# Low-bit-rate computer-generated color Fresnel holography with compression ratio of over 1600 times using vector quantization [Invited]◇

Peter Tsang,[1,*] Kayton W. K. Cheung,[1] and Ting-Chung Poon[2,3]

[1]Department of Electronic Engineering, City University of Hong Kong, 83 Tat Chee Avenue, Kowloon, Hong Kong, China

[2]The Bradley Department of Electrical and Computer Engineering, Virginia Tech, Blacksburg, Virginia 24061, USA

[3]Shanghai Institute of Optics and Fine Mechanics, Chinese Academy of Sciences,
P.O. Box 800-211, Shanghai 201800, China

*Corresponding author: eewmtsan@cityu.edu.hk

We propose a method for compressing a digital color Fresnel hologram based on vector quantization (VQ). The complex color hologram is first separated into three complex holograms, each representing one of the primary colors. Subsequently, each hologram is converted into what we call a real Fresnel hologram and compressed with VQ based on a universal codebook. Experimental evaluation reveals that our scheme is capable of attaining a compression ratio of over 1600 times and still preserving acceptable visual quality on the reconstructed images. Moreover, the decoding process is free from computation and highly resistant to noise contamination on the compressed data.   © 2011 Optical Society of America

*OCIS codes:*    090.0090, 090.1705, 090.1760, 090.1995.

◇**Datasets associated with this article are available at** http://hdl.handle.net/10376/1576. **Links such as "View 1" that appear in figure captions and elsewhere will launch custom data views if ISP software is present.**

## 1. Introduction

In the past two decades, investigation of computer-generated holography (CGH) [1] has become an area of immense interest. This approach enables holograms of both real and synthetic subjects to be implemented without the requirement of expensive and precision setup of optical equipments. Fueled by the rapid advancement of computing technology, it is likely that CGH could become a potential solution for three-dimensional (3D) televisions [2] in the foreseeable future. An effective means in CGH is the point light concept, which is a numerical realization of the zone-plate approach to holography as explained in [3]. In brief, an object is assumed to be composed of self-illuminated points, each contributing to an elementary hologram known as a Fresnel zone plate (FZP). The hologram of a 3D object is generated by summing the FZP corresponding to each individual object points. A hologram generated with this method, which simulates a Fresnel hologram, is hereafter refer as a generated digital Fresnel hologram (or simply a generated digital hologram). Given a set of 3D object points, $O(x,y,z) = [o_0(x_0,y_0,z_0), o_1(x_1,y_1,z_1), ....., o_{N-1}(x_{N-1},y_{N-1},z_{N-1})]$, the complex amplitude of the object beam $H(x,y)$ is given by

$$H(x,y) = \sum_{j=0}^{N-1} \frac{a_j}{r_j} \exp(ikr_j), \qquad (1)$$

where $k = 2\pi/\lambda$ is the wave number of the laser light with $\lambda$ being the wavelength of the laser. $a_j$ represents the amplitude of the "$j$th" object point,

and $r_j = \sqrt{(x_j - x)^2 + (y_j - y)^2 + z_j^2}$ is the distance between the object point $(x_j, y_j, z_j)$ and a point $(x, y)$ on the hologram. Without loss of generality, the hologram is positioned on the $x$–$y$ plane (i.e., on the $z = 0$ plane), and the object point is at a distance $z_j$ away from the hologram. The method can be directly extended to the generation of a color Fresnel hologram. A color 3D scene is first decomposed into the red, green, and blue primary color components, and each of them is converted into a generated digital Fresnel hologram based on Eq. (1). Subsequently, the source 3D scene can be reconstructed by integrating the three holograms into a suitable format (e.g., a bitmap or a JPEG file) and displaying it on a high-resolution color display [e.g., the liquid crystal on silicon (LCOS) device]. One of the major concerns associated with CGH is the enormous amount of data required to generate a hologram. The problem is obviously more serious with a color hologram, where the data size is three times more than that of a monochrome hologram. In the past, a number of solutions have been proposed to provide moderate compression on monochrome holograms. This includes, but is not limited to, the early works of Yoshikawa [4] and, more recently, the use of quantization [5], Fresnelets [6], lossless and lossy compression [7], companding [8], the histogram approach [9], and vector quantization (VQ) [10,11]. However, compression of generated digital color Fresnel holograms at low bit rates is seldom mentioned. In this paper, we propose a method on very low bit rate encoding of the computer-generated color Fresnel hologram with the following advantages over existing state-of-the-art approaches:

a. high compression ratio (CR) of over 1600 times,
b. computation-free and memory-less decompression, and
c. high resistance to noise contamination on the encoded data.

Our method can be divided into two parts. First, the complex color hologram is separated into three holograms, each corresponding to one of the primary colors (red, green, and blue). Each hologram is converted into a real hologram by adding an off-axis reference beam. This process effectively combines the real and imaginary components, reducing the representation by two times. Moreover, with the addition of the off-axis reference beam, the twin image and the zeroth-order diffraction are separated from the main image in the reconstruction. Second, each real hologram is compressed with VQ. We shall illustrate, in the later part of this paper, that the color hologram can be easily recovered (i.e., decoded) from the compressed data in a computation-free manner. Subsequently, the source color 3D scene represented by the decompressed hologram is optically reconstructed on a color LCOS device.

Organization of the paper is given as follows. In Section 2, a brief outline on the principles of vector quantization is given. The proposed method is reported in Section 3. We shall describe the conversion of the generated digital complex color hologram into three real holograms (each corresponding to one of the three primary colors) and compression of each hologram with VQ. Evaluation on the proposed method and comparison with the use of JPEG 2000 scheme are reported in Section 4. This is followed by a conclusion summarizing the essential findings.

## 2. Vector Quantization

VQ [12] can be interpreted as a mapping $G : R^k \rightarrow C$ converting each vector in a $k$-dimensional Euclidean space $R^k$ to the nearest member in a finite set $C = \{c_i \in R^k : i = [0, N-1]\}$ known as the codebook. $c_i$, $i$, and $N$ are generally referred to as the codevector, the transmission label (or index), and the codebook size, respectively. In the compression of a given set of source data $S = \{s_0, s_1, \dots\dots s_{M-1}\}_{s_j \in R^k|_{0 \le j < M}}$, each element $s_j$ is mapped to the nearest codevector $c_q$ in $C$ according to the minimum distortion criteria, i.e., $d(s_j, c_q) \le d(s_j, c_i)$ for $i = [0, N-1]$. The term $d(s_j, c_q)$ represents the Euclidean distance given by

$$d(s_j, c_q) = \|s_j - c_q\| = \sqrt{\sum_{m=0}^{k-1}(s_j^m - c_q^m)^2}, \qquad (2)$$

where $s_j^m$ and $c_q^m$ represent the $m$th dimension of vector $s_j$ and $c_q$, respectively.

In the encoding process, each of the $k$-dimensional vector $s_j$ in set $S$ will be replaced by a scalar index $i$ and transmitted to the decoder. In the latter, assuming codebook $C$ is available, the transmission label $i$ is used as the index to extract the corresponding vector $c_q$ in the codebook to approximate the input vector $s_j$. Suppose the value of each dimension of an input vector is quantized with $Q$ bits, the data size of the vector is $k \times Q$ bits. After encoding with VQ, this vector will be mapped to the nearest vector in the codebook and represented by the corresponding index. Assuming that there are $N$ vectors in the codebook (i.e., codebook size $= N$), where $N$ is a power of 2 (i.e., $N = 2^m$), the CR is given by

$$CR = \frac{k \times Q}{m} = \frac{k \times Q}{\log_2 N}. \qquad (3)$$

Compressing a set of vectors $S$ with VQ requires a codebook $C$. This process is usually conducted with the iterative algorithm reported in [13] by clustering a set of vectors $T = \{t_0, t_1, \dots\dots t_{N-1}1\}$ into $N$ code vectors (also referred to as the centroids). The algorithm is commonly known as the "LBG algorithm," which is derived from the first character of the names of its three inventors: Linde, Buzo, and Gray. As the details have been reported in [13] and various literature, we shall only show the essential steps as listed

**Table 1.   LBG Algorithm in Codebook Generation [13]**

| Step | Operation |
|------|-----------|
| 1 | Initialize an initial codebook $C = [c_0, c_1, \ldots, c_{N-1}]$ by randomly selecting $N$ vectors from $T$. |
| 2 | Determine the centroid in the codebook that is nearest to each vector in $T$ according to the Euclidean distance given in Eq. (2). |
| 3 | Update each centroid by averaging the vectors in $T$ that are nearest to it. |
| 4 | Calculate the mean square error (MSE) between the centroids obtained in step 3 and the nearest vectors that are associated with each of them. |
| 5 | If the MSE is lower than a certain threshold, stop and output codebook $C$. |
| 6 | Go to step 2. |



Fig. 1.   Partitioning of the generated digital Fresnel hologram into nonoverlapping blocks.

in Table 1. If $C$ is generated from the image to be compressed (i.e., $T = S$), the codebook is termed the in-training set. Otherwise, it will be referred to as the out-training set.

## 3.   Hologram Compression

### A.   Converting a Complex Color Fresnel Hologram into a Real Color Fresnel Hologram

To begin with, let the color complex hologram be denoted by $H(x,y)$, where $x$ and $y$ are the horizontal and vertical axes. The red, green, and blue color planes of $H(x,y)$ are represented by $H_0(x,y)$, $H_1(x,y)$, and $H_2(x,y)$, respectively. Each of the hologram planes is generated, based on Eq. (1), from the corresponding color component of the source 3D image. For example, the red component of the source image will be taken to generate the red hologram $H_0(x,y)$. According to Eq. (1), the generated digital complex hologram, $H(x,y)$, as well as its constituting color hologram, are composed of a real component and an imaginary component. Both components have to be preserved to reconstruct the source object; otherwise the twin and the main images will overlap with each other. However, from the practical point of view, a complex hologram is not desirable, as its data size is twice that of a real hologram and also difficult to display with optical means. To overcome this dilemma, we propose to include an off-axis planar reference wave $B(y)$ to $H_0(x,y)$, $H_1(x,y)$, and $H_2(x,y)$, and we take the real part of the result to give

$$H'_i(x,y) = \mathrm{RE}[B(y)H_i(x,y)], \quad i = 0, 1, \quad \text{or} \quad 2, \quad (4)$$

where $\mathrm{Re}[\cdot]$ represents the real part of the quantity being bracketed. $H'_i(x,y)$ is a real hologram that is only half the data size compared with that of $H_i(x,y)$. When $H'_i(x,y)$ is illuminated with an off-axis reconstruction planar beam, a virtual image and a real twin image will be reconstructed. However, these two images will be separated by an angle due to the use of the off-axis reference beam. The generated digital holograms will now be compressed with VQ.
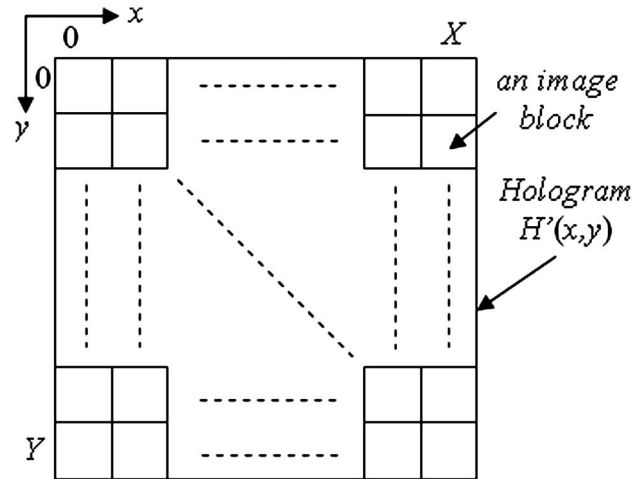
### B.   Compression of Generated Digital Fresnel Holograms with Vector Quantization

In this section, we explore the feasibility of applying VQ in compressing the three generated digital holograms obtained in the last section. To start with, each hologram $H'_i(x,y)$ is first partitioned into nonoverlapping square blocks of size $b \times b$, as shown in Fig. 1. In brevity, we shall drop the subscript $i$ in $H'_i(x,y)$, with the understanding that $H'(x,y)$ refers to one of the three color holograms. An enlarged view of the $j$th image block in the hologram is shown in Fig. 2. The pixels within the block are arranged sequentially in a row-by-row manner, starting from the first element $p_{j;0}$ to the last element $p_{j;b^2-1}$.

A one-dimensional vector $s_j$ is composed to represent the image block as $s_j^m = p_{j;m}|_{0 \leq m < b^2}$. Hence, the dimension of $s_j$, $k$, is equivalent to the number of pixels within a square block as given by $k = b^2$. A collection of all the block vectors forms the dataset $S = \{s_0, s_1, \ldots s_{M-1}\}_{s_j \in R^k|_{0 \leq j < M}}$. Suppose the horizontal and vertical extent of the hologram is $X$ and $Y$, respectively, there will be $Y/b$ blocks in each row and $X/b$ blocks in each column, constituting a total of $M = XY/b^2$ vectors in $S$. If $H'(x,y)$ is utilized to train a codebook, its corresponding set of vectors $S$ will be passed to the LBG algorithm described in Table 1 to generate codebook $C$.

On the other hand, if $H'(x,y)$ is the hologram to be compressed, the Euclidean distances between each vector in $S$ and all the code vectors in $C$ are computed with Eq. (2). Subsequently, each vector will be labeled by the index of code vector $C$, which is nearest in terms of the distance amongst all the members in the codebook. This results in a sequence of integer labels $L = \{l_0, l_1, \ldots l_{M-1}\}$, each corresponding to a unique block vector in $S$. The sequence of labels will be transmitted to the decoder to recover the hologram.

At the decoder, each received label $l_j$ will be taken to extract the corresponding vector in codebook $C$. The
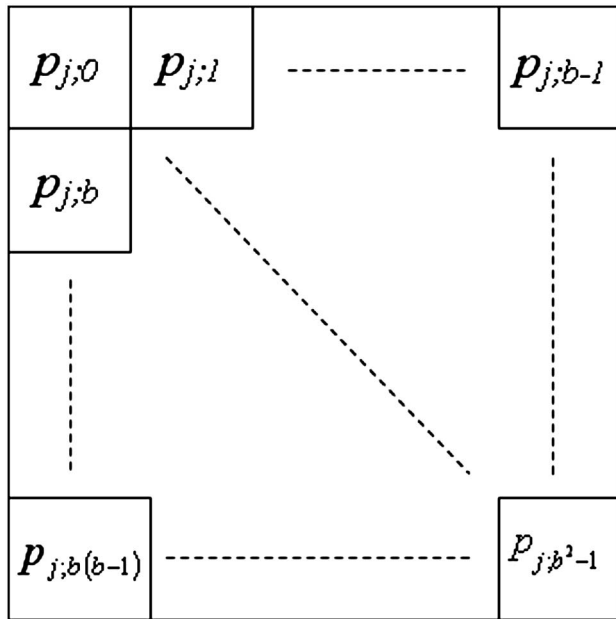
Fig. 2. Enlarged view of an image block in $H'(x,y)$.

retrieved vectors reconstruct, with a certain amount of distortion, the block vectors in the original dataset $S$. In other words, each decoded block vector in $S$ substitutes an image block in the decompressed hologram. As the decoding process only involves a lookup table operation, the amount of computation loading as well as the memory requirement, are negligible. However, it is mandatory that the codebook be present in both the encoder and the decoder. As the size of the codebook is large, it is undesirable to send it along with the compressed data. Consequently, we have adopted a fixed universal, out-training set codebook, which is present in both the encoder and the decoder.

## 4. Experimental Results

A color object image, shown in Fig. 3a, which is located at a distance of $z = 0.3\,\text{m}$ from the hologram, is employed to demonstrate the feasibility of the proposed method. A complex Fresnel color hologram $H(x,y)$ is generated based on Eq. (1) and with the settings listed in Table 2. From $H(x,y)$, three generated digital complex Fresnel color holograms are extracted and converted into real holograms by adding a reference beam $B(y)$ based on Eq. (4). Each pixel in a hologram plane is quantized to 8 bits (i.e., $Q = 8$). We employ a holographic display, which is modified from the Sony VPL-HW15 Bravia color projector

**Table 2. Settings of the Hologram Generation Process**

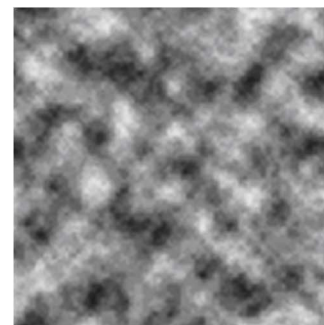| | |
|---|---|
| Hologram size | $2048 \times 2048$ pixels |
| Pixel size of the hologram | $7\,\mu\text{m}$ square |
| Quantization of pixel intensity in the hologram | 16 bit (8 bit for the real part and 8 bit for the imaginary part) |
| Wavelength of light | 650 nm |
| Angle of the off-axis plane reference wave $B(y)$ | $1.2°$ along the $y$ direction |

to display the color hologram. The projector is composed of three LCOS chips, each displaying one of the three primary color components (the red, green, and blue colors) of an input image. Each LCOS has a size of 1920 (horizontal) ×1080 (vertical) pixels, and a dot pitch of $7\,\mu\text{m}$. The projector lamp is removed, and the illumination is replaced with a white LED. In addition, the focusing lens is also removed so that the three LCOS displays can be viewed directly through a set of inherently built beam splitters. The optical reconstructed image of the hologram is shown in Fig. 3b. We observed that there is a certain blurriness and a change in color, which is mainly caused by the imperfection of the LCOS devices, the wide spectrum of the LED, misalignment of the optical axes of the three LCOS devices, and partly by the recording camera. In addition, there are some horizontal
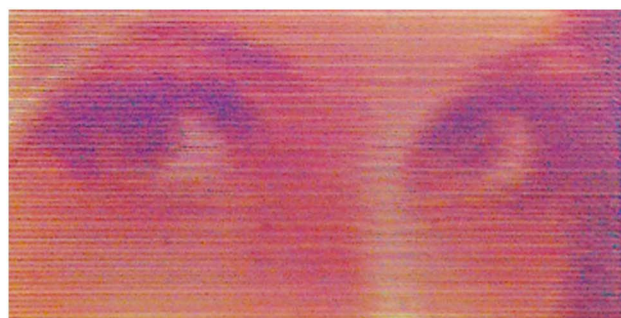


a



b



c

Fig. 3. (Color online) a, Color image positioned at $z = 0.3\,\text{m}$ from the hologram. b, Optical reconstructed image at $z = 0.3\,\text{m}$ with a LCOS display of the generated digital Fresnel color hologram representing the source image in Fig. 3a. c, Random color image for training the codebook.
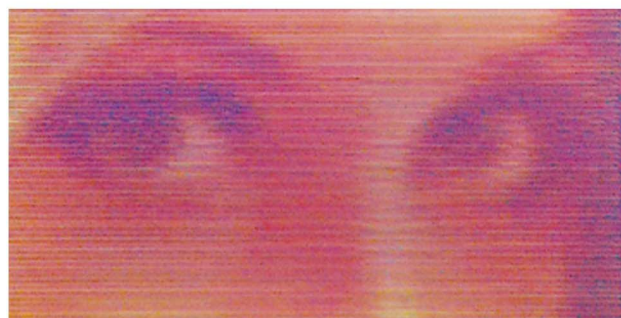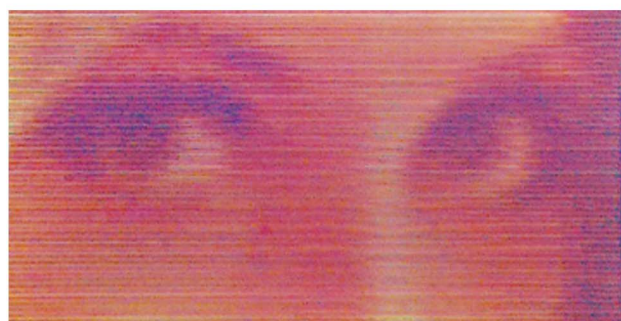
streaks generated by the partial interference of the remote part of the first-order diffraction. Apart from these artifacts, the reconstructed image is similar to the original one. Next, a Fresnel hologram representing the randomly generated image shown in Fig. 3c is generated based on identical optical settings. The random image is positioned at 0.3 m from the hologram. A codebook $C$, comprised of $N = 1024$ centroids, is trained with the hologram of the random image. Note that the same codebook $C$ is applied to compress all the three color holograms.

VQ is applied to encode the three real holograms of the color image in Fig. 3a with block height (and width) of $b = 16$, $b = 24$ and 32 (i.e., $k = b^2$), resulting in CRs of 204, 460, and 819 times, respectively. However, as the real hologram is only half the data size of the original complex hologram, the actual CR is doubled. The optical reconstructed images of the three compressed holograms on the LCOS display are shown in Figs. 4a–4c. It can be seen from the figures that, despite the large CR, in general the color and visual quality of the reconstructed images are
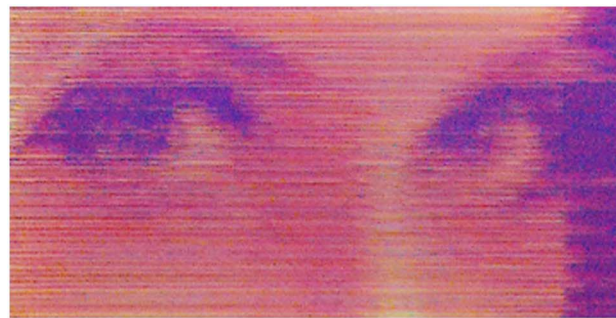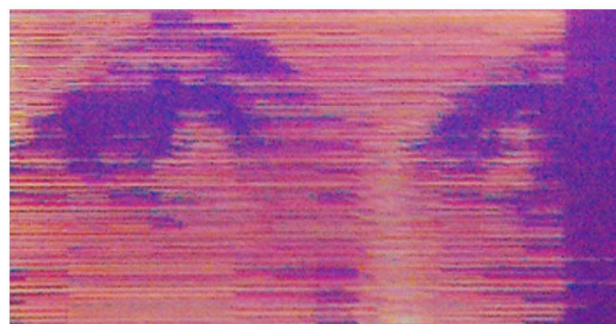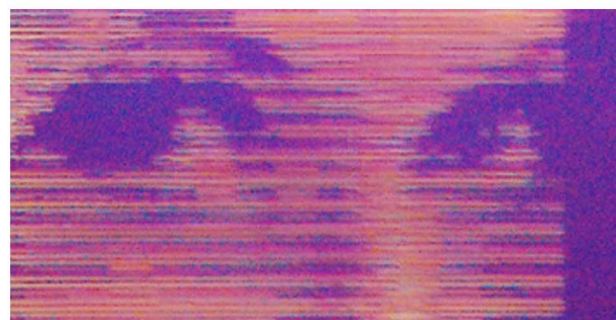


a



b



c

Fig. 4. (Color online) Optical reconstructed image of the generated digital Fresnel color holograms representing the image in Fig. 3a. Each of the holograms is compressed by 204 times (408 times with respect to the original complex hologram) b, Optical reconstructed image of the generated digital Fresnel color holograms representing the image in Fig. 3a. Each of the holograms is compressed by 460 times (920 times with respect to the original complex hologram). c, Optical reconstructed image of the generated digital Fresnel color holograms representing the image in Fig. 3b. Each of the holograms is compressed by 819 times (1638 times with respect to the original complex hologram).



a



b



c

Fig. 5. (Color online) a, Optical reconstructed image of the generated digital Fresnel color holograms representing the image in Fig. 3a. Each of the holograms is compressed with JPEG 2000 by 199 times (398 times with respect to the original complex hologram). b, Optical reconstructed image of the generated digital Fresnel color holograms representing the image in Fig. 3a. Each of the holograms is compressed with JPEG 2000 by 443 times (886 times with respect to the original complex hologram). c, Optical reconstructed of the generated digital Fresnel color holograms representing the image in Fig. 3b. Each of the holograms is compressed with JPEG 2000 by 820 times (1640 times with respect to the original complex hologram).

close to the original one shown in Fig. 3b. As the CR reaches 819 times, a slight change in the color and more blurriness (e.g., the circumference of the eyeball) are noted.

Our method is compared with holograms compressed with the JPEG 2000 technique. Reconstructed images from the color holograms (representing the color image in Fig. 3a), which have been compressed by JPEG 2000 by 419, 800, and 1600 times (with respect to the original complex hologram) are shown in Figs. 5a–5c, respectively.

The results in Fig. 5a show that at a CR of around 400 times, the quality of the reconstructed images is slightly inferior to that obtained with our proposed method. However, the degradation is more severe as the CR increases. When the CR reaches 920 and 1638 times, significant distortions are noted in the reconstructed images.

Next, the effect of noise contamination in our method is illustrated by randomly changing, for each of the VQ compressed holograms, on average 1 out of every 16 transmission indices (equivalent to 6% of the noise). The set of reconstructed images is illustrated in Figs. 6a–6c. It can be seen that, despite the noise contamination, the reconstructed images are still acceptable in both samples. Small patches of erroneous regions can be seen on the images (especially in Figs. 6b and 6c, where a larger block size is used), but these artifacts do not impose severe degradation to the visual quality of the images. As for the JPEG 2000 compression method, the compressed file cannot be decoded even if it is contaminated with a small amount of noise.

To further illustrate our proposed method, we have generated a sequence of holograms, each corresponding to a unique view of a globe that is rendered with a more complicated texture of the Earth image. The radius of the globe is around 0.005 m, and the front tip of the globe is located at 0.3 m from the hologram. A sequence of generated digital color Fresnel holograms, each corresponding to a unique view of the globe, is generated, and the optically reconstructed image of a single frame excerpt of the hologram sequence is shown in Fig. 7. Subsequently, we applied our proposed method to compress the image sequence by 408, 920, and 1638 times (with respect to the original size of the complex holograms) based on the same codebook generated from the hologram in Fig. 3c, and a single-frame excerpt for each case is shown in Figs. 8a–8c. A complete display of different views of the globe image corresponding to the original and the VQ compressed holograms can be found in View 1, View 2, View 3, and View 4. It can be seen from the excerpt, as well as in the reconstructed images, that, despite the complexity of the texture, the Earth images on the globe are similar to the original reconstructed images. However, with a higher CR of over 408 times, a mild increase in the blurriness as well as a slight fading of the color are noted.
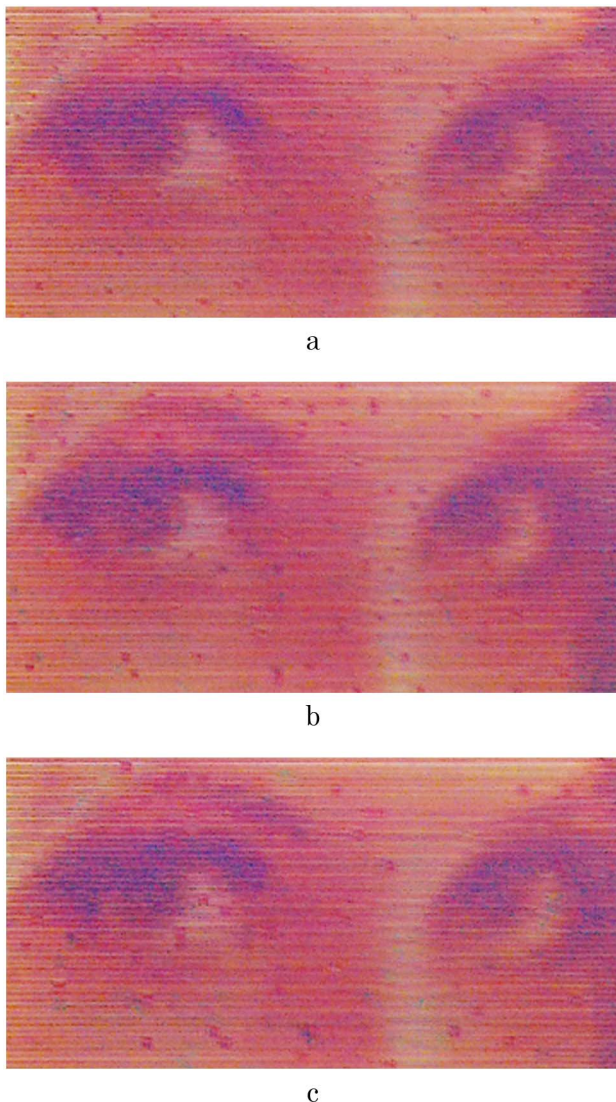


a

b

c

Fig. 6. (Color online) a, Image reconstruction of the holograms that have been compressed with the proposed method by 408 times, and contaminated with 6% noise on the encoded data. b, Image reconstruction of the holograms that have been compressed with the proposed method by 920 times, and contaminated with 6% noise on the encoded data. c, Image reconstruction of the holograms that have been compressed with the proposed method by 1638 times, and contaminated with 6% noise on the encoded data.
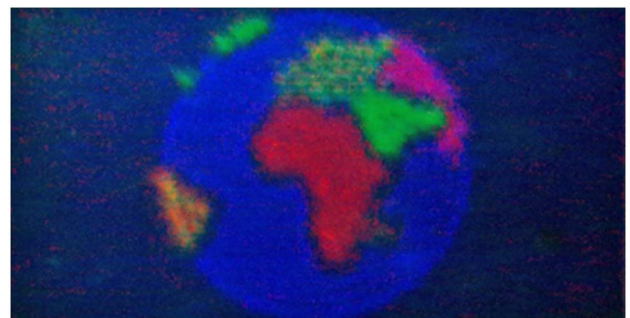


Fig. 7. (Color online) Optical reconstructed image with a LCOS display of an excerpt frame of the generated digital Fresnel color hologram sequence representing the globe image. Reconstructed images of other views of the hologram are shown in View 1.
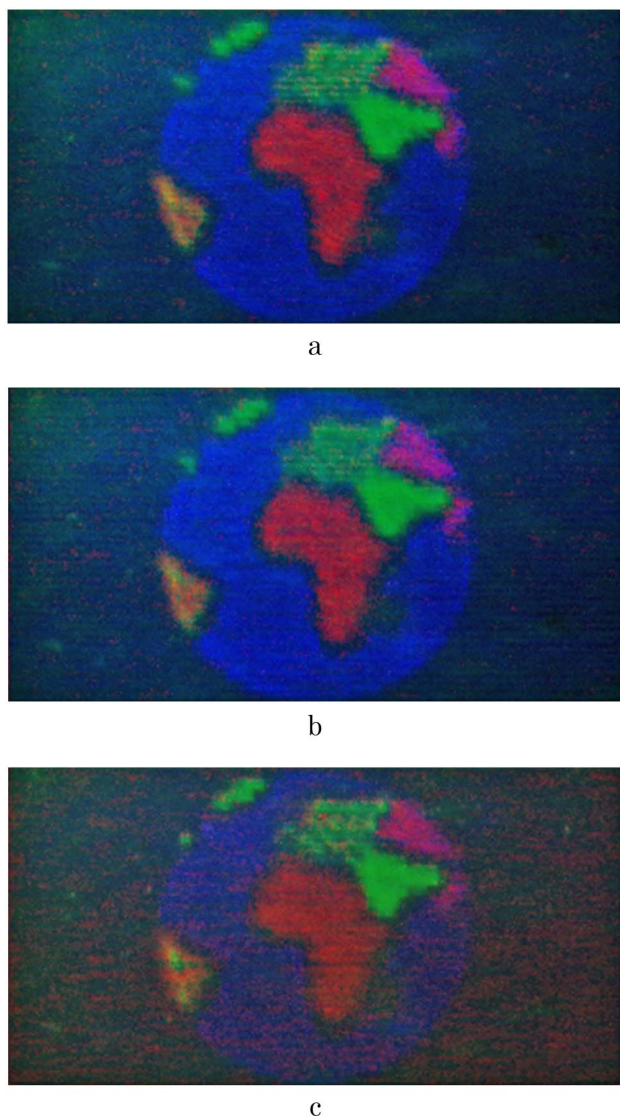
Fig. 8. (Color online) a, Optical reconstructed image of the generated digital Fresnel color hologram representing the globe image. Each of the generated digital Fresnel holograms is compressed by 204 times (408 times with respect to the original complex hologram). Reconstructed images of other views of the compressed holograms are shown in View 2. b, Optical reconstructed image of the generated digital Fresnel color hologram representing the globe image. Each of the generated digital Fresnel holograms is compressed by 460 times (920 times with respect to the original complex hologram plane). Reconstructed images of other views of the compressed holograms are shown in View 3. c, Optical reconstructed image of the generated digital Fresnel color hologram. Each of the generated digital Fresnel holograms is compressed by 819 times (1638 times with respect to the original complex hologram plane). Reconstructed images of other views of the compressed holograms are shown in View 4.

## 5. Conclusion

We have presented a method for reducing the data size of a generated digital complex color Fresnel hologram. The complex hologram is first separated into three holograms, each corresponding to one of the primary colors. Each hologram is then converted into a real hologram, resulting in a compression of two times. Subsequently, the generated digital holograms are encoded with VQ based on an out-training universal codebook that is present in both the transmission and the receiving ends. Experimental evaluations on the optical reconstructed images of the holograms compressed by our proposed method demonstrate a number of advantages over existing approaches. First, high compression of over 1600 times on the original complex Fresnel hologram can be realized with acceptable visual qualities. Second, the decoding process is computation free. Third, our method is highly resistant to noise contamination compared to the compression scheme by JPEG. We have also compared the visual quality of the reconstructed images of our proposed method with that of the JPEG 2000 technique, and we found that the degradation of JPEG 2000 is severe, especially at higher CRs. In addition, the compressed file resulting from JPEG 2000 cannot be decoded even if contaminated with small amount of noise. These favorable outcomes suggested that our proposed method could be applied in color holographic video systems where continuous sequences of generated digital color Fresnel holograms have to be compressed and transmitted at a low bit rate and eventually reconstructed with low-complexity hardware or software decoders at the receiving end.

In this paper, we only focused on the compression of amplitude-type digital generated holograms, which are displayed on an amplitude-only spatial light modulator (SLM). However, we believe that the method could be extended, with suitable modifications, to the compression of a phase-type hologram. As a phase-only SLM could improve the quality of the optical reconstructed images, compression of the phase-type hologram based on our proposed method is a potential research area for further investigation.

## References

1. T.-C. Poon, ed., *Digital Holography and Three-Dimensional Display: Principles and Applications* (Springer, 2006).
2. T.-C. Poon, "Three-dimensional television using optical scanning holography," J. Info. Disp. **3**, 12–16 (2002).
3. T.-C. Poon, "On the fundamentals of optical scanning holography," Am. J. Phys. **76**, 738–745 (2008).
4. K. Sasaki, E. Tanji, and H. Yoshikawa, "Data compression for holographic 3D image," J. Inst. Tele. Eng. Japan **48**, 1238–1244 (1994).
5. G. A. Mills and I. Yamaguchi, "Effects of quantization in phase-shifting digital holography," Appl. Opt. **44**, 1216–1225 (2005).
6. E. Darakis and J. J. Soraghan, "Use of Fresnelets for phase-shifting digital hologram compression," IEEE Trans. Image Proc. **15**, 3804–3811 (2006).
7. T. J. Naughton, Y. Frauel, B. Javidi, and E. Tajahuerce, "Compression of digital holograms for three-dimensional object reconstruction and recognition," Appl. Opt. **41**, 4124–4132 (2002).

8. A. E. Shortt, T. J. Naughton, and B. Javidi, "A companding approach for nonuniform quantization of digital holograms of three-dimensional objects," Opt. Express **14**, 5129–5134 (2006).

9. A. E. Shortt, T. J. Naughton, and B. Javidi, "Histogram approaches for lossy compression of digital holograms of three-dimensional objects," IEEE Trans. Image Proc. **16**, 1548–1556 (2007).

10. A. E. Shortt, T. J. Naughton, and B. Javidi, "Vector quantisation compression of digital holograms of three-dimensional objects," Proc. SPIE **5827**, 265–273 (2005).

11. P. W. Tsang, W. K. Cheung, and T.-C. Poon, "Low bit-rate compression of computer-generated Fresnel holograms based on vector quantization," in *Digital Holography and Three-Dimensional Imaging*, OSA Technical Digest (CD) (Optical Society of America, 2011), paper DTuD1.

12. R. M. Gray, "Vector quantization," IEEE ASSP Mag. **1**(2), 4–29 (1984).

13. Y. Linde, A. Buzo, and R. M. Gray, "An algorithm for vector quantizer design," IEEE Trans. Commun. **28**, 84–95 (1980).