

Compressing Double-Phase Holograms using 2D Gaussians

Xiaoyue Fan^{1*} and Yicheng Zhan¹ and Amrita Mazumdar² and Kaan Akşit¹

¹University College London, United Kingdom

²NVIDIA, New York, United States of America

Abstract

Effective compression of double-phase holograms remains an unresolved challenge due to their high-frequency nature, impeding the practicality of holographic displays. To address this challenge, we propose a hologram compression method by modifying the GaussianImage. Our method decomposes phase-only holograms into two components based on their intrinsic checkerboard pattern, separately optimizing each with a reduced set of 2D Gaussians. Our best case reduces the primitive count to only 3% of the baseline, achieving a compression ratio of 26% while preserving Mean PSNR = 43.39 dB in the reconstructed scenes.

Keywords: Computational Holography, Double-Phase Hologram, Hologram Compression, Gaussian Splatting

CCS Concepts

• Computing methodologies → Image compression; Reconstruction; • Theory of computation → Data compression;

1. Introduction

Holographic displays [KSU*23] represent 3D scenes by reshaping the light wavefront using dynamic, programmable phase-only holograms. Although recent deep learning approaches have shown improvements in phase-only hologram compression [PZSA25], these methods still exhibit limited compression efficiency.

We introduce a decomposition-based hologram compression framework built on our modified GaussianImage (GI) [ZGX*24] algorithm, tailored for double-phase holograms [HS78]. We leverage their intrinsic checkerboard pattern by partitioning them into complementary components, each optimized using a reduced set of 2D Gaussian primitives, resulting in a compressed representation. Unlike conventional learned methods [PZSA25] treating holograms as regular images, operating on 64×64 patches (~10s each) and requiring ~26 minutes for 512×512 holograms, our approach enables efficient sequential compression of two components in only ~4 minutes regardless of resolution, without compromising reconstruction fidelity. This is achieved by a refined rendering constraint in our modified GI preserves the underlying distribution and further suppresses artifacts. As shown in Fig. 1, our approach achieves 40.10 dB PSNR with a 26% compression ratio, whereas our baseline, GI, suffers from severe distortions and low perceptual quality.

2. Methods

Our baseline compresses double-phase holograms using a patch-based framework built on GI [ZGX*24], a Gaussian Splatting (GS)-based image representation method where each Gaussian is defined by its position, covariance, color coefficients, and opacity.

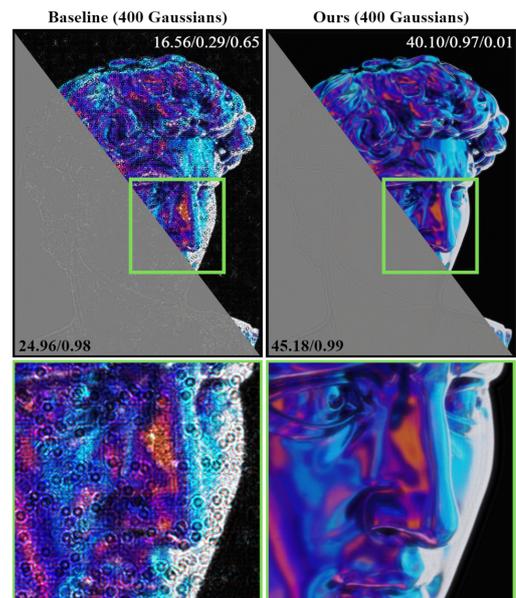


Figure 1: Comparison between our method and the baseline without decomposition. Each set presents the compressed double-phase holograms with PSNR and SSIM metrics (lower triangle) and the corresponding 3D reconstruction evaluated by PSNR, SSIM, and LPIPS (upper triangle) (Source Image: SIMON LEE).

To ensure valid covariance matrices during training, two factorizations are adopted: Cholesky, encoding the lower triangular el-

ements, and Rotation-and-Scale (RS), encoding the rotation angle with scaling factors. Rendering is simplified to a weighted summation $C_i = \sum_{n \in \mathcal{N}} c'_n \cdot \exp(-\sigma_n)$, where c'_n denotes the learnable color-opacity coefficients and σ_n controls the Gaussian falloff. The representation is optimized by the L2 loss. However, hard clamping during rendering distorts the underlying distribution, causing blob artifacts, as illustrated in the lower-left corner of Fig. 1. Our framework replaces it with a sinusoidal constraint, which smoothly controls the value range, yielding artifact-free perceptual quality.

Table 1: Comparison between the baseline with our work (both using RS factorization), with horizontal or vertical decomposition, in hologram fidelity (left) and reconstructed image quality (right) using PSNR and SSIM under equivalent compression ratios (CR).

	Patch Size	Gauss.	PSNR (dB) ↑	SSIM ↑	CR ↓
Baseline	64 × 64	400	23.40/14.93	0.98/0.16	26%
	128 × 128	1,600	24.97/15.49	0.98/0.16	26%
	256 × 256	8,000	26.89/17.50	0.99/0.22	33%
	512 × 512	32,000	27.92/18.45	0.99/0.25	33%
Ours-Hor.	64 × 64	400	49.18/43.39	0.99/0.97	26%
	128 × 128	1,600	48.98/42.80	0.99/0.96	26%
	256 × 256	8,000	48.88/42.39	0.99/0.96	33%
	512 × 512	32,000	48.41/41.58	0.99/0.95	33%
Ours-Ver.	64 × 64	400	49.08/42.84	0.99/0.97	26%
	128 × 128	1,600	48.93/42.22	0.99/0.96	26%
	256 × 256	8,000	48.82/41.72	0.99/0.96	33%
	512 × 512	32,000	48.48/41.22	0.99/0.95	33%

Vertical or Horizontal Decomposition The high-frequency nature of double-phase holograms poses challenges for Gaussian-primitive representation, as Gaussians struggle to capture rapid phase oscillations across pixels, leading to suboptimal results. To improve this, we employ directional decomposition as a preprocessing step that separates high- and low-value components, allowing the model to bypass the high-frequency pattern from the outset. The decomposition is performed along vertical and horizontal directions, dividing checkerboard pixels into four groups by value:

$$P_{h1} = P[0:2, 0:2], P_{h2} = P[1:2, 1:2], \quad (1)$$

$$P_{l1} = P[0:2, 1:2], P_{l2} = P[1:2, 0:2]. \quad (2)$$

For high-value components P_{h1}, P_{h2} , the vertical decomposition reorganizes sparse pixels into denser images by interleaving columns:

$$P_{\text{high}}^{(v)}[:, 2j] = P_{h1}[:, j], P_{\text{high}}^{(v)}[:, 2j+1] = P_{h2}[:, j], \quad (3)$$

while horizontal decomposition interleaves the rows:

$$P_{\text{high}}^{(h)}[2i, :] = P_{h1}[i, :], P_{\text{high}}^{(h)}[2i+1, :] = P_{h2}[i, :]. \quad (4)$$

The low-value components P_{l1}, P_{l2} follow the same procedure. Consequently, each decomposed image preserves half the spatial resolution on one axis, allowing us to train the GS-based model for each decomposed image independently. The final compressed image is recombined by reversing the decomposition process, restoring the complete double-phase hologram.

3. Results

We evaluate the PSNR and SSIM for compressed holograms utilizing Cholesky or RS factorization, experimenting with patch sizes from 64×64 to 512×512 to find the optimal balance between the compression ratio and fidelity. We employ the free-space light propagator from *odak* [AK23] to numerically reconstruct holograms with a volume depth of 5 mm across three depth layers. The consequent evaluation of 3D reconstructions averages over focal planes quantified by PSNR, SSIM, and LPIPS. With both decomposition strategies, our method achieves an average PSNR improvement of 26 dB over the baseline at the same compression ratio in reconstructed scenes. RS factorization shows significant gains, although it originally performs notably worse than Cholesky decomposition. As shown in Table. 1, the optimal 64×64 patch size yields the highest fidelity via horizontal decomposition using RS factorization, reaching an average of 49.18 dB PSNR and 0.99 SSIM, whose reconstructions further yield 43.39 dB PSNR, 0.97 SSIM, and 0.016 LPIPS on average. Both vertical and horizontal decompositions perform comparably, indicating that the primary improvement arises from the decomposition strategy itself rather than the choice of direction. Fig. 2 illustrates reconstruction comparisons across methods and compression ratios. Compared with the baseline [ZGX*24], our methods effectively eliminate blob artifacts and inter-patch boundary lines, as shown in Fig. 1, while introducing moderate computational overhead, doubling the runtime. Extensive results and visualizations are provided in the Supplementary.

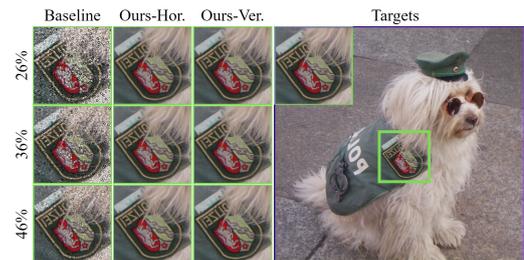


Figure 2: Comparison of numerical reconstructions using baseline and our framework (Source Image: *exfordy*).

References

- [AK23] AKŞIT K., KAVAKLI K.: Flexible modeling of next-generation displays using a differentiable toolkit. In *Practical Holography XXXVII: Displays, Materials, and Applications* (2023), vol. 12445, SPIE, pp. 131–132. 2
- [HS78] HSUEH C.-K., SAWCHUK A. A.: Computer-generated double-phase holograms. *Applied optics* 17, 24 (1978), 3874–3883. 1
- [KSU*23] KAVAKLI K., SHI L., UREY H., MATUSIK W., AKŞIT K.: Multi-color holograms improve brightness in holographic displays. In *SIGGRAPH ASIA 2023 Conference Papers* (New York, NY, USA, Dec. 2023), SA '23, ACM. 1
- [PZSA25] PENG Z., ZHAN Y., SPJUT J., AKŞIT K.: Assessing learned models for phase-only hologram compression. In *SIGGRAPH 2025 Posters (SA Posters '25)* (New York, NY, USA, August 10–14 2025), ACM, p. 2. 1
- [ZGX*24] 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 *Computer Vision - ECCV 2024* (2024), Springer, pp. 327–345. 1, 2