

HoloPathTracer: Fast and Accurate Wave Path Tracing for Holography

WENBIN ZHOU*, The University of Hong Kong, Hong Kong SAR
 XIANGYU MENG*, The University of Hong Kong, Hong Kong SAR
 JIANKAI XING, The University of Hong Kong, Hong Kong SAR and Tsinghua University, China
 XIN LIU, The University of Hong Kong, Hong Kong SAR
 SUYEON CHOI, Stanford University, USA and Seoul National University, Republic of Korea
 YIFAN PENG, The University of Hong Kong, Hong Kong SAR

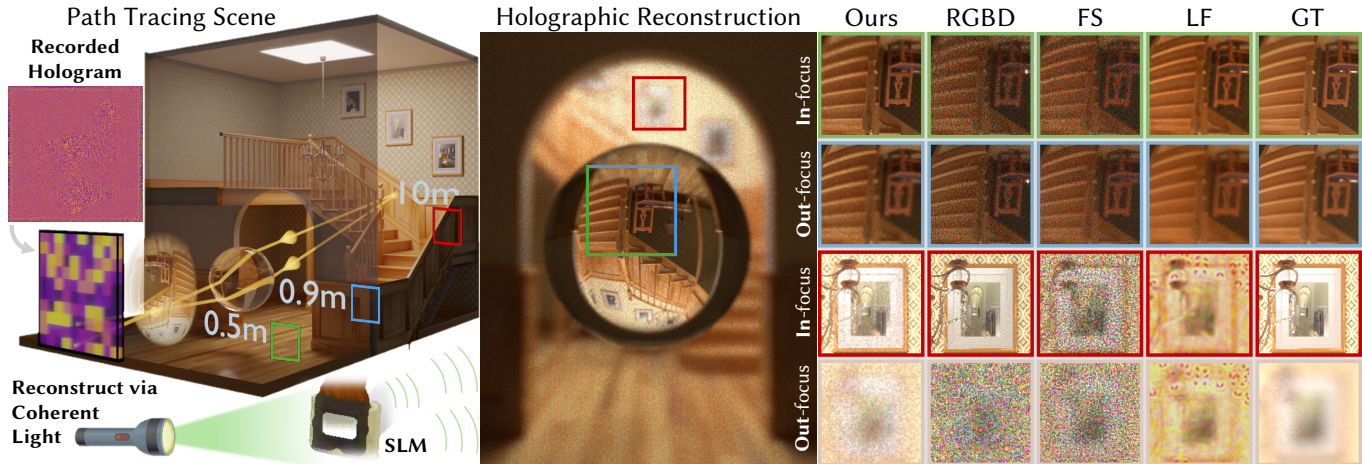


Fig. 1. Technical footprint and reconstruction results of our HoloPathTracer. **Left:** Wave field rendering from a target scene to the hologram recording plane, with physically accurate visual cues encoded. The resulting hologram is displayed on an SLM and illuminated by coherent light for optical reconstruction. **Center:** Reconstruction (simulation) via propagating the hologram-modulated wavefront to the observation plane, with viewpoint near-focused at the imagery of the stairs via the glass sphere. **Right:** Close-up comparisons of our method and existing CGH baselines, validating superior visual fidelity (e.g., natural defocus, glossy highlight) at varying physically accurate depths (0.5m, 0.9m, and 10m). FS: Focal stack; LF: Light field; GT: Ground truth from Mitsuba3.

Holography offers unique advantages for delivering perceptual realism while preserving compact form factors in VR/AR. Its perceptual quality, however, hinges on encoding rich wavefronts of photorealistic scenes into interference patterns and then incoherently multiplexing the resulting wave fields for perception. Existing CGH paradigms decouple radiance estimation from wave propagation by pre-rendering radiance on discretized scene sectors. This separation between radiometric and wave-optical computation inherently limits the range of focus cues and visual effects that can be faithfully reproduced, including depth- and view-continuity, and physically based material behaviors such as glossy or mirror-like reflection and refraction.

*denotes equal contribution.

Authors' Contact Information: Wenbin Zhou, The University of Hong Kong, Hong Kong SAR, zhouwb@connect.hku.hk; Xiangyu Meng, The University of Hong Kong, Hong Kong SAR, mengxy22@connect.hku.hk; Jiankai Xing, The University of Hong Kong, Hong Kong SAR and Tsinghua University, China, xjk21@mails.tsinghua.edu.cn; Xin Liu, The University of Hong Kong, Hong Kong SAR, liuxin24@hku.hk; Suyeon Choi, Stanford University, USA and Seoul National University, Republic of Korea, suyeon@stanford.edu; Yifan Peng, The University of Hong Kong, Hong Kong SAR, evanpeng@hku.hk.



This work is licensed under a Creative Commons Attribution 4.0 International License.
 © 2026 Copyright held by the owner/author(s).
 ACM 1557-7368/2026/7-ART39
<https://doi.org/10.1145/3811351>

We present a *physically accurate yet computationally efficient* wave optics rendering framework leveraging path tracing to encode full 3D visual cues into phase holograms. Specifically, we employ a Monte Carlo method to solve both the rendering equation and the Rayleigh–Sommerfeld integral simultaneously. Our algorithm is fully compatible with modern graphics techniques and can generate multiple time-multiplexed random holograms with minimal additional time cost via *Path Reuse*. By employing a fast approximation with an ambient radiance cache, we realize an order of magnitude convergence speed improvement. The resulting coherent wave fields that inherently encode comprehensive visual effects are converted into phase-only holograms under complex-amplitude supervision. Through extensive simulations and experimental validations on a spatial light modulator-based display prototype, we demonstrate faithful holographic reconstructions of natural 3D cues and complex materials, including realistic defocus blur, view-dependent effects, as well as appearance highlights and reflections.

CCS Concepts: • **Computing methodologies** → **Rendering**; **Mixed / augmented reality**; **Virtual reality**; **3D imaging**; • **Applied computing** → **Physics**; • **Hardware** → **Emerging technologies**.

Additional Key Words and Phrases: computational displays, holography, path tracing, wave optics

ACM Reference Format:

Wenbin Zhou, Xiangyu Meng, Jiankai Xing, Xin Liu, Suyeon Choi, and Yifan Peng. 2026. HoloPathTracer: Fast and Accurate Wave Path Tracing for Holography. *ACM Trans. Graph.* 45, 4, Article 39 (July 2026), 12 pages. <https://doi.org/10.1145/3811351>

1 Introduction

The objective of holographic displays is to reconstruct a wavefront indistinguishable from those emitted by real objects, thereby enhancing the immersive capabilities of VR/AR [Jang et al. 2024; Zhou et al. 2025a]. By leveraging light diffraction and interference, such displays typically employ a spatial light modulator (SLM) with a synthesized 3D hologram to manipulate coherent wave fields [Blanche 2021]. Realizing *Perceptual Realism*—with natural depth and other visual cues—motivates mimicking incoherent emission via multiplexing coherent behaviors inherent in holography (see also Sec. 3.3), a direction further enabled by emerging high-speed SLMs [Choi et al. 2022]. Complementing photonics hardware advances [Yin et al. 2022], computer-generated holography (CGH) algorithms compute the wave fields on the hologram plane from a 3D scene, which are then encoded into phase-only holograms through either iterative or direct methods [Choi et al. 2021] for optical reconstruction.

Exploring a **physically accurate yet computationally efficient rendering framework for complex-valued wave fields** is essential for realistic handling of specular reflection, glossy surfaces, and transparent materials under in-focus and defocused conditions. Yet, this rendering, grounded in Rayleigh–Sommerfeld (RS) and Huygens principles [Buitrago-Duque and Garcia-Sucerquia 2019; Shen and Wang 2006], is computationally demanding due to stringent sampling requirements and the heavy load of multi-view propagation. To remain tractable, existing numerical diffraction computations for CGH “practically” abandon fully physically accurate wave propagation models in favor of *approximate* formulations.

Wave propagation formulated with point-based models [Tsang et al. 2018], while physically expressive, demands dense sampling and becomes prohibitively expensive when rendering coherent fields for complex scattering and involving multiple viewing angles. Alternative image/layer-based models (e.g., focal stacks, holographic stereograms, light fields [Choi et al. 2022; Schiffers et al. 2023]) partition the scene into discretized layers and leverage FFT to accelerate intra-layer propagation. Yet, the inter-layer *discretization* compromises perceptual quality (e.g., incorrect scene occlusions), and runtime still scales linearly with the number of layers. Polygon-based paradigms [Wang et al. 2023] similarly adopt layer-based propagation, where high-fidelity configurations are achieved, yet computation scales linearly with the facet count. Recently, Gaussian primitives-based propagation enables fast neural rendering of 3D cues [Choi et al. 2025a]. This concurrent direction remains an *approximation* to full physical wave propagation and generally yields smooth-like patterns, trading reconstruction fidelity (e.g., correct refraction/reflection) for computational efficiency, akin to prior formulations. More broadly, overly smooth patterns may suppress the stochastic nature of holograms, which is important for resilience and perceptual realism, as discussed in Sec. 2.

Path tracing is a physically accurate framework for modeling light transport, as it characterizes the propagation of light energy in a scene by stochastically sampling light paths via sequences of surface interactions [Steinberg et al. 2024a; Yu et al. 2023]. Integrating such a rendering paradigm with wave optics in an efficient manner is of significance for advancing CGH [Watanabe et al. 2024]. Notably, conventional path tracing pipelines render only radiance (amplitude) while discarding the phase information carried by light.

Consequently, they do not “encode” the full complex wavefront essential for holographic reconstruction. Yet, naively computing the complex wavefront via explicit wave propagation remains computationally prohibitive [Blinder et al. 2021; Chen et al. 2023], due to multi-bounce light transport and high sampling rates demanded.

This work employs RS integral-grounded path tracing to render complex-valued wave propagation, without relying on predefined, discretized geometric primitives. The resulting wave field naturally encodes physically based rendering effects such as view-dependent highlights, specular reflections, transmissive materials, and global illumination, enabling holographic reconstructions with photorealism comparable to that of the Mitsuba renderer [Jakob et al. 2022]. To validate the rendered wave field, we adopt a stochastic gradient descent-based phase-encoding algorithm supervised in the complex-valued domain, and build a holographic display prototype using a phase-only SLM-based light engine.

In summary, this work makes the following contributions:

- We propose an explicit wave path tracing rendering framework for CGH, *HoloPathTracer*, that computes physically accurate complex-valued wave fields for 3D scenes. Our pipeline decomposes physically based illumination into coherent and incoherent components, pre-optimizing the incoherent part, to preserve physical fidelity while realizing practical computational efficiency.
- We develop a parallel-like rendering pipeline that efficiently generates random-like instances of the complex wave field to support multi-frame time-multiplexed holographic displays.
- Our method exhibits sub-linear scaling with respect to the number of primitives in 3D scenes. Consequently, we successfully conduct simulations on highly complex physics-based rendering mesh scenes with up to *two million* geometric elements, maintaining photorealism comparable to Mitsuba path-traced renderings.
- We build a holographic display prototype, optimize phase-only holograms, and experimentally show holographic images with natural defocus cues.

Source code is available on our GitHub repository¹.

2 Related Work

This work builds on recent progress in CGH algorithms, but focuses on exploring a physically accurate yet computationally efficient wave field rendering framework for CGH. For a more comprehensive overview of holographic displays, we refer the readers to [Blanche 2021; Chang et al. 2020; Javidi et al. 2021].

Wave Field Rendering and Hologram Recording. Prior CGH pipelines first render scene radiance on *discrete* geometric primitives, and then apply numerical wave propagation to the hologram plane. Point clouds are widely employed due to their implementation simplicity [Chakravarthula et al. 2019; Tsang et al. 2018], but they typically require extensive sampling. Polygon meshes [Matsushima and Nakahara 2009; Wang et al. 2023] can be processed analytically, offering high computational speed at the cost of lacking detailed texture representation; or numerically, they resemble point-based methods and

¹github.com/zhoubw/HoloPathTracer

Table 1. **Summary of CGH algorithms w.r.t. computational efficiency and reconstruction visual fidelity.** Compared frameworks include the point-based, polygon-based, CNN w/ RGB-D input, SGD-based w/ focal stacks and light fields, and Gaussian-based wave splatting methods. Herein, $\blacktriangle/\checkmark/\times$ denote partial/yes/no for each item, and green/yellow/red indicate good/moderate/poor capability. Our approach uses readily available 3D scene data in graphics with less dependency on pre-defined geometric primitives, and does not require pre-training any networks tied to a specific configuration. Importantly, the path tracing-empowered wave optics rendering naturally supports advanced shading, occlusion, and view-dependent effects. Representative methods: 1 [Maimone et al. 2017], 2 [Wang et al. 2023], 3 [Shi et al. 2021], 4 [Choi et al. 2022], 5 [Kim et al. 2024], 6 [Choi et al. 2025a].

	Computational Efficiency				Visual Fidelity				
	Fast inference	Efficient multi-frame	Data availability	Pretraining required	Reflection/Refraction	Complex material/GI	View-dependent	Depth-continuity	Natural defocus
Point-based ¹	\blacktriangle	\checkmark	\checkmark	\times	\times	\blacktriangle	\checkmark	\checkmark	\times
Polygon-based ²	\times	\times	\checkmark	\times	\blacktriangle	\blacktriangle	\checkmark	\checkmark	\checkmark
CNN w/ RGBD ³	\checkmark	\blacktriangle	\checkmark	\checkmark	\times	\blacktriangle	\times	\times	\times
SGD w/ Focal Stack ⁴	\times	\blacktriangle	\times	\times	\blacktriangle	\blacktriangle	\times	\blacktriangle	\checkmark
SGD w/ Light Field ⁵	\times	\blacktriangle	\times	\times	\blacktriangle	\blacktriangle	\checkmark	\blacktriangle	\checkmark
Gaussians-based ⁶	\blacktriangle	\times	\checkmark	\checkmark	\times	\blacktriangle	\checkmark	\checkmark	\times
Path Tracing (ours)	\blacktriangle	\checkmark	\checkmark	\times	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark

demand significant sampling during baking process. Layer (RGBD)-based sampling methods are computationally attractive as they operate on a finite set of pre-rendered 2D images [Choi et al. 2021; Kavakli et al. 2023], but they can struggle to provide continuous depth cues. Light fields can be transformed via short-time Fourier transform (STFT) into holographic stereograms [Chakravarthula et al. 2022a; Padmanaban et al. 2019], which encode both depth- and view-dependent effects but remain computationally expensive. Recently, Choi et al. [2025a] utilize 2D Gaussians for CGH with enhanced computational efficiency over classical meshes-based methods. Despite progress, initiating wave propagation from discretized, pre-rendered primitives inevitably introduces approximations that degrade wave-field accuracy and hinder faithful reproduction of certain optical/shading effects (see Table 1). This work integrates these two stages—radiance rendering and hologram wave recording—to directly derive target wave fields.

Ray Tracing with Wave Optics. Ray tracing has been widely adopted in optical system analysis and lens design [Steinberg and Yan 2021; Wang et al. 2022]. Recently, wave-optical behaviors have been incorporated into Monte-Carlo rendering frameworks [Liu et al. 2025a; Magallón et al. 2021; Steinberg et al. 2022] to model effects beyond the reach of geometric optics alone, including diffraction and high-frequency interference patterns [Bar et al. 2019; Kim et al. 2025; Steinberg and Yan 2022], as well as wave-informed BRDFs [Wei et al. 2025; Zeng et al. 2025a]. As a representative, Steinberg et al. [2024a] introduce generalized rays for backward wave-optical light transport, enabling sensor-to-source path sampling while retaining the coherence information needed for rendering diffraction and interference effects.

Particularly relevant to our work, Blinder et al. [2021] introduce a point-based CGH framework rendered via path tracing; however, their method primarily targets global illumination on point primitives and does not model coherent propagation through complex materials. Chen et al. [2023] improve BSDFs and ray samplings to accelerate computation. Notably, as these methods leverage path tracing solely to acquire physically accurate radiance without carefully incorporating “long-path” wave propagation influence, they offer limited support for encoding full visual cues via transparent or glossy objects and demonstrate results on a narrow range of reconstruction scenes. We primarily validate the proposed framework using the Disney principled BSDF [Burley 2015]; nevertheless,

established path integral regimes that exploit wave-optical light transport and microstructure-based BxDFs [Jakob et al. 2014; Steinberg et al. 2024b; Yan et al. 2018] are, in principle, compatible with our formulation and can be incorporated.

Smooth-Random Phase Trade-off and Photorealism Matter. The spatio-angular behavior of resulting reconstructions can depend on the wavefront’s phase profile: smooth-phase versus random-phase holograms [Choi et al. 2025b; Yoo et al. 2021]. The former type [Shimobaba et al. 2015] typically yields reasonable 2D image-like quality for a few fixed viewpoints/depths; however, the relative energy concentration in low angular frequencies inherently limits the eye-box size and produces unnatural defocus cues [Tseng et al. 2024; Xia et al. 2025]. By contrast, constructing a random-like phase for a given 3D scene better utilizes the available spatio-angular bandwidth of holographic displays [Schiffers et al. 2023], a prerequisite for perceptual realism. This benefit comes at a cost: rapid phase variations across pixels tend to induce speckle noise and other visual artifacts [Goodman 2007], which can become more prominent when enforcing pixelated SLMs affected by crosstalk [Ban et al. 2026a]. Consequently, additional speckle reduction and hardware calibration strategies are typically required [Peng et al. 2021].

Faithfully reproducing visual effects requires considering both amplitude and phase during hologram recording. Yet, established solutions, including aforementioned layers- and light field-based ones, primarily enforce constraints on amplitude to approximate wavefront continuity in 3D space [Shi et al. 2021], thereby only *indirectly* shaping phase during modeling and ultimately achieving “partial” photorealism. Prior work that integrates path tracing concepts into CGH [Blinder et al. 2021; Sun et al. 2020] has targeted either computational acceleration or vanilla global illumination, that are, again, not inherently phase-centric. This work departs from existing wave presentations: we develop a microfacet-extended formulation and a full wave field rendering framework from light sources to hologram recording planes (see Sec. 4.1). This design provides fine-grained control over the hologram “spectrum” in both amplitude and phase, further enabling *multiple* time-multiplexing frames to be rendered in parallel with minimal extra cost via path reuse, thereby supporting a richer set of 3D visual effects at reduced computational expense.

Computational Techniques for Hologram Encoding. Direct methods (e.g., DPAC and its variants [Shi et al. 2022]) are highly efficient, analytically decomposing a complex-amplitude into a coherent sum

of two phase-only patterns, but at the cost of reduced light efficiency and additional opto-mechanical filtering. Deep learning offers an alternative direct-inference route via inverse networks [Liu et al. 2025b; Zhou et al. 2025b], some with support for compression through compact latent representations [Ban et al. 2026b]. Yet, many models tend to generate phase patterns with checkerboard-like structures, an artifact that may degrade holographic reconstructions. Most hologram encoding schemes tend to overfit the amplitude of supervised pixels [Choi et al. 2021], rather than faithfully reconstructing the underlying 3D wave field. Our framework is compatible with mainstream phase-only hologram encoding paradigms and we adopt SGD-based optimization under complex-amplitude supervision [Chen et al. 2021]. By explicitly constraining both amplitude and phase of the target field, we mitigate the tendency of overfitting to pixel-wise intensity. We also introduce a phase-shift-invariant scale factor to handle global and local phase ambiguities for better convergence.

3 Preliminaries of Path Tracing and Wave Propagation

Path Tracing and RS integral. For path tracing, one can employ the rendering equation to solve for radiance under global illumination [Schmidt et al. 2013], expressed as follows:

$$L_o(\mathbf{x}, \omega_o) = L_e(\mathbf{x}, \omega_o) + \int_{\Omega} f_r(\mathbf{x}, \omega_i, \omega_o) L_i(\mathbf{x}, \omega_i) \cos(\mathbf{n}, \omega_i) d\omega_i, \quad (1)$$

where $L_o(\mathbf{x}, \omega_o)$ and $L_e(\mathbf{x}, \omega_o)$ denote the outgoing and emitted radiance at point \mathbf{x} with the direction ω_o . $L_i(\mathbf{x}, \omega_i)$ indicates the incoming radiance with the direction ω_i . $f_r(\mathbf{x}, \omega_i, \omega_o)$ represents the bidirectional reflectance distribution function (BRDF). The integral sums contributions of all possible incoming directions. To account for both amplitude (intensity) and phase information for physical realism, one can incorporate the RS integral [Buitrago-Duque and Garcia-Sucerquia 2019; Shen and Wang 2006] within the path-tracing framework to describe wavefront propagation, as:

$$u(\mathbf{x}) = \frac{1}{i\lambda} \iint_A u(\mathbf{x}_i) \frac{1}{r} e^{i\frac{2\pi}{\lambda}r} \cos(\mathbf{n}, \mathbf{r}) dS, \quad (2)$$

where $u(\mathbf{x})$ is the complex amplitude at point \mathbf{x} , λ is the wavelength, r is the distance between \mathbf{x} and \mathbf{x}_i , and \mathbf{r} denotes the vector from \mathbf{x} to \mathbf{x}_i yields $r = |\mathbf{r}|$. A represents the illuminated aperture.

The RS integral can be interpreted as the wave-optics counterpart of Eq. 1, with Huygens' principle enforcing isotropic radiation from each point. While this integral could be solved via a path tracing framework, it is often computationally prohibitive due to the combined costs of multi-bounce light transport and dense sampling rates inherent to wave optics. As such, we extend the microfacet model and decompose diffuse reflection, glossy reflection, mirror reflections, and transmissions into coherent (solved by RS-integral) and incoherent components (solved by path tracing). We then record the resulting BRDF into a complex amplitude on the surface, thereby accelerating sampling for holographic scene reconstruction.

Angular Spectrum of Wave Field. For a complex amplitude $u(x, y)$, its *angular spectrum* is defined as its Fourier transform $U(f_x, f_y)$, which gives the coefficients of its plane-wave decomposition. For detailed derivations, we refer readers to Supplementary Sec. S1.1.

The key insight is that the coefficient $U(f_x, f_y)$ corresponds to the plane-wave component propagating along direction (θ_x, θ_y) , where $(\theta_x, \theta_y) = (\arcsin(\lambda f_x), \arcsin(\lambda f_y))$. Its magnitude $|U(f_x, f_y)|$ therefore measures the strength of the wave component traveling in that direction [Matsushima et al. 2003], which we refer to as the *angular amplitude* in this work.

Notably, an ideal plane wave extends infinitely in space and lacks spatial localization, while an ideal point source is perfectly localized but exhibits a spatially uniform spectrum, limiting its ability to represent anisotropic scattering. We instead model light as **coherent plane-wave bundles** emitted from mesh facets that are locally small on the macroscopic scale yet much larger than the wavelength. Each bundle is assumed to have a finite divergence angle, allowing efficient ray-based tracing [Steinberg et al. 2022]. This intermediate representation, lying between the extremes of an ideal point source and an ideal plane wave, can be propagated through the scene via path tracing, thereby capturing aspects of the wave-particle duality.

Coherent versus Incoherent Propagation. Notably, a hologram with a smooth phase yields speckle-free images under coherent illumination, but concentrates its angular-spectrum energy near zero frequency. Consequently, after propagating through the optical system it undergoes pronounced diffraction, resulting in smaller caustics with ringing artifacts. In real-world scenarios, light sources are generally incoherent, even if monochromatic. Objects illuminated by such sources scatter complex waves, whose initial phases fluctuate randomly over time. The human visual system integrates and averages these fluctuations over tens of milliseconds [Goodman 2007; Schiffrers et al. 2025]. Thus, the angular-amplitude of an incoherent light beam tends to uniformly occupy the available bandwidth at the receiving plane, leading to natural defocus blur after propagation. We leverage this context to synthesize perceptually realistic incoherent blur via time-multiplexed holograms (see also Sec. 4.3), an increasingly practical strategy as high-speed SLMs are gradually entering the market [Choi et al. 2022].

However, most aforementioned CGH algorithms are designed to generate a *single* frame per inference, suggesting a substantial increase in computation time for multi-frame generation for a given 3D scene.

4 Wave Field Rendering via Path Tracing

Figure 2 overviews our *HoloPathTracer*, whose rendering pipeline consists of nine key stages. A ray is emitted from a point on the SLM (recording plane) in hologram space with a random initial direction (s1). This ray is then transformed into world space (s2) and traced through the scene (s3). Upon intersecting a facet, scattering is evaluated according to its BSDF, where different lobes are sampled probabilistically (s4). The first non-delta BSDF sample is defined as the real image point \mathbf{P} in world space. Before reaching \mathbf{P} , the coherent plane-wave bundle is traced with OPL_{world} accumulated. While after reaching \mathbf{P} , several incoherent radiance rays are traced to estimate the amplitude of the incident plane-wave bundle (s5). Then, the virtual image point \mathbf{P}' is calculated and mapped back into hologram space \mathbf{Q} via a coordinate transformation (s6). Based on the position of \mathbf{Q} , the OPL_{Holo} is determined (s7) and several random phase samples are selected from the pre-computed Gaussian random fields (s8) to

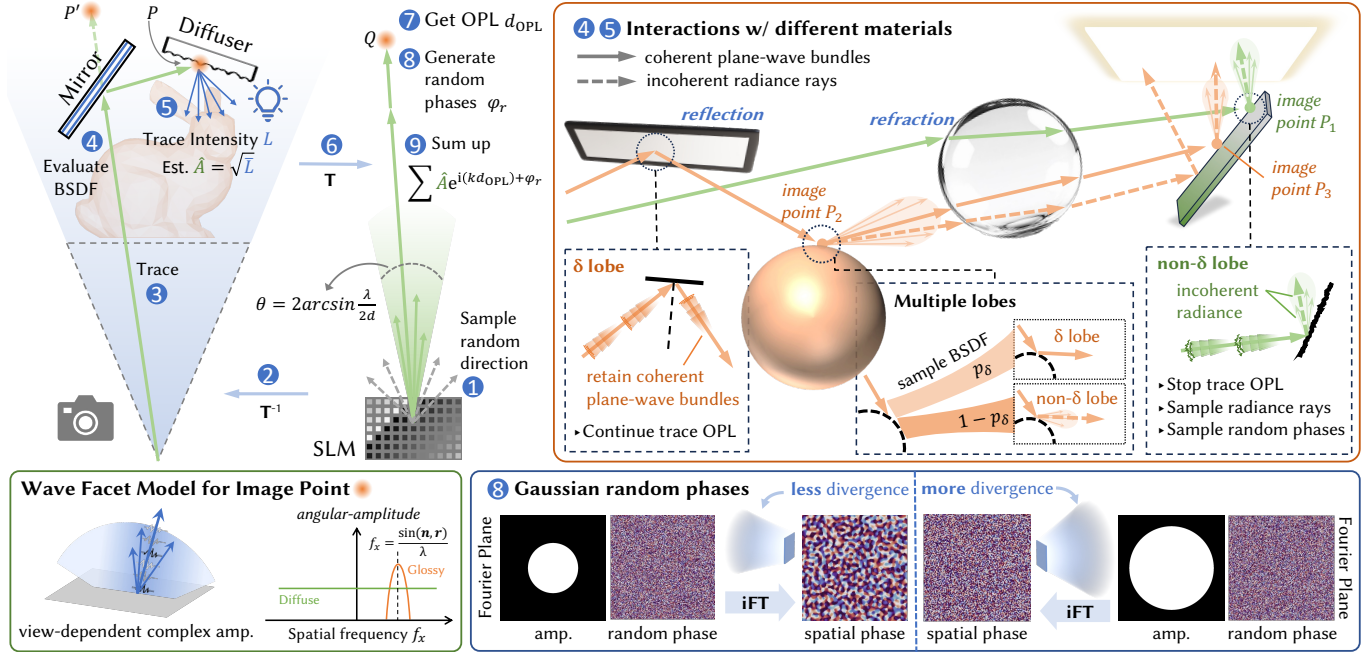


Fig. 2. **HoloPathTracer processing pipeline.** **Top-left:** the proposed method can be divided into 9 Stages including ray sampling and tracing (1, 3), projection transformation (2, 6), BxDF evaluation (4, 5), phase sampling (8), and complex amplitude accumulation (7, 9). Step-by-step details are presented in Sec. 4. **Top-right:** the wave field propagation undergoes various materials where we evaluate the BxDF of the intersected facet (point) to determine which condition, i.e., δ or non- δ , to activate before reaching the SLM, inherently encoding these complex visual cues into the hologram recording process. **Bottom-left:** we model image points as the view-dependent wave-facet to record certain angular-amplitude distribution obtained via path tracing. **Bottom-right:** wave facets are initialized with random phases sampled from Gaussian random fields to spread their energy over the desired divergence range.

model the incoherent scattering at P . Finally, the contribution of each plane-wave bundle is calculated by combining the amplitude (from s5) and phase (from s7 and s8), and then coherently accumulated on the hologram recording plane (s9).

In the following, we introduce the wave-optics-based scattering model: wave facet (Sec. 4.1), wave path tracing scheme for coherent plane-wave bundles (Sec. 4.2), parallel wave-field rendering and Gaussian random fields for multi-frame holography (Sec. 4.3), and strategies for accelerated sampling (Sec. 4.4). More algorithmic details are presented in Supplementary Material.

4.1 Wave Facet: Extended Microfacet Modeling

The microfacet model, commonly used in physically based rendering, provides an effective description of incoherent irradiance and is formulated as a BxDF in conventional path tracing. However, for wave field rendering grounded in coherent scalar diffraction, an extended microfacet model must also incorporate phase information, not just intensity. In this work, we denote our facet model for CGH rendering as **wave facet**.

We refer to a small macroscopic surface patch (characteristic size $M \gg \lambda$) on the recording plane or object surface as a facet, with light wavelength λ . The recording plane physically corresponds to the hologram (i.e., SLM) plane. Due to the finite feature size of the hologram, the bandwidth of a facet on the recording plane is limited: each facet can record only incident waves with angles smaller than $\arcsin[\lambda/(2d)]$ without aliasing, where d is the hologram feature

size. When a coherent plane-wave bundle interacts with a facet, the scattered field may contain both coherent and incoherent components. For coherent scattering (e.g., specular reflection or ideal refraction), the wavefront remains phase-consistent in accordance with Huygens' principle. It is therefore necessary to continue tracing the subsequent OPL to preserve correct depth cues and multi-view consistency for objects observed through specular surfaces. In contrast, for diffuse and glossy surfaces, photons undergo path-length variations due to stochastic scattering events. We model this incoherent behavior by applying random phase perturbations at the surface, effectively decorrelating the outgoing field. In subsequent tracing we omit further phase (OPL) accumulation—as the phase has already been randomized—and only record the ray-carried amplitude (i.e., the square root of average radiance).

Matsushima et al. [2009] show that, for both diffuse and glossy reflectance, assigning a certain random phase pattern can effectively modulate the propagation directions of the resulting wave field. Motivated by this, we associate each rough facet in the scene (e.g., diffuse and/or glossy) with a random phase sampled from a Gaussian random field (see Sec. 4.3 for details). This stochastic phase perturbation broadens the angular bandwidth of the emitted wave, yielding a statistically uniform yet random *angular-amplitude*. The degree of randomness provides an explicit control to balance angular spread against speckle noise. A subsequent path tracer then resolves the geometric component, determining how radiant energy is distributed across outgoing directions from each rough facet. In




Fig. 3. **Simulated angular amplitude on each wave facet**, constructed by combining the amplitude from path tracing and the phase from OPL, facilitating variable-frame-multiplexing (right).

Fig. 3 we analyze the *angular-amplitude* computed by our framework, which characterizes the distribution of light’s propagation directions on the wave facet. Specifically, for light incident on a rough-like surface, the recorded *angular-amplitude* closely follows the PDFs of outgoing directions for different materials, validating that the proposed wave facet model can faithfully reconstruct the material PDFs by combining the “random-like” phase structure and view-dependent amplitude distribution of the wave field.

4.2 Wave Path Tracing

The recording plane resides in *hologram space*, where physical reconstruction is ultimately realized. Accordingly, all diffraction accumulation for solving the RS integral must be performed in hologram space. In contrast, path tracing is carried out in *world space*, which hosts the virtual scene and is required to recover physically correct light transport. The transformation from world space to hologram space can be achieved by $\mathbf{x}_{\text{holo}} = \mathbf{T}\mathbf{x}_{\text{world}}$, where \mathbf{T} is defined as:

$$\mathbf{T} = (\mathbf{P}_{\text{holo}}\mathbf{V}_{\text{holo}}\mathbf{M}_{\text{holo}})^{-1}\mathbf{P}_{\text{world}}\mathbf{V}_{\text{world}}\mathbf{M}_{\text{world}}, \quad (3)$$

where \mathbf{M} , \mathbf{V} , \mathbf{P} denote model, view, and projection transformations.

Rays are randomly sampled within a limited divergence cone at the recording plane in hologram space, transformed into the world space, and then traced. Upon intersecting a wave facet, the path tracer, i.e., Mitsuba, samples one lobe of the material’s BSDF according to its probability. Each sampled lobe is marked as either δ or non- δ (see top-right in Fig. 2): δ lobes represent specular reflection or ideal refraction, where the outgoing direction is described by a delta function; non- δ lobes correspond to diffuse or glossy reflection, where the outgoing direction follows a PDF determined by material properties such as metallicity and roughness.

When a ray encounters a δ lobe, we continue to propagate that path coherently by accumulating its OPL, yielding $\text{OPL} = \sum_i L_i n_i$, where n_i denotes the refractive index of the medium and L_i specifies the associated segment length. In contrast, for non- δ lobes we only compute the subsequent ray intensity, since rough surfaces induce stochastic path-length variations. Equivalently, the outgoing field can be regarded as the superposition of many coherent plane-wave bundles with random phases, making phase tracking unnecessary; therefore, we switch to incoherent path tracing to estimate the remaining intensity transport. This design improves efficiency by omitting wave sampling requirements in subsequent transport.

Specifically, we define the point at which a ray first hits a non- δ facet as the *real image point* P . The amplitude traced beyond this point is accumulated at P , assigned several random phases (for time-multiplexed reconstruction, see Sec. 4.3) to model interactions with




Fig. 4. **Gaussian random field sampling on facet** in UV and Holo spaces.

rough surfaces, and then coherently propagated back to the recording plane. We assume the accumulated OPL through all δ -lobes prior to P to be d_{OPL} . The *virtual image point* for perception is therefore $P' = O + d_{\text{OPL}}\mathbf{v}$, where O and \mathbf{v} denote ray origin and direction. To compute the wave perturbation induced by P' in hologram space and accumulate its contribution on the recording plane, we transform P' into hologram space using Eq. 3, yielding $Q = \mathbf{T}P'$. After assigning the corresponding phase and amplitude, the resulting wave is propagated to the recording plane. Summing the contributions of all sampled rays on the recording plane yields a solution to the RS integral.

Even after leveraging incoherent light transport to accelerate random wave propagation, the wave path tracer may still converge slowly (see Fig. 5) due to multi-bounce evaluation for amplitude and dense RS-integral sampling for phase. To reduce amplitude variance, we adopt a texture-baking strategy [Knodt et al. 2023] that caches view-independent ambient illumination (Supplementary Sec. S1.4); when a ray hits the ambient-light lobe, its intensity is directly fetched from the cache. To further accelerate phase accumulation, we introduce a *two-stage* strategy that substantially reduces the RS-integral sampling requirements (refer to Sec. 4.4).

4.3 Multi-frame Wave Field Rendering

Recent advances in digital holography with SLMs leverage multi-frame time-multiplexing to realize temporally incoherent superposition, thereby reducing coherent speckle and improving resilience to artifacts [Choi et al. 2022; Chu et al. 2025]. To support such a multi-frame holography regime, we efficiently render multiple complex wave fields of the same scene in parallel, each differing only in the random phase assigned at non- δ facets. We note that the random phase should not arbitrarily perturb the wavefront in object-space parameterizations; instead, it should redistribute energy within the angular bandwidth physically recordable by the hologram. Let $\theta_{\text{max}} = \arcsin[\lambda/(2d)]$ denote the maximum incident angle supported by the recording plane. For the m -th frame, we therefore precompute a band-limited Gaussian random field (GRF) on the hologram plane as:

$$\begin{aligned} r_m(x, y) &= \iint_{\Omega_f} e^{i\xi_m(f_x, f_y)} e^{2\pi i(f_x x + f_y y)} df_x df_y \\ &= \mathcal{F}^{-1} \left\{ \chi_{\Omega_f}(f_x, f_y) e^{i\xi_m(f_x, f_y)} \right\}, \end{aligned} \quad (4)$$

where $\Omega_f = \{(f_x, f_y) \mid f_x^2 + f_y^2 \leq (\sin \theta_{\text{max}}/\lambda)^2\}$ indicates the admissible angular-spectrum support, χ_{Ω_f} is the corresponding circular passband, and $\xi_m(f_x, f_y) \sim \mathcal{U}[0, 2\pi]$ are independent random




Fig. 5. **Convergence behavior of wave field rendering schemes.** Our method adopts a two-stage scheme that first renders the complex-valued wave field to a wave recording plane, then uses ASM to continue propagation to the SLM (hologram) plane; we further accelerate this via texture-baking (“fast” variant). A naïve baseline directly renders from the scene to the (typically distant) SLM plane, leading to higher computational complexity. x -axis: samples per pixel (SPP); y -axis: number of time-multiplexed frames. Convergence rate is measured by the ratio of reconstruction PSNR against that of 10^7 spp. Note, white dashed lines indicate the 80% convergence.

spectral phases. This formulation can be interpreted as a superposition of many plane-wave components with random initial phases, while confining all energy to the desired divergence cone. Although the Fourier-domain phases are sampled uniformly in our implementation, the inverse transform sums many independent phases, so r_m is approximately Gaussian in the spatial domain by the central limit theorem; we therefore refer to it as a GRF. We then use $\phi_m(x, y) = \arg(r_m(x, y))$ as the random phase map of the m -th frame.

Herein we consider an *image point* in world space that is transformed to a point $Q(x_Q, y_Q, z_Q)$ in hologram space. For the m -th frame, the phase assigned to this point is sampled from the GRF at its projected location, i.e., $\phi_{Q,m} = \phi_m(x_Q, y_Q)$. If the traced path associated with Q yields amplitude A_Q and optical path length $d_{\text{OPL},Q}$, its contribution to the recording plane is

$$U_m \leftarrow U_m + A_Q \exp \left[i \left(\frac{2\pi}{\lambda} d_{\text{OPL},Q} + \phi_{Q,m} \right) \right]. \quad (5)$$

This yields spatially and angularly consistent random phases across the scene, as shown in Fig. 4. Compared with assigning phases via facet UV coordinates, our strategy keeps the sampling density and effective bandwidth tied to hologram space, avoiding distortions caused by changes in facet normals or UV scaling.

Moreover, only $\phi_{Q,m}$ depends on the frame index m ; the traced geometry, amplitudes, and OPL terms are shared by all frames. The same ray paths can therefore be reused across time-multiplexed holograms with negligible additional overhead (see Fig. 6). Further discussion on path reuse is provided in Supplementary Sec. S1.5.

4.4 Wave Sampling with Rays

In FFT-based diffraction, the sampling requirements of the quadratic phase term depend on both propagation distance and angle [Wei et al. 2023], and must be performed on a specific plane. If the pre-defined spatial or angular sampling intervals are coarser than required, numerical diffraction produces artifacts. In contrast, our path tracing-based approach avoids predefined sampling grids by employing Monte Carlo ray sampling, equivalent to randomly




Fig. 6. **Runtime ablation** on representative CGH frameworks. Colors indicate varying-frame-multiplexing: 1, 3, 8, 24 frames. The proposed methods generate multiple frames with only marginal computational overhead.

sampling the sub-hologram of each object point. This strategy suppresses artifacts of undersampling and naturally supports heuristic, spatially varying non-uniform sampling schemes.

Converting rays to wave fields on the hologram plane resembles an integral-based diffraction operation with computational complexity $\mathcal{O}(N_{p,0}N_{\text{ray},0})$, where $N_{p,0}$ and $N_{\text{ray},0}$ denote the pixel samples and ray samples per pixel (SPP) on the hologram plane, respectively. Here, SPP corresponds to the number of samples in the angular domain and scales quadratically with the hologram area, which itself grows quadratically with propagation distance. Consequently, the overall complexity of this Monte Carlo integral increases approximately with the fourth power of propagation distance. To obtain a high-fidelity hologram, it is essential to satisfy these sampling requirements during path tracing. In our setting, the hologram plane is placed 85 mm away from the farthest scene point. To reduce the computational cost, we adopt a *two-stage* strategy inspired by the wave recording plane (WRP) technique [Shimobaba et al. 2009]. Specifically, we first convert traced rays into wave fields on a WRP placed 5 mm away from the farthest scene point, then propagate this field from WRP to the hologram plane using ASM. The computational complexity of this procedure is $\mathcal{O}(N_p N_{\text{ray}} + 2N_p \log N_p)$, where N_p and N_{ray} denote the counterparts on WRP. The factor 2 accounts for double zero-padding in FFT-based ASM to avoid aliasing. In practice, the ASM overhead is negligible because both N_{ray} and N_p remain much smaller than their counterparts $N_{\text{ray},0}$ and $N_{p,0}$ on the distant hologram plane. Therefore, as illustrated in Fig. 5, compared to direct conversion on the hologram plane, this WRP-based scheme substantially reduces the required samples, which is particularly beneficial when enforcing the multi-frame wave field rendering in practical CGH.

5 Experiments and Results

5.1 Scene and System Configurations

We use readily available 3D scenes [Bitterli 2016] from Mitsuba, which we further edit in Blender to show complex rendering effects. Their visualization is presented in Supplementary Sec. S4. The wavelengths for simulations are set to 640.0 nm (red), 516.5 nm (green), and 455.4 nm (blue), with a pixel pitch of $8 \mu\text{m}$ and a propagation distance of 80 mm (corresponding to virtual infinity). Detailed setup specifications are provided in Supplementary Sec. S5. Baseline methods for assessment are described in Supplementary Sec. S2.

5.2 Rendering Results

Figure 7 compares rendering results of different CGH frameworks, including two HoloPathTracer variants, RGBD, focal stack (FS), light field (LF), and ground truth from Mitsuba3 (GT). “Fast” denotes our renderings employing a texture-baking strategy for accelerated computation. Similarly, Fig. 8 compares rendering results of Ours, Gaussian wave splatting (GWS), and mesh-based (Mesh). For each set, we present close-up regions with the viewpoint focused at front, middle, and back, as well as the left/right view-dependent effects.

Across evaluated scenes, RGBD exhibits noise-corrupted defocus regions (e.g., when focused at front) due to depth-layer discontinuities in supervision. FS improves the defocus consistency but softens details in focused regions (see red boxes for middle-focus). LF provides the highest visual quality among image/layer-based baselines because of dense angular supervision, yet noticeable artifacts remain (e.g., third-column close-ups), likely due to overfitting to predefined viewpoints and pupil sizes. Both GWS and Mesh struggle to reproduce accurate mirror-like visual cues (see pink boxes where Lego truck imagery is missing in the mirror). Our method closely matches GT under arbitrary viewpoints, refocus distances, and pupil sizes, yielding smooth focus transitions, and accurately reproducing lens/mirror-formed imagery. “Fast” variant preserves the same focus cues while being faster than baselines, with only minor detail/color drifts mainly due to baking approximations. Please refer to the Supplementary Material for results on additional scenes, and to the Supplementary Video, where the view-dependent effects and focus transitions are more clearly visible.

5.3 Experimental Results

Phase-only Hologram Encoding. We adopt the established SGD-based optimization following recent literature [Choi et al. 2021]. Importantly, to better align with the scope of efficiently and accurately encoding both amplitude and phase, we supervise the optimization directly in the complex-valued domain, rather than supervising only the amplitude as in the prior work. In such a way, one would only need to supervise on one single reconstruction plane without losing both depth and view continuity, significantly reducing optimization time. Further details of the phase-only optimization are described in Supplementary Sec. S3.2.

Physical Configuration. Our holographic display prototype closely follows the simulation, as shown in Supplementary Fig. S4. Synthesized holograms are quantized and displayed on a phase-only SLM, illuminated by a fiber-coupled RGB laser. The beam is directed through a neutral density filter, a collimating lens, a linear polarizer, and a beam splitter before incident onto the SLM. The modulated wavefronts transmit through an eyepiece and a lens before being recorded by a sensor. We configure an iris mounted on translation mechanics in between the eyepiece and the camera lens to mimic the pupil-filtering. For multi-depth holographic reconstruction, we propagate to 5 depth planes, uniformly spaced at 1.25 mm intervals within 80–85 mm. The camera focus can be adjusted using an Arduino micro-controller to support re-focusing on multiple pre-calibrated planes. At this proof-of-concept stage, color and time-multiplexing results are obtained via post-image processing (e.g.,

24 frames). Detailed specifications of devices/parts are presented in Supplementary Sec. S5.

Acquired Display Results. We present captured 3D focal stacks of holograms generated and displayed on our experimental prototype in Fig. 9, comparing ours with the representative focal stack (FS) baseline. For each set, we present close-up regions highlighting depth-dependent effects. We observe across three scenes that ours produced sharper in-focus details and more natural defocus effects.

Noteworthy, although the displayed image contrast may appear low, it remains comparable to that reported in recent CGH systems [Chakravarthula et al. 2022b; Chu et al. 2025; Kuo et al. 2023] built upon similar light-engine setups without extensive filters. In addition, random-phase holograms are known to be more sensitive to SLM inter-pixel crosstalk, which can further reduce image contrast. Possible remedies include learned camera-in-the-loop calibration [Choi et al. 2021], crosstalk-aware optimization [Markley et al. 2023], and higher-fidelity SLM hardware. Improving display contrast is therefore largely orthogonal to the main scope of this work, which is to establish a fast wave path tracing framework that faithfully reproduces 3D visual cues.

6 Discussion and Conclusion

We develop a path tracing-based wave optics rendering paradigm for CGH that encodes physically accurate 3D scene attributes into holograms while maintaining reasonably fast computation. In particular, our method delivers physically accurate and spatially continuous depth cues, even in the presence of curved reflective and/or refractive surfaces. We further achieve substantial computation speedups via coherency decomposition, ambient-light caching, and a two-stage sampling strategy. We envision this novel wave field rendering framework paving the way towards efficient, multi-frame, random hologram generation in large-scale, complex 3D scenes. We validate our HoloPathTracer framework through simulations and experiments, demonstrating photorealistic reconstruction quality comparable to state-of-the-art scene rendering pipelines (e.g., the Mitsuba renderer) on off-the-shelf scene datasets.

Limitations and Future Work. While our wave field rendering execution has been demonstrated to efficiently support multi-frame hologram generation with only marginal computational overhead, its runtime is presently bounded by the inherent sampling cost of wave propagation, and can be further accelerated by training advanced network models to represent/infer the propagation (rendering) [Qu et al. 2025; Zeng et al. 2025b]. We also observe minor visual discrepancies in brightness (dynamic range) and color details between our reconstruction and ground-truth, especially for the “fast” variant tailored for higher computational efficiency.

For scenes illuminated by incoherent light sources, the time-multiplexed random phases reconstruction adopted in this work is an unbiased estimator of the target light-field intensity distribution in expectation. The remaining discrepancies are therefore mainly tied to modeling approximations in the implementations. In particular, the binary δ /non- δ treatment can introduce errors for low-roughness glossy materials that still preserve limited imaging ability, and using a single wave recording plane may cause color




Fig. 7. Rendering results of different CGH frameworks. *Fast* denotes our rendering using texture-baking for computation acceleration.

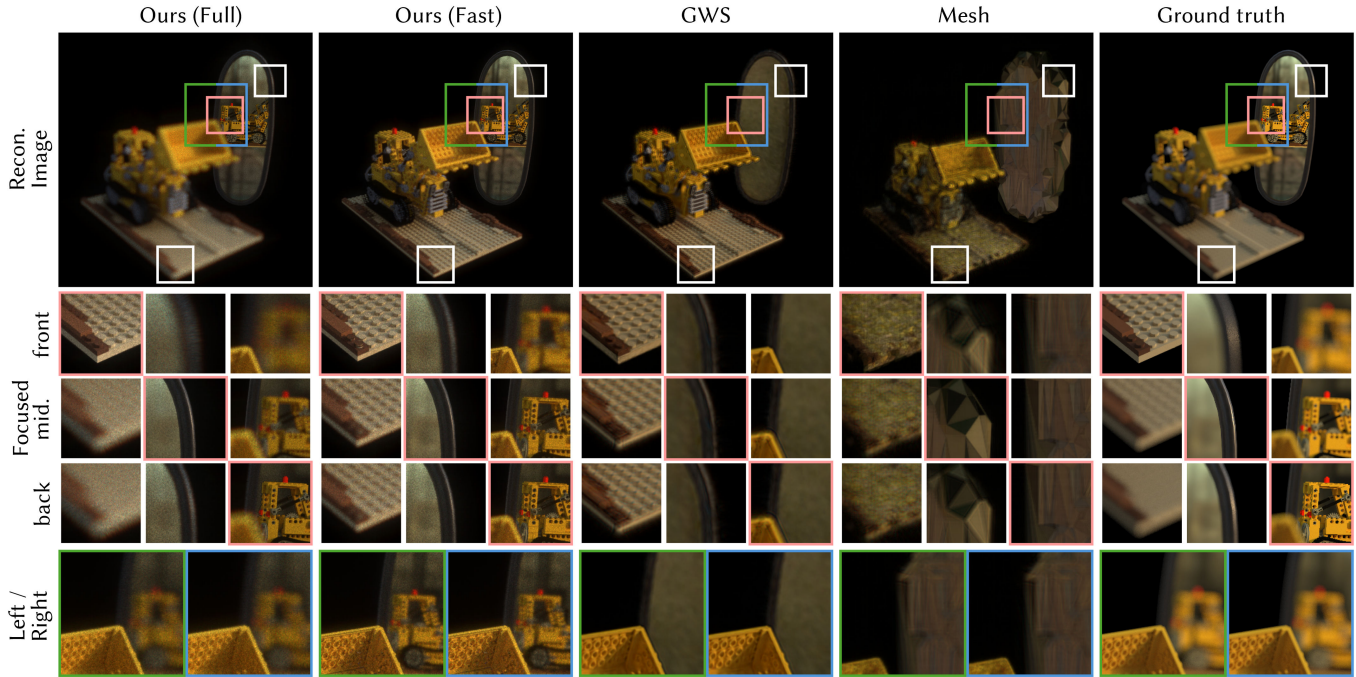


Fig. 8. **Rendering results** of our two variants, Gaussian wave splatting, and mesh-based. We refer readers to “imagery” differences on the *mirror*.

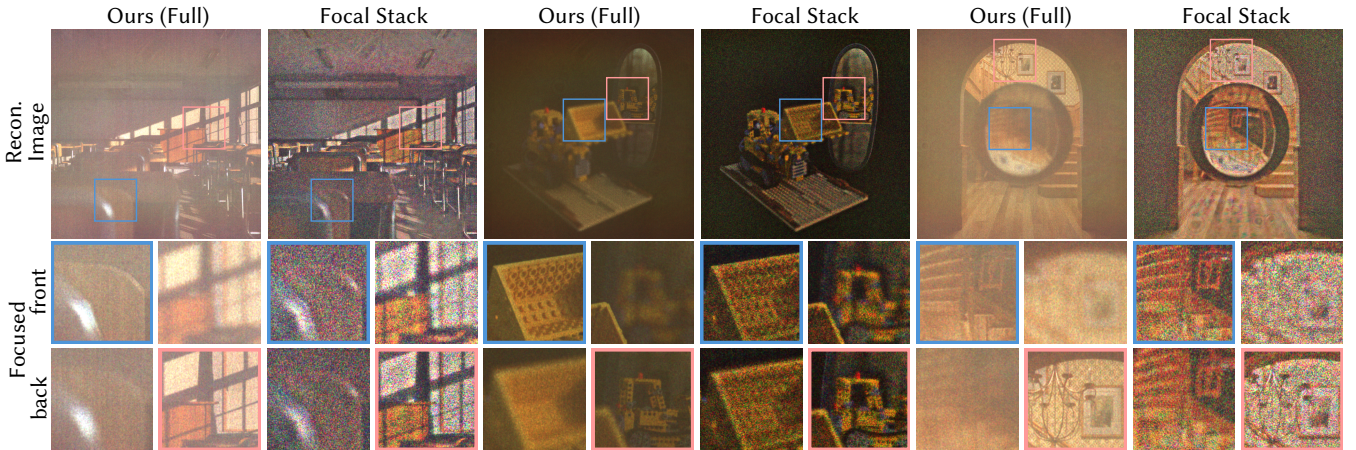


Fig. 9. **Acquired display results** comparing Ours with Focal Stack, with front- and back-focused close-ups. Our method produces sharper in-focus details and more natural defocus. The observed contrast loss primarily stems from the higher sensitivity of our random-phase holograms to SLM crosstalk compared with smoother, single-view-supervised image-based holograms, while better supporting multi-view reconstruction.

drift for objects placed very close to it due to the limited spatial resolution for recording view dependence. Roughness-aware coherence modeling and multiple recording planes are promising directions to mitigate these effects, as further discussed in Supplementary Sec. S1.3.

In addition, our experimental holographic display quality is bounded by the restricted étendue of the SLM and residual degradations from imperfect optics. The étendue bottleneck is a long-standing challenge in CGH, and emerging photonics-based SLMs with much smaller pixel pitch [Kaczorowski et al. 2024] promise to lift this limitation. A camera-calibrated, learned propagation model, similar

to those explored in prior work [Choi et al. 2021], could correct for most residual degradations along the optics path. We regard such calibration as a valuable future add-on.

Acknowledgments

We thank Jinfan Lu and Li Liao for fruitful discussions. This work was partially supported by the National Natural Science Foundation of China (62322217), the Innovation and Technology Fund of Hong Kong (MHP/313/24), and the Research Grants Council of Hong Kong (GRF 17208023).

References

- Hyunmin Ban, Wenbin Zhou, Xiangyu Meng, and Yifan Evan Peng. 2026b. Multi-color compressive hologram synthesis with learned wave propagation. *Computational Visual Media* 12, 2 (2026), 435–448. doi:10.26599/CVM.2025.9450497
- Hyunmin Ban, Wenbin Zhou, and Yifan Peng. 2026a. Towards Edge Holography via Implicit Neural Representation and Compression. *IEEE Transactions on Visualization and Computer Graphics* (2026), 1–10. doi:10.1109/TVCG.2026.3679061
- Chen Bar, Marina Alterman, Ioannis Gkioulekas, and Anat Levin. 2019. A Monte Carlo framework for rendering speckle statistics in scattering media. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–22. doi:10.1145/3306346.3322950
- Benedikt Bitterli. 2016. Rendering resources. <https://benedikt-bitterli.me/resources/>.
- Pierre-Alexandre Blanche. 2021. Holography, and the future of 3D display. *Light: Advanced Manufacturing* 2, 4 (2021), 446–459. doi:10.37188/lam.2021.028
- David Blinder, Maksymilian Chlipala, Tomasz Kozacki, and Peter Schelkens. 2021. Photorealistic computer generated holography with global illumination and path tracing. *Optics Letters* 46, 9 (2021), 2188–2191. doi:10.1364/OL.422159
- Carlos Buitrago-Duque and Jorge Garcia-Sucerquia. 2019. Non-approximated Rayleigh–Sommerfeld diffraction integral: advantages and disadvantages in the propagation of complex wave fields. *Applied Optics* 58, 34 (2019), G11–G18. doi:10.1364/AO.58.000G11
- Brent Burley. 2015. Extending the Disney BRDF to a BSDF with integrated subsurface scattering. *SIGGRAPH Course: Physically Based Shading in Theory and Practice*. ACM, New York, NY 19, 7 (2015), 9.
- Praneeth Chakravarthula, Seung-Hwan Baek, Florian Schiffrers, Ethan Tseng, Grace Kuo, Andrew Maimone, Nathan Matsuda, Oliver Cossairt, Douglas Lanman, and Felix Heide. 2022a. Pupil-Aware Holography. *ACM Transactions on Graphics (TOG)* 41, 6 (2022), 1–15. doi:10.1145/3550454.3555508
- 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. doi:10.1145/3355089.3356539
- Praneeth Chakravarthula, Ethan Tseng, Henry Fuchs, and Felix Heide. 2022b. Hogel-Free Holography. *ACM Transactions on Graphics (TOG)* 41, 5 (2022), 1–16. doi:10.1145/3516428
- Chenliang Chang, Kiseung Bang, Gordon Wetzstein, Byoungcho 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. doi:10.1364/OPTICA.406004
- Chun Chen, Byoungcho Lee, Nan-Nan Li, Minseok Chae, Di Wang, Qiong-Hua Wang, and Byoungcho Lee. 2021. Multi-depth hologram generation using stochastic gradient descent algorithm with complex loss function. *Optics Express* 29, 10 (May 2021), 15089–15103. doi:10.1364/OE.425077
- Xindi Chen, Aiming Ge, Dongsheng Zhu, Qiuyang Wang, Jiangbo Wu, and Shuo Liu. 2023. Photorealistic rapid computer-generated holography employing an enhanced path tracing technique with sequence generated trial. *IEEE Photonics Journal* 15, 5 (2023), 1–12. doi:10.1109/JPHOT.2023.3314525
- Suyeon Choi, Brian Chao, Jacqueline Yang, Manu Gopakumar, and Gordon Wetzstein. 2025a. Gaussian wave splatting for computer-generated holography. *ACM Transactions on Graphics (TOG)* 44, 4 (2025), 1–13. 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*. 1–9. 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. doi:10.1145/3478513.3480542
- Seungmi Choi, TaeHwa Lee, Jun Yeong Cha, Suhyun Jo, Hyunmin Ban, Kwan-Jung Oh, Hyunseok Ko, and Hui Yong Kim. 2025b. Phase Distribution Matters: On the Importance of Phase Distribution Alignment (PDA) in Holographic Applications. In *Proceedings of the 33rd ACM International Conference on Multimedia*. 8930–8938. doi:10.1145/3746027.3755797
- Victor Chu, Oscar Pueyo-Ciudad, Ethan Tseng, Florian Schiffrers, Grace Kuo, Nathan Matsuda, Alberto Redo-Sanchez, Douglas Lanman, Oliver Cossairt, and Felix Heide. 2025. Artifact-Resilient Real-Time Holography. *ACM Transactions on Graphics (TOG)* 44, 6 (2025), 1–14. doi:10.1145/3763361
- Joseph W Goodman. 2007. *Speckle phenomena in optics: theory and applications*. Roberts and Company Publishers. doi:10.1117/3.2548484
- Wenzel Jakob, Eugene d’Eon, Otto Jakob, and Steve Marschner. 2014. A comprehensive framework for rendering layered materials. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–14. doi:10.1145/2601097.2601139
- Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, Merlin Nimier-David, Delio Vicini, Tizian Zeltner, Baptiste Nicolet, Miguel Crespo, Vincent Leroy, and Ziyi Zhang. 2022. *Mitsuba 3 renderer*. <https://mitsuba-renderer.org>.
- Changwon Jang, Kiseung Bang, Minseok Chae, Byoungcho Lee, and Douglas Lanman. 2024. Waveguide Holography for 3D Augmented Reality Glasses. *Nature Communications* 15, 1 (2024), 66. doi:10.1038/s41467-023-44032-1
- Bahram Javidi, Artur Carnicer, Arun Anand, George Barbastathis, Wen Chen, Pietro Ferraro, J_W Goodman, Ryoichi Horisaki, Kedar Khare, Malgorzata Kujawinska, et al. 2021. Roadmap on digital holography. *Optics express* 29, 22 (2021), 35078–35118. doi:10.1364/OE.435915
- Andrzej Kaczorowski, Edward Buckley, Richard Stahl, Dmitri Choutov, Theodore Marescaux, and Michael Noonan. 2024. Sub-Wavelength Spatial Light Modulator Enabling Compact Holographic Augmented Reality. In *Frontiers in Optics + Laser Science 2024 (FiO, LS)*. *Frontiers in Optics + Laser Science 2024 (FiO, LS)*, FTu6A.2. doi:10.1364/LS.2024.FTu6A.2
- Koray Kavaklı, Yuta Itoh, Hakan Urey, and Kaan Akşit. 2023. Realistic defocus blur for multiplane computer-generated holography. In *2023 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*. IEEE, 418–426. doi:10.1109/VR55154.2023.00057
- 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. doi:10.1145/3658168
- Juhyeon Kim, Craig Benko, Magnus Wrenninge, Ryusuke Villemin, Zeb Barber, Wojciech Jarosz, and Adithya Pediredla. 2025. A Monte Carlo Rendering Framework for Simulating Optical Heterodyne Detection. *ACM Transactions on Graphics (TOG)* 44, 4 (2025), 1–19. doi:10.1145/3731150
- Julian Knodt, Zherong Pan, Kui Wu, and Xifeng Gao. 2023. Joint UV optimization and texture baking. *ACM Transactions on Graphics (TOG)* 43, 1 (2023), 1–20. doi:10.1145/3617683
- Grace Kuo, Florian Schiffrers, Douglas Lanman, Oliver Cossairt, and Nathan Matsuda. 2023. Multisource holography. *ACM Transactions on Graphics (TOG)* 42, 6 (2023), 1–14. doi:10.1145/3618380
- Ninghe Liu, Kexuan Liu, Yixin Yang, Yifan Peng, and Liangcai Cao. 2025b. Propagation-adaptive 4K computer-generated holography using physics-constrained spatial and Fourier neural operator. *Nature communications* 16, 1 (2025), 7761. doi:10.1038/s41467-025-62997-z
- Zhengze Liu, Yuchi Huo, Yifan Peng, and Rui Wang. 2025a. A Fully-statistical Wave Scattering Model for Heterogeneous Surfaces. *ACM Transactions on Graphics (TOG)* 44, 4 (2025), 1–17. doi:10.1145/3730828
- Juan Magallón, Alfonso Blesa, and Francisco Serón. 2021. SLM simulation and Monte-Carlo path tracing for computer-generated holograms. *SN Computer Science* 2, 233 (2021), 1–16. doi:10.1007/s42979-021-00632-6
- 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. doi:10.1145/3072959.3073624
- Eric Markley, Nathan Matsuda, Florian Schiffrers, Oliver Cossairt, and Grace Kuo. 2023. Simultaneous Color Computer Generated Holography. In *SIGGRAPH Asia 2023 Conference Papers*. ACM, Sydney NSW Australia, 1–11. doi:10.1145/3610548.3618250
- 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. doi:10.1364/AO.48.000H54
- 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. *Journal of the Optical Society of America A* 20, 9 (2003), 1755–1762. doi:10.1364/JOSAA.20.001755
- 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. doi:10.1145/3355089.3356517
- 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. doi:10.1126/sciadv.abg5040
- Feifan Qu, Wenbin Zhou, Xiangyu Meng, Zhenyang Li, Hyunmin Ban, and Yifan Peng. 2025. 83-4: Invited Paper: Filter-Free 3D HoloNet with Hardware-Aware Calibration. In *SID Symposium Digest of Technical Papers*, Vol. 56. Wiley Online Library, 1143–1146. doi:10.1002/sdtp.18380
- Florian Schiffrers, 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. doi:10.1109/ICCP56744.2023.10233716
- Florian Andreas Schiffrers, Grace Kuo, Nathan Matsuda, Douglas Lanman, and Oliver Cossairt. 2025. HoloChrome: Polychromatic Illumination for Speckle Reduction in Holographic Near-Eye Displays. *ACM Transactions on Graphics (TOG)* 44, 3 (2025), 1–18. doi:10.1145/3732935
- Thorsten-Walther Schmidt, Jan Novak, Johannes Meng, Anton S Kaplanyan, Tim Reiner, Derek Nowrouzezahrai, and Carsten Dachsbacher. 2013. Path-space manipulation of physically-based light transport. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 1–11. doi:10.1145/2461912.2461980
- Fabin Shen and Anbo Wang. 2006. Fast-Fourier-transform based numerical integration method for the Rayleigh–Sommerfeld diffraction formula. *Applied Optics* 45, 6 (2006), 1102–1110. doi:10.1364/AO.45.001102
- 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. doi:10.1038/s41586-020-03152-0

- 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. doi:10.1038/s41377-022-00894-6
- Tomoyoshi Shimobaba, Takashi Kakue, and Tomoyoshi Ito. 2015. Review of fast algorithms and hardware implementations on computer holography. *IEEE Transactions on Industrial Informatics* 12, 4 (2015), 1611–1622. doi:10.1109/TII.2015.2509452
- Tomoyoshi Shimobaba, Nobuyuki Masuda, and Tomoyoshi Ito. 2009. Simple and Fast Calculation Algorithm for Computer-Generated Hologram with Wavefront Recording Plane. *Optics Letters* 34, 20 (2009), 3133–3135. doi:10.1364/OL.34.003133
- Shlomi Steinberg, Ravi Ramamoorthi, Benedikt Bitterli, Eugene D'EON, Ling-Qi Yan, and Matt Pharr. 2024a. A generalized ray formulation for waveoptical light transport. *ACM Transactions on Graphics (TOG)* 43, 6 (2024), 1–15. doi:10.1145/3687902
- Shlomi Steinberg, Ravi Ramamoorthi, Benedikt Bitterli, Arshiya Mollazainali, Eugene D'Eon, and Matt Pharr. 2024b. A Free-Space Diffraction BSDF. *ACM Transactions on Graphics (TOG)* 43, 4 (2024), 1–15. doi:10.1145/3658166
- Shlomi Steinberg, Pradeep Sen, and Ling-Qi Yan. 2022. Towards practical physical-optics rendering. *ACM Transactions on Graphics (TOG)* 41, 4 (2022), 1–24. doi:10.1145/3528223.3530119
- Shlomi Steinberg and Ling-Qi Yan. 2021. A generic framework for physical light transport. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–20. doi:10.1145/3450626.3459791
- Shlomi Steinberg and Ling-Qi Yan. 2022. Rendering of subjective speckle formed by rough statistical surfaces. *ACM Transactions on Graphics (TOG)* 41, 1 (2022), 1–23. doi:10.1145/3472293
- Minyuan Sun, Yuan Yuan, Yong Bi, Shuo Zhang, Jianying Zhu, and Wenping Zhang. 2020. Acceleration and expansion of a photorealistic computer-generated hologram using backward ray tracing and multiple off-axis wavefront recording plane methods. *Optics Express* 28, 23 (2020), 34994–35005. doi:10.1364/OE.410314
- PWM Tsang, T-C Poon, and YM Wu. 2018. Review of fast methods for point-based computer-generated holography. *Photonics Research* 6, 9 (2018), 837–846. doi:10.1364/PRJ.6.000837
- Ethan Tseng, Grace Kuo, Seung-Hwan Baek, Nathan Matsuda, Andrew Maimone, Florian Schiffrers, 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. doi:10.1038/s41467-024-46915-3
- Congli Wang, Ni Chen, and Wolfgang Heidrich. 2022. do: A differentiable engine for deep lens design of computational imaging systems. *IEEE Transactions on Computational Imaging* 8 (2022), 905–916. doi:10.1109/TCI.2022.3212837
- Fan Wang, Tomoyoshi Ito, and Tomoyoshi Shimobaba. 2023. High-speed rendering pipeline for polygon-based holograms. *Photonics Research* 11, 2 (2023), 313–328. doi:10.1364/PRJ.474158
- Keita Watanabe, Keigo Yamauchi, and Yuji Sakamoto. 2024. Realistic rendering method for specular reflections with continuous motion parallax in a computer-generated hologram. *Applied Optics* 63, 7 (2024), B126–B133. doi:10.1364/AO.506341
- Haoyu Wei, Xin Liu, Xiang Hao, Edmund Y. Lam, and Yifan Peng. 2023. Modeling Off-Axis Diffraction with the Least-Sampling Angular Spectrum Method. *Optica* 10, 7 (2023), 959–962. doi:10.1364/OPTICA.490223
- Haoyu Wei, Xin Liu, Yuhui Liu, Qiang Fu, Wolfgang Heidrich, Edmund Y Lam, and Yifan Peng. 2025. Learned off-aperture Encoding for Wide Field-of-view RGBD Imaging. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 01 (2025), 1–13. doi:10.1109/TPAMI.2025.3598340
- Xinxing Xia, Pengfei Mi, Yiqing Tao, Xiangyu Meng, Wenbin Zhou, Yingjie Yu, and Yifan Peng. 2025. Multi-illumination-interfered Neural Holography with Expanded Eyebow. *IEEE Transactions on Visualization and Computer Graphics* 31, 11 (2025), 9667–9676. doi:10.1109/TVCG.2025.3616793
- Ling-Qi Yan, Miloš Hašan, Bruce Walter, Steve Marschner, and Ravi Ramamoorthi. 2018. Rendering specular microgeometry with wave optics. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–10. doi:10.1145/3197517.3201351
- Kun Yin, En-Lin Hsiang, Junyu Zou, Yannanqi Li, Zhiyong Yang, Qian Yang, Po-Cheng Lai, Chih-Lung Lin, and Shin-Tson Wu. 2022. Advanced liquid crystal devices for augmented reality and virtual reality displays: principles and applications. *Light: Science & Applications* 11, 1 (2022), 161. doi:10.1038/s41377-022-00851-3
- 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. doi:10.1364/OL.437375
- Yunchen Yu, Mengqi Xia, Bruce Walter, Eric Michielssen, and Steve Marschner. 2023. A full-wave reference simulator for computing surface reflectance. *ACM Transactions on Graphics (TOG)* 42, 4 (2023), 1–17. doi:10.1145/3592414
- Chong Zeng, Yue Dong, Pieter Peers, Hongzhi Wu, and Xin Tong. 2025b. Renderformer: Transformer-based neural rendering of triangle meshes with global illumination. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers*. 1–11. doi:10.1145/3721238.3730595
- Yixin Zeng, Kiseok Choi, Hadi Amata, Kaizhang Kang, Wolfgang Heidrich, Hongzhi Wu, and Min H Kim. 2025a. Designing and Fabricating Color BRDFs with Differentiable Wave Optics. *ACM Transactions on Graphics (TOG)* 44, 6 (2025), 1–17. doi:10.1145/3763275
- Wenbin Zhou, Xiangyu Meng, Li Liao, and Yifan Peng. 2025a. Empowering Head-up AR: Leveraging Holographic Display Engine, Geometry Optical Combiner, and Learned Calibration. In *Proceedings of the SIGGRAPH Asia 2025 Emerging Technologies*. 1–3. doi:10.1145/3757373.3763776
- Wenbin Zhou, Feifan Qu, Xiangyu Meng, Zhenyang Li, and Yifan Peng. 2025b. 3D-HoloNet: fast, unfiltered, 3D hologram generation with camera-calibrated network learning. *Optics Letters* 50, 4 (2025), 1188–1191. doi:10.1364/OL.544816