

Hogel-Free Holography

PRANEETH CHAKRAVARTHULA, Princeton University and UNC Chapel Hill

ETHAN TSENG, Princeton University

HENRY FUCHS, UNC Chapel Hill

FELIX HEIDE, Princeton University

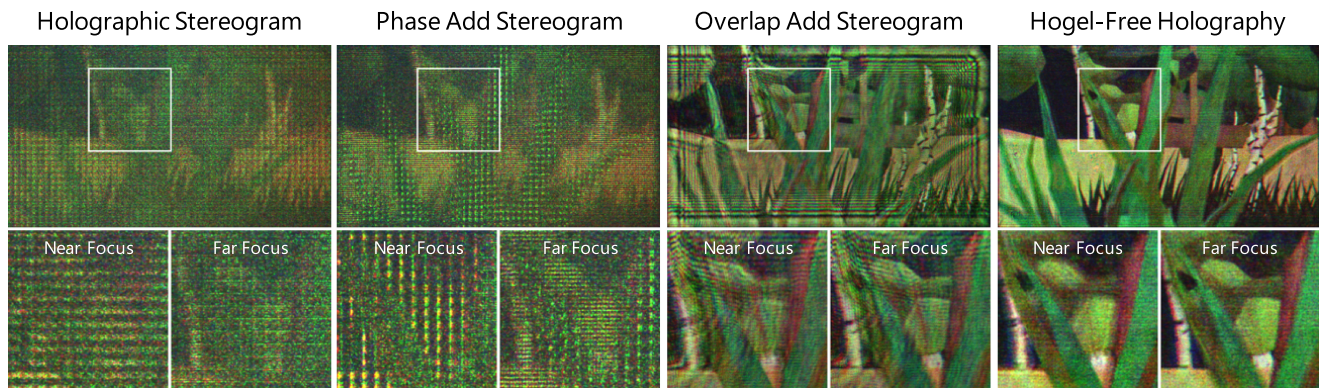


Fig. 1. Hogel-free Holography. We propose a hogel-free approach to high-quality 3D holographic display with accurate depth- and view-dependent effects, and without sacrificing spatio-angular resolution. Existing approaches rely on computing sub-holograms, so-called hogels, for achieving such effects. However, the chosen size of hogels is typically scene dependent and follows a tradeoff between angular and spatial resolution of holographic imagery. We lift these limitations by formulating a holographic forward model and phase retrieval that takes RGB-D light fields as input and directly optimize the target phase, without spatial segmentation into hogels or phase encoding approaches. The proposed method achieves high-quality 3D holograms with accurate parallax and depth focus effects.

Holography is a promising avenue for high-quality displays without requiring bulky, complex optical systems. While recent work has demonstrated accurate hologram generation of 2D scenes, high-quality holographic projections of 3D scenes has been out of reach until now. Existing multiplane 3D holography approaches fail to model wavefronts in the presence of partial occlusion while holographic stereogram methods have to make a fundamental tradeoff between spatial and angular resolution. In addition, existing 3D holographic display methods rely on heuristic encoding of complex amplitude into phase-only pixels which results in holograms with severe artifacts. Fundamental limitations of the input representation, wavefront modeling, and optimization methods prohibit artifact-free 3D holographic projections in today's displays.

To lift these limitations, we introduce hogel-free holography which optimizes for true 3D holograms, supporting both depth- and view-dependent

effects for the first time. Our approach overcomes the fundamental spatio-angular resolution tradeoff typical to stereogram approaches. Moreover, it avoids heuristic encoding schemes to achieve high image fidelity over a 3D volume. We validate that the proposed method achieves 10 dB PSNR improvement on simulated holographic reconstructions. We also validate our approach on an experimental prototype with accurate parallax and depth focus effects.

CCS Concepts: • **Hardware** → **Displays and imagers**;

Additional Key Words and Phrases: Computational optics, computational displays, holography

ACM Reference format:

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

1 INTRODUCTION

Augmented reality displays are emerging technologies, and, while existing displays achieve convincing 2D overlays, practical 3D displays for AR are an open challenge. Holographic displays offer perhaps the most promising avenue toward achieving such 3D displays of the future. These displays rely on **spatial light modulators (SLMs)** that *theoretically* offer fine-grained modulation of light waves in a lightweight and compact form factor, allowing one to produce 3D volumetric imagery. However, today's SLMs unfortunately cannot modulate both amplitude and phase component of light simultaneously, although a phase

The authors are grateful to Joel S. Kollin and Bernard Kress for lending the SLM, Roarke Horstmeyer for lending the laser diode controller, and Jim Mahaney for hardware support during the COVID-19 pandemic. This research is supported by NSF Grants no. 1840131 and no. 1405847. Felix Heide was supported by an NSF CAREER Award (2047359), a Sony Young Faculty Award, and a Project X Innovation Award. Authors' addresses: P. Chakravarthula, E. Tseng, and F. Heide, 35 Olden St, Princeton, NJ 08544; emails: cpk@cs.unc.edu, efteng@cs.princeton.edu, fheide@cs.princeton.edu; H. Fuchs, 216 Fred Brooks CS Building, 201 S. Columbia St, Chapel Hill, NC 27599-3175; email: fuchs@cs.unc.edu.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2022 Copyright held by the owner/author(s).
 0730-0301/2022/10-ART178
<https://doi.org/10.1145/3516428>

modulating SLM is generally preferred for its diffraction efficiency. As a result, sophisticated algorithms are required for determining the appropriate phase modulation on the SLM for producing a desired 3D holographic projection.

Recently, **computer generated holography (CGH)** techniques have demonstrated success in optimizing high-fidelity phase-only holograms for 2D imagery [Chakravarthula et al. 2019, 2020; Peng et al. 2020]. However, algorithms for generating accurate high-resolution 3D holograms have not been demonstrated so far. Existing approaches to 3D holography compute a full Fresnel integral *complex hologram* for every point in the 3D scene. These point-based methods [Maimone et al. 2017] impose a handcrafted target phase to improve the 3D image quality, but this prohibits end-to-end phase optimization and restricts existing methods to one-step propagation followed by heuristic encoding into a phase-only hologram. Such one-step *point-based integration* approaches have only been successfully demonstrated for very sparse scenes with no more than a single 3D object. Very recently, Shi et al. [2021] demonstrated that point-based integration methods can be extended to more complicated 3D scenes using deep learning and double-phase amplitude encoding. Their method takes RGB-D image data as input which do not accurately model a 3D scene in terms of parallax, especially at object interfaces where the light propagates in an off-axis fashion around occlusion edges. This leads to physically inaccurate modeling of defocus and occlusion effects, as well as visible light leakage and ringing at the occlusion boundaries. As such, when a scene distribution is specified in a 3D domain, conventional techniques tend to decompose it into independent discrete scatterers: planes, segments, and points, which are treated separately. If the axial extent of this distribution is increased, intolerable cross talk between the independent wave fronts will occur [Piestun et al. 1997, 1996].

In contrast to RGB-D images, a light field provides both depth and parallax information at the object surface, representing a physically accurate description of the scene light transport. However, existing light field-based approaches rely on computing non-overlapping hologram patches of elemental light field images, called *hogels*. These hogels capture the effects of the light field arriving from the target scene [Kang et al. 2016], providing an accurate representation of defocus and occlusion effects. However, the choice of the hogel size is akin to the size of a microlens in a lenticular array-based light field display [Lanman and Luebke 2013], and subject to a fundamental spatio-angular resolution tradeoff of the final 3D holographic projection. Recently, researchers have proposed **overlap-add stereograms (OLAS)** with the goal to overcome this tradeoff by using overlapping hogels [Padmanaban et al. 2019]. Unfortunately, this method is also fundamentally limited by the underlying hogel representation which prohibits direct optimization of the SLM phase for producing true 3D phase-only holograms. As such, OLAS is confined to non-iterative one-step propagation followed by double-phase encoding of complex amplitude into a phase-only hologram, resulting in reduced spatial resolution and ringing—especially around parallax regions that are the challenge at hand.

In this work, we propose **Hogel-Free Holography (HFH)** as the first approach for computing physically accurate 3D holograms that lifts these existing limitations by departing from

explicit hogel assumptions and, as such, allows for high-quality artifact-free true 3D holograms for natural scenes (Figure 1). Specifically, in contrast to existing approaches that use hogels to approximate the full wavefront with a discrete set of plane waves, we directly compute the SLM phase pattern for a *full wavefront* of the 3D scene. To this end, we pose the SLM phase synthesis as a phase-relaxed optimization over the *continuous volume*, which we solve using first-order solvers. With this approach, we are able to lift the requirements of handcrafted object phase constraints, hogels, or heuristic encoding schemes that limited all previous 3D holography methods. Moreover, the proposed approach accurately models per-pixel defocus and occlusion effects unlike existing point-based methods and overcomes the spatio-angular resolution tradeoff native to other light field-based approaches. By exploiting known wave propagation physics, we are able to *optimize for the entire continuous depth volume*. We validate the proposed method in simulation and experimentally, and we demonstrate that it achieves accurate occlusion and defocus cues which was not previously possible with existing methods.

Specifically, we make the following contributions:

- We introduce a method for computing physically accurate true 3D holograms via RGB-D light fields without using hogels.
- Our method admits direct optimization of the SLM phase for a 3D scene using first-order solvers. We achieve this by posing 3D phase retrieval as a phase-relaxed continuous volume optimization, without the need for heuristic phase encoding schemes.
- We validate the method in simulation and experimentally using a prototype near-eye holographic display. We achieve an order of magnitude lower reconstruction error than previous methods throughout the depth volume of all tested scenes.

Overview of Limitations. While the proposed method is capable of achieving unprecedented image quality of 3D holograms, the proposed method does not generate these holograms at real-time framerates but requires 300 s per hologram. We envision that further research into phase optimization, perhaps leveraging feed-forward neural network predictions as initialization, could remedy this in the future. The SLM used in our experimental setup is currently unsupported due to its age and could not be recalibrated which resulted in some degradation of the experimental image quality. While orthogonal to the proposed method, recent hardware-in-the-loop calibration techniques [Chakravarthula et al. 2020; Peng et al. 2020] may mitigate such issues with affordable legacy hardware in the future.

2 RELATED WORK

Our work is related to a large body of display and holography research. In this section, we summarize the most relevant display technologies and CGH algorithms which our work builds upon.

2.1 Holographic Near-Eye Displays

Holographic displays have demonstrated powerful features such as compact form factor, variable focus control, and aberration correction, all of which are essential for eye-wear displays. In order

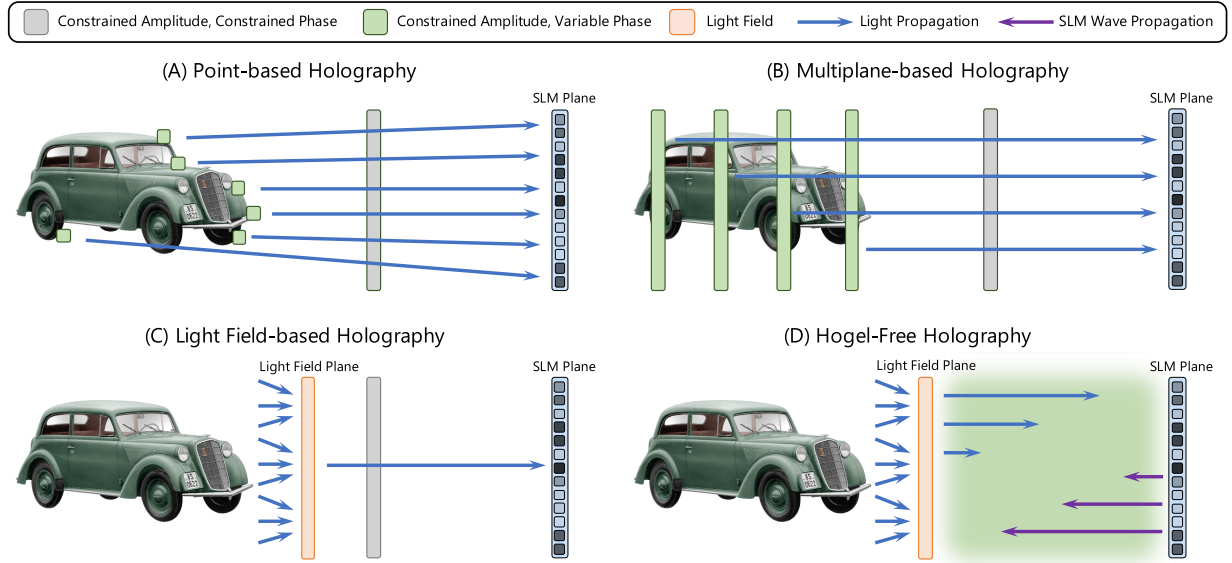


Fig. 2. Comparison of 3D CGH Algorithms. Early point-based (a) and layer-based (b) approaches specified the desired amplitude measurement at specific points or layers, but this imposes a strict constraint on the phase and amplitude of the wavefront as illustrated in the figure. This constraint prohibits optimization of the SLM phase so these methods employ heuristic amplitude phase encodings [Maimone et al. 2017] which result in low-quality imagery. Furthermore, these methods do not generate accurate occlusion and defocus effects. (c) Recent approaches [Padmanaban et al. 2019] have utilized light fields to generate 3D holograms with correct parallax; however, these methods still suffer from the joint phase and amplitude constraint. (d) Hogel-Free Holography is the first technique to enable optimization of the SLM phase by translating the joint constraint into an amplitude-only volume constraint that lives continuously throughout the volume.

to reduce the form factor, researchers have used multi-functional **Holographic Optical Elements (HOEs)** [Jang et al. 2019, 2017; Li et al. 2016; Maimone et al. 2017] and waveguides [Yeom et al. 2015] to relay the projected imagery into the eye. Holographic displays typically suffer from a tiny eyebox and poor image quality. However, significant advances are already being made in recent years to solve these problems. Recent approaches such as increasing the eyebox size by using eyetracking [Jang et al. 2019, 2017] and employing specially designed HOEs [Jang et al. 2019; Xia et al. 2020] or scattering elements [Kuo et al. 2020; Park et al. 2019] show tremendous promise for achieving practical displays in the future. In addition, employing eyetracking [Lu et al. 2020] might further decrease computational burden [Deng et al. 2021] and facilitate reflectance, focus control, and aberration correction [Chakravarthula et al. 2021], overcoming challenges related to the limited space-bandwidth product of existing SLMs. While recent algorithmic advances have made it possible to generate high-quality 2D holograms [Chakravarthula et al. 2020; Peng et al. 2020], modest image quality of 3D holographic projections with visible artifacts [Shi et al. 2021] remains as a critical limitation of today's holographic displays.

2.2 2D Holographic Phase Retrieval

Complex holograms for a given target image can be computed by propagating the light waves emanating from the scene to the SLM (hologram) plane, where they are interfered with the reference beam [Benton and Bove Jr 2008]. While this hologram field on the SLM is *complex* and has both amplitude and phase, existing SLMs can only modulate either amplitude or phase but not both, simultaneously. As phase-only SLMs are typically preferred

for better efficiency, this necessitates the computation of *phase-only holograms* that are capable of producing the diffraction field that most closely mimics the target image. In order to do so, a variety of techniques have been proposed which can be broadly categorized into direct and iterative phase retrieval methods. Direct methods attempt to encode the complex wavefield at the SLM into a phase-only representation using amplitude-phase coding [Burckhardt 1970; Hsueh and Sawchuk 1978; Lee 1970] but typically result in reduced image quality. Iterative methods, on the other hand, improve the image quality by minimizing the reprojection error of the target intensity distribution using iterative optimization but are usually slow [Chakravarthula et al. 2019; Gerchberg 1972]. Remarkably, high-quality 2D holograms have been recently demonstrated on experimental prototypes using hardware-in-the-loop optimization [Chakravarthula et al. 2020; Peng et al. 2020].

2.3 3D Holographic Phase Retrieval

Algorithms for 3D computer-generated holography have been the focus of active research for several years. The most popular existing methods use RGB-D or multiplane images as shown in Figure 2 to represent 3D objects. These methods fail to accurately model wavefronts in the presence of occlusions in the scene. We overcome this challenge by representing the 3D scene as a light field which implicitly models occlusions and parallax, thereby also modeling the wavefronts emanating from occluded scene parts. In this section, we discuss existing 3D holographic phase retrieval methods as summarized in Figure 2.

Point Cloud and Polygonal Mesh-Based Methods. Researchers have leveraged 3D scene representations such as point clouds or polygonal meshes for computing holograms [Leseberg and Frère

1988; Waters 1966], where diffraction is simulated for every scene point or polygon (see Figure 2(a)). Computing such holograms is generally computationally expensive and results in a complex field at the SLM plane. This complex field is then converted into a phase-only hologram using heuristic encodings such as double-phase amplitude encoding [Hsueh and Sawchuk 1978; Maimone et al. 2017]. Similarly, a look-up table of elemental fringes can also be precomputed [Lucente 1993; Shi et al. 2017]. Existing encoding schemes have been shown to result in reduced spatial resolution and image quality [Chakravarthula et al. 2019].

Multiplane Images and Focal Stack-Based Methods. Instead of computing the wave propagation for millions of points, a 3D object can be represented as a stack of intensity layers [Choi et al. 2021; Eybposh et al. 2020; Zhang et al. 2016; Zhao et al. 2015] (see Figure 2(b)). Wave propagation methods such as the inverse Fresnel transform or angular spectrum propagation are typically used for propagating the waves from several layers of the 3D scene toward the SLM plane, where they are interfered to produce a complex hologram. Although this approach can be implemented efficiently, it cannot support continuous focus cues and accurate occlusion due to discrete plane sampling [Kuo et al. 2020; Makowski et al. 2007; Peng et al. 2020; Zhang et al. 2017]. A further approximation to the layer-based methods is to determine the focal depth of the user (i.e., distance of the object to which the user fixates) via an eyetracker and adjusting the focal plane of the 2D holographic projection to match the user's focal distance [Chakravarthula et al. 2020; Maimone et al. 2017]. While emulating a 3D scene by adaptively shifting a 2D holographic projection in space is computationally more efficient, operating in a varifocal mode under-utilizes the capabilities of a holographic display. Moreover, achieving natural focus cues and physically accurate occlusion effects still remains a challenge.

Light Field-Based Methods. To support occlusion and depth-dependent effects, a light field can be encoded into a hologram partitioned spatially into elementary hologram patches, called “hogels” [Zhang et al. 2015], similar to elementary images in a light field (see Figure 2(c)). These hogels produce local ray distributions that reconstruct multiple (light field) views [Lucente and Galyean 1995; Smithwick et al. 2010; Yamaguchi et al. 1993]. Such holograms which encode a light field are dubbed “holographic stereograms.” Conventional stereograms, where hogels are out of phase with each other, suffer from a lack of focus cues and limited depth of field [Lucente and Galyean 1995]. To keep the hogels of a holographic stereogram in phase across the hologram, researchers have introduced an additional phase factor to calculate what is called a **phase-added stereogram (PAS)** [Yamaguchi et al. 1993]. However, akin to a microlens array-based light field display [Lanman and Luebke 2013], stereograms suffer from the fundamental spatio-angular resolution tradeoff: a larger hogel size leads to a decreased spatial resolution. This fundamental limitation does not allow for holographic stereogram projections of high spatial resolution. However, recent methods have attempted to overcome this tradeoff [Blinder and Schelkens 2018; Padmanaban et al. 2019; Park and Askari 2019] via **Short-Time Fourier Transform (STFT)** inversion. These methods do not match the image quality achieved for 2D holograms [Chakravarthula et al. 2020; Peng et al. 2020]




Fig. 3. Multiplane vs. RGB-D vs. Proposed. A multiplane image stack results in aberrations due to leaked wavefronts from depth discontinuities (top). A single RGB-D image does not contain parallax information and hence cannot correctly model the wavefronts emanating from occluded sources (middle). Using an RGB-D light field, we are able to model physically accurate wavefronts arising from occluded points (bottom). Note that parts of the scene background that are occluded in one camera view are revealed in the other views of the light field.

and suffer from severe artifacts around object discontinuities due to sub-optimal inversion of the STFT. In this work, we lift these limitations and demonstrate an iterative optimization method for 3D holographic phase retrieval (see Figure 2(d)).

3 HOGL-FREE HOLOGRAPHY

In this section, we present our method for computing 3D phase holograms supporting depth and view-dependent effects, such as defocus, parallax, and occlusion. We first describe RGB-D light fields as an input representation before explaining the proposed method. To describe our method, we start by formulating the SLM phase optimization as a consensus minimization across light field angular views, which we recast as a single term minimization problem in the wavefront domain (Section 3.2) that is solved continuously in the volume (Section 3.5).

3.1 Input Scene Representation

We use an RGB-D light field as input to our method because this representation contains parallax and mutual occlusion information. As depicted in the bottom row of Figure 3, light field arrays consist of multiple views of the scene which allows us to capture




Fig. 4. Holographic Stereograms vs. Light Fields. Holographic stereograms partition the SLM phase into discrete components, each of which corresponds to an angular view of a light field coming from the scene. This is akin to a microlens array which measures the light field using discretized lenslet arrays and suffers from a fundamental spatio-angular resolution tradeoff. We depart from hogel-based stereogram approaches to compute the full wavefield and demonstrate high-quality true 3D holograms.

missing scene content from partial occlusions for a single camera view. In contrast, commonly employed scene representations such as RGB-D images (Figure 3, middle row) and multiplane images (Figure 3, top row) insufficiently describe the scene due to depth discontinuities and missing information in the presence of occlusions. Specifically, for computing the 3D hologram, every point on the visible scene is assumed to be a *spherical wave emitter* whose diffraction pattern is independently superposed on the SLM. As a result, wavefronts in the presence of occluders are inaccurately modeled resulting in visible artifacts at the occlusion boundaries (see Figure 3). Existing methods relying on single-view RGB-D images or multiplane images as inputs ignore *diffraction* at depth discontinuities, causing severe ringing artifacts and noise. Note that focal stacks suffer from the same limitations.

3.2 Consensus 3D Holograms from Light Fields

Our aim is to directly optimize the SLM phase pattern for a *full propagating wavefront* of a 3D scene as opposed to existing light field-based stereogram approaches which solve *piecewise planar approximations* of the wavefronts represented by hogels, as depicted in Figure 4. In other words, for a given light field $L(x, y, \theta_x, \theta_y)$ parameterized by positions x, y and angles θ_x, θ_y , we want to find a phase-only SLM pattern Φ that minimizes the following objective:

$$\Phi_{\text{OPT}} = \min_{\Phi} \sum_{\{\theta_x, \theta_y\}} \left\| |\sigma_{(\theta_x, \theta_y)}\{\mathcal{P}(e^{j\Phi}, z_{\text{WR}})\}|^2 - L(\theta_x, \theta_y) \right\|, \quad (1)$$

where $U = e^{j\Phi}$ is the wavefield at the SLM, $\mathcal{P}(U, z_{\text{WR}})$ is a wave propagation kernel, such as angular spectrum propagation, which propagates the input field U to the **wave recorder (WR)** plane at the distance z_{WR} where the light field is measured, and

$$\sigma_{(\theta_x, \theta_y)} : U(x, y) \rightarrow L(x, y, \theta_x, \theta_y)$$

is the wavefront sampling operator parameterized by the light field angular views θ_x, θ_y . The positions x, y are ignored in Equation (1) for brevity. The operator $\sigma_{(\theta_x, \theta_y)}$ and its inverse $\sigma_{(\theta_x, \theta_y)}^{-1}$ are detailed in the following sections.

For an $M \times N$ light field input, performing the above optimization in the light field domain requires optimizing over MN summation terms. While such an objective may be minimized using

consensus optimization [Boyd et al. 2004], we instead solve it in the wave domain in order to pose the objective as a single term. Specifically, the objective in the wave domain becomes

$$\Phi_{\text{OPT}} = \min_{\Phi} \left\| \mathcal{P}(e^{j\Phi}, z_{\text{WR}}) - \sum \sigma_{(\theta_x, \theta_y)}^{-1} \left\{ \sqrt{L(\theta_x, \theta_y)} e^{j\Psi(\theta_x, \theta_y)} \right\} \right\|, \quad (2)$$

where $\sigma_{(\theta_x, \theta_y)}^{-1}$ maps discrete light field angular views to their corresponding complex wavefront description. Note that light fields are intensity-only images and the associated phase $\Psi(\theta_x, \theta_y)$ corresponding to each angular view of the light field is typically missing, which makes solving the above problem by inverting σ ill-posed.

Furthermore, besides being ill-posed, notice that Equation (2) is defined in the complex domain. As most commodity SLMs only support phase modulation, existing 3D holography methods have to resort to heuristic phase encoding schemes that turn the complex solution into a phase-only solution.

To overcome both challenges in solving Equation (2), we ameliorate the complex domain objective by relaxing the target phase constraints in Equation (2) into a continuous volume amplitude constraint and solve it using stochastic gradient descent to find a phase-only true 3D hologram. We describe this approach in detail in the following sections (along with operators $\sigma_{(\theta_x, \theta_y)}$, $\sigma_{(\theta_x, \theta_y)}^{-1}$). Specifically, Section 3.3 expresses σ in Equation (1) as a **Windowed Fourier Transform (WFT)** and Section 3.4 describes how we cure the ill-posed inversion of σ described in Equation (2). Finally, Section 3.5 describes how we relax the phase constraints in Equation (2) and how we pose the 3D phase retrieval as an optimizing amplitude across a continuous 3D volume.

3.3 Light Field Forward Model

We will describe how a source wavefield $U_{\text{SRC}}(u, v; 0)$ is related to a given light field $L(x, y, \theta_x, \theta_y)$ measured on a certain wave recorder plane at a distance z_{WR} :

$$U_{\text{SRC}}(u, v; 0) \rightarrow U_{\text{WR}}(x, y; z_{\text{WR}}) \longleftrightarrow L(x, y, \theta_x, \theta_y), \quad (3)$$

where $U_{\text{WR}}(x, y; z_{\text{WR}})$ is the underlying complex wavefront on the wave recorder plane. We formalize the forward pass $U_{\text{SRC}}(u, v) \rightarrow L(x, y, \theta_x, \theta_y)$, that is, how a source field is propagated from a scene to a wave recorder and captured as a light field. To this end, let us recall angular spectrum decompositions [Goodman 2005]: The angular spectrum propagation of a source field $U_{\text{SRC}}(u, v; 0)$ by some distance z is described by

$$U_{\text{WR}}(x, y; z) = \mathcal{F}^{-1}(\mathcal{F}\{U_{\text{SRC}}(u, v; 0)\} \circ H(f_x, f_y; z)), \quad (4)$$

where U_{WR} is the field on a WR plane where the light field is measured and H is the angular spectrum propagation kernel given by

$$H(f_x, f_y; z) = \begin{cases} \exp \left[j2\pi \frac{z}{\lambda} \sqrt{1 - (\lambda f_x)^2 - (\lambda f_y)^2} \right], & \sqrt{f_x^2 + f_y^2} < \frac{1}{\lambda} \\ 0, & \text{otherwise} \end{cases} \quad (5)$$

with z being the propagation distance of the modulated monochromatic wave of wavelength λ with spatial frequencies f_x, f_y in the x and y directions, respectively. The Fourier transform $\mathcal{F}(\cdot)$ in the above Equation (4) decomposes the source wave field into its component plane waves, that is, the angular spectrum.

The spatial frequencies (f_x, f_y) present in the angular spectrum relate to the direction cosine/light field camera view angles (θ_x, θ_y) [Goodman 2005] as

$$f_{\{x,y\}} = \frac{\sin(\theta_{\{x,y\}})}{\lambda}. \quad (6)$$

Therefore, to model wavefronts in the presence of partial occlusions via light fields, it is sufficient to model *local* plane waves with corresponding spatial frequencies as described by Equation (6). However, every spatial frequency component of the angular spectrum extends over the entire (x, y) -domain, but only sparse discrete views are available from the input light field. However, note that a light field is nothing but a coerced representation of a continuous wave field. As a result, the underlying holographic wavefield is capable of representing continuous novel views and at any focal plane, these views are aggregated to form an appropriately defocused image.

Wavefront View Sampling. To tackle the mismatch between discrete light field views and the dense frequency support, locally sampling the propagated wavefront U_{WR} offers a solution. The spatially localized frequency distribution corresponding to light field angular views can be obtained by *spatially limiting the Fourier spectrum to only a constrained set of frequencies*, which we do by using a windowing function $w(x, y)$ that is non-zero only for a limited support [Ziegler et al. 2007]. The resulting local frequency spectrum can be computed by a WFT as

$$S(x, y, f_x, f_y) = \sigma_{(\theta_x, \theta_y)} \{U_{WR}(x, y)\} \times \iint_{-\infty}^{\infty} U_{WR}(x', y') w(x' - x, y' - y) e^{-j2\pi(f_x x' + f_y y')} dx' dy'. \quad (7)$$

Multiplying the field on the wave recorder U_{WR} with the window function w (e.g., Hamming window) suppresses the field outside the window, causing localization. The light field L measured on the wave recorder plane can then be described as the squared magnitude of the angular distributions of the local frequency spectrum S , that is,

$$L(x, y, \theta_x, \theta_y) = |S(x, y, f_x, f_y)|^2 = |\text{WFT}\{U_{WR}(x, y)\}(x, y, f_x, f_y)|^2. \quad (8)$$

In other words, we can formulate our RGB-D light field forward model from Equation (1) with σ being the WFT.

3.4 Ill-Posed Light Field Inversion

After having modeled the input light field in terms of the underlying wavefield, we now describe the inversion of the forward model $\sigma^{-1} : L(x, y, \theta_x, \theta_y) \rightarrow U_{WR}(x, y)$ that was shown in Equation (2). It can be seen from Equation (7) that the windowed Fourier transform of a wavefield results in a complex spectrum containing both amplitude and phase. The wavefield can be recovered by naively inverting the discretely sampled spectrum




Fig. 5. Evolution of the Complex Wavefield. Given an intensity source image at the origin (shown in the bottom right), the associated initial phase profile of the wave dictates the diffractive propagation and hence observed intensity image throughout the depth field. Here, we show examples of how different initial wave phase profiles affect the evolution of the wave intensity through a continuous depth volume. Our method aims at optimizing the phase pattern to produce the desired intensity image throughout a *continuous* depth volume.

from Equation (7) in the frequency domain as

$$U_{WR}(x, y) = \iint_{-\infty}^{\infty} w(x - x', y - y') \iint_{-\infty}^{\infty} L(x', y', \theta_x, \theta_y) \times e^{j(2\pi(f_x x' + f_y y') + \Phi(x, y, f_x, f_y))} df_x df_y dx' dy' \quad (9) \\ = \text{WFT}^{-1} \left\{ \sqrt{L(x, y, \theta_x, \theta_y)} e^{j\Psi(x, y, f_x, f_y)} \right\},$$

where $\Psi(x, y, f_x, f_y)$ is the associated phase and w is the window function as described in Equation (7). Note that this relation can also be obtained via Wigner distribution function [Zhang and Levoy 2009] and solved via stereogram approaches [Padmanaban et al. 2019].

This reveals a problem, as shown in Equation (8); a captured or rendered light field contains only amplitude information and does not describe the phase Ψ . This makes the inversion from

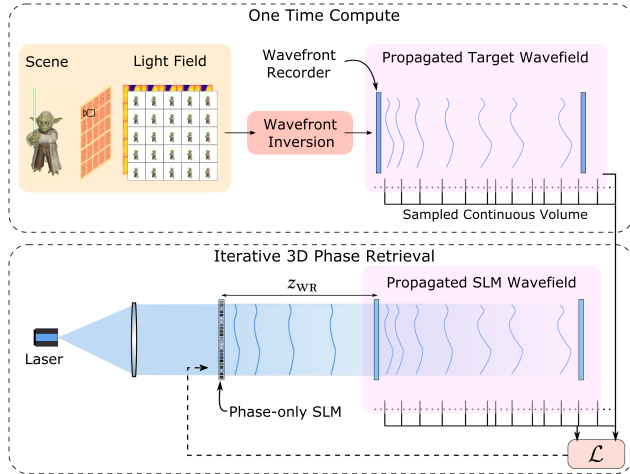


Fig. 6. Overview of Hogel-Free Holography. We start from the target 3D scene that we want to display and we convert that scene into a light field by computing the spectrogram at different viewpoints. We then convert this spectrogram into a wave representation without relying on hogels. This wave is propagated throughout a 3D volume to provide amplitude targets. The phase modulation of the SLM Φ is iteratively optimized so that the evolved wavefront produces the desired intensity at continuously sampled locations throughout the volume.

Equation (9) ill-posed. We alleviate this by defining the missing phase using the depth maps associated with the RGB-D light field array. Specifically, we evaluate the path delays along each individual angular light field direction. These path delays are obtained via the depth fields $D(x, y, \theta_x, \theta_y)$ associated with the light field intensities described by $L(x, y, \theta_x, \theta_y)$ [Yamaguchi et al. 1993]. We then set the missing phase information Ψ to

$$\frac{2\pi}{\lambda} D(x, y, \theta_x, \theta_y). \quad (10)$$

We plug the target phase from Equation (10) into Equation (9) so that we can perform the light field inversion described in Equation (2). And in doing so, we obtain the target wavefield U_{WR} . This resolves the ill-posed inversion.

3.5 Phase-Relaxed 3D Holography

With the forward and inverse of σ formulated above, the complex domain objective in Equation (2) is fully defined. However, the optimization objective in Equation (2) penalizes both amplitude and phase and formulating an appropriate penalty on phase is still a challenge. Recall that the phase of a wavefront at any given distance directly relates to the amplitude of that wavefront over a continuous 3D volume. For example, Figure 5 shows the amplitude of the propagated wave with different initial phases (shown in the inset). The wavefronts with a uniform initial phase on the aperture plane (first row) evolve differently from that of a linear (second row) or random phase (last two rows). This dependence of the wave propagation on the initial phase allows us to convert a phase objective into an amplitude objective. Specifically, we pose the objective as an amplitude optimization over the continuous 3D

ALGORITHM 1: Hogel-Free Holography Algorithm for Generating True 3D Phase-Only Holograms.

```

1: Inputs:
    $L$  // Light field
    $D$  // Depth field
    $N$  // Number of samples across 3D volume
    $T$  // Number of optimization iterations
2: Outputs:
    $\Phi_{OPT}$ 
3:  $\Phi = \text{Initial Phase}$ 
4:  $U_{WR} = \text{WFT}^{-1}(L, D)$ 
5: for  $t = 1$  to  $T$  do
6:    $z_1, \dots, z_N = \text{Random } N \text{ samples}$ 
7:    $\mathcal{E}(\Phi) = \sum_{i=1}^N \mathcal{L}(|\mathcal{P}(e^{j\Phi}, z_{WR} + z_i)|, |\mathcal{P}(U_{WR}, z_i)|)$ 
8:   Update( $\Phi, \partial \mathcal{E}(\Phi) / \partial \Phi$ )
9: end for
10:  $\Phi_{OPT} = \Phi$ 

```

volume as

$$\Phi_{OPT} = \min_{\Phi} \int_V \mathcal{L}(|\mathcal{P}(e^{j\Phi}, z_{WR} + z)|, |\mathcal{P}(U_{WR}, z)|) dz, \quad (11)$$

where \mathcal{L} is a custom penalty function, described further below, which we evaluate on the amplitudes over the continuous 3D volume V , Φ is the SLM phase, and U_{WR} is the complex wavefront measured at the wavefront recorder plane where the light field is measured as described in Section 3.4.

With the objective from Equation (11), we can optimize 3D holograms over a continuous volume, for the first time, to the best of our knowledge. Note that this approach does not explicitly calculate *hogels* corresponding to the input light field in the final optimized hologram, making it hogel-free. Also observe that our method can support any depth range.

Continuous Volume Optimization. Algorithm 1 reports our optimization method to solve Equation (11). See Figure 6 for an illustration of the algorithm. Our algorithm solves the objective via stochastic gradient descent. Specifically, an RGB-D light field is provided as input and the optimizer is initialized with a random phase (see line 3). Line 4 then applies the inverse of σ described in Section 3.4 to convert the light field into the underlying wavefront U_{WR} on the wave recorder. Line 6 randomly samples depths from a uniform distribution for every iteration within the continuous volume. We propagate the target complex field U_{WR} throughout a continuous depth volume using an angular spectrum wave propagation function $\mathcal{P}(U_{WR}, z)$. We then propagate the phase-only hologram field at the SLM $U_{SLM} = e^{j\Phi}$ (initialized with random phase) as $\mathcal{P}(U_{SLM}, z_{WR} + z)$ toward where the target amplitude from the complex wavefield U_{WR} is evaluated. As shown in line 7, the penalty is computed only on the amplitude of the wavefronts at these sampled depths. Then the error is backpropagated into the phase, as shown in line 8, where the SLM phase is updated until convergence.

Penalty Function. The custom penalty function \mathcal{L} can be freely defined as a function of observed image amplitude or intensity to facilitate optimization of perceptually appealing 3D holograms. To

that end, we use a weighted combination of ℓ_2 penalty \mathcal{L}_{ℓ_2} , SSIM $\mathcal{L}_{\text{SSIM}}$, perceptual penalty based on VGG-19 $\mathcal{L}_{\text{PERC}}$ [Johnson et al. 2016], and Watson FFT $\mathcal{L}_{\text{WFFT}}$ [Czolbe et al. 2020]:

$$\mathcal{L} = \lambda_{\ell_2} \mathcal{L}_{\ell_2} + \lambda_{\text{SSIM}} \mathcal{L}_{\text{SSIM}} + \lambda_{\text{PERC}} \mathcal{L}_{\text{PERC}} + \lambda_{\text{WFFT}} \mathcal{L}_{\text{WFFT}}. \quad (12)$$

We use a least-square penalty \mathcal{L}_{ℓ_2} for per-pixel accuracy in the reconstruction and $\mathcal{L}_{\text{SSIM}}$ as a handcrafted perceptual quality function. The perceptual penalty $\mathcal{L}_{\text{PERC}}$ compares the image features from activation layers in a pre-trained VGG-19 neural network, that is,

$$\mathcal{L}_{\text{PERC}} = \sum_l v_l \|\phi_l(x) - \phi_l(y)\|_1, \quad (13)$$

where ϕ_l is the output of the l -th layer of the pre-trained VGG-19 network and v_l are the corresponding penalty-balancing weights. Specifically, we use the outputs of ReLU activations just before the first two maxpool layers, i.e., `relu1_2` and `relu2_2`. Adopting $\mathcal{L}_{\text{PERC}}$ therefore helps achieve finer details in the reconstructed image. However, note that the VGG-19 network [Simonyan and Zisserman 2015] is optimized for classification and detection tasks, and is robust to the perceptual influence of artifacts such as noise. As a result, using a VGG-based perceptual penalty does not always guarantee perceptually high-quality results. Therefore, we further improve the reconstruction quality by adopting the Watson FFT error function which is crafted specifically for human visual system, based on Watson’s visual perception model [Watson 1993]. The combination of the above losses helps steer the optimization toward holograms that are perceptually pleasing for the human visual system. The perceptual losses mitigate perceptually apparent reconstruction noise across the 3D volume.

4 SETUP AND IMPLEMENTATION

We assess the 3D holograms generated by the proposed hogel-free holography method both in simulation and experimentally on a hardware display prototype. We discuss the specific software and hardware implementation details here.

4.1 Implementation

We test our hogel-free holography approach in simulation and implement the entire optimization framework within PyTorch running on a single Nvidia GeForce GTX 1080 GPU providing 8 GB of memory. PyTorch now provides complex gradients within its auto-differentiation modules making implementation of our optimization scheme with state-of-the-art first-order optimizers straightforward. We notice that different optimizers result in slightly different reconstruction quality. We optimize for $1,080 \times 1,920$ phase-only holograms, same as our SLM resolution, using the Adam optimizer running for 800 iterations with a learning rate of 0.1. This optimization currently takes about 300 s. Our 3D scene dataset consists of light field and depth field images rendered at 1,080p in Unity as well as those from Padmanaban et al. [2019] to accommodate the SLM maximum diffraction angle. For computing the WFT we use a 9×9 Hamming window to match the number of light field views. The depth range of our 3D scenes extends throughout a 13 mm volume; however, note that this depth range is not fundamental to the method.




Fig. 7. Hogel-Free Holography Experimental Prototype. Our benchtop prototype holographic display for validating true 3D holograms computed with our hogel-free method. For producing color holograms, we couple RGB lasers into a single-mode fiber which illuminates the reflective LCoS SLM displaying the phase patterns. Unwanted diffraction orders arising due to double-phase amplitude coding of previous methods are filtered using an iris. The modulated wave is then measured using a camera.

4.2 Experimental Prototype

To validate our simulations, we built an experimental benchtop prototype with a HOLOEYE LETO 1,080p phase-only LCoS SLM with a pixel pitch of $6.4 \mu\text{m}$ (Figure 7). We illuminated the SLM using a collimated beam from a single optical fiber that is coupled to three laser diodes emitting at nominal wavelengths of 630 nm, 520 nm, and 450 nm. The laser light was linearly polarized, matching the requirements of the SLM. The SLM was connected to a PC running on 64-bit Windows OS. However, the look-up-table adjustments to the SLM were made from a 32-bit Windows machine due to the age of our SLM. We use a telescopic optical relay system to image the holographic 3D projections with a Canon CMOS APS-C sensor. Additionally, we also placed an iris at an intermediate image plane to discard unwanted diffraction orders arising from the double phase encoding of OLAS [Padmanaban et al. 2019]. The RGB holographic images were captured in a color sequential manner with the camera settings unchanged but laser powers tuned to approximately white balance the illumination.

5 ANALYSIS

In this section, we perform quantitative analysis in simulation to validate Hogel-Free Holography for generating 3D holograms compared to existing methods. Holograms from all other methods are computed using the code provided by the authors.

5.1 Comparison with Existing Stereogram Approaches

We evaluate the proposed Hogel-Free Holography in simulation and we demonstrate a significant improvement over existing light field based state-of-the-art stereogram methods. Specifically, we compare against the holographic stereograms [Yatagai 1976], the phase-added stereograms [Yamaguchi et al. 1993], and the overlap add stereograms [Padmanaban et al. 2019]. We see in Figure 8 that our method closely matches the target 3D scene *despite modulating only the phase of a simulated laser beam*. Qualitative

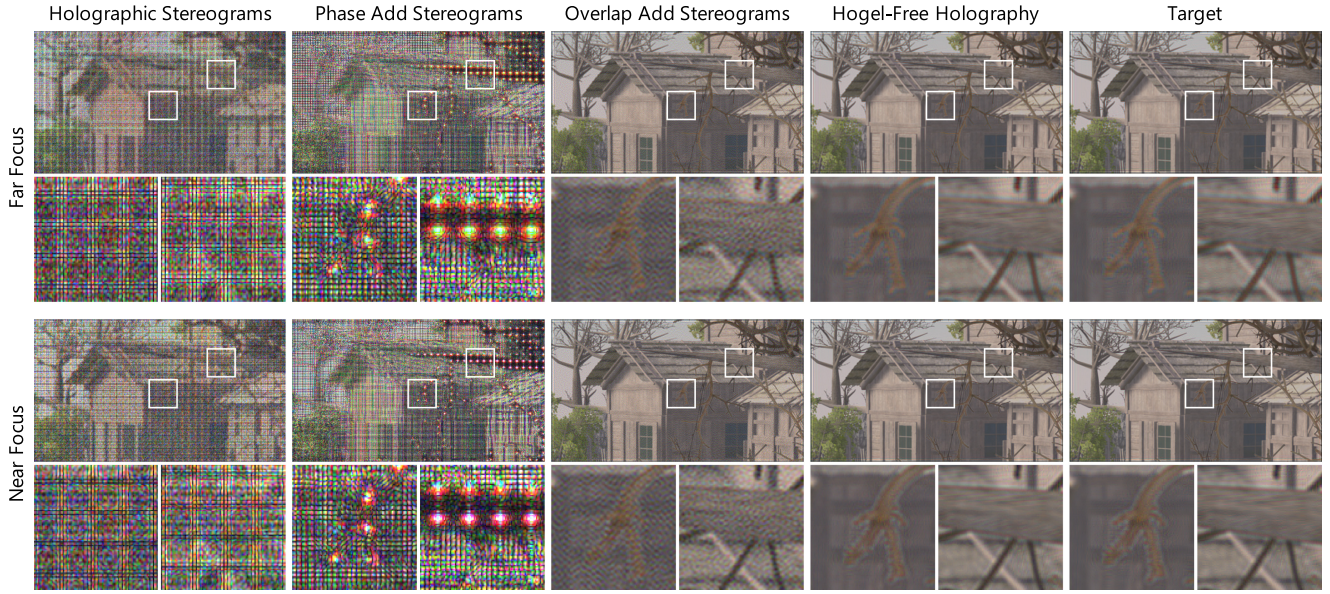


Fig. 8. Synthetic Evaluation. We compare our hogel-free holography to state-of-the-art holographic stereogram methods in simulated experiments. Holographic Stereograms [Kang et al. 2016] and Phase-Added Stereograms [Shi et al. 2017] both demonstrate low quality and visible tessellation artifacts caused by the spatio-angular tradeoff incurred by hogels. Overlap-Add Stereograms [Padmanaban et al. 2019] removes the tessellation artifacts but there is noticeable ringing caused by their usage of double-phase encoding. Our method demonstrates the highest quality at both near and far focus. The holograms generated by HFH strongly match the target 3D holograms which were generated with a theoretical perfect SLM that can modulate both amplitude and phase.

comparisons against baseline stereogram methods are shown in Figure 8 and we observe visible improvement at both near and far focus.

For quantitative evaluation, we compute the PSNR, SSIM, and LPIPS [Zhang et al. 2018] metric scores on the depth planes within each 3D scene. The average metric scores computed at depths sampled across the continuous volume are reported in Table 1. We observe over 10 dB improvement in PSNR, significantly outperforming the existing methods. Also, note that the LPIPS perceptual similarity metric shows significant improvement compared to previous methods, validating perceptually high-quality optimization of 3D holograms using our method.

5.2 Comparison with Occlusion-Aware Tensor Holography

Here, we compare our approach with the tensor holography method [Shi et al. 2021]. Tensor holography is a real-time neural network approach that implements a ray visibility-based occlusion-aware point-based method. Figure 10 shows qualitative comparisons of our method with holograms produced using tensor holography. As discussed in Section 3, ray visibility-based methods do not model diffraction due to hard occlusion edges and depth discontinuities in a physically accurate manner. As a result of this inaccurate occlusion modeling, light leaks from the background into the foreground severely contaminating the image with visible ringing artifacts, as can be seen in Figure 10. For example, observe the edges of the blades of grass where ringing artifacts can be clearly seen.

The physically inaccurate model in tensor holography also causes *incorrect defocus effects*, as shown in Figure 10. Since tensor holography generates holograms based on a single RGB-D image, the single perspective depth information is not always sufficient for computing accurate 3D holograms. As an example, we compare the 3D holographic reconstructions of a 3D scene visible from a non-Lambertian mirror surface in Figure 9. Note that the depth map of the mirror is rendered as a planar surface, without providing any detail of the 3D scene shown within the mirror. As a result, reconstructions produced by tensor holography do not exhibit accurate defocus effects. Also, ringing due to light leakage from the background can be seen near the candles. The proposed hogel-free holography produces accurate defocus effects in this challenging scene as the defocus and occlusion cues are implicitly encoded from the light field.

5.3 Robustness to Hardware Non-Idealities

Holographic displays unfortunately possess several deviations from the ideal light transport model [Chakravarthula et al. 2020; Peng et al. 2020]. This results in substantial disparity between simulated reconstructions and real holographic display output. We demonstrate that holograms generated by our HFH are more robust to these non-idealities than the previous state-of-the-art OLAS. Figure 11 shows an example where we simulate the effect of a non-linear SLM look-up table by introducing a gamma factor of 0.8 on the simulated phase. This phase non-linearity incurs a significant drop in quality for OLAS, both qualitatively and quantitatively, while our proposed HFH is much

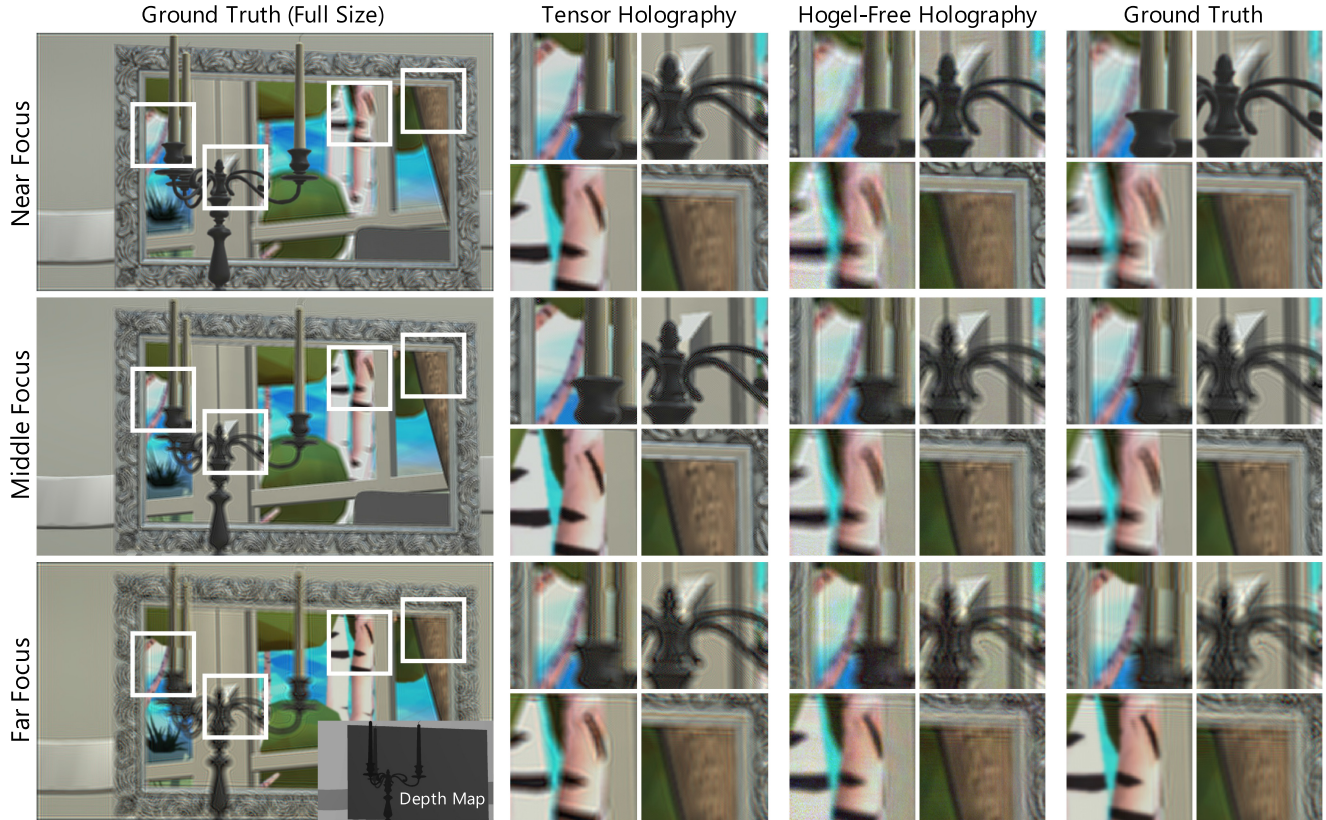


Fig. 9. Comparison with Tensor Holography. We compare our hogel-free holography to state-of-the-art Tensor holography [Shi et al. 2021] for computing 3D holograms. Tensor holography takes in a single RGB-D image as input and supports only a shallow depth range. As a result, for scenes like this which demonstrate a non-Lambertian surface reflection, the depth map does not contain the depth information of the scene within the mirror. This inaccurate depth map results in physically inaccurate defocus effects. On the other hand, our method utilizes an RGB-D light field which implicitly contains accurate representation of both view- and depth-dependent effects. As a result, our hogel-free holography eliminates visible ringing artifacts in tensor holography due to depth discontinuities at the occlusion edges and results in physically accurate parallax and defocus effects.

Table 1. Quantitative Results for Synthetic Evaluation of 3D Hologram Generation Methods

	PSNR (dB) \uparrow	SSIM \uparrow	LPIPS \downarrow
Hogel-Free Holography	40.2	0.964	0.140
Overlap Add Stereogram [2019]	30.0	0.810	0.417
Phase Added Stereogram [2017]	12.6	0.123	0.804
Holographic Stereogram [2016]	13.0	0.130	0.793

We evaluate the generated 3D hologram for each method across several planes within the depth field for each scene and we report the average quantitative score.

more tolerant of the deviation, validating the robustness of our method. Additional details can be found in the Supplementary Information.

6 ASSESSMENT

In this section, we validate the proposed hogel-free holography method by comparing with 3D holograms generated by existing methods and displayed on a real prototype holographic display.

6.1 Experimental Validation

Figures 12 and 13 show experimentally acquired results from the prototype display described in Section 4. The results in Figure 12 demonstrate that the proposed approach mitigates most of the severe artifacts present in existing light field-based stereogram methods. Specifically, spatio-angular resolution tradeoff due to the use of hogels in computing holographic stereograms limited the achieved spatial resolution of previous methods. While the overlap-add stereograms attempt to overcome this tradeoff, the usage of **double-phase amplitude coding (DPAC)** leads to visible artifacts and a loss of spatial resolution. Furthermore, DPAC leads to a checkerboard pattern on the hologram phase which results in the unwanted effect of a noticeable portion of light escaping the image window into higher diffraction orders. This creates multiple higher-order ghost images which leads to contamination of the holographic projections. Furthermore, OLAS produces ringing artifacts around high-contrast edges and in the periphery leading to image quality degradation. The proposed hogel-free holography suppresses severe ringing and reconstruction noise, achieving significantly higher spatial resolution, light efficiency, and contrast than all of the prior methods.




Fig. 10. Comparison with Tensor Holography. The state-of-the-art 3D Tensor holography [Shi et al. 2021] does not model physically accurate occlusion. As a result, the holographic projections show visible light leakage from the background into the foreground, and results in ringing artifacts at depth discontinuities and occlusion edges, such as at the blades of grass in the current scene. Furthermore, this physically inaccurate modeling results in incorrect defocus effects. Our method eliminates such artifacts by accurately modeling wave propagation at depth discontinuities at the occlusion edges and results in appropriate parallax and defocus effects.

The full color results as seen in Figures 12 and 13 confirm the improvements as seen in simulations. Our method reduces ringing and reconstruction noise present in existing methods and overcomes the fundamental spatio-angular resolution tradeoff. As a result, fine details such as the blades of grass are accurately captured in 3D holograms generated using HFH, along with their associated defocus effects as the user (camera) focus changes. On the other hand, it can be clearly seen that the previous methods which relied on single-step propagation and heuristic encoding such as double-phase amplitude encoding sacrifice spatial resolution. For example, observe that the cattails near the pond are very well reconstructed using hogel-free holography, whereas a significant spatial resolution drop can be observed with the other methods. While some artifacts do remain in our hardware captures due to imperfections of the SLM, uncalibrated phase look-up tables, and DC unmodulated light, these artifacts can be eliminated by tightly calibrating the hardware with recent hardware-in-the-loop optimization schemes [Chakravarthula et al. 2020; Peng et al. 2020]. We note that previous methods could not be remedied with hardware-in-the-loop refinement due to their restriction on both amplitude and phase, and non-iterative nature of hologram

generation; our hogel-free holography is the first to bridge this gap.

6.2 Depth of Field Effects

We validate the ability of our method to produce high-quality true 3D holograms by measuring and evaluating the defocus effects of holographic projections. To ensure tight focusing on the projected 3D scene, we first display a calibration spokes wheel image at different focus distances to adjust the camera focus to the desired distance. The true 3D hologram projection is then imaged with the calibrated camera focus, with the camera settings kept constant across all focus distances. Our method produces appropriate defocus effects as seen from Figure 12. Specifically, observe the defocus effects between the grass blades in the foreground and the white tree in the background for 3D scenes at the top. Similarly, the cattails at the pond go out of focus as the camera is changed from far focus to near focus. In Figure 13, we evaluate the defocus effects on a more challenging scene where the depth map is insufficient to describe the depth of the scene, and much of the 3D scene is visible due to the reflection from a mirror. Our method is able to produce appropriate defocus




Fig. 11. Robustness to SLM Phase Non-Linearities and Model Mismatches. A real-world SLM does not impute the exact phase modulation generated by CGH due to hardware imperfections. Nevertheless, the phases obtained with HFH are highly robust to errors caused by phase non-linearities and other model mismatches. In contrast, the performance of OLAS rapidly degrades as more hardware non-idealities are factored in. Here, we compare the performance under ideal conditions versus non-ideal conditions with a 0.8 gamma non-linearity on the SLM phase.

effects when the camera is focused on objects reflected from the mirror.

6.3 Parallax and Occlusion Effects

We study the parallax and occlusion effects produced by holograms optimized using hogel-free holography. To capture these view-dependent effects, an aperture is placed on a translating stage in the Fourier plane of our setup to filter the angular views. We use a camera viewpoint baseline of 4 mm with an identical sized aperture in the Fourier plane. As the eye-pupil is typically about 4 mm [Kahneman and Beatty 1966], this baseline supports intra-ocular occlusions. Figures 14 and 15 validate that our method produces true 3D holograms with accurate parallax and occlusion effects. In Figure 14, observe the white tree in the background of the leftmost scene which is visible in the left view but nearly disappears in the right view. Similarly, observe the blades of grass that appear in the left view but disappear in the right view. The proposed method produces a wavefield that continuously encodes the angular views of the light field. Therefore, Figure 15 shows similar parallax effects for the trees reflected in the mirror despite inaccurate cues from the depth map (also see Figure 9). Specifically, observe how more of the blue sky is revealed in one view than the other. We refer to the Supplementary Video for a dynamic illustration of the parallax effects.

Comparison to Tensor Holography. The tensor holography [Shi et al. 2021] neural network works only for a specific *HoloEye Pluto SLM* which has a pixel pitch of $8\mu\text{m}$, significantly different compared to our *HoloEye Leto SLM* with a $6.4\mu\text{m}$ pixel pitch. As a result, while we provided comparisons with tensor holography in simulation (see Section 5.2), we are unable to compare the method on a real holographic display. However, we note that our method significantly outperforms tensor holography as validated in simulations. The 3D volume supported by tensor holography is within a modest 6 mm depth range which is magnified by an eyepiece, whereas our continuous volume optimization method supports significantly larger depth ranges. Moreover, due to its dependence on a single RGB-D image, an insufficient depth map results in failure of producing true 3D holograms, such as the case described in Figure 9. Additionally, due to inaccurate and insufficient modeling of occlusions, tensor holography *does not* support parallax, and results in visible ringing artifacts and light leakage at depth discontinuities (see Figure 10).

7 DISCUSSION AND CONCLUSION

We introduced Hogel-Free Holography for computing high-quality 3D holograms that reproduce accurate defocus and parallax effects. To achieve this, we invert an input RGB-D light field into a complex wavefront and let it evolve by propagating the wavefront throughout a 3D volume. We then use the intensity distribution in this continuous 3D volume as the target, and optimize the SLM phase until the light modulated by the SLM produces the desired target intensity distribution throughout a continuous 3D volume. Through this technique, we have achieved high-quality 3D image quality with all depth and view-dependent effects incorporated. Specifically, we achieved physically accurate defocus and parallax effects, eliminated the background light leakage and ringing artifacts at depth discontinuities and hard occlusion edges that are commonly seen in existing methods, and we provided the first iterative optimization technique for 3D holograms that significantly improves over existing heuristic encoding methods. Furthermore, we also evaluated the proposed technique extensively in simulation and by using a real-world experimental prototype.

The proposed method qualitatively and quantitatively improves 3D holographic projections compared to existing methods. Although the experimental results exhibit residual aberrations, we note that these are not caused by the method but are instead due to an aged unsupported SLM which does not support custom look-up table calibration (see the Supplementary Information for further details). Although optimization has enabled high-quality results over single-step methods, our current implementation does not support real-time hologram generation. The proposed method may be accelerated using off-the-shelf technologies such as trained neural networks as demonstrated by recent methods. We envision that the proposed optimization approach and the results validated on real hardware will pave the way for future research in real-time 3D holography phase estimation methods that can take advantage of the first-order optimization capability that HFH provides.




Fig. 12. Experimental Validation of Our Proposed Method. We demonstrate high-quality 3D holograms using our experimental prototype which validate that the proposed method achieves observable improvements over the previous methods.

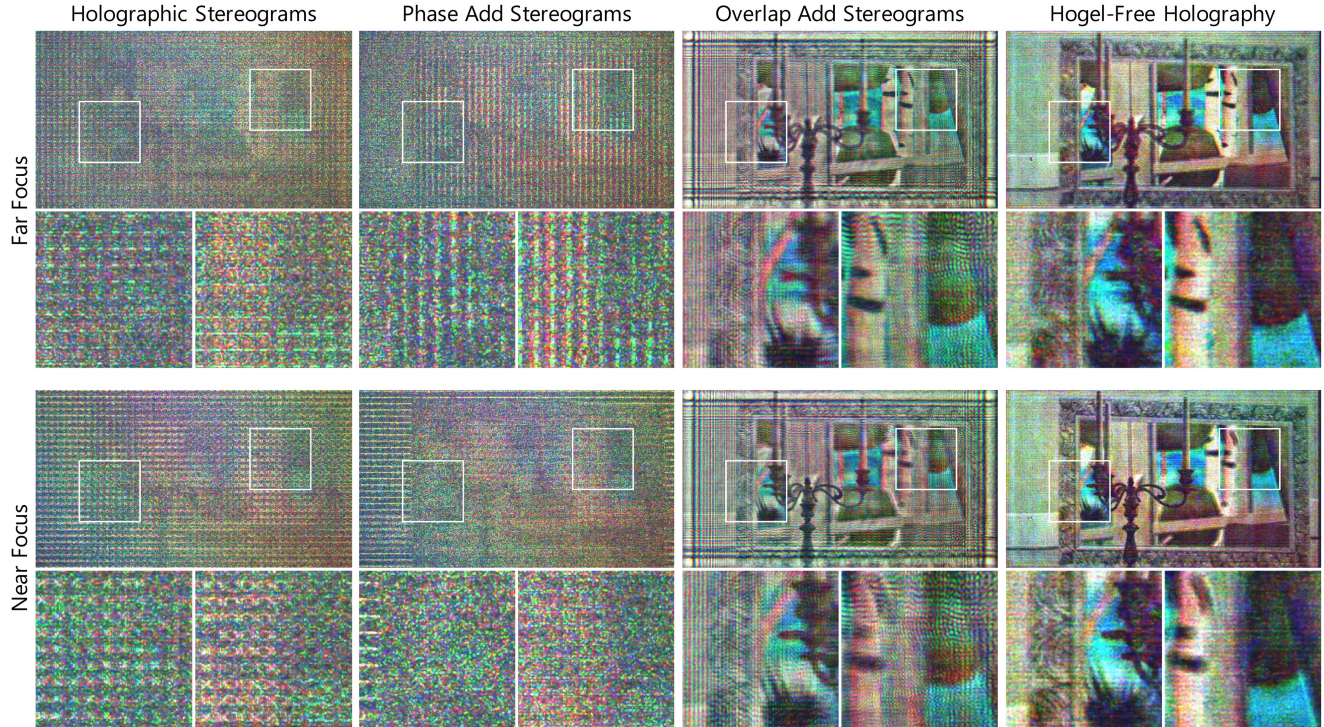


Fig. 13. Experimental Validation of Our Proposed Method. Here we demonstrate experimental results from a scene where the focusing effects are visible in the mirror despite it having a flat depthmap. This is possible as the said HFH method encodes the light field into the hologram wavefield. Therefore, although the correct depth map is lacking, our method is able to focus on objects within the mirror. Scene adapted from Padmanaban et al. [Padmanaban et al. 2019].




Fig. 14. Parallax in Hogel-Free Holograms. Our hogel-free holography method produces true 3D holograms with accurate parallax effects. For this forest scene, we observe a change in position of the background trees relative to the foreground grass (orange inset) and a change in the occlusion of the blades of grass (red inset). Please see the Supplementary Video for further visualization of the parallax effects.




Fig. 15. Parallax in Hogel-Free Holograms. Our hogel-free holography method produces true 3D holograms with accurate parallax effects. We used a camera viewpoint baseline of 4 mm. For this mirror scene, we observe a change in the occlusion of the blue sky. Specifically, the right view shows more sky for the orange inset, whereas the left view shows more sky for the red inset. Please see the Supplementary Video for further visualization of the parallax effects.

REFERENCES

- Stephen A. Benton and V. Michael Bove, Jr. 2008. *Holographic Imaging*. John Wiley & Sons.
- David Blinder and Peter Schelkens. 2018. Accelerated computer generated holography using sparse bases in the STFT domain. *Optics Express* 26, 2 (2018), 1461–1473.
- Stephen Boyd, Stephen P. Boyd, and Lieven Vandenbergh. 2004. *Convex Optimization*. Cambridge University Press.
- C. B. Burckhardt. 1970. A simplification of Lee's method of generating holograms by computer. *Applied Optics* 9, 8 (1970), 1949–1949.
- 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), 213.
- Praneeth Chakravarthula, Ethan Tseng, Tarun Srivastava, Henry Fuchs, and Felix Heide. 2020. Learned hardware-in-the-loop phase retrieval for holographic near-eye displays. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 186.
- Praneeth Chakravarthula, Zhan Zhang, Okan Tursun, Piotr Didyk, Qi Sun, and Henry Fuchs. 2021. Gaze-contingent retinal speckle suppression for perceptually-matched foveated holographic displays. *IEEE Transactions on Visualization and Computer Graphics* 27, 11 (2021), 4194–4203.
- 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.
- Steffen Czolbe, Oswin Krause, Ingemar Cox, and Christian Igel. 2020. A loss function for generative neural networks based on Watson's perceptual model. In *Advances in Neural Information Processing Systems*, Vol. 33. Curran Associates, Inc., 2051–2061.
- Nianchen Deng, Zhenyi He, Jiannan Ye, Praneeth Chakravarthula, Xubo Yang, and Qi Sun. 2021. Foveated neural radiance fields for real-time and egocentric virtual reality. arXiv:2103.16365. <https://arxiv.org/abs/2103.16365>.
- M. Hossein Eybposh, Nicholas W. Caira, Mathew Atisa, Praneeth Chakravarthula, and Nicolas C. Pégard. 2020. DeepCGH: 3D computer-generated holography using deep learning. *Optics Express* 28, 18 (2020), 26636–26650.
- 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.
- C. K. Hsueh and A. A. Sawchuk. 1978. Computer-generated double-phase holograms. *Applied Optics* 17, 24 (1978), 3874–3883.
- Changwon Jang, Kiseung Bang, Gang Li, and Byoungcho Lee. 2019. Holographic near-eye display with expanded eye-box. *ACM Transactions on Graphics (TOG)* 37, 6 (2019), 195.
- Changwon Jang, Kiseung Bang, Seokil Moon, Jonghyun Kim, Seungjae Lee, and Byoungcho Lee. 2017. Retinal 3D: Augmented reality near-eye display via pupil-tracked light field projection on retina. *ACM Transactions on Graphics (TOG)* 36, 6 (2017), 190.
- Justin Johnson, Alexandre Alahi, and Li Fei-Fei. 2016. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision (ECCV'16)*.
- Daniel Kahneman and Jackson Beatty. 1966. Pupil diameter and load on memory. *Science* 154, 3756 (1966), 1583–1585.
- Hoonjong Kang, Elena Stoykova, and Hiroshi Yoshikawa. 2016. Fast phase-added stereogram algorithm for generation of photorealistic 3D content. *Applied Optics* 55, 3 (2016), A135–A143.
- 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), 66.
- Douglas Lanman and David Luebke. 2013. Near-eye light field displays. *ACM Transactions on Graphics (TOG)* 32, 6 (2013), 220.
- Wai Hon Lee. 1970. Sampled Fourier transform hologram generated by computer. *Applied Optics* 9, 3 (1970), 639–643.
- Detlef Leseberg and Christian Frère. 1988. Computer-generated holograms of 3-D objects composed of tilted planar segments. *Applied Optics* 27, 14 (1988), 3020–3024.
- Gang Li, Dukho Lee, Youngmo Jeong, Jaebum Cho, and Byoungcho Lee. 2016. Holographic display for see-through augmented reality using mirror-lens holographic optical element. *Optics Letters* 41, 11 (2016), 2486–2489.
- Conny Lu, Praneeth Chakravarthula, Yujie Tao, Steven Chen, and Henry Fuchs. 2020. Improved vergence and accommodation via Purkinje Image tracking with multiple cameras for AR glasses. In *2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR'20)*. IEEE, 320–331.
- Mark Lucente and Tinsley A. Galyean. 1995. Rendering interactive holographic images. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques*. ACM, 387–394.
- Mark E. Lucente. 1993. Interactive computation of holograms using a look-up table. *Journal of Electronic Imaging* 2, 1 (1993), 28–34.
- Andrew Maimone, Andreas Georgiou, and Joel S. Kollin. 2017. Holographic near-eye displays for virtual and augmented reality. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 85.
- Michał Makowski, Maciej Sypek, Andrzej Kolodziejczyk, Grzegorz Mikula, and Jarosław Suszek. 2007. Iterative design of multiplane holograms: Experiments and applications. *Optical Engineering* 46, 4 (2007), 1–6.
- 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), 214.
- Jongchan Park, KyeoReh Lee, and YongKeun Park. 2019. Ultrathin wide-angle large-area digital 3D holographic display using a non-periodic photon sieve. *Nature Communications* 10, 1 (2019), 1–8.
- Jae-Hyeung Park and Mehdi Askari. 2019. Non-hogel-based computer generated hologram from light field using complex field recovery technique from Wigner distribution function. *Optics Express* 27, 3 (2019), 2562–2574.
- Yifan Peng, Suyeon Choi, Nitish Padmanaban, and Gordon Wetzstein. 2020. Neural holography with camera-in-the-loop training. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 185.
- Rafael Piestun, Joseph Shamir, Bernhard Weßkamp, and Olof Bryngdahl. 1997. On-axis computer-generated holograms for three-dimensional display. *Optics Letters* 22, 12 (June 1997), 922–924.
- Rafael Piestun, Boris Spektor, and Joseph Shamir. 1996. Wave fields in three dimensions: Analysis and synthesis. *Journal of the Optics Society of America* 13, 9 (Sept. 1996), 1837–1848.
- Liang Shi, Fu-Chung Huang, Ward Lopes, Wojciech Matusik, and David Luebke. 2017. Near-eye light field holographic rendering with spherical waves for wide field of view interactive 3D computer graphics. *ACM Transactions on Graphics (TOG)* 36, 6 (2017), 236.
- 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.
- Karen Simonyan and Andrew Zisserman. 2015. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*.

- Quinn Y. J. Smithwick, James Barabas, Daniel E. Smalley, and V. Michael Bove. 2010. Interactive holographic stereograms with accommodation cues. In *Practical Holography XXIV: Materials and Applications*, Vol. 7619. International Society for Optics and Photonics, 761903.
- James P. Waters. 1966. Holographic image synthesis utilizing theoretical methods. *Applied Physics Letters* 9, 11 (1966), 405–407.
- Andrew B. Watson. 1993. DCT quantization matrices visually optimized for individual images. In *Human Vision, Visual Processing, and Digital Display IV*, Vol. 1913. International Society for Optics and Photonics, 202–216.
- Xinxing Xia, Yunqing Guan, Andrei State, Praneeth Chakravarthula, Tat-Jen Cham, and Henry Fuchs. 2020. Towards eyeglass-style holographic near-eye displays with statically. In *2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR'20)*. IEEE, 312–319.
- Masahiro Yamaguchi, Hideshi Hoshino, Toshio Honda, and Nagaaki Ohyama. 1993. Phase-added stereogram: Calculation of hologram using computer graphics technique. In *Proceedings of the SPIE*, Vol. 1914. 25–31.
- Toyohiko Yatagai. 1976. Stereoscopic approach to 3-D display using computer-generated holograms. *Applied Optics* 15, 11 (1976), 2722–2729.
- Han-Ju Yeom, Hee-Jae Kim, Seong-Bok Kim, Huijun Zhang, BoNi Li, Yeong-Min Ji, Sang-Hoo Kim, and Jae-Hyeung Park. 2015. 3D holographic head mounted display using holographic optical elements with astigmatism aberration compensation. *Optics Express* 23, 25 (2015), 32025–32034.
- Hao Zhang, Yan Zhao, Liangcai Cao, and Guofan Jin. 2015. Fully computed holographic stereogram based algorithm for computer-generated holograms with accurate depth cues. *Optics Express* 23, 4 (2015), 3901–3913.
- Hao Zhang, Yan Zhao, Liangcai Cao, and Guofan Jin. 2016. Layered holographic stereogram based on inverse Fresnel diffraction. *Applied Optics* 55, 3 (2016), A154–A159.
- Jingzhao Zhang, Nicolas Pégard, Jingshan Zhong, Hillel Adesnik, and Laura Waller. 2017. 3D computer-generated holography by non-convex optimization. *Optica* 4, 10 (2017), 1306–1313.
- Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR'18)*.
- Zhengyun Zhang and Marc Levoy. 2009. Wigner distributions and how they relate to the light field. In *2009 IEEE International Conference on Computational Photography (ICCP'09)*. IEEE, 1–10.
- Yan Zhao, Liangcai Cao, Hao Zhang, Dezhao Kong, and Guofan Jin. 2015. Accurate calculation of computer-generated holograms using angular-spectrum layer-oriented method. *Optics Express* 23, 20 (2015), 25440–25449.
- Remo Ziegler, Simon Bucheli, Lukas Ahrenberg, Marcus Magnor, and Markus Gross. 2007. A bidirectional light field-hologram transform. In *Computer Graphics Forum*, Vol. 26. Wiley Online Library, 435–446.

Received 30 August 2021; revised 22 December 2021; accepted 31 January 2022