Instant Random Phase Holography with Parallax and Defocus Cues

OSCAR PUEYO-CIUTAD, Universidad de Zaragoza - I3A, Spain
VICTOR CHU, University of Toronto, Canada
CONGLI WANG, Princeton University, USA
ETHAN TSENG, Princeton University, USA
FLORIAN SCHIFFERS, Northwestern University, USA
ALBERT REDO-SANCHEZ, Universidad de Zaragoza - I3A, Spain
DIEGO GUTIERREZ, Universidad de Zaragoza - I3A, Spain
FELIX HEIDE, Princeton University, USA

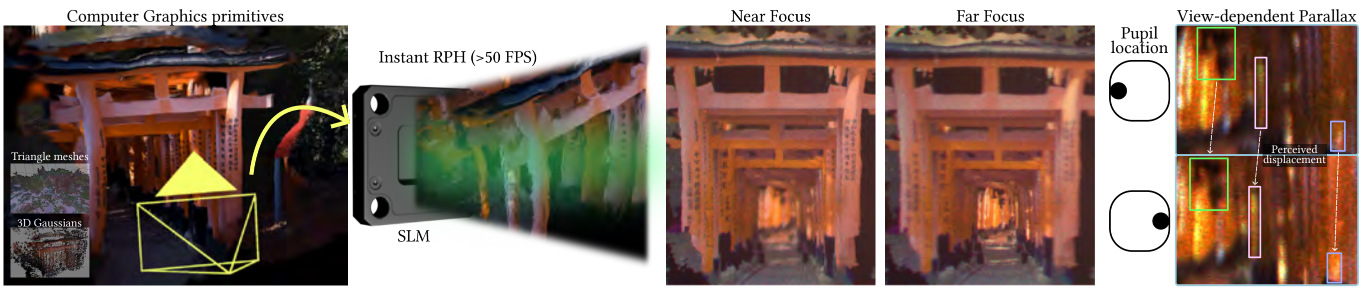

Fig. 1. **Instant Random Phase Holography.** Holographic displays promise immersive, true 3D visuals within a compact form factor. However, existing phase retrieval methods struggle to support modern graphics primitives, fail to achieve real-time framerates, or rely on brittle "smooth phase" profiles that succumb to dynamic obstructions [Chu et al. 2025]. We introduce Instant Random Phase Holography (Instant RPH), the first method for high-quality, real-time random phase synthesis for holographic displays. By rendering arbitrary graphics primitives into a wave-optics-friendly, diopter-space Multiplane Image (MPI) and reparameterizing phase retrieval via Linear Amplitude Encoders (LAE), our method efficiently synthesizes artifact-resilient phase patterns at 54 FPS. Instant RPH natively reproduces natural defocus cues, handles view-dependent parallax, and remains temporally stable, effectively bridging the gap between modern real-time graphics and holographic hardware.

Computer-generated holography (CGH) has the potential to become a practical medium for immersive AR/VR, offering true volumetric depth, real-time performance, and robustness to optical deviations. However, existing CGH methods are forced to compromise between performance, volumetric accuracy, and robustness. Iterative solvers achieve high contrast but lack interactive speed. Recent neural methods achieve speed but converge on "smooth phase" solutions that are brittle to eye movement, aberrations, and fail to produce natural defocus blur and parallax effects.

While random phase holography leverages the full angular spectrum of a display to generate natural visual cues and artifact resilience, current real-time random phase methods suffer from severe speckle noise and contrast loss. In this work, we introduce Instant Random Phase Holography (Instant RPH), the first random phase CGH method that simultaneously achieves real-time performance, high image quality, and true volumetric 3D. Our key insight is that real-time random phase holography via amplitude discard produces *predictable* speckle noise and degraded image quality. Therefore, we reparameterize phase retrieval as a target amplitude modulation problem, designing Linear Amplitude Encoders (LAE) to counteract this degradation. Our method generates time-multiplexed holograms that are collaboratively trained to actively suppress speckle and maximize image quality. To extend this formulation to high-quality 3D, we leverage diopter-space Multiplane Images (MPI) as an efficient intermediate representation that decouples rendering cost from geometric complexity. By replacing global wave propagation with compact Rayleigh-Sommerfeld Diffraction (C-RSD) convolutions, we simulate light transport between layers with correct occlusion in real-time.

Instant RPH generates artifact-resilient holograms by construction, naturally encoding accurate occlusion, view-dependent parallax, and optical blur in a single forward pass. Natively supporting arbitrary 3D representations, from standard triangle meshes to 3D Gaussian Splatting (3DGS), we achieve high-fidelity, photorealistic 3D content streaming at 54 FPS (> 400× speed-up over light-field synthesis), effectively transforming the holographic display into a robust, plug-and-play monitor.

CCS Concepts: • **Hardware** → **Emerging optical and photonic technologies**; **Emerging interfaces**.

Additional Key Words and Phrases: computational display, holographic display, AR/VR

ACM Reference Format:
Oscar Pueyo-Ciutad, Victor Chu, Congli Wang, Ethan Tseng, Florian Schiffers, Albert Redo-Sanchez, Diego Gutierrez, and Felix Heide. 2026. Instant

Authors' Contact Information: Oscar Pueyo-Ciutad, Universidad de Zaragoza - I3A, Zaragoza, Spain; Victor Chu, University of Toronto, Toronto, Canada; Congli Wang, Princeton University, Princeton, NJ, USA; Ethan Tseng, Princeton University, Princeton, NJ, USA; Florian Schiffers, Northwestern University, Evanston, IL, USA; Albert Redo-Sanchez, Universidad de Zaragoza - I3A, Zaragoza, Spain; Diego Gutierrez, Universidad de Zaragoza - I3A, Zaragoza, Spain; Felix Heide, Princeton University, Princeton, NJ, USA.

SA Conference Papers '26, Kuala Lumpur, Malaysia
 The definitive Version of Record was published in *SIGGRAPH Asia 2026 Conference Papers (SA Conference Papers '26), December 01–04, 2026, Kuala Lumpur, Malaysia*, https://doi.org/10.1145/3829340.3842293.

Random Phase Holography with Parallax and Defocus Cues. In *SIGGRAPH Asia 2026 Conference Papers (SA Conference Papers '26), December 01–04, 2026, Kuala Lumpur, Malaysia.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3829340.3842293

1 Introduction

Conventional holograms aim to reproduce the full waveform of an illuminated scene, inherently capturing light-transport effects, accurate depth cues and view-dependent parallax and occlusions. In the context of near-eye displays, holographic techniques promise a wearable form factor for dynamic content [Akşit and Itoh 2023; Chakravarthula et al. 2019; Choi et al. 2025b; Kim et al. 2024, 2022b,a; Maimone et al. 2017].

However, despite recent progress on image quality [Choi et al. 2025b; Shi et al. 2022a], existing rendering pipelines remain an obstacle to practical, immersive holographic displays. Because of the 3D nature of holographic light transport, popular methods relying on 2D RGB-D maps [Chu et al. 2025; Shi et al. 2021] or focal stacks [Chakravarthula et al. 2019; Kuo et al. 2023] only offer a partial solution, sacrificing occlusion and parallax accuracy to simplify computation. Conversely, dense light field methods [Chakravarthula et al. 2019, 2022; Kim et al. 2024; Padmanaban et al. 2019; Schiffers et al. 2023] incorporate these effects, but at the cost of massive memory and rendering requirements, disconnecting holography from interactive 3D content.

Solving the non-convex phase retrieval problem for high-fidelity holography demands massive computational resources. Iterative optimization and Fourier calculations [Chakravarthula et al. 2022; Kavakli et al. 2023; Kim et al. 2024; Padmanaban et al. 2019; Yang et al. 2022] are simply too slow for interactive scenes. Recent 3D Gaussian splatting methods [Choi et al. 2025a; Kerbl et al. 2023] are similarly prohibitive, as the number of Fourier operations scales with scene complexity. Neural methods achieving real-time performance [Chen et al. 2025; Choi et al. 2021; Peng et al. 2020; Shi et al. 2021, 2022b] typically converge on *smooth phase* solutions. While fast, smooth phase holograms are notoriously brittle and restricted to a narrow eyebox [Kuo et al. 2020; Tseng et al. 2024]. They struggle to reproduce realistic defocus and parallax cues, and they lack the stability to handle common optical artifacts like eyelashes or lens dust [Chu et al. 2025].

In contrast, random phase holograms spread intensity over the eyebox, enabling natural defocus, parallax, and artifact resilience [Kim et al. 2024]. However, current real-time random phase methods suffer from severe contrast loss and heavy speckle noise. Recent attempts at content-adaptive pseudo-random phase generation [Chu et al. 2025] attempt to suppress this noise but are restricted to 2.5D RGB-D data and minimal depth resolution, creating holograms that show incorrect parallax cues with visible artifacts. Parallax cues in near-eye holographic displays are subtle, but crucial for an immersive experience [Kim et al. 2024].

In this work, we introduce Instant RPH, a novel real-time method for high-quality volumetric 3D holograms that bridges the gap between modern graphics and holographic hardware. To achieve it, we observe that within a One-Step Phase Retrieval (OSPR) framework, which aims at predicting the phase pattern to be displayed in one forward step, naive amplitude discard produces structured, predictable image degradation. To leverage this observation, we first reparameterize OSPR in 2D as a target amplitude modulation problem, jointly training time-multiplexed Linear Amplitude Encoders (LAEs) to actively counteract speckle noise and contrast loss without imposing any iterative overhead. Since the target plane amplitude and SLM phase are strictly coupled via wave propagation, we can encode target amplitude —which would otherwise be lost during amplitude discard— directly into the spatial phase profile. Then, to scale this approach to full 3D, we leverage a perceptually-driven diopter-space Multiplane Image (MPI) representation that decouples the holographic phase-retrieval cost from geometric complexity [Möller et al. 2018], replace costly global Fourier Transforms (FFTs) [Goodman 2005] with highly localized spatial-domain convolutions, and train a LAE per layer to maximize image quality.

Fig. 2. **Real-Time Random Phase Holography.** Existing random phase holography methods are bounded by a strict trade-off between speed and image quality. To achieve interactive rates, iterative solvers significantly sacrifice quality, yielding very poor results (top). The recent ARH [Chu et al. 2025] struggles in both image quality and runtime when scaling up focal resolution (middle). Our Instant RPH overcomes this trade-off, synthesizing high-quality 3D holograms with natural defocus and parallax (bottom).

To ensure a fair performance comparison, our reported frame rates are measured in the critical execution-path regime, consistent with previous state-of-the-art evaluations [Chu et al. 2025]. Under this standard RGB regime, Instant RPH runs at 54 FPS. More impressively, when scaling our pipeline to maximum-quality settings, processing full 3D RGB volumes with 24× time-multiplexing, our method sustains interactive rates of 5 FPS. In contrast, the recent ARH method [Chu et al. 2025] offers significantly less quality at just 7.6 FPS in the standard regime, and 1 FPS under ARH's maximum-quality settings (see Figure 2). This fundamental efficiency leap was achieved using standard PyTorch JiT; integrating custom CUDA kernels, model quantization, or deploying on next-generation hardware can be explored to make our method even faster in future work.

We hope our work helps unlock the potential of real-time random phase holography. The source code and data for this project can be found on our GitHub repository[1].

2 Related Work

In the following, we review the related work on Computer-Generated Holography (CGH), covering both phase retrieval and scene representation. We also provide a summary in Table 1.

2.1 Computer-Generated Holography

The goal of CGH is to retrieve the phase pattern to display it, typically, on a phase-only Spatial Light Modulator (SLM), which will propagate it to form a target 3D field. Classical iterative solvers can converge to high-fidelity solutions [Chakravarthula et al. 2019; Gerchberg 1972], although the requisite iteration count precludes real-time operation for dynamic scenes.

Real-time approaches lean towards One-Step Phase Retrieval (OSPR) methods, where preliminary techniques relied on analytical schemes that preserve amplitude at the cost of halving spatial resolution, sacrificing image quality or increasing latency [Maimone et al. 2017; Tsang and Poon 2013]. To achieve high quality at interactive rates, recent research has pivoted to Neural OSPR with deep learning [Choi et al. 2021; Peng et al. 2020; Shi et al. 2021, 2022a] to predict phase patterns directly from intensity inputs. A structural limitation of these methods is their reliance on intensity-based loss functions (e.g., MSE), which bias networks toward "smooth phase" solutions. While smooth phase holograms achieve high-fidelity, they concentrate energy into narrow diffraction angles. This lack of angular diversity results in fragile holograms that are often disrupted by ocular occlusions, such as eyelashes or floaters [Chu et al. 2025], and fail to reproduce correct defocus blur and parallax.

In contrast, *random phase* holograms distribute energy across the full aperture, providing inherent resilience to occlusion and physically accurate defocus blur [Kim et al. 2022b]. While recent works, such as Real-Time ARH [Chu et al. 2025], attempt to regularize HoloNet-like backbones toward random phase, these methods remain strictly bound by the limitations of modulating the target phase pattern and complex-to-phase conversion. Furthermore, they are bounded by RGB-D input data and scale poorly with depth resolution, typically restricted to a sparse approximation (e.g., three planes) to maintain interactive frame rates.

In our work, we reparameterize neural OSPR to achieve high-quality 3D holography. To leverage the full angular diversity of an SLM, we use a fixed random phase initialization. We observe that this initialization, combined with amplitude discard, produces a predictable increase in speckle noise and a decrease in contrast. To address this, we train a time-multiplexing-aware encoder that collaboratively counteracts these predictable artifacts. Then, we construct a graphics-driven pipeline based on multiplane images to leverage the full capabilities of random phase holography, naturally reproducing defocus and parallax cues.

[1] https://github.com/opueyociutad/InstantRPH

Table 1. **Near-eye Computer-Generated Holography.** Computer-generated holograms for near-eye displays must provide high image quality, view-dependent effects, and retinal blur, while being robust to artifacts and computationally lightweight in order to be used in head-mounted displays. We propose a method that addresses all of these aspects simultaneously.

Data type	RGB	RGB-D		Focal stack	Light field			MPI	
Method	HoloNet [2020]	TensorV1 [2021]	ARH [2025]	MSH [2023]	OLAS [2019]	SLFH [2023]	HP [2024]	TensorV2 [2022a]	Ours
Quality	✓	✗	✓	✓	✗	✓	✓	✓	✓
Depth res.	✗	✓	✗	✓	✓	✓	✓	✓	✓
Parallax	✗	✗	✗	✗	✗	✓	✓	✗	✓
Retinal blur	✗	✗	✓	✓	✗	✓	✓	✗	✓
Artifacts	✗	✗	✓	✓	✗	✓	✓	✗	✓
Random phase	✗	✗	✓	✓	✗	✓	✓	✗	✓
Single-step	✓	✓	✓	✗	✓	✗	✗	✓	✓
Real-time	✓	✓	✓	✗	✗	✗	✗	✓	✓

2.2 Efficient Volumetric Holography

Moving beyond RGB-D inputs to full 3D content introduces a major challenge: computational scaling with scene complexity. Early point-based methods [Maimone et al. 2017] offered precision but suffered from linear scaling with the number of primitives. To utilize modern rasterization hardware, recent work has formalized diffraction for surface primitives [Liu et al. 2023; Yang et al. 2014; Zhang et al. 2018] and GPU culling techniques [Fricke et al. 2025; Wang et al. 2023].

Despite these optimizations, primitive-based approaches scale poorly with the number of primitives, thus prohibiting rendering high-fidelity scenes with millions of polygons or Gaussian splats [Choi et al. 2025a] directly to a hologram. We use *Multiplane Image (MPI)* [Zhou et al. 2018] as an intermediate representation widely used in view synthesis [Tucker and Snavely 2020; Wizadwongsa et al. 2021] and multifocal displays [MacKenzie et al. 2010; Narain et al. 2015]. Several holography works also leverage MPIs to achieve occlusion effects for efficient holography [Zhan et al. 2026; Zhang et al. 2017; Ziegler et al. 2008]. By rasterizing geometry into a fixed-resolution MPI, the cost of holographic generation is decoupled from the geometric complexity of the 3D scene.

The wave field propagation over the MPI has been made through wave recording planes (WRP) [Li et al. 2025; Shimobaba et al. 2009; Symeonidou et al. 2015], with learned wave propagation between layers for 3D holography [Choi et al. 2021; Shi et al. 2022b], or using Split-Lohmann transformations [Ding et al. 2024; Lagrange et al. 2025; Zhou et al. 2025] or depth-of-focus pruning [Ma et al. 2024].

However, these methods either rely on slow ASM propagation between planes or on learned propagators. In contrast, we derive our MPI layers from the human depth-of-field [Campbell 1957; Marcos et al. 1999], and leverage that in a holographic setup these layers are physically close. This allows us to replace costly ASM propagation between planes with a spatially Compact Rayleigh-Sommerfeld Diffraction (C-RSD) convolution, enabling us to synthesize our perception-derived MPI of 21 layers at real-time rates.

3 Background

Before describing the proposed method, we briefly formalize the background on holographic wave propagation.

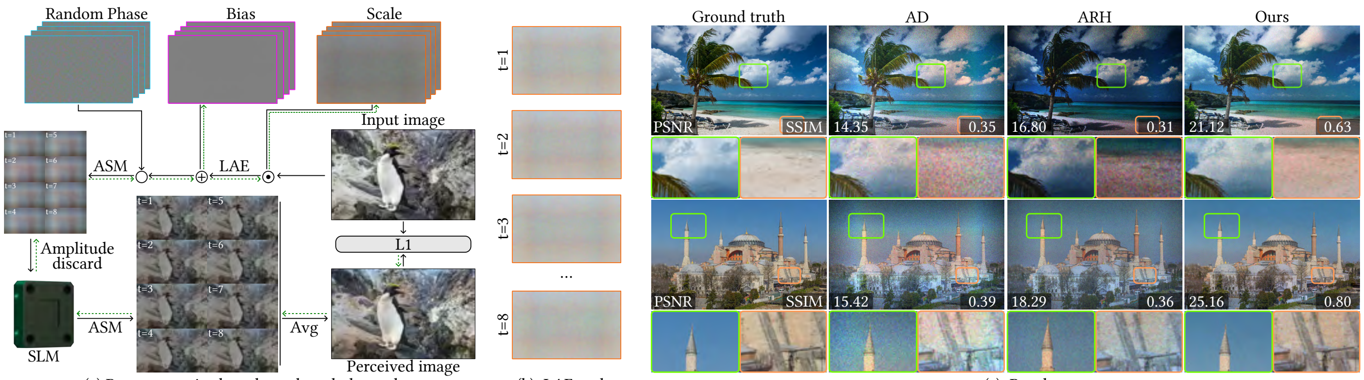

(a) Reparameterized random-phase holography (b) LAE scale (c) Results

Fig. 3. **One-Step Phase Retrieval (OSPR) Reparameterization.** (a) We reparameterize Neural OSPR methods by jointly learning a Linear Amplitude Encoder (LAE) per multiplexed frame ($T = 8$ in this example) to effectively counteract the loss of image quality and speckle from traditional random-phase holography. (b) Our jointly learned LAEs resemble a speckle-like pattern, hinting at how they are counteracting speckle in the result. (c) Our Neural OSPR Reparameterization significantly enhances image quality and suppresses speckle, achieving a PSNR increase up to 10 dB over uncorrected OSPR (Amplitude Discard, AD), and up to 7 dB over Real-Time ARH [Chu et al. 2025], the state of the art in real-time random-phase holography.

3.1 Wave Propagation Model

Light transport between parallel planes can be computed using the Angular Spectrum Method (ASM) [Delen and Hooker 1998; Goodman 2005; Matsushima and Shimobaba 2009] by decomposing the complex field into plane waves and propagating its Fourier components. Given $U(x, y; z)$, denoting the complex field at a distance z, its propagation to a new plane at a distance $z + \Delta z$ is expressed, in the frequency domain, as a multiplication by the transfer function H as

$$U(x, y; z + \Delta z) = \mathcal{F}^{-1}\left[\mathcal{F}\left(U(x, y; z)\right) \cdot H(f_x, f_y; \Delta z)\right] , \quad (1)$$

where $\mathcal{F}$ and $\mathcal{F}^{-1}$ denote the Fourier Transform and its inverse.

The transfer function H for a propagation distance Δz is

$$H(f_x, f_y; \Delta z) = \begin{cases} \exp\left(j\frac{2\pi}{\lambda}\Delta z\sqrt{1 - \lambda^2({f_x}^2 + {f_y}^2)}\right) , \\ \quad \text{if } f_x^2 + f_y^2 < \lambda^{-2} \\ 0 , \quad \text{otherwise} \end{cases} , \quad (2)$$

where (f_x, f_y) are the spatial frequencies of the plane waves, and λ is the wavelength of light. The ASM allows us to simulate diffraction effects between discrete layers. From now on, we define the function ASM (U, z) to propagate the complex field U by a distance z. Its spatial-domain counterpart, the Rayleigh Sommerfeld Diffraction (RSD), computes this propagation as a convolution. We will leverage Compact RSD (C-RSD) convolutions to accelerate hologram synthesis by avoiding global Fourier transformations.

3.2 Propagating Waves over MPI Layers

MPIs have been used to render holograms [Shi et al. 2022a; Zhang et al. 2017; Ziegler et al. 2008] by simulating the light propagation alongside the amplitude and phase-based occlusions of each layer. Holograms are commonly synthesized by compositing the MPI layers from the last layer ($i = N$) to the front ($i = 1$).

Let U_{i+1} be the complex field from the layer $i + 1$, propagated with the ASM method. When this field encounters the layer i, it is partially occluded by the layer's occluder o_i, with its own emission U_i considered. We model this interaction as

$$U_i' = \underbrace{\text{ASM}(U_{i+1}, \Delta z_i) \cdot o_i}_{\text{Occluded Field}} + \underbrace{U_i \cdot (1 - o_i)}_{\text{Emission}} . \quad (3)$$

This occlusion o_i can either modulate an opaque or a thin transparent object via alpha blending, where $o_i = \alpha_i$, and α_i is the transparency of layer i. The term α_i acts as an amplitude mask, attenuating both the real and imaginary components of the incoming background wave, thereby correctly simulating occlusion in the holographic reconstruction.

4 One-Step Phase Retrieval Reparameterization

Random phase holography leverages the full angular spread of light afforded by a holographic display, achieving resilience to unpredictable artifacts and natural focus cues. However, real-time methods suffer from significant contrast loss and speckle noise. In this section, we first introduce the background formulation in Section 4.1. Then, given our observation that converting a complex hologram into a phase-only hologram via amplitude discard produces a structured, predictable loss of quality when the propagation distance and phase initialization are fixed, we present our reparameterization of One-Step Phase Retrieval (OSPR) for 2D holography to counteract this degradation (Fig. 3a,b) and significantly enhance hologram quality (Fig. 3c). This reparameterization will be extended to MPIs to generate high-quality 3D holograms in Section 5, becoming the high-quality cornerstone of Instant RPH.

4.1 2D Random Phase Holography

The goal of 2D random-phase holography is to generate T time-multiplexed phases $\Phi_{\text{SLM}} = \{\phi^1_{\text{SLM}}, \dots, \phi^T_{\text{SLM}}\}$ that are multiplexed on an SLM to resemble a target image I at distance z_{SLM} as

$$\underset{\Phi_{\text{SLM}}}{\arg\min}\, \mathcal{L}(I, \hat{I}),\ \hat{I} = \frac{1}{T}\sum_{t=1}^{T}\left|\text{ASM}\left(e^{i\phi^t_{\text{SLM}}}, z_{\text{SLM}}\right)\right|^2, \tag{4}$$

where $\mathcal{L}$ is a loss function such as mean absolute error (MAE). Solving this via gradient descent is too slow for real-time applications.

Real-time OSPR combines a target image with a generated phase pattern, propagates it to the SLM plane, and converts it to a phase-only hologram. Amplitude Discard (AD) yields T independent wave fields with T phase patterns $\phi^t_{\text{u}} \sim \mathcal{U}(0, 2\pi)$ and discards the amplitude at the SLM plane as

$$\phi^t_{\text{SLM}} = \angle\text{ASM}\left(A\, e^{i\phi^t_{\text{u}}}, -z_{\text{SLM}}\right), \tag{5}$$

where $A = \sqrt{I}$ is the amplitude of the target image and $\angle$ is the angle operator. Unfortunately, as shown in Fig. 3c, this strategy causes severe contrast loss, border artifacts, and speckle noise.

4.2 Random Phase Reparameterization

To invert the artifacts introduced by random-phase and amplitude discard, we propose to reparameterize random-phase holography by modifying the target image amplitude A. While prior work [Chu et al. 2025] trained convolutional networks to generate pseudo-random phase holograms, these architectures struggle to preserve the extremely high-frequency, spatially-dependent features. Rather than performing complex-to-phase conversions at the SLM plane, we shift the optimization entirely to the *target amplitude plane*, whose modulation will flow into the displayed SLM phase.

We propose Linear Amplitude Encoders (LAEs), a lightweight, real-time approach to collaboratively counteract the quality loss. Our LAE g_t applies a learned per-pixel scale s and bias b to the target amplitude for each time-multiplexed frame t as

$$A_e(t, x, y) = g_t(A(x, y)) = s(t, x, y)A(x, y) + b(t, x, y), \tag{6}$$

where s and b counteract the specific artifacts from the fixed random phase ϕ^t_{u}. For clarity, g_t denotes this per-frame modulation, whereas g denotes modulation across all T time multiplexing frames.

We jointly optimize these time-multiplexed LAEs g_t as depicted in Fig. 3a, to resemble the target image I:

$$\underset{\{g_1,\dots,g_T\}}{\arg\min}\, \mathcal{L}\left(I, \frac{1}{T}\sum_{t=1}^{T}\left|\text{ASM}\left(e^{i\phi^t_{\text{SLM}}}, z_{\text{SLM}}\right)\right|^2\right), \tag{7}$$

$$\phi^t_{\text{SLM}} = \angle\text{ASM}\left(g_t(A)e^{i\phi^t_{\text{u}}}, -z_{\text{SLM}}\right). \tag{8}$$

Our LAEs (Fig. 3b) are well-suited for real-time holography since, in inference, they only involve a per-frame multiplication and an addition to the target amplitude. Furthermore, our method is the first to generate time-multiplexed holograms in real-time that are collaboratively trained to maximize image quality. See the Supplemental Material for complete training details.

4.3 Results

We evaluate our LAE reparameterization ($T = 8$) against naive Amplitude Discard and Real-Time ARH [Chu et al. 2025]. Our LAEs maximize image quality for the specific configuration used ($z_{\text{SLM}} =$ 100 mm), requiring retraining when the optical configuration is modified. As shown in Fig. 3c, our collaboratively learned LAEs effectively correct color shifts and suppress speckle noise, improving image quality by up to 10 dB over naive amplitude discard and up to 7 dB over Real-Time ARH. This unlocks real-time, high-quality 2D random phase holography and serves as the foundation for our volumetric approach.

5 Real-Time Volumetric Random Phase Hologram

Given our LAE for 2D random phase holography, we develop Instant Random Phase Holography (Instant RPH), our real-time 3D hologram generator with natural defocus and parallax cues. We extend our 2D formulation to 3D by leveraging Multiplane Images (MPIs) as a natural backbone for efficient and true 3D holography. However, in 2D, LAEs only correct intra-layer speckle (from other regions of the image), whereas in 3D, there is also inter-layer speckle, arising from the interference between MPI layers. Furthermore, real-time performance and perceptual accuracy are often thought of as a trade-off. We draw relations to perception theory to achieve a perceptually-driven MPI that can be efficiently synthesized into a hologram.

As illustrated in Fig. 4, our Instant RPH synthesizes holograms through three primary stages: (a) scene representation via a diopter-space Multiplane Image (MPI) capable of rasterizing both meshes and 3DGS; (b) complex field initialization with LAE and fixed random phase, propagated using efficient wave propagation via Compact Rayleigh Sommerfeld Diffraction convolutions (C-RSD); and (c) phase-only hologram formation through amplitude discard.

We choose MPIs as our intermediate representation. We design our slicing of the MPI to match the optical characteristics of the human eye and the holographic display. We then leverage the underlying optical properties of our slicing strategy to use a fast diffraction operator to propagate the complex field between layers. This enables real-time holographic rendering with a depth layer resolution that accurately samples the eye's accommodation range. Key to the efficiency of our method is to rely on our learned amplitude correction

Algorithm 1. **Instant RPH** generates high-quality holograms with natural defocus and parallax cues in real-time in three main steps: (1) diopter-space MPI rendering to extract scene amplitudes and occluders, (2) layer-wise wave composition using efficient occlusion-aware propagation (*Compact RSD*, C-RSD), where amplitudes are modulated by our jointly-optimized LAEs (g) and combined with a fixed random phase (ϕ), and (3) final phase-only conversion via Amplitude Discard ($\angle$).

def Instant RPH(scene, $(g, \phi)[L, T, C, X, Y]$)**:**
    $A[L, C, X, Y], o[L, X, Y] \leftarrow$ **RenderDiopterMPI**(scene)
    $U'[T, C, X, Y] \leftarrow 0$
    $A_e[L, T, C, X, Y] \leftarrow g(A)$
    for $l = L \rightarrow 1$ **do**
        $o_l \leftarrow o[l, \dots]$
        $U_l \leftarrow A_e[l, \dots] \cdot e^{i\phi[l, \dots]}$
        $U' \leftarrow \text{C-RSD}(U', \Delta z) \cdot o_l + U_l(1 - o_l)$
    $\Phi_{\text{SLM}} \leftarrow \angle\text{ASM}\left(U', -z_{\text{SLM}}\right)$
    return Φ_{SLM}

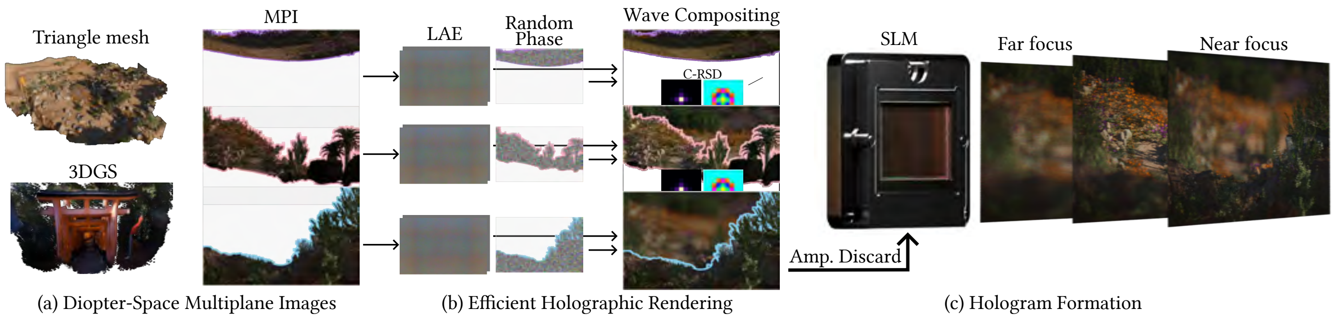

Fig. 4. **Overview of Instant RPH.** Our framework acts as a bridge between modern graphics engines and holographic hardware. (a) We take in arbitrary 3D content, such as meshes or 3D Gaussian Splats, and render it into a diopter-space Multiplane Image (MPI), decoupling scene complexity from rendering cost. We use 21 layers in our method to match the depth-of-field of the human eye (we only display three for simplicity). (b) We convert the MPI to a complex field with a fixed random phase and our Linear Amplitude Encoders (LAE) to enhance hologram quality and reduce speckle noise. To avoid costly FFT-based ASM propagation, we rely on complex-field propagation between layers with efficient Compact RSD (C-RSD) convolutions. (c) The complex hologram is displayed on an SLM through amplitude discard, allowing real-time streaming of artifact-resilient 3D content with natural defocus and parallax cues.

through a per-pixel scale and bias, and the use of Amplitude Discard [Maimone et al. 2017] for phase retrieval. We allow for a direct, one-step hologram generation that maintains high reconstructive fidelity without expensive iterative solvers.

Instant RPH is a *differentiable, real-time 3D hologram generator*, allowing us to jointly optimize the LAEs per time-multiplexed frames and per-layer, to maximize focal stack quality. We provide the complete hologram generation details of Instant RPH in Algorithm 1.

5.1 Diopter-space MPI in CGH

Human accommodation is characterized in diopters ($D = d^{-1}$, with d in meters), where the depth-of-field (DoF) is roughly constant in diopter-space. Vision science indicates that the DoF for the human eye is $0.3D$ [Campbell 1957; Marcos et al. 1999], and has been leveraged in multifocal displays [Narain et al. 2015] to achieve correct perceptual visual cues. Let $\mathcal{D}$ be the total diopter range of the display ($\mathcal{D} = 6$ in our implementation). We ensure our holographic display provides continuous accommodation cues without perceptible stepping artifacts by setting the sampling interval $\Delta d = 0.31D$ and, therefore, the number of required layers N as

$$N = \lceil \frac{\mathcal{D}}{\Delta d} \rceil + 1 \, . \tag{9}$$

This intermediate MPI representation decouples the render time from the geometries and holograms.

Furthermore, according to the thin-lens equation, for a standard eyepiece ($f = 25$ mm), the propagation distance between adjacent MPI layers ($\Delta d = 0.3D$) corresponds to a physically small spacing of ~ 0.25 mm. Evaluating Eq. (3) using the standard FFT-based ASM Eq. (1) in holography requires two costly FFTs per layer. For such short propagation distances, the diffraction kernel in the frequency domain is highly oscillatory, but its spatial impulse response (the convolution kernel) is compact. We replace the costly global FFT with a Compact RSD (C-RSD) convolution with small kernel size $K = 5$, which reduces the complexity of the propagation step from $O(HW \log(HW))$ to $O(HWK^2)$ and allows real-time hologram synthesis.

We refer the reader to the Supplemental Material for details on diopter-space MPI rendering, as well as an analysis of the efficiency of Compact RSD convolutions.

5.2 End-to-end, jointly-optimized Instant RPH

Instant RPH is naturally differentiable, which allows us to perform end-to-end training of our LAEs from Section 4.2. To account for the extra dimension (number of layers L), we also add this extra dimension in our fixed random phases and amplitude correctors (s and b) to effectively account for the number of layers, as shown in Algorithm 1. By training our LAEs in this manner, we are able to counteract not only intra-layer speckle, but also inter-layer speckle.

We generate a dataset consisting of randomly spawned objects and camera positions. We render a diopter-space MPI matching our 6 diopters range, 21 layers, and an image resolution of 1280×800 pixels. We refer to the Supplemental Material for details on dataset generation and data augmentation strategies. We co-optimize our LAEs by computing the L1 loss over all the distances of the target focal stack f as

$$\begin{aligned} \mathcal{L} &= \frac{1}{L} \sum_{l=1}^{L} \left| \hat{f}_l \frac{\mu(f_l)}{\mu(\hat{f}_l)} - f_l \right|, \\ \hat{f}_l &= \frac{1}{T} \sum_{t=1}^{T} \left| \mathrm{ASM}\left(e^{i\phi^t_{\mathrm{SLM}}}, z_{\mathrm{SLM}} - l\Delta z \right) \right|^2, \end{aligned} \tag{10}$$

where the phase patterns $\Phi_{\mathrm{SLM}} = \{\phi^1_{\mathrm{SLM}}, \phi^2_{\mathrm{SLM}}, \dots, \phi^T_{\mathrm{SLM}}\}$ are obtained as described in Algorithm 1, and we color-balance adjusting the means (μ) with the corresponding image from the focal stack. We use $T = 24$ for maximum quality in our results while keeping the real-time performance. The final LAEs, g, are composed of textures s and b of size $L \times T \times C \times H \times W$, which in our settings adds up to a total of 3 billion parameters. We distribute the training of each wavelength across four NVIDIA H100 GPUs. We train our method for 20 epochs, which takes around twelve hours.

6 Assessment

We assess the merits of our method with both high-fidelity simulations and experimental results. Specifically, we validate the following three key claims: (1) our LAEs improve image quality in 3D Holography, (2) Instant RPH produces optically correct accommodation defocus, and (3) parallax cues that are lost in standard RGB-D methods. Results are generated for a holographic display (both in simulation and on the experimental prototype) with a pixel pitch of 10.8 μm and a resolution of 1280×800, targeting a physical distance to the SLM of 95 mm to 100 mm, which corresponds to a perceived range of 0 to 6 diopters.

We compare our Instant RPH with *Artifact-Resilient Holography* (ARH) [Chu et al. 2025], the recent, state-of-the-art real-time random phase method. Instant RPH not only reaches a 7× depth resolution, but a 4.15× speed-up as well. We thus can afford to time-multiplex 24 phase patterns [Schiffers et al. 2025], as opposed to the eight phase patterns in ARH. On an NVIDIA L40 GPU, the measured FPS for three wavelengths (RGB) and 21 focal distances/images in the measured critical execution-path regime ($T = 1$) is 7.6 FPS for Real-Time ARH and 54 FPS for our method. In the maximum quality settings shown in the results, Real-Time ARH ($T = 8$) runs at 1 FPS, and Instant RPH ($T = 24$) runs at 5 FPS. We have developed our method with PyTorch JiT, so it could be further optimized using specific CUDA kernels.

Hardware Prototype. Our experimental prototype uses a 4-bit SLM (Texas Instruments DLP670S with DLPC900 EVM) in a 4F optical relay configuration. A three-wavelength fiber-coupled laser source (wavelengths 450 nm, 520 nm, 660 nm) illuminated the SLM via a beam splitter. The modulated wavefront was relayed through a 4F system (f=100 mm and f=80 mm) to a monochrome sensor, enabling spatial filtering at the Fourier plane to suppress undiffracted light and higher diffraction orders. For focal sweep imaging, the sensor was mounted on a motorized translation stage. For parallax imaging, we added an eyepiece (f=75 mm) and camera lens (f=12.5 mm). Holographic images were captured sequentially at each wavelength and digitally combined to generate full-color reconstructions. See the Supplemental Material for full details.

6.1 Accommodation and Defocus Cues

We compare our method in *simulation and experimentally* in three different scenes made up of triangle meshes [Haven 2026] and one 3DGS scene [Kerbl et al. 2023]. We compare our results in Fig. 5 with ARH [Chu et al. 2025], the state-of-the-art in real-time, random phase holography. ARH operates at a lower rate of 7.6 FPS (RGB, $T = 1$) to process 21 distances from an RGB-D input, due to their neural architecture and ASM propagation of each layer, preventing it from achieving real-time performance. Our LAEs significantly improve color quality and speckle reduction, mostly noticeable in homogeneous areas, such as the Bulldozer's shovel. Furthermore, ARH does not handle occlusion between layers, a problem that is aggravated when considering parallax in the next section; this leads to dark artifacts surrounding different depth layers.

In contrast, our Instant RPH improves up to 54 FPS (RGB, $T = 1$), handling the same number of layers (21) from our diopter-space MPI. Moreover, due to our efficient propagation between layers, our Instant RPH shows artifact-free, correct occlusions, while enabling natural, continuous focus cues within the limits of human depth perception.

When the virtual camera is focused on the foreground (Near Focus), the background naturally diffuses into a bokeh-like blur. Conversely, focusing on the background renders the foreground blurred. The amount of blur scales with the dioptric distance from the focal plane (we show three of the 21 layers), showing how the combination of our MPI slicing strategy and our efficient per-layer propagation effectively reconstructs the light field in real time. The high-frequency details of the in-focus regions are well preserved, eliminating the need for iterative Gerchberg-Saxton-style refinement [Choi et al. 2022; Gerchberg 1972; Kim et al. 2024]. We refer the reader to the Supplemental Material for a comparison with iterative methods, ablation studies on end-to-end training, speckle improvement from LAE in real captures, and a video comparison of the results.

6.2 View-Dependent Parallax

A critical limitation of current real-time random phase holographic pipelines [Chu et al. 2025] is their reliance on RGB-D (2.5D) data, which lacks information about occluded geometry [Matsushima et al. 2014]. Therefore, these holograms show incorrect parallax cues with visible artifacts and produce incorrect results when observed from a different, non-centered perspective. Parallax cues in near-eye holographic displays are subtle, but crucial for an accurate immersive experience [Kim et al. 2024].

To demonstrate how our Instant RPH presents correct parallax cues and improves previous methods, we compare it in simulation and experimentally in Fig. 6 by recreating a pupil moving from left to right. In Fig. 6a, we show a basic test on parallax, where a set of 3×3 pillars are axially aligned. Our method leverages the MPI to yield correct view-dependent parallax, so the hidden pillars are shown when observed from the side. Then, in Fig. 6b we show the parallax effect in a more complex scene, when the pupil location affects the separation between objects at different layers (depths). Please refer to the Supplemental Material for dynamic visualizations of perceived parallax.

We then compare our results against ARH [Chu et al. 2025] in an experimental capture in Fig. 6c. ARH relies on input RGB-D data, yielding dark fringes even when the view is centered. These dark fringes appear from the lack of information from occluded geometry and are view-dependent; thus, they become even more visible when the geometry is observed from the side (left and right pupil). In contrast, our Instant RPH encodes the necessary, view-dependent parallax cues, so that, as the pupil moves from left to right, the background pattern is correctly revealed without artifacts. This confirms that our method enables view-dependent parallax effects that are not possible with previous real-time random phase methods. See the Supplemental Material for video visualizations.

7 Conclusion and Future Work

We first presented our random phase holography reparameterization to collaboratively generate high-quality 2D holograms in real-time, by training a set of LAEs that encode the target image to instantly generate holograms that minimize speckle noise and maximize image quality. Then, we presented Instant RPH. By extending our reparameterization to 3D holography through diopter-space MPIs to learn the amplitude encoding of each MPI layer, and combining it with efficient propagation between layers, we instantly generate high-quality holograms that represent the volumetric effects of the scene. We achieve high-quality, immersive 3D visuals with accurate accommodation and defocus cues, and view-dependent parallax at 54 FPS by shifting the holographic synthesis paradigm from iterative optimization to a direct, graphics-aware pipeline based on human perception.

Random phase holograms naturally preserve continuous defocus cues and parallax, and spread intensity over the eyebox, allowing them to remain robust against the dynamic obstructions and optical aberrations common in wearable displays [Chu et al. 2025]. We rely on temporal multiplexing and our reparameterization to mitigate speckle noise. However, there are other sources of image degradation, such as mismatches in the laser power distribution, SLM field-fringing, or imperfections in the optical system, which we do not currently include in our forward model. We believe future work could explore combining our reparameterization with hardware-aware models [Choi et al. 2022; Peng et al. 2020] to maximize experimental image quality. Other alternatives to explore include: leveraging 3D networks to analyze the inter-layer speckle and adjust the LAEs, using partially coherent light sources [Peng et al. 2021], developing phase initializations that leverage temporal multiplexing, or combining our method with recent work on multiplexing of sources [Kuo et al. 2023] or wavelengths [Schiffers et al. 2025]. Another promising direction is to integrate our method with simultaneous color CGH, which removes the frame-rate penalty of time-sequential RGB illumination and preserves the full SLM frame rate [Markley et al. 2023]. The recovered temporal bandwidth can then be spent on eye-box expansion via pupil steering/pupil scanning [Jang et al. 2018].

We hope that our work helps take a step towards transforming holographic displays into robust, plug-and-play monitors, enabling fully interactive and immersive spatial displays.

Acknowledgments

We thank German Marin for his help with initial prototyping. This work has been supported by grant PID2025-169453NB-I00, funded by MICIU/AEI/10.13039/501100011033 and by ERDF, EU. Oscar Pueyo-Ciutad was supported by the FPU22/02432 predoctoral grant. Ethan Tseng was supported by a Google PhD Fellowship. This research is funded in part by the Gordon and Betty Moore Foundation. Felix Heide was supported by an NSF CAREER Award (2047359), a Packard Foundation Fellowship, a Sloan Research Fellowship, a Sony Young Faculty Award, a Project X Innovation Award, a Amazon Science Research Award, and a Bosch Research Award. Felix Heide is a co-founder of Algolux (now Torc Robotics), Head of AI at Torc Robotics, and a co-founder of Cephia AI.

References

Kaan Akşit and Yuta Itoh. 2023. HoloBeam: Paper-Thin Near-Eye Displays. In *2023 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*. IEEE, Shanghai, China, 581–591. doi:10.1109/VR55154.2023.00073

F.W. Campbell. 1957. The Depth of Field of the Human Eye. *Optica Acta: International Journal of Optics* 4, 4 (1957), 157–164. arXiv:https://doi.org/10.1080/713826091 doi:10.1080/713826091

Praneeth Chakravarthula, Yifan Peng, Joel Kollin, Henry Fuchs, and Felix Heide. 2019. Wirtinger Holography for Near-Eye Displays. *ACM Trans. Graph.* 38, 6, Article 213 (nov 2019), 13 pages. doi:10.1145/3355089.3356539

Praneeth Chakravarthula, Ethan Tseng, Henry Fuchs, and Felix Heide. 2022. Hogel-Free Holography. *ACM Trans. Graph.* 41, 5, Article 178 (Oct. 2022), 16 pages. doi:10.1145/3516428

Kenneth Chen, Anzhou Wen, Yunxiang Zhang, Praneeth Chakravarthula, and Qi Sun. 2025. View synthesis for 3D computer-generated holograms using deep neural fields. *Opt. Express* 33, 9 (May 2025), 19399–19408. doi:10.1364/OE.559364

Suyeon Choi, Brian Chao, Jacqueline Yang, Manu Gopakumar, and Gordon Wetzstein. 2025a. Gaussian Wave Splatting for Computer-Generated Holography. *ACM Trans. Graph.* 44, 4, Article 57 (July 2025), 13 pages. doi:10.1145/3731163

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* (Vancouver, BC, Canada) *(SIGGRAPH '22)*. Association for Computing Machinery, New York, NY, USA, Article 32, 9 pages. doi:10.1145/3528233.3530734

Suyeon Choi, Manu Gopakumar, Yifan Peng, Jonghyun Kim, and Gordon Wetzstein. 2021. Neural 3D holography: Learning accurate wave propagation models for 3D holographic virtual and augmented reality displays. *ACM Transactions on Graphics (TOG)* 40, 6 (2021), 1–12.

Suyeon Choi, Changwon Jang, Douglas Lanman, and Gordon Wetzstein. 2025b. Synthetic aperture waveguide holography for compact mixed-reality displays with large étendue. *Nature Photonics* 19, 8 (01 Aug 2025), 854–863. doi:10.1038/s41566-025-01718-w

Victor Chu, Oscar Pueyo-Ciutad, Ethan Tseng, Florian Schiffers, Grace Kuo, Nathan Matsuda, Albert Redo-Sanchez, Lanman Douglas, Oliver Cossairt, and Felix Heide. 2025. Artifact-Resilient Real-Time Holography. *ACM Trans. Graph.* 44, 6, Article 219 (Dec. 2025), 13 pages. doi:10.1145/3763361

N. Delen and B. Hooker. 1998. Free-space beam propagation between arbitrarily oriented planes based on full diffraction theory: a fast Fourier transform approach. *J. Opt. Soc. Am. A* 15, 4 (Apr 1998), 857–867. doi:10.1364/JOSAA.15.000857

X. Ding, C. Chang, B. Dai, Q. Wang, D. Zhang, and S. Zhuang. 2024. Real-time holographic 3D display using Split-Lohmann Fresnel computer-generated hologram. *Optics Express* 32, 23 (2024), 40175–40189. doi:10.1364/OE.534359

Simon Fricke, Christian Rautenberg, Rafał K. Mantiuk, Jan Błaszczyk, Karol Myszkowski, and Oliver Bimber. 2025. Real-Time Rendering Framework for Holography. *Computer Graphics Forum* 44 (2025), e70058. doi:10.1111/cgf.70058

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

Joseph W Goodman. 2005. *Introduction to Fourier optics.* Roberts and Company publishers, Greenwood Village, CO.

Poly Haven. 2026. Poly Haven. https://polyhaven.com/ Accessed: 2026-01-21.

Changwon Jang, Kiseung Bang, Gang Li, and Byoungho Lee. 2018. Holographic near-eye display with expanded eye-box. *ACM Transactions on Graphics (TOG)* 37, 6 (2018), 1–14.

Koray Kavakli, Yuta Itoh, Hakan Urey, and Kaan Aksit. 2023. Realistic Defocus Blur for Multiplane Computer-Generated Holography . In *2023 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*. IEEE Computer Society, Los Alamitos, CA, USA, 418–426. doi:10.1109/VR55154.2023.00057

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkuehler, and George Drettakis. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Trans. Graph.* 42, 4, Article 139 (July 2023), 14 pages. doi:10.1145/3592433

Dongyeon Kim, Seung-Woo Nam, Suyeon Choi, Jong-Mo Seo, Gordon Wetzstein, and Yoonchan Jeong. 2024. Holographic Parallax Improves 3D Perceptual Realism. *ACM Trans. Graph.* 43, 4, Article 68 (July 2024), 13 pages. doi:10.1145/3658168

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 Transactions on Graphics (TOG)* 41, 4 (2022), 1–15.

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. doi: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.

Grace Kuo, Laura Waller, Ren Ng, and Andrew Maimone. 2020. High-Resolution Étendue Expansion for Holographic Displays. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 1–9. https://dl.acm.org/doi/abs/10.1145/3386569.3392414

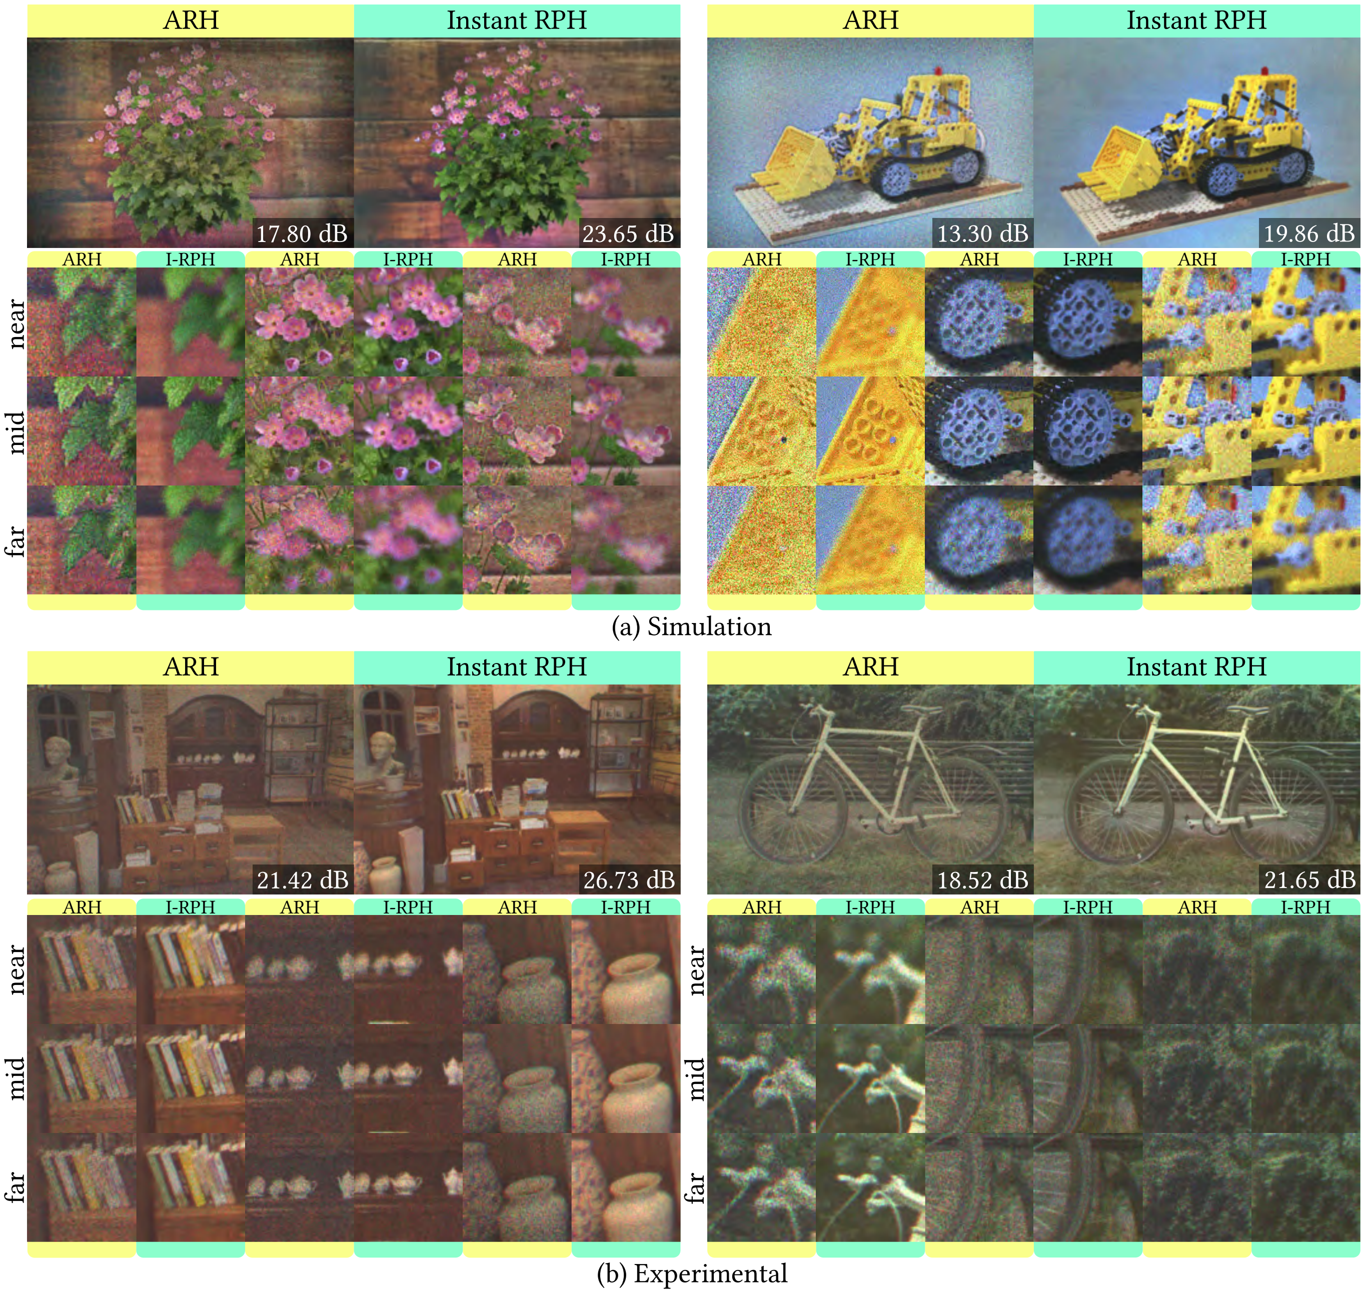

Fig. 5. **Defocus Fidelity and Performance Comparison [Simulation (a) and Experimental (b)].** We compare the reconstruction quality and runtime speed of our Instant RPH against ARH [Chu et al. 2025], the state-of-the-art real-time random phase method. We display the field of view for four complex scenes and show insets visually refocused at *near*, *mid*, and *far* depths (independent of horizontal position). While ARH fails to scale to dense optical cues of 21 layers in quality and real-time (7.6 FPS for $T = 1$, 1 FPS for their maximum quality settings of $T = 8$), our method supports dense distinct layers (21 layers) at real-time rates (**54** FPS for $T = 1$, 5 FPS for our maximum quality settings of $T = 24$), improving PSNR by an average of 5–6 dB. Note how our approach yields smoother optical blur in out-of-focus regions and reduced speckle noise compared to the baseline, confirming the effectiveness of our diopter-space MPI for immersive 3D displays and the speed to enable 24× time multiplexing.

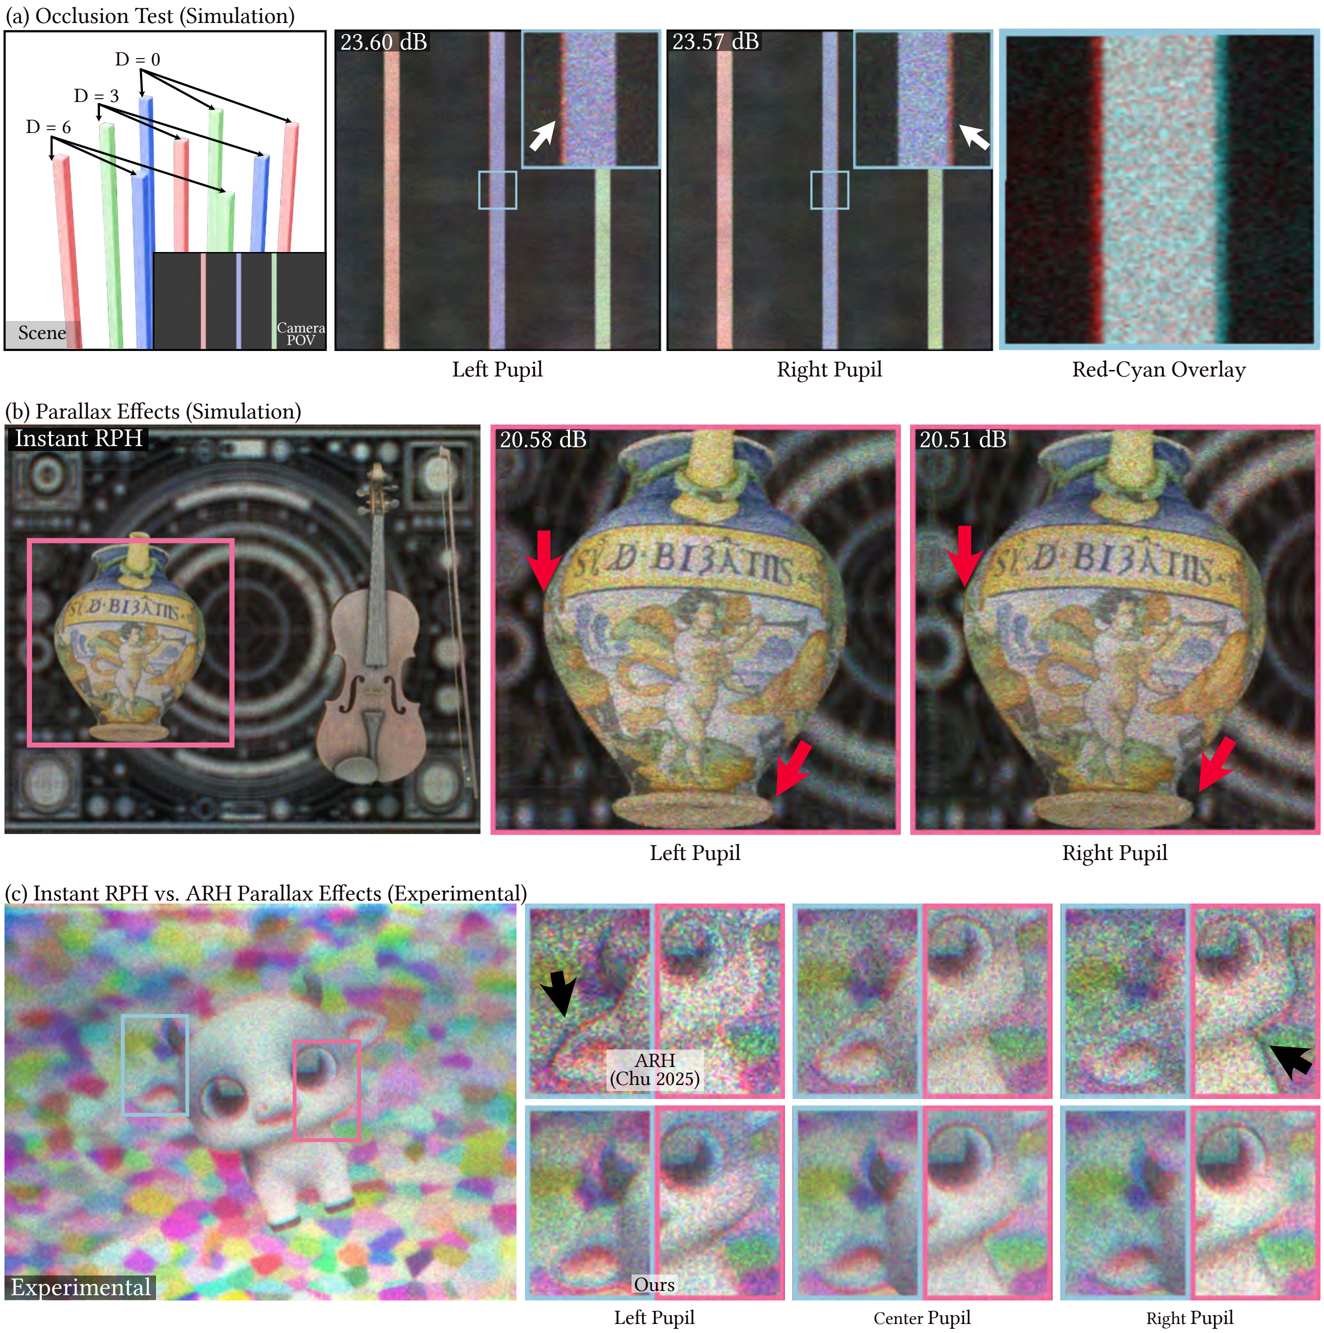

Fig. 6. **Parallax and Occlusion Validation [Simulation (a and b) and Experimental Capture (c)].** (a: Simulation): "Pillars" deep occlusion test. The schematic (left) shows varying diopters (D). We visualize disparity using a red-cyan overlay, where the red channel encodes the left pupil perspective and the cyan channel encodes the right pupil perspective. Regions visible in only one view appear as pure red or cyan. White arrows highlight red/green pillars revealed behind the blue foreground by the shifting viewpoints. (b: Simulation): Instant RPH shows accurate parallax. Red arrows mark view-dependent geometry (e.g., stars behind vase) consistent across pupil views. (c: Experimental): "Cow" comparison against ARH [Chu et al. 2025]. Crops show our MPI approach yields higher fidelity. Note ARH's incorrect border artifacts (due to RGB-D input) vs. our clean occlusion boundaries.

Antoine Lagrange, Antonin Gilles, Kevin Heggarty, Bruno Fracasso, Patrick Gioia, and Vincent Savaux. 2025. Real-time calculation of Fresnel holograms using the split-Lohmann method. *Optics Letters* 50, 8 (2025), 2554–2557. doi:10.1364/OL.554228

Yilong Li, Haokun Xiong, Zhiling Guo, Jie Ding, and Di Wang. 2025. An Efficient Hologram Generation Method via Multi-Layer WRPs and Optimal Segmentation. *Electronics* 14, 23 (2025), 4591. doi:10.3390/electronics14234591

Jiawei Liu, Jiabao Chen, and Xiang Peng. 2023. Hybrid hologram generation using point-based and polygon-based computer-generated holography. *Optics Letters* 48, 12 (2023), 3339–3342. doi:10.1364/OL.492745

Xiandong Ma, Jinbin Gui, Junchang Li, and Qinghe Song. 2024. A Layered Method Based on Depth of Focus for Rapid Generation of Computer-Generated Holograms. *Applied Sciences* 14, 12 (2024), 16 pages. doi:10.3390/app14125109

Kevin J MacKenzie, David M Hoffman, and Simon J Watt. 2010. Accommodation to multiple-focal-plane displays: Implications for improving stereoscopic displays and for accommodation control. *J Vis* 10, 8 (July 2010), 22.

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), 85.

Susana Marcos, Esther Moreno, and Rafael Navarro. 1999. The depth-of-field of the human eye from objective and subjective measurements. *Vision Research* 39, 12 (1999), 2039–2049. doi:10.1016/S0042-6989(98)00317-4

Eric Markley, Nathan Matsuda, Florian Schiffers, Oliver Cossairt, and Grace Kuo. 2023. Simultaneous Color Computer Generated Holography. In *SIGGRAPH Asia 2023 Conference Papers* (Sydney, NSW, Australia) *(SA '23)*. Association for Computing Machinery, New York, NY, USA, Article 22, 11 pages. doi:10.1145/3610548.3618250

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. *Opt. Express* 22, 20 (Oct 2014), 24450–24465. doi:10.1364/OE.22.024450

Kyoji Matsushima and Tomoyoshi Shimobaba. 2009. Band-Limited Angular Spectrum Method for Numerical Simulation of Free-Space Propagation in Far and Near Fields. *Opt. Express* 17, 22 (Oct 2009), 19662–19673. doi:10.1364/OE.17.019662

T. Möller, E. Haines, and N. Hoffman. 2018. *Real-time Rendering*. CRC Press, Boca Raton, FL. https://books.google.es/books?id=ia9LvgAACAAJ

Rahul Narain, Rachel A. Albert, Abdullah Bulbul, Gregory J. Ward, Martin S. Banks, and James F. O'Brien. 2015. Optimal presentation of imagery with focus cues on multi-plane displays. *ACM Trans. Graph.* 34, 4, Article 59 (July 2015), 12 pages. doi:10.1145/2766909

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.

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 Trans. Graph.* 39, 6, Article 185 (nov 2020), 14 pages. doi:10.1145/3414685.3417802

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 Computer Society, Los Alamitos, CA, USA, 1–12. doi:10.1109/ICCP56744.2023.10233716

Florian Andreas Schiffers, Grace Kuo, Nathan Matsuda, Douglas Lanman, and Oliver Cossairt. 2025. HoloChrome: Polychromatic Illumination for Speckle Reduction in Holographic Near-Eye Displays. *ACM Trans. Graph.* 44, 3, Article 26 (May 2025), 18 pages. doi:10.1145/3732935

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. 2022a. End-to-end learning of 3D phase-only holograms for holographic display. *Light Sci Appl* 11, 1 (Aug. 2022), 247.

Liang Shi, Richard Webb, Lei Xiao, Changil Kim, and Changwon Jang. 2022b. Neural compression for hologram images and videos. *Optics Letters* 47, 22 (2022), 6013–6016.

Tomoyoshi Shimobaba, Nobuyuki Masuda, and Tomoyoshi Ito. 2009. Simple and fast calculation algorithm for computer-generated hologram with wavefront recording plane. *Opt Lett* 34, 20 (Oct. 2009), 3133–3135.

Athanasia Symeonidou, David Blinder, Adrian Munteanu, and Peter Schelkens. 2015. Computer-generated holograms by multiple wavefront recording plane method with occlusion culling. *Opt. Express* 23, 17 (Aug 2015), 22149–22161. doi:10.1364/OE.23.022149

P.W.M. Tsang and T. C. Poon. 2013. Novel method for converting digital Fresnel hologram to phase-only hologram based on bidirectional error diffusion. *Opt. Express* 21, 20 (Oct 2013), 23680–23686. doi:10.1364/OE.21.023680

Ethan Tseng, Grace Kuo, Seung-Hwan Baek, Nathan Matsuda, Andrew Maimone, Florian Schiffers, Praneeth Chakravarthula, Qiang Fu, Wolfgang Heidrich, Douglas Lanman, et al. 2024. Neural étendue expander for ultra-wide-angle high-fidelity holographic display. *Nature communications* 15, 1 (2024), 2907.

Richard Tucker and Noah Snavely. 2020. Single-View View Synthesis With Multiplane Images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Seattle, WA, USA, 548–557. doi:10.1109/CVPR42600.2020.00063

Fan Wang, Tomoyoshi Ito, and Tomoyoshi Shimobaba. 2023. High-speed rendering pipeline for polygon-based holograms. *Photonics Research* 11, 2 (2023), 313–321. doi:10.1364/PRJ.474158

Suttisak Wizadwongsa, Pakkapon Phongthawee, Jiraphon Yenphraphai, and Supasorn Suwajanakorn. 2021. NeX: Real-time View Synthesis with Neural Basis Expansion. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Nashville, TN, USA, 8530–8539. doi:10.1109/CVPR46437.2021.00843

Daeho Yang, Wontaek Seo, Hyeonseung Yu, Sun Il Kim, Bongsu Shin, Chang-Kun Lee, Seokil Moon, Jungkwuen An, Jong-Young Hong, Geeyoung Sung, and Hong-Seok Lee. 2022. Diffraction-engineered holography: Beyond the depth representation limit of holographic displays. *Nature Communications* 13, 1 (12 Oct 2022), 6012. doi:10.1038/s41467-022-33728-5

Fan Yang, Andrzej Kaczorowski, and Timothy D. Wilkinson. 2014. Fast precalculated triangular mesh algorithm for 3D binary computer-generated holograms. *Applied Optics* 53, 35 (2014), 8261–8267. doi:10.1364/AO.53.008261

Yicheng Zhan, Dong-Ha Shin, Seung-Hwan Baek, and Kaan Akşit. 2026. Complex-Valued Holographic Radiance Fields. *ACM Transactions on Graphics* 45, 3 (April 2026), 1–16. doi:10.1145/3804450

Hao Zhang, Liangcai Cao, and Guofan Jin. 2017. Computer-generated hologram with occlusion effect using layer-based processing. *Appl. Opt.* 56, 13 (May 2017), F138–F143. doi:10.1364/AO.56.00F138

Ya-Ping Zhang, Fan Wang, Ting-Chung Poon, Shuang Fan, and Wei Xu. 2018. Fast generation of full analytical polygon-based computer-generated holograms. *Optics Express* 26, 15 (2018), 19206–19224. doi:10.1364/OE.26.019206

Jie Zhou, Shuyang Xie, Lei Jiang, Yimou Luo, Yang Wu, and Jun Wang. 2025. Lateral-shift-free Split-Lohmann computer-generated holography. *Optics Letters* 50, 21 (2025), 6558–6561. doi:10.1364/OL.573768

Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. 2018. Stereo magnification: learning view synthesis using multiplane images. *ACM Trans. Graph.* 37, 4, Article 65 (July 2018), 12 pages. doi:10.1145/3197517.3201323

Remo Ziegler, Simone Croci, and Markus Gross. 2008. Lighting and Occlusion in a Wave-Based Framework. *Computer Graphics Forum* 27, 2 (2008), 211–220. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-8659.2008.01118.x doi:10.1111/j.1467-8659.2008.01118.x