

Artifact-Resilient Real-Time Holography

- Supplemental Document -

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, I3A, Spain

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

FELIX HEIDE, Princeton University, USA

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

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

ACM Reference Format:

Victor Chu, Oscar Pueyo-Ciudad, Ethan Tseng, Florian Schiffrers, Grace Kuo, Nathan Matsuda, Albert Redo-Sanchez, Douglas Lanman, Oliver Cossairt, and Felix Heide. 2025. Artifact-Resilient Real-Time Holography - Supplemental Document -. *ACM Trans. Graph.* 44, 6, Article 219 (December 2025), 14 pages. <https://doi.org/10.1145/3763361>

This supplementary document provides additional detail and further experimental results in support of the findings from the main manuscript. Specifically, we provide the following material in this document.

CONTENTS

Contents	1
1 Additional Simulation Details	1
2 Additional Experiment Details	2
3 Additional Simulation Results	2
4 Additional Experimental Results	4
References	4

1 Additional Simulation Details

Additional Neural Network Training Details. We trained all REAL-TIME ARH models on a NVIDIA L40 for 40 epochs over the DIV2K training subset with a learning rate of $1e-3$. As mentioned in the main paper, the first epoch is trained without our Rayleigh Distance

Authors' Contact Information: Victor Chu, Princeton University, Princeton, NJ, USA; Oscar Pueyo-Ciudad, Universidad de Zaragoza, I3A, Zaragoza, Spain; Ethan Tseng, Princeton University, Princeton, NJ, USA; Florian Schiffrers; 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>

regularizer in order to make training more stable. Our holograms have resolution 1280×800 .

For our regularizer parameters, we set the size of the Gaussian blur kernel to be 11 with $\sigma_K = 0.1$. We use $\mu_R = 0.2$ and $\sigma_R = 0.1$.

We find that runs that diverge (produce random noise outputs at the object plane) within the first few epochs do not converge again. Our training took roughly 8 hours per model.

Vanilla random phase optimization (VRO) algorithmic details. We used the Adam optimizer [Kingma and Ba 2015] for all holograms generated via vanilla random phase optimization seen in the paper and in the supplement. We initialized the phase to be uniform random values between 0 and 2π . The learning rate was set to 0.1 and we used MSE loss.

Parameters for modeling floaters. Eye floaters are often small aggregations of collagen proteins. These type II collagen proteins have a complex index of refraction of about $\sim 1.49 + j0.001$ relative to the surrounding fluid of the vitreous humor, which has a refractive index of about 1.335 [Harmer et al. 2022]. Using the relative index of refraction, we can calculate both the amplitude attenuation and the phase response of the floaters on our holographic display. We use the following model to calculate the resultant amplitude A after passing through a medium with complex index of refraction $\tilde{n} = n + j\kappa$ and input amplitude A_0 , that is

$$A = A_0 \exp(-\alpha L) = A_0 \exp\left(-\frac{2\pi}{\lambda} \kappa L\right) \quad (1)$$

for a given thickness L and wavelength λ . The attenuation factor is thus $\exp\left(-\frac{2\pi}{\lambda} \kappa L\right)$. To calculate the phase shift ϕ of a medium with complex index of refraction $\tilde{n} = n + j\kappa$, we use the following

$$\phi = \frac{2\pi}{\lambda} nL \quad (2)$$

for a given thickness L and wavelength λ . This means that the real component n of the relative refractive index between the aqueous humor and the collagen proteins that make up the floater is $1.49/1.335 \approx 1.116$. Using Equations 1 and 2, we can calculate the respective attenuation factor and phase shift of the floaters as $\exp\left(-\frac{2\pi}{\lambda} (0.001)L\right)$ and $\frac{2\pi}{\lambda} (1.116)L$, respectively. Floaters generally range from a few micrometers to several tens of micrometers [Karitans et al. 2021].

We then use the respective attenuation and phase shift to generate a complex modulation floater mask. For floaters, we set $z_{\text{Post-Pupil}}$ to

be 5-15 mm to place the floater mask between the retina and the pupil.

Parameters for Modeling Eyelashes. As eyelashes range between 8-12 mm [Thibaut et al. 2010], we limit the eyebox propagation distance (z_{Eyelash}) to 4-8 mm. Since eyelashes are fully opaque, M is an intensity mask with value 1 where there are no eyelashes and 0 where there is an eyelash.

2 Additional Experiment Details

Additional details for the spatial light modulator. We performed our experiments using a TI DLP6750Q1EVM. This SLM exhibits a 1440 Hz frame rate. We target a performance of 60 Hz, similar to many commercial displays. When running at 60 Hz we have a budget of 24 frames for rendering the full color hologram. We budget 8 frames towards each color. The phase patterns that we optimize for are encoded following the DLP's prescribed encoding system. Our time multiplexing scheme facilitates temporal denoising by averaging in manner similar to existing work on temporal multiplexing for holography [Choi et al. 2022]. Different phase patterns are generated for each color and different wavelength laser diodes are correspondingly switched on and off for each phase pattern.

The phase response of the DLP is non-linear and so we first calibrate for this non-linearity by determining the response function of the DLP to different wavelengths. This is done by displaying random phase patterns at different resolutions (binning the pixels of the SLM by factors of 1, 5, 10, etc.) in order to obtain images with interpretable speckle to learn from. We manually search for a look up table scaling factor that minimizes the gap between the simulated results and experimental captures. As we know the rough propagation distance and that the look up table scaling factor should be between 1.0 and 2.5, this is not an unreasonably large search space. After obtaining a reasonable initialization, we learn the look up table scaling factor, propagation distance, and a homography using gradient descent to minimize differences between simulations and experimental measurements. No additional calibration is performed.

Additional optical hardware configuration details. We use a 100 mm eyepiece lens together with a 80 mm pupil lens. This results in a FOV of 9° and an eyebox size of 4 mm when using the full étendue of the SLM. See Fig. 1 for how our hardware phantoms are aligned in our setup and Fig. 5 to see the effect caused by these phantoms under uniform illumination.

This particular hardware configuration is only used to validate our key insight: that random phase holograms with larger eyeboxes exhibit greater resilience to artifacts. The étendue of the system can easily be redistributed in a practical device that is viewed by the eye pupil while preserving appropriate eye relief.

Varying Eyebox Size. We utilize the existing bandwidth of our hardware to investigate the impact of a changing eyebox size and we disregard the need for a large field-of-view. Specifically, we use the étendue of our SLM by setting the field-of-view to be at 9° and the eyebox to be at 4 mm. To emulate a smaller eyebox system we enforce learned phase patterns to have a 1.33 mm eyebox. We do this by Fourier filtering everything outside of the central 1.33 mm while optimizing the phase pattern. Then, when displaying these

Table 1. **Quantitative Improvement from Time Multiplexing.** We show how the image quality (PSNR) of our models and HoloNet [Peng et al. 2020] changes with differing levels of time multiplexing (1×, 2×, 4×, and 8× frames).

	1× TMX PSNR (dB) ↑	2× TMX PSNR (dB) ↑	4× TMX PSNR (dB) ↑	8× TMX PSNR (dB) ↑
Ours: REAL-TIME ARH	15.7	16.7	17.4	17.7
HoloNet	20.4	19.9	19.7	20.3

patterns in hardware, we ensure that the iris aperture in our setup is set to be 1.33 mm, thereby reducing the eyebox size while keeping the field-of-view and available degrees of freedom unchanged. We choose these settings because it allows us to assess the cross-over point between an eyebox that is too small to provide resilience to artifacts, even when using random phase, and an eyebox that is sufficiently large to provide the needed resilience.

3 Additional Simulation Results

Comparison against other real-time random phase methods. While other random-phase methods have been shown to be successful, we find that these methods sacrifice image contrast and speckle reduction for speed. Specifically in Fig. 2, we see that our method shows better image contrast simple amplitude discard (performing one step of Gerchberg Saxton [Gerchberg 1972] by initializing the object field with random phase, propagating it back to the SLM plane, and using the phase of the propagated wavefront on the SLM) while still only requiring one ASM propagation step.

Additional comparisons of CGH algorithms. We show the phase patterns generated by different CGH algorithms in Fig. 3. We observe that both DPAC [Maimone et al. 2017] and Tensor Holography [Shi et al. 2021] produce smooth phase holograms that visually resemble the target images. While sideband encoding with an amplitude SLM [Kim et al. 2024] is capable of producing a wavefront with random phase characteristics, the method requires that half of the eyebox to be blocked. Vanilla random phase optimization produces random phase holograms while utilizing the full étendue of the display.

Analyzing Time-Multiplexing. In the main text, we primarily compare single frame smooth-phase HoloNets against our eight frame Real-Time ARH models. We found this to be a fair comparison because single frame smooth-phase HoloNets is what is proposed in the original HoloNet paper [Peng et al. 2020]. Additionally, we verified that using multiple smooth-phase HoloNets for 8-frame time multiplexing yields no improvement versus a 2 dB gain in our Real-Time ARH networks using our ensembled time-multiplexing approach. We find our comparisons to be fair. These results are detailed in table 1.

Given that our model is lightweight and does not saturate the computational capacity of the L40 GPU, we report the single-frame latency (85 FPS) as it represents the theoretical throughput in a fully parallelized system. This treatment differs from other methods for fundamental reasons: VRO requires a joint optimization across all frames, DPAC (as described in [Maimone et al. 2017]) is not




Fig. 1. **Experimental Setup.** Schematic details on how both pre- and post-pupil interference (eyelashes and floaters) were emulated in our experimental captures. After the eyepiece lens and prior to the eyebox plane, pieces of human hair obstruct the wavefront. The yellow box shows our pre-pupil obstruction (pieces of human hair) as well as the obstruction under red light. Similarly, after the pupil lens, we place a piece of glass with several dots of expo marker applied to emulate a complex of index of refraction within the eye (floaters). The green box images both the glass slide as well as the obstructions under white light.




Fig. 2. **Amplitude Discard vs REAL-TIME ARH.** We compare the visual quality and PSNR in dB (bottom right) of three images obtained with the Amplitude Discard method and our REAL-TIME ARH, both without time multiplexing. Our method obtains higher visual quality and PSNR values than the standard amplitude discard.

amenable to time multiplexing since its deterministic operation on a

constant initialization produces no temporal variation, and smooth-phase neural networks do not gain a significant quality benefit from

multiplexing, making their reported latencies for the full sequence the most direct comparison. While optimizations such as quantization could potentially allow the complete multi-frame pipeline to run concurrently on a single GPU, we leave this implementation as an avenue for future work.

4 Additional Experimental Results

Additional 3D REAL-TIME ARH captures. Fig. 4, we provide additional experimental captures of REAL-TIME ARH for holograms supervised on RGBD input.

Floaters and Eyelashes. We show additional experimental results in support of the main manuscript. Images of the eyelash and floater phantoms under red light plane wave illumination are shown in Fig. 5. Figures 6, 7, 8, 9, and 10 show experimentally captured 2D holograms when different phantoms are applied. We ablate between no obstructions, applying only the floater phantom, applying only the eyelash phantom, and applying both phantoms simultaneously. We observe that in all cases holograms generated using our vanilla random phase optimization outperforms the smooth phase holograms. Furthermore, the results also show that the highest reconstruction fidelity is achieved when using a large eyebox (4mm).

Alternative phantom design (smudges and scratches). Fig. 12 shows experimentally captured 2D holograms when applying a third type of phantom. This phantom is constructed by taking a thin plastic lid and scratching the top surface with a dull knife. The phantom was then placed in the imaging path causing the distortion seen in Fig. 11. It exhibits a stronger phase distortion response and a smaller amplitude reduction than the other two phantoms. We again observe that random phase holograms generated with vanilla random phase optimization provides superior resistance to lens artifacts – in this case both scratches and smudges on either user glasses or the display optics – than smooth phase holograms.

Additional 3D results. We also show an additional 3D hologram result generated via vanilla random phase optimization in Figures 13 and 14.

References

- 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
- Ralph W Gerchberg. 1972. A practical algorithm for the determination of the phase from image and diffraction plane pictures. *Optik* 35 (1972), 237–246.
- Stuart W Harmer, Andrew J Luff, and Giampaolo Gini. 2022. Optical scattering from vitreous floaters. *Bioelectromagnetics* 43, 2 (2022), 90–105.
- Varis Karitans, Sergejs Fomins, and Maris Ozolinsh. 2021. Phase retrieval for studying the structure of vitreous floaters simulated in a model eye. *Journal of Modern Optics* 68, 15 (2021), 792–797.
- 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
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*. <https://openreview.net/forum?id=8gmWwjFyLj>
- 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.
- 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
- 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.
- S. Thibaut, E. De Becker, L. Caisey, D. Baras, S. Karatas, O. Jammayrac, P.J. Pisella, and B.A. Bernard. 2010. Human eyelash characterization. *British Journal of Dermatology* 162, 2 (02 2010), 304–310. arXiv:<https://academic.oup.com/bjd/article-pdf/162/2/304/47588218/bjd0304.pdf> doi:10.1111/j.1365-2133.2009.09487.x




Fig. 3. **Comparison of Phase Patterns Generated by CGH Algorithms:** Above are phase patterns generated by DPAC [Maimone et al. 2017], Tensor Holography [Shi et al. 2021], single sideband [Kim et al. 2024], and vanilla random phase optimization for the red channel of the target amplitude (shown in the first row). DPAC and Tensor Holography generate smooth phase patterns, while our Real-Time ARH generate random phase patterns.




Fig. 4. **HOLONET vs our REAL-TIME ARH.** Experimental capture of the holograms generated by HOLONET and Real-Time ARH under obstructions. For each row, we show the results in three focal planes. HOLONET results are distorted for all images, while our REAL-TIME ARH is robust to artifacts.




Fig. 5. **Eyelash and Floater Phantoms.** We show captures of the eyelash and floater phantoms being illuminated by collimated red light (by displaying constant phase on the SLM). These phantoms positions are used for the following captures. The eyelash phantoms are placed right before the eyebox plane. The floater phantoms are placed between the pupil lens and the detector.




Fig. 6. **Experimental Captures of Pre-Pupil Interference (2D).** We show experimental captures where we display a smooth phase pattern alongside two patterns using vanilla random phase optimization – one with a 1.3 mm eyebox and the other with a 4 mm eyebox. For this figure, only the pre-pupil obstruction (eyelash phantom) is applied. Experiments reveal that even a small dense eyebox (1.3 mm) can effectively mitigate the impact of eyelashes. In contrast, smooth phase holograms, as shown in simulations, fail to guard against obstructions occurring before the pupil.




Fig. 7. **Experimental Captures of Post-Pupil Interference (2D).** This figure shows the same three configurations as Fig. 6; however, here, only the post-pupil obstruction (floater phantom) is applied. Experimental validation mirrors our simulation results, showing that larger eyeboxes enhance resilience to artifacts. Following our simulation results, smooth phase holograms provide insufficient protection against these post-pupil obstructions.




Fig. 8. **Experimental Captures of Artifacts (2D).** Additional experimental captures of artifacts are provided. The columns show our three eyebox configurations with smooth phase, vanilla random phase optimization (VRO) (1.3 mm), and VRO (4 mm). The rows show the different viewing conditions: unobstructed, only floater, only eyelash, and eyelash and floater interference together. Observe that these holograms also are affected by the presence of dust particles in our experimental setup (see the ringing artifacts appearing throughout the unobstructed smooth phase example). These dust artifacts disappear with VRO even with a 1.3 mm eyebox.




Fig. 9. **Experimental Captures of Artifacts (2D).** Additional experimental captures of artifacts are provided.




Fig. 10. **Experimental Captures of Artifacts (2D).** Additional experimental captures of artifacts are provided.




Fig. 11. **Construction of Scratch and Smudge Phantom.** We create a phantom that emulates scratches and smudges in the optical path of a holographic setup using a thin plastic lid (middle). We show both an unobstructed image of our collimated red source (left) as well as passing through our setup as well as the same setup with the scratch and smudge phantom placed inside the optical path (right). The red arrows note places where the scratches from the phantom interfere with the image.




Fig. 12. **Experimental Captures of Scratch and Smudge Artifacts (2D).** Using our scratch and smudge phantom, we capture holograms with a smooth phase pattern and holograms generated using vanilla random phase optimization (in both the 1.3 mm Eyebox and 4 mm Eyebox configuration). When interfered with our scratches and smudges phantom, we show that smooth phase holograms are severely degraded as compared to random phase holograms.




Fig. 13. **Experimental 3D Results.** This figure features 3D holograms spanning 3 diopter values, comparing the performance of smooth phase and random phase holograms in the presence of artifacts. In the scene, the background is placed at 0 diopters (far focus), the violin at 1 diopters (middle focus) and the vase is at 2 diopters (near focus). For the interfered cases (left and middle), both pre and post pupil obstructions are applied in the imaging path. The unobstructed holograms generated with vanilla random phase optimization is provided as a baseline for image contrast. Random phase holograms with a 4mm eyebox show impressive capabilities to resist artifacts as well as have accurate defocus cues.




Fig. 14. **Experimental 3D Results.** An additional 3D holograms spanning 3 diopter values, comparing the performance of smooth phase and random phase holograms in the presence of artifacts is provided. Each number is placed at its respective diopter with 0D representing the 0th diopter and so on.