Physically Grounded Graphics Pipelines for Full-3D Holography

Wenbin Zhou[1,*], Xiangyu Meng[1,*], Jiankai Xing[1,2], Xin Liu[1], Suyeon Choi[3,4], and Yifan Peng[1]
[1]The University of Hong Kong; [2]Tsinghua University; [3]Stanford University; [4]Seoul National University

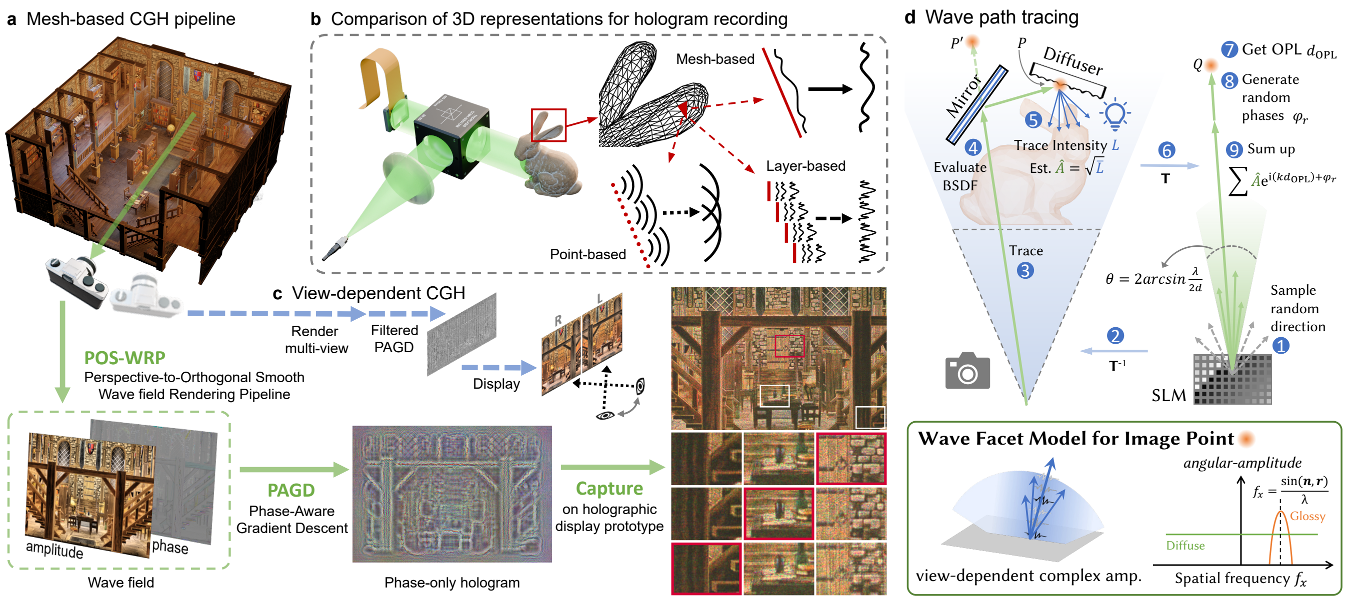

Figure 1: Overview of physically grounded holographic graphics pipelines. (a) Mesh-based CGH converts a textured 3D scene into a displayable hologram. (b) Mesh-, point-, and layer-based representations differ in how they preserve geometry and wave information for hologram recording. (c) Mesh Holography supports continuous focus cues and view-denpendent effects simultaneously. (d) HoloPathTracer uses wave path tracing and a wave facet model to encode photorealistic effects in complex target fields.

Abstract

Holographic displays offer natural 3D visual cues by reconstructing complete light fields through optical interference. Recent advances in computer-generated holography (CGH) have improved display quality, expanded etendue, and reduced computation time. However, most CGH pipelines still rely on image-based supervision, separating rendering with wave propagation, which cannot preserve complete light transport with phase information. As a result, they often lose important 3D cues such as continuous focus, parallax, occlusion, and material-dependent visual effects.

Mesh representations, which are widely studied in computer graphics, naturally retain continuous geometry, visibility, texture, and material information. In this extended abstract, we discuss two recent mesh-based holography pipelines that aim to address these limitations. In the first, we develop a camera-calibrated wave propagation model tailored to mesh-derived wave fields, together with complex-valued optimization, enabling optical reconstructions that preserve continuous depth and view-dependent structure. In the second, we further combine path tracing with Rayleigh–Sommerfeld integration to capture complete wave-field transport in photorealistic scenes with reflective, refractive, glossy, and globally illuminated effects. Together, our pipelines bridge CGH and computer graphics, advancing CGH toward more faithful full-3D visual reconstruction.

Keywords

computer-generated holography, wave optics, physically based rendering, calibration, computational displays

1 Motivation

Holography records the complete information of object waves, including both amplitude and phase, in an interference pattern. Through optical reconstruction, it can reproduce the original wave field such that the reconstructed light is, in principle, indistinguishable from that emitted by the real object. Recent CGH systems have improved display quality through camera-in-the-loop calibration, time multiplexing, additional light sources [Chao et al. 2024; Choi et al. 2022; Peng et al. 2020; Xia et al. 2025]; and reduced computational resources using advanced deep learning models and implicit neural representations [Ban et al. 2026; Zhou et al. 2025]. However, many existing computer-generated holography pipelines still rely on discretized image-based scene representations, such as RGB-D images, focal stacks, or light fields, as algorithmic inputs [Chakravarthula et al. 2022; Padmanaban et al. 2019; Shi et al. 2021]. In such pipelines, intensity-domain graphics rendering is first performed, and the resulting images are then converted into

*Equal contribution. Correspondence: evanpeng@hku.hk.

holograms through approximate layer- or light-field-based wave propagation. This rendering-then-propagation paradigm separates graphics rendering from wave-optical field formation and therefore cannot fully preserve the complete wave-field information of the original 3D scene, limiting depth continuity, view-dependent effects, occlusion handling, and material fidelity, especially for scenes involving reflective, refractive, glossy, or globally illuminated light transport.

In our two works, we address this separation by keeping mesh-based graphics representations coupled to complex wave-field formation [Meng et al. 2026; Zhou et al. 2026]. In Mesh Holography, we convert mesh scenes into complex wave fields through a perspective-to-orthogonal smooth wave-field rendering pipeline, and then optimize phase-only holograms with complex-valued supervision. We further unveil Flying Planes, a mesh-wave-field-hologram dataset tailored for true 3D holographic display calibration, and use it to train a camera-calibrated propagation model that improves display quality under real SLM constraints. Together, these components enable holograms to support continuous depth cues and view dependence. In HoloPathTracer, we focus on complex-field rendering. We integrate path tracing with Rayleigh–Sommerfeld propagation, bringing photorealistic light transport effects, including reflection, refraction, glossy appearance, and global illumination, directly into the hologram patterns. Together, our works move CGH from image-only supervision toward mesh-based wave-field modeling, where amplitude and phase remain tied to the 3D scene through optical reconstruction.

2 Mesh Holography

In Mesh Holography, we address the representation gap between graphics scenes and holographic modulation [Meng et al. 2026]. Instead of flattening a 3D scene into RGB-D layers, focal stacks, or sampled light fields, we start from textured meshes, where geometry, visibility, texture, and shading remain tied to the original scene. This is useful because meshes preserve continuous surfaces and explicit occlusions, while also matching the data structures used in modeling, simulation, and physically based rendering. The key step in our pipeline is POS-WRP, a perspective-to-orthogonal smooth wave-field rendering pipeline. POS-WRP preserves perspective cues, maps the visible mesh geometry into the SLM coordinate system, and propagates tilted surface fields to the hologram recording plane. Our target is therefore not only a rendered intensity image, but a mesh-derived complex wave field whose amplitude and phase are organized by the underlying 3D scene.

This mesh-centered formulation keeps visibility, surface orientation, and texture sampling explicit before hologram synthesis, preserving depth ordering and pupil-dependent views in the complex target field.

Since phase-only spatial light modulators (SLMs) are widely adopted for their high diffraction efficiency, we convert the target complex field into a displayable phase-only hologram. Our Phase-aware Graduate Descent (PAGD) optimizer supervises reconstruction in the complex domain, enabling phase and amplitude to be optimized jointly rather than discarding phase after rendering. This is crucial because defocus, off-axis views, and hidden-surface behavior depend on phase consistency beyond the supervised focal planes. We further account for hardware mismatch through the Flying Planes dataset, which provides paired mesh scenes, wave fields, RGB-D renderings, and holograms, together with camera-calibrated propagation that adapts synthetic wave propagation to real SLM behavior. Our optical validation demonstrates refocusing, pupil-dependent parallax, and view-dependent occlusion changes, showing how a standard mesh scene can be converted into a calibrated hologram for physical holographic display.

3 HoloPathTracer

HoloPathTracer asks a complementary question: if the target scene contains reflections, refractions, glossy materials, and indirect illumination, how should CGH generate the complex wave field before phase-only encoding [Zhou et al. 2026]? Instead of first rasterizing the scene into RGB-D layers, focal stacks, or light fields, we formulate target-field generation as wave path tracing. A sampled path follows ordinary graphics light transport through mesh geometry and BSDF interactions, but we accumulate its contribution to the hologram plane as a complex wave term. In this way, we solve radiance estimation and Rayleigh–Sommerfeld propagation in the same Monte Carlo process, so path length, visibility, material scattering, amplitude, and phase remain coupled.

We further introduce a wave-facet formulation to connect local surface interactions with wave propagation. Each sampled surface event can be viewed as a small wave-emitting facet whose contribution is propagated to the hologram recording plane. This makes our estimator compatible with standard physically based rendering ingredients, including glossy reflection, mirror reflection, transparent refraction, and indirect bounces, while still producing a target complex field rather than only an image. As a result, we encode effects that are normally difficult for layer-based CGH, such as view-dependent highlights, refractive imagery, continuous defocus, and globally illuminated appearance, before hologram optimization.

We handle efficiency by reusing the structure of sampled light paths. Path Reuse produces multiple random-like complex fields for time-multiplexed display with little additional path-sampling cost, which is important because perceptual holography benefits from incoherently averaging several coherent reconstructions. We also use an ambient radiance cache to approximate indirect components and improve convergence. After wave-field rendering, we convert the coherent target fields to phase-only holograms with complex-amplitude supervision, keeping the optimization aligned with the wave-optical objective instead of matching only intensity. In this sense, HoloPathTracer extends our mesh-based CGH work from calibrated geometric reconstruction to photorealistic holography with material-dependent and global light transport effects.

4 Conclusion

Together, our Mesh Holography and HoloPathTracer works show a graphics-to-wave route for CGH. In Mesh Holography, we make mesh-derived fields displayable with high fidelity through calibrated propagation and complex-valued optimization, while in HoloPathTracer, we make target fields photorealistic through path-traced wave transport. Overall, our work highlights the value of treating geometry, materials, phase, calibration, and phase-only encoding as a single coupled holographic pipeline.

References

Hyunmin Ban, Wenbin Zhou, and Yifan Peng. 2026. Towards Edge Holography via Implicit Neural Representation and Compression. *IEEE Transactions on Visualization and Computer Graphics* (2026).

Praneeth Chakravarthula, Ethan Tseng, Henry Fuchs, and Felix Heide. 2022. Hogel-Free Holography. *ACM Trans. Graph.* 41, 5 (2022), 1–16. doi:10.1145/3516428

Brian Chao, Manu Gopakumar, Suyeon Choi, Jonghyun Kim, Liang Shi, and Gordon Wetzstein. 2024. Large étendue 3D holographic display with content-adaptive dynamic Fourier modulation. In *SIGGRAPH Asia 2024 Conference Papers*. 1–12.

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 (SIGGRAPH '22)*. 1–9. doi:10.1145/3528233.3530734

Xiangyu Meng, Wenbin Zhou, and Yifan Peng. 2026. Mesh-represented and learning-empowered hologram synthesis for full 3D holographic displays. *Nature Communications* (2026).

Nitish Padmanaban, Yifan Peng, and Gordon Wetzstein. 2019. Holographic Near-Eye Displays Based on Overlap-Add Stereograms. *ACM Trans. Graph.* 38, 6 (2019), 1–13. doi:10.1145/3355089.3356517

Yifan Peng, Suyeon Choi, Nitish Padmanaban, and Gordon Wetzstein. 2020. Neural Holography with Camera-in-the-Loop Training. *ACM Trans. Graph.* 39, 6 (2020), 1–14. doi:10.1145/3414685.3417802

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

Xinxing Xia, Pengfei Mi, Yiqing Tao, Xiangyu Meng, Wenbin Zhou, Yingjie Yu, and Yifan Peng. 2025. Multi-illumination-interfered Neural Holography with Expanded Eyebox. *IEEE Transactions on Visualization and Computer Graphics* (2025).

Wenbin Zhou, Xiangyu Meng, Jiankai Xing, Xin Liu, Suyeon Choi, and Yifan Peng. 2026. HoloPathTracer: Fast and Accurate Wave Path Tracing for Holography. *ACM Transactions on Graphics* 45, 4 (2026). doi:10.1145/3811351

Wenbin Zhou, Feifan Qu, Xiangyu Meng, Zhenyang Li, and Yifan Peng. 2025. 3D-HoloNet: fast, unfiltered, 3D hologram generation with camera-calibrated network learning. *Optics Letters* 50, 4 (2025), 1188–1191.

‘’