Perceptually guided Computer-Generated Holography

Kaan Akşit, University College London (UK).

Abstract

Inventing immersive displays that can attain realism in visuals is a long standing quest in the optics, graphics and perception fields. As holographic displays can simultaneously address various depth levels, experts from industry and academia often pitch these holographic displays as the next-generation display technology that could lead to such realism in visuals. However, holographic displays demand high computational complexity in image generation pipelines and suffer from visual quality-related issues.

This talk will describe our research efforts to combine visual perception related findings with Computer-Generated Holography (CGH) to achieve realism in visuals and derive CGH pipelines that can run at interactive rates (above 30 Hz). Specifically, I will explain how holographic displays could effectively generate three-dimensional images with good image quality and how these images could be generated to match the needs of human visual perception in resolution and statistics. Furthermore, I will demonstrate our CGH methods running at interactive rates with the help of learning strategies. As a result, we provide a glimpse into a potential future where CGH helps to replace two-dimensional images generated on today's displays with authentic three-dimensional visuals that are perceptually realistic.

Introduction

Today's conventional displays can paint two-dimensional pixelated images. However, the scenes humans view in their day-to-day lives can be viewed from various perspectives, leading to a continuous viewing frustum with no pixelation. Thus, replicating the three-dimensional visuals could be vital in attaining realism in next-generation displays [1].

Computer-Generated Holography (CGH) [2] offers a computational method to reconstruct light fields in three dimensions in a programmable manner. CGH generates these light fields by interfering with light from diffracted pixels on some image plane. Generated images are then a result of a superposition of continuous-wave functions. Unlike conventional displays, these images are not pixelated and are focused on various depth levels. Thus, CGH is often pitched by industry and academia as the enabling technology in next-generation displays that could lead to lifelike three-dimensional images [3].

We take this opportunity to introduce our CGH related research works aiming for perceptual realism in visuals. We believe that there are three components to be discussed within this scope. These components are (1) an accurate definition of a standard holographic display, (2) light transport models that could account for the imperfections in holographic displays, and (3) perceptual guidance that could potentially help to ease the computational complexity of hologram generation routines in the future. Furthermore, I will connect on optimising phase-only holograms that lead to three-dimensional images using CGH.

Overview of State of the art

CGH techniques and holographic displays have seen a tremendous amount of development in the last two to three years time frame. Here we briefly review the most noticeable developments and advancements in the field.

Standard holographic display. Inspired by the holography pipeline from Maimone et al. [4], the work by Shi et al. [5] introduces a new Convolutional Neural Network (CNN) that could generate phase-only holograms with occlusion support. This CNN is unique as it bypasses classical rendering pipelines and generates high-quality three-dimensional images at various depths. The work by Shi et al. [5] could be interpreted as the state of the art when it comes to image quality in CGH. The standard display definition that we will be provided in the next section follows a similar implementation to these works [4, 5].

Modelling holographic displays. Researchers have recently looked into advancing the modelling of light transport in holographic displays. This way, imperfections in optical components could be modelled effectively. Such modelling can enable the usage of the resources in the correct way leading to higher quality image generation. These efforts often involved building dedicated CNNs, trained using data from captured photographs, for two-dimensional images [6, 7] and three-dimensional images [8]. We have recently extended these efforts by improving the classical light transport models with learned variables [9]. This way, we trust that we mitigate issues related to small receptive fields typically found in CNNs, leading to physically informed and accurate light transport models that accounts for hardware related imperfections (e.g., optical abberations or nonlinearities).

Perceptual guidance in holography. Classical CGH pipelines optimize or predict holograms that aim to match a certain visual quality across all the parts of a target image or a light field. Aiming for perfection in every target corner may lead to moderate image quality as the density of points reconstructed in a target light field increases. This aim often degrades images with novel noise patterns known as Speckle [10] in the literature. Researchers have recently looked into gaze-contingency in holography by foveating visuals [11], which resulted in improving foveal visuals while potentially offering speckle noise suppression. We extend the idea of gaze contingency for CGH from Chakravarthula et al. [11] by incorporating novel loss functions that statistically match peripheral images rather than pixel by pixel accuracy [12]. Our work [13] generates peripheral visuals with speckle noise such that they are statistically represented correctly. Thus, it distinguishes itself as the first that tries to take advantage of speckle rather than suppress it. Our work is also unique as it generates varifocal images leading to perceptually accurate depth representations.

Standard Holographic Display

Like any other CGH work, our work relies on a holographic display with phase-only modulation. We rely on the most common implementation of a holographic display where a slightly diverging light source approaches a pixelated phase-only modulator. In our implementation, we use a Jasper Display JD7714

Spatial Light Modulator (SLM) as our pixelated phase-only modulator. This specific SLM runs at 30 Hz and provides 360 degrees phase delays and 4094 by 2400 pixels. The SLM's modulated beam goes through a 4f imaging system with a pinhole filter on a Fourier plane. We used 100 mm and 35 mm focal length plano-convex lenses with a one-inch diameter and purchased them from Thorlabs. Finally, a bare imaging sensor is placed near this specific 4f imaging system to capture the images generated by the holographic display from various depths. We use a standard linear stage to move this image sensor back and forth to get it to the correct focal plane. A photograph from our experimental setup can be seen in Figure 1.

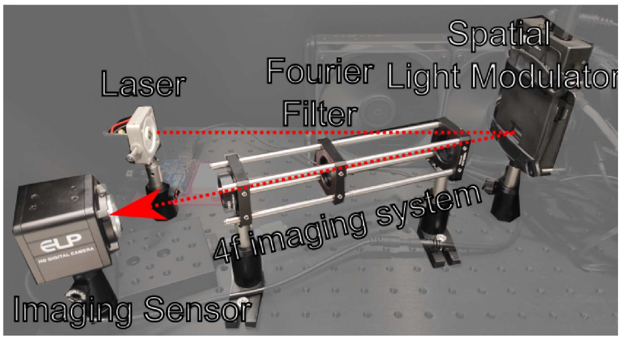

Figure 1. *Standard Holographic Display. Our implementation of a standard holographic display is photographed, showing pieces such as laser light source, spatial light modulator, 4f imaging system, fourier filter, imaging sensor and a linear stage.*

Holographic Light Transport

The light transport for coherent beams is often described using Rayleigh-Sommerfeld diffraction integrals [14], the first solution of the Rayleigh-Sommerfeld integral, also known as the Huygens-Fresnel principle [15], described as follows:

$$u(x,y) = \frac{1}{j\lambda} \iint u_0(x,y) \frac{e^{jkr}}{r} cos(\theta) dxdy, \tag{1}$$

where the field at a target image plane, $u(x,y)$, is calculated by integrating over every point of the input complex field, $u_0(x,y)$. Note that, for the above equation, r represents the optical path between a selected point over an initial complex field and a selected point in the image plane, θ represents the angle between these two points, k represents the wavenumber ($\frac{2\pi}{\lambda}$) and λ represents the wavelength of light. The described holographic light transport model is often simplified into a single convolution with a fixed spatially invariant complex kernel, $h_z(x,y)$ [16].

$$\begin{aligned} u(x,y) &= u_0(x,y) * h_z(x,y) \\ &= \mathscr{F}^{-1}\{\mathscr{F}\{u_0(x,y)\} \cdot \mathscr{F}\{h_z(x,y)\}\} \end{aligned} \tag{2}$$

In our case, we choose to use the Fresnel approximation to the Rayleigh-Sommerfeld which is a commonly used form of h described as

$$h_z(x,y) = \frac{e^{jkz}}{j\lambda z} e^{\frac{jk}{2z}(x^2+y^2)}, \tag{3}$$

where z represents the distance between an initial complex field plane and a target image plane. We identify a new h_z by training with holograms and their corresponding reconstructed images captured as photographs [9]. This way, a more accurate form of h_z could be established for a given hardware, leading to improvements in image quality as sampled in Figure 2.

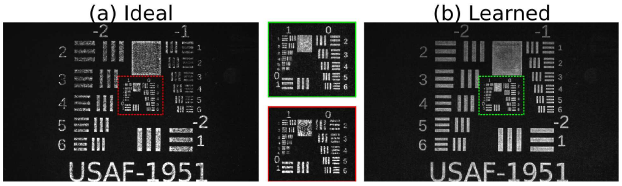

Figure 2. *Learned holographic light transport. (a) Ideal light transport models do not account for imperfections in optical hardware, often times leading to degraded image quality. (b) Rather than using an ideal light transport, learning light transport in the form of a single convolutional kernel can lead to improving image quality. Training process of this learned kernel requires training with an input hologram and a corresponding photograph capturing reconstructed images [9].*

Perceptually guided holograms

Our most recent work provides a varifocal and foveated hologram generation pipeline using the metamers in the peripheral vision, images that look statistically correct but not pixel correct; such images are perceived as identical when viewed in peripheral vision [12] (see Figure 3. We have reformulated our previous work [12] in the form of a differentiable loss, function to optimize holograms with a Stochastic Gradient-based optimizer. This way, we can generate plausible images at the peripheral and focus on the depth that a user is interested in observing. For more, please consult our technical paper [13].

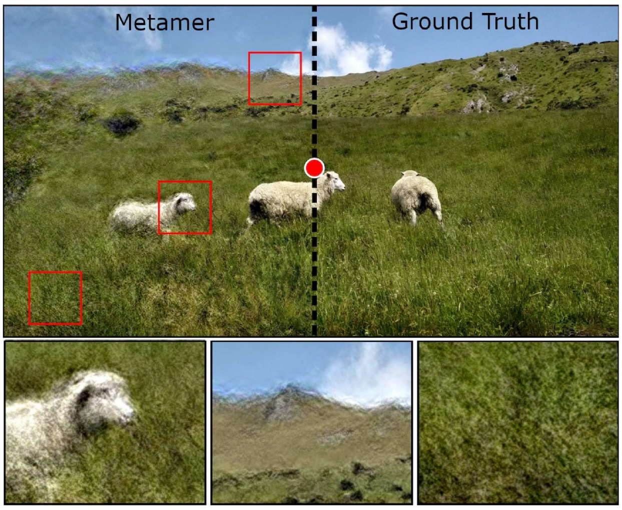

Figure 3. *A sample metamer display by following the work by Walton et al. [12] using a gaze location at the center of the image (red dot at the center).*

Both our loss function and optimizers are readily available in our code library available publicly [17].

Conclusion

CGH offers a unique graphics system that can provide genuinely three-dimensional images. This work discusses a base holographic display, which we call a standard holographic display. We also briefly introduce new tools that can help improve the visual quality of holographic displays. We believe these works formulate a new base for achieving "perceptual" realism in holograms in the future.

Acknowledgments

The author would like to thank Özgür Yöntem and Rafał Mantiuk for their invitation to the London Imaging Conference

2022; all active members of the computational light laboratory, specifically David Walton and Koray Kavaklı, for their support in the development and research introduced in this talk.

References

[1] George Alex Koulieris, Kaan Akşit, Michael Stengel, Rafał K Mantiuk, Katerina Mania, and Christian Richardt. Near-eye display and tracking technologies for virtual and augmented reality. In *Computer Graphics Forum*, volume 38, pages 493–519. Wiley Online Library, 2019.

[2] Chris Slinger, Colin Cameron, and Maurice Stanley. Computer-generated holography as a generic display technology. *Computer*, 38(8):46–53, 2005.

[3] Koray Kavaklı, Yuta Itoh, Hakan Urey, and Kaan Akşit. Realistic defocus blur for multiplane computer-generated holography. *arXiv preprint arXiv:2205.07030*, 2022.

[4] Andrew Maimone, Andreas Georgiou, and Joel S Kollin. Holographic near-eye displays for virtual and augmented reality. *ACM Transactions on Graphics (Tog)*, 36(4):1–16, 2017.

[5] Liang Shi, Beichen Li, Changil Kim, Petr Kellnhofer, and Wojciech Matusik. Towards real-time photorealistic 3d holography with deep neural networks. *Nature*, 591(7849):234–239, 2021.

[6] Yifan Peng, Suyeon Choi, Nitish Padmanaban, and Gordon Wetzstein. Neural holography with camera-in-the-loop training. *ACM Transactions on Graphics (TOG)*, 39(6):1–14, 2020.

[7] Praneeth Chakravarthula, Ethan Tseng, Tarun Srivastava, Henry Fuchs, and Felix Heide. Learned hardware-in-the-loop phase retrieval for holographic near-eye displays. *ACM Transactions on Graphics (TOG)*, 39(6):1–18, 2020.

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

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

[10] Joseph W Goodman. Some fundamental properties of speckle. *JOSA*, 66(11):1145–1150, 1976.

[11] Praneeth Chakravarthula, Zhan Zhang, Okan Tursun, Piotr Didyk, Qi Sun, and Henry Fuchs. Gaze-contingent retinal speckle suppression for perceptually-matched foveated holographic displays. *IEEE Transactions on Visualization and Computer Graphics*, 27(11):4194–4203, 2021.

[12] David R Walton, Rafael Kuffner Dos Anjos, Sebastian Friston, David Swapp, Kaan Akşit, Anthony Steed, and Tobias Ritschel. Beyond blur: Real-time ventral metamers for foveated rendering. *ACM Transactions on Graphics*, 40(4):1–14, 2021.

[13] David R Walton, Koray Kavaklı, Rafael Kuffner dos Anjos, David Swapp, Tim Weyrich, Hakan Urey, Anthony Steed, Tobias Ritschel, and Kaan Akşit. Metameric varifocal holography. March 2022.

[14] John C Heurtley. Scalar rayleigh–sommerfeld and kirchhoff diffraction integrals: a comparison of exact evaluations for axial points. *JOSA*, 63(8):1003–1008, 1973.

[15] Joseph W Goodman. Introduction to fourier optics, roberts & co. *Publishers, Englewood, Colorado*, 2005.

[16] Maciej Sypek. Light propagation in the fresnel region. new numerical approach. *Optics communications*, 116(1-3):43–48, 1995.

[17] Kaan Akşit, Ahmet Serdar Karadeniz, Praneeth Chakravarthula, Wang Yujie, Koray Kavaklı, Yuta Itoh, and David Robert Walton. Odak, September 2021. If you use this software, please cite it as below.

Author Biography

Kaan Akşit is an Associate Professor at University College London, where he leads the Computational Light Laboratory. Kaan's research works are widely known among the optics and graphics community for his contributions to display technologies dedicated to virtual reality, augmented reality, and three-dimensional displays. For more on his research, please refer to https://kaanaksit.com and https://complightlab.com.