Random-phase Gaussian Wave Splatting for Computer-generated Holography

BRIAN CHAO*, Stanford University, USA
JACQUELINE YANG*, Stanford University, USA
SUYEON CHOI, Stanford University, USA
MANU GOPAKUMAR, Stanford University, USA
RYOTA KOISO, Stanford University, USA and KDDI Research, Japan
GORDON WETZSTEIN, Stanford University, USA

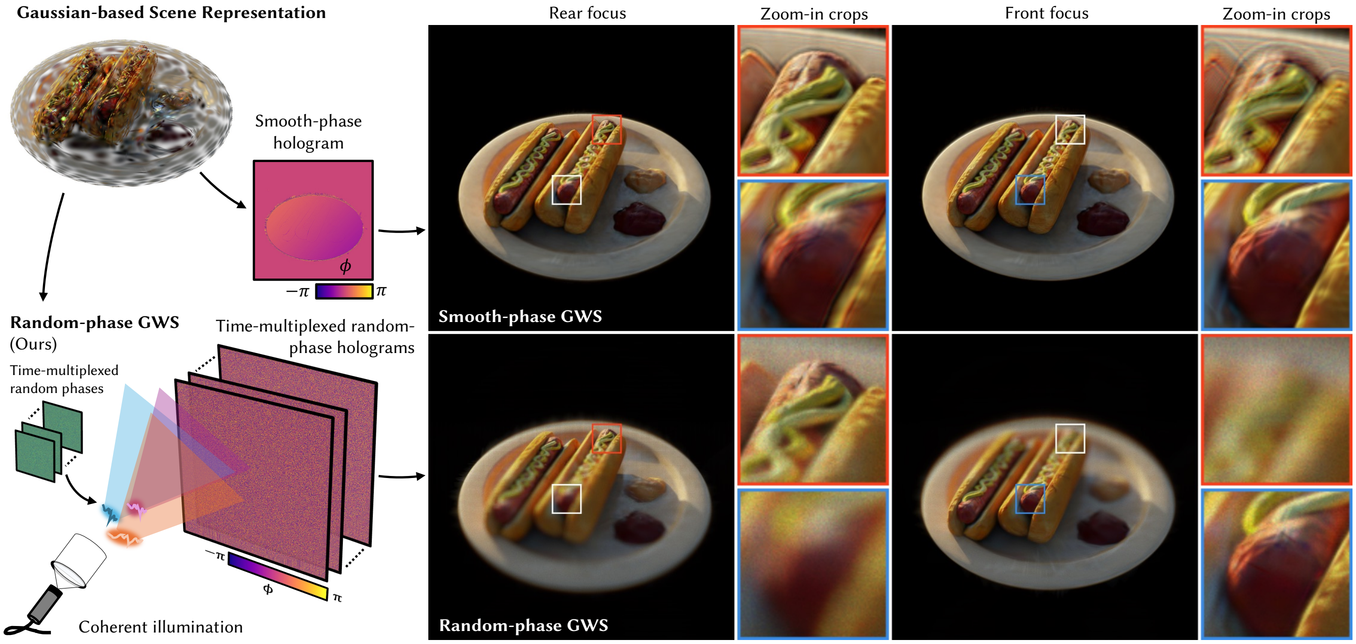

Fig. 1. Random-phase Gaussian Wave Splatting (GWS-RP) takes Gaussian-based scene representations (top left) as input and converts them into holograms (bottom left). Prior Gaussian-based CGH algorithms are unable to accurately reconstruct natural defocus blur due to the smooth-phase nature of the synthesized holograms, resulting in limited blur variation across depths and severe ringing artifacts (top right). Our GWS-RP algorithm, on the other hand, maximally utilizes the bandwidth of the spatial light modulator and reconstructs accurate parallax and defocus blur across the eyebox, as seen in the simulated focal stacks of the GWS-RP holograms on the bottom right.

Holographic near-eye displays offer ultra-compact form factors for virtual and augmented reality systems, but rely on advanced computer-generated holography (CGH) algorithms to convert 3D scenes into interference patterns that can be displayed on spatial light modulators (SLMs). Gaussian Wave Splatting (GWS) has recently emerged as a powerful CGH paradigm that allows for the conversion of Gaussians—a state-of-the-art neural 3D representation—into holograms. However, GWS assumes smooth-phase distributions over the Gaussian primitives, limiting their ability to model view-dependent effects and reconstruct accurate defocus blur, and severely under-utilizing the space–bandwidth product of the SLM. In this work, we propose random-phase GWS (GWS-RP) to improve bandwidth utilization, which has the effect of increasing eyebox size, reconstructing accurate defocus blur and parallax, and supporting time-multiplexed rendering to suppress speckle artifacts. At the core of GWS-RP are (1) a fundamentally new wavefront compositing procedure and (2) an alpha-blending scheme specifically designed for random-phase Gaussian primitives, ensuring physically correct color reconstruction and robust occlusion handling. Additionally, we present the *first formally derived algorithm* for applying random phase to Gaussian primitives, grounded in rigorous statistical optics analysis and validated through practical near-eye display applications. Through extensive simulations and experimental validations, we demonstrate that these advancements, collectively with time-multiplexing, uniquely enables full-bandwith light field CGH that supports accurate accurate parallax and defocus, yielding state-of-the-art image quality and perceptually faithful 3D holograms for next-generation near-eye displays.

---

*denotes equal contribution.

Authors' Contact Information: Brian Chao, brianchc@stanford.edu, Stanford University, USA; Jacqueline Yang, jyang01@stanford.edu, Stanford University, USA; Suyeon Choi, suyeon@stanford.edu, Stanford University, USA; Manu Gopakumar, manugopa@stanford.edu, Stanford University, USA; Ryota Koiso, ryotak@stanford.edu, Stanford University, USA and KDDI Research, Japan; Gordon Wetzstein, gordon.wetzstein@stanford.edu, Stanford University, USA.

1 Introduction

Holographic near-eye displays offer eyeglasses-like device form factors with unprecedented compactness for virtual and augmented reality display systems [Gopakumar et al. 2024; Jang et al. 2024; Kim et al. 2022a; Maimone et al. 2017]. Holographic displays, however, require advanced computer-generated holography (CGH) algorithms to convert a target image or 3D scene into the interference pattern that must be displayed on a spatial light modulator (SLM) [Matsushima 2020; Pi et al. 2022]. Gaussian Wave Splatting (GWS) is an emerging CGH paradigm that unlocks neural rendering [Tewari et al. 2020, 2022] for holographic displays with photorealistic content [Choi et al. 2025]. For this purpose, a few unstructured photographs are distilled into a set of Gaussian primitives [Huang et al. 2024; Kerbl et al. 2023], which are efficiently converted to a hologram using a wave-based CGH algorithm [Choi et al. 2025].

The recently introduced GWS algorithm makes the assumption that the phase distribution over the surfaces of the Gaussian primitives is near constant or *smooth*. This smooth-phase assumption is common in the CGH literature as it typically leads to high image quality (e.g., [Maimone et al. 2017; Shi et al. 2021]), but it is well known that natural reflectance properties of real objects can only be achieved using random phase distributions in holography [Goodman 2015; Matsushima 2020; St-Hilaire 1995]. The use of random phase distributions in coherent optics is analogous to microfacet models used in computer graphics [Pharr et al. 2023]; the geometric structure of surfaces at the micro-scale is described by a specific phase distribution in a hologram. Neglecting to model these phase distributions results in an under-utilization of the inherent space–bandwidth product, or étendue, of an SLM, resulting in a small eyebox providing poor user experiences [Kim et al. 2024, 2022b] and unnatural defocus effects and small parallax, among other artifacts.

Indeed, the discussion of smooth vs. random-phase holograms has been prevalent throughout the history of CGH [Lohmann and Paris 1967]. The use of non-smooth phase distributions has been proposed by many works [Pi et al. 2022], including those using polygon primitives [Matsushima 2020]. Typically, random-phase distributions lead to an increased amount of objectionable speckle, but this can be mitigated using time multiplexing [Choi et al. 2022] or other approaches.

In this work, we propose an algorithm that converts off-the-shelf, pre-optimized Gaussian splats to a set of time-multiplexed random-phase holograms, which we call Random-phase Gaussian Wave Splatting (GWS-RP). We introduce a novel wave compositing procedure and alpha blending algorithm that is specifically designed for random-phase Gaussians for accurate color reconstruction and robust occlusion handling. Additionally, we formally describe a way of applying random phase to Gaussians that allows for arbitrary control over the spatial and angular emission profiles of the resulting random-phase GWS hologram. For the first time, we provide rigorous and extensive proofs grounded in statistical optics that validate the mathematical correctness of the heuristic outlined in the GWS supplemental materials. We demonstrate practical benefits of our algorithmic approach in better bandwidth utilization of the SLM, more natural defocus behavior of the produced holograms, and an increased eyebox size that supports larger parallax, which is preferred by users in practice [Kim et al. 2024] and required to drive their accommodation [Kim et al. 2022b]. These advancements, together with time-multiplexing, uniquely enables accurate light field and focal stack reconstruction, speckle reduction, and offers increased image quality and perceptual realism of the produced 3D holograms.

Specifically, our contributions include:

- Design of an exact wave-splatting and alpha-blending algorithm to composite Gaussian wavefronts *that works exclusively with random-phase Gaussians* (Sec. 3.2).
- *First statistical optics-grounded analysis* of a principled framework to apply random phase and time multiplexing to Gaussians to flexibly control the energy distribution of the resulting complex hologram spectrum (Sec. 3.3).
- Demonstration of the benefits of Random-phase GWS in image quality, natural defocus behavior, and large parallax effects, among others in both simulation and experiments (Sec. 4.2).

Source code and example datasets will be made public.

2 Related Work

Our work builds on a large body of research on CGH algorithms, which we review below. For a more comprehensive overview of holographic displays, we refer the reader to [Chang et al. 2020; Javidi et al. 2021; Park 2017; Pi et al. 2022; Yaraş et al. 2010].

CGH Algorithms. Holograms create a visible image or 3D scene indirectly by displaying an interference pattern, i.e., the hologram, on a 2D amplitude- or phase-only SLM. Methods that convert a target intensity distribution into a hologram are called CGH algorithms. These algorithms have been developed to accommodate a wide variety of input 3D representations, including point clouds [Chen and Wilkinson 2009a; Lucente 1993], meshes [Ahrenberg et al. 2008; Matsushima et al. 2003], wireframes [Blinder et al. 2021], light fields [Choi et al. 2022; Kang et al. 2008; Padmanaban et al. 2019; Park and Askari 2019; Zhang et al. 2019], image layers [Chen and Chu 2015; Shi et al. 2022], and most recently, Gaussians [Choi et al. 2025]. Classic or deep-learning-based *direct* CGH methods typically involve several steps. First, the desired intensity distribution is encoded into a 2D complex-valued wavefront, which typically requires an assumption on the phase distribution of the wavefront. Subsequently, this wavefront is numerically propagated to the SLM plane using wave propagation models such as the angular spectrum method [Goodman 2005; Pellat-Finet 1994]. The resulting complex-valued field is then encoded into a phase- or amplitude-only pattern, depending on the type of SLM used [Maimone et al. 2017; Qi et al. 2016; Tsang and Poon 2013]. In contrast, *iterative* CGH methods make use of iterative optimization to achieve a better image quality, albeit at the cost of increased runtime [Chakravarthula et al. 2019; Fienup 1980; Gerchberg 1972; Peng et al. 2020; Zhang et al. 2017].

Primitive-based CGH Algorithms. Among the various CGH algorithms described above, polygon-based CGH algorithms that use meshes as the input 3D format are closely related to our approach [Askari et al. 2017; Chen and Wilkinson 2009b; Matsushima 2005a,b; Matsushima and Nakahara 2009; Matsushima et al. 2014;

Matsushima and Sonobe 2018]. These methods, in principle, support view-dependent effects such as parallax through the use of random phase maps applied to the primitives. However, the wavefront of polygon primitives varies with their geometry, requiring substantial memory for look-up table approaches and suffers from numerical stability issues. Furthermore, these methods struggle to render highly detailed appearances unless an excessive number of tiny triangles are used as pointed out in prior work [Choi et al. 2025]. Gaussian Wave Splatting [Choi et al. 2025], on the other hand, has emerged as a new primitives-based CGH algorithm that converts Gaussian-based scene representations [Huang et al. 2024; Kerbl et al. 2023] to complex holograms, greatly outperforming traditional polygon-based CGH in terms of image quality and computing time due to the compact Gaussians representation compared to meshes and the numerical stability of Gaussians. However, GWS does not support view-dependent effects and natural defocus due to its smooth-phase nature. Additionally, the alpha-blending and wave-splatting procedure described in GWS does not naturally extend to random-phase Gaussians.

In this work, we show that for random-phase Gaussian wavefronts, alpha blending is exact in expectation in the intensity domain, in contrast to the amplitude domain for smooth-phase Gaussians. We then draw inspiration from polygon-based CGH methods and devise a novel wavefront composition and alpha blending technique that specifically work with random-phase Gaussians (Sec. 3.2) for accurate color reconstruction and occlusion handling. As such, our algorithm generates holograms with high in-focus image quality using fewer primitives by fully leveraging the compact Gaussian representation while synthesizing accurate defocus blur and wide parallax via random phase and time-multiplexing.

Phase Distributions of Holograms. Although phase is not directly observable, the phase profile of a wavefront plays a crucial role in determining the spatio-angular behavior of the observable light field [Chakravarthula et al. 2022; Kim et al. 2022b; Schiffers et al. 2023; St-Hilaire 1995]. For this reason, two popular heuristics have been developed that are widely used in CGH literature: *smooth-phase* and *random-phase* holograms [Maimone et al. 2017; Yoo et al. 2021]. Smooth-phase, sometimes called random-phase-free, holograms [Shimobaba and Ito 2015] achieve high image quality that can be demonstrated with relatively simple experimental setups [Choi et al. 2025; Maimone et al. 2017; Peng et al. 2020; Shi et al. 2021]. The main drawback of smooth-phase distributions, however, is that they concentrate energy in the low frequencies of the angular spectrum, resulting in a severely restricted eye box size, limited defocus effects, and increased sensitivity to pupil position. These effects limit the perceptual realism and overall user experience of the produced holograms [Kim et al. 2024] as well as the support for perceptually important focus cues [Kim et al. 2022b].

On the other hand, random-phase holograms are capable of reconstructing larger parallax and natural defocus blur, which is necessary for a perceptually-realistic and natural viewing experience [Kim et al. 2024]. However, prior random-phase CGH algorithms can only mimic the behavior of Lambertian, i.e. fully diffuse, object reflectance [Amako et al. 1995; Lohmann and Paris 1967; Matsushima 2020] and are unable to achieve arbitrary angular emission profiles. Furthermore, rapid phase variations between adjacent pixels introduce unwanted speckle noise created by constructive and destructive interference [Goodman 2007]. Speckle reduction techniques often utilize some form of partial coherence, introduced by partially coherent or multiple coherent light sources, or by time multiplexing [Chao et al. 2024; Choi et al. 2022; Curtis et al. 2021; Kuo et al. 2023; Lee et al. 2020; Peng et al. 2021].

In this work, we formally introduce a principled way of controlling the phase randomness of individual Gaussian primitive wavefronts by convolving the spectrum of the wavefronts with a random spectral kernel (Sec. 3.3), a procedure first introduced as a experimental heuristic sketch in the supplemental materials of Gaussian Wave Splatting [Choi et al. 2025]. This differs from prior work, where random phase is applied to wavefronts in the spatial domain, which in turn can only model fully diffuse reflectance properties. The introduction of random phase to GWS greatly improves its defocus blur quality and supports view-dependent effects. Additionally, we can assign arbitrary Fourier amplitude distributions, or angular emission profiles, to the synthesized holograms, introducing benefits such as direct control over depth of field. Finally, we for the first time prove that this algorithm is exact using rigorous statistical optics analysis, providing the mathematical foundation for the re-emerging field of primitives-based CGH.

3 Gaussian-based Computer-generated Holography

Our approach takes as input a set of multi-view images that are turned into 2D Gaussians representing a 3D scene. These Gaussians are then converted to a hologram that can optionally be time multiplexed. We briefly review the relevant background on Gaussian-based scene representations and existing work on *smooth-phase* Gaussian wave splatting, before introducing our *random-phase* CGH algorithm, which uniquely enables a time-multiplexed image formation. Moreover, we derive bounds on the achieved bandwidth utilization and defocus behavior to formalize the benefits of our approach.

3.1 Background

3.1.1 Gaussian-based Scene Representations. 2D Gaussians have been established as a common scene representation as they can be distilled from a few unstructured photographs or renderings of a scene [Huang et al. 2024]. Each of these $i = 1 \dots N$ Gaussians is described by its mean $\boldsymbol{\mu}_i \in \mathbb{R}^3$, 3D covariance $\boldsymbol{\Sigma}_i = \mathbf{R}_i\mathbf{S}_i\mathbf{S}_i^\top\mathbf{R}_i^\top \in \mathbb{R}^{3\times3}$ that can be factorized into a rotation matrix $R_i \in \mathbb{R}^{3\times3}$ and a scaling matrix $S_i \in \mathbb{R}^{3\times3}$, opacity $o_i \in \mathbb{R}$, and color $c_i \in \mathbb{R}$ (for a single color channel).

Any Gaussian splatting approach requires the N Gaussians representing a scene to be depth sorted from *front to back* based on the z value with respect to the camera position, or in a holographic display setup, the SLM plane. We closely follow the *holographics* pipeline described in GWS [Choi et al. 2025] that transforms these primitives into an adequate hologram space for CGH calculation.

3.1.2 Gaussian Wave Splatting. Gaussian Wave Splatting (GWS) [Choi et al. 2025] is a CGH method capable of computing holograms that accurately represent 3D scenes from collections of 2D Gaussians [Huang et al. 2024] extracted from any off-the-shelf optimized

2DGS models, e.g., models optimized using the `gsplat` library [Ye et al. 2024]. GWS first analytically determines the spectrum of each Gaussian described by its mean $\boldsymbol{\mu}_i$ and 3D covariance Σ_i and computes the wavefront $u_i(\mathbf{x}) = a_i(\mathbf{x})e^{ikz_i}$, where $z_i = (\boldsymbol{\mu}_i)_z$ is the Gaussian object depth and $k = \frac{2\pi}{\lambda}$. Then, each wavefront $u_i(\mathbf{x})$ is propagated using the angular spectrum propagation operator $\mathcal{P}(\cdot; z)$ [Goodman 2005; Matsushima and Shimobaba 2009] and alpha blended from *front to back* using the opacity o_i and color c_i associated with each Gaussian to get the final composited wavefront profile at the SLM, given by Eqs. 1 and 2. We refer to this process as *alpha wave blending*:

$$u_{\text{SLM}}(\mathbf{x}) = \sum_{i}^{N} \mathcal{P}\Big(c_i o_i |u_i(\mathbf{x})| \mathcal{T}_i(\mathbf{x})\, e^{ikz_i}; -z_i\Big), \tag{1}$$

$$\mathcal{T}_i(\mathbf{x}) = \prod_{j=1}^{i-1} (1 - o_j |u_j(\mathbf{x})|). \tag{2}$$

GWS inherits the ability of Gaussian splatting to seamlessly merge large numbers of Gaussians for high-quality reconstruction. GWS further collapses to the original ray-based alpha blending [Huang et al. 2024; Kerbl et al. 2023] and volume rendering [Kajiya and Von Herzen 1984] if we ignore the wave propagation operator and match the phase of all wavefronts at all depths (i.e., $\angle u_i = kz_i$) such that the composited wavefront at the SLM plane achieves a smooth or near-constant phase profile. Therefore, this formulation of GWS inherently generates “smooth-phase” holograms. Although GWS has demonstrated the potential to recreate sharp details with photorealistic image quality [Choi et al. 2025], prior works in holography [Choi et al. 2022; Kim et al. 2022b; Lee et al. 2022; Schiffers et al. 2023; Shi et al. 2024] have pointed out that smooth-phase holograms are undesirable due to their poor SLM bandwith utilization, unnatural defocus blur, large depth of field (i.e., small blur variation across different depths), and floater artifacts. In the following sections we refer to GWS as smooth-phase GWS, or GWS-SP.

3.2 Random-phase Gaussian Wave Splatting

To maximize SLM bandwidth utilization, we propose *Random-phase Gaussian Wave Splatting* (GWS-RP) where each Gaussian wavefront is effectively modulated by a random phase map $\phi_i(\mathbf{x})$ to scatter the light passing through each Gaussian in the scene away from the optical axis. Gaussians are then alpha blended and composited from *back to front* with respect to the SLM, and multiple such wavefronts with individually sampled random phases are time multiplexed to achieve the desired intensity distribution. The novel alpha blending and wave compositing procedure in GWS-RP is specifically designed to work with random-phase Gaussians with non-binary alpha values, and greatly outperforms prior alpha blending schemes for complex wavefronts [Choi et al. 2025; Yanagiya and Matsushima 2019] as we demonstrate in Section 4.2.

In Eq. 3, let $g_i(\mathbf{x})$ denote the *back-to-front* composited wavefront starting from the N^{th} primitive up to the i^{th} primitive and $M_i(\mathbf{x})$ be the transmittance mask of the i^{th} primitive. The wavefront at the $i-1^{\text{th}}$ parallel plane where the next primitive is located, which is a propagation distance $\Delta z = z_{i-1} - z_i$ away, is given by:

$$g_{i-1}^{(t)}(\mathbf{x}) = \mathcal{P}\Big(M_i(\mathbf{x})\; g_i(\mathbf{x}) + \sqrt{c_i}\sqrt{o_i|u_i(\mathbf{x})|}\;\, e^{i\angle u_i(\mathbf{x}) + \phi_i^{(t)}(\mathbf{x})}; \Delta z\Big), \tag{3}$$

$$M_i(\mathbf{x}) = \sqrt{(1 - o_i|u_i(\mathbf{x})|)}, \tag{4}$$

where $t, 1 \leq t \leq T$ is the index of the time-multiplexed frame, T is the total number of multiplexed frames, and $\phi_i^{(t)}$ is the sampled random phase for the i^{th} Gaussian at the t^{th} frame. The final composited wavefront at the SLM plane that is located at $z_0 = 0$ is simply defined by $u_{\text{SLM}}^{(t)}(\mathbf{x}) = g_0^{(t)}(\mathbf{x})$.

Note that for GWS-RP, we use Gaussian splats optimized to reconstruct the *square* (i.e., intensity) of the target scene, instead of the amplitude like in GWS-SP, hence the square root of the color in Eq. 3. This is because alpha blending is exact in the intensity domain for random-phase wavefronts, instead of in the amplitude domain for smooth-phase wavefronts. Please refer to the supplemental materials for a formal mathematical analysis of why this subtle distinction significantly affects the color of the rendered image and the accuracy of occlusion handling.

The intensity of reconstructed images of the time-multiplexed hologram flexibly described by the operator $O(\cdot;\cdot)$ can therefore be described by:

$$I(\mathbf{x}) = \frac{1}{T}\sum_{t=1}^{T} \left| O\left(u_{\text{SLM}}^{(t)}(\mathbf{x}); \cdot\right)\right|^2, \tag{5}$$

where the operator $O(\cdot;\cdot)$ could describe a single propagation that reconstructs a 2D image, multiple propagations that reconstructs a 3D focal stack, or the Short-time Fourier Transform (STFT) that reconstructs a 4D light field [Choi et al. 2022].

The approach of multiplying the wavefronts of individual primitives directly by random phase maps $\phi_i^{(t)}$ (Eq. 3) is widely used in prior works. Albeit straightforward, this is a heuristic and does not allow for flexible control over the Fourier amplitude distribution of the final synthesized hologram. Instead, we leverage an exact, principled formulation of applying random phase to Gaussians, which we call *structured random phase*, and describe in detail in Section 3.3.

With this random phase modulation, the bandwidth of the SLM can be maximally utilized to achieve several benefits including large eyebox and parallax, shallow depth of field (i.e., large blur variation across different depths), natural defocus blur, as we will illustrate in the following sections.

3.2.1 Bandwidth Utilization. Consider the wavefront of a single Gaussian $u(\mathbf{x}) = a(\mathbf{x})e^{i\phi(\mathbf{x})}, \mathbf{x} \in \mathbb{R}^2$. For smooth-phase Gaussians, we assume that $\phi(\mathbf{x}) = \text{const}$. For random-phase Gaussians, we assume that $\phi^i(\mathbf{x}) \overset{\text{iid}}{\sim} \mathcal{U}(-\pi, \pi)$ over all spatial locations i. To formally quantify the bandwidth characteristics of $u(\mathbf{x})$, we can analyze the *power spectral density*, or the intensity of the Fourier transform of the wavefront $\hat{u}(\mathbf{k}), \mathbf{k} \in \mathbb{R}^2$, represented by $|\hat{u}(\mathbf{k})|^2 = \hat{u}(\mathbf{k})\hat{u}^*(\mathbf{k})$.

For Gaussians with constant phase, we can show from theorems in statistical optics [Goodman 2015] that the power spectral density is a delta function with a small spatial support $|\hat{u}(\mathbf{k})|^2 \propto \delta(\mathbf{k})$, leading

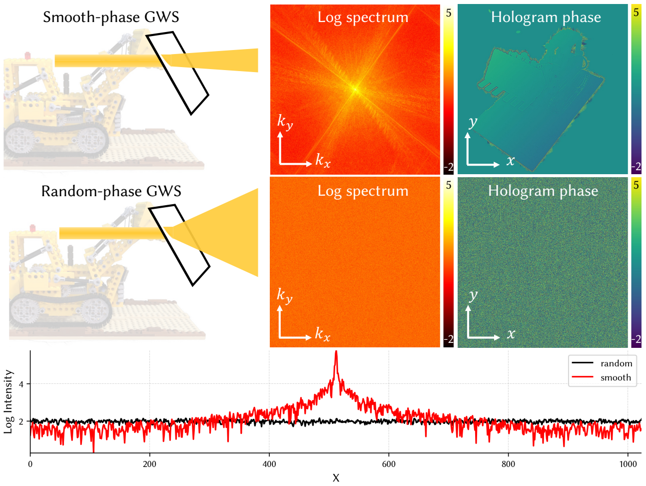

Fig. 2. **Spectral analysis of smooth-phase and random-phase GWS holograms.** We show the log spectrum and the phase of the complex-valued smooth-phase and random-phase GWS holograms. Smooth-phase GWS diffracts light into small angles, leading to a smooth phase profile and highly concentrated spectrum. Random-phase GWS, on the other hand, diffracts light into a much wider cone of angles, resulting in a random phase profile and a flat spectrum with energy evenly distributed across the eyebox, as shown in the cross section figure. Random-phase GWS therefore naturally supports larger and accurate parallax across the eyebox and high-quality defocus blur if paired with time multiplexing.

to a small eyebox, while for random-phase Gaussians the *expected value* of the power spectral density is flat with a large spatial support $\mathbb{E}\left[|\hat{u}(\mathbf{k})|^2\right] = \text{const.} \in \mathbb{R}$, leading to a large eyebox. This result is demonstrated with simulated experiments in Fig. 2.

This power spectral density analysis suggests that (time-multiplexed) random-phase holograms maximally utilize the SLM bandwidth while constant or smooth-phase holograms severely underutilize it. Random-phase holograms generated using GWS-RP can therefore naturally support a large eyebox and parallax, shallow depth of field, and natural defocus blur, as we show in the next section.

3.2.2 Defocus Blur. Random-phase holograms naturally support shallow depths of field due to the larger spatial spread across depths compared to smooth-phase holograms. This allows the hologram to create larger blur variations across different depths, which provides a depth cue crucial to human vision. We illustrate this phenomenon with the wavefront of a single Gaussian propagating to different depths shown in Fig. 3. With additional time-multiplexing to reduce speckles, the defocus blur of random phase holograms at different depths closely emulates incoherent defocus blur.

To formally quantify the spatial spread of a wavefront $u(\mathbf{x})$ propagated to a specific depth $z \in \mathbb{R}$, we can analyze the *variance of the intensity* of the propagated field $\mathcal{P}(u(\mathbf{x}), z)$:

$$\sigma^2(z) = \iint_{\mathbb{R}^2} |\mathbf{x}|^2 \, |\mathcal{P}(u(\mathbf{x}), z)|^2 \, d\mathbf{x}. \tag{6}$$

Since variance by definition quantifies how spread out a distribution is, the variance of the intensity of the propagated wavefront naturally characterizes the spatial spread of the wavefront.

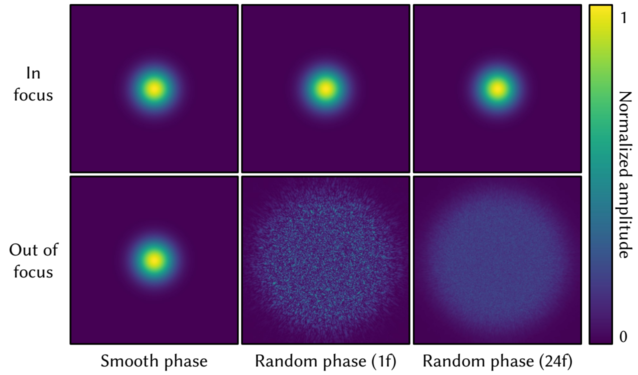

Fig. 3. **Depth of field comparisons of smooth-phase and random-phase Gaussians.** We simulate the propagation of a single Gaussian primitive with different phase profiles to various distances and show the resulting amplitude after propagation. For smooth- or constant-phase Gaussians, their amplitude profile does not change after propagation since light is traveling in approximately the same direction, resulting in a large depth of field. The amplitude profiles of random-phase Gaussians spread out as they propagate, leading to larger blur and shallower depth of field. With additional time multiplexing, the speckles in the propagated field are greatly reduced and can thus be used to correctly synthesize incoherent blur to support natural accommodation cues.

Eq. 6 can be further rewritten as:

$$\sigma^2(z) \approx \iint_{\mathbb{R}^2} \mathbf{x}^2 |u(\mathbf{x})|^2 \, d\mathbf{x} + \left(\frac{z}{k}\right)^2 \cdot \iint_{\mathbb{R}^2} |\mathbf{k}|^2 \, |\widehat{u}(\mathbf{k})|^2 \, d\mathbf{k}, \tag{7}$$

where the first term is the variance of the initial intensity field (which are the same for smooth-phase and random-phase wavefronts that have the same amplitude) and the second term includes the variance of the power spectral density (defined in 3.2.1) that quantifies the angular spread. Note that the second term increases quadratically as propagation distance z increases. Please refer to the supplemental material for a detailed derivation of Eq. 7.

Since the variance of the power spectral density of a random phase wavefront is much larger than that of a smooth-phase wavefront, random-phase holograms exhibit much larger spatial spread across depths. This property allows random-phase holograms to accurately create different levels of defocus blur, which is a crucial depth cue for human vision.

3.3 Structured Random Phase with Partial Coherence

Prior works have pointed out that a single coherent wavefront can only represent light fields with rank-1 mutual intensity [Choi et al. 2022; Hamann et al. 2018; Zhang 2011]. Due to this constraint, a single coherent wavefront modulated by a band-limited SLM cannot simultaneously achieve both the desired spatial and angular amplitude distributions. However, it has been demonstrated that through partial coherence in the form of time multiplexing, high-quality light-field holograms of higher ranks can be reconstructed [Choi et al. 2022; Kim et al. 2024]. Inspired by these works, we discuss a principled way of applying random phase to Gaussians that allows for arbitrary control over the spatial and angular emission profiles of the resulting random-phase GWS hologram, which we call *structured random phase with partial coherence.* Specifically, we

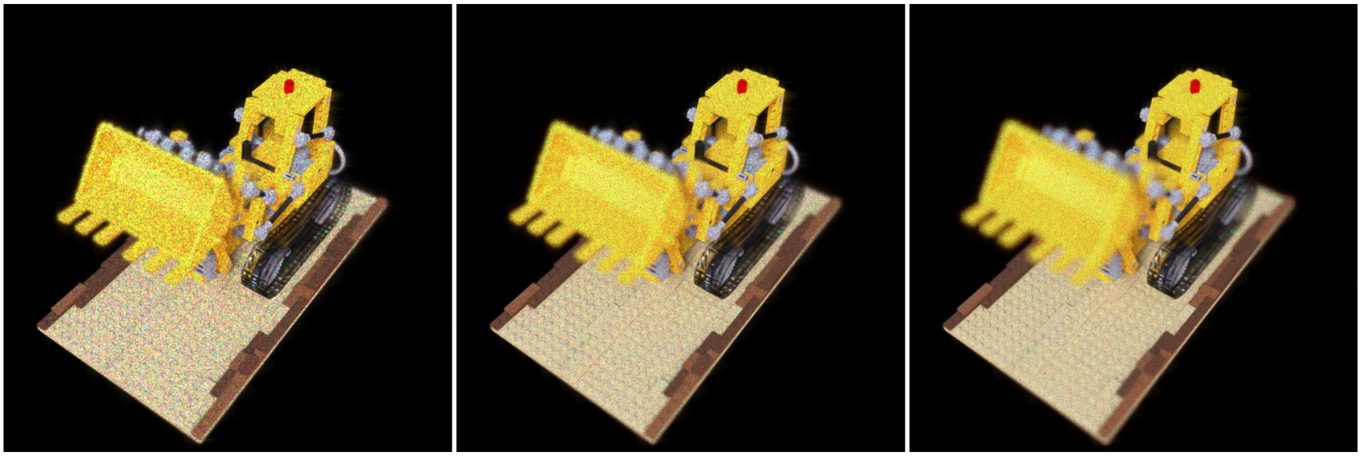

$r = 0.25$ $r = 0.5$ $r = 1.0$

Fig. 4. **Programmatically controlling depth of field of GWS-RP holograms using structured random phase.** By setting different r values in the pupil function $P(\mathbf{k}, r)$, we can programmatically control the depth of field of the synthesized GWS-RP holograms. With smaller r, the pupil is smaller and filters out high frequency components in the hologram spectrum, resulting in small blur and large depth of field. With large r, the hologram exhibits shallow depth of field and large blur.

Algorithm	Previous	Ours & Choi et al. [2025]
Modulated wavefront	$u(\mathbf{x}) \cdot e^{i\phi(\mathbf{x})}$	$u(\mathbf{x}) \cdot \mathcal{F}^{-1}\{Q(\mathbf{k})e^{i\phi(\mathbf{k})}\}$
Spatial amplitude (TM)	$\lvert u(\mathbf{x})\rvert$	$\lvert u(\mathbf{x})\rvert$
Fourier amplitude	**constant.**	$Q(\mathbf{k})$

Table 1. **Comparison of random phase modulation algorithms.** Here, we compare the spatial and Fourier amplitude of the random phase-modulated wavefronts using different algorithms. Previous random-phase modulation algorithims apply random-phase in the spatial domain, resulting in a constant Fourier amplitude that can only represent diffuse surfaces, significantly limiting the angular expressivity of individual primitive wavefronts. Conversely, we convolve the *spectrum* of the Gaussian wavefront with random phase kernels in *structured random phase*, which preserves the amplitude distribution $|u(\mathbf{x})|$ (expected value, or time-multiplexed (TM)) of the Gaussians in the spatial domain while achieving arbitrary Fourier amplitude distributions or angular emission profiles $Q(\mathbf{k})$ in the Fourier domain.

convolve the angular spectrum of each Gaussian primitive with a complex spectrum $Q(\mathbf{k})e^{j\phi^{(t)}(\mathbf{k})}$ and time multiplex multiple such wavefronts $1 \leq t \leq T$, to achieve the desired angular emission profile $Q(\mathbf{k})$ in the *Fourier domain* through partial coherence, as described in Table 1.

The phase profile $\phi(\mathbf{k})$ of the angular emission kernel $Q(\mathbf{k})e^{j\phi(\mathbf{k})}$, on the other hand, plays a crucial role in correctly reconstructing the Gaussian amplitude in the *spatial domain*. We prove that this can be achieved through sampling multiple phases for $\phi(\mathbf{k})$ from the uniform distribution $\mathcal{U}(-\pi, \pi)$. Through time multiplexing multiple such wavefronts, the expected value of the amplitude of the inverse Fourier transform of $Q(\mathbf{k})e^{j\phi(\mathbf{k})}$ is a constant in the spatial domain, allowing us to assign arbitrary amplitude distributions $Q(\mathbf{k})$ to the angular emission kernel in the Fourier domain while retaining the correct Gaussian amplitude profile of each Gaussian primitive in the spatial domain. Although this algorithm was briefly mentioned as a heuristic sketch in the supplemental materials of prior relevant work Gaussian Wave Splatting [Choi et al. 2025], we for the first time formally prove that this algorithm is exact based on rigorous derivations grounded in statistical optics. Please refer to the supplemental materials for the detailed mathematical derivations and analysis regarding the structured random phase algorithm.

One example of the amplitude of the angular emission profile includes the use of spherical harmonics $Q(\mathbf{k}) = Y_m^l(\mathbf{k})$ to model specular colors, where $Y_m^l(\mathbf{k})$ is the is the spherical harmonics function Y_l^m of degree l and order m evaluated at the normalized direction $\hat{k} = \mathbf{k}/\|\mathbf{k}\|$ of the wave vector $\mathbf{k}$. Another example of $Q(\mathbf{k})$ is $Q(\mathbf{k}) = \mathbf{1}$, representing fully diffuse primitives where the angular emission profile is uniform across all angles, or the circular binary pupil function $Q(\mathbf{k}) = P(\mathbf{k}, r)$:

$$P(\mathbf{k}, r) = \begin{cases} 1, & \text{if } \|\mathbf{k}\| \leq r \\ 0, & \text{otherwise} \end{cases}, \tag{8}$$

where $P(\mathbf{k}, r)$ filters out higher frequency components in the Fourier domain, which is useful for programmatically controlling the depth of field of the reconstructed focal stack, which we demonstrate in Fig. 4.

4 Experiments

4.1 Implementation Details

4.1.1 Datasets and 3D Scene Representations. We generate 3D holograms of selected scenes from the synthetic Blender dataset [Mildenhall et al. 2020] and the MipNeRF-360 dataset [Barron et al. 2022] using different 3D scene representations, including 2D Gaussians and textured meshes. We follow GWS [Choi et al. 2025] to prepare our 3D scene representations. We use the open-source Gaussian splatting software library `gsplat` [Ye et al. 2024] to optimize 2DGS models. For textured meshes, we use the optimized textured mesh models released by the NeRF2Mesh [Tang et al. 2023]. Please refer to the supplemental materials for more data preparation details.

4.1.2 Algorithm Implementations. We implement GWS-RP and all baseline methods in PyTorch. Please refer to the supplementary materials for pseudocode and implementation details.

4.1.3 Experimental Setup. GWS-RP generates complex-valued holograms. We perform SGD [Choi et al. 2022; Peng et al. 2020] with a complex field supervision loss [Chen et al. 2021] to synthesize phase-only holograms that accurately reconstruct the target complex wavefront at a fixed propagation distance (4cm) from the phase-only SLM. We use a 1080p HOLOEYE Pluto-2.1 phase-only SLM with $8\mu m$ pixel pitch and a FISBA READYBeam fiber-coupled laser module as the illumination source. We use a FLIR GS3-U3-123S6C-C color camera placed on a Thorlabs MTS25-Z8 25mm translation stage to capture focal slices at different distances from the SLM.

We refer readers to the supplementary materials for more details on phase-only optimization for complex valued GWS-RP holograms and the experimental setup.

4.2 Simulation and Experimental Results

4.2.1 Baseline Comparisons with Simulation Results. We compare GWS-RP with several primitive-based CGH baselines, including random-phase polygon-based CGH (Polygons-RP) using the silhouette method [Matsushima 2005b; Matsushima and Nakahara 2009; Matsushima et al. 2003] and its time-multiplexed variants (1 frame, 8 frames, and 24 frames) as well as Gaussian Wave Splatting (GWS) [Choi et al. 2025].

Fig. 6 shows the simulated focal stack comparisons of the synthesized complex holograms. Since GWS generates smooth-phase holograms, there is little variation in blur size across depths and the blur shape is unnatural with ringing artifacts (coherent blur), leading to incorrect retinal blur cues for a human observer. Single-frame random phase holograms (Polygon-RP and GWS-RP) suffer from prominent speckle artifacts but image quality improves as time multiplexing increases to 8 and 24 frames. Random-phase polygon-based CGH generates much more natural blur, but the in-focus image quality is inherently limited by the underlying low-quality per-face textured mesh 3D scene representation. GWS-RP achieves the most natural defocus blur that accurately emulates incoherent blur while achieving similar image quality in in-focus regions compared to GWS.

Fig. 7 shows the simulated light-field comparisons of the synthesized complex holograms in terms of individual light-field views and epipolar images. Since the energy of the wavefront is greatly concentrated around the center of the eyebox for smooth-phase GWS, light-field views at the periphery of the eyebox are severely degraded. For GWS-RP, the energy of the wavefront is evenly spread out across the eyebox, which results in accurate light-field view reconstruction at all pupil locations within the eyebox. The image quality of random-phase polygon-based CGH is again limited by the coarse per-face textured mesh 3D representation. We additionally show quantitative comparisons of dense 10×10 light-field reconstruction performance in terms of various standard image quality metrics between different CGH methods in Table 2.

CGH Algorithm	Blender	Mip-NeRF 360
Polygons-RP (1 frame)	17.02 / 0.14 / 0.75	12.77 / 0.14 / 0.76
Polygons-RP (8 frames)	20.27 / 0.30 / 0.65	15.63 / 0.27 / 0.69
GWS-RP (1 frame)	18.11 / 0.15 / 0.76	14.17 / 0.17 / 0.76
GWS-RP (8 frames)	23.42 / 0.36 / 0.64	19.65 / 0.36 / 0.65
GWS-RP (24 frames)	**24.87** / **0.49** / **0.56**	**21.21** / **0.48** / **0.58**
GWS-SP [Choi et al. 2025]	14.35 / 0.18 / 0.63	10.50 / 0.17 / 0.60

Table 2. **Quantitative light-field reconstruction performance of different CGH algorithms.** We evaluate the image quality of 10×10 dense light fields reconstructed from the simulated holograms generated using different CGH methods in terms of PSNR (↑) / SSIM (↑) / LPIPS (↓). The best performing metrics for all CGH baselines are boldfaced. Our method achieves the best image light-field reconstruction performance and eyebox uniformity among all CGH baselines.

4.2.2 Time-multiplexing Discussions. While time-multiplexing reduces the effective frame rate of our holographoc display, our framework supports both single- and multi-frame modes, where we validate our method up to 24 frames time-multiplexing. The multi-frame results merely represent an upper bound, and ablations in Table 2 show exceptional performance even with fewer frames (single frame, 8 frames) compared to smooth-phase GWS and polygon-based CGH baseline methods. GWS-RP can continue to benefit from the rapid advanvements in next-generation SLM technologies, such as the 3600 fps FLCoS used in Holograhic Parallax [Kim et al. 2024] and recent developements in the 5760 fps Texus Instruments MEMS-based PLM.

4.2.3 Experimentally Captured Focal Stack Results. We experimentally capture 3D focal stacks of GWS-RP and other CGH baselines encoded using SGD [Peng et al. 2020] with complex wavefront supervision to demonstrate real-world refocusing capabilities of the synthesized phase-only holograms. Individial frames and color channels are captured independently and merged in post-processing.

The experimentally captured focal stacks match the simulation results well, as shown in Fig. 8. GWS-RP achieves the best balance between high in-focus image quality and natural defocus blur. Smooth-phase GWS achieves good in-focus image quality but suffers from ringing coherent blur artifacts and limited blur size in defocus regions. The in-focus image quality of polygon-based CGH is inherently limited by the per-face colored mesh 3D representation.

4.2.4 Ablation Study on Alpha Blending Operations. Several methods have been proposed to address occlusion handling of translucent primitives in CGH. We compare several of these methods in Fig. 5. Yanagiya and Matsushima [2019] proposed an extended traditional silhouette-based methods to handle primitives with non-binary opacity values. However, this method was not specifically designed for Gaussian splats, resulting in poor reconstruction quality for even in-focus regions. GWS [Choi et al. 2025] extended alpha blending in ray-based Gaussian splats rendering [Kerbl et al. 2023] to work with coherent and smooth-phase Gaussian wavefronts, where they refer to it as *alpha wave blending*. This yields somewhat plausible results when applied to random-phase Gaussians. However, the resulting focal stacks exhibit prominent dark halo artifacts at depth discontinuities. In light field reconstruction, this manifests as dark borders

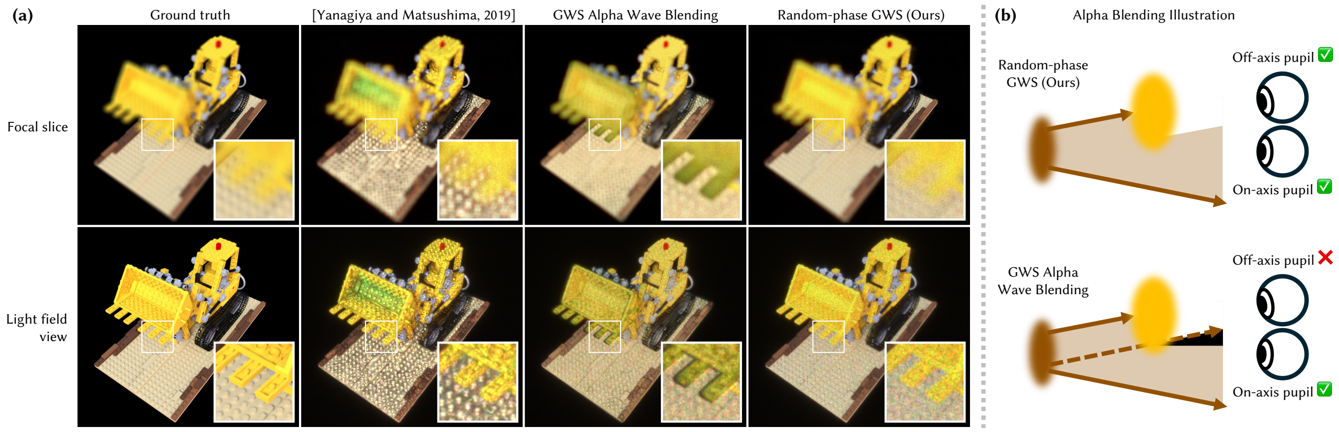

Fig. 5. **Ablation study on different alpha-blending methods.** We compare our method with two other algorithms that perform alpha blending on wavefronts. The algorithm proposed by Yanagiya and Matsushima [2019] simply does not work with optimized Gaussian splats, resulting in incorrectly reconstructed images. Alpha wave blending in GWS-SP [Choi et al. 2025] reconstructs plausible results, but suffers from dark halo and edge artifacts at depth discontinuities and occlusion borders due to incorrect occlusion handling for random-phase Gaussians. From Fig. (b), we see that light rays in the dark region cannot reach the off-axis pupils since light is assumed to be only traveling on-axis given the smooth-phase assumption and the way the transmittance maps are calculated in GWS-SP. Our algorithm reconstructs the most accurate focusing and occlusion behavior.

around foreground occluders. These artifacts arise from blurred alpha map edges across depth layers that lead to imprecise blending, and from incorrect occlusion of off-axis diffracted light. GWS-SP assumes purely on-axis light propagation from the smooth-phase Gaussians and thus the transmittance maps are only considered for the on-axis pupil, resulting in its occlusion model erroneously blocking light for random-phase Gaussians where light travels at steeper angles, as shown in Fig. 5 (b).

Our proposed alpha-blending formulation together with a back-to-front compositing procedure, achieves the most accurate focal stack reconstruction and parallax rendering results, completely eliminating artifacts at depth discontinuities and occlusion borders while accurately reconstructing image content.

4.2.5 Comparisons with STFT-based Light Field CGH. Despite belonging to fundamentally different algorithmic classes, we delineate the key distinctions between GWS-RP and STFT-based light field CGH methods [Choi et al. 2022; Kim et al. 2024]. STFT-based CGH methods require densely sampled per-view light field targets (therefore all prior STFT-based light field CGH typically use synthetic datasets), per-view iterative optimization, and suffer from spatio-angular tradeoffs and poor generalizability in intermediate views due to the window function used in STFT. In contrast, our direct, non-iterative method operates directly on *pre-optimized, off-the-shelf* Gaussian splats, enabling arbitrary-view rendering without re-optimization. While we do not claim superiority, GWS-RP addresses a fundamentally different CGH regime that is optimized for real-world capture and forward compatibility with state-of-the-art Gaussians-based 3D representations, benefiting from any ongoing advances in Gaussian splatting.

5 Discussion

Limitations and Future Work. In GWS-RP, partial coherence in the form of time multiplexing is used to reduce noise introduced by random phase, yet many time-multiplexed frames are required to reconstruct a noise-free focal stack. In practice, time-multiplexing capabilities are limited by the speed of the SLM. Therefore, optimal phase distributions that reduce the number of time-multiplexed frames while providing noise-free images is an interesting future direction. Dealing with quantization in fast SLMs would also be an interesting future direction. Due to the use random phase, our experimental results suffer from reduced contrast, but learning-based model calibration techniques [Choi et al. 2022; Peng et al. 2020] can potentially be used to further improve the image quality. Extending these models to directly reconstruct complex fields [Jang et al. 2024] could be useful for our complex GWS-RP holograms. Our implementation of the GWS-RP algorithm does not run in real time. Although hologram synthesis using deep neural networks have been an ongoing research topic [Peng et al. 2020; Shi et al. 2021], these designs are often restricted to synthesize smooth-phase holograms. Designing neural networks that can directly synthesize *random-phase holograms* would be immensely useful.

Conclusion. Random-phase Gaussian Wave Splatting synthesizes photorealistic holograms from state-of-the-art Gaussian-based scene representations and reconstructs accurate defocus blur and large parallax, unlocking the full potential of next-generation holographic displays for perceptually realistic and comfortable VR experiences.

References

Lukas Ahrenberg, Philip Benzie, Marcus Magnor, and John Watson. 2008. Computer generated holograms from three dimensional meshes using an analytic light transport model. *Applied optics* 47, 10 (2008), 1567–1574.

Jun Amako, Hirotsuna Miura, and Tomio Sonehara. 1995. Speckle-noise reduction on kinoform reconstruction using a phase-only spatial light modulator. *Applied optics* 34, 17 (1995), 3165–3171.

Mehdi Askari, Seong-Bok Kim, Kwang-Soo Shin, Seok-Bum Ko, Sang-Hoo Kim, Dae-Youl Park, Yeon-Gyeong Ju, and Jae-Hyeung Park. 2017. Occlusion handling using angular spectrum convolution in fully analytical mesh based computer generated hologram. *Optics Express* 25, 21 (2017), 25867–25878.

Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. 2022. Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields. In *CVPR*.

David Blinder, Takashi Nishitsuji, and Peter Schelkens. 2021. Real-time computation of 3D wireframes in computer-generated holography. *IEEE Transactions on Image Processing* 30 (2021), 9418–9428.

Praneeth Chakravarthula, Seung-Hwan Baek, Florian Schiffers, Ethan Tseng, Grace Kuo, Andrew Maimone, Nathan Matsuda, Oliver Cossairt, Douglas Lanman, and Felix Heide. 2022. Pupil-aware Holography. *ACM Transactions on Graphics (TOG)* 41, 6, Article 212 (2022).

Praneeth Chakravarthula, Yifan Peng, Joel Kollin, Henry Fuchs, and Felix Heide. 2019. Wirtinger holography for near-eye displays. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–13.

Chenliang Chang, Kiseung Bang, Gordon Wetzstein, Byoungho Lee, and Liang Gao. 2020. Toward the next-generation VR/AR optics: a review of holographic near-eye displays from a human-centric perspective. *Optica* 7, 11 (2020), 1563–1578.

Brian Chao, Manu Gopakumar, Suyeon Choi, Jonghyun Kim, Liang Shi, and Gordon Wetzstein. 2024. Large Étendue 3D Holographic Display with Content-adaptive Dynamic Fourier Modulation. In *SIGGRAPH Asia 2024 Conference Papers*. 1–12.

Chun Chen, Byounghyo Lee, Nan-Nan Li, Minseok Chae, Di Wang, Qiong-Hua Wang, and Byoungho Lee. 2021. Multi-depth hologram generation using stochastic gradient descent algorithm with complex loss function. *Opt. Express* 29, 10 (May 2021), 15089–15103. https://doi.org/10.1364/OE.425077

J-S Chen and DP Chu. 2015. Improved layer-based method for rapid hologram generation and real-time interactive holographic display applications. *Optics express* 23, 14 (2015), 18143–18155.

Rick H.-Y. Chen and Timothy D Wilkinson. 2009a. Computer generated hologram from point cloud using graphics processor. *Applied optics* 48, 36 (2009), 6841–6850.

Rick H-Y Chen and Timothy D Wilkinson. 2009b. Computer generated hologram with geometric occlusion using GPU-accelerated depth buffer rasterization for three-dimensional display. *Applied optics* 48, 21 (2009), 4246–4255.

Suyeon Choi, Brian Chao, Jacqueline Yang, Manu Gopakumar, and Gordon Wetzstein. 2025. Gaussian Wave Splatting for Computer Generated Holography. *ACM Transactions on Graphics (TOG)* 44, 4 (2025), 1–11.

Suyeon Choi, Manu Gopakumar, Yifan Peng, Jonghyun Kim, Matthew O'Toole, and Gordon Wetzstein. 2022. 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*. 1–9.

Vincent R Curtis, Nicholas W Caira, Jiayi Xu, Asha Gowda Sata, and Nicolas C Pégard. 2021. DCGH: dynamic computer generated holography for speckle-free, high fidelity 3D displays. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*. IEEE, 1–9.

JR Fienup. 1980. Iterative method applied to image reconstruction and to computer-generated holograms. *Optical Engineering* 19, 3 (1980), 297–305.

Ralph W Gerchberg. 1972. A practical algorithm for the determination of plane from image and diffraction pictures. *Optik* 35, 2 (1972), 237–246.

Joseph W Goodman. 2005. *Introduction to Fourier optics*. Roberts and Company publishers.

Joseph W Goodman. 2007. *Speckle phenomena in optics: theory and applications*. Roberts and Company Publishers.

Joseph W. Goodman. 2015. *Statistical Optics* (2nd ed.). Wiley.

Manu Gopakumar, Gun-Yeal Lee, Suyeon Choi, Brian Chao, Yifan Peng, Jonghyun Kim, and Gordon Wetzstein. 2024. Full-colour 3D holographic augmented-reality displays with metasurface waveguides. *Nature* (2024), 1–7.

Stephen Hamann, Liang Shi, Olav Solgaard, and Gordon Wetzstein. 2018. Time-multiplexed light field synthesis via factored Wigner distribution function. *Optics letters* 43, 3 (2018), 599–602.

Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2024. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 Conference Papers*. 1–11.

Changwon Jang, Kiseung Bang, Minseok Chae, Byoungho Lee, and Douglas Lanman. 2024. Waveguide holography for 3D augmented reality glasses. *Nature Communications* 15, 1 (2024), 66.

Bahram Javidi, Artur Carnicer, Arun Anand, George Barbastathis, Wen Chen, Pietro Ferraro, JW Goodman, Ryoichi Horisaki, Kedar Khare, Malgorzata Kujawinska, et al. 2021. Roadmap on digital holography. *Optics Express* 29, 22 (2021), 35078–35118.

James T Kajiya and Brian P Von Herzen. 1984. Ray tracing volume densities. *ACM SIGGRAPH computer graphics* 18, 3 (1984), 165–174.

Hoonjong Kang, Takeshi Yamaguchi, and Hiroshi Yoshikawa. 2008. Accurate phase-added stereogram to improve the coherent stereogram. *Applied optics* 47, 19 (2008), D44–D54.

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM TOG* 42, 4 (July 2023).

Dongyeon Kim, Seung-Woo Nam, Suyeon Choi, Jong-Mo Seo, Gordon Wetzstein, and Yoonchan Jeong. 2024. Holographic parallax improves 3D perceptual realism. *ACM Transactions on Graphics (TOG)* 43, 4 (2024), 1–13.

Dongyeon Kim, Seung-Woo Nam, Byounghyo Lee, Jong-Mo Seo, and Byoungho Lee. 2022b. Accommodative holography: improving accommodation response for perceptually realistic holographic displays. *ACM Trans. Graph.* 41, 4, Article 111 (July 2022), 15 pages. https://doi.org/10.1145/3528223.3530147

Jonghyun Kim, Manu Gopakumar, Suyeon Choi, Yifan Peng, Ward Lopes, and Gordon Wetzstein. 2022a. Holographic Glasses for Virtual Reality. In *ACM SIGGRAPH 2022 Conference Proceedings* (Vancouver, BC, Canada) *(SIGGRAPH '22)*. Association for Computing Machinery, New York, NY, USA, Article 33, 9 pages. https://doi.org/10.1145/3528233.3530739

Grace Kuo, Florian Schiffers, Douglas Lanman, Oliver Cossairt, and Nathan Matsuda. 2023. Multisource holography. *ACM Transactions on Graphics (TOG)* 42, 6 (2023), 1–14.

Byounghyo Lee, Dongyeon Kim, Seungjae Lee, Chun Chen, and Byoungho Lee. 2022. High-contrast, speckle-free, true 3D holography via binary CGH optimization. *Scientific Reports* 12, 1 (2022), 2811. https://doi.org/10.1038/s41598-022-06405-2

Seungjae Lee, Dongyeon Kim, Seung-Woo Nam, Byounghyo Lee, Jaebum Cho, and Byoungho Lee. 2020. Light source optimization for partially coherent holographic displays with consideration of speckle contrast, resolution, and depth of field. *Scientific reports* 10, 1 (2020), 18832.

Adoph W Lohmann and DP Paris. 1967. Binary Fraunhofer holograms, generated by computer. *Applied optics* 6, 10 (1967), 1739–1748.

Mark E Lucente. 1993. Interactive computation of holograms using a look-up table. *Journal of Electronic Imaging* 2, 1 (1993), 28–34.

Andrew Maimone, Andreas Georgiou, and Joel S Kollin. 2017. Holographic near-eye displays for virtual and augmented reality. *ACM Transactions on Graphics (Tog)* 36, 4 (2017), 1–16.

Kyoji Matsushima. 2005a. Computer-generated holograms for three-dimensional surface objects with shade and texture. *Applied optics* 44, 22 (2005), 4607–4614.

Kyoji Matsushima. 2005b. Exact hidden-surface removal in digitally synthetic full-parallax holograms. In *Practical Holography XIX: Materials and Applications*, Vol. 5742. SPIE, 25–32.

Kyoji Matsushima. 2020. *Introduction to Computer Holography: Creating Computer-Generated Holograms as the Ultimate 3D Image*. Springer Nature.

Kyoji Matsushima and Sumio Nakahara. 2009. Extremely high-definition full-parallax computer-generated hologram created by the polygon-based method. *Applied optics* 48, 34 (2009), H54–H63.

Kyoji Matsushima, Masaki Nakamura, and Sumio Nakahara. 2014. Silhouette method for hidden surface removal in computer holography and its acceleration using the switch-back technique. *Optics express* 22, 20 (2014), 24450–24465.

Kyoji Matsushima, Hagen Schimmel, and Frank Wyrowski. 2003. Fast calculation method for optical diffraction on tilted planes by use of the angular spectrum of plane waves. *JOSA A* 20, 9 (2003), 1755–1762.

Kyoji Matsushima and Tomoyoshi Shimobaba. 2009. Band-limited angular spectrum method for numerical simulation of free-space propagation in far and near fields. *Optics express* 17, 22 (2009), 19662–19673.

Kyoji Matsushima and Noriaki Sonobe. 2018. Full-color digitized holography for large-scale holographic 3D imaging of physical and nonphysical objects. *Applied Optics* 57, 1 (2018), A150–A156.

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *ECCV*.

Nitish Padmanaban, Yifan Peng, and Gordon Wetzstein. 2019. Holographic near-eye displays based on overlap-add stereograms. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–13.

Jae-Hyeung Park. 2017. Recent progress in computer-generated holography for three-dimensional scenes. *Journal of Information Display* 18, 1 (2017), 1–12.

Jae-Hyeung Park and Mehdi Askari. 2019. Non-hogel-based computer generated hologram from light field using complex field recovery technique from Wigner distribution function. *Optics express* 27, 3 (2019), 2562–2574.

Pierre Pellat-Finet. 1994. Fresnel diffraction and the fractional-order Fourier transform. *Optics Letters* 19, 18 (1994), 1388–1390.

Yifan Peng, Suyeon Choi, Jonghyun Kim, and Gordon Wetzstein. 2021. Speckle-free holography with partially coherent light sources and camera-in-the-loop calibration. *Science advances* 7, 46 (2021), eabg5040.

Yifan Peng, Suyeon Choi, Nitish Padmanaban, and Gordon Wetzstein. 2020. Neural holography with camera-in-the-loop training. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–14.

Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2023. *Physically based rendering: From theory to implementation*. MIT Press.

Dapu Pi, Juan Liu, and Yongtian Wang. 2022. Review of computer-generated hologram algorithms for color dynamic holographic three-dimensional display. *Light: Science & Applications* 11, 1 (2022), 231.

Yijun Qi, Chenliang Chang, and Jun Xia. 2016. Speckleless holographic display by complex modulation based on double-phase method. *Optics express* 24, 26 (2016), 30368–30378.

Florian Schiffers, Praneeth Chakravarthula, Nathan Matsuda, Grace Kuo, Ethan Tseng, Douglas Lanman, Felix Heide, and Oliver Cossairt. 2023. Stochastic light field holography. In *2023 IEEE International Conference on Computational Photography (ICCP)*. IEEE, 1–12.

Liang Shi, Beichen Li, Changil Kim, Petr Kellnhofer, and Wojciech Matusik. 2021. Towards real-time photorealistic 3D holography with deep neural networks. *Nature* 591, 7849 (2021), 234–239.

Liang Shi, Beichen Li, and Wojciech Matusik. 2022. End-to-end learning of 3d phase-only holograms for holographic display. *Light: Science & Applications* 11, 1 (2022), 247.

Liang Shi, DongHun Ryu, and Wojciech Matusik. 2024. Ergonomic-Centric Holography: Optimizing Realism, Immersion, and Comfort for Holographic Display. *Laser & Photonics Reviews* n/a, n/a (2024), 2300651. https://doi.org/10.1002/lpor.202300651

Tomoyoshi Shimobaba and Tomoyoshi Ito. 2015. Random phase-free computer-generated hologram. *Optics express* 23, 7 (2015), 9549–9554.

Pierre St-Hilaire. 1995. Phase profiles for holographic stereograms. *Optical Engineering* 34, 1 (1995), 83–89.

Jiaxiang Tang, Hang Zhou, Xiaokang Chen, Tianshu Hu, Errui Ding, Jingdong Wang, and Gang Zeng. 2023. Delicate Textured Mesh Recovery from NeRF via Adaptive Surface Refinement. In *ICCV*.

Ayush Tewari, Ohad Fried, Justus Thies, Vincent Sitzmann, Stephen Lombardi, Kalyan Sunkavalli, Ricardo Martin-Brualla, Tomas Simon, Jason Saragih, Matthias Nießner, et al. 2020. State of the art on neural rendering. In *Computer Graphics Forum*, Vol. 39. Wiley Online Library, 701–727.

Ayush Tewari, Justus Thies, Ben Mildenhall, Pratul Srinivasan, Edgar Tretschk, Wang Yifan, Christoph Lassner, Vincent Sitzmann, Ricardo Martin-Brualla, Stephen Lombardi, et al. 2022. Advances in neural rendering. In *Computer Graphics Forum*, Vol. 41. Wiley Online Library, 703–735.

Peter Wai Ming Tsang and T-C Poon. 2013. Novel method for converting digital Fresnel hologram to phase-only hologram based on bidirectional error diffusion. *Optics express* 21, 20 (2013), 23680–23686.

Taichi Yanagiya and Kyoji Matsushima. 2019. Rendering of translucent models with alpha blending in polygon-based computer-generated holograms. *Proceedings of the 3D Image Conference* 27 (2019), 73–76. https://doi.org/10.60374/sanjigen.27.0_73

Fahri Yaraş, Hoonjong Kang, and Levent Onural. 2010. State of the art in holographic displays: a survey. *Journal of display technology* 6, 10 (2010), 443–454.

Vickie Ye, Ruilong Li, Justin Kerr, Matias Turkulainen, Brent Yi, Zhuoyang Pan, Otto Seiskari, Jianbo Ye, Jeffrey Hu, Matthew Tancik, and Angjoo Kanazawa. 2024. gsplat: An Open-Source Library for Gaussian Splatting. *arXiv preprint arXiv:2409.06765* (2024). arXiv:2409.06765 [cs.CV] https://arxiv.org/abs/2409.06765

Dongheon Yoo, Youngjin Jo, Seung-Woo Nam, Chun Chen, and Byoungho Lee. 2021. Optimization of computer-generated holograms featuring phase randomness control. *Optics Letters* 46, 19 (2021), 4769–4772.

Hao Zhang, Liangcai Cao, and Guofan Jin. 2019. Three-dimensional computer-generated hologram with Fourier domain segmentation. *Optics express* 27, 8 (2019), 11689–11697.

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.

Zhengyun Zhang. 2011. *Analysis and synthesis of three-dimensional illumination using partial coherence.* Stanford University.

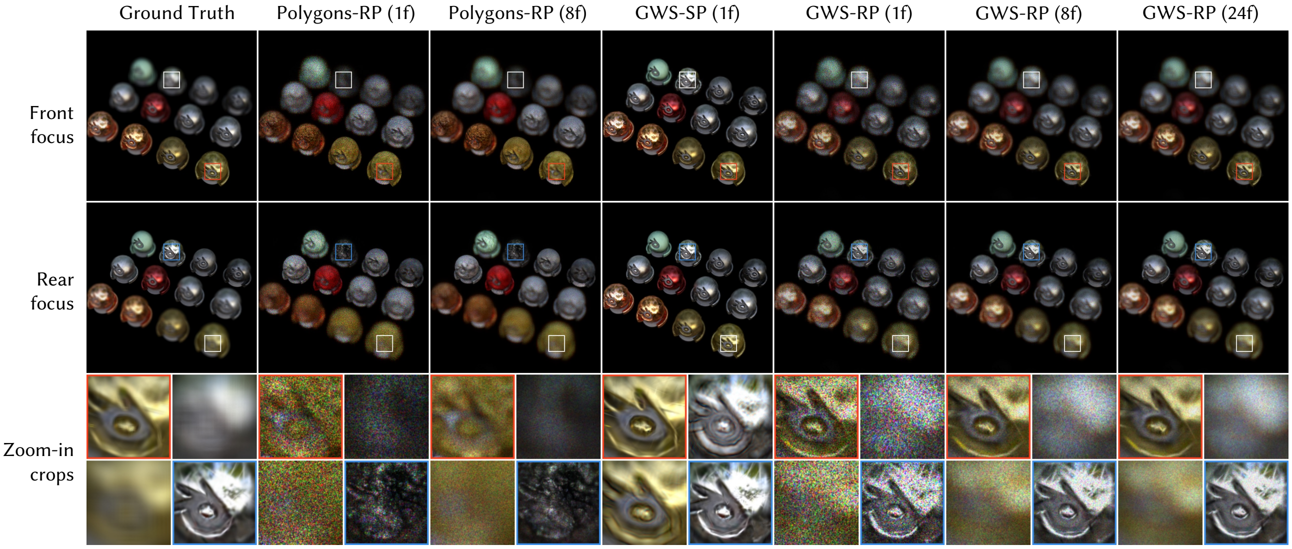

Fig. 6. **Simulated 3D focal stacks reconstructed from various baseline CGH algorithms.** The image quality of random-phase polygons-based CGH (Polygons-RP) is inherently limited by the coarse per-face textured mesh representation, resulting in poor image quality even in in-focus regions. Smooth-phase GWS (GWS-SP) [Choi et al. 2025] reconstructs sharp details at in-focus regions, but suffers from large depth of field and unnatural ringing artifacts. Our method (GWS-RP) generates sharp content at focused regions and the resulting hologram has shallow depth of field, reconstructing natural defocus blur across different depths. With additional time-multiplexing, the image quality of GWS-RP significantly improves.

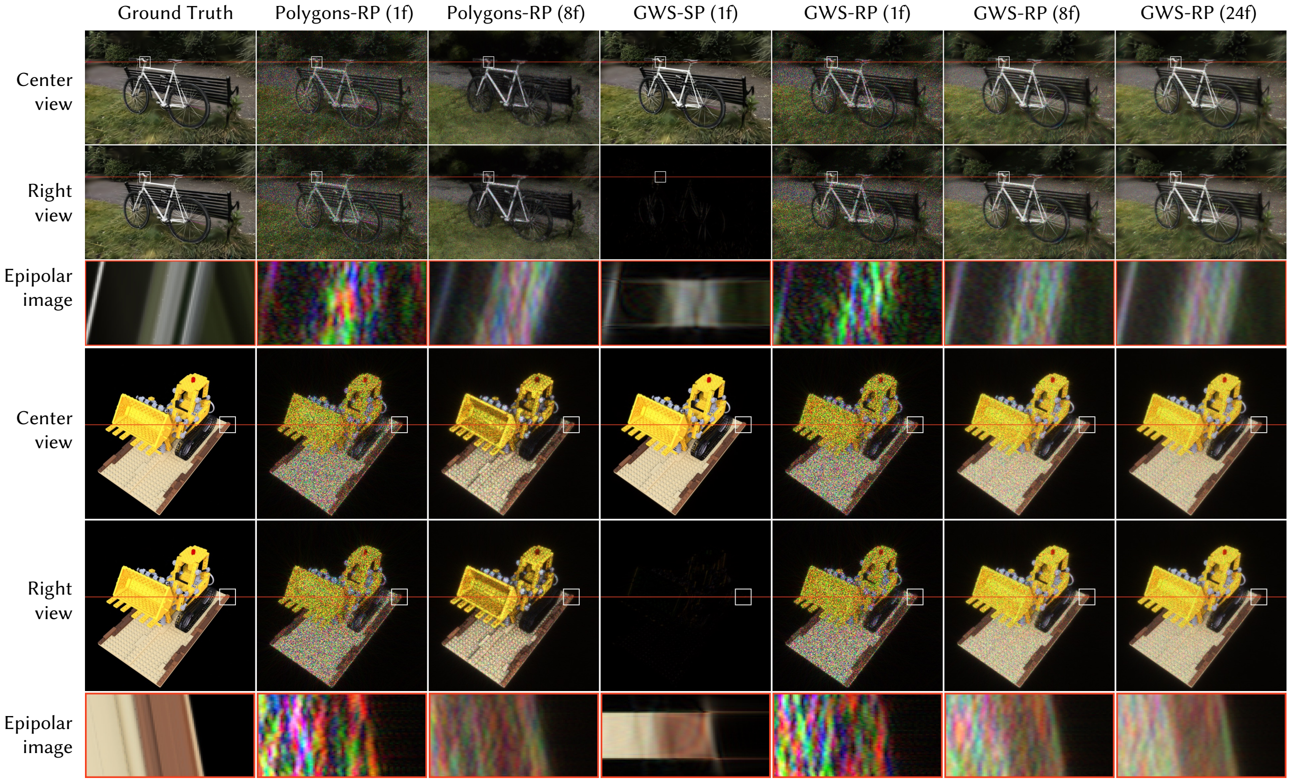

Fig. 7. **Simulated light fields reconstructed from various baseline CGH algorithms.** The spectrum of of smooth-phase GWS (GWS-SP) holograms [Choi et al. 2025] is highly concentrated, resulting in a tiny eyebox. As the pupil moves to the edge of the eyebox, the light field views completely disappear, as shown in the dark edges in the horizontal epipolar images. Conversely, energy is evenly distributed in the spectrum of random-phase holograms (Polygons-RP and GWS-RP), thus accurate parallax can be reconstructed across the full eyebox.

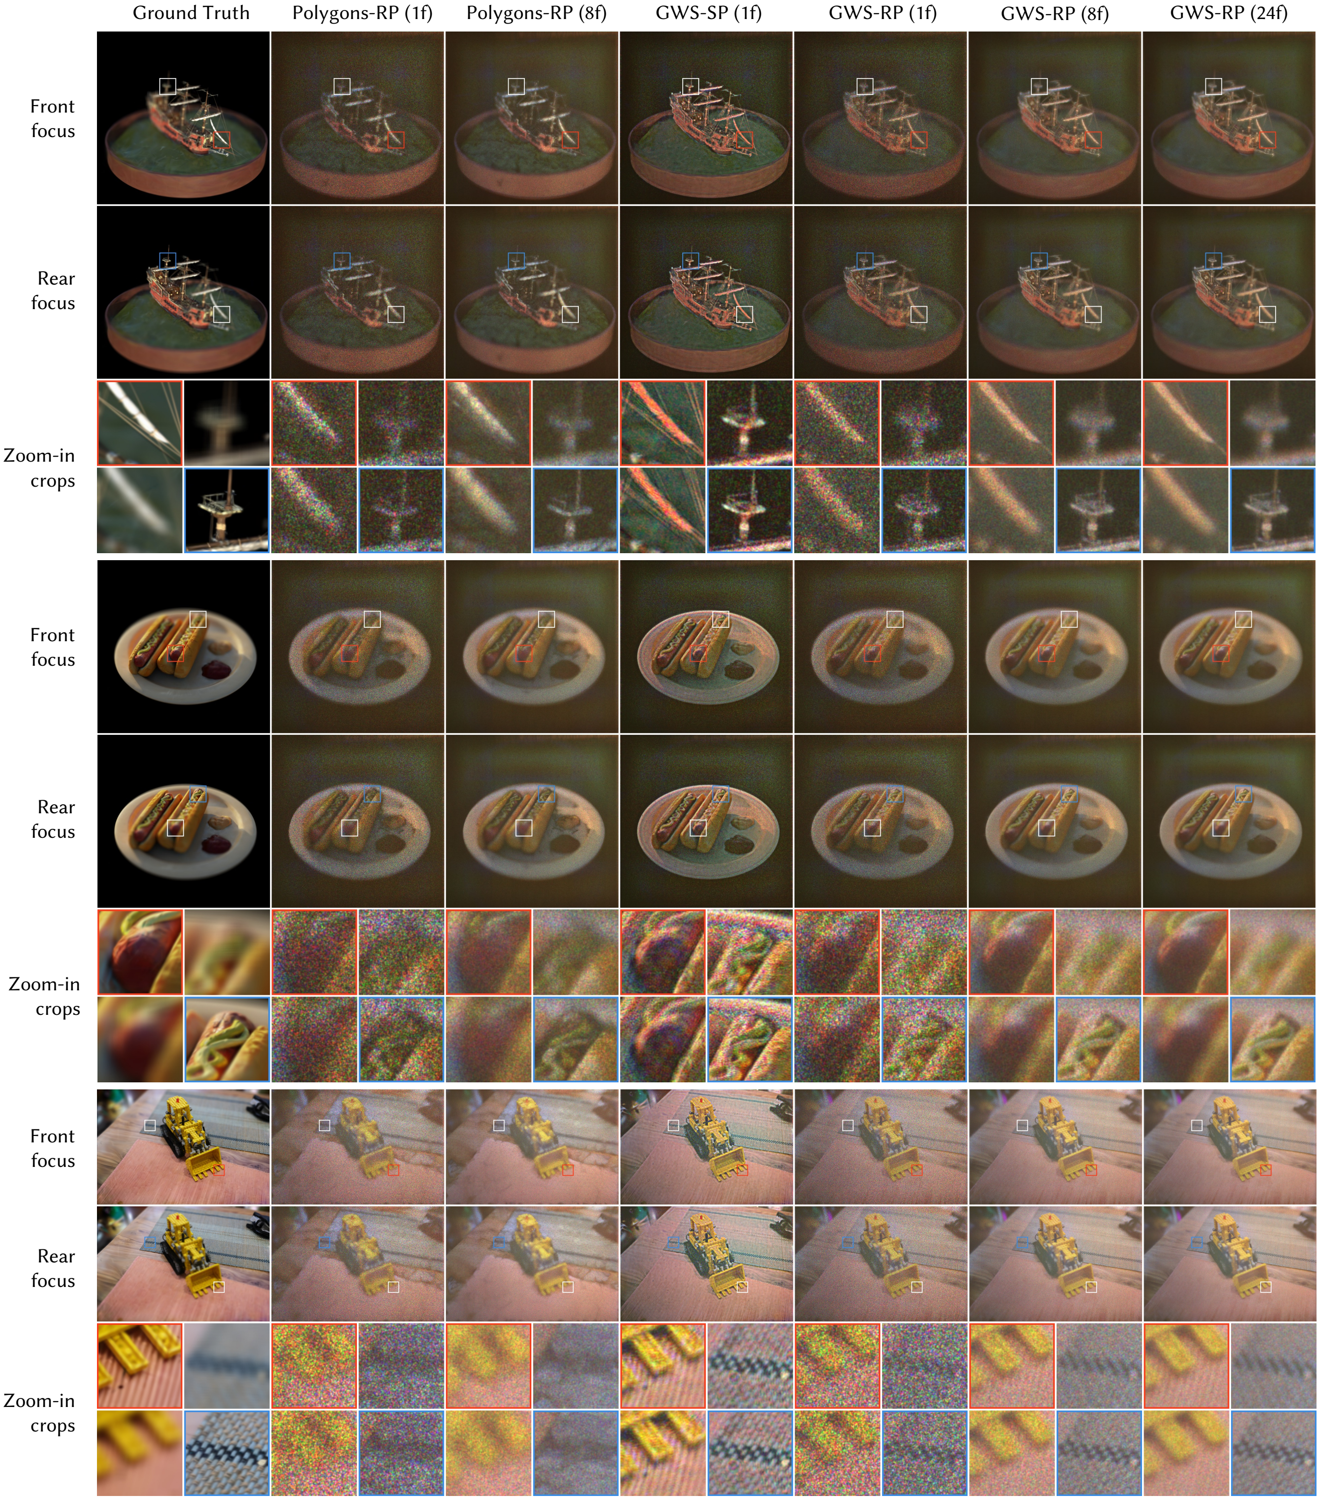

Fig. 8. **Experimentally captured 3D focal stacks of holograms generated using different CGH algorithms.** Polygon-based CGH (Polygons-RP) [Matsushima 2005a; Matsushima and Nakahara 2009; Matsushima et al. 2014] achieve low image quality even in in-focus regions due to the low quality the underlying textured mesh 3D representation. GWS-SP [Choi et al. 2025] generates smooth-phase holograms, resulting in limited defocus blur with unnatural ringing artifacts which are most prominent in regions with thin structures such as the rope on the ship. Our method, GWS-RP, achieves good image quality in in-focus regions and reconstructs natural incoherent blur in defocus regions. With 24 frames time-multiplexing, GWS-RP achieves near speckle-free results.