

Compression of digital holograms for three-dimensional object reconstruction and recognition

Thomas J. Naughton, Yann Frauel, Bahram Javidi, and Enrique Tajahuerce

We present the results of applying lossless and lossy data compression to a three-dimensional object reconstruction and recognition technique based on phase-shift digital holography. We find that the best lossless (Lempel–Ziv, Lempel–Ziv–Welch, Huffman, Burrows–Wheeler) compression rates can be expected when the digital hologram is stored in an intermediate coding of separate data streams for real and imaginary components. The lossy techniques are based on subsampling, quantization, and discrete Fourier transformation. For various degrees of speckle reduction, we quantify the number of Fourier coefficients that can be removed from the hologram domain, and the lowest level of quantization achievable, without incurring significant loss in correlation performance or significant error in the reconstructed object domain. © 2002 Optical Society of America

OCIS codes: 100.6890, 090.1760, 100.2000, 100.0100, 100.5010, 110.4280.

1. Introduction

Digital holography has been used for three-dimensional (3D) measurement and inspection. Recently, the concept of two-dimensional (2D) pattern recognition^{1–3} has been extended to 3D objects,^{4–11} with digital holograms, created through phase-shift interferometry (PSI),^{12,13} providing the means for one such set of 3D object recognition techniques.^{9,10} Each digital hologram encodes multiple views of the object from a small range of angles. A particular view of the object can be constructed by extracting the appropriate window of pixels from the hologram and applying a numerical propagation technique.^{9,13} These real-valued views could be combined as a composite filter¹⁰ or in a filter bank. In advance of knowing which of the 4 M pixels are required for particular views, each hologram requires 65 Mbytes of storage in its native double precision format (5 s of transmission time over a 100 Mbit/s network connection). This is too slow for real-time object reconstruction or

recognition, and is impractical for any type of holographic video streaming. We would like to compress¹⁴ these holograms for more efficient storage and transmission. Because one of our primary applications for digital holography is 3D pattern recognition,^{9,10} we choose normalized cross-correlation as one of our metrics for reconstruction integrity. In anticipation of wider uses for digital holography, the error in the reconstructed object is also measured. We are not directly interested in compression noise or artifacts (such as blocking effects) that appear in the decompressed hologram, only how compression losses affect object reconstruction.

When either of the two steps in holography, recording or reconstruction, are performed digitally, the process has been referred to as computer holography. Synthesis of holograms by computer,^{15,16} digital reconstruction of optically recorded objects,^{17,18} and both steps performed together as part of a simulation study¹⁶ have been demonstrated. In this paper, our holograms are created through PSI. We also call these digital holograms, and introduce a third step, that of digital compression and decompression. This work combines aspects of 3D object recognition, correlation performance under lossy compression conditions, speckle noise compression, and compression of digital holograms. Hologram compression differs from image compression principally because our holograms store 3D information in complex-valued pixels, and secondly because of the inherent speckle content.

Algorithms for 3D data compression do exist, such

T. J. Naughton (tom.naughton@may.ie) is with the National University of Ireland, Department of Computer Science, Maynooth, Republic of Ireland. Y. Frauel and B. Javidi (bahram@engr.uconn.edu) are with the University of Connecticut, Electrical and Computer Engineering Department, U-157, Storrs, Connecticut 06269-2157. E. Tajahuerce is with the Universitat Jaume, Departament de Ciències Experimentals, E-12080 Castelló, Spain.

Received 24 September 2001; revised manuscript received 4 February 2002.

0003-6935/02/204124-09\$15.00/0

© 2002 Optical Society of America

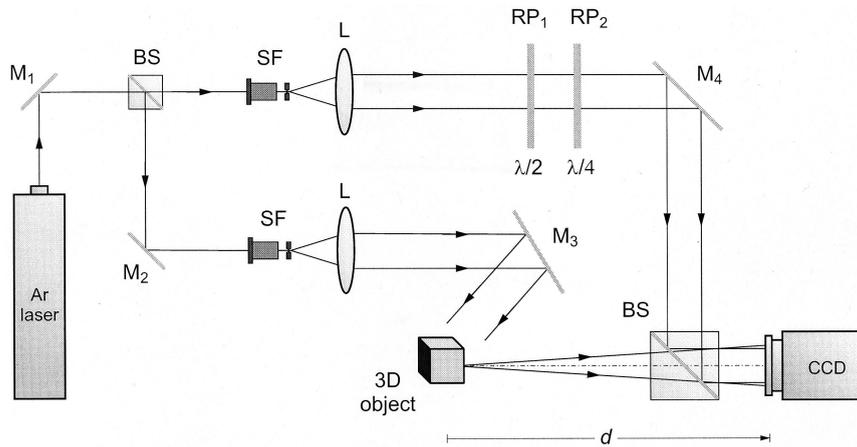


Fig. 1. Experimental setup for PSI: M, mirror; BS, beam splitter; SF, spatial filter; L, lens; RP, retardation plate.

as the wavelet technique of Bilgin *et al.*¹⁹ or the vector-quantization technique of Qian *et al.*²⁰ However, these techniques are not suitable. Our hologram is actually a 2D (complex-valued) image and differs from a multispectral image, for example, in that there is little correlation between the real and imaginary components of each pixel.

Some work has been done on object recognition under lossy compression conditions. Ewing and Woodruff²¹ have examined the human ability to recognize objects that have undergone lossy JPEG (Joint Photographic Experts Group) and fractal-based compression, and found JPEG superior. Similar subjective tests have been conducted by Morioka *et al.*²² in a medical application where JPEG outperformed wavelet-based compression. Farn and Goodman²³ have examined degradation in correlation performance due to phase quantization, and Mahalanobis and Daniell²⁴ have shown how to combine wavelet compression and correlation-filtering processes.

Compression of noise, including speckle noise, has been investigated. Vago *et al.*²⁵ have studied the compression of data from a speckle-interferometry application and found that a low-pass filter in the Fourier domain served to selectively reduce unwanted speckle noise. Shahnaz *et al.*²⁶ have applied baseline JPEG (the standard JPEG implementation) compression to images with speckle noise and found that its compression performance suffers greatly in the presence of speckle. These techniques have been applied to real-valued images with simulated speckle noise. Murtagh *et al.*²⁷ have investigated compression of real data from astronomical images. By modeling and removing noise they have been able to increase compression rates by a factor of six over baseline JPEG. Wyrowski and Bryngdahl²⁸ have investigated removing speckle from digital holograms. We have found that median filtering provides a good tradeoff between image detail and robustness to speckle noise.

Holographic data compression is a new field of research. Nomura *et al.*²⁹ investigated quantiza-

tion of real-valued (rather than complex-valued) holograms. Compression rates of 4.0 [see Eq. (3) for definition of compression rate] for real-valued holograms of binary 2D inputs, with minimal reconstruction error, were achieved. They found that baseline JPEG performs poorly for such holograms. However, previous studies with 2D images have shown that careful manipulation of the quantized cosine coefficients can improve image quality over baseline JPEG.³⁰ We adopt this strategy and also apply quantization directly to the complex-valued holographic pixels. Phase quantization^{31,32} has been applied successfully to Fourier and holographic data in the past.

In Section 2 we describe phase-shift digital holography and present our experimental setup. In Section 3 we apply standard lossless data compression techniques to the digital holograms. The simplest form of lossy compression, that of resampling or hologram resizing, is examined in Section 4, and quantization is examined in Section 5. In Section 6 we use a discrete Fourier transform (DFT) technique to selectively remove Fourier coefficients from 8×8 pixel nonoverlapping blocks of the hologram.

2. Phase-Shift Digital Holography

We record digital holograms with an optical system based on a Mach-Zehnder interferometer (see Fig. 1). A linearly polarized Argon ion (514.5 nm) laser beam is divided into object and reference beams, both of which are expanded and spatially filtered. The first beam illuminates a reference object placed at a distance $d = 350$ mm from a 10-bit 2028×2044 pixel Kodak Megaplus CCD camera. We refer to the complex amplitude distribution in the plane of the object as $U_0(x, y)$. The reference beam passes through half-wave plate RP_1 and quarter-wave plate RP_2 . The linearly polarized beam can be phase modulated by rotating the two retardation plates. Through permutation of the fast and slow axes of the plates we can achieve phase shifts of $0, \pi/2, \pi,$ and $3\pi/2$. The reference beam combines with the light diffracted

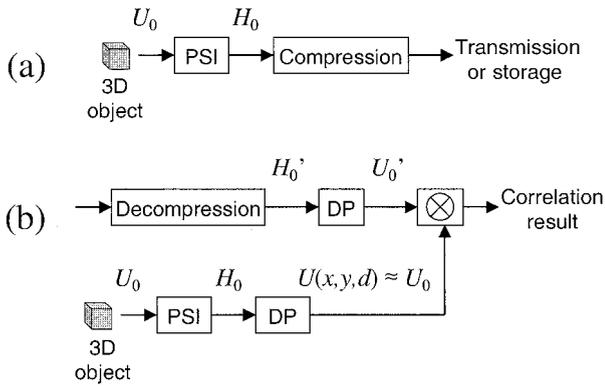


Fig. 2. Illustration of the problem statement: (a) digital hologram H_0 must be compressed and transmitted such that (b) decompressed and reconstructed U_0' compares closely with the approximation of the original complex amplitude distribution $U(x, y, d)$. PSI, image capture and interferometry stage; DP, digital propagation (reconstruction) stage; \otimes , normalized cross-correlation operation.

from the object and forms an interference pattern in the plane of the camera. At each of the four phase shifts we record an interferogram. We use these four real-valued images to compute the camera-plane complex field $H_0(x, y)$ by PSI.¹² We call this computed field a digital hologram. Our holograms have dimensions 2028×2044 pixels and are originally in floating point representation with 8 bytes of amplitude information and 8 bytes of phase information for each pixel.

A digital hologram $H_0(x, y)$ contains sufficient amplitude and phase information to reconstruct the complex field $U(x, y, z)$ in a plane in the object beam at any distance z from the camera.^{9,13} This can be calculated by use of the Fresnel-Kirchhoff formula

$$U(x, y, z) = H_0(x, y) \star h(x, y, z), \quad (1)$$

where

$$h(x, y, z) = -\frac{i}{\lambda z} \exp\left(i \frac{2\pi}{\lambda} z\right) \exp\left[i\pi \frac{(x^2 + y^2)}{\lambda z}\right] \quad (2)$$

is the point spread function for free space, λ is the wavelength of the illumination, and \star denotes a convolution operation. At $z = d$, and ignoring errors in digital propagation due to discrete space (pixelation) and rounding, the digital reconstruction $U(x, y, z)$ approximates $U_0(x, y)$. Furthermore, as with conventional holography,³³ a windowed subset of the hologram can be used to reconstruct a particular view of the object.

The problem we address can be stated as follows: A digital hologram H_0 of some object U_0 is to be compressed and transmitted from sender to receiver (see the illustration in Fig. 2). At the receiver, the hologram is decompressed as H_0' and an object U_0' reconstructed by numerical propagation. We wish to determine the most effective way of compressing H_0 such that U_0' is reconstructed with minimal rms difference and with high correlation with U_0 . Full holograms will be compressed, without windowing.

3. Lossless Data Compression

Lossless data compression techniques are used in situations where the data must be faithfully decompressed, such as in text compression. If we use lossless techniques we are assured that U_0' will be identical to U_0 , apart from rounding and pixelation errors. The set of 3D objects used in the compression experiments is shown in Fig. 3. The digital holograms were treated as binary data streams. Four of the most common industry-standard compression techniques were chosen: Huffman coding,³⁴ Lempel-

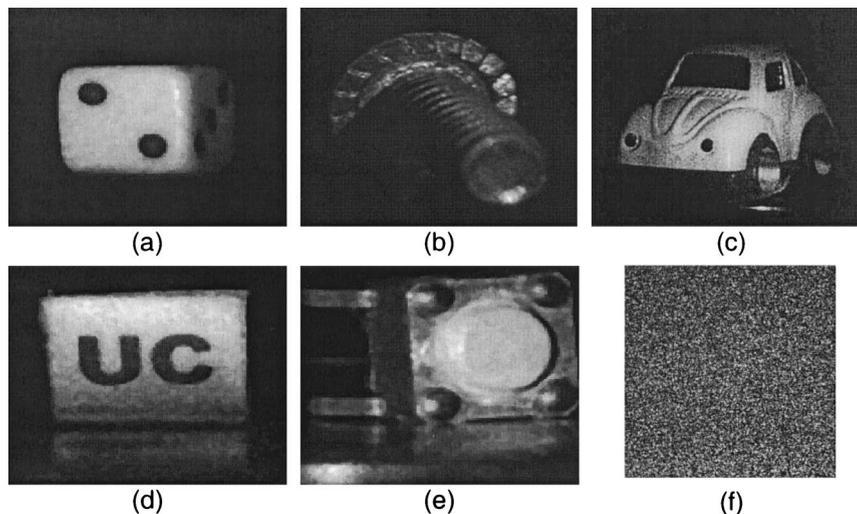


Fig. 3. Set of holograms used in these experiments: (a) through (e) are the amplitudes of the reconstructed wavefronts for holograms No. 1 through No. 5, respectively. Image (f) shows the amplitudes of an example 512×512 subset of digital hologram No. 1.

Table 1. Compression with LZ77, LZW, Huffman, and BW^a

Hol. No.	Size (kB)	LZ77 (kB)	LZW (kB)	Huff. (kB)	BW (kB)	LZ77 c.r. ^b	LZW c.r.	Huff. c.r.	BW c.r.
1	64,769	52,038	64,769	62,236	37,020	1.24	1.00	1.04	1.75
2	64,769	62,353	64,769	62,298	63,309	1.04	1.00	1.04	1.02
3	64,642	32,718	54,766	61,784	17,698	1.98	1.18	1.05	3.65
4	64,769	54,923	64,769	62,262	40,596	1.18	1.00	1.04	1.60
5	64,769	53,608	64,769	62,267	37,923	1.21	1.00	1.04	1.71
Averages:						1.33	1.04	1.04	1.95

^aThe hologram is treated as a single binary data stream with alternating amplitude and phase angle components.

^bc.r., compression rate.

Ziv coding³⁵ (LZ77), Lempel–Ziv–Welch coding³⁶ (LZW), and Burrows–Wheeler coding³⁷ (BW).

Huffman coding,³⁴ an entropy-based technique, is one of the oldest and most widely used compression methods. It replaces each symbol in the input by a codeword, assigning shorter codewords to more frequent symbols. The LZ77 algorithm³⁵ takes advantage of repeated substrings in the input data. In contrast to Huffman coding, a variable length string of input symbols is replaced by a fixed-size codeword (a reference to the previous occurrence of that string). LZW³⁶ is a refinement of LZ77. It maintains a dictionary (or lookup table) of variable-sized codewords and is less biased toward local redundancy. The more recent BW algorithm³⁷ transforms its input through a sorting operation into a format that can be compressed very effectively using standard techniques (in our case, Huffman coding).

The result of using these lossless algorithms is shown in Table 1. Treating the holograms (each a sequence of pairs of amplitude and phase values) as binary data streams achieves compression rates in the range [1.0, 3.65], where compression rate r is calculated from

$$r = \frac{\text{uncompressed size}}{\text{compressed size}}, \quad (3)$$

and where a rate of 1.0 was used when no compression was achieved, or when the processed hologram was actually larger in size. By the term compression rate we indicate the number of bits of uncompressed data that are effectively communicated with a single bit of compressed data. As shown in Table 1, on average, each bit of compressed data encodes 1.33 bits of uncompressed holographic data with

LZ77, 1.04 bits with LZW and Huffman, and 1.95 bits with BW. Given their size, it is worrying how poorly the holograms compress using these techniques. Their performance is possibly due to the noisy influence of speckle in the hologram [see Fig. 3(f)].

For some compression algorithms, the way data is represented can have a great impact on how well it will be compressed. We investigate three additional representations (or intermediate codings) for the digital holographic data. First, the pixel values are divided into separate amplitude and phase data streams to exploit the possible redundancy between neighboring amplitude values. This produced slight but consistent improvements in compression rate for each hologram. For the other representations, the holographic data was transformed from (amplitude, phase) domain to the equivalent, neglecting rounding errors, (real, imaginary) domain. Treating the hologram as a single binary data stream of pairs of real and imaginary values, or as separate streams for real and imaginary values, results in a further improvement. Results for the latter case are shown in Table 2, where it is evident that both the real and imaginary streams are equally difficult to compress. If a higher rate is required then a lossy form of compression will have to be applied. Lossy systems are investigated in Section 4.

4. Compression by Resampling

The criteria for grading the performance of lossy compressors will not be the same as in conventional image compression. The errors introduced into the digital hologram as a result of lossy compression are not of direct concern; errors in the reconstructed object, loss of viewing angle, and so on, are of most

Table 2. Compression with LZ77, LZW, Huffman, and BW^a

Hol. No.	Size (kB)	LZ77 (kB)	LZW (kB)	Huff. (kB)	BW (kB)	LZ77 c.r. ^b	LZW c.r.	Huff. c.r.	BW c.r.
1	64,769	9356 + 9127	8931 + 8791	13,899 + 13,079	6389 + 6263	3.50	3.65	2.40	5.12
2	64,769	29,890 + 29,752	32,385 + 32,385	30,987 + 31,009	24,413 + 24,163	1.09	1.00	1.04	1.33
3	64,642	7254 + 8180	7747 + 8318	27,154 + 27,282	4512 + 5199	4.19	4.02	1.19	6.66
4	64,769	9465 + 9125	9047 + 8805	13,815 + 13,027	6512 + 6283	3.48	3.63	2.41	5.06
5	64,769	9227 + 9061	8839 + 8718	13,439 + 13,082	6336 + 6265	3.54	3.69	2.44	5.14
Averages:						3.16	3.20	1.90	4.66

^aThe hologram is treated as two separate real and imaginary data streams.

^bc.r., compression rate.

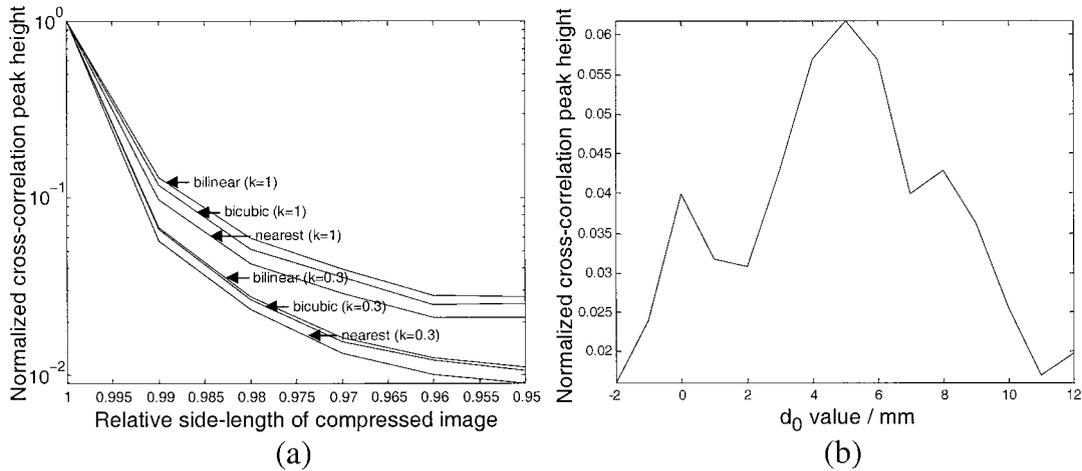


Fig. 4. Resampling of the digital hologram: (a) a plot of hologram side length (relative to the original side length) against linear and nonlinear correlation performance for three interpolation strategies; (b) searching the z axis for an appropriate d_0 offset for a hologram resize of 0.97 and bilinear interpolation.

interest. The simplest and most common form of lossy compression is that of resampling. We find that digital holograms are unduly sensitive to resampling.

In Fig. 4(a), hologram No. 1 is resized by use of three different interpolation strategies prior to reconstruction of U_0' . The reconstructed object U_0' is then correlated with U_0 . The figure shows normalized cross-correlation peak height on a \log_{10} scale plotted against hologram side length (relative to the original side length). The plot shows peak heights for both linear and nonlinear correlation (where $k = 0.3$ is the k th law nonlinearity³⁸). Resizing the hologram to 0.97 of its previous side length causes the normalized cross-correlation peak height to fall dramatically to the order of 10^{-2} of its former height.

Resizing the hologram will cause the object to be formed at a different distance from the hologram plane. To compensate for this, z in Eq. (2) becomes $z = d + d_0$, where d , as before, is the original object distance from the camera and d_0 is an offset. Figure 4(b) shows the result of a search for an appropriate d_0 value. In this plot, a resizing to 0.97 of the original image side length requires a d_0 offset of +5.0 mm for maximum correlation. Even then, the correlation is poor (0.06 when normalized). This sensitivity is due in some part to speckle. The unique speckle pattern constructed in the object plane by the hologram causes a large correlation normalization factor. With the slightest modification to the digital hologram a completely different speckle pattern is reconstructed, and the value for normalized cross correlation drops considerably. In order to reduce the effect of speckle we discard the phase information in the reconstructed object wavefront and apply a median filtering operation. Our justification for keeping only the object plane amplitude information is based on this information's dependency on both the amplitude and phase of the hologram plane: If the amplitude information in the object plane has been

reconstructed correctly, this indicates that sufficient amounts of both amplitude and phase information were preserved in the hologram plane during compression. The appropriate level of median filtering will be dependent on the application. As the level of median filtering increases the gross structure of the object is enhanced and detail is removed.

To avoid the need to search for an appropriate d_0 offset every time, holograms are returned to full size after the loss due to compression is introduced. (We assume that this resizing operation does not introduce additional error.) The reconstructed object amplitudes were evaluated in terms of normalized cross-correlation peak height (with $k = 1$) and normalized rms difference D , calculated from

$$D = \frac{1}{P_{U_0}} \left[\frac{1}{N_x N_y} \sum_{m=0}^{N_x-1} \sum_{n=0}^{N_y-1} |U_0(m, n) - U_0'(m, n)|^2 \right]^{1/2}, \quad (4)$$

where (m, n) are discrete spatial coordinates in the object plane, and N_x and N_y are the number of samples in the x and y directions, respectively. P_{U_0} is the power, per pixel, in the uncompressed object amplitudes and is defined as

$$P_{U_0} = \left[\frac{1}{N_x N_y} \sum_{m=0}^{N_x-1} \sum_{n=0}^{N_y-1} |U_0(m, n)|^2 \right]^{1/2}. \quad (5)$$

Figure 5 contains plots of normalized RMS difference and normalized cross-correlation peak height for hologram No. 1, for three different interpolation strategies (nearest neighbor, bilinear, and bicubic), and with or without 11×11 pixel median filtering. Comparing with Fig. 4(a), these plots show that when only the amplitude of the reconstructed object is taken into account, the hologram is more tolerant to resampling. Resampling, if it is acceptable for a

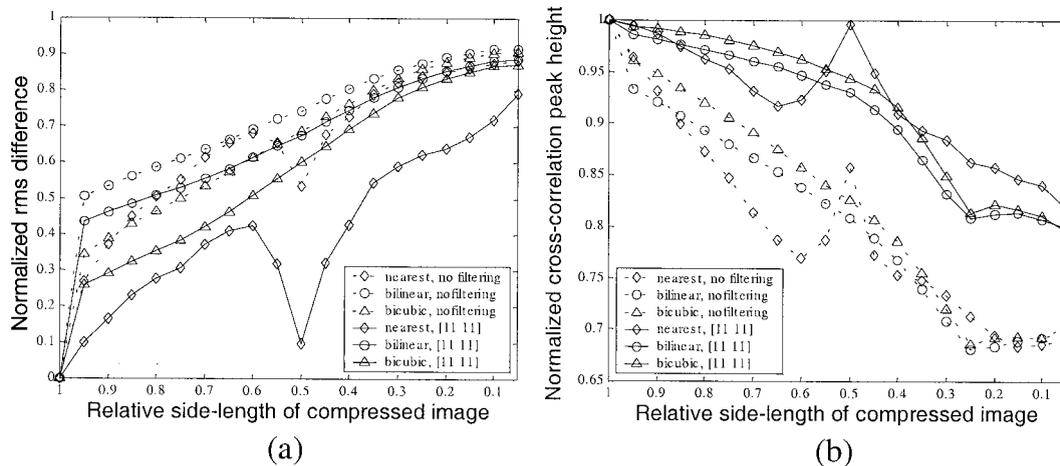


Fig. 5. Resizing hologram No. 1 with three different interpolation strategies and then using only the amplitude information in the reconstructed object plane. Plots for (a) normalized RMS difference, and (b) normalized cross-correlation peak height, show the effect of resizing both without filtering and with 11×11 pixel median filtering.

particular application, can therefore improve the compression rate. Furthermore, because resampling simply reduces the number of pixels and (possibly) changes their values rather than compresses the underlying data format, the lossless techniques of Section 3 can be applied as a final step.

The large artifacts for nearest-neighbor interpolation in each plot of Fig. 5 occur at exactly a side-length resizing of 0.5 (when 4 pixels are compressed to 1). Figure 5(a) shows that even in the presence of measures introduced to reduce the effects of speckle, such as abandoning the phase and median filtering, the rms errors grow quickly with resizing. Nearest-neighbor interpolation is found to introduce the fewest errors in reconstruction. For correlation, a hybrid bicubic-nearest-neighbor strategy would seem to achieve best performance for all levels of median filtering. Taking a quantitative example from Fig. 5, if a normalized cross correlation of at least 0.98 is required (and 11×11 pixel median filtering is acceptable), then bicubic interpolation can resize to a

side length of 0.75. This reduces the number of pixels by a factor of 1.78, which when combined with BW gives an average compression rate of 8.29. This compression does, however, come at the cost of an unsatisfactory 0.38 normalized rms error. The technique of resampling achieves its best performance at a side-length resizing of 0.5, 11×11 -pixel filtering, and nearest-neighbor interpolation, where a compression rate of 18.6 can be averaged.

These results show that arbitrary resampling (as used, for example, by video streaming applications to maintain frame rates) is not an effective technique for digital holograms. This result is not surprising because in digital holography systems the holographic microfringes are usually of roughly the same size as the camera pixels. A downsampling thus easily results in an undersampling.

5. Quantization

The technique of pixel-value quantization is analogous to nearest-neighbor resampling in the spatial

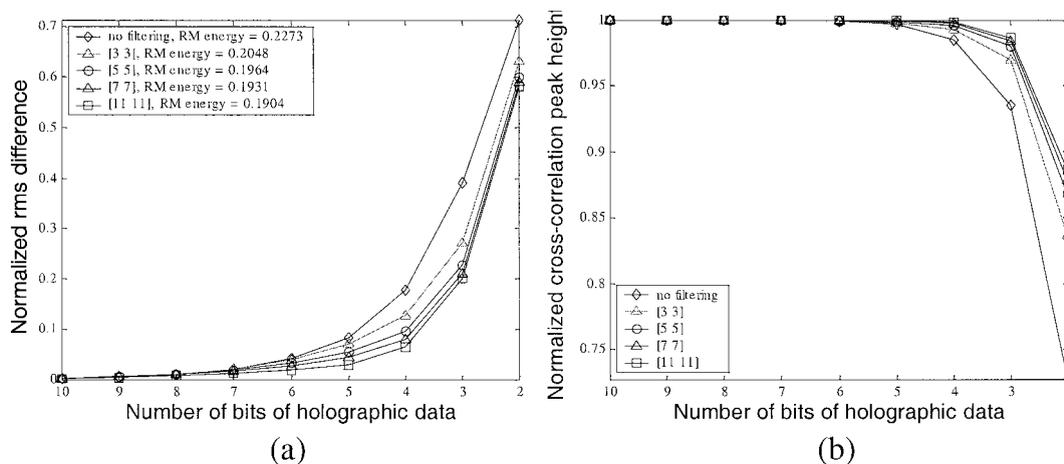


Fig. 6. Quantization with hologram No. 1: (a) normalized rms difference, and (b) normalized cross-correlation peak height, plotted against the number of bits in each of the real and imaginary values.

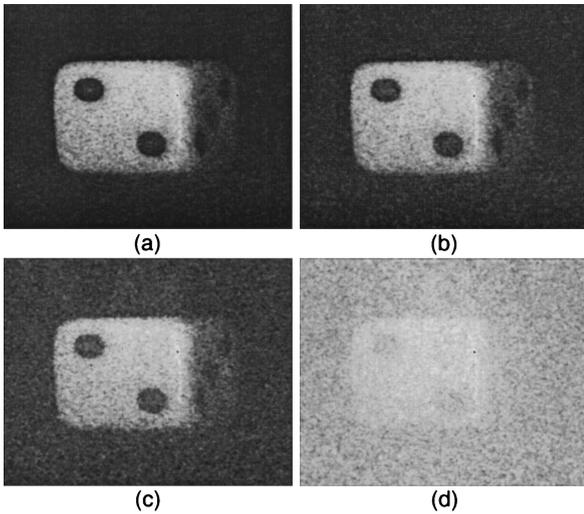


Fig. 7. Reconstructed amplitudes for hologram No. 1 for various numbers of quantization levels and with 11×11 pixel median filtering: (a) 4 bits (15 quantization levels); (b) 3 bits (7 quantization levels); (c) 5 quantization levels; (d) 2 bits (3 quantization levels).

domain. By reducing the number of possible values (or levels) available to each pixel we reduce the number of bits required to describe it. Such a technique was anticipated to provide a compression factor of at least 6.4 (8 bytes/10 bits) with minimal losses because images with only 10 bits of intensity resolution were used in the PSI stage. Figure 6(a) shows a plot of normalized rms difference against number of bits per data value for hologram No. 1, and for each of five median filtering neighborhoods from 1×1 (no filtering) to 11×11 pixels. The digital holograms were stored in real-imaginary format; each holographic pixel requires two such data values. Quantization levels were chosen to be symmetrical about zero; as a result n bits encode $2^n - 1$ levels. For example, two bits encode levels $\{-1, 0, 1\}$, 3 bits encode levels $\{-3, -2, -1, 0, 1, 2, 3\}$, and so on. Figure 6(b) shows a

plot of normalized cross-correlation peak height ($k = 1$) and Fig. 7 shows the reconstructed object amplitudes for selected numbers of quantization levels.

For pattern recognition, as few as 4 or 5 bits (compression rates of 16.0 or 12.8, respectively) need to be retained in each of the real and imaginary components. If errors in visual appearance are of primary concern, that threshold should be raised to 6 or 7 bits (compression rates of 10.7 or 9.1, respectively) for low (<0.02) reconstruction losses. For example, with 4 bits, and with moderate amounts of median filtering, a normalized cross-correlation value of greater than 0.98 and a normalized rms error of less than 0.1 can be achieved, corresponding to a compression rate of 16.

6. DFT-based Compression

Quantization appears to be a promising technique for the compression of digital holograms. The JPEG algorithm also performs most of its compression during a quantization stage. It uses the discrete cosine transform to allow it to perform quantization in the spatial frequency domain. The JPEG standard is defined for real-valued images only, but we adapt it to holograms, taking instead the DFT of each nonoverlapping block of 8×8 pixels. Rounding errors aside, this will not reduce the amount of information in the image but will tend to concentrate the majority of the hologram information into a few DFT coefficients in each block. By quantizing (or setting to zero, in our case) particular coefficients in each block we reduce the length of its bit description and thus allow an entropy coder (such as Huffman or a run-length technique) to further compress the coefficients. In our technique, we set a fixed number of the smallest DFT coefficients in each block to zero. This involves sorting the values in each 8×8 DFT and setting the $(64 - n)$ lowest-valued coefficients to zero, where n is a positive integer in the range $[1, 64]$ denoting the number of DFT coefficients to be retained in each block. Sorting of the DFT coefficients was per-

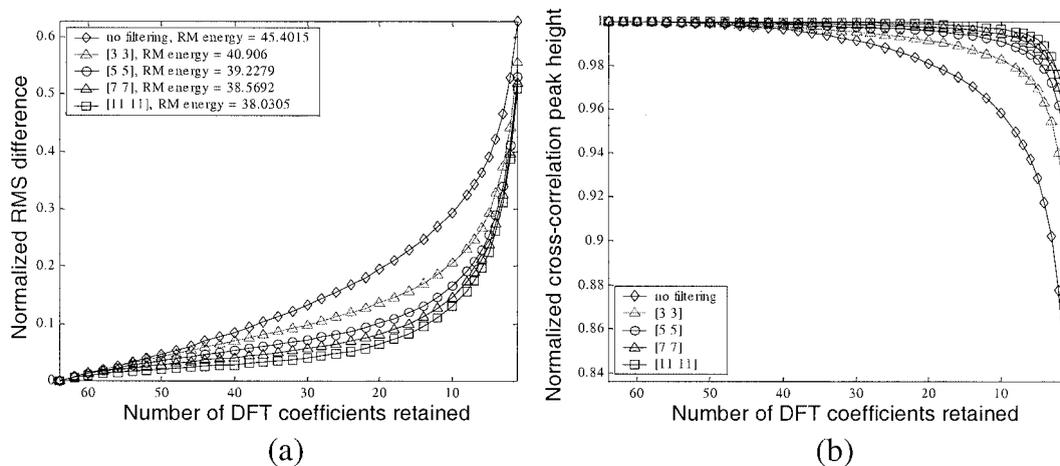


Fig. 8. Removing Fourier coefficients from hologram No. 1: (a) normalized rms difference, and (b) normalized cross-correlation peak height, as functions of DFT coefficients retained, for various median filter neighborhoods.

formed by amplitude first, and then (if necessary) by phase angle. Upon decompression (a blockwise inverse DFT) the hologram was used to construct object U_0' for comparison with U_0 . Linear correlation ($k = 1$) was also employed. Once again, the phase of the object wavefront was discarded and the amplitude median filtered prior to comparison to lessen the effects of speckle.

Figure 8 shows normalized rms difference and normalized cross-correlation peak height for hologram No. 1, for various values of n , and for different median filtering neighborhoods. With 11×11 filtering, as many as 92% of the DFT coefficients can be removed with minimal loss in correlation performance (less than 1%). This corresponds to a compression rate of 12.8. In this case, the normalized rms error would be 0.22. For rms errors of less than 0.1 no more than 78% of the coefficients can be removed giving a compression rate of 4.6. The remaining nonzero-valued DFT coefficients are each stored with 8 bytes, and so the quantization technique of Section 5 or the lossless techniques of Section 3 could be used to further increase these compression rates.

7. Conclusion

We have investigated various techniques for the compression of digital holograms created by PSI. With industry-standard lossless data compression techniques an average lossless compression rate of 4.66 can be expected. This rate was achieved by an intermediate coding of separated real and imaginary components before application of the BW algorithm. Lossy resampling techniques, combined with phase-removal and median filtering to lessen the effects of speckle, were examined in terms of reconstruction error and normalized correlation peak height. Hologram resampling resulted in a high degradation in reconstructed image quality, but for resizing to a side length of 0.5, and in the presence of a high degree of median filtering, a compression rate of 18.6 could be achieved. If the metric is correlation peak height, a hybrid bicubic-nearest-neighbor interpolation strategy would seem preferable, while nearest-neighbor interpolation alone provides lowest (albeit still high) reconstruction errors. Quantization proved to be a very effective technique. Each real and imaginary component could be reduced from its original 8 bytes to 4 bits while maintaining a high correlation peak and an acceptable reconstruction error, resulting in a compression rate of 16. The technique based on the removal of DFT coefficients achieves approximate compression rates of up to 12.8 for good cross correlation, and up to 4.6 for reasonable reconstruction integrity. It is anticipated that this can be improved further by applying lossless compression or quantization to the remaining DFT coefficients. Based on a compression rate of 10.7 (6-bit quantization), and without exploiting inter-frame redundancy, complex-valued holographic video frames of dimensions 640×640 pixels could be streamed over a 100 Mbit/s connection at a rate of 20 Hz or frames with 1024×1024 pixels at 8 Hz.

Compression will permit more efficient storage of digital holograms. In order to be useful for a real-time object recognition system, our compression strategies will have to be shown to admit efficient algorithms that make it advantageous to spend time compressing and decompressing rather than transmitting the original data. Efficient solutions must also be developed for each of the other stages, namely the recording of four interferograms (during which the object must be stable), phase-shift interferometry, reconstruction, and recognition.

The authors wish to thank the reviewers for their helpful and insightful suggestions. The first author acknowledges support from Enterprise Ireland, the Department of Electrical and Computer Engineering at the University of Connecticut and the Department of Computer Science at the National University of Ireland at Maynooth.

References

1. A. Vander Lugt, "Signal detection by complex spatial filtering," *IEEE Trans.* **IT-10**, 139–145 (1964).
2. J. W. Goodman, *Introduction to Fourier Optics*, 2nd ed. (McGraw-Hill, New York, 1996).
3. A. D. McAulay, *Optical Computer Architectures* (Wiley, New York, 1991).
4. A. Pu, R. F. Denkwalter, and D. Psaltis, "Real-time vehicle navigation using a holographic memory," *Opt. Eng.* **36**, 2737–2746 (1997).
5. R. Bamler and J. Hofer-Alfeis, "Three- and four-dimensional filter operations by coherent optics," *Opt. Acta* **29**, 747–757 (1982).
6. J. Rosen, "Three-dimensional joint transform correlator," *Appl. Opt.* **37**, 7538–7544 (1998).
7. J. J. Esteve-Taboada, D. Mas, and J. García, "Three-dimensional object recognition by Fourier transform profilometry," *Appl. Opt.* **38**, 4760–4765 (1999).
8. J. Guerrero-Bermúdez, J. Meneses, and O. Gualdrón, "Object recognition using three-dimensional correlation of range images," *Opt. Eng.* **39**, 2828–2831 (2000).
9. B. Javidi and E. Tajahuerce, "Three-dimensional object recognition by use of digital holography," *Opt. Lett.* **25**, 610–612 (2000).
10. Y. Frauel, E. Tajahuerce, M.-A. Castro, and B. Javidi, "Distortion-tolerant three-dimensional object recognition with digital holography," *Appl. Opt.* **40**, 3887–3893 (2001).
11. Ph. Réfrégier and F. Goudail, "Statistical processing of polarization diversity images," in *Optoelectronic Information Processing: Optics for Information Systems*, Ph. Réfrégier, B. Javidi, C. Ferreira, S. Vallmitjana, eds., *Proc. SPIE CR81-13*, 262–288 (2001).
12. J. H. Bruning, D. R. Herriott, J. E. Gallagher, D. P. Rosenfeld, A. D. White, and D. J. Brangaccio, "Digital wavefront measuring interferometer for testing optical surfaces and lenses," *Appl. Opt.* **13**, 2693–2703 (1974).
13. I. Yamaguchi and T. Zhang, "Phase-shifting digital holography," *Opt. Lett.* **22**, 1268–1270 (1997).
14. M. Rabbani, *Selected Papers on Image Coding and Compression*, SPIE Milestone Series **MS48** (SPIE Press, 1992).
15. L. P. Yaroslavsky and N. S. Merzlyakov, *Methods of Digital Holography* (Consultants Bureau, New York, 1980).
16. O. Bryngdahl and F. Wyrowski, "Digital holography—computer-generated holograms," *Prog. Opt.* **28**, 1–86 (1990).
17. J. W. Goodman and R. W. Lawrence, "Digital image formation

- from electronically detected holograms," *Appl. Phys. Lett.* **11**, 77–79 (1967).
18. J. W. Goodman and A. Silvestri, "Digital reconstruction of holographic images," *Nerem Record* **10**, 118 (1968).
 19. A. Bilgin, G. Zweig, and M. W. Marcellin, "Three-dimensional image compression with integer wavelet transforms," *Appl. Opt.* **39**, 1799–1814 (2000).
 20. S. Qian, A. B. Hollinger, D. J. Williams, and D. Manak, "Fast three-dimensional data compression of hyperspectral imagery using vector quantization with spectral-feature-based binary coding," *Opt. Eng.* **35**, 3242–3249 (1996).
 21. G. J. Ewing and C. J. Woodruff, "Comparison of JPEG and fractal-based image compression on target acquisition by human observers," *Opt. Eng.* **35**, 284–288 (1996).
 22. C. A. Morioka, M. P. Eckstein, J. L. Bartroff, J. Hausleiter, G. Aharanov, and J. S. Whiting, "Observer performance for JPEG vs. wavelet image compression of x-ray coronary angiograms," *Opt. Express* **5**, 8–19 (1999).
 23. M. W. Farn and J. W. Goodman, "Bounds on the performance of continuous and quantized phase-only matched filters," *J. Opt. Soc. Am. A* **7**, 66–72 (1990).
 24. A. Mahalanobis and C. Daniell, "Data compression and correlation filtering," in *Smart Imaging Systems*, B. Javidi, ed. SPIE **PM91** (SPIE Press, 2001), pp. 111–132.
 25. J. Vago, H. Vermeulen, and A. Verga, "Fast Fourier transform based image compression algorithm optimized for speckle interferometer measurements," *Opt. Eng.* **36**, 3052–3063 (1997).
 26. R. Shahnaz, J. F. Walkup, and T. F. Krile, "Image compression in signal-dependent noise," *Appl. Opt.* **38**, 5560–5567 (1999).
 27. F. Murtagh, J.-L. Starck, and M. Louys, "Very-high-quality image compression based on noise modeling," *Int. J. Imaging Syst. Technol.* **9**, 38–45 (1998).
 28. F. Wyrowski and O. Bryngdahl, "Speckle-free reconstruction in digital holography," *J. Opt. Soc. Am. A* **6**, 1171–1174 (1989).
 29. T. Nomura, A. Okazaki, M. Kameda, Y. Morimoto, and B. Javidi, "Digital holographic data reconstruction with data compression," *Algorithms and Systems for Optical Information Processing V*, B. Javidi and D. Psaltis, ed., Proc. SPIE **4471**, (2001).
 30. E. Y. Lam and J. W. Goodman, "Discrete cosine transform domain restoration of defocused images," *Appl. Opt.* **37**, 6213–6218 (1998).
 31. J. W. Goodman and A. M. Silvestri, "Some effects of Fourier domain phase quantization," *IBM J. Research and Dev.* **14**, 478–484 (1970).
 32. W. J. Dallas and A. W. Lohmann, "Phase quantization in holograms," *Appl. Opt.* **11**, 192–194 (1972).
 33. H. J. Caulfield, *Handbook of Optical Holography* (Academic Press, San Diego, Calif., 1979).
 34. D. A. Huffman, "A method for the construction of minimum redundancy codes," *Proc. IRE* **40**, 1098–1101 (1952).
 35. J. Ziv and A. Lempel, "A universal algorithm for sequential data compression," *IEEE Trans.* **IT-23**, 337–343 (1977).
 36. T. A. Welch, "A technique for high performance data compression," *IEEE Computer* **17**, 8–19 (1984).
 37. M. Burrows and D. J. Wheeler, "A block-sorting lossless data compression algorithm," Digital SRC Report 124 (1994).
 38. B. Javidi, "Nonlinear joint power spectrum based optical correlation," *Appl. Opt.* **28**, 2358–2367 (1989).