Artifact-Resilient Real-Time Holography

VICTOR CHU, Princeton University, USA OSCAR PUEYO-CIUTAD, Universidad de Zaragoza, I3A, Spain ETHAN TSENG, Princeton University, USA

FLORIAN SCHIFFERS, GRACE KUO, and NATHAN MATSUDA, Reality Labs Research, Meta, USA

ALBERT REDO-SANCHEZ, Universidad de Zaragoza, 13A, Spain

DOUGLAS LANMAN and OLIVER COSSAIRT, Reality Labs Research, Meta, USA

FELIX HEIDE, Princeton University, USA

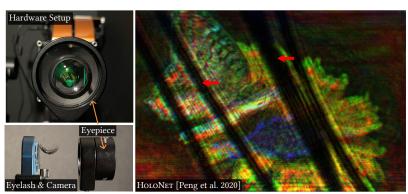




Fig. 1. Holographic Display with Eyelash Obstructions. Holographic near-eye displays show promising immersive visuals with rich depth cues and fine detail at a compact form factor. However, dynamic artifacts such as eyelashes and floaters inside the eyeball can be detrimental to the perceived image quality of holographic displays. We propose a new differentiable metric that can be used to distribute the eyebox for both iterative and neural methods to create artifact resilient holograms. By applying our metric, we are able to retrain off-the-shelf phase retrieval neural networks to generate high-quality artifact-resilient holograms in real time (85 fps). (Left) Holographic displays show images to a viewer by directing an engineered beam of light into the pupil of the viewer's eye. We simulate the effect of dynamic artifacts by placing phantoms for occluders (such as the eyelashes shown in the bottom left) in the display path. (Middle) Dynamic artifacts can obstruct the displayed image if the light beam exhibits a smooth phase profile, which is typical of several of today's state-of-the-art holographic displays [Choi et al. 2022; Peng et al. 2020]. The red markers indicate where interference with human eyelashes occurs. (Right) We find that our proposed method for real-time Artifact-Resilient Holography (ARH) is able to compute holograms that are robust to such distortions.

Holographic near-eye displays promise unparalleled depth cues, high-resolution imagery, and realistic three-dimensional parallax at a compact form factor, making them promising candidates for emerging augmented and virtual reality systems. However, existing holographic display methods often assume ideal viewing conditions and overlook real-world factors such as eye floaters and eyelashes—obstructions that can severely degrade perceived image quality.

In this work, we propose a new metric that quantifies hologram resilience to artifacts and apply it to computer generated holography (CGH) optimization. We call this Artifact Resilient Holography (ARH). We begin by introducing a simulation method that models the effects of pre- and postpupil obstructions on holographic displays. Our analysis reveals that eyebox

Authors' Contact Information: Victor Chu, Princeton University, Princeton, NJ, USA; Oscar Pueyo-Ciutad, Universidad de Zaragoza, I3A, Zaragoza, Spain; Ethan Tseng, Princeton University, Princeton, NJ, USA; Florian Schiffers; Grace Kuo; Nathan Matsuda, Reality Labs Research, Meta, Seattle, WA, USA; Albert Redo-Sanchez, Universidad de Zaragoza, I3A, Zaragoza, Spain; Douglas Lanman; Oliver Cossairt, Reality Labs Research, Meta, Seattle, WA, USA; Felix Heide, Princeton University, Princeton, NJ, USA.



This work is licensed under a Creative Commons Attribution 4.0 International License. @ 2025 Copyright held by the owner/author(s). ACM 1557-7368/2025/12-ART219 https://doi.org/10.1145/3763361

regions dominated by low frequencies—produced especially by the smooth-phase holograms broadly adopted in recent holography work—are vulnerable to visual degradation from dynamic obstructions such as floaters and eyelashes. In contrast, random phase holograms spread energy more uniformly across the eyebox spectrum, enabling them to diffract around obstructions without producing prominent artifacts.

By characterizing a random phase eyebox using the Rayleigh Distribution, we derive a differentiable metric in the eyebox domain. We then apply this metric to train a real-time neural network—based phase generator, enabling it to produce artifact-resilient 3D holograms that preserve visual fidelity across a range of practical viewing conditions—enhancing both robustness and user interactivity.

CCS Concepts: • Hardware \rightarrow Emerging optical and photonic technologies; Emerging interfaces.

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

1 Introduction

Realizing the transformative potential of AR/VR headsets depends heavily on the performance of their near-eye displays, which must excel in field of view, image quality, and depth perception [Itoh et al. 2021]. Holographic displays offer a compelling solution, combining a

compact form factor with a wide field of view, rich color gamut, high spatial resolution, and realistic depth cues to meet these demanding requirements and unlock new possibilities for augmented reality applications. Existing extended reality (XR) devices [Cakmakci and Rolland 2006] rely on OLED/LCD displays paired with compound optics, which inherently limits them in their ability to reproduce focus cues and parallax effects [Matsuda et al. 2017]. These limitations are central to the challenge of solving the vergence-accommodation conflict (VAC), a significant issue in today's commercial VR systems. Holographic displays are emerging as a promising solution for creating 3D visual experiences. Employing spatial light modulators (SLMs), which allow for precise modulation of light waves, holographic displays have the potential to produce 3D volumetric images in a compact form [An et al. 2020; Kim et al. 2022a; Maimone and Wang 2020; Xia et al. 2020].

Although artifacts from obstructions like dust particles or eyelashes are not detrimental when viewing the LCD displays found in current commercial VR headsets, these artifacts are amplified by the coherent nature of holographic light, potentially rendering the display useless (see Fig. 1). In the past five years, researchers have designed camera-in-the-loop systems (CITL) to effectively learn a model of the physical light propagation in order to optimize phase patterns that invert the aberrations that come from dust particles or lens imperfections and create high-quality 3D holograms with accurate parallax and defocus effects [Chakravarthula et al. 2020; Peng et al. 2020]. These camera-in-the-loop methods have been validated both in simulations and in physical hardware [Chakravarthula et al. 2020; Peng et al. 2020].

However, camera-in-the-loop techniques cannot handle dynamic obstructions, such as eyelashes and eye floaters (see Fig. 2), potentially severely impacting the viewing experience of these state-of-the-art holographic displays. When working in the near-field configuration, previously, smooth phase patterns have been preferred as they have maximal perceptual quality with minimal speckle and high contrast. While these smooth phase solutions combined with CITL can generate high quality volumetric holograms in controlled settings, smooth phase holograms have very small eyeboxes and do not make full use of the angular spread of light afforded by the SLM, making them incapable of diffracting around unanticipated temporally varying obstructions.

In this work, we define methods to generate holograms that make full use of the angular spread of light afforded by the SLM. Consequently, we generate holograms capable of being viewed through a wide range of scatterers from imperfect optics to dynamic obstructions, such as eyelashes. These artifact resilient holograms (ARH) not only improve the display's artifact resilience capabilities, but also achieve accurate defocus blur. To this end, we apply our understanding of how obstructions interact with holographic displays to retrain neural networks to generate real-time artifact-resilient holograms. Prototype experiments confirm that the proposed holographic display method is robust to these dynamic distortions.

Specifically, in this paper, we make the following contributions:

 We introduce a technique for simulating pre and post-pupil obstructions, namely eyelashes and vitreous floaters, to more accurately evaluate how a holographic display performs when

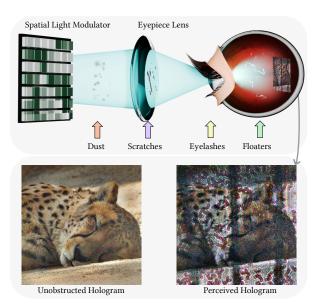


Fig. 2. **Holographic Artifacts Model**. To view a near-field holographic display (top), the pupil of an observer's eye takes in the propagated SLM field after the eyepiece lens. During free space propagation, the hologram is interfered by a variety of obstructions including, and not limited to, dust particles, scratches on lenses, eyelashes, and eye floaters. Artifacts from these obstructions interfere with image content reaching the retina, rendering the display useless (bottom).

perceived in practice. We verify our simulation method experimentally.

- Relying on the findings in simulation, we devise a differentiable metric to quantify holographic artifact resilience by mapping it to the underlying Rayleigh Distribution. We find these metrics can be applied to scenarios that have biases towards smooth phase holograms.
- Using our differentiable metric, we are able to train a real-time neural network to create pseudo-random phase holograms that are able to inherently diffract around a wide range of unknown, dynamic obstructions without the need for a camera in the loop. With our artifact resilient network, we are able to generate holograms at 85 FPS.

We hope this work not only highlights the limitations of using smooth phase holograms in practical holographic displays, but also demonstrates that many existing holographic methods can be readily adapted to produce artifact-resilient holograms in real time. The source code for this project is available on our GitHub repository¹.

2 Related Work

Computer-generated holography (CGH) employs algorithms to find complex-valued signals that will generate a target wavefront after the signal has traveled some distance. In its most general form, the amplitude and phase for each pixel location on a spatial light modulator need to be determined correctly so that the desired target can be produced. Hardware-specific constraints can narrow the

 $^{^{1}} https://github.com/princeton-computational-imaging/ArtifactResilientHolography$

domain of solutions so that only the phase component or only the amplitude component is variable. In this work, we compare and analyze the following two different categories of holograms, smooth and random phase holograms, that can produce similar wavefronts in ideal settings but have vastly different qualities when viewed by an observer.

2.1 Smooth Phase Holograms

We refer to the first category of holograms as "smooth phase". These holograms have the property that the phase component varies slowly over propagation distance, resulting in a holographic phase pattern that strongly resembles the target amplitude component. The smoothness of the phase component results in optical coherence within local patches of the hologram. This often results in better-quality holograms in ideal experimental settings because the coherence provides signal redundancy. However, in this work, we show that a drawback of the coherence is that these types of holograms are less robust to dynamic artifacts.

Recent works that use neural networks as the CGH engine for near field holography produce this category of holograms [Choi et al. 2022, 2021a; Horisaki et al. 2018; Li et al. 2025; Markley et al. 2023; Peng et al. 2020; Shi et al. 2021, 2022a,b; Wang et al. 2022]. The most popular architecture for neural network-based CGH relies on convolutional neural networks that map target amplitudes and phases to a hologram to be displayed on an SLM. This process bears similarity to adjacent work in the image translation and style transfer space [Zhu et al. 2017]. As such, the predicted holograms tend to have a smooth phase component that strongly resembles the input amplitude, resulting in a highly concentrated eyebox.

Smooth phase holograms can also occur even if neural networks are not used for the CGH step. Several works have proposed methods for performing phase retrieval using optimization and learned wave propagation models [Chakravarthula et al. 2019, 2022b, 2020; Choi et al. 2021b; Gopakumar et al. 2024; Peng et al. 2021]. Because the optimization is influenced only by the image fidelity defined in intensity space and there is no constraint on the hologram phase, these optimization procedures also tend to produce smooth phase holograms that maximize reconstruction scores.

Non-optimization CGH methods can also produce smooth phase holograms [Padmanaban et al. 2019]. The work of Maimone et al. [2017] encodes a low-frequency phase map into a complex-valued signal using double-phase amplitude coding. This technique requires that the input phase map has an upper bound on the frequency and thus works best with smooth phase holograms.

2.2 Random Phase Holograms

The other category of holograms are "random phase holograms". These holograms exhibit pseudo-random phase components and are they are more similar to diffuse surfaces that are common in everyday life [Goodman 2005]. This results in noisier reconstructions because of cross-talk between the wavefronts. However, the advantage of random phase holograms is that due to randomness in the phase component, the hologram is more uniformly distributed across many propagation directions, reducing the impact of an occlusion at any one point in the volume.

Recent work on accommodative holography utilized random phase holograms to generate natural defocus effects to drive the accommodation response of the human visual system [Kavakli et al. 2023; Kim et al. 2024, 2022b; Yoo et al. 2021]. Parallel work on pupilaware holography has also promoted the benefits of random phase holograms [Chakravarthula et al. 2022a; Schiffers et al. 2023].

Contrast decrease in random phase holograms has been investigated as well [Kabuli et al. 2025] and can be mitigated using a second, low-resolution modulator.

All these works argue in favor of random phase holograms since they produce a more uniform energy distribution within the eyebox, allowing for improved robustness to pupil motion.

Holographic displays that generate the image content at the far field typically utilize random phase holograms [Eybposh et al. 2020; Kuo et al. 2020; Tseng et al. 2024]. Monin et al. [2022] analyzed and compared random phase versus smooth phase holograms in the context of étendue expansion, finding that although smooth phase holograms produce higher quality reconstructions, they do not produce a usable eyebox.

Random phase holograms are typically noisy, as speckles at the pixel level occur inherently, reducing the visual image quality. Therefore, noise reduction techniques are required.

Temporal-multiplexing [Choi et al. 2022] is a robust method for reducing speckle that is well understood and will also be used in our analysis in this paper.

Other systems exploit redundancies in different domains to combat speckle, e.g., using diversity in either spatial [Kuo et al. 2023] or wavelength [Schiffers et al. 2025] space.

2.3 Artifacts in Holography

A recent trend in the holographic research community has been to calibrate out the hardware imperfections using feedback from the imaging camera [Chakravarthula et al. 2020; Choi et al. 2021a; Gopakumar et al. 2021, 2024; Peng et al. 2020] as "camera-in-theloop" (CITL) calibration techniques. The general theme of these approaches is to learn to reproduce the discrepancy between what the imaging camera sees and the result that one achieves when using an ideal simulation.

While some sources of artifacts are shared between random and smooth phase holography, for example, ripple effects caused by dust and scratches on the lenses, there are several key causes of image quality degradations that are not shared between the two types of holograms. Most concerning is that hardware-in-the-loop calibration methods have not been devised for dynamic artifacts (e.g., ocular obstructions such as eyelashes and floaters). Because these artifacts are dynamically moving during the viewing process, it is challenging to calibrate these artifacts using an imaging camera. As such, there is no existing CITL technique that can calibrate for dynamic artifacts to the best of our knowledge.

Holography in the presence of Obstructions

In this section, we first describe image formation for conventional near-eye holography without obstructions before detailing our model

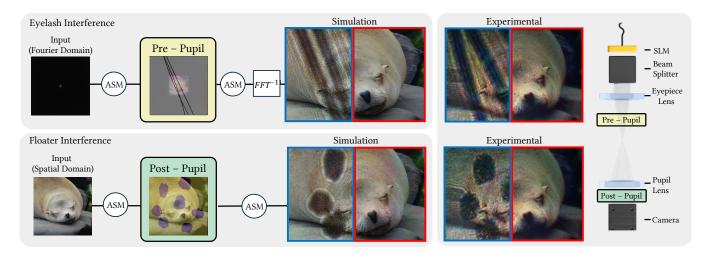


Fig. 3. **Verifying our Pre- and Post-Pupil Artifact Simulation Method.** Using our simulation method, we are able to accurately model the influence of obstructions that we place before and after the pupil lens in hardware. On the right, we show a schematic describing how we create both eyelash and floater phantoms. At the pre-pupil location, we place pieces of human hair to replicate the effect of an eyelash. Similarly, at the post-pupil location, we place a piece of glass with several dots of marker applied to emulate a complex of index of refraction within the eye (floaters). With the proposed simulation method, we are able to leverage ASM to propagate the wavefront to these same pre and post-pupil locations and apply complex obstruction masks. We show that our method accurately models artifacts caused by pre- and post-pupil obstructions for both smooth phase (blue boxes) and random phase holograms (red boxes).

that incorporates realistic occluders (see Fig. 2). We then experimentally validate our model and use it to analyze the distortions on holograms of varying eyebox size.

3.1 Vanilla Holography without Dynamic Obstructions

The angular spectrum method (ASM) [Goodman 2005] allows us to accurately model the light propagation for our holographic display in simulation with sampling constraints [Matsushima and Shimobaba 2009]. In our case, we use ASM as a function applied to the field modulated by the SLM that travels from the SLM towards the eye (see Fig. 2).

We express the propagated field $U(x, y, \lambda; z)$ at a distance z from input source field $U(x, y, \lambda; 0)$ using the ASM transfer function \mathcal{H} [Goodman 2005], as

$$U(x, y, \lambda; z) = \mathcal{F}^{-1}(\mathcal{F}(U(x, y, \lambda; 0)) \cdot \mathcal{H}(f_x, f_y, \lambda; z))$$

= ASM $(U(x, y, \lambda; 0), \mathcal{H}(f_x, f_y, \lambda; z))$. (1)

 $\mathcal F$ and $\mathcal F^{-1}$ are the Fourier transform and the inverse Fourier transforms respectively and both operate over their respective spatial or frequency domains. The transfer function $\mathcal H$ is defined here with the Fourier frequencies f_x and f_y in the x and y directions respectively and λ is the wavelength of light being propagated as

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

We define as shorthand the function ASM (U, z) to refer to propagation of the wavefront U by z distance.

3.2 Modeling Obstructions

Building upon vanilla holography, we define f_{Pre} for modeling obstructions before the pupil and f_{Post} for modeling obstructions after the pupil on any generic wavefront U. In short, we can model any interfered wavefront U' as

$$U' = \left(f_{\text{Post}}^M \circ \mathcal{F}^{-1} \circ f_{\text{Pre}}^N \circ \mathcal{F} \right) (U), \tag{3}$$

where $f_{\operatorname{Pre}}^N = f_{\operatorname{Pre},1}(M_1, z_1) \circ \cdots \circ f_{\operatorname{Pre},N}(M_N, z_N)$ and $f_{\operatorname{Post}}^M = f_{\operatorname{Post},1}(M_1, z_1) \circ \cdots \circ f_{\operatorname{Post},M}(M_M, z_M)$ and M_i is a complex obstruction mask located at distance z_i for a total of N pre-pupil occluders and M post-pupil occluders.

3.2.1 Modeling Pre-Pupil Obstructions. Obstructions that occur between the eyepiece lens and the pupil lens (e.g. eyelashes) can be detrimental to the viewing experience of holographic displays. To model this correctly, we first take the field at the eyebox plane (the field that the pupil takes in) then propagate it back towards the eyepiece lens.

When passing the propagated SLM field through the eyepiece with focal length f, we are left with a field of different dimensions from the SLM. Explicitly, given a wavelength λ and SLM pixel pitch $p_{\rm SLM}$, the extent of the first order in the Fourier domain L_f is [Goodman 2005]

$$L_f = \frac{\lambda f}{p_{\text{SLM}}}. (4)$$

We assume the field at the imaging plane is U(x,y;z) and define E(x,y) to be the eyebox plane. As we propagate the eyebox field back to the eyepiece plane, one would expect to see the image grow back to the full SLM extent at the eyepiece plane corresponding to the target image. To simulate this effect, prior to propagating

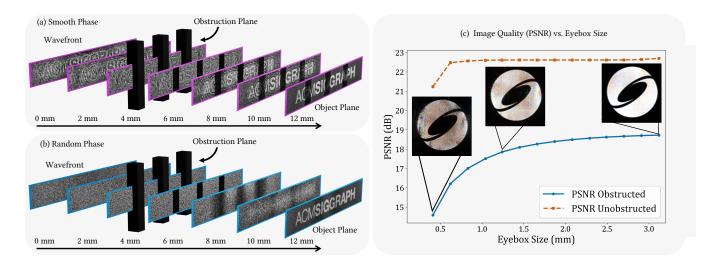


Fig. 4. Large Eyeboxes Create Stronger Artifact Resilience. We create optimized smooth and random phase holograms that form the SIGGRAPH logo at the object plane. We then show how both holograms interact with three rectangular opaque obstructions that are located 8 mm away from the imaging plane. (a) Over short propagation distances, smooth phase holograms scatter light to a very narrow range of angles, resulting in the image being easily compromised by a variety of obstructions. (b) In contrast, image content in random phase holograms form close to the imaging plane, increasing its robustness to artifacts. (c) To further analyze how different eyebox sizes affect artifact resilience, we optimize holograms over a large range of étendue configurations to keep the field of view constant at 78° and we sweep the effective eyebox size from 0.42 mm to 3.13 mm. Leveraging our proposed artifact simulation method, we are then able to model the effect of dynamic obstructions on these holograms. Above we plot the associated average PSNR over 6 holograms for each étendue configuration both with and without obstructions. Please refer to the supplemental materials for experimental validation of this trend.

the eyebox field, we pad the eyebox to the SLM area by using the padding function P that takes in SLM sizes l_x and l_y in x and y dimensions, respectively, that is

$$E(f_x, f_y) = P(\mathcal{F}\{U(x, y; z)\}, l_x, l_y). \tag{5}$$

Once padded, we propagate the padded eyebox via ASM to the pre-pupil plane (we assume this to be a singular plane for simplicity) and apply a mask M_{Pre} at this plane before propagating it back to the eyebox plane.

$$f_{\text{Pre}}(E, M_{\text{Pre}}, z_{\text{Pre}}) = \text{ASM}(M_{\text{Pre}} \text{ ASM}(E, -z_{\text{Pre}}), z_{\text{Pre}}).$$
 (6)

Parameters for modeling eyelashes accurately can be found in the supplemental materials.

3.2.2 Modeling Post-Pupil Obstructions. Obstructions that are located after the pupil lens and before the retina, particularly floaters (small particles within the vitreous humor of the eye) and dust, are significant contributors to visual degradation as noted by clinical studies [Hammer et al. 2024; Kennelly et al. 2015; Paniagua-Diaz et al. 2024]. To model the affect of these obstructions on a holographic display, we can generate a complex mask M_{Post} that can be applied throughout the interior of the eye to simulate collagen floaters, or other particles, being distributed within the vitreous humor. To apply the mask, we take the field at the retina U and use ASM to propagate it to some intermediate distance z_{Post} , apply the mask, and then propagate the masked field back to the retina via

$$f_{\text{Post}}(U, M_{\text{Post}}, z_{\text{Post}}) = \text{ASM}(M_{\text{Post}} \text{ ASM}(U, -z_{\text{Post}}), z_{\text{Post}}).$$
 (7)

Parameters for modeling floaters accurately can be found in the supplemental materials.

While we focus on inter eve obstructions in this section, as this operates on the spatial domain, our post-pupil model can be generalized to all obstructions before the eyepiece and after the pupil lens, including dust and lens scratches as seen in Fig. 2, by setting $z_{\rm Post}$ accordingly.

3.3 Experimentally Emulating Pre and Post-Pupil Artifacts

Since we are unable to electronically read out the response of the human eye, we choose to construct a hardware analogue shown in Fig. 3 that models the same effects to validate our simulation model. To this end, we use the lenses in the 4F relay system to act as the eyepiece lens and the human viewer's eye lens. The bare sensor acts as the retina of the pupil.

We then construct imaging phantoms that emulate eyelashes and floaters to act as our pre and post-pupil obstructions. The eyelash phantom is constructed by attaching pieces of human hair to a lens mount, chosen for its similarity to real eyelashes. For the floater phantom, we applied viscous liquids such as marker ink, which exhibit both phase and amplitude responses, onto a microscope slide. Photographs of both phantoms are included in the Supplemental Materials.

The eyelash phantom is mounted after the eyepiece lens and before the Fourier space, emulating how eyelashes interact with the wavefront before entering the viewer's eye lens. Similarly, the floaters phantoms are mounted after the system's pupil lens to emulate how the floaters reside within the human eyeball. Since floaters can exist anywhere throughout the human eyeball the floaters phantom can be placed anywhere after the second lens of the second 4F relay system. We choose to place the floaters close to the imaging sensor. In Fig. 3, we also show simulated and experimental captures of both random phase and smooth phase holograms in the presence of our experimental phantoms. Our simulations closely match what we see in our real captures. Please refer to the Supplemental Information for more details on the hardware construction.

3.4 Examining Artifact Resilience with our Artifact Model

Using our experimentally validated artifact model, we examine in Fig. 4 the impact of eyebox size on artifact resilience. Conceptually, in order for a hologram to resist an obstruction, scattering is needed for the hologram to "bypass" the obstruction. As the eyebox size bounds the maximal angular spread of light within the system, a certain eyebox requirement is needed regardless of the CGH method.

Our analysis indicates that larger eyeboxes promotes artifact resilience. We find that an eyebox of approximately 3 mm is necessary for achieving high robustness against obstructions. Additional experimental results of a variety of pre- and post-pupil artifacts on varying eyebox sizes can be found in the supplemental materials.

4 Quantifying Phase Randomness

While several recent methods have also advocated for random phase holography [Kuo et al. 2023; Schiffers et al. 2023, 2025], none have formally defined or quantified just how random the phase of a hologram needs to be in order to reap the benefits of random phase. The lack of a quantitative measure also makes it difficult to compare the degree of randomness of holograms generated by different random phase CGH methods. For example, how does SGD with random phase initialization [Kuo et al. 2023] compare to SGD with stochastic pupil sampling [Schiffers et al. 2023]?

4.1 Eyebox Distribution Matching

We propose a deterministic metric for quantifying the degree of phase randomness of a hologram. Our metric is defined such that holograms can be ordered on a scale from completely smooth (constant phase) to completely random. In this section, we motivate the metric by drawing connections to known properties of complex stochastic processes. Akin to a tunable knob, the amount of smoothness and the amount of randomness can be adjusted to prioritize either image quality or artifact resilience, and we show how the metric facilitates this in developing real time artifact resilient CGH methods in the next section. We choose to define this metric for phase-only SLMs, but the derivation of the metric can be generalized to complex SLMs that modulate both amplitude and phase.

From our analysis in the previous section we found that large eyeboxes that have energy distributed evenly throughout the eyebox provide the best resilience against dynamic and static artifacts. We can measure the difference of a given eyebox E is from a flat energy distribution for all Fourier frequencies as a first-order approximation for quantifying phase randomness. Parseval's theorem states that the Fourier transform preserves energy. Therefore, we can relate

the energy at the SLM plane to the energy at the eyebox plane as

$$\sum_{x,y} |U(x,y;0)|^2 = \sum_{f_x,f_y} |E(f_x,f_y)|^2.$$
 (8)

Since we are using phase-only SLMs we know that $|U(x,y;0)| = 1 \ \forall x,y$. Thus, we can define a uniform distance measure \mathcal{L}_{UD} for measuring the distance between an eyebox's energy from a flat energy spectrum as

$$\mathcal{L}_{\text{UD}}(E) = \sum_{f_x, f_y} (|E(f_x, f_y)|^2 - 1)^2.$$
 (9)

Although the metric \mathcal{L}_{UD} is necessary for phase randomness, unfortunately we found that it is not a sufficient for optimization. We show in Fig. 5 that there exists many holograms that exhibit smooth phase while still being able to minimize \mathcal{L}_{UD} . We observed that the error incurred by a single concentrated energy outlier within the eyebox, which is a property shared by smooth phase holograms, can be offset by the other Fourier frequencies that exhibit low energy.

We thus modify the first-order approximation metric by incorporating second-order terms. Specifically, we do not only want to quantify the mean of the distribution of energy within the eyebox of a random phase hologram but also the variance. In order to determine the variance of this distribution we perform a statistical analysis of complex stochastic processes.

Suppose that x[n] is a random vector of length n where each entry has a random phase and unitary magnitude, that is,

$$|x[i]| = 1$$
 and $\angle x[i] \sim \text{Uniform}(0, 2\pi) \ \forall i \in [n].$

Then we have

$$\mu_{\text{Re}(x)} = \mu_{\text{Im}(x)} = 0 \text{ and } \sigma_{\text{Re}(x)}^2 = \sigma_{\text{Im}(x)}^2$$

where $\mu_{(\cdot)}$, $\sigma_{(\cdot)}$ are the statistical mean and standard deviation operators, and $\text{Re}(\cdot)$, $\text{Im}(\cdot)$ are the real and imaginary component operators. It is known that complex stochastic processes with zero mean and equal real and imaginary variances have certain properties [Richards 2013]. In particular, if $X = \mathcal{F}(x)$ then the distributions of Re(X) and Im(X) are the same and follow a Gaussian distribution with the probability distribution function

$$\mathcal{P}(\operatorname{Re}(X)) = \frac{1}{\sqrt{2\pi\sigma_{\operatorname{Re}(X)}^2}} \exp\left(-\frac{\operatorname{Re}(X)^2}{2\sigma_{\operatorname{Re}(X)}^2}\right). \tag{10}$$

Furthermore, if the real and imaginary components are identically distributed Gaussian distributions then the magnitude component |X| follows a scaled Rayleigh distribution with a probability distribution function that is given by

$$\mathcal{P}(|X|) = \begin{cases} \frac{|X|}{\sigma_{\text{Re}(x)}^2} \exp\left(-\frac{|X|^2}{2\sigma_{\text{Re}(x)}^2}\right) & : 0 \le |X|, \\ 0 & : \text{otherwise.} \end{cases}$$
(11)

and the mean and variance of this distribution are given by

$$\mu_{|X|} = \sigma_{\mathrm{Re}(x)} \sqrt{\frac{\pi}{2}} \text{ and } \sigma_{|X|}^2 = \frac{4-\pi}{2} \sigma_{\mathrm{Re}(x)}^2.$$

We now have a closed form solution for the mean and variance of the intensity of the eyebox of a random hologram. With this variance in hand, we can now assess how the means and variances

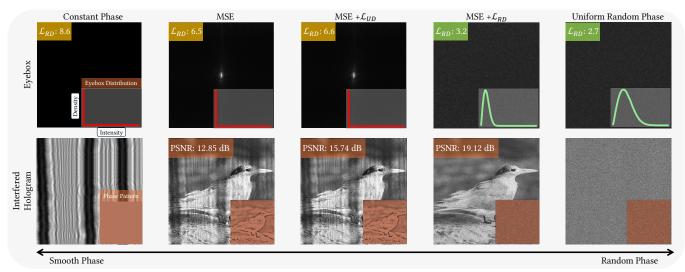


Fig. 5. Smooth to Random Holograms Continuum. From left to right we show a constant phase hologram, a hologram optimized using constant phase initialization and MSE loss, a hologram optimized using constant phase and MSE loss and \mathcal{L}_{UD} , a hologram optimized using constant phase initialization and MSE loss and \mathcal{L}_{RD} , and an unoptimized hologram where the phase values are uniformly randomly initialized. We observe that our metric \mathcal{L}_{RD} can guide the optimization towards solutions that exhibit greater phase randomness and wider energy distributions within the eyebox, even if the initialization was a constant phase hologram. Furthermore, we observe that total randomness is not necessary for artifact resilience and that partial randomness can also provide strong protection against artifacts, as shown in the second column from the right.

of the eyeboxes of our generated holograms compare to the target mean and variance.

To compute the distance of a given hologram's eyebox *E* from the target Rayleigh distribution we need to compute the variance of E. Instead of computing the variance over the entire eyebox, we instead choose to compute local patchwise sample variances in a manner akin to structural similarity [Wang et al. 2004]. This computation is performed by first defining a Gaussian blur kernel K with standard deviation σ_K . The local means and local variances of E can then be calculated as

$$\mu'_{F} = E \circledast K \text{ and } \sigma'_{F}^{2} = E^{2} \circledast K - {\mu'}_{F}^{2}$$

where ® denotes convolution. We can now define our Rayleigh Distance, \mathcal{L}_{RD} , which quantifies the distance between the eyebox E and the target Rayleigh distribution as

$$\mathcal{L}_{RD}(E) = \sum_{f_x, f_y} \frac{(2\mu_E'(f_x, f_y)\mu_R + c_1)c_2}{(\mu_E'^2(f_x, f_y) + \mu_R^2 + c_1)(\sigma_E'^2(f_x, f_y) + \sigma_R^2 + c_2)},$$
(12)

where μ_R and σ_R are the mean and variance of the target Rayleigh distribution, which can be derived from Eq. (11). The constants, c_1 and c_2 , are to prevent division by zero. We do not use a covariance term since the reference Rayleigh distribution is fixed and the eyebox is dependent on the hologram.

While higher order measures of distribution differences can be incorporated such as skewness and kurtosis, we have found that the second order suffices for our purposes. Furthermore, the connection with the Rayleigh distribution only relies on the SLM entries being independently randomly drawn from a complex random distribution that has equal real and imaginary variances. The connection holds even if the mean is non-zero because a non-zero mean can

be treated as a constant offset to the distribution. Thus, our metric. which we call the Rayleigh Distance (RD), for phase-only SLMs can be extended to complex SLMs that modulate both phase and amplitude by deriving the corresponding real and imaginary variances. We validate the efficacy of the Rayleigh Distance in the following sections.

Validation of Metric

We validate the efficacy of our metric by demonstrating that it can be used to optimized for artifact-resilient holograms. Conventional methods for obtaining holograms with random phase set the initial phase values using random uniform initialization [Kuo et al. 2023]. Due to the presence of local minima around the initial starting point the initialization alone is capable of facilitating random phase holography. However, for the same reason, if the initial starting point is a constant smooth phase then conventional optimization will struggle to find random phase solutions.

We perform an illustrative experiment in Fig. 5, optimizing for a continuum of holograms spanning smooth to random phases. We launch three optimizations where all of the optimizations use constant phase as the initial iterate, and the loss functions \mathcal{L}_{MSE} , $\mathcal{L}_{MSE} + \mathcal{L}_{UD}$, and $\mathcal{L}_{MSE} + \mathcal{L}_{RD}$. We observe that the optimization with the proposed loss metric is the only optimization that is capable of pushing smooth phase initializations towards random phase solutions. Since the Rayleigh Distance is agnostic to eyepiece focal length, it encourages even light distribution across the viewing zone, enabling practical use across display configurations with sufficient eyebox size (see Fig. 4). Furthermore, we also find that the randomness of the solution when using our proposed metric is not as random as a hologram that exhibits true uniform random phase.

These findings suggests that there is a range of degrees of phase randomness that can facilitate artifact resilience. The proposed metric is the first metric, to the best of our knowledge, that facilitates the investigation of the threshold at which holograms become artifact resilient.

While we could now use an iterative method such as the one from above and optimize for phase patterns with the proposed metric, this unfortunately does not allow for real-time throughput. Moreover, in the case of training neural networks for real time CGH [Peng et al. 2020; Shi et al. 2021], it is not clear how to initialize the optimization problem, specifically the weights of the neural network, such that random phase holograms are the starting point. Indeed, all existing real time neural network based CGH methods for near field holography, such as HOLONET and TENSOR HOLOGRAPHY, produce smooth phase holograms. In the next section we will demonstrate that the metric can be used as a training loss to bias the outputs of a real time CGH neural network towards random phase.

5 Real-Time, Artifact Resilient-Holography

Real-time, high-quality random phase CGH is still an unsolved problem, as existing real-time methods currently only produce smooth-phase holograms, exhibiting high-quality in-focus imagery, but lacking the desirable properties of random phase holograms such as artifact resiliency or natural defocus blur. On the other hand, there is no method that generates high-quality, random phase holograms in real time. In this section, we apply our Rayleigh distance metric to bridge this gap.

Previous methods for real time holography trained deep neural networks to generate the phase pattern given a target image. As these neural methods are only trained on the likelihood loss (i.e. mean squared error or perceptual), they learn to generate smooth-phase holograms since the image features are directly translated into the phase pattern, converging into a local minimum. By incorporating our metric for randomness into the loss function, we are able to successfully train an off-the-shelf neural network, HoloNet [Peng et al. 2020], to generate artifact resilient holograms. Our artifact resilient holograms maintain many of the desirable properties of random phase holograms, such as artifact resilience and natural defocus blur.

First, we show how to include our phase randomness metric that was described in Section 4 to train our Artifact Resilient Holography Network (Real-Time ARH) along with the associated impact on generating holograms that are robust to artifacts. Then, we explain how to enhance inference with time multiplexing through model ensembling. Finally, we show our Real-Time ARH extension to generate holograms with realistic 3D effects from RGB-D inputs, without any retraining or fine-tuning, showing high image quality, artifact resilience and natural defocus blur, while operating in real time.

5.1 Artifact Resilient Network

In this section, we show that our metric can be used to train neural networks for real time, artifact resilient CGH. We choose to demonstrate this by using an off-the-shelf neural network method, but note that the metric can be applied to any CGH optimization. HoloNet,

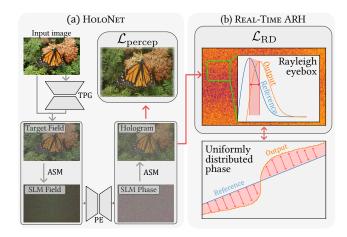


Fig. 6. Learning to generate random phase holograms. We train an off-the-shelf neural network for CGH [Peng et al. 2020] combining a perceptual loss function with our proposed artifact resilience regularizer. (a) Holonet overview. The Target-Phase Generator (TPG) produces a phase pattern at the target plane from the input image. Then, the wavefront at image plane, composed of the image and the generated target phase, is propagated to the SLM. Then, the Phase Encoder (PE) converts the complex wavefront into a phase-only hologram. Finally, the model is trained through a perceptual loss on the propagated phase pattern. (b) Real-Time Artifact-Resilient Holography. We incorporate our random phase metric into the training as a regularizer. As described in Section 4, \mathcal{L}_{RD} ensures that the eyebox possesses a spread out distribution of energy that is correlated with a uniform distribution of phase at the SLM plane.

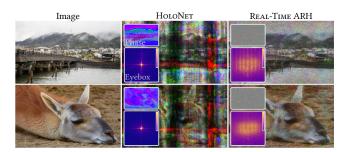


Fig. 7. **HOLONET versus our Real-Time ARH**. We compare the visual quality of holograms produced by Holonet and our Real-Time ARH, without time multiplexing, in the presence of artifacts with their eyeboxes (clipped for visualization). The vanilla Holonet produces smooth phase holograms with small eyebox, which are distorted by the artifacts. In contrast, our Real-Time ARH generates random phase holograms with large eyeboxes, leading to artifact resilient holograms.

shown in Fig. 6a, is trained only with the perceptual loss $\mathcal{L}_{\text{percep}}$, which leads to smooth phase holograms. It generates holograms by combining two networks: the Target-Phase Generator and the Phase Encoder, combined through a physical propagator (ASM in our case). The Target-Phase Generator creates a phase pattern at the target plane, which is combined with the image and propagated to the SLM plane. Then, the Phase Encoder converts the complex

Table 1. Quantitative Quality and Runtime Evaluation of CGH Methods in the Presence of Dynamic Obstructions. Our dynamic obstruction simulation framework and our proposed metric are used to assess the quality of different CGH methods: Vanilla Random Phase Optimization (VRO), DPAC, Tensor Holography, Vanilla HoloNet, and our Real-Time ARH. For each metric, the arrow indicates whether higher is better or lower is better. We find that our REAL-TIME ARH achieves a high performance in the presence of dynamic obstructions with a fast inference time. FPS of each CGH method was evaluated on an NVIDIA L40. All holograms are 1280 × 800 with a pixel pitch of 10.8 um. For Tensor Holography [2021] we use their reported FPS. For our REAL-TIME ARH we use 8 frame time multiplexing and report the FPS of one model. We decouple latency from the number of models involved since each model generates an independent frame, hence they can be run in parallel. The supplemental materials provide further details on time-multiplexing and include a selection of phase patterns generated with these real-time CGH methods.

	PSNR (dB) ↑	PSNR (dB) ↑	$\mathcal{L}_{ ext{RD}}\downarrow$	FPS ↑
	w/o obstructions	w/ obstructions		
VRO (10 iter) [2023]	21.6	18.1	2.92	3
VRO (300 iter) [2023]	69.0	20.1	2.81	0.1
DPAC [2017]	17.0	12.5	8.15	370
Tensor Holography [2021]	18.1	11.2	8.30	>60
Vanilla HoloNet [2020]	18.9	14.8	8.26	85
Ours: Real-Time ARH	17.7	17.1	3.45	85

wavefront into a phase-only hologram that produces the image when displayed in the SLM.

To allow the network to learn how to generate robust holograms, we combine the likelihood, perceptual loss with our Rayleigh distance metric from Section 4 as shown in Fig. 6b, acting as a regularizer. Our loss function \mathcal{L} is

$$\mathcal{L} = \mathcal{L}_{percep} (X, |ASM (U, z)|) + \rho \mathcal{L}_{RD} (\mathcal{F} (ASM (U, z))), \quad (13)$$

where X is the target image at distance z, U is the wavefront at the SLM plane, z is the propagation distance from the SLM to the target plane, and \mathcal{L}_{RD} is our Rayleigh Distance regularizer for randomness and eyebox dispersion as defined in Eq. 12.

First, we train the network one epoch with the perceptual VGG loss [Johnson et al. 2016] as \mathcal{L}_{percep} ($\rho = 0$), then continue with our \mathcal{L}_{RD} regularizer ($\rho = 1$) for random phase convergence. We perform this initialization epoch to help stabilize training. We do not rely on MSE loss as the likelihood loss function since we want the hologram to encode the target image but we are not constrained to having the exact same pixel values. We tested different ρ values, but stick to $\rho = 1$ since for $\rho \ll 1$ the optimizer favors sharp, smooth-phase outputs and for large $\rho \gg 1$ it promotes randomness, producing incoherent results (see rightmost column of Fig. 5). More details on training our real-time ARH can be found in the supplemental materials.

Our random phase metric teaches the model to generate random phase holograms. In Fig. 7, the vanilla HoloNet produces smoothphase holograms, that are completely distorted by artifacts. Our model, trained with the randomness regularizer \mathcal{L}_{RD} , generates artifact-resilient holograms.

5.2 Time Multiplexing

Random phase holograms produce robust results, but they suffer from speckle noise. This problem is commonly solved by time multiplexing different holograms of the same target to preserve image structure while reducing speckle noise.

Our current model is deterministic, therefore it only generates one random phase pattern for a given target image. However, due to the stochastic training (data shuffling, stochastic gradient descent, ...) we can interpret that each trained model is sampling one possible phase pattern per target. Therefore, we implement time multiplexing through model ensembling. By combining different trained models, we can sample random phase holograms, achieving time-multiplexing.

Random phase 3D Holograms in Real Time

In this section, we extend our pre-trained random phase model to generate random phase holograms with realistic 3D effects, without any retraining or fine-tuning, by incorporating depth information into inference. Opposed to HoloNet, we find that our Real-Time ARH generalizes for different depths z due to the random phase

We compute holograms with 3D effects from RGBD images as depicted in Fig. 8. The complex wavefront at the SLM plane U is constructed by layering the image into N depths at different propagation distances z_i , and propagate each slice independently as

$$U = \sum_{i=1}^{N} ASM \left(U_{3D} \odot \delta(D - z_i), -z_i \right), \tag{14}$$

where D denotes the layered depth input for each pixel of the 3D wavefront U_{3D} , and \odot is the Hadamard product.

We report in Fig. 8e simulated results of our 3D extension of our REAL-TIME ARH. In both motorcycle and castle scene, we discretize the depth input into three layers and place them at z = (98, 100, 102)mm, and show the results for all three target planes. These z propagation distances are relative to the SLM, not the viewer, and a 4 mm SLM distance range can span over 10 diopters. The vanilla HOLONET produces strong chromatic aberration effects, as it does not generalize to 3D scenes. Our REAL-TIME ARH, however, naturally generalizes to 3D, producing high-quality in focus imagery, and natural defocus blur.

5.4 Experimental Results

We validate that our additional regularizer unlocks the artifact resilience capabilities of HoloNet by simulating how both holograms generated via our real-time ARH method and HOLONET interact with pre and post-pupil obstructions using our simulation method and hardware phantoms. We ensemble 8 different Real-Time ARH models for 8× time multiplexing, and find that unlike HoloNet, our Real-Time ARH produces holograms that are robust to both eyelashes and floaters in both simulation and experimental captures. As shown in Fig. 9, we see our obstructions cast dark shadows on the image content from holograms generated via vanilla HoloNet. In contrast, holograms generated with our Real-Time ARH leave little to no indication of any obstruction being present.

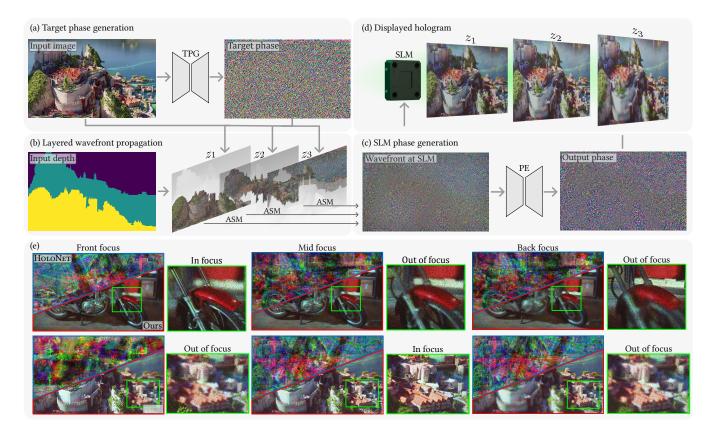


Fig. 8. 3D Holography with Real-Time ARH. We extend our pre-trained Real-Time ARH with 2D images for real-time 3D computed generated holography from RGB-D inputs. We do not finetune the network for 3D scenes, but it naturally generalizes to them. (a) We input the captured 2D image into the Target-Phase Generator (TPG) to sample a target phase. (b) We combine the amplitude (image) and phase and split it into three depth layers based on the depth layer mask. Then, we propagate each wavefront layer a distance z_i independently with the ASM and merge the wavefronts at the SLM plane. (c) We convert the complex wavefront into a phase-only hologram with our Phase Encoder (PE). (d) The phase-only hologram is displayed on a Spatial Light Modulator (SLM), presenting high-quality in-focus image quality, and natural defocus blur. (e) Simulated results. We test our method on two RGB-D scenes, without using time multiplexing, and compare the results of HoloNet with our Real-Time ARH. In both scenes, HoloNet produces incorrect results, with strong chromatic aberration effects due to its lack of generalization to different propagation distances z. In contrast, our Real-Time ARH naturally generalizes to 3D, showing high quality in-focus imagery, while showing natural defocus blur.

We also observe that our Real-Time ARH produces favorable image quality compared to HoloNet *without* obstructions, as the latter requires camera-in-the-loop calibration to overcome the gap between simulation and reality. Our Real-Time ARH proves to be simultaneously more robust to imperfections in hardware and more resilient to dynamic ocular artifacts.

We also tested the artifact resilience of our 3D Real-Time ARH within realistic viewing settings. To this end, we capture these results using a smartphone as an approximator of how the human eye will perceive our ARH holograms. We use the camera focus settings to accommodate to the front and back of the hologram. The vanilla Holoret results exhibit distortions ranging from the lack of natural 3D effects to the visible shadows cast by the eyelashes. In contrast, our Real-Time ARH performs admirably, demonstrating high-quality in focus imagery and realistic defocus blur, while also being artifact resilient.

6 Discussion

We expect that this work inspires the community to pursue random phase holograms as viable alternatives to smooth phase holograms. These types of holograms are inherently artifact resilient and naturally eliminate the need for lengthy camera-in-the-loop calibration procedures that cannot handle dynamic obstructions. These holograms also exhibit more natural defocus cues than their smooth phase counterparts. We also hope our proposed metric for measuring phase randomness and our simulation framework for modeling dynamic artifacts can be used to assess and compare current, past, and future holographic display systems.

The ability of the trained neural network to generate fully random phase holograms can be ascribed to our deterministic metric for measuring phase randomness, or equivalently, the energy distribution of the eyebox. Unlike prior works that affect the eyebox energy distribution indirectly through pupil sampling or by random initialization, we directly measure the distributional distance between the

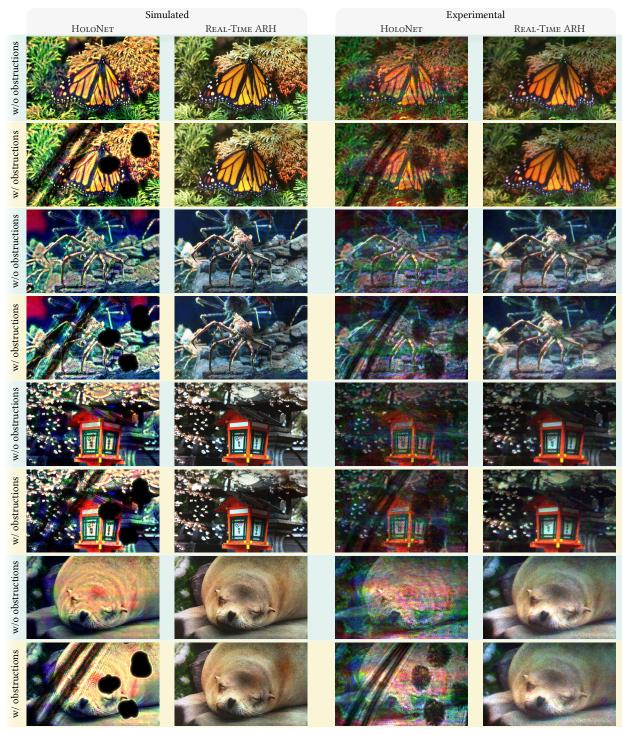


Fig. 9. **HOLONET vs. our Real-Time ARH.** We compare HoloNet vs. our Real-Time ARH in three images, both in simulation and in experiment, with a quantized phase light modulator (TI DLP6750Q1EVM). Without obstructions, our Real-Time ARH produces better visual quality, as HoloNet requires camera in the loop calibration to overcome the hardware imperfections of the real setup. Our Real-Time ARH, as it relies on random phase holography, is naturally more robust to hardware imperfections. With obstructions (eyelashes, floaters), HoloNet holograms are clearly distorted by the obstructions. On the other hand, our Real-Time ARH is artifact resilient, since it relies on random phase holography and light scatters around obstructions. These experimental captures use 8-frame time multiplexing per channel.

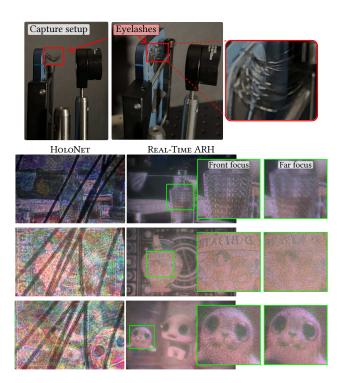


Fig. 10. Artifact Resilience with 3D Effects. We compute holograms from RGB-D inputs using 8-frame time multiplexing per channel, and capture them on a smartphone to test for artifact resilience with eyelashes in front of the camera sensor. HoloNet produces distorted images by the eyelashes, since it is not artifact resilient. It also does not reproduce any defocus blur effects, as it does not generalize to 3D. In contrast, our Real-Time ARH produces holograms that scatter around the obstructions, and are not captured in the final image, proving the artifact-resilience of our method. Also, as our network naturally generalizes to 3D, we not only produce the correct hologram at the target plane, but also produce natural defocus blur.

generated eyebox and the target eyebox. We derive the target mean and variance that we want to match from known properties of the Rayleigh distribution. This metric not only allowed us to train a real time random phase generating neural network, it also allowed us to find a range of pseudo-random holograms that still offered artifact resilience even though they have noticeably less phase randomness than total random holograms.

There are several future directions of investigation that could build on this work. One direction is in the generation of random phase 3D holograms from light fields or other volumetric representations [Choi et al. 2025]. In this work we primarily focused on multiplane RGB-D holograms, it remains an open question what extensions need to be incorporated to generalize real-time random phase holography to focal stacks and light fields. Another direction is the generation of random holograms within the context of pupil under sampling. In this work we assume that eye tracking is applied and that the pupil always fully samples the eye box. Random phase holograms should also support pupil under sampling, but

perhaps more innovations need to be applied to support both pupil invariance and high image quality.

7 Conclusion

This work introduces the first method for real-time artifact-resilient holography. Our method for generating holograms are naturally artifact resilient because of their phase randomness. We depart from existing works that exploit phase smoothness to improve image quality. Instead, we exploit phase randomness to provide resilience against both static and dynamic causes of artifacts, for example, floaters, eyelashes, and dust. The proposed method operates at high frame rates of approximately 85 fps. We assess all methods using our method for simulating dynamic obstructions and we validate the results with experimental captures.

Acknowledgments

Oscar Pueyo-Ciutad was supported by the FPU22/02432 predoctoral grant. Ethan Tseng was supported by a Google PhD Fellowship. Felix Heide was supported by an NSF CAREER Award (2047359), a Packard Foundation Fellowship, a Sloan Research Fellowship, a Sony Young Faculty Award, a Project X Innovation Award, an Amazon Science Research Award, and a Bosch Research Award.

References

Jungkwuen An, Kanghee Won, Young Kim, Jong-Young Hong, Hojung Kim, Yongkyu Kim, Hoon Song, Chilsung Choi, Yunhee Kim, Juwon Seo, et al. 2020. Slim-panel holographic video display. *Nature communications* 11, 1 (2020), 5568.

Ozan Cakmakci and Jannick Rolland. 2006. Head-Worn Displays: A Review. Display Technology, Journal of 2 (10 2006), 199 – 216. doi:10.1109/JDT.2006.879846

Praneeth Chakravarthula, Seung-Hwan Baek, Florian Schiffers, Ethan Tseng, Grace Kuo, Andrew Maimone, Nathan Matsuda, Oliver Cossairt, Douglas Lanman, and Felix Heide. 2022a. Pupil-Aware Holography. ACM Trans. Graph. 41, 6, Article 212 (Nov. 2022), 15 pages. doi:10.1145/3550454.3555508

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

Praneeth Chakravarthula, Ethan Tseng, Henry Fuchs, and Felix Heide. 2022b. Hogel-Free Holography. ACM Trans. Graph. 41, 5, Article 178 (Oct. 2022), 16 pages. doi:10. 1145/3516428

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 Trans. Graph. 39, 6, Article 186 (Nov. 2020), 18 pages. doi:10.1145/ 3414685.3417846

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

Suyeon Choi, Manu Gopakumar, Yifan Peng, Jonghyun Kim, Matthew O'Toole, and Gordon Wetzstein. 2022. Time-multiplexed Neural Holography: A Flexible Framework for Holographic Near-eye Displays with Fast Heavily-quantized Spatial Light Modulators. In ACM SIGGRAPH 2022 Conference Proceedings (Vancouver, BC, Canada) (SIGGRAPH '22). Association for Computing Machinery, New York, NY, USA, Article 32, 9 pages. doi:10.1145/3528233.3530734

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

Suyeon Choi, Jonghyun Kim, Yifan Peng, and Gordon Wetzstein. 2021b. Optimizing image quality for holographic near-eye displays with Michelson holography. Optica 8, 2 (2021), 143–146.

M. Hossein Eybposh, Nicholas W. Caira, Mathew Atisa, Praneeth Chakravarthula, and Nicolas C. Pégard. 2020. DeepCGH: 3D computer-generated holography using deep learning. Opt. Express 28, 18 (Aug 2020), 26636–26650. doi:10.1364/OE.399624

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

Manu Gopakumar, Jonghyun Kim, Suyeon Choi, Yifan Peng, and Gordon Wetzstein. 2021. Unfiltered holography: Optimizing high diffraction orders without optical filtering for compact holographic displays. Optics Letters 46, 23 (2021), 5822–5825.

- Manu Gopakumar, Gun-Yeal Lee, Suyeon Choi, Brian Chao, Yifan Peng, Jonghyun Kim, and Gordon Wetzstein. 2024. Full-colour 3D holographic augmented-reality displays with metasurface waveguides. Nature 629, 8013 (2024), 1-7.
- Maximilian Hammer, Marcel Muuss, Sonja Schickhardt, Alexander Scheuerle, Ramin Khoramnia, Grzegorz Labuz, Philipp Uhl, and Gerd Uwe Auffarth. 2024. Forward Light Scattering of the Vitreous Gel After Enzymatic Aging: An In Vitro Model to Study Vitreous Opacification. Investigative Ophthalmology & Visual Science 65, 3
- Ryoichi Horisaki, Ryosuke Takagi, and Jun Tanida. 2018. Deep-learning-generated holography. Applied optics 57, 14 (2018), 3859-3863.
- Yuta Itoh, Tobias Langlotz, Jonathan Sutton, and Alexander Plopski. 2021. Towards Indistinguishable Augmented Reality: A Survey on Optical See-through Headmounted Displays. ACM Comput. Surv. 54, 6, Article 120 (July 2021), 36 pages. doi:10.1145/3453157
- Justin Johnson, Alexandre Alahi, and Li Fei-Fei. 2016. Perceptual Losses for Real-Time Style Transfer and Super-Resolution. In Computer Vision - ECCV 2016, Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling (Eds.). Springer International Publishing, Cham, 694-711.
- Leyla Kabuli, Oliver Cossairt, Florian Schiffers, Nathan Matsuda, and Grace Kuo. 2025. High contrast holography through dual modulation. Scientific Reports 15, 1 (21 May 2025), 17615. doi:10.1038/s41598-025-00459-8
- Koray Kavakli, Yuta Itoh, Hakan Urey, and Kaan Aksit. 2023. Realistic Defocus Blur for Multiplane Computer-Generated Holography . In 2023 IEEE Conference Virtual Reality and 3D User Interfaces (VR). IEEE Computer Society, Los Alamitos, CA, USA, 418-426. doi:10.1109/VR55154.2023.00057
- Kevin Patrick Kennelly, James Plunkett Morgan, David Jude Keegan, and Paul Patrick Connell. 2015. Objective assessment of symptomatic vitreous floaters using optical coherence tomography: a case report. BMC ophthalmology 15 (2015), 1-3.
- Dongyeon Kim, Seung-Woo Nam, Suyeon Choi, Jong-Mo Seo, Gordon Wetzstein, and Yoonchan Jeong. 2024. Holographic Parallax Improves 3D Perceptual Realism. ACM Trans. Graph. 43, 4, Article 68 (July 2024), 13 pages. doi:10.1145/3658168
- Dongyeon Kim, Seung-Woo Nam, Byounghyo Lee, Jong-Mo Seo, and Byoungho Lee. 2022b. Accommodative holography: improving accommodation response for perceptually realistic holographic displays. ACM Transactions on Graphics (TOG) 41, 4 (2022), 1-15.
- Jonghyun Kim, Manu Gopakumar, Suyeon Choi, Yifan Peng, Ward Lopes, and Gordon Wetzstein. 2022a. Holographic Glasses for Virtual Reality. In ACM SIGGRAPH 2022 Conference Proceedings (Vancouver, BC, Canada) (SIGGRAPH '22). Association for Computing Machinery, New York, NY, USA, Article 33, 9 pages. doi:10.1145/3528233.
- Grace Kuo, Florian Schiffers, Douglas Lanman, Oliver Cossairt, and Nathan Matsuda. 2023. Multisource holography. ACM Transactions on Graphics (TOG) 42, 6 (2023),
- Grace Kuo, Laura Waller, Ren Ng, and Andrew Maimone. 2020. High-Resolution Étendue Expansion for Holographic Displays. ACM Transactions on Graphics (TOG) 39, 4 (2020), 1–9. https://dl.acm.org/doi/abs/10.1145/3386569.3392414
- Zhao-Song Li, Chao Liu, Xiao-Wei Li, Yi Zheng, Qian Huang, Yi-Wei Zheng, Ye-Hao Hou, Chen-Liang Chang, Da-Wei Zhang, Song-Lin Zhuang, Di Wang, and Qiong-Hua Wang. 2025. Real-time holographic camera for obtaining real 3D scene hologram. Light: Science & Applications 14, 1 (08 Feb 2025), 74. doi:10.1038/s41377-024-01730-9
- Andrew Maimone, Andreas Georgiou, and Joel S Kollin. 2017. Holographic near-eye displays for virtual and augmented reality. ACM Transactions on Graphics (TOG) 36,
- Andrew Maimone and Junren Wang. 2020. Holographic optics for thin and lightweight virtual reality. ACM Transactions on Graphics (TOG) 39, 4 (2020), 67-1.
- Eric Markley, Nathan Matsuda, Florian Schiffers, Oliver Cossairt, and Grace Kuo. 2023. Simultaneous Color Computer Generated Holography. In SIGGRAPH Asia 2023 Conference Papers (Sydney, NSW, Australia) (SA '23). Association for Computing Machinery, New York, NY, USA, Article 22, 11 pages. doi:10.1145/3610548.3618250
- Nathan Matsuda, Alexander Fix, and Douglas Lanman. 2017. Focal surface displays. ACM Transactions on Graphics (TOG) 36, 4 (2017), 1-14.
- Kyoji Matsushima and Tomoyoshi Shimobaba. 2009. Band-Limited Angular Spectrum Method for Numerical Simulation of Free-Space Propagation in Far and Near Fields. Optics express 17 (10 2009), 19662-73. doi:10.1364/OE.17.019662
- Sagi Monin, Aswin C. Sankaranarayanan, and Anat Levin. 2022. Analyzing phase masks for wide étendue holographic displays . In 2022 IEEE International Conference on Computational Photography (ICCP). IEEE Computer Society, Los Alamitos, CA, USA, 1-12. doi:10.1109/ICCP54855.2022.9887757
- 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.
- Alba M Paniagua-Diaz, Justin H Nguyen, Pablo Artal, Wei Gui, and J Sebag. 2024. Light Scattering by Vitreous of Humans With Vision Degrading Myodesopsia From Floaters. Investigative Ophthalmology & Visual Science 65, 5 (2024), 20-20.
- Yifan Peng, Suyeon Choi, Jonghyun Kim, and Gordon Wetzstein. 2021. Speckle-free holography with partially coherent light sources and camera-in-the-loop calibration.

- Science Advances 7, 46 (2021), eabg5040.
- Yifan Peng, Suyeon Choi, Nitish Padmanaban, and Gordon Wetzstein. 2020. Neural Holography with Camera-in-the-Loop Training. ACM Trans. Graph. 39, 6, Article 185 (nov 2020), 14 pages. doi:10.1145/3414685.3417802
- Mark A. Richards. 2013. The Discrete-Time Fourier Transform and Discrete Fourier Transform of Windowed Stationary White Noise. Technical Memorandum. Georgia Institute of Technology. 1-24 pages
- Florian Schiffers, Praneeth Chakravarthula, Nathan Matsuda, Grace Kuo, Ethan Tseng, Douglas Lanman, Felix Heide, and Oliver Cossairt. 2023. Stochastic Light Field Holography . In 2023 IEEE International Conference on Computational Photography (ICCP). IEEE Computer Society, Los Alamitos, CA, USA, 1-12. doi:10.1109/ICCP56744.
- Florian Andreas Schiffers, Grace Kuo, Nathan Matsuda, Douglas Lanman, and Oliver Cossairt. 2025. HoloChrome: Polychromatic Illumination for Speckle Reduction in Holographic Near-Eye Displays. ACM Trans. Graph. 44, 3, Article 26 (May 2025), 18 pages. doi:10.1145/3732935
- Liang Shi, Beichen Li, Changil Kim, Petr Kellnhofer, and Wojciech Matusik. 2021. Towards real-time photorealistic 3D holography with deep neural networks. Nature 591, 7849 (2021), 234-239.
- Liang Shi, Beichen Li, and Wojciech Matusik. 2022a. End-to-end learning of 3D phaseonly holograms for holographic display. Light Sci Appl 11, 1 (Aug. 2022), 247.
- Liang Shi, Richard Webb, Lei Xiao, Changil Kim, and Changwon Jang. 2022b. Neural compression for hologram images and videos. Optics Letters 47, 22 (2022), 6013-6016.
- Ethan Tseng, Grace Kuo, Seung-Hwan Baek, Nathan Matsuda, Andrew Maimone, Florian Schiffers, Praneeth Chakravarthula, Qiang Fu, Wolfgang Heidrich, Douglas Lanman, et al. 2024. Neural étendue expander for ultra-wide-angle high-fidelity holographic display. Nature communications 15, 1 (2024), 2907.
- Yujie Wang, Praneeth Chakravarthula, Qi Sun, and Baoquan Chen. 2022. Joint neural phase retrieval and compression for energy- and computation-efficient holography on the edge. ACM Trans. Graph. 41, 4, Article 110 (July 2022), 16 pages. doi:10.1145/ 3528223.3530070
- Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. IEEE transactions on image processing 13, 4 (2004), 600-612.
- 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). IEEE Computer Society, Los Alamitos, CA, USA, 312-319. doi:10.1109/ ISMAR50242.2020.00057
- 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.
- Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. 2017. Unpaired Imageto-Image Translation Using Cycle-Consistent Adversarial Networks . In 2017 IEEE International Conference on Computer Vision (ICCV). IEEE Computer Society, Los Alamitos, CA, USA, 2242-2251. doi:10.1109/ICCV.2017.244