Considering the representational capacity, we present a novel primitive: Complex-Valued Quadratic Phase Gaussian (CVQPG). Our design incorporates the Quadratic Phase Factor (QPF) to explicitly modulate field curvature with a minimal overhead of one extra parameter per primitive. Conventionally, we modulate a Gaussian primitive with a standard QPF,

$$Q_{std}(\mathbf{p}) = \exp\left(j\frac{k}{2f}\|\mathbf{p} - \mathbf{x}\|^2\right), \quad (3)$$

where k is the wavenumber, and f is the focal length. However, the standard QPF models an isotropic circular wavefront, leading to a mismatch with the Gaussian’s anisotropic amplitude envelope (Fig. 2). To resolve this, we utilize the squared Mahalanobis distance $D_{mah}^2(\mathbf{p}) = (\mathbf{p} - \mathbf{x})^\top \Sigma^{-1}(\mathbf{p} - \mathbf{x})$ as a unified spatial metric. Therefore, the shape-adaptive Mahalanobis QPF is formulated as

$$Q_{mah}(\mathbf{p}) = \exp\left(j\frac{k\gamma \det(\mathbf{S})}{2}D_{mah}^2(\mathbf{p})\right), \quad (4)$$




Fig. 1. Complex-Valued Quadratic Phase Gaussian (CVQPG) represents full complex holograms, preserving high-frequency details in reconstructions. CVQPG adds an extra learnable curvature parameter γ , and optimizes the hologram by comparing the reconstructed intensity images with the synthesized targets. (Source Image: [8])

where $\gamma = 1/f$ is the learnable curvature parameter representing optical power, and $\det(S)$ is the determinant of the scaling matrix S . Consequently, each CVQPG primitive is parameterized as $\mathcal{G} = \{\alpha, c, x, s, \theta, \phi, \gamma\}$, with an added scalar parameter $\gamma \in \mathbb{R}$. Hence, the contribution of a CVQPG primitive at pixel coordinate $\mathbf{p}$ is given by

$$\mathcal{G}(\mathbf{p}) = \alpha \cdot c \cdot \exp\left(-\frac{1}{2}D_{mah}^2(\mathbf{p})\right) \cdot \exp(j\phi) \cdot Q_{mah}(\mathbf{p}). \quad (5)$$

The configuration of a CVQPG primitive is equivalent to the combination of a planar Gaussian light source and a virtual lens with optical power γ . During optimization, we symmetrically bound γ within $[-\gamma_{\max}, \gamma_{\max}]$, which enables a continuous wavefront transformation from divergence ($\gamma > 0$) to convergence ($\gamma < 0$). We can represent a planar wavefront by $\gamma = 0$, avoiding numerical instability caused by the focal length $f \rightarrow \pm\infty$.




Fig. 2. Comparison of QPFs on a Gaussian primitive. (a) Shared amplitude envelope. Unlike (b) the standard QPF with misaligned circular fringes, (c) the Mahalanobis QPF yields elliptical fringes perfectly aligned with the envelope.

Scale-Adaptive Curvature Control. The quadratic phase profile acts as a linear chirp signal, providing high-frequency components for detailed hologram reconstruction. However, to prevent aliasing artifacts caused by the physical bandwidth limitations of the Spatial Light Modulator (SLM), we impose a frequency constraint on each primitive. By restricting the chirp evaluation to the 3σ footprint of each primitive, the maximum instantaneous spatial frequency is

$$\nu_{\max} = \frac{\gamma \cdot 3\sigma \cdot (\Delta x)^2}{\lambda}, \quad (6)$$

where $\sigma = \max(s)$ is the equivalent standard deviation, $\lambda = 2\pi/k$ is the wavelength, and Δx is the pixel pitch of the SLM. Constraining this peak frequency below the SLM's Nyquist limit ($\nu_{\max} \leq 0.5$ cycles/pixel) directly yields the scale-adaptive curvature bound for each primitive:

$$\gamma_{\max} = \eta \cdot \frac{\lambda}{6\sigma \cdot (\Delta x)^2}, \quad (7)$$

where $\eta = 0.9$ is a relaxation factor.

Hologram Rendering and Optimization. In the rendering process, rasterization is applied to accumulate all primitive contributions (Eq. 5) across the pixel grid to form the hologram,

$$\mathbf{H} = \left\{ \sum_{n=1}^N \mathcal{G}_n(\mathbf{p}) \mid \mathbf{p} \in [1, W] \times [1, H] \right\}, \quad (8)$$

where $H \times W$ is the size of the hologram. We propagate the hologram via the conventional Band-limited Angular Spectrum Method (BLASM) to reconstruct the field intensity $I_m = |\text{BLASM}_m(\mathbf{H})|^2$ at plane m . At each plane, we compare the reconstructed intensity I_m with the synthesized target intensity $\hat{I}_m$, minimizing the visual difference between them. The per-plane loss function consists of a reconstruction loss $\mathcal{L}_{recon}$ proposed by Kavaklı et al. [4] and a similarity loss $\mathcal{L}_{SSIM} = (1 - \text{SSIM})$. To accurately simulate depth-dependent focus and defocus effects, the total training loss $\mathcal{L}$ is summed over all depth planes as

$$\mathcal{L} = \sum_{m=1}^M (\mathcal{L}_{recon}^m(I_m, \hat{I}_m) + w_1 \cdot \mathcal{L}_{SSIM}^m(I_m, \hat{I}_m)), \quad (9)$$

where $w_1 = 0.005$ is the SSIM loss weight.

Curvature Warm-Up Strategy. We optimize the learnable curvature parameter γ via a two-stage strategy. We initialize and lock the curvature at $\gamma = 0$ for all primitives, utilizing the planar Gaussians to optimize the low-frequency patterns. In the refinement stage, the curvature parameter is unlocked and optimized with other parameters. This coarse-to-fine approach stabilizes early-stage training and allows the model to refine the residual high-frequency details.

3 Evaluation and Discussion

Implementation. We utilize a $3.74\ \mu\text{m}$ pixel pitch and RGB wavelengths of 639 nm, 532 nm, and 473 nm to propagate the holograms to two depth planes at 1 mm and 5 mm. Shared parameters follow the same initialization schemes as the baseline model [11]. $\gamma_{\max}$ is calculated from the shortest wavelength and bound γ via a tanh activation. We train the models using identical PyTorch pipelines and the Adan optimizer [9]. The training runs for 2000 steps, with a 400-step warm-up stage. All experiments are conducted on a single NVIDIA TITAN RTX GPU. **We display the reconstructed holograms on our holographic display prototype (see Supplementary Section 1).**

Quantitative Evaluations. The visual quality of reconstructed images is evaluated using five metrics, including Peak Signal-to-noise Ratio (PSNR), Structural Similarity (SSIM), Perceptual Similarity Metric (LPIPS) [14], FLIP [2], and ColorVideoVDP (CVVDP) [5]. Our dataset is derived from both DIV2K [1] and the Real Forward-Facing dataset [6]. Then we generate the depth map using Depth Anything V2 [10]. We first compare CVQPG against the planar Gaussian baseline on ten scenes at a resolution of 640×480 with compression ratios of 10% and 20%. Under equal primitive counts, CVQPG outperforms the baseline across all metrics, achieving mean PSNR gains of +0.19 dB in RGB and +0.33 dB in grayscale (Tab. 1). Since CVQPG adds one trainable scalar per primitive that brings minor parameter overhead, we further conduct an equal-parameter evaluation by enlarging the baseline’s primitive count by $13/12\times$ in RGB and $9/8\times$ in grayscale (see Baseline⁺ rows in Tab. 1). **Per-scene results are given in Section 3 of the supplementary material.** To assess improvements in the frequency domain, we measure the reconstruction MSE at each spatial frequency against the target. Evaluated on ten scenes at an equal primitive count (Fig. 3), the MSE reductions of CVQPG over the baseline are negligible near zero frequency, peak in the mid-to-high frequency band (11% RGB, 22% grayscale), and vanish near the Nyquist limit. **Training time and peak GPU memory are reported in Section 2 of the supplementary material.**

Table 1. Quantitative comparison averaged over ten scenes and two depth planes. We compare CVQPG against the standard Baseline (equal primitive count N) and the enlarged Baseline⁺ (equal parameter count). Δ rows show CVQPG’s improvement over the standard Baseline. Best in **bold**.

Setting	Method	N	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	FLIP $\downarrow$	CVVDP $\uparrow$
RGB 20%	Baseline	30,720	29.265	0.8307	0.2895	0.1144	9.168
	Baseline ⁺	33,280	29.386	0.8336	0.2853	0.1134	9.186
	CVQPG	30,720	29.450	0.8364	0.2828	0.1125	9.203
	Δ		+0.185	+0.0057	−0.0066	−0.0019	+0.035
RGB 10%	Baseline	15,360	28.034	0.7992	0.3292	0.1292	8.917
	Baseline ⁺	16,640	28.179	0.8034	0.3241	0.1271	8.950
	CVQPG	15,360	28.225	0.8053	0.3217	0.1265	8.966
	Δ		+0.191	+0.0060	−0.0076	−0.0026	+0.049
Grayscale 20%	Baseline	15,360	29.504	0.8340	0.3813	0.0753	8.935
	Baseline ⁺	17,280	29.753	0.8394	0.3766	0.0733	8.988
	CVQPG	15,360	29.834	0.8443	0.3694	0.0725	9.019
	Δ		+0.331	+0.0103	−0.0119	−0.0028	+0.085
Grayscale 10%	Baseline	7,680	28.076	0.7964	0.4204	0.0888	8.548
	Baseline ⁺	8,640	28.317	0.8030	0.4134	0.0862	8.621
	CVQPG	7,680	28.398	0.8079	0.4069	0.0850	8.658
	Δ		+0.322	+0.0115	−0.0135	−0.0038	+0.109




Fig. 3. Frequency-resolved ΔMSE of CVQPG over the baseline for ten scenes. The gains cover the mid-to-high frequency band that contains fine details of natural images.

Qualitative Observations. To further explore the high-frequency feature of CVQPG, we reconstruct a single-plane, grayscale Siemens star. With equal primitives, CVQPG achieves a significant PSNR gain (+3.04 dB) with a distinct visual improvement in the central region (Fig. 5a), demonstrating its capacity to resolve high-frequency details. CVQPG exhibits ringing artifacts under zero-padding due to the abrupt truncation of its chirp phase, which is avoided by the baseline (Fig. 5b). We suppress this by applying a cosine-tapered window to smooth the hologram edges before propagation.

Ablation study. The ablation study covers two components of our method: the Scale-Adaptive Curvature Control and the Mahalanobis QPF. Starting from full CVQPG, we compare five scenes at the RGB 10% budget (Tab. 2). Replacing the scale-adaptive bound with a constant bound significantly reduces the PSNR gain (+0.185 $\rightarrow$ +0.048 dB), identifying it as the decisive component. Replacing the Mahalanobis QPF with the standard QPF costs 0.036 dB on average, which confirms the benefit of the anisotropic phase profile. Finally, a vanilla model without the two components retains only a gain of +0.046 dB. This performance floor indicates that simply appending a standard QPF fails to capture the full potential of phase modulation.

Table 2. Component ablation averaged over five scenes (RGB, 10%, $N=15,360$). Δ denotes the mean PSNR gain over the baseline. Best in **bold**.

Variant	Δ	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	FLIP $\downarrow$	CVVDP $\uparrow$
Baseline	—	29.809	0.8419	0.3119	0.1128	9.095
Full (CVQPG)	+0.185	29.994	0.8471	0.3041	0.1109	9.137
w/o Adaptive Bound	+0.048	29.857	0.8438	0.3103	0.1122	9.108
w/o Mahalanobis QPF	+0.149	29.958	0.8462	0.3044	0.1108	9.134
None (Vanilla)	+0.046	29.855	0.8437	0.3098	0.1118	9.109

Conclusion and Discussion. Our approach improves reconstruction quality across all metrics and settings without per-scene hyperparameter tuning, demonstrating its inherent robustness. Under equal-parameter evaluation, CVQPG retains an advantage over the baseline, showing that optimizing a curvature parameter per primitive is more effective than scaling up the count of planar Gaussians. Our experiments also show that the additional curvature




Fig. 4. Visual comparison on simulated holographic reconstructions (RGB, 20%, $N=30,720$) at the front (1 mm) and rear (5 mm) depth planes. CVQPG provides improved visual quality over the baseline. (Source Image: [6])




Fig. 5. Qualitative comparisons. (a) CVQPG outperforms the baseline by +3.04 dB PSNR on the Siemens star target, providing distinct high-frequency details in the central region. (b) Abrupt boundary truncation induces ringing artifacts in CVQPG, which is absent in the baseline.

parameter achieves its full potential only when strictly bounded by the per-primitive sampling limit. On the other hand, results in Tab. 1 reveal a chromatic limitation. The PSNR gains in grayscale (+0.331 dB and +0.322 dB) notably exceed those in RGB (+0.185 dB and +0.191 dB) because a shared curvature cannot be simultaneously optimal for all wavelengths. We also discover that CVQPG specifically preserves mid-to-high frequencies of natural images, demonstrating a content-dependent nature in reconstructions. **In future work, we could explore the wavelength-based curvature setting to address the chromatic limitation of our method.** We also aim to explore advanced phase formulations and further investigate the representational capacity of explicit hologram representations.

References

- [1] Eirikur Agustsson and Radu Timofte. 2017. NTIRE 2017 Challenge on Single Image Super-Resolution: Dataset and Study. In *2017 IEEE CVPRW*. 1122–1131.

- doi:10.1109/CVPRW.2017.150
- [2] Pontus Andersson, Jim Nilsson, Tomas Akenine-Möller, Magnus Oskarsson, Kalle Åström, and Mark D. Fairchild. 2020. FLIP: A Difference Evaluator for Alternating Images. *Proc. ACM Comput. Graph. Interact. Tech.* 3, 2, Article 15 (Aug. 2020), 23 pages. doi:10.1145/3406183
- [3] Suyeon Choi, Brian Chao, Jacqueline Yang, Manu Gopakumar, and Gordon Wetzstein. 2025. Gaussian Wave Splatting for Computer-Generated Holography. *ACM TOG* 44, 4, Article 57 (July 2025), 13 pages. doi:10.1145/3731163
- [4] Koray Kavaklı, Yuta Itoh, Hakan Urey, and Kaan Akşit. 2023. Realistic Defocus Blur for Multiplane Computer-Generated Holography. In *2023 IEEE VR*. 418–426. doi:10.1109/VR55154.2023.00057
- [5] Rafal K. Mantiuk, Param Hanji, Maliha Ashraf, Yuta Asano, and Alexandre Chapiro. 2024. ColorVideoVDP: A visual difference predictor for image, video and display distortions. *ACM TOG* 43, 4, Article 129 (July 2024), 20 pages. doi:10.1145/3658144
- [6] Ben Mildenhall, Pratul P. Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. 2019. Local light field fusion: practical view synthesis with prescriptive sampling guidelines. *ACM TOG* 38, 4, Article 29 (July 2019), 14 pages. doi:10.1145/3306346.3322980
- [7] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. 2020. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems* 33 (2020), 7462–7473.
- [8] Bernard Spragg. 2013. Colorful windmills Shenzhen China. Flickr.
- [9] Xingyu Xie, Pan Zhou, Huan Li, Zhouchen Lin, and Shuicheng Yan. 2024. Adan: Adaptive Nesterov Momentum Algorithm for Faster Optimizing Deep Models. *IEEE TPAMI* 46, 12 (2024), 9508–9520. doi:10.1109/TPAMI.2024.3423382
- [10] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. 2024. Depth Anything V2. arXiv:2406.09414 [cs.CV] https://arxiv.org/abs/2406.09414
- [11] Yicheng Zhan, Xiangjun Gao, Long Quan, and Kaan Akşit. 2026. Complex-Valued 2D Gaussian Representation for Computer-Generated Holography. In *Computer Vision – ECCV 2026*, Paolo Favaro, Zuzana Kukelova, Atsuto Maki, Anna Rohrbach, Konrad Schindler, and Federico Tombari (Eds.). Springer Nature Switzerland, Cham, 557–576. doi:10.1007/978-3-032-37574-2_31
- [12] Yicheng Zhan, Dong-Ha Shin, Seung-Hwan Baek, and Kaan Akşit. 2026. Complex-Valued Holographic Radiance Fields. *ACM TOG* 45, 3, Article 31 (April 2026), 16 pages. doi:10.1145/3804450
- [13] Jingzhao Zhang, Nicolas Pégard, Jingshan Zhong, Hillel Adesnik, and Laura Waller. 2017. 3D computer-generated holography by non-convex optimization. *Optica* 4, 10 (2017), 1306–1313.
- [14] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *2018 IEEE/CVF CVPR*. 586–595. doi:10.1109/CVPR.2018.00068
- [15] Xinjie Zhang, Xingtong Ge, Tongda Xu, Dailan He, Yan Wang, Hongwei Qin, Guo Lu, Jing Geng, and Jun Zhang. 2024. GaussianImage: 1000 FPS Image Representation and Compression by 2D Gaussian Splatting. arXiv:2403.08551 [eess.IV] https://arxiv.org/abs/2403.08551