



Visual perception of noise in a simulated holographic display—A user study[☆]

Andreas Georgiou^{a,*}, Joel Kollin^b, Charlie Hewitt^a, Praneeth Chakravarthula^a, Brian Guenter^b

^a Interactive Media Group, Microsoft Research, Cambridge, CB1 2FB, UK

^b Interactive Media Group, Microsoft Research, Redmond, WA 98052, USA

ARTICLE INFO

MSC:

41A05

41A10

65D05

65D17

Keywords:

Holography

Foveation

Head mounted displays

Computer Generated Holograms

ABSTRACT

This work explores ways to bypass the fundamental image quality limitations of displays using Computer Generated Holograms (CGHs) and, specifically, the high-frequency noise associated with phase-only holograms. Although there is a wealth of literature on building experimental holographic systems, there are no user studies to assess the performance of a holographic system projecting a dynamic two-dimensional image. In this study, 18 participants blindly compared three groups of images displayed on a conventional monitor. The first group contained the original image, the second the simulated holographic reconstruction of the original image, and the third group had the foveated reconstruction of the original image, based on the pupil position. Holograms in the second group were computed using the Fienup algorithm and the third group using the Fienup with Perceptual Don't Care Areas (FiPDc) algorithm, a novel algorithm that uses eye tracking to optimize image quality in CGHs. The aim of the study was to find out if the holographic display, assuming an ideal hardware, can be as good as a conventional display and whether eye tracking can help this goal. Most participants distinguished between the original image and the un-foveated simulated reconstruction. However, the participants could not differentiate between the original image and the foveated reconstruction. Thus, foveation may be essential in designing and building the first commercial holographic displays.

1. Introduction

1.1. Motivation

Holography has excellent potential as a display technology. Its ability to manipulate the phase of the light can enable compact 2D projectors, truly 3D displays [1] and ultra-thin Head Mounted Display (HMD). In the case of HMD, holography's properties make it highly desirable: a wafer-thin hologram in front of the eye can create an image without requiring any bulky imaging optics. The image only exists in the software and the user's retina, thus, making the entire system extremely compact.

The scientific community has made significant progress towards improving image quality, [2] hologram design algorithms [3] and phase modulators [1,4–6]. In most cases, the proposed goal is a holographic system that will replace or complement existing display technologies, either as an HMD, [2] a 3D display, [1] a HUD [7] or simply a conventional display.

Nevertheless, holograms *cannot* create noise-free images. All images reconstructed from phase-only or amplitude-only holograms contain unwanted orders and may render holographic displays inferior to conventional ones. The noise is due to the nature of phase-only holograms,

and it is irrespective of the algorithm, phase-modulating device or optical setup. For example, the Fourier hologram forming a two-pixel image is a sine-wave, a function that cannot be rendered on a phase-only or amplitude-only hologram. Rendering a sine-wave requires both amplitude (from 0 to 1) and phase (from 0π to 1π) modulation, something that existing Spatial Light Modulators (SLMs) cannot perform. Even by optimizing the phase of the two spots (i.e., pixels), the combined amplitude-phase modulation on the hologram plane remains, and so is image noise.

Amplitude-only holograms also have higher orders but less than phase-only holograms. For this reason, amplitude-only holograms have been used by some systems where low noise levels are necessary [1,8]. Nevertheless, they still create unwanted orders. For example, displaying a single on-pixel requires a blazed grating, a phase-only hologram, which cannot be rendered on an amplitude only device. Amplitude holograms have low efficiency (11%), have a large undiffracted central (aka zero) order and a mirror image. These drawbacks make them unattractive in a HMD where power efficiency is important, and a spatial filter cannot be implemented.

There is a need for a better understanding of the *theoretical* noise levels of a holographic system and whether it will be acceptable to

[☆] This paper was recommended for publication by Prof G. Guangtao Zhai.

* Corresponding author.

E-mail address: ag245@cantab.net (A. Georgiou).

the user. Though the exact user's tolerance to the noise will depend on the application, experimental results are unlikely to be better than simulations. Therefore, simulations can provide us with a "best-case scenario". A user study of this best-case scenario can educate the researcher on whether an application could be feasible in the future. Without user studies, there is a risk that experiments will improve asymptotically towards unacceptable image quality, to an image quality limited by the theoretical bounds of the hologram and the demands of the human visual perception.

1.2. Novelty

The main novelty of this work is the objective comparison of holographic reconstructions (foveated and not) and real-life target images in a controlled user study. Our understanding is that this is the first user study to use an eye tracker in conjunction with the foveated design of holograms. This study becomes further relevant for future holographic systems by proposing a manufacturable optical headset and using target images where high-frequency noise will be visible.

While there are some excellent user studies in the field of perceptual image quality [9,10], they focus in imperfections associated with capture, compression, transmission or reproduction of images in conventional display systems. They do not address the unique noise characteristics of images projected from phase-only holograms. A recent user study, does address the perceived image quality in holographic reconstructions [11]. In this two different imperfect natural (i.e., photographs) images are compared. The first image is the holographic reconstruction and the second a degraded version of the original. The comparison is not made with the original image as done in this study and foveation was not used. It also used primarily photographs while this study also compared text, geometrical shapes and lines.

The user study in this paper puts the simulated holographic image in context with a real-life scenario. Users looked at the simulated reconstructions from the same distance, with the same size and context as they would do in their daily life. Therefore, they had similarly high expectations in terms of image quality.

This paper uses images that are difficult to render with holograms. Many experimental works use images that hide the high-frequency noise generated by the hologram. Such images used include cartoon characters (the human brain is not accustomed to the "ground truth"), images with a bright background, low contrast, Augmented Reality (AR) images where contrast is not an issue or images with high-frequency components. Unlike, this work uses a range of challenging-to-display images for a holographic projector, including pure white areas, text, human faces and high contrast images.

Thirdly, the user study is associated with a practical and manufacturable optical system, and all parameters used are directly linked to current technology solutions. While the aim was to show the feasibility of the application rather than show experimental results, experimental demonstrations that are not compatible with an HMD were excluded. Systems that were excluded due to practical considerations include systems requiring:

- a 4f optical relay with Fourier filters due to their size.
- pixel sizes below what has been commercially demonstrated.
- free-space propagation between a projector and a reflector in the space in front of the eyes due to form factor and eye-lashes interference

In contrast, the optical system proposed here has the potential to become a practical and manufacturable HMD with existing technology.

Beyond the user study, this work does not include any experimental work. The main contribution of this work is to compare the best case scenario holographic system with a conventional display. Comparing an actual holographic reconstruction with an experimental result would confound two factors: the theoretical errors introduced by the hologram's phase-only (or amplitude-only) nature and the experimental

errors due to optical imperfections. Existing optical systems are far from ideal. Only the current SLMs suffer from a series of significant imperfections, including (a) non-flat silicon devices (b) variable thickness of the liquid crystal layer (c) temporal phase errors due to backplane addressing scheme or liquid crystal response (d) inter-pixel fringing fields that low-pass the phase profile (e) polarization modulation between pixels (f) inter-pixel absorption and scattering (g) illumination wavefront phase and amplitude uniformity (h) laser stability and coherence (i) pixel shape. It is currently impossible to model or correct all these imperfections, and any experimental system would introduce some of these noise sources making the comparison meaningless.

1.3. Paper structure

This paper starts by proposing a feasible optical architecture for a holographic HMD. This architecture uses commercially available optical components, and its thickness is expected to be below 5 mm. The appropriate CGH algorithm for this architecture is then described. A crucial element to the CGH algorithm is eye-tracking, which is also used in the study. Then a detailed description of the study setup is given, including the physical setup and how the comparison between the holographic reconstruction and the conventional monitor is made as objective as possible. The Results section presents the reconstructions of the holographic simulations paying particular attention to the distribution of noise in the foveated reconstructions case. It also shows the detailed results of the user study. Discussion elaborates on the performance of the foveated CGH algorithm and how this is reflected in the results. Particular attention is given to the type of images used and how this reflects the study results. Finally, the Conclusions are presented. The authors' views are given on how this work may impact future research, specifically future applications.

2. Method

2.1. Image quality and hologram computation

In a holographic display, image quality is affected by the CGH algorithm and the quality of the optical setup (primarily the SLM). Non-optimal algorithms and imperfect SLMs create high-frequency noise, reduce contrast and reduce the resolution of the reconstructed image. High-frequency noise is the most common imperfection associated with holography, and reduced contrast is usually a consequence of the high noise floor and scattering in the optical system.

Resolution is, in fact, not a significant issue in holographic displays as they are *diffraction limited*. Any resolution loss is usually down to insufficient pixel spacing on the image plane such that the Point Spread Function (PSF) representing the image pixel is wider than the pixel spacing itself.

The PSF representing the image pixel size depends on the device dimensions only (that is, assuming flat illumination). By increasing the number of pixels on the device but keeping the device size the same, the number of image pixels increases but their spacing remains the same i.e., the FOV of the image increases. Thus, the relative spacing between neighboring pixels and their PSF remains the same irrespective of device size and number of pixels. With this dense pixel spacing, there is interference between the PSF of adjacent pixels, and the resolution degrades significantly, typically by a factor of $\times 2$. Increasing the unused space between adjoining image pixels can eliminate this resolution loss. For example, a single pixel from the original image may be replaced with four pixels in a 2×2 arrangement, where three are set to zero. This process is described in [12] as "zero circumscribing" of image pixels. Alternative, replicating the hologram twice in each dimension achieves the same image pixel spacing with less computation. However, the latter approach assumes uniform illumination, which is not always an accurate assumption.

The random noise associated with the algorithm is a consequence of the phase-only nature of the hologram [13]. This noise, often inaccurately referred to as speckle, is the focus of this work. Noise formed due to experimental imperfections (e.g., poor characterization of the SLM) will *not* be considered in this work. Separating the two sources of noise, i.e., noise arising from the hologram computation and noise arising from experimental imperfections helps with the scientific process of identifying and evaluating the causes and impacts of each noise source. Of course, the need for better understanding and characterization of phase devices remains, and reduction of the mismatch between theoretical simulations and real-world SLMs remains very important.

The simplest CGH [14] is a far-field hologram, where the image formed is the Fourier transform of the hologram. Real holograms, displayed on SLMs, are usually phase-only functions with the amplitude of the modulation restricted to unity. On the reconstruction plane, the amplitude of the electric field is also fixed and set by the target image. Obtaining a zero-error reconstruction is theoretically impossible [13]. However, it is possible to push the noise below a perceivable level. The purpose of any CGH algorithm is to minimize the amplitude error in the reconstruction relative to the target image by using the hologram plane phase as a free variable (or the respective phase of the reconstruction).

2.2. The optical architecture

The purpose of this work is to explore ways to bypass the fundamental image quality limitations of holographic displays and, specifically, the high-frequency noise associated with CGHs. This exploration is meaningful only in conjunction with a realistic and practical optical system. There is a plethora of proposed holographic display systems in the literature. In this work, a holographic HMD was chosen as an exemplar application for two main reasons. First, holographic HMDs can have significant advantages over their conventional counterparts. They can be slim, lightweight, do not require prescription glasses and avoid the vergence-accommodation conflict; all highly desirable characteristics by the industry and consumers. Second, the fixed position of the eye relative to the display and the presence of eye-tracking makes the optical system simpler.

The proposed optical architecture referred to as a *Direct-View Holographic Display*, is shown in Fig. 1. Similar to previously proposed architectures, [1,15] it uses a steerable back-light to increase the space-bandwidth of a slow-changing dynamic hologram. It consists of an SLM, a two-dimensional rolling shutter, a waveguide and a two-dimensional scanning mirror. Except for the two-dimensional scanner, all the components can have a sub-millimeter thickness, and a two-dimensional scanner can fit in a sub-5 mm cube. While the proposed architecture can provide all the key functionality expected from a holographic display, including truly three-dimensional projection, this work will only consider the projection of a two-dimensional image at infinity. In many aspects, three-dimensional perception of content is more complex. Parameters like iris diameter (affected by the room's lighting), accommodation (affected by age) and even inter-pupil distance influence depth perception. These parameters are not easily controlled in the user study. And, of course, the existing technology for three-dimensional displays is relatively poor — it is impossible to obtain hardware to display good quality, image points that lie on multiple planes in space.

An HMD must perform two functions. First, create an image and second, create an “eye box”. Image forming corresponds to sending rays into the eye that converge to a virtual point in three-dimensional space. Here, an image at infinity is assumed; therefore, these rays will be parallel. When the image is at infinity, the angular extent of rays defines the size of the FOV and the image in the angle-space. The eye box, a less familiar term, corresponds to the area where the eye can be and still observe the entire image. Therefore, the rays forming all the image pixels should exist in the eye-box. As the eye rotates, the pupil moves within the eye-box and therefore, rays must exist over this area,




Fig. 1. The architecture of the Direct-View Holographic Display. A waveguide together with a two-dimensional scanning mirror create a scanning backlight. The shutter is synchronized with the mirror and only opens when the transmitted light illuminates the center of the user's pupil (equivalent to the zero-order hitting the center of the pupil). The dynamic hologram creates the image by modulating the phase of the collimated beam of light. Here the hologram is shown to form two pixels only.

usually around a square centimeter. In the Direct-View Holographic Display architecture, the eye box is formed with the aid of eye-tracking via *pupil steering*.

In its simplest form, as shown in Fig. 1, the Direct-View Holographic Display consists of three components. First, a waveguide with a scan the laser illumination that illuminates the entire FOV of the display. Second, a shutter that at any moment blocks most of the light and creates narrow beam that illuminates the hologram. Third, the dynamic hologram which modulates this narrow beam of light and encodes the image information on it.

The shutter and the directional backlight create a pupil steering arrangement. Without the hologram, a single beam is made from each shutter opening that would arrive in the user's pupil. It would appear as an array of dots to the user, with each dot corresponding to a single shutter opening.

The purpose of the hologram is to split that single beam, or single dot, created by the pupil steering and form part of the image. It modulates the flat wavefront emerging from the waveguide and splits it into multiple parallel beams, each beam corresponding to an image pixel. The low-resolution high-speed binary amplitude shutter is used to illuminate only a small part of the hologram (referred to as a hologram tile) and thus remove the requirement for a very fast dynamic hologram. Color can be achieved by time-sequential switching red, green and blue lasers while updating the dynamic hologram with the corresponding color sub-frame. Time-sequential color increases the speed requirement of the dynamic hologram by a factor of $\times 3$. However, this is not an issue as many commercial LCOS devices can support sequential color projection.

The Direct-View Holographic Display uses the scanning mirror to create the large angular deflections in the FOV. Therefore, this architecture suffers much less from chromatic dispersion and coherence effects compared to other approaches where the entire image is formed by the hologram. The Direct-View Holographic Display only requires enough coherence and spectral width to create a single tile of the image, which is only a few degrees.




Fig. 2. A comparison of the Direct-View Holographic Display and the study arrangement geometry. (a) A single hologram tile must create a FOV $2\theta_h$ that is larger than the tile's subtended angle, $2\phi_h$. (b). The user kept their head in position using a headrest, and their gaze was recorded using an eye tracker attached to the table. The emulated hologram tile is in front of the user's eye.

Note that the dynamic hologram here in Fig. 1, for clarity, is shown as a transmissive device. Waveguides can also be transparent, like in commercially available AR headsets. Light emerges from the waveguide away from the eye in a reflective arrangement and illuminates the hologram. Light reflected from the hologram then passes through the waveguide again to reach the eye. The shutter can be placed between the waveguide and the reflective hologram. When the reflective device is used, it is essential to ensure light only exists from one waveguide's sides.

There are four degrees of freedom to consider at any moment in time: the two-dimensional position of a pixel in angle-space and the two-dimensional position of the users' eye in xy -space. With the xy -position of the eye known, backwards ray tracing from the user's pupil to the hologram will identify the part of the hologram illuminating the eye at a given time. Therefore, the purpose of the shutter is to break the hologram into smaller tiles. Each hologram tile is illuminated only at the moment in time when the angle of the beam is such that when passing undiffracted through it (or reflected from it), the beam would arrive in the center of the user's pupil. When image information is encoded on the hologram tile, a small part of the FOV is created. This small section of the image is called an "image tile".

The size of the image tile, in angle space, is determined by the diffraction equation and is equal to

$$\theta_h = \pm \arcsin\left(\frac{\lambda}{2\Delta_h}\right) \quad (1)$$

where λ is the wavelength of light and Δ_h the pixel pitch of the hologram. This is also shown in Fig. 2(a). The center of the image tile, again in angle space, is determined by the position of the hologram tile relative to the user's pupil. The subtended angle of the hologram tile is how large it looks (again in angle space) from a point in the eye box. When the center of the tile and the point in the eye-box are on the same optical axis, the subtended angle is given by:

$$\phi_h = \pm \arctan\left(\frac{N_h \times \Delta_h}{2 \times d_h}\right) \quad (2)$$

where N_h is the pixel count of the hologram tile and d_h is the eye relief (see Fig. 2(a)). Assuming paraxial approximation, the subtended angle ϕ_h of a tile is also approximately the separation (in angle space) of two tiles' central pixels (i.e., the undiffracted light). Therefore, two neighboring tiles must create sufficiently large FOV that there is minimal overlap between adjacent tiles. Otherwise, there will be gaps in the image's FoV. For a continued image to be formed, $\phi_h < \theta_h$ and assuming paraxial approximation, the pixel pitch must be

$$\Delta_h < \sqrt{\frac{d_h \lambda}{N_h}} \quad (3)$$

A final consideration of the image and hologram tile sizes is the spacing between two adjacent image pixels formed by the hologram and the total number of image pixels relative to hologram pixels. Interference between neighboring pixels can be avoided by doubling (in each dimension) the hologram resolution relative to the image. Therefore, a hologram tile with N_h pixels creates $N_h/2$ image pixels over an angle of $2\theta_h$. For an angle separation of $\delta\theta$ in radians between two neighboring pixels (assuming paraxial approximation)

$$L_h > \frac{2\lambda}{\delta\theta} \quad (4)$$

where L_h is the hologram tile size and equal to $N_h \Delta_h$. This equation can also be obtained from the Rayleigh criterion but with a different scaling factor due to the square aperture (1.22 becomes 1) and more stringent spot separation due to the coherent nature of laser light (1 becomes 2). The resolution can also be defined in Pixels Per Degree (ppd) so that $\delta\theta = (1/\text{ppd}) \times (\pi/180)$. Therefore the minimum hologram dimensions become:

$$L_h > \frac{360 \times \text{ppd} \times \lambda}{\pi} \quad (5)$$

2.3. Foveation

Exploiting the perceptive field of the eye to improve apparent image quality is an established technique in conventional computer graphics. Particularly useful is the concept of foveation. In the center of the user's vision (the fovea), images are rendered at high resolution, while in the periphery are rendered in low resolution. This variation in resolution matches the perceptive field of the eye [16] such that, when combined with eye tracking, the user perceives only the highest quality image. This technique has been used previously in the rendering of three-dimensional scenes [17] and the context of gaze-contingent displays [18,19]. There has also been increased interest in foveation more recently concerning HMDs [20–23].

The concept of foveation in holography is a relatively new topic. Some work has been done in determining the density with which to sample three-dimensional point clouds to reduce computation times [24–26]. Foveated display of holograms has also been proposed [27], with lower resolution hologram tiles used in the periphery compared with the center of the vision. This technique reduces the total computational load but does not enable more significant optimization of the hologram itself and brings no improvement in the reconstructed image quality. Recently, a weighting function that takes into account the foveated nature of human vision was suggested by [28] in the context of eye box expansion.

In this work, uniform sampling is used for the image. However, the eye's visual acuity for the different areas of the image is considered. By doing so, the hologram design algorithm is allowed more flexibility to shift the noise into the periphery. The freedom to form additional noise in the periphery creates flexibility in the hologram design, making the central parts of the image significantly better quality. The specific algorithm is compatible with the Direct View Display architecture.

2.4. The tile approach

At any moment in time, the shutter only allows light from a part of the hologram (i.e., a hologram tile) to reach the user's eye and to form only a part of the image (i.e., an image tile). The scanning of the hologram is equivalent to breaking the large hologram into many smaller hologram tiles and the image into many smaller image tiles.

Light from the hologram tile enters the user's pupil and is focused by the user's lens onto their retina. The center of the image tile corresponds to the undiffracted light from the hologram tile. Then, the user's lens brings the far-field of the hologram to their retina. Therefore, the Fourier transform relates the hologram tile and image tile. Note that by breaking the large hologram into smaller tiles, computation is also reduced because FFT computation (essential in a CGH algorithm) scales hyper-linearly.

Changing the hologram tile boundaries on a frame-by-frame basis ensures an even distribution of the hologram tiles and pixels around the fovea. Moving the boundary so that the fovea always sees the intersection of four hologram tiles ensures that these four tiles evenly contribute towards the high-fidelity foveated area. This tile arrangement places the center of the eye's FOV and the center of the foveated area in the *corner* of four image tiles (and not in the center of a single tile). Therefore, the tile segmentation moves together with the eye. At any moment, the fovea points to the corners of four hologram tiles. The right image in Fig. 4 shows how the target image is split into four tiles with the fovea (here a blue dotted square) always pointing into the corner of the four tiles' intersection.

2.5. Computing a Fraunhofer hologram

Many algorithms attempt to reduce noise in the far-field of a holographically reconstructed image. Gerchberg-Saxton (GS) [29] is probably the most widely known algorithm, and most modern CGH algorithms are based on it. Fienup [30] is an evolution of GS algorithm where feedback is used between successive iterations. It accelerates computation and reduces final noise. Other known algorithms include the Fienup with Don't Care Areas (FiDoC) [12], Dual-Pixel [31] and Error Diffusion (ED) [32] algorithms.

These algorithms and many more [3] are unsuitable for the Direct-View Holographic Display architecture as they create noisy areas that require a physical stop on the image plane, something which cannot exist in this architecture, or most HMDs. Note that this is why binary phase or amplitude-only holograms are not suitable for this application. For example, FiDoC requires a rectangular aperture to block noise. ED and Dual-Pixel have some similarities with the proposed algorithm; noise increases progressively from the center towards the outer parts of the reconstruction. However, there is no control over the noise distribution. For example, the algorithm cannot control the diameter or position of the foveated region. It is also impossible to have the foveated area on the edge of the reconstruction.

2.6. Fienup with perceptual don't care areas

Here the FiPDoC algorithm is introduced, based on the FiDoC method. FiDoC uses a binary mask, M , to define areas of interest ($M = 1$) and don't care areas ($M = 0$). This flexibility in amplitude within the don't care areas adds free variables that make the optimization easier. By distributing noise in the don't care areas, there is the freedom to

further optimize the area of interest and improve quality. The image defines the target function in the areas of interest, while in the don't care areas, the target function is set to zero. The combined image and don't care areas are defined by G_0 , i.e., the target intensity profile of the reconstruction. After each iteration, the target value is updated to a new value, G_n , that takes into account the error of the previous reconstruction, G_{n-1} .

Below is the pseudo-code of the FiDoC algorithm.

$$H_{n+1} = \mathcal{F}^{-1}(G_n e^{i\angle G_n}) \quad (6)$$

$$h_{n+1} = H_{n+1} / |H_{n+1}| \quad (7)$$

$$g_{n+1} = \mathcal{F}(h_{n+1}) \quad (8)$$

$$G_{n+1} = M [G_0 + (G_0 - |g_{n+1}|) \kappa] + \gamma(1 - M) * |g_{n+1}| \quad (9)$$

where H_n is the ideal phase-amplitude hologram after the n^{th} iteration, h_n is the actual phase-only hologram, g_n the actual reconstruction and G_n the amplitude of the next target reconstruction. The mask, M , is set to unity at the image area and zero at the don't care areas. The parameter γ is the noise suppression parameter and is somewhere between zero and unity. It determines the freedom of the pixel to have the wrong intensity (i.e. how little the algorithm cares about the pixel's error). When γ is set to zero, the effect of the don't care areas disappears and the noise in the don't care areas is suppressed like the rest of the image (thus no additional freedom in the hologram design). When γ is set to unity, the pixel can take any value (which may lead to all the energy arriving in the don't care areas).

The FiPDoC algorithm works in the same way as FiDoC, but the mask is a continuous function from 0 to 1. Where the fovea is centered, the mask has a value of 1 and gradually decreases towards 0. Unlike FiDoC, in this arrangement, even the periphery (i.e., when M is small) is of interest, so the mask value will not drop completely to 0. The mask's weight is determined by the gaze direction, with the highest value where the user's gaze is centered.

The segmentation of the hologram into tiles moves on a frame-by-frame basis, and so does the mask position relative to the image. Therefore, the mask is always centered at the intersection of four tiles. As discussed above, this approach evenly distributes the foveated part of the image between the four tiles. At the same time, it creates ample space in the tile, especially in the diagonally opposite corner to the foveated region, to allow space for the noise to shift.

The function used for the mask is

$$M = (1 - \epsilon) \frac{\zeta}{\zeta + 1} + \epsilon \quad (10)$$

where

$$\zeta = e^{n(R_0 - r)/R_0} \quad (11)$$

with R_0 a constant controlling the size of the foveated region, and n (where $n \geq 0$) is a constant controlling the roll-off in image quality. ϵ controls the minimum value of the mask and so imposes a limit on the reconstructed image quality in the periphery. The parameter r is the pixel's distance from the fovea center. The pixel's distance from the fovea is equal to $r = \sqrt{x^2 + y^2}$ where x and y are the pixel coordinates, with the first pixel being $(0, 0)$.

The mask's purpose is to control the image quality but inadvertently also affects the energy distribution across the image. The variation of energy across the image can be seen by looking at Eq. (9). For example, when all the image pixels are the same, the mask will introduce an amplitude bias as a function of r , the distance from the fovea. It was attempted to reduce these intensity variations by using more complex mask functions. However, the power distribution was image content dependent, and uniformity remained poor.

An alternative and more effective approach were to calculate this slow changing intensity variation across the image and correct it in the FiPDoC algorithm. This correction has to be done in a way that does not affect the control of the high-frequency noise. For this purpose, the

Table 1

This table compares the parameters of the proposed holographic display and the user study setup.

Optical system (Proposed phase-only LCOS device)		User study (Dell 4k monitor)	
Parameter	Value	Parameter	Value
Eye relief (d_h)	15.0 mm	User-monitor separation (d_m)	433 mm
Bragg angle (θ_h)	5.4°	Image tile subtended angle (θ_m)	5.5°
Hologram tile size (N_h)	1024 pix	–	–
Hologram tile pixel pitch (Δ_h)	2.8 μ m	Image tile pixel pitch (Δ_m)	161.4 μ m
Hologram tile subtended angle (ϕ_h)	5.4°	–	–
Hologram tile size ($N_h \Delta_h$)	2.9 mm	Image tile size ($N_m \Delta_m$)	82.6 mm
Image tile FOV ($2\theta_h$)	10.9°	Image tile FOV ($2\theta_m$)	10.9°
Image tile resolution ($N_h/2$)	512 pix	Image tile resolution (N_m)	512 pix
Target image FOV ($4\theta_h$)	21.8°	Target image FOV ($4\theta_m$)	21.8°
Target image resolution ($2(N_h/2)$)	1024 pix	Target image resolution ($2N_m$):	1024 pix
Pixels per degree	47 pix/°	Pixels per degree	47 pix/°

reconstructed image after each iteration, g_{n+1} , and the original target image, G_0 , where low-pass filtered with a Gaussian filter with σ of about a third of the fovea size. Their ratio,

$$f_{n+1} = \left(\frac{\text{LP}(G_0)}{\text{LP}(g_{n+1})} - 1 \right) \mu + 1 \quad (12)$$

is then introduced in Eq. (9) to correct for the error so that

$$G_{n+1} = M [G_0 f_{n+1} + (G_0 f_{n+1} - g_{n+1})\kappa] + \gamma(1 - M) * |g'_{n+1}| \quad (13)$$

3. Perceptual study

The FiPDoC algorithm heavily relies on the perceptual characteristics of the eye and hence the human visual system. Factors like the spatial extent of the fovea, the spatial and temporal response of the peripheral vision, and the user's gaze influence the apparent perception of imagery. Assessing these factors via modeling is not practical for the present. Understanding and modeling the dynamic behavior of the eye is something that is still an active research topic. Instead, a computer monitor displayed the reconstructed holographic images (calculated using the FFT), and subsequently, their perceptual quality was assessed by real human subjects in a user study. The study was designed to reflect the characteristics of a practical implementation of the Direct-View Holographic Display. Key factors, like the pixel pitch of the Liquid Crystal Over Silicon (LCOS) device, eye relief, and shutter speed, were taken from existing commercial devices.

3.1. User study setup

The study used an eye tracker, a chin rest where the user places their head, and a monitor to display the simulated reconstruction of the hologram. The purpose of this setup was to emulate a holographic HMD attached to the user's head. Fig. 2 shows the evaluation setup used in the study and how it emulates the Direct-View Holographic Display. The eye tracker was internally developed and was used to track the user's gaze at a frame rate of 60fps. The eye tracker resembles a pair of spectacle frames and is using a small IR camera positioned on the left and right of the nose, next to the left and right nose rests. The display (28-inch 4k Dell) displayed the simulated reconstruction. A camera was used to take high-magnification photographs of various test patterns on the monitor. These photographs ensured that the displayed image was not post-processed by the video card or the monitor. The eye tracker and the chin rest were fixed to the desk. The monitor's position and size of the simulated reconstruction on the monitor were arranged so that the FOV corresponds to that of a practical holographic HMD. Table 1 shows the parameters used for the reconstruction evaluation setup.

3.2. Image displayed

Both the dynamic hologram and image are segmented into tiles. This segmentation into tiles is a function of the eye position and changes dynamically as the user moves their gaze. The user's gaze always points to the corner of four image tiles. As the monitor has a finite size, only the *equivalent area* of 2×2 tiles is shown on the screen, and this area had a fixed position on the monitor. However, given the image (and hologram) tiles' segmentation changes, 3×3 image (and hologram) tiles are required to cover the *equivalent area* of 2×2 tiles on the monitor. The additional tiles create a variable margin around the displayed reconstruction (the size of this margin depends on the gaze position). Note that all 3×3 image tiles will be displayed in an experimental implementation of this system. Beyond that, a foveated hologram is not necessary due to the poor resolution of the periphery.

At any moment, the dynamic hologram created an image of 3×3 tiles with each image tile of 512×512 pixels; i.e., the equivalent of a 1548×1548 image. This image was cropped to a size of 1024×1024 , which represents the image displayed on the monitor. As the image was over-sampled by a factor of two, the actual resolution of the hologram has to be twice that, thus making the hologram tile 1024×1024 and the total number of hologram pixels 3072×3072 . The image displayed on the screen was not over-sampled and kept at 1024×1024 as it would not have any meaningful impact on a conventional (i.e., non-diffractive) display.

The images presented to the users were color. Time-sequential color is assumed with the same sampling points for all three colors. In reality, to achieve the same sampling points for red, green, and blue, their respective holograms must have different sizes so that the image pixel spacing is the same. The pixel spacing is kept constant by keeping the term λ/N_h constant for all three colors. The variable hologram size will make the red hologram tile and image tile larger, meaning relatively smaller foveated area, and better image quality. On the other side, the blue color will have worse performance. However, the human eye has significantly lower resolution in blue compared to red and green; this is due to increased chromatic aberrations of the eye's lens, defocus, and reduced cones. Some studies put the eye's contrast sensitivity in blue almost an order of magnitude lower than in green and red [33] (that is why prescription glasses are only adjusted for red and green colors). Therefore the marginal increase in noise in blue is unlikely to impact our evaluation.

The holograms for all the target images and multiple gaze positions were computed off-line, and their reconstructions were stored on the computer's hard drive. The hologram tiles were calculated using the FiPDoC algorithm in 128 iterations. The characteristics of the mask parameters were established empirically by the authors. The simulated reconstructions were 1548×1548 pixels and were cropped to 1024×1024 . The monitor displayed the appropriate reconstruction according to the user's gaze direction. The eye tracker reported the user's eye position, and then the correct reconstruction was shown on




Fig. 3. The five images used for the user study: (a) bars (b) face (c) C++ (d) excel and (e) tiger. Each image is 1024×1024 pixels which is equivalent in size to 2×2 image tiles.

display. The reconstructions for all gaze positions for a specific image were kept in the computer's RAM to minimize latency. While the fovea can point at any point on the image, it was impossible to store many reconstructions in RAM simultaneously. Therefore, only reconstructions for a grid of 9×9 gaze positions were computed, keeping the number of images stored to only 81.

3.3. Target images

In total, this study used five target images which are shown in Fig. 3. The image selection represents the variety of images shown on conventional displays. The aim was to challenge the algorithm (and the user) by choosing images where random noise cannot hide. Therefore, four out of five images (bars, face, C++ and excel) contain large areas of uniform or nearly uniform light distribution. tiger was a more natural image, and noise is less likely to be perceivable. The most challenging image is expected to be bars. The gray background was chosen so there is no clipping of any noise due to the limited dynamic range of the monitor. In parallel, the bars, black and white, are expected to create higher orders that may be visible in the gray background. face is also challenging due to the cognitive ability of humans to recognize fine detail in faces. C++ and excel were used to identify any issues with color uniformity and character rendering. Finally, tiger was used as an example of a natural image with lots of high-frequency components.

3.4. Reconstruction evaluation

The user study required subjects to sit in front of the monitor displaying the holographic reconstruction. Their head was inside the harness of the eye tracker, and their chin was leaning on the chin rest. Five target images were used, each having three versions: (a) original target image, (b) Fienup reconstruction and (c) FiPDoC reconstruction. Fig. 3 shows the original target images. The user was presented with two of these three seemingly identical images for 20 s each. At the end of the 2×20 s, the user was asked to vote whether one image was better visually than the other or if both images were the same. An image voted as better received two points, if identical, one point, and if worse, zero points. Each comparison uses two images, and there are three versions of each image. Therefore, for each image, there are three unique comparison pairs. As the images were presented in sequence, each pair-comparison was shown twice but with their order reversed, making six comparisons per target image. Since there are five images, the total number of comparisons is 30. The order of the pairs was randomized for each user. In addition, the study used some dummy comparisons at the beginning to train the user and avoid any learning errors for the core comparisons. In total, 18 subjects took part in the study, with none aware of the detailed purpose of the investigation.

4. Results

4.1. Simulation results

Fig. 4 shows the simulated FiPDoC reconstruction compared with the target image and the Fienup reconstruction for bars. The hologram of this image was the most challenging one to compute. For presentation simplicity, the fovea is assumed to look exactly at the center of the image. So each image shown here consists of exactly 2×2 tiles; thus, no cropping occurs (for all other cases, an array of 3×3 tiles is reconstructed, and the image is cropped to the equivalent area of 2×2 tiles). The dashed lines show the extent of the four tiles. The solid red line and blue dotted line squares show the periphery and central part of the image in detail. In the magnified squares, the top left triangle shows the actual reconstruction. The bottom right triangle shows only the noise (i.e., the difference between original and reconstruction) amplified by a factor of 5 to make it clearer.

Fig. 5 shows a more detailed comparison of the simulated reconstructions of the image face using Fienup and FiPDoC. For presentation simplicity (as in Fig. 4), the fovea is placed in the center of the image, and the image consists of 2×2 tiles. In this comparison, the Peak Signal-to-Noise Ratio (PSNR) is plotted for a horizontal window in the center of the image. The extent of the window is shown surrounded by a solid black line. The approximate area where FiPDoC outperforms Fienup is enclosed in the dashed-lined circle.

4.2. Study results

The scores of the five images in the study are shown in Table 2. If two images were identical, and for a large sample, users would have voted the two images equally, and the comparison score would be 50%. Therefore, the comparison was made simpler by showing the score as an offset from 50%, making the ideal score for a perfect reconstruction $50\%+0\%$. Each subject made a comparison twice for each pair, and there were 18. Therefore, each comparison performance was averaged across 36 votes.

The statistical significance of the results is evaluated by how likely it is that the event happened by accident. If the voting were done randomly (for example, if the images were identical), the score distribution for 36 votes would be Gaussian-like with a standard deviation, σ of $\sigma = 6.8\%$ and a mean value of 50%. Therefore in Table 2 the results in parenthesis show how far the result is from the mean value normalized to the standard deviation of the distribution. There are many tests to define if an event is biased or not, and it highly depends on the application. Here results with less than 10% probability ($\approx 1.3\sigma$) are considered biased; i.e., the two images were not identical.

5. Discussion

The key to this work is the Direct-View Holographic Display architecture. This arrangement avoids pupil replication; therefore, full control of each ray's angle and position is achieved. Pupil steering




Fig. 4. Here the results for the Fienup and FiPDoC algorithms are compared. The zoomed regions in the periphery (solid red line) and fovea (dashed blue line) show the reconstruction in magnification (top left triangle) and the absolute error ($\times 5$). Low-intensity noise is uniformly perceptible using Fienup. No noise is perceptible using FiPDoC. The dashed lines show the extent of each tile.

Table 2

The scores of the five images for the Fienup and FiPDoC algorithms.

	Users identifying Original better than Fienup 50%+	Users identifying Original better than FiPDoC 50%+
bars	44.4% (6.5 σ)	23.6% (3.5 σ)
C++	12.5% (1.8 σ)	2.8% (0.4 σ)
excel	26.4% (3.9 σ)	1.4% (0.2 σ)
face	9.7% (1.4 σ)	1.4% (0.2 σ)
tiger	2.8% (0.4 σ)	0% (0.0 σ)

enables prescription-glasses-free operation, three-dimensional content, and no conflict between vergence and accommodation. Conventional waveguides can only support two dimensions in angle; positional information is lost. The Direct-View Display architecture retains control of positional information of the rays by encoding the image information *after* the waveguide. Instead, the waveguide here encodes only the pupil positional information.

5.1. Simulation results discussion

Figs. 4 and 5 show clearly how the noise is pushed into the periphery of the image. For both images, the central part of the FiPDoC reconstruction is virtually identical to the original image. Only when the noise is amplified by a factor of 5 is it just visible. At the same time, the noise in the periphery is significantly higher using FiPDoC compared to Fienup.

In Fig. 5, the circle shows how the noise is effectively pushed into the four corners of the four tiles. The width of each tile is $\sim 11^\circ$, so the total diameter of the circle is $\sim 20^\circ$. This value was set empirically and was a compromise between having a sufficiently large area for the fovea and a smooth transition between low-noise and high-noise areas. The dashed circle roughly matches the size of the macula of the eye while the fovea has approximately a quarter of this diameter at $\sim 5^\circ$.

5.2. Study results discussion

The study results show that the foveated simulated reconstructions appeared almost as good as the original images. In four out of five images, the users voted them as being almost identical. There were 18 subjects, two comparisons for each pair, and two points per comparison. Therefore the 1.4% (for excel, face and tiger) and 2.8% (for C++) correspond to only one and two comparison bias towards the original out of the 36 comparisons.

From the user study results, only bars was identified by the users as being of lower image quality than the original. Though not evaluated directly in the study, the bars image had some low-frequency intensity variation that appeared as a flicker to the users. The low-spatial intensity correction (as described in Section 2.6) could cause these low-frequency temporal intensity variations. The low pass spatial filtering was not able to perform very well when the image also contains low spatial frequencies, here in the form of black, white, and gray bars. Therefore, as the eye moves, different foveated images appear with the wrong (and different between them) DC component at the edges of the image; to the user, this seems as temporal variations. The aim is that future versions of the algorithm will adapt the low-pass filtering process with a more image content-based approach to correct these low-frequency variations.

Interestingly, the tiger image, with its primarily high-frequency content, was able to hide the noise using the Fienup algorithm effectively. This image shows how important it is to choose a wide range of images when evaluating a hologram design algorithm. Selecting an image with high-frequency content can hide many of the issues associated with the main limitation of holograms: the high-frequency noise.

5.3. Engineering considerations

This work concentrated on a simplified optical system to identify the theoretical limitations of holograms and physiological parameters of the human visual system. Some of the aspects of an experimental system would be more challenging; the implementation of the dynamic




Fig. 5. Reconstructed images examples using Fienup and FiPDoC algorithms. Reconstructed regions in the fovea and extreme periphery, respectively (with half the region showing absolute error $\times 5$) and plots of PSNR across each image are shown below. PSNR is measured using a 32-pixel sliding window on the central horizontal stripe. Note that PSNR is image content dependent hence the slight variations across the image. Both images are 1024×1024 pixels and optimized over 128 iterations.

hologram, for example. Other elements may bring some advantages to the system.

For example, the foveation points can take any position and are not restricted to the 9×9 grid, like in this work. Instead, the hologram (and appropriate mask M) will be computed for the exact point of the fovea. This non-discrete approach to foveation will reduce the required diameter of the foveated region and the corresponding size of the image tile (in angle space). Consequently, it can help increase the dynamic hologram's pixel pitch. Currently, the spacing between two foveated positions is $\sim 2.72^\circ$ meaning that the foveated region must be at least $\sim 1.36^\circ$ larger than required.

At the same time, faster eye-tracking could reduce the angular distance the eye travels between successive frames. The eye tracker in this study was running at 60fps maximum speed. Faster eye tracking (currently commercially available) and better synchronization between the monitor and eye tracker can further reduce the latency.

Finally, this study used the same image reconstructions for each fovea position. In other words, if the fovea looks at the same position, the identical foveated reconstruction is presented for multiple display frames. In contrast, a real holographic display will compute a different hologram for each successive image frame. Therefore, the noise distribution will change sufficiently fast that the temporal integration of the eye will further reduce the perceived noise intensity.

6. Conclusions

This work aimed to propose an optical architecture and the associated hologram design algorithm that could enable future holographic HMDs. Both the Direct-View Holographic Display architecture and the FiPDoC algorithm are tools that can help meet the expectations of future users while also having a practical implementation. In some respects, holographic HMDs are the most straightforward application of holographic displays, with a single eye viewing the holographic display

within a small eye box. This work further simplifies this task. Firstly, pupil steering reduces the actual size of the eye box. Secondly, by using foveation, the image fidelity observed by the user can match the quality of a conventional display, as demonstrated by the user study conducted here. Specifically, when foveation was used, only a single image out of the five targets shown to the users was identifiable as being worse than the original, only from about half of the users (23.6%). In contrast, four out of five images were identifiable as being worse when a non-foveated algorithm was used and by many more users.

The first commercially viable holographic display must compromise in form factor, image quality, cost, power consumption, and functionality. While this work looked at a potential solution for form factor and image quality, many more challenges must be resolved before the first holographic display is mass-produced. Bringing a holographic display into the market is a massive task spanning many disciplines and requiring the collaboration of multiple organizations.

CRedit authorship contribution statement

Andreas Georgiou: Conceptualization, Software, Methodology, Validation, Writing – original draft, Writing – review & editing, Visualization, Supervision, Project administration, Funding acquisition. **Joel Kollin:** Conceptualization, Methodology, Writing – original draft, Writing – review & editing, Supervision, Project administration, Funding acquisition. **Charlie Hewitt:** Software, Methodology, Writing – original draft, Validation. **Praneeth Chakravarthula:** Software, Methodology, Writing – original draft, Validation. **Brian Guenter:** Conceptualization, Supervision, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

References

- [1] J. An, K. Won, Y. Kim, J.-Y. Hong, H. Kim, Y. Kim, H. Song, C. Choi, Y. Kim, J. Seo, A. Morozov, H. Park, S. Hong, S. Hwang, K. Kim, H.-S. Lee, Slim-panel holographic video display, *Nature Commun.* 11 (1) (2020) 5568.
- [2] A. Cem, M.K. Hedili, E. Ulusoy, H. Urey, Foveated near-eye display using computational holography, *Sci. Rep.* 10 (1) (2020) 14905.
- [3] L. Shi, B. Li, C. Kim, P. Kellnhofer, W. Matusik, Towards real-time photorealistic 3D holography with deep neural networks, *Nature* 591 (7849) (2021) 234–239.
- [4] J.-H. Yang, J.H. Choi, J.-E. Pi, C.-Y. Hwang, G.H. Kim, W.-J. Lee, H.-O. Kim, K. Choi, Y.-H. Kim, C.-S. Hwang, High-resolution spatial light modulator on glass for digital holographic display, in: S. Miyata, T. Yatagai, Y. Koike (Eds.), *Ultra-High-Definition Imaging Systems II*, Vol. 10943, SPIE, International Society for Optics and Photonics, 2019, pp. 25–31.
- [5] D. Smalley, Q. Smithwick, V.M. Bove Jr., J. Barabas, S. Jolly, Anisotropic leaky-mode modulator for holographic video displays, *Nature* 498 (2013) 313–317.
- [6] C.-S. Hwang, Y.-H. Kim, J.-H. Yang, J.H. Choi, J.-E. Pi, G.H. Kim, C.-Y. Hwang, K. Choi, H.-O. Kim, W.-J. Lee, S.-H.K. Park, J. Kim, 46-2: Invited paper: Ultimate resolution active matrix display with oxide TFT backplanes for electronic holographic display, *SID Symposium Digest of Technical Papers*, Vol. 49, (1) 2018, pp. 610–612.
- [7] B. Mullins, P. Greenhalgh, J. Christmas, 59-5: Invited paper : The holographic future of head up displays, *SID Symposium Digest of Technical Papers*, Vol. 48, 2017, pp. 886–889.
- [8] C. Slinger, C. Cameron, M. Stanley, Computer-generated holography as a generic display technology, *Computer* 38 (8) (2005) 46–53.
- [9] X. Min, K. Gu, G. Zhai, X. Yang, W. Zhang, P.L. Callet, C.W. Chen, Screen content quality assessment: Overview, benchmark, and beyond, *ACM Comput. Surv.* 54 (2022) 1–36.
- [10] G. Zhai, X. Min, Perceptual image quality assessment: a survey, *Sci. China Inf. Sci.* 63 (2020).
- [11] M.H. Eybposh, A. Moossavi, J. Xu, N.W. Caira, N.C. Pegard, Perceptual quality assessment in holographic displays with a semi-supervised neural network, in: *Digital Holography and 3-D Imaging*, (Th1A.6) Optica Publishing Group, 2022.
- [12] A. Georgiou, J. Christmas, N. Collings, J. Moore, W.A. Crossland, Aspects of hologram calculation for video frames, *J. Opt. A: Pure Appl. Opt.* 10 (3) (2008) 035302.
- [13] A. Georgiou, Noise formation in Fourier phase-only holograms, *J. Opt. Soc. Amer. B* 27 (12) (2010) 2677–2686.
- [14] H. Dammann, K. Görtler, High-efficiency in-line multiple imaging by means of multiple phase holograms, *Opt. Commun.* 3 (5) (1971) 312–315.
- [15] N. Leister, A. Schwerdtner, G. Fütterer, S. Buschbeck, J.-C. Olaya, S. Flon, Full-color interactive holographic projection system for large 3D scene reconstruction, in: L.-C. Chien (Ed.), *Emerging Liquid Crystal Technologies III*, Vol. 6911, SPIE, International Society for Optics and Photonics, 2008, pp. 202–211.
- [16] H. Strasburger, I. Rentschler, M. Jüttner, Peripheral vision and pattern recognition: A review, *J. Vis.* 11 (5) (2011) 13.
- [17] B. Guenter, M. Finch, S. Drucker, D. Tan, J. Snyder, Foveated 3D graphics, *ACM Trans. Graph.* 31 (6) (2012) 164:1–164:10.
- [18] A.T. Duchowski, A. Çöltekin, Foveated gaze-contingent displays for peripheral LOD management, 3D visualization, and stereo imaging, *ACM Trans. Multimed. Comput. Commun. Appl.* 3 (4) (2007) 6.
- [19] E.M. Reingold, L.C. Loschky, G.W. McConkie, D.M. Stampe, Gaze-contingent multiresolutional displays: An integrative review, *Hum. Factors* 45 (2) (2003) 307–328.
- [20] A. Patney, M. Salvi, J. Kim, A. Kaplanyan, C. Wyman, N. Benty, D. Luebke, A. Lefohn, Towards foveated rendering for gaze-tracked virtual reality, *ACM Trans. Graph.* 35 (6) (2016) 179:1–179:12.
- [21] A. Patney, J. Kim, M. Salvi, A. Kaplanyan, C. Wyman, N. Benty, A. Lefohn, D. Luebke, Perceptually-based foveated virtual reality, in: *ACM SIGGRAPH Emerging Technologies*, ACM, 2016, p. 17.
- [22] R.C. Olmo, K. Ahmed, Foveated displays for virtual and augmented reality, 2018, US Patent App. 15/461, 215.
- [23] J. Kim, Y. Jeong, M. Stengel, K. Akşit, R. Albert, B. Boudaoud, T. Greer, J. Kim, W. Lopes, Z. Majercik, et al., Foveated AR: dynamically-foveated augmented reality display, *ACM Trans. Graph.* 38 (4) (2019) 1–15.
- [24] J. Hong, Y. Kim, S. Hong, C. Shin, H. Kang, Gaze contingent hologram synthesis for holographic head-mounted display, *Proc. SPIE* 9771 (2016) 97710K–97710K–6.
- [25] Y. Ju, J. Park, Fast generation of mesh based CGH in head-mounted displays using foveated rendering technique, in: *Digital Holography and Three-Dimensional Imaging*, OSA, 2018, pp. DTu5F–6.
- [26] C. Chang, W. Cui, L. Gao, Foveated holographic near-eye 3D display, *Opt. Express* 28 (2) (2020) 1345–1356.
- [27] J.S. Kollin, A. Georgiou, A. Travis, D.C. Burger, Holographic near-eye display, 2016, US Patent App. 14/754, 451.
- [28] G. Kuo, L. Waller, R. Ng, A. Maimone, High resolution Étendue expansion for holographic displays, in: *ACM Transactions on Graphics*, ACM, 2020, p. 14.
- [29] R.W. Gerchberg, W.O. Saxton, A practical algorithm for the determination of the phase from image and diffraction plane pictures, *Optik* 35 (1972) 237–246.
- [30] J.R. Fienup, Reconstruction of a complex-valued object from the modulus of its Fourier transform using a support constraint, *J. Opt. Soc. Amer. A* 4 (1) (1987) 118–123.
- [31] Y. Qi, C. Chang, J. Xia, Speckleless holographic display by complex modulation based on double-phase method, *Opt. Express* 24 (26) (2016) 30368–30378.
- [32] P.W.M. Tsang, T.-C. Poon, Novel method for converting digital Fresnel hologram to phase-only hologram based on bidirectional error diffusion, *Opt. Express* 21 (20) (2013).
- [33] K.T. Mullen, The contrast sensitivity of human colour vision to red-green and blue-yellow chromatic gratings, *J. Physiol.* 359 (1) (1985) 381–400.