Simultaneous Color Computer Generated Holography

Eric Markley
emarkley@berkeley.edu
Reality Labs Research, Meta
USA

Nathan Matsuda
nathan.matsuda@meta.com
Reality Labs Research, Meta
USA

Florian Schiffers
fschiffers@meta.com
Reality Labs Research, Meta
USA

Oliver Cossairt
ocoissart@meta.com
Reality Labs Research, Meta
USA

Grace Kuo
gracekuo@meta.com
Reality Labs Research, Meta
USA

Figure 1: Simultaneous color holograms captured in experiment. Traditionally, color holograms are illuminated sequentially with a unique spatial light modulator (SLM) pattern for each color channel. In this work we outline a flexible framework that enables the use of a single SLM pattern for red-green-blue (RGB) holograms using simultaneous RGB illumination. We validate this framework experimentally on a simple and compact optical setup.

ABSTRACT

Computer generated holography has long been touted as the future of augmented and virtual reality (AR/VR) displays, but has yet to be realized in practice. Previous high-quality, color holographic displays have made either a 3× sacrifice on frame rate by using a sequential color illumination scheme or used more than one spatial light modulator (SLM) and/or bulky, complex optical setups. The reduced frame rate of sequential color introduces distracting judder and color fringing in the presence of head motion while the form factor of current simultaneous color systems is incompatible with a head-mounted display. In this work, we propose a framework for simultaneous color holography that allows the use of the full SLM frame rate while maintaining a compact and simple optical setup. Simultaneous color holograms are optimized through the use of a perceptual loss function, a physics-based neural network wavefront propagator, and a camera-calibrated forward model. We measurably improve hologram quality compared to other simultaneous color methods and move one step closer to the realization of color holographic displays for AR/VR.

SA Conference Papers '23, December 12–15, 2023, Sydney, NSW, Australia

ACM ISBN 979-8-4007-0315-7/23/12.
https://doi.org/10.1145/3610548.3618250

CCS CONCEPTS

• **Hardware** → **Emerging optical and photonic technologies**; • **Computing methodologies** → **Modeling and simulation**; • **Human-centered computing** → **Virtual reality**; **Mixed / augmented reality**.

KEYWORDS

holography, neural networks, augmented reality, virtual reality, computational displays

ACM Reference Format:
Eric Markley, Nathan Matsuda, Florian Schiffers, Oliver Cossairt, and Grace Kuo. 2023. Simultaneous Color Computer Generated Holography. In *SIGGRAPH Asia 2023 Conference Papers (SA Conference Papers '23), December 12–15, 2023, Sydney, NSW, Australia.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3610548.3618250

1 INTRODUCTION

Holographic displays are a promising technology for augmented and virtual reality (AR/VR). Such displays use a spatial light modulator (SLM) to shape an incoming coherent wavefront so that it appears as though the wavefront came from a real, three-dimensional (3D) object. The resulting image can have natural defocus cues, providing a path to resolve the vergance-accommodation conflict of stereoscopic displays [Kim et al. 2022b]. Additionally, the fine-grain control offered by holography can also correct for optical aberrations, provide custom eyeglass prescription correction in software, and enable compact form-factors [Maimone et al. 2017], while improving light efficiency compared to traditional LCD or

OLED displays [Yin et al. 2022]. Recent publications have demonstrated significant improvement in hologram image quality [Choi et al. 2021a; Maimone et al. 2017; Peng et al. 2020] and computation time [Eybposh et al. 2020; Shi et al. 2021], but color holography for AR/VR has remained an open problem.

Traditionally, red-green-blue (RGB) holograms are created through *field sequential color*, where a separate hologram is computed for each of the three wavelengths; these are displayed in sequence and synchronized with the color of the illumination source. Due to persistence of vision, this appears as a single full color image if the update is sufficiently fast, enabling color holography for static displays. However, in a head-mounted AR/VR system displaying world-locked content, frame rate requirements are higher to prevent noticeable judder [Van Waveren 2016]. In fact, all modern VR displays are "low persistance" meaning the image content is only displayed for a fraction of the frame time (usually about 10%) and no content is shown during the rest of the frame [Zielinski et al. 2015]. This is usually achieved by strobing the illumination, but if one wished to display three sequential color frames all within a 10% persistence time, it would require the display to update 30× faster than the effective frame rate. Without low persistence, field sequential color leads to strong color fringing (visible spatial separation of the colors) particularly when the user rotates their head while tracking a fixed object with their eyes [Riecke et al. 2006].

Low frame rate displays exacerbate these artifacts, and the most common SLM technology for holography, liquid-crystal-on-silicon (LCoS), is quite slow due to the physical response time of the liquid crystal (LC) layer [Zhang et al. 2014]. Although most commercial LCoS SLMs can be driven at 60 Hz, at that speed the SLM will have residual artifacts from the prior frames [Haist and Osten 2015]. High speed SLMs based on micro-electro-mechanical system (MEMS) [Choi et al. 2022; Duerr et al. 2021] or dual-frequency LCoS [Serati et al. 2003] are becoming more widely available, but even with these devices, simultaneous color is desirable since it eliminates color fringing, enables low persistence, and frees temporal bandwidth for other uses, such as increasing the effective etendue by scanning the field of view or eyebox position [Lee et al. 2020].

In this work, we aim to display RGB holograms using only a single SLM pattern, enabling a 3× increase in frame rate compared to sequential color and completely removing color fringing artifacts. Our compact setup does not use a physical filter in the Fourier plane or bulky optics to combine color channels. Instead, the full SLM is simultaneously illuminated by an on-axis RGB source, and we optimize the SLM pattern to form the full color image. We design a flexible framework for end-to-end optimization of the digital SLM input from the target RGB intensity, allowing us to optimize for SLMs with extended phase range, and we develop a color-specific perceptual loss function which further improves color fidelity. Our method is validated experimentally on 2D and 3D content.

Specifically, we make the following contributions:

- We introduce a novel algorithm for generating simultaneous color holograms which takes advantage of the extended phase range of the SLM in an end-to-end manner and uses a new loss function based on human color perception.
- We analyze the "depth replicas" artifact in simultaneous color holography and demonstrate how these replicas can be mitigated with extended phase range.
- We demonstrate experimental simultaneous color holograms in both 2D and 3D using a custom camera-calibrated model.

2 RELATED WORKS

Field Sequential Color. The vast majority of color holographic displays use field sequential color in which the SLM is sequentially illuminated by red, green, and blue sources while the SLM pattern is updated accordingly [Chakravarthula et al. 2022, 2019, 2020; Choi et al. 2021a,b; Jang et al. 2018; Li et al. 2016; Maimone et al. 2017; Peng et al. 2021, 2020; Shi et al. 2021; Yang et al. 2022]. Field sequential color is effective at producing full color holograms but reduces frame rate by a factor of 3× and creates color fringing artifacts in the presence of head motion. These limitations are not alleviated by recent work where the color of each sub-frame is manipulated to increase peak brightness [Kavaklı et al. 2023], and they present a particular challenge for LCoS SLMs where refresh rate is severely limited by the LC response time [Zhang et al. 2014]. Although, SLMs based on MEMS technology can run at high frame rates in the kilohertz range [Duerr et al. 2021], so far these modulators are maximum 4-bit displays, with most being binary [Choi et al. 2022; Kim et al. 2022b; Lee et al. 2022]. Even with emerging 8-bit high frame rate modulators [Serati et al. 2003], it may be worthwhile to maintain the full temporal bandwidth, since the extra bandwidth can be used to address other holography limitations. For example, speckle can be reduced through temporal averaging [Choi et al. 2022; Kim et al. 2022b; Lee et al. 2022], and limited etendue can be mitigated through pupil scanning [Jang et al. 2018; Kim et al. 2022a].

Spatial Multiplexing. An alternate approach is spatial multiplexing, which maintains the native SLM frame rate by using different regions of the SLM for each color. Most prior works in this area use three separate SLMs and an array of optics to combine the wavefronts [Nakayama et al. 2010; Shiraki et al. 2009; Yaraş et al. 2009]. Although this method produces high quality holograms, the resulting systems are bulky, expensive, and require precise alignment, making them poorly suited for near-eye displays. Spatial multiplexing can also be implemented with a single SLM split into sub-regions [Makowski et al. 2011, 2009]; while less expensive, this approach still requires bulky combining optics and sacrifices space-bandwidth product (SBP), also known as etendue. Etendue is already a limiting factor in holographic displays [Kuo et al. 2020], and further reduction limits the range of viewing angles or display field-of-view.

Frequency Multiplexing. Rather than split the physical extent of the SLM into regions, frequency multiplexing assigns each color a different region in the frequency domain, and the colors are separated with a physical color filter at the Fourier plane of a $4f$ system [Lin et al. 2019; Lin and Kim 2017; Makowski et al. 2010]. A variation on this idea uses different angles of illumination for each color so that the physical filter in Fourier space is not color-specific [Xue et al. 2014]. Frequency multiplexing can also be implemented with white light illumination, which reduces speckle noise at the

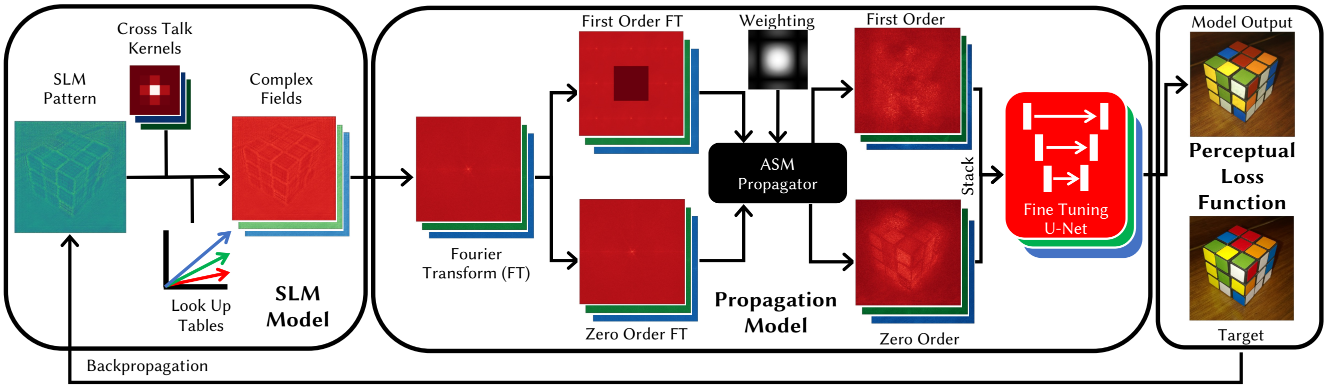

Figure 2: Hologram optimization framework. This figure illustrates the three key components of the simultaneous color optimization framework: an SLM model, a propagation model, and a perceptual loss function. The SLM model maps voltage values to a complex field using a learned cross-talk kernel and a linear lookup table. The complex wavefront from the SLM is then propagated to the sensor plane using a modified version of the model proposed by Gopakumar et al. [2021], which separates the zeroth and first diffraction orders and combines them through a U-Net. The output is then fed into the perceptual loss function, and gradients are calculated using PyTorch's autograd implementation. The SLM voltages are then updated using these gradients. Rubik's cube source image by Iwan Gabovitch (CC BY 2.0).

cost of resolution [Kozacki and Chlipala 2016; Yang et al. 2019]. However, all of these techniques involve filtering in Fourier space, which sacrifices system etendue and requires a bulky $4f$ system.

Depth Division and Bit Division for Simultaneous Color. The prior methods most closely related to our work also use simultaneous RGB illumination over the SLM, maintain the full SLM etendue, and don't require a bulky $4f$ system [Pi et al. 2022]. We refer to the first method as *depth division multiplexing* which takes advantage of the ambiguity between color and propagation distance (explained in detail in Sec. 3.1) and assigns each color a different depth [Makowski et al. 2010, 2008]. After optimizing with a single color for the correct multiplane image, the authors show they can form a full color 2D hologram when illuminating in RGB. However, this approach does not account for wavelength dependence of the SLM response, and since it explicitly defines content at multiple planes, it translates poorly to 3D.

Another similar approach is *bit division multiplexing*, which takes advantage of the extended phase range of LCoS SLMs [Jesacher et al. 2014]. The authors calibrate an SLM lookup-table consisting of phase-value triplets (for RGB) as a function of digital SLM input, and they note that SLMs with extended phase range (up to 10π) can create substantial diversity in the calibrated phase triplets. After pre-optimizing a phase pattern for each color separately, the lookup-table is used on a per-pixel basis to find the digital input that best matches the desired phase for all colors. In our approach, we also use an extended SLM phase range for the same reason, but rather than using a two-step process, we directly optimize the output hologram. This flexible framework also allows us to incorporate a perceptual loss function to further improve perceived image quality.

Algorithms for Hologram Generation. Our work builds on a body of literature applying iterative optimization algorithms to holographic displays. Perhaps most popular is the Gerchberg-Saxton (GS) method [Gerchberg 1972], which is effective and easy to implement, but does not have an explicitly defined loss function, making it challenging to adapt to specific applications. Zhang et al. [2017] and Chakravarthula et al. [2019] were the first to explicitly formulate the hologram generation problem in an optimization framework. This framework has been very powerful, enabling custom loss functions [Choi et al. 2022] and flexible adaptation to new optical configurations [Choi et al. 2021b; Gopakumar et al. 2021]. In particular, perceptual loss functions can improve the perceived image by taking aspects of human vision into account, such as human visual acuity [Kuo et al. 2020], foveated vision [Walton et al. 2022], and sensitivity to spatial frequencies during accommodation [Kim et al. 2022b]. Like these prior works, we use an optimization-based framework which we adapt to account for the wavelength dependence of the SLM; this also enables our new perceptual loss function for color, which is based on visual acuity difference between chrominance and luminance channels.

Camera-Calibration of Holographic Displays. Mismatch between the computational model and physical system creates artifacts in experimental holograms. Recently, several papers have addressed this issue using measurements from a camera in the system for calibration. These approaches use pairs of SLM patterns and camera captures to estimate the learnable parameters in a model, which is then used for offline hologram generation. Learnable parameters can be physically-based [Chakrabarti 2016; Kavaklı et al. 2022; Peng et al. 2020], black box CNNs [Choi et al. 2021a], or a combination of both [Choi et al. 2022]. The choice of learnable parameters effects the ability of the model to match the physical system; we introduce a new parameter for modeling SLM cross talk and tailor the CNN architecture for higher diffraction orders from the SLM.

3 SIMULTANEOUS COLOR HOLOGRAPHY

A holographic image is created by a spatially coherent illumination source incident on an SLM. The SLM imparts a phase delay on the electric field; after light propagates some distance, the intensity of the electric field forms an image. Our goal in this work is to compute a single SLM pattern that simultaneously creates an RGB hologram. For instance, when the SLM is illuminated with a red source, the SLM forms a hologram of the red channel of an image; with a green source the same SLM pattern forms the green channel; and with the blue source it creates the blue channel.

We propose a flexible optimization-based framework (Fig. 2) for generating simultaneous color holograms. We start with a generic model for estimating the hologram from the digital SLM pattern, s, as a function of illumination wavelength, λ:

$$g_\lambda = e^{i\phi_\lambda(s)} \tag{1}$$

$$I_{z,\lambda} = \left| f_{\text{prop}}\left(g_\lambda, z, \lambda\right)\right|^2 . \tag{2}$$

Here, ϕ_λ is a wavelength-dependent function that converts the 8 bit digital SLM pattern to a phase delay, g_λ is the electric field coming off the SLM, f_{prop} represents propagation of the electric field, and $I_{z,\lambda}$ is the intensity a distance z from the SLM.

To calculate the SLM pattern, s, we can solve the following optimization problem

$$\underset{s}{\operatorname{argmin}} \sum_z \mathcal{L}\left(\hat{I}_{z,\lambda_r}, I_{z,\lambda_r}\right) + \mathcal{L}\left(\hat{I}_{z,\lambda_g}, I_{z,\lambda_g}\right) + \mathcal{L}\left(\hat{I}_{z,\lambda_b}, I_{z,\lambda_b}\right), \tag{3}$$

where $\hat{I}$ is the target image, $\mathcal{L}$ is a pixel-wise loss function such as mean-square error, and $\lambda_r, \lambda_g, \lambda_b$ are the wavelengths corresponding to red, green, and blue respectively. Since the model is differentiable, we solve Eq. 3 with gradient descent.

3.1 Color-Depth Ambiguity

A common model for propagating electric fields is Fresnel propagation[1] [Goodman 2005], which can be written in Fourier space as

$$f_{\text{fresnel}}(g, z, \lambda) = \mathcal{F}^{-1}\left\{\mathcal{F}\{g\} \cdot H(z, \lambda)\right\} \tag{4}$$

$$H(z, \lambda) = \exp\left(i\pi\lambda z\left(f_x^2 + f_y^2\right)\right) \tag{5}$$

where $\mathcal{F}$ is a 2D Fourier transform, H is the Fresnel propagation kernel, and f_x, f_y are the spatial frequency coordinates. In Eq. 5, note that λ and z appear together, creating an ambiguity between wavelength and propagation distance.

To see how this ambiguity affects color holograms, consider the case where ϕ_λ in Eq. 1 is independent of wavelength ($\phi_\lambda = \phi$). For example, this would be the case if the SLM had a linear phase range from 0 to 2π at every wavelength. Although this is unrealistic for most off-the-shelf SLMs, it is a useful thought experiment. Note that if ϕ is wavelength-independent, then so is the electric field off the SLM ($g_\lambda = g$). In this scenario, assuming $f_{\text{prop}} = f_{\text{fresnel}}$, the Frensel kernel is the only part of the model affected by wavelength.

[1] Fresnel propagation is the paraxial approximation to the popular angular spectrum method (ASM). Since most commercials SLMs have pixel pitch greater than 3 µm, resulting in a maximum diffraction angle under 5° (well within the small angle approximation), Fresnel and ASM are almost identical for holography.

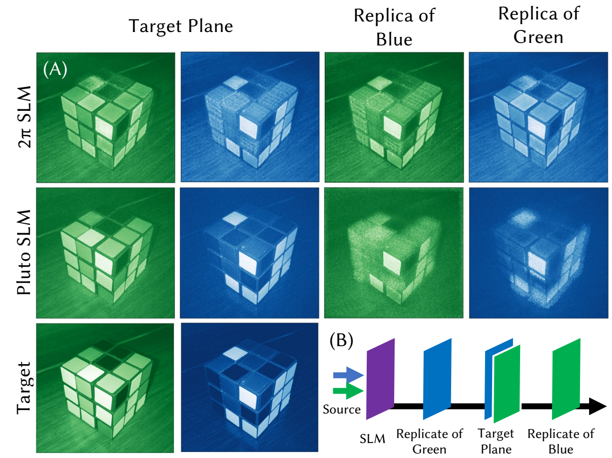

Figure 3: Extended phase range reduces depth replicas in simulation. (A) Using an SLM with a uniform 2π phase range across all channels leads to strong depth replicas (top row), which reduce image quality at the target plane compared to the target (bottom row) and add in-fcous content at depths that should be defocused. By using the extended phase Holoeye Pluto-2.1-Vis-016 SLM (with Red: 2.4π, Green: 5.9π, Blue: 7.4π phase ranges), depth replicas are significantly reduced (middle row), improving the quality of target plane holograms and creating defocused content at other depths. (B) The illumination schematic illustrates the positions of the replicate planes and target plane. See Supplement for the three-color version of this figure. Rubik's cube source image by Iwan Gabovitch (CC BY 2.0).

Now assume that the SLM forms an image at distance z_0 under red illumination. From the ambiguity in the Frensel kernel, we have the following equivalence:

$$H(z_0, \lambda_r) = H\left(\frac{\lambda_g}{\lambda_r} z_0, \lambda_g\right) = H\left(\frac{\lambda_b}{\lambda_r} z_0, \lambda_b\right). \tag{6}$$

This means the *same* image formed in red at z_0 will also appear at $z = z_0\lambda_g/\lambda_r$ when the SLM is illuminated with green and at $z = z_0\lambda_b/\lambda_r$ when the SLM is illuminated with blue. We refer to these additional copies as "depth replicas," and this phenomena is depicted in Fig. 3. Note that depth replicas do not appear in sequential color holography since the SLM pattern optimized for red is never illuminated with the other wavelengths.

If we only care about the hologram at the target plane z_0, then the depth replicas are not an issue. In fact, we can take advantage of the situation for hologram generation: The SLM pattern for an RGB hologram at z_0 is equivalent to the pattern that generates a three-plane red hologram where the RGB channels of the target are each at a different depth ($z0$, $z_0\lambda_r/\lambda_g$, and $z_0\lambda_r/\lambda_b$ for RGB respectively). This is the basis of the depth division multiplexing approach of Makowski et al. [2010, 2008], where the authors optimize for this three-plane hologram in red, then illuminate in RGB. Although this makes the assumption that ϕ does not depend on λ, this connection between simultaneous color and multi-plane holography suggests simultaneous color should be possible for a single plane, since multi-plane holography has been successfully demonstrated in prior work.

However, the ultimate goal of holography is to create 3D imagery, and the depth replicas could prevent us from placing content arbitrarily over the 3D volume. In addition, in-focus images can appear at depths that should be out-of-focus, which may prevent the hologram from successfully driving accommodation [Kim et al. 2022b]. We propose taking advantage of SLMs with extended phase range to mitigate the effects of depth replicas.

3.2 SLM Extended Phase Range

In general, the phase ϕ_λ of the light depends on its wavelength, which was not considered in Sec. 3.1. Perhaps the most popular SLM technology today is LCoS, in which rotation of birefringent LC molecules causes a change in refractive index. The phase of light traveling through the LC layer is delayed by

$$\phi_\lambda = \frac{2\pi d}{\lambda} n(s, \lambda), \tag{7}$$

where d is the thickness of the LC layer, and its refractive index, n, is controlled with the digital input s. n also depends on λ due to dispersion [Jesacher et al. 2014].

The wavelength dependence of ϕ_λ presents an opportunity to reduce or remove the depth replicas. Even if the propagation kernel H is the same for several (λ, z) pairs, if the phase, and therefore the electric field off the SLM, changes with λ, then the output image intensity at the replica plane will also be different. As the wavelength-dependence of ϕ_λ increases, the replicas are diminished.

We can quantify the degree of dependence on λ by looking at the derivative $d\phi/d\lambda$ which informs us that larger n will give λ more influence on the SLM phase. However, the final image intensity depends only on relative phase, not absolute phase; therefore, for the output image to have a stronger dependence on λ, we desire larger $\Delta n = n_{\max} - n_{\min}$. In addition, $d\phi/d\lambda$ increases with $-dn/d\lambda$, suggesting that more dispersion is helpful for simultaneous color. Although $d\phi/d\lambda$ also depends on the absolute value of λ, we have minimal control over this parameter since there are limited wavelengths corresponding to RGB. In summary, this means we can reduce depth replicas in simultaneous color with larger phase range on the SLM and higher dispersion.

However, there is a trade-off: As the range of phase increases, the limitations of the bit depth of the SLM become more noticeable, leading to increased quantization errors. We simulate the effect of quantization on hologram quality and find that PSNR and SSIM are almost constant for 6 bits and above (see Supplement). This suggests that each 2π range should have at least 6 bits of granularity. Therefore, we think that using a phase range of around 8π for an 8-bit SLM will be the best balance between replica reduction and maintaining accuracy for hologram generation. Figure 3 simulates the effect of extended phase range on depth replica removal. While holograms were calculated on RGB images, only two color channels are shown for simplicity (see Supplement for full color version).

In the first row of Fig. 3, we simulate an SLM with no wavelength dependence to ϕ (i.e. 0 - 2π phase range for each color). Consequently, perfect copies appear at the replica planes. In the second row, we simulate using the specifications from an extended phase range SLM (Holoeye Pluto-2.1-Vis-016), which has 2.4π range in red, 5.9π range in green, and 7.4π range in blue demonstrating that replicas are substantially diminished with an extended phase range. By reducing the depth replicas, the amount of high frequency out-of-focus light at the sensor plane is reduced, leading to improved hologram quality.

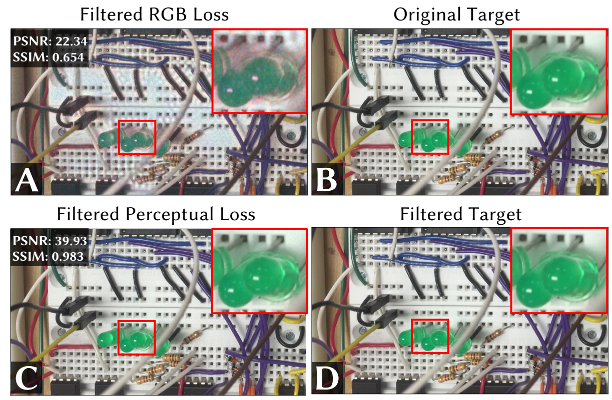

Figure 4: Perceptual loss improves color fidelity and reduces noise in simulation. The first column of this figure depicts simulated holograms optimized with an RGB loss function (A) and our perceptual loss function (C). The same filters for the perceptual loss function then were applied to both of these simulated holograms as well as the target image. Image metrics were calculated between the filtered holograms and the filtered target image (D). All image metrics are better for the perceptually optimized hologram (C). One should also note that the filtered target (D) and original target (B) are indistinguishable suggesting our perceptual loss function only removes information imperceptible by the human visual system.

3.3 Perceptual Loss Function

Creating an RGB hologram with a single SLM pattern is an overdetermined problem as there are 3× more output pixels than degrees of freedom of the SLM. As a result, it may not be possible to exactly match the full RGB image, which can result in color deviations and de-saturation. To address this, we take advantage of color perception in human vision. There's evidence that the human visual systems converts RGB images into a luminance channel (a grayscale image) and two chrominance channels, which contain information about the color [Wandell 1995]. The visual system is only sensitive to high resolution features in the luminance channel, so the chrominance channels can be lower resolution with minimal impact on the perceived image [Wandell 1995]. This observation is used in JPEG compression [Pennebaker and Mitchell 1992] and subpixel rendering [Platt 2000], but to our knowledge, it has never been applied to holographic displays. By allowing the unperceived high frequency chrominance and extremely high frequency luminance features to be unconstrained, we can better use the the degrees of freedom on the SLM to faithfully represent the rest of the image.

Our flexible optimization framework allows us to easily change the RGB loss function in Eq. 3 to a perceptual loss. For each depth, we transform the RGB intensities of both $\hat{I}$ (the target image) and I (the simulated hologram) into opponent color space as follows:

$$\begin{aligned} O_1 &= 0.299 \cdot I_{\lambda_r} + 0.587 \cdot I_{\lambda_g} + 0.114 \cdot I_{\lambda_b} \\ O_2 &= I_{\lambda_r} - I_{\lambda_g} \\ O_3 &= I_{\lambda_b} - I_{\lambda_r} - I_{\lambda_g} \end{aligned} \tag{8}$$

where O_1 is the luminance channel, and O_2, O_3 are the red-green and blue-yellow chrominance channels, respectively. We can then update Eq. 3 to

$$\begin{aligned} \underset{s}{\operatorname{argmin}} \sum_z \Big[\mathcal{L}\left(\hat{O}_1 * k_1, O_1 * k_1\right) + \mathcal{L}\left(\hat{O}_2 * k_2, O_2 * k_2\right) + \\ \mathcal{L}\left(\hat{O}_3 * k_3, O_3 * k_3\right) \Big], \end{aligned} \tag{9}$$

where $*$ represents a 2D convolution with a low pass filter ($k_1 \ldots k_3$) for each channel in opponent color space . $\hat{O}_i$ and O_i are the i-th channel in opponent color space of $\hat{I}$ and I, respectively. In order to mimic the contrast sensitivity functions of the human visual system, we implement filters in the Fourier domain by applying a low-pass filter of 45% of the width of Fourier space to the chrominance channels (O_2, O_3) and a filter of 75% of the width of Fourier space to the luminance channel (O_1). In a system with a 36.6 mm focal length eye piece, these cutoffs correspond to 30 cycles/deg and 18 cycles/deg in luminance and chrominance respectively, approximately matched to human vision [Mullen 1985].

By de-prioritizing high frequencies in chrominance and extremely high frequencies in luminance, the optimizer is able to better match the low frequency color. This low frequency color is what is perceivable by the human visual system. Figure 4 highlights the improvement provided by our perceptual loss function, comparing perceptually filtered versions of simulated holograms generated using an RGB loss function (Fig 4A) and our perceptual loss function (Fig 4B). The original unfiltered target image (Fig 4C) and the perceptually filtered target image (Fig 4D) are nearly indistinguishable, indicating that our perceptual filter choices align well with the human visual system. The PSNR and SSIM values are higher for the perceptually optimized hologram (Fig. 4C), which is visually less noisy with better color fidelity. This suggests that the loss function has effectively shifted most of the error into imperceptible regions of the opponent color space. We see an average PSNR increase of 6.4 dB and average increase of 0.266 in SSIM across a test set of 294 images.

3.4 Simulation Comparisons

We compare the performance of our method to the depth and bit division approaches [Jesacher et al. 2014; Makowski et al. 2010], which, like our method, use only a single SLM, make use of the full SLM space-time-bandwidth, and contain no bulky optics or filters (see Supplement for implementation details). The holograms simulated with depth and bit division, shown in Fig. 5, are much noisier and have lower color fidelity than our proposed method. Depth division has the worst color fidelity due to to the replica planes discussed in Sec. 3.1 contributing defocused light at the target plane. Our approach directly optimizes the simultaneous color hologram using our perceptual loss function, resulting in less noise and better color fidelity compared to these other indirect optimization approaches.

4 CAMERA-CALIBRATED MODEL

We've demonstrated that our algorithm can generate simultaneous color holograms in simulation. However, experimental holograms frequently do not match the quality of simulations due to mismatch between the physical system and the model used in optimization (Eqs. 1, 2). Therefore, to demonstrate simultaneous color experimentally, we need to calibrate the model to the experimental system.

To do this, we design a model based on our understanding of the system's physics, but we include several learnable parameters representing unknown elements. To fit the parameters, we capture a dataset of SLM patterns and camera captures and use gradient descent to estimate the learnable parameters based on the dataset. Next we explain the model which is summarized in Fig. 2.

Lookup Table. A key element in our optimization is ϕ_λ which converts the digital SLM input into the phase coming off the SLM. It's important this function accurately matches the behavior of the real SLM. Many commercial SLMs ship with a lookup-table (LUT) describing ϕ_λ; however, this LUT is generally only calibrated at a few discrete wavelengths. Consequently, we learn a LUT for each color channel's wavelength as part of the model. Based on a pre-calibration of the LUT using the approach of Yang et al. [2015], we observe the LUT is close to linear; we therefore parameterize the LUT with a linear model to encourage physically realistic solutions.

SLM Crosstalk. SLMs are usually modeled as having a constant phase over each pixel with sharp transitions at boundaries. However, in LCoS SLMs, elastic forces in the LC layer prevent sudden spatial variations, and the electric field that drives the pixels changes gradually over space. As a result, LCoS SLMs suffer from crosstalk, also called field fringing, in which the phase is blurred [Apter et al. 2004; Moser et al. 2019; Persson et al. 2012]. We model crosstalk with a convolution on the SLM phase. Combined with our linear LUT described above, we can describe the phase off the SLM as

$$\phi_\lambda(s) = k_{\text{xt}} * (a_1 \cdot s + a_2) \tag{10}$$

where a_1, a_2 are the learn parameters of the LUT, and k_{xt} is a learned 5×5 convolution kernel representing crosstalk. Separate values of these parameters are learned for each color channel.

Propagation with Higher Diffraction Orders. The discrete pixel structure of the SLM creates higher diffraction orders that are not modeled by ASM or Fresnel propagation. With the use of a $4f$ system, a physical aperture at the Fourier plane of the SLM can be used to block higher orders. However, this adds significant size to the optical system, reducing the practicality for head-mounted displays. Therefore, we chose to avoid additional lenses after the SLM and instead account for higher orders in the propagation model.

We adapt the higher-order angular spectrum model (HOASM) of Gopakumar et al. [2021]. The zero order diffraction, $G_0(f_x, f_y)$, and first order diffraction, G_1, patterns are propagated with ASM to the plane of interest independently. Then the propagated fields are stacked and passed into a U-net, which combines the zero and

Figure 5: Comparison of bit division, depth division and our method of simultaneous color holography in simulation. Bit division (Col. 1) is noisier than our method (Col. 3) but achieves comparable color fidelity, although more washed out. The depth division method (Col. 2) is also noisier than our method and has inferior color fidelity. Our method matches the target image (Col. 4) well. Our method uses our perceptual loss function and a high order angular spectrum propagation model with no learned components. Further implementation details for each method are available in the supplement.

first orders and returns the image intensity:

$$f_{\mathrm{ASM}}(G, z) = \mathcal{F}^{-1}\left\{G \cdot H_{\mathrm{ASM}}(z)\right\} \tag{11}$$

$$I_z = \mathrm{Unet}\left(f_{\mathrm{ASM}}(G_0, z),\ f_{\mathrm{ASM}}(G_1, z)\right), \tag{12}$$

where $H_{\mathrm{ASM}}(z)$ is the ASM kernel. The U-Net architecture is detailed in the supplement; a separate U-net for each color is learned from the data. The U-Net helps to address any unmodeled aspects of the system that may affect the final hologram quality such as source polarization, SLM curvature, and beam profiles, and the U-net better models superposition of higher orders, allowing for more accurate compensation in SLM pattern optimization.

5 IMPLEMENTATION

Experimental Setup. Our system starts with a fiber-coupled RGB source ($\lambda_r = 636$ nm, $\lambda_g = 512$ nm, $\lambda_b = 453$ nm), collimated with a 400 mm lens. The beam is aligned using two mirrors, passes through a linear polarizer and beamsplitter, reflects off the SLM (Holoeye-2.1-Vis-016), and passes through the beamsplitter a second time before directly hitting the color camera sensor with Bayer filter (FLIR GS3-U3-123S6C). As seen in Fig. 9, there's no $4f$ system between the SLM and camera, which allows the setup to be compact, but requires modeling of higher diffraction orders. The camera sensor is on a linear motion stage, enabling a range of propagation distances from $z = 80$ mm to $z = 130$ mm.

For our source, we use a superluminescent light emitting diode (SLED, Exalos EXC250011-00) rather than a laser due to its lower coherence, which has been demonstrated to reduce speckle in holographic displays [Deng and Chu 2017]. Although previous work showed state-of-the-art image quality by modeling the larger bandwidth of the SLED as a summation of coherent sources [Peng et al. 2021], we found the computational cost to be prohibitively high for our application due to GPU memory constraints. We achieved sufficient image quality while assuming a fully coherent model, potentially due to the U-net which is capable of simulating the additional blur we expect from a partially coherent source.

Our experimental system directly forms the hologram on a bare sensor, but for a human-viewable system, an eyepiece is necessary between the image plane and the user's eye. See Supplement for details on how the eyepiece effects the depth replicas.

Calibration Procedure. We learn parameters in our model (Eqs. 10 - 12) using a dataset captured on the experimental system. We pre-calculate 882 SLM patterns from a personally collected dataset of images using the ASM propagation model. Each SLM pattern is captured in 10 mm increments from $z = 90$ mm to 120 mmThe camera data is debayered and an affine transform is applied to align the image with the SLM (see Supplement for details). Model fitting is implemented in PyTorch using an L1 loss function between the model output and camera capture. To account for the camera color balance, we additionally learn a 3×3 color calibration matrix. We train until convergence, which is typically reached in 2-3 days on Nvidia A6000 GPU.

Hologram Generation. After training, we can generate holograms by solving Eq. 9 using the trained model for $I_{z,\lambda}$, implemented with PyTorch's native autodifferentiation. The SLM pattern, s, is constrained to the range where the LUT is valid (for example, 0 - 255); the values outside that range are wrapped after every optimization step. On the Nvidia A6000 GPU, it takes about two minutes to optimize a 2D hologram. Computation time for the optimization of a 3D hologram scales proportionally to the number of depth planes.

6 EXPERIMENTAL RESULTS

2-Dimensional Holograms. We validate our simulation results by capturing holograms in experiment. For simultaneous color, the SLM patterns were optimized for a propagation distance of 120 mm using our perceptual loss function described in Section 3.3. A white border was added to each target image to improve the color fidelity by encouraging a proper white balance. After each hologram is captured, debayering is performed and a homography is applied to map from camera space to SLM space.

Figure 6 compares the simultaneous color capture using a single frame (B) to sequential color using 3 frames (C). Unlike the simultaneous color version, which was captured in one shot with RGB illumination, the sequential color was captured with only the red light source (due to a failure of the green channel in the SLED), and the correct color was assigned in software. Although the sequential captures are higher contrast than our simultaneous results, we'd like to emphasize that our approach uses 3× fewer degrees of freedom and can still produce full color images. In addition, the simulation output from our model (D) shows color fidelity on par with the sequential capture; the difference between the simulation output and experimental capture can be attributed to model mismatch. This suggests improvements to the calibration pipeline could enable experimental results with the quality of the simultaneous model.

3-Dimensional Holograms. A major appeal of holography is the ability to solve the vergence-accommodation conflict, so we also validate our method for 3D scenes. A 4-plane focal stack was rendered with 0.5 pixels blur radius per millimeter depth. Holograms were captured at distance from 90 mm to 120 mm in 10 mm increments. The results are displayed in Fig. 7, and once again pseudo-color sequential images (B), which use 3× the number of frames, are shown for comparison. Although model mismatch creates some color shift in the experimental captures (C), the simultaneous model output (D) shows what the results could look like with improved calibration. We note that 3D hologram generation is not as well-posed as 2D; despite this, our results demonstrate the ability to form 3D color holograms with natural defocus blur from a single SLM frame.

7 DISCUSSION

While our method improves hologram quality for simultaneous illumination and is compatible with VR/AR displays, it does have limitations. First, our method is not equally effective for all images. Natural images with high levels of texture work best, as they have similarly structured color channels and contain high frequency color information that is perceptually suppressible by our loss function. Images with large flat areas may exhibit noticeable artifacts due to the more difficult task of determining an SLM pattern that produces 3 largely unique holograms (see Supplement Fig. S4).

SLMs with large phase range can be slower than their short phase range counterparts. Although our SLM has 7.4π phase range in blue, we show in the Supplement that we can achieve reasonable quality with only a 4π range, opening the possibility for simultaneous color with a wider variety of SLMs.

Calculating a single SLM pattern for a 2D image using an Nvidia A6000 takes minutes with our method, inhibiting real-time displays. Neural nets can generate SLM patterns in real-time while retaining quality, suggesting a potential future solution for simultaneous color holography [Eybposh et al. 2020; Shi et al. 2021; Yang et al. 2022].

In summary, we developed a framework for high-quality color holograms using simultaneous RGB illumination in a compact setup, featuring a camera-calibrated, differentiable model and custom loss functions.

REFERENCES

Boris Apter, Uzi Efron, and Eldad Bahat-Treidel. 2004. On the fringing-field effect in liquid-crystal beam-steering devices. *Applied optics* 43, 1 (2004), 11–19.

Ayan Chakrabarti. 2016. Learning sensor multiplexing design through back-propagation. *Advances in Neural Information Processing Systems* Nips (2016), 3089–3097.

Praneeth Chakravarthula, Seung-Hwan Baek, Ethan Tseng, Andrew Maimone, Grace Kuo, Florian Schiffers, Nathan Matsuda, Oliver Cossairt, Douglas Lanman, and Felix Heide. 2022. Pupil-aware Holography. *arXiv preprint arXiv:2203.14939* (2022).

Praneeth Chakravarthula, Yifan Peng, Joel Kollin, Henry Fuchs, and Felix Heide. 2019. Wirtinger holography for near-eye displays. *ACM Transactions on Graphics* 38, 6 (2019). https://doi.org/10.1145/3355089.3356539

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* 39, 6 (2020). https://doi.org/10.1145/3414685.3417846

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. (2022), 1–9. https://doi.org/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* 40, 6 (2021). https://doi.org/10.1145/3478513.3480542

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.

Yuanbo Deng and Daping Chu. 2017. Coherence properties of different light sources and their effect on the image sharpness and speckle of holographic displays. *Scientific Reports* 7, 1 (2017), 1–12. https://doi.org/10.1038/s41598-017-06215-x

Peter Duerr, Andreas Neudert, Christoph Hohle, Hagen Stolle, Johannes Pleikies, and Hagen Sahm. 2021. MEMS Spatial Light Modulators for Real Holographic 3D Displays. In *MikroSystemTechnik Congress 2021; Congress.* 1–4.

M Hossein Eybposh, Nicholas W Caira, Praneeth Chakravarthula, Mathew Atisa, and Nicolas C Pégard. 2020. High-speed computer-generated holography using convolutional neural networks. In *Optics and the Brain.* Optica Publishing Group, BTu2C–2.

Ralph W Gerchberg. 1972. A practical algorithm for the determination of plane from image and diffraction pictures. *Optik* 35, 2 (1972), 237–246.

Joseph W Goodman. 2005. *Introduction to fourier optics.* Roberts & Company Publishers.

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. https://doi.org/10.1364/ol.442851

Tobias Haist and Wolfgang Osten. 2015. Holography using pixelated spatial light modulators—part 1: theory and basic considerations. *Journal of Micro/Nanolithography, MEMS, and MOEMS* 14, 4 (2015), 041310.

Changwon Jang, Kiseung Bang, Gang Li, and Byoungho Lee. 2018. Holographic near-eye display with expanded eye-box. *ACM Transactions on Graphics (TOG)* 37, 6 (2018), 1–14.

Alexander Jesacher, Stefan Bernet, and Monika Ritsch-Marte. 2014. Colour hologram projection with an SLM by exploiting its full phase modulation range. *Optics Express* 22, 17 (8 2014), 20530. https://doi.org/10.1364/oe.22.020530

Koray Kavaklı, Liang Shi, Hakan Ürey, Wojciech Matusik, and Kaan Akşit. 2023. HoloHDR: Multi-color Holograms improve Dynamic Range. *arXiv preprint arXiv:2301.09950* (2023).

Koray Kavaklı, Hakan Urey, and Kaan Akşit. 2022. Learned holographic light transport. *Applied Optics* 61, 5 (2022), B50–B55.

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.* 1–9.

Tomasz Kozacki and Maksymilian Chlipala. 2016. Color holographic display with white light LED source and single phase only SLM. *Optics Express* 24, 3 (feb 2016), 2189. https://doi.org/10.1364/oe.24.002189

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), 66–1.

Byounghyo Lee, Dongyeon Kim, Seungjae Lee, Chun Chen, and Byoungho Lee. 2022. High-contrast, speckle-free, true 3D holography via binary CGH optimization. *Scientific reports* 12, 1 (2022), 1–12.

Byounghyo Lee, Dongheon Yoo, Jinsoo Jeong, Seungjae Lee, Dukho Lee, and Byoungho Lee. 2020. Wide-angle speckleless DMD holographic display using structured illumination with temporal multiplexing. *Optics letters* 45, 8 (2020), 2148–2151.

Gang Li, Dukho Lee, Youngmo Jeong, Jaebum Cho, and Byoungho Lee. 2016. Holographic display for see-through augmented reality using mirror-lens holographic optical element. *Optics letters* 41, 11 (2016), 2486–2489.

Shu-Feng Lin, Hong-Kun Cao, and Eun-Soo Kim. 2019. Single SLM full-color holographic three-dimensional video display based on image and frequency-shift multiplexing. *Opt. Express* 27, 11 (may 2019), 15926–15942. https://doi.org/10.1364/OE.27.015926

Shu-Feng Lin and Eun-Soo Kim. 2017. Single SLM full-color holographic 3-D display based on sampling and selective frequency-filtering methods. *Opt. Express* 25, 10 (may 2017), 11389–11404. https://doi.org/10.1364/OE.25.011389

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.

Michal Makowski, Izabela Ducin, Karol Kakarenko, Jaroslaw Suszek, Maciej Sypek, Andrzej Kolodziejczyk, P.-H Yao, C.-H Chen, J.-N 4 Kuo, and H.-W Wu. 2011. *Simple holographic projection in color.* Technical Report. 7–9 pages. www.osiris-project.eu

Michal Makowski, Izabela Ducin, Maciej Sypek, Agnieszka Siemion, Andrzej Siemion, Jaroslaw Suszek, and Andrzej Kolodziejczyk. 2010. *Color image projection based on Fourier holograms.* Technical Report 8. 1227 pages.

Michal Makowski, Mciej Sypek, Izabela Ducin, Agnieszka Fajst, Andrzej Siemion, Jaroslaw Suszek, and Andrzej Kolodzijczyk. 2009. Experimental evaluation of a full-color compact lensless holographic display. *Opt. Express* 17, 23 (2009), 20840–20846. https://doi.org/10.1364/OE.17.020840

Michal Makowski, Maciej Sypek, and Andrzej Kolodziejczyk. 2008. *Colorful reconstructions from a thin multi-plane phase hologram.* Technical Report.

Simon Moser, Monika Ritsch-Marte, and Gregor Thalhammer. 2019. Model-based compensation of pixel crosstalk in liquid crystal spatial light modulators. *Optics express* 27, 18 (2019), 25046–25063.

Kathy T Mullen. 1985. The contrast sensitivity of human colour vision to red-green and blue-yellow chromatic gratings. *The Journal of physiology* 359, 1 (1985), 381–400.

Hirotaka Nakayama, Naoki Takada, Yasuyuki Ichihashi, Shin Awazu, Tomoyoshi Shimobaba, Nobuyuki Masuda, and Tomoyoshi Ito. 2010. *Real-time color electroholography using multiple graphics processing units and multiple high-definition liquid-crystal display panels.* Technical Report.

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). https://doi.org/10.1126/sciadv.abg5040

Yifan Peng, Suyeon Choi, Nitish Padmanaban, and Gordon Wetzstein. 2020. Neural holography with camera-in-the-loop training. *ACM Transactions on Graphics* 39, 6 (11 2020). https://doi.org/10.1145/3414685.3417802

William B Pennebaker and Joan L Mitchell. 1992. *JPEG: Still image data compression standard.* Springer Science & Business Media.

Martin Persson, David Engström, and Mattias Goksör. 2012. Reducing the effect of pixel crosstalk in phase only spatial light modulators. *Optics express* 20, 20 (2012), 22334–22343.

Dapu Pi, Juan Liu, and Yongtian Wang. 2022. Review of computer-generated hologram algorithms for color dynamic holographic three-dimensional display. https://doi.org/10.1038/s41377-022-00916-3

John C Platt. 2000. Optimal filtering for patterned displays. *IEEE Signal Processing Letters* 7, 7 (2000), 179–181.

Bernhard E Riecke, Hans-Günther Nusseck, and Jörg Schulte-Pelkum. 2006. Selected technical and perceptual aspects of virtual reality displays. (2006).

Steve Serati, Xiaowei Xia, Owais Mughal, and Anna Linnenberger. 2003. High-resolution phase-only spatial light modulators with submillisecond response. In *SPIE Proceedings*, Vol. 5106. Boulder Nonlinear Systems, Inc., 450 Courtney Way, #107, Lafayette, Colorado 80026.

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.

Atsushi Shiraki, Naoki Takada, Masashi Niwa, Yasuyuki Ichihashi, Tomoyoshi Shimobaba, Nobuyuki Masuda, Tomoyoshi Ito, P S Hilaire, S A Benton, M Lucente, M L Jepsen, J Kollin, H Yoshikawa, and J Underkoffler. 2009. *Simplified electroholographic color reconstruction system using graphics processing unit and liquid crystal display projector References and links.* Technical Report. http://www.opticsinfobase.org/oe/abstract.cfm?id=81092.

JMP Van Waveren. 2016. The asynchronous time warp for virtual reality on consumer hardware. In *Proceedings of the 22nd ACM Conference on Virtual Reality Software and Technology.* 37–46.

David R Walton, Koray Kavaklı, Rafael Kuffner Dos Anjos, David Swapp, Tim Weyrich, Hakan Urey, Anthony Steed, Tobias Ritschel, and Kaan Akşit. 2022. Metameric Varifocal Holograms. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces (VR).* IEEE, 746–755.

Brian A Wandell. 1995. *Foundations of vision.* Sinauer Associates.

Gaolei Xue, Juan Liu, Xin Li, Jia Jia, Zhao Zhang, Bin Hu, and Yongtian Wang. 2014. Multiplexing encoding method for full-color dynamic 3D holographic display. *Opt. Express* 22, 15 (Jul 2014), 18473–18482. https://doi.org/10.1364/OE.22.018473

Daeho Yang, Wontaek Seo, Hyeonseung Yu, Sun Il Kim, Bongsu Shin, Chang-Kun Lee, Seokil Moon, Jungkwuen An, Jong-Young Hong, Geeyoung Sung, et al. 2022. Diffraction-engineered holography: Beyond the depth representation limit of holographic displays. *Nature Communications* 13, 1 (2022), 1–11.

Lei Yang, Jun Xia, Chenliang Chang, Xiaobing Zhang, Zhiming Yang, and Jianhong Chen. 2015. Nonlinear dynamic phase response calibration by digital holographic microscopy. *Applied optics* 54, 25 (2015), 7799–7806.

Xin Yang, Ping Song, HongBo Zhang, and Qiong-Hua Wang. 2019. Full-color computer-generated holographic near-eye display based on white light illumination. *Optics Express* 27, 26 (2019), 38236–38249.

Fahri Yaraş, Hoonjong Kang, and Levent Onural. 2009. *Real-time phase-only color holographic video display system using LED illumination.* Technical Report.

Kun Yin, En Lin Hsiang, Junyu Zou, Yannanqi Li, Zhiyong Yang, Qian Yang, Po Cheng Lai, Chih Lung Lin, and Shin Tson Wu. 2022. Advanced liquid crystal devices for augmented reality and virtual reality displays: principles and applications. *Light: Science and Applications* 11, 1 (2022). https://doi.org/10.1038/s41377-022-00851-3

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.

Zichen Zhang, Zheng You, and Daping Chu. 2014. Fundamentals of phase-only liquid crystal on silicon (LCOS) devices. *Light: Science & Applications* 3, 10 (2014), e213–e213.

David J Zielinski, Hrishikesh M Rao, Mark A Sommer, and Regis Kopper. 2015. Exploring the effects of image persistence in low frame rate virtual environments. In *2015 IEEE Virtual Reality (VR).* IEEE, 19–26.

Figure 6: Experimentally captured 2D holograms. For each target image (A), we show (B) the experimental capture with sequential pseudo-color, (C) our experimental capture with full simultaneous color, and (D) the simulated model output for simultaneous color. Recall that our simultaneous color results (C) use 3× fewer degrees of freedom than the sequential capture (B). Although some color fidelity is lost in experiment (C), the simulated model output (D) shows good color quality, demonstrating that accurate color is possible with our method and improvements to the calibration.

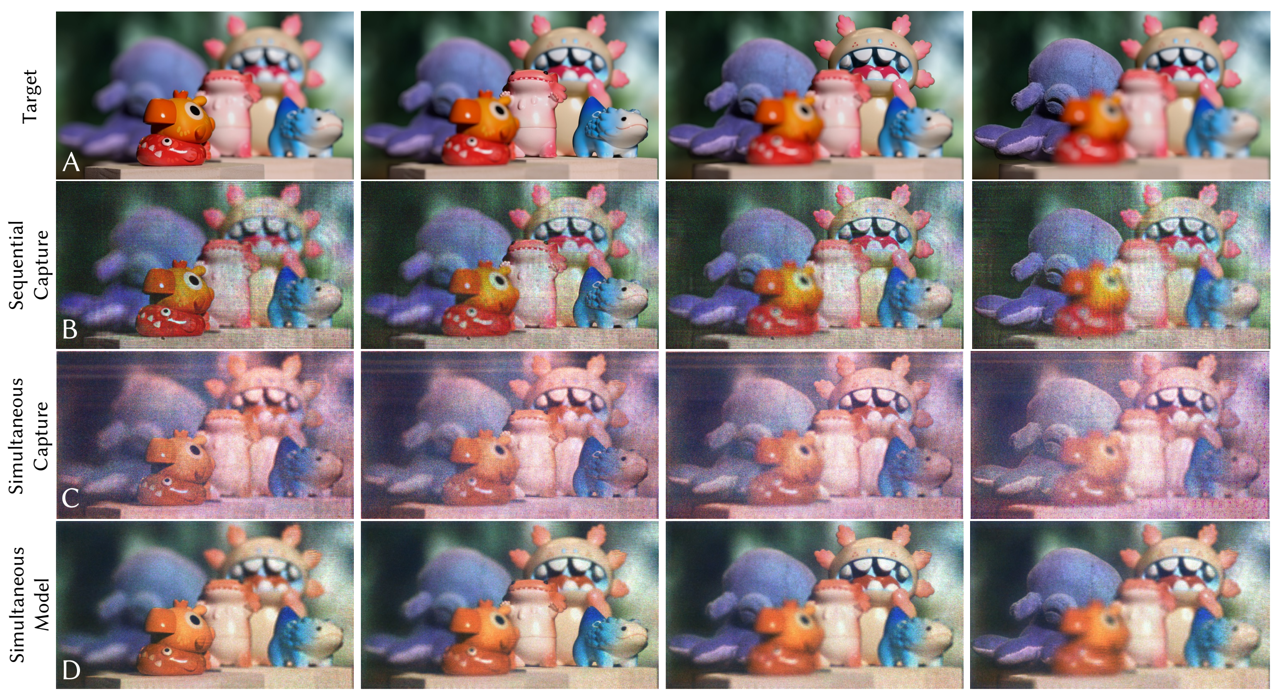

Figure 7: Experimentally captured focal stack. This figure displays a focal stack, with the target shown in (A), captured from 90 mm **to** 120 mm **in** 10 mm **increments. We compare (B) the sequential pseudo-color experimental capture with (C) the experimental capture of the simultaneous full color hologram and (D) the simulated model output for simultaneous color. Although model mismatch creates some deviations between the simultaneous capture (C) and the target (A), the simulated model (D) is representative of the color fidelity we expect from our method with improvements to the system calibration.**

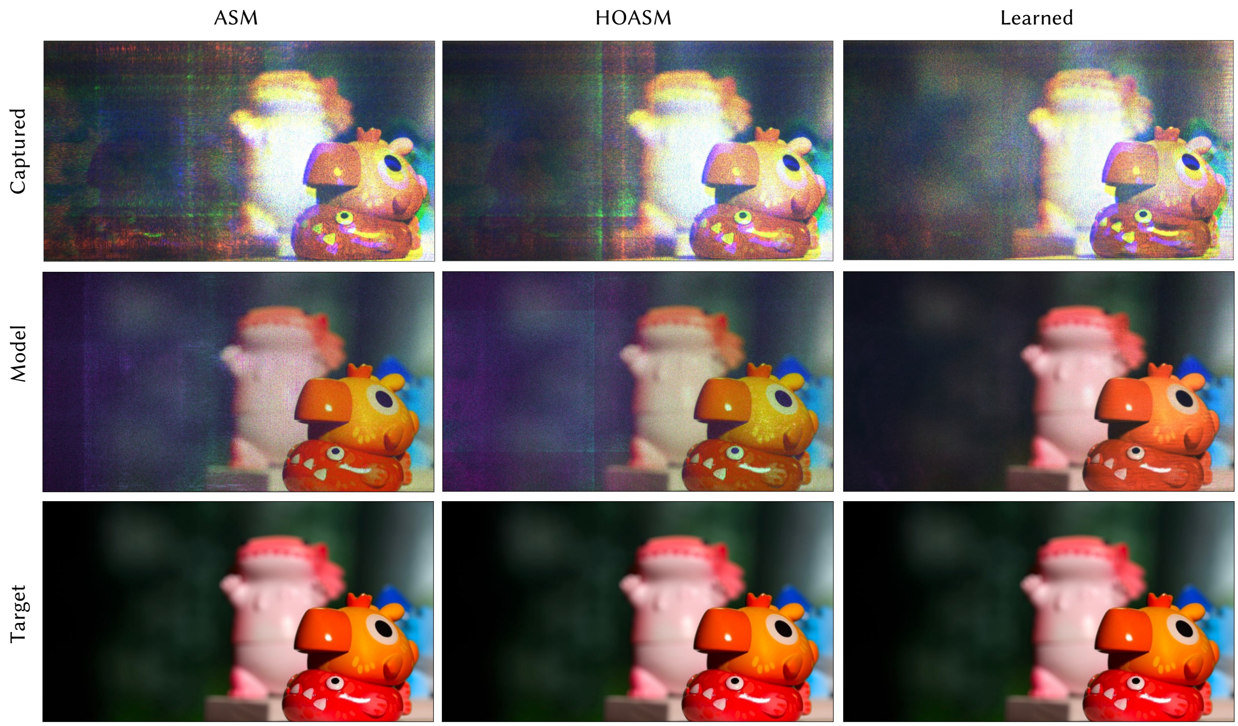

Figure 8: Comparison of different propagation methods for suppressing higher diffraction orders. The first column shows the results obtained using the traditional angular spectrum method (ASM) which doesn't model higher diffraction orders. The second column shows the results obtained using HOASM which reduces the visibility of higher orders but fails to completely suppress them. The third column shows the results obtained using our proposed learned propagation method that includes a U-net, which largely suppresses the higher diffraction orders and results in a hologram with the fewest artifacts, suggesting the learned propagation model best matches the physical propagation.

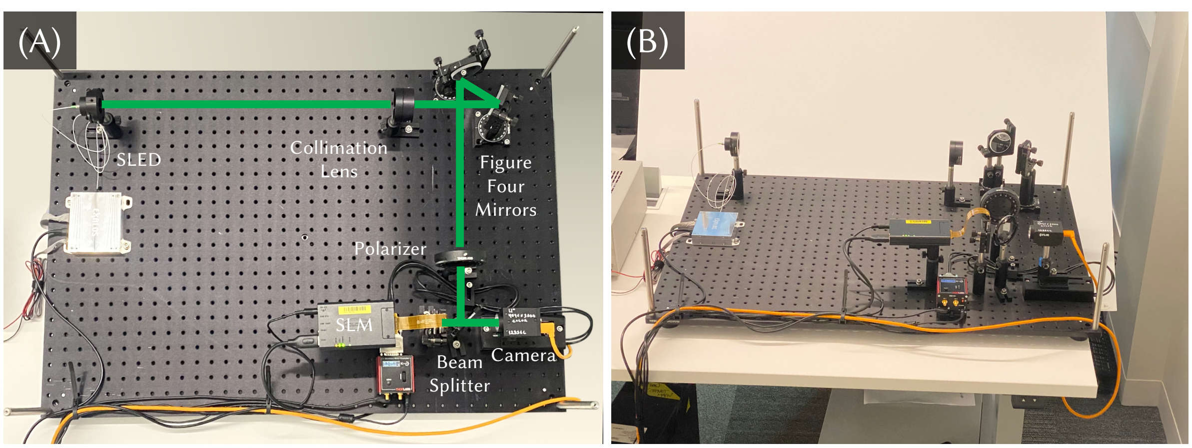

Figure 9: Experimental setup A top view of our system with labeled components and an approximate beam path drawn in green is depicted in (A). A side-view of the system is provided by (B). Note that the hologram is formed directly on the bare camera sensor with no lens or eyepiece between. This configuration allows us to validate our method, but for a human-viewable system, an eyepiece must be added between the hologram plane and the user's eye.