Optical Engineering

OpticalEngineering.SPIEDigitalLibrary.org

JPEG 2000-based compression of fringe patterns for digital holographic microscopy

David Blinder
Tim Bruylants
Heidi Ottevaere
Adrian Munteanu
Peter Schelkens

SPIE.

JPEG 2000-based compression of fringe patterns for digital holographic microscopy

David Blinder,^{a,b,*} Tim Bruylants,^{a,b} Heidi Ottevaere,^c Adrian Munteanu,^{a,b} and Peter Schelkens^{a,b}

^aVrije Universiteit Brussel, Department of Electronics and Informatics (ETRO), Pleinlaan 2, 1050 Brussels, Belgium

^bMinds, Department of Multimedia Technologies (MMT), Gaston Crommenlaan 8 (box 102), 9050 Ghent, Belgium

^cVrije Universiteit Brussel, Brussels Photonics Team (B-PHOT), Pleinlaan 2, 1050 Brussels, Belgium

Abstract. With the advent of modern computing and imaging technologies, digital holography is becoming widespread in various scientific disciplines such as microscopy, interferometry, surface shape measurements, vibration analysis, data encoding, and certification. Therefore, designing an efficient data representation technology is of particular importance. Off-axis holograms have very different signal properties with respect to regular imagery, because they represent a recorded interference pattern with its energy biased toward the high-frequency bands. This causes traditional images' coders, which assume an underlying $1/f^2$ power spectral density distribution, to perform suboptimally for this type of imagery. We propose a JPEG 2000-based codec framework that provides a generic architecture suitable for the compression of many types of off-axis holograms. This framework has a JPEG 2000 codec at its core, extended with (1) fully arbitrary wavelet decomposition styles and (2) directional wavelet transforms. Using this codec, we report significant improvements in coding performance for off-axis holography relative to the conventional JPEG 2000 standard, with Bjøntegaard delta-peak signal-to-noise ratio improvements ranging from 1.3 to 11.6 dB for lossy compression in the 0.125 to 2.00 bpp range and bit-rate reductions of up to 1.6 bpp for lossless compression. © The Authors. Published by SPIE under a Creative Commons Attribution 3.0 Unported License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI. [DOI: 10.1117/1.OE.53.12.123102]

Keywords: wavelet; JPEG 2000; holography; packet decomposition; directional transform; image compression.

Paper 141105P received Jul. 11, 2014; accepted for publication Nov. 11, 2014; published online Dec. 5, 2014.

1 Introduction

Fringe patterns occur in many applications across different scientific disciplines and are mainly used for the characterization of object shapes and dimensions. Applications include interferometric synthetic aperture radar, holographic interference microscopy, photoelasticity measurements, digital holography, and even noninterferometric patterns such as structured imaging applications. However, such fringe patterns do not possess the typical frequency distributions found in natural photographic imagery, which results in suboptimal compression performance when using common image compression coder-decoder architectures (codecs). This paper focuses on one particular application, namely digital holographic microscopy (DHM), which is essentially digital holography applied to microscopy.

Holography, first discovered in the late 1940s by Dennis Gabor, measures the full wavefront of a scene by capturing both the amplitude and the phase information. However, practical applications for optical holography only started to appear in the 1960s, mainly due to the development of the laser. For many years, the recording of holograms was only possible by means of analog high-resolution light-sensitive film (like photographic film). Digital representation and reconstruction of holograms were already proposed in the 1970s,¹ but the lack of adequate computing power and digital imaging devices made it impractical. Over time, however, with the advent of high-resolution digital image sensors

(like CCDs) and increasingly faster computers, practical implementations and applications of digital holography started to appear in the 1990s. As a consequence, DHM has been successfully utilized for many different purposes such as the analysis of biological samples² and materials, characterization of lenses,³ and tomography.

Because the digital holography is getting more widespread and applied in an increasing number of scientific fields, the need for an efficient digital representation technology is growing. Given the multidisciplinary nature of holography, various techniques have been experimented with. Earlier experiments involved the use of histogram-based approaches⁴ or relied solely on quantization principles.⁵ Over time, more advanced codecs were evaluated such as directly coding phase-shifted holograms with the AVC and HEVC codecs⁶ or the JPEG and JPEG 2000 codecs.⁷ More recently, compression frameworks have been proposed which were tailored for digital holography such as computer-generated holograms from encoded multiview video streams⁸ and a vector-lifting scheme for phase-shifted holograms.⁹ Alternatively, some proposals even designed specialized transforms for the efficient compression of holograms. One such example is Fresnelets,¹⁰ which can be interpreted as Fresnel-transformed B-spline wavelets that asymptotically converge to Gabor functions.¹¹ Although they have several interesting properties, Fresnelets are not suitable for our present coding requirements. The reasons are (1) the lack of support for integer taps that enable lossless compression, (2) the inherent presence of complex-valued coefficients, even for real-valued off-axis holograms, (3) the occasional requirement of zero-padding to correctly simulate Fresnel

*Address all correspondence to: David Blinder, E-mail: dblinder@etro.vub.ac.be

propagation in order to avoid aliasing, and (4) the need to predetermine a focus depth parameter prior to compression of the data. Several other analysis functions have been investigated as well; one notable example is the use of Gabor wavelets¹² for hologram coding: their excellent time-frequency uncertainty bounds and orientability allow for effective descriptions of localized frequency content. Unfortunately, Gabor wavelets are not reversible and form an overcomplete representation, requiring some coefficient selection method (such as using the ridge of the Gabor wavelet transform¹²).

Moreover, the wide range of applications and recording setup parameters of DHM typically make the compression requirements dependent on the use case. In contrast, we aim to propose a generic, modality-independent coding architecture, targeted primarily at lossless and near-lossless compressions. Doing so enables us to provide a framework suited for archiving the raw hologram data, while still allowing postprocessing algorithms (such as speckle filtering, unwanted order removal, extended focusing, and so on) after compression. As such, we propose a framework based on JPEG 2000, with specific additional extensions that significantly improve the off-axis DHM data representation.

This paper is an extension of our work published in two different conference papers.^{13,14} We extend our previous work by (1) thoroughly presenting and comparing our coding architecture against the state-of-the-art, (2) providing the necessary technical details in order to practically realize our coding system and ensuring the reproducibility of the results, (3) providing experimental results on a larger and more varied collection of holograms, and (4) evaluating more coding technologies (such as JPEG and JPEG-LS) and more wavelet decompositions (5-level wavelet decompositions, with and without directional wavelets and/or packet decompositions).

This paper is structured as follows. In Sec. 2, we briefly explain the principle of off-axis holography and discuss the data characteristics of such digital holograms. Subsequently, we introduce JPEG 2000 in Sec. 3 and explain how it can be efficiently configured to improve the compression performance for digital holographic imagery. Section 4 then discusses our proposed extensions to the JPEG 2000 standard to further improve the compression efficiency by optimally

exploiting the specific data characteristics of hologram recordings. We then report on the experiments in Sec. 5. Finally, we present the conclusions in Sec. 6.

2 Coding of Holograms

Digital holography is a measurement technique based on the interference of electromagnetic waves. This technology allows for the recovery of both the amplitude and the phase of the wavefront and enables the full description and visualization of three-dimensional objects. Besides the attractiveness for entertainment purposes, this is an extremely useful property for many measurement and visualization applications. In particular, the use of digital holography offers many advantages for microscopic applications. Regular microscopes only provide a two-dimensional (2-D) snapshot of the intensity with a single-focal plane, while holographic microscopes, on the other hand, capture the full wavefront emanating from the sample. This offers several benefits and substantially expands the available tools for data analysis. For example, the phase data give quantifiable information about optical distance and topographical information, enabling postcapture digital refocusing capabilities. Moreover, holographic microscopy has no image forming lens and will not suffer from typical optical aberrations caused by intrinsic lens imperfections of regular microscopes.

Many methods with varying properties, degrees of quality, and feasibility exist for recording holograms¹⁵ such as Fourier holography, Gabor holography, phase-shifting digital holography, and off-axis (Fresnel) holography. The holograms used in this paper were recorded using the off-axis configuration, as shown in Fig. 1, also known as Leith–Upatnieks holography.¹⁶ Such a configuration allows one to capture a single real-valued recording from which the sought wavefront can be subsequently extracted. Basically, the CCD sensor captures the amplitude of the interference pattern that results from the superposition of a reference beam and an object beam. The expression for the detected irradiance I , with R and O representing the amplitudes of the reference and object beams, respectively, and where $*$ is the complex conjugate operator, is then given by

$$I = (R + O)^2 = |R|^2 + |O|^2 + R \cdot O^* + R^* \cdot O. \quad (1)$$

Fig. 1 Simplified representation of the setup for off-axis hologram acquisition (transmission imaging).

Reproduction of the hologram using the same reference beam R to illuminate the recorded hologram effectively requires modulating the irradiance² and can be described by

$$\Psi = R \cdot I = R \cdot (|R|^2 + |O|^2) + R^2 \cdot O^* + |R|^2 \cdot O. \quad (2)$$

The term $|R|^2 \cdot O$ in Eq. (2), or the real image, is directly proportional to the sought object beam. However, the same equation shows that the reproduced hologram also contains a number of additional undesired terms. The first term represents the zero-order diffraction and the second term is a so-called twin image. With off-axis holography, the reference beam will reach the CCD at an offset angle $\theta \neq 0$ deg instead of being collinear with the object-beam axis. The detected reference wave will, therefore, approximately be a tilted plane wave, denoted by $Re^{i\omega x}$. The spatial frequency ω depends on the incidence angle θ $\omega = (2\pi/\lambda) \sin(\theta/2)$.¹⁷ The resulting irradiance will now be the following:

$$I = (Re^{i\omega x} + O)^2 = |R|^2 + |O|^2 + Re^{i\omega x} \cdot O^* + R^* e^{-i\omega x} \cdot O. \quad (3)$$

Equation (3) shows that the terms can be effectively separated in the frequency domain. In principle, a sufficiently high carrier frequency ω allows the object beam to be recovered unambiguously. However, in practice, large overlaps are often still present in the frequency domain, especially with the zero-order term: the total term separation constraints are often too stringent, leaving limited spectral support for the real image. Moreover, directly extracting the object wave field is not straightforward as it requires additional nontrivial processing steps. In fact, various techniques have been proposed in the literature for extracting the real image,¹⁸ all with their respective advantages and disadvantages. Examples of such methods are (1) the basic spatial filtering techniques^{19,20} that have the drawback of altering the reconstructed signal as a consequence of their global filtering effect,²¹ or the more advanced (2) wavelet-domain coefficient selection,²² (3) (linear and nonlinear) Fresnel filtering,¹⁸ or (4) nonlinear cepstrum filtering.^{23,24}

Moreover, the quality of the resulting real image depends on additional factors for which no general solution exists: (1) Scenes with objects at multiple depths have to be modeled using postprocessing such as extended focusing,²⁵ (2) the applied quality metric is extremely use case-specific, as it will, e.g., have to determine the relative importance between the amplitude and the phase information, and (3) distortions caused by the setup's nonidealities, such as aberrations in the microscope objective and in the wave planarity, have to be taken into account and compensated for.²⁶

The selection of specific methods for preprocessing the recorded image before compression will thus inevitably limit the scope of an applied compression algorithm with respect to the range of supported modalities. Thus, our proposal instead uses a modality-independent compression architecture, intended to allow for compressing the entire interferogram I in a progressive lossy-to-lossless manner; that is, the proposed coding architecture controls the losses incurred by coding, eventually offering lossless decoding of the input data when this is needed.

Due to the nature of off-axis holography, the omnipresent high-frequency components manifest themselves as salient fringes in the hologram and cause the power spectrum distribution to significantly deviate from the typical $1/f^2$

distribution found in regular imagery (Fig. 2). The important basis functions, therefore, have to consist of well-oriented high-frequency components, as they will largely represent the real image to be viewed. We confirmed this¹³ by using independent component analysis to formally characterize the nature of the information content in multivariate data. The twin image is also contained in these high-frequency components. However, this is not an issue in the case of off-axis holography compression, because the twin image is the complex conjugate of the real image in the frequency domain. Thus, no extra information will be coded in this respect. Some recent publications also have confirmed the importance of orientation and high frequencies in holography by using Gabor wavelets, evaluated at multiple orientations,¹² or by using the wavelet-bandelets transform.²⁷ However, these publications mainly used coefficient thresholding and did not provide complete coding frameworks. Finally, the good space localization properties of wavelets are of larger importance in DHM data, because the shallow focus distances result in spatially localized structures in the hologram (see Fig. 7).

In the following section, we will first concisely introduce the JPEG 2000 coding architecture on which we will build our system.

3 JPEG 2000

3.1 Introduction

JPEG 2000 is a scalable wavelet-based still image coding system. The JPEG 2000 standard represents a family of standards, where Part 1 describes the core coding technology. The other parts define the extensions by amending Part 1

Fig. 2 Spatial versus frequency domain representations of a natural image and a hologram. (a) Lena, a natural image; (b) Lena, in Fourier frequency domain; (c) Coin, a hologram; (d) Coin, in Fourier frequency domain.

with new features or capabilities, thus making JPEG 2000 modular by design. Its core technology, as defined in Part 1, offers a rich set of features such as native tiling support, resolution and quality scalability, progressive decoding, lossy and lossless compressions, region-of-interest coding, error resilience, true random access in the code-stream, and so on. Moreover, JPEG 2000 natively supports various color formats, such as RGB, $Y'C_bC_r$, and arbitrary n -channel, at bit depths ranging from 1 to 38 bits per channel. Especially important is the rate-distortion (RD) optimization capability that is inherent to the design of JPEG 2000 and allows for optimal control of the bit rate of the produced code-stream while minimizing the overall distortion in a lossy compression scenario.

Both its modular and extendable designs and its excellent RD characteristics make JPEG 2000 a well-suited codec for the compression of many types of holographic images.

3.2 Architecture

This section gives a brief description of the core coding technology of JPEG 2000. As shown in Fig. 3, the core architecture of JPEG 2000 can be roughly divided into two main parts: (1) the discrete wavelet transform (DWT), and (2) the two-tiered embedded block coding by optimized truncation (EBCOT).^{28,29}

The first step in encoding an image with JPEG 2000 consists of a multilevel 2-D DWT using the Mallat dyadic decomposition structure,³⁰ where only the low-pass sub-bands are further decomposed in the subsequent resolution levels. JPEG 2000 employs two wavelet kernels, both synthesized using the lifting scheme: (1) the integer 5×3 kernel for lossless coding and (2) the floating-point 9×7 kernel for

lossy to near-lossless codings. Both kernel implementations are strictly defined by the JPEG 2000 specification. Because the lifting coefficients of the 5×3 kernel are rational numbers with denominators of powers of 2, and in order to be able to guarantee lossless reconstruction at the decoder side, the specification restricts its implementation to integer calculus only using well-defined rounding rules. The 9×7 kernel, on the other hand, offers much better energy compaction over the 5×3 kernel. However, due to its coefficients being real numbers, it cannot be easily fit into an integer-based calculus system without severely sacrificing the energy compaction performance. For this reason, JPEG 2000 specifies this kernel using floating-point calculus, making it an irreversible transform. Extensive results on natural data³¹ show that the lossy 9×7 wavelet kernel performs better in the RD sense than the 5×3 wavelet kernel; however, due to its pure-integer implementation, the 5×3 kernel is better suited for lossless compression.

After the wavelet decomposition step, the resulting sub-bands are further entropy encoded with the two-tiered EBCOT. For each of the sub-bands, EBCOT Tier-1 starts out by grouping the wavelet coefficients into equally sized rectangular areas, so-called code-blocks. Then, it performs entropy coding on each of these code-blocks by employing context-based binary arithmetic coding. The wavelet coefficients are scanned per bit-plane, starting with the most significant bit-plane, using three types of coding passes in alternating order, namely the significance coding pass, the magnitude refinement coding pass, and the cleanup pass. As such, every coefficient bit becomes a member of exactly one of these three coding passes and gets encoded into the respective code-block bit-stream. The end of every code-pass and inherently also the end of every bit-plane scan, marks

Fig. 3 JPEG 2000 encoder schematic.

Fig. 4 Diagrams of the tested 4-level decomposition structures. (a) 4-level Mallat decomposition style; (b) 4-level partial packet decomposition style; and (c) 4-level full packet decomposition style.

a potential truncation point in the resulting bit-stream. Along with each of these truncation points, Tier-1 also estimates the associated mean square error (MSE) distortion reduction values that will drive the Tier-2 RD optimization process. Thus, after Tier-1 is done, every code-block is represented by an independently compressed bit-stream and an associated table of truncation points with distortion reduction estimates per truncation point.

EBCOT Tier-2 represents the actual RD optimization and packetization process, responsible for generating the final JPEG 2000 code-stream. Given the rate and/or quality constraints, pieces of the individual code-block bit-streams from Tier-1 are selected and recombined into larger packets to form the final JPEG 2000 code-stream. The selection of the bit-stream pieces is performed in an RD optimal manner by prioritizing bit-stream chunks based on their respective RD costs over less important chunks, while still maintaining causality—i.e., by maintaining the information dependency between chunks to allow for correct decoding. The RD optimization stops when the rate and/or quality constraints are met, or when all bit-stream chunks are included in the final code-stream. Finally, the necessary JPEG 2000 headers and markers are appended in order to signal the required decoding options.

3.3 Full Packet Decomposition with JPEG 2000

As stated before, due to the nature of off-axis holography, a significant part of the important information in these recordings is contained in the high-frequency bands. This contrasts with natural images where most of the visually meaningful information resides in the lower-frequency bands. Thus, replacing the Mallat dyadic wavelet transform with a full packet wavelet transform allows for further decomposition of the high-pass sub-bands to improve the compression efficiency.

By default, JPEG 2000 Part 1 features only the Mallat dyadic decomposition. However, JPEG 2000's Part 2 Arbitrary Decomposition (AD) extension enables the use of alternative decomposition structures, signaled within two additional marker segments in the code-stream. As such, using this extension enables the configuration of various packet decomposition structures. The AD extension specifies a decomposition structure in two parts: (1) an underlying decomposition to generate the resolution levels and (2) per resolution level, the extra sublevel decomposition of the respective high-pass sub-bands.

The resolution levels are defined similarly to Part 1 with the Mallat dyadic decomposition, with the difference that the splitting of the low-pass sub-band at each level can be either in both horizontal and vertical directions or only in one of the two directions. This sequence of resolution reduction splits is signaled in the down-sampling factor styles (DFS) marker in the code-stream, represented as an array, *Ddfs*, containing two-bit symbols ("1" = both rows and columns, "2" = rows only, and "3" = columns only). In the absence of a DFS marker, the decoder assumes both-ways splitting as the default, which is the compliant case with a Part 1 code-stream.

Subsequently, one or more AD Style (ADS) markers can be used to signal the sublevel decomposition of high-pass sub-bands. Unfortunately, according to the standard, the ADS syntax only allows for two additional decompositions of the high-pass sub-bands, inherently limiting the possible wavelet packet transforms. The ADS marker contains two arrays: (1) *DOads* specifies the maximum number of split levels per resolution using entries of two-bit values, and (2) *DSads* specifies the type of extra split using two-bit symbols (0 = no extra split, 1 = both rows and columns, 2 = rows only, and 3 = columns only). *DOads* entries with a value of 1 indicate that no extra high-pass decompositions are required, while values 2 and 3 indicate one and two extra decompositions, respectively. The *DSads* array, on the other hand, describes the depth-first traversal of the decomposition tree, with the sub-bands ordered from the highest resolution to the lowest and within each level as *HH-LH-HL-LL*, *HX-LX*, or *XH-XL*, depending on the applied split type.

Hence, the AD extension of JPEG 2000 limits full packet decompositions to $N_L = 3$ levels. As such, the application of four or more levels in such a full packet decomposition structure, as illustrated in Fig. 4(c), is not possible without modification of the standard. Figure 4(b) visualizes the closest matching decomposition style to a full packet with four levels that can be described by the AD extension (designated as the "partial packet decomposition"). Finally, to illustrate how the AD extension works, we show in Table 1 some of the more commonly known decomposition structures and how to signal them. The last column lists the code-stream signaling cost in bits.

4 Proposed Extensions for JPEG 2000

This section discusses the extensions to JPEG 2000 that can significantly enhance the compression efficiency of holographic microscopy images. Figure 5 shows that our

Table 1 Various well-known decomposition structures, using the arbitrary decomposition (AD) extension syntax (signaling cost reflects additional required bits, excluding the cost for N_L).

Decomposition style	N_L	Ddfs	DOads	DSads	Signaling cost
3-level full packet	3	111	321	18 1's	80 bits
4-level partial packet	4	1111	3321	33 1's	112 bits
4-level full packet	NA	NA	NA	NA	NA
5-level full packet	NA	NA	NA	NA	NA
5-level federal bureau of investigation (FBI)	5	11111	2321	1110111111111111	88 bits

Fig. 5 Our extended encoder schematic, based on JPEG 2000.

two proposed extensions are part of the transform phase of the codec and replace the default DWT block in the JPEG 2000 encoder scheme of Fig. 3.

4.1 Truly Arbitrary Packet Decompositions

As explained in Sec. 3.3, JPEG 2000's AD extension can only handle two additional decompositions within high-pass sub-bands. To overcome the limitation of the AD extension, we designed our codec to employ an alternative code-stream syntax that is able to truly describe arbitrary wavelet decompositions¹⁴ including full packet decompositions containing more than three levels.

The proposed syntax describes a decomposition structure as an ordered array of split-operations that work on a stack of available sub-bands. The split-operations are each represented as a tuple (s, m, r) in the array:

1. Symbol s represents the split type (XY = both rows and columns, $X-$ = rows only, $-Y$ = columns only, and $--$ = termination). Furthermore, the selected split type s also determines the number of bits required to signal m and their relation with the generated sub-bands, as given in Table 2.
2. A binary pattern mask m indicates which of the resulting sub-bands will be added to the stack after splitting for further processing: a bit-value of $m_n = 1$ marks the respective sub-band for further processing, whereas $m_n = 0$ signals termination. When $s = --$, m contains no bits, meaning that in this case m is simply not signaled.
3. A positive integer value $r \in \mathbb{N}$ is reserved for which of its functional definitions depends on the actual applied split type in s and the associated value of mask m :

Table 2 The relation between the applied split type, signaled by s , and the definition of mask bits m_n , indicating the termination of the decomposition ($m_n = 1$ marks the respective sub-band for further processing, whereas $m_n = 0$ signals termination). It also specifies the bit-stream encoding for r , with L_{sbs} equals to the sub-band stack size just before processing the respective tuple.

Split type s	Code-words for s	Generated sub-bands (in order)	Associated mask bits	Bit coding for r	No. bits for r
XY	11_b	HH, LH, HL, LL	m_3, m_2, m_1, m_0	Golomb	$r + 1$
X-	10_b	HX, LX	m_1, m_0	Golomb	$r + 1$
-Y	01_b	XH, XL	m_1, m_0	Golomb	$r + 1$
--	00_b	None	None	Raw	$(\log_2 L_{\text{sbs}})$

- If $s \neq --$ and $m \neq 0$, then r specifies the number of times to recursively repeat the split operation onto all of the respectively generated sub-bands.
- If $s \neq --$ and $m = 0$, then the value of r is not signaled, because none of the created sub-bands will be further processed anyway.
- On the other hand, if $s = --$, then r specifies the number of times to repeat this termination operation to the stack of sub-bands (thus, one termination is applied with $r = 0$). Applying the termination operation on the sub-band stack is identical to removing the top element.

The decomposition process starts with LL_0 (i.e., the original image) as the only available sub-band on the stack. Subsequently, the process iterates over the list of (s, m, r) -tuples and with each tuple the according split-operation is executed. Generated sub-bands that result from a split operation are immediately pushed onto the sub-band stack, following $HH-LH-HL-LL$, $HX-LX$, or $XH-XL$ ordering. The process ends when all tuples on the list are processed (left-over sub-bands on the stack are not further processed).

Table 3 shows how to specify some commonly used decomposition styles. Please note that although possible, in practice the extended syntax will not be used to specify a Mallat dyadic decomposition style, as this is the JPEG 2000 default anyway.

Signaling of the decomposition style in the final code-stream happens via a newly proposed XAD marker that

encapsulates the binary representation of the array of split-operation tuples, padded to the byte boundary with 0 bits. The marker length field L_{XAD} determines the number of elements in the array. It is valid in the main and tile-component headers to allow defining different decomposition structures per tile-component.

4.2 Directional Adaptive Discrete Wavelet Transform

A salient feature of off-axis digital holographic images is the strongly oriented interference fringes. This hints that the use of directional wavelet transforms can improve the compression performance, as they are able to align with the directional features of the data.¹³ For that reason, we show how the JPEG 2000 architecture can easily be extended to include the block-based directional adaptive DWT (DA-DWT).^{32–34}

We employ a separable lifting scheme, similar to that of JPEG 2000 DWT, but with modified prediction and update functions that are no longer confined only to the horizontal direction (1,0) for row-based splits and the vertical direction (0,1) for column-based splits. Doing so enables the directional DWT to adapt to local geometric features by adjusting its operational direction. However, all the applied direction vectors are also required at the decoder side in order to perform the inverse DA-DWT operation. From a compression performance point of view, it is evident that the unavoidable increment in rate for signaling these directions to the decoder should not jeopardize the rate reduction brought by the improved energy compaction of the transform. Thus, in practice, the adaptability of the directional DWT is restricted by (1) allowing the selection of only one vector per block of samples for the row and column splits, 2) limiting the directions to a discrete set of vectors, and (3) only performing the DA-DWT on low-pass sub-bands (i.e., LL , XL , or LX bands).

With the first restriction, DA-blocks are defined in an identical way as JPEG 2000's code-blocks and precincts. Per decomposition level, they represent a grid of equally sized rectangles that anchor at (0,0) with their dimensions restricted to powers of 2. The width and height parameters are signaled in a dedicated marker segment, which we label XDA, for the DA-DWT. The smallest possible DA-block is 4×4 .

Second, the proposed extension uses the following set of (x, y) vectors for row-based splits $\{(1,0), (3,1), (3,2), (1,1), (1,2), (1,3), (-1,3), (-1,2), (-1,1), (-3,2), (-3,1)\}$ and orthogonally $\{(0,1), (1,3), (2,3), (1,1), (2,1), (3,1), (3,-1), (2,-1), (1,-1), (2,-3), (1,-3)\}$ for column-based

Table 3 Various well-known decomposition structures using our proposed syntax.

Decomposition style	Operation tuples	Signaling cost (bits)
3-level full packet	$(XY, 1111_b, 2)$	9
4-level partial packet	$(XY, 1111_b, 2),$ $(XY, 0000_b, 0)$	15
4-level full packet	$(XY, 1111_b, 3)$	10
5-level full packet	$(XY, 1111_b, 4)$	11
5-level FBI	$(XY, 1111_b, 1), (-, 13),$ $(XY, 1111_b, 1), (-, 16),$ $(XY, 1111_b, 1), (-, 16),$ $(XY, 1111_b, 1), (-, 15),$ $(XY, 0000_b, 0)$	65

splits (for more information on direction vectors and the associated lifting schemes, we refer to Chang and Girod³²). Note that the inclusion of the vectors (1,0) and (0,1) allows the DA-DWT to fall back to the classic DWT in the case where no dominant direction is present. Thus, each vector can be represented as an index in the set of available vectors. Moreover, it is also possible to use different direction vectors sets, depending on the use case and specific image characteristics. Per DA-DWT level, we use the JPEG 2000 tag-tree system to encode the two grids of direction indexes (one for the row-based split and one for the column-based split). The actual tag-tree values (two for every DA-block) are coded in synchronization with the first instance of any of the possibly associated code-blocks—i.e., depending on the chosen dimensions of code-blocks and DA-blocks, each DA-block can relate to one or more code-blocks. This also means that, at very low bit-rate constraints, it can happen that no code-block contributions exist whatsoever for a specific DA-block. In such a case, the associated direction indices are simply skipped and not encoded in the tag-trees. The resulting encoded bit-stream is signaled in an XDA marker segment.

Third, the restriction to allow only the DA-DWT on the low-pass sub-bands does not negatively impact the compression performance capabilities of the coding framework. As demonstrated in Fig. 6, an intrinsic property of the directional DWT causes resulting high-pass coefficients to be already horizontally or vertically aligned after the directional wavelet prediction. As such, a single parameter (i.e., one byte) in the XDA marker segment signals the number of decomposition levels that use the DA-DWT, starting at LL_0 .

Technically, an encoder implementation is free to use any type of direction vector selection mechanism to drive the forward DA-DWT. To avoid being trapped in local minima, our framework takes a full-search approach by trying all directions and selecting the one that minimizes the L_1 -norm of the high-pass coefficients.

5 Experiments

5.1 Test Data

The experiments in this paper make use of 12 off-axis holographic test images, courteously provided by the Lyncée Tec

Fig. 6 Realignment of the directional features with the image axes, after using the noninterpolated 1-D-discrete wavelet transform (DWT_D). (a) Original image with distinct directional features; and (b) Resulting LL , HL , LH , and HH sub-bands after a 1-level directional adaptive DWT (DA-DWT), using direction (1,1).

SA, the Microgravity Research Center (ULB), and Nicolas Pavillon. For the acquisition of microscopic off-axis holographic images, two typical setups exist. One setup uses transmission imaging, which is well suited for transparent specimens such as biological cells or lenses. Another setup uses reflection imaging, which is mainly useful for capturing opaque objects such as surface measurements. Figure 7 shows the thumbnail versions of the holographic images with their specifications given in Table 4. All images contain 8 bpp samples.

5.2 Objective Quality Metrics

The experiments in this paper report on both lossless and lossy compression performances. In the case of lossless compression, where the input signal and the reconstructed signal are identical, the compression performance is quantified in terms of average bit rate in bits per pixel. In order to facilitate easier comparisons between different compression strategies, we present relative bit rates with respect to a common reference, which is JPEG 2000 with a 4-level Mallat wavelet decomposition structure.

In the case of lossy compression, we calculate for a given set of bit rates the respective reconstructed quality (or distortion) as the peak signal-to-noise ratio (PSNR). The PSNR is basically a logarithmic representation of the MSE between the original signal I and the reconstructed signal R and is defined as

$$\text{PSNR} = 10 \cdot \log \left(\frac{I_{\max}^2}{\text{MSE}} \right) \quad \text{with} \quad \text{MSE} = \frac{1}{n} \sum_{i=1}^n (I_i - R_i)^2, \quad (4)$$

where $I_{\max}$ represents the maximum signal value (255 for 8-bit data), and n is the total pixel count in I and R .

In lossy compression, the paper reports summarized RD results using the Bjøntegaard delta PSNR metric (BD-PSNR),³⁵ which is a commonly accepted objective metric for image compression performance evaluations. The BD-PSNR methodology calculates the difference between two such RD curves as the surface area size between the curves within the operating bit-range divided by the integration interval (see Fig. 8). The bit-rates where PSNR differences are measured for the BD-PSNR metric³⁵ in this paper are taken between 0.125 and 2.00 bpp.

5.3 Decomposition Structures and JPEG 2000 Settings

The purpose of these experiments is to assess whether JPEG 2000 or the proposed, extended JPEG 2000 compatible coding architecture can be used to efficiently compress off-axis holographic image data. As such, it is evident to include configurations that are fully compliant with the current JPEG 2000 standard, thus including the configurations that rely on the AD extension of JPEG 2000 Part 2. More specifically, we test the “Mallat dyadic,” the “3-level full packet” and the “4-level partial packet” decomposition structures as listed in Table 1. On the other hand, by using our proposed extended syntax for the decomposition structures, we also test the “4-level full packet” and the “5-level full packet”

Fig. 7 Set of off-axis holographic images. (a) Neuron; (b) Erythrocyte; (c) Microlenses; (d) Ball; (e) Scratch; (f) Seaweed 1; (g) Seaweed 2; (h) Seaweed 3; (i) Coin; (j) Mirror; (k) Sine; and (l) Pollen.

Table 4 Specification of the 8-bit images that were used for the experiments.

Image	Provider	Content description	Dimensions	Imaging type
Neuron	Lyncée Tec	Slice of neuronal tissue	1024 × 1024	Transmissive
Erythrocyte	Lyncée Tec	Erythrocytes from a blood sample	512 × 512	Transmissive
Microlenses	Lyncée Tec	Array of microlenses	1024 × 1024	Transmissive
Ball	Lyncée Tec	Surface of a rough microball	1024 × 1024	Reflective
Scratch	Lyncée Tec	Scratch in brittle material	1024 × 1024	Reflective
Seaweed 1	MRC/ULB	Green seaweed specimen	1280 × 1024	Transmissive
Seaweed 2	MRC/ULB	Green seaweed specimen	2048 × 2048	Transmissive
Seaweed 3	MRC/ULB	Green seaweed specimen	1280 × 1024	Transmissive
Coin	N. Pavillon	Speckle hologram of a coin	512 × 512	Reflective
Mirror	N. Pavillon	Scratch on a mirror	512 × 512	Reflective
Sine	N. Pavillon	Artificial object with sinusoidal amplitudes at multiple frequencies and orientations	512 × 512	Transmissive
Pollen	N. Pavillon	Solution of yew pollens	512 × 512	Transmissive

Fig. 8 Example of the Bjøntegaard delta peak signal-to-noise ratio metric (BD-PSNR).

decompositions, as mentioned in Table 3. Still, given the relative image dimensions, a decomposition tree of typically four levels does suffice to reach optimal energy compaction.

In addition to the wavelet packet transform, our framework also provides support for directional wavelets, for which results are also presented. In combination with the described packet decomposition structures, we include the

results using a DA-DWT for the first two decomposition levels, applied only on the low-pass sub-bands with DA-blocks of 32×32 . Please note that the results include the extra overhead cost for signaling the direction vectors, using the described tag-tree encoding methodology.

For the lossless compression experiments, we make use of the standard integer-based 5×3 wavelet kernel. For the lossy compression results, we rely on the more efficient, but inherently lossy, 9×7 kernel. All the experiments use code-blocks of 32×32 , and precincts and tiling are disabled.

5.4 Results

In order to give an indication of the expected compression performance on holographic images using common JPEG 2000 compression settings and in comparison to regular images, such as Lena, Barbara, and Mandril, we first present results using a conventional 4-level Mallat decomposition. These results, as shown in Table 5, indicate that such a regular Mallat wavelet decomposition performs similarly well for off-axis holographic recordings as for regular images. In fact, all subsequently reported results will be determined relative to these figures.

Table 6 summarizes the obtained lossless compression results, presented as the bit-rate gains relative to the lossless rate when using the 5×3 wavelet kernel in a default 4-level Mallat decomposition mode. These results clearly show that in most cases, the largest compression efficiency gain is obtained by enabling the DA-DWT transform and using a conventional Mallat decomposition structure. A notable exception is the Seaweed recordings that benefit from the

Table 5 Lossless compression rates and peak signal-to-noise ratio (PSNR) results on holographic and natural imageries at 2 bpp down to 0.125 bpp JPEG 2000, when applying a 4-level Mallat wavelet decomposition structure.

Image	Lossless (bpp)	2 bpp (dB)	1 bpp (dB)	0.5 bpp (dB)	0.25 bpp (dB)	0.125 bpp (dB)
Lena	4.35	44.76	40.23	37.07	33.10	29.78
Barbara	4.81	43.05	37.08	32.07	27.71	24.83
Mandrill	6.14	34.61	28.97	25.44	22.73	21.18
Neuron	4.67	43.47	37.62	33.36	28.44	25.80
Erythrocyte	4.91	42.89	37.49	29.93	24.85	22.17
Microlenses	4.12	47.07	41.11	35.06	30.60	27.05
Ball	5.67	36.81	29.50	25.24	22.66	20.88
Scratch	4.44	43.74	38.05	35.76	34.26	32.75
Seaweed 1	3.69	48.69	43.49	37.12	32.04	25.64
Seaweed 2	4.16	45.32	40.90	37.57	31.90	27.33
Seaweed 3	3.97	46.61	42.13	38.95	32.70	29.58
Coin	4.52	43.35	37.88	35.04	33.33	32.07
Mirror	6.14	35.20	26.92	22.14	19.26	18.05
Pollen	5.26	39.74	32.93	29.53	27.86	26.34
Sine	5.44	38.72	31.13	27.88	26.12	24.71

Table 6 Results for lossless compression where the values represent bit-rate reductions (in Δ bpp) in comparison to the standard 4-level Mallat decomposition. The second column shows the bitrates obtained using the default JPEG 2000 configuration with a 4-level Mallat decomposition. The third column is the default results obtained with JPEG-LS, while the other columns report the results obtained with lossless JPEG 2000, using the 5×3 wavelet kernel. Column-notations use abbreviations for Mallat (M), partial Packet (PP), full Packet (FP), and DA-DWT enabled (+DA), preceded by the number of decomposition levels. Only the columns marked with an asterisk are JPEG 2000 Part 1/Part 2 compliant. The last row shows the averages for the holographic images.

Lossless 5×3	Orig. (4M)	JPEG-LS	3FP*	4PP*	4FP	5M*	5FP	3FP + DA	4M + DA	5M + DA	4PP + DA	4FP + DA	5FP + DA
Barbara	4.69	-0.05	-0.18	-0.23	-0.36	0.00	-0.79	-0.11	0.08	0.08	-0.16	-0.28	-0.67
Lena	4.35	0.10	-0.19	-0.24	-0.32	0.00	-0.64	-0.18	0.00	0.00	-0.23	-0.31	-0.62
Mandrill	6.14	0.10	-0.15	-0.17	-0.23	0.00	-0.52	-0.13	0.02	0.02	-0.15	-0.22	-0.49
Neuron	4.67	0.20	0.31	0.34	0.37	0.00	0.32	0.41	0.45	0.45	0.42	0.39	0.32
Erythrocyte	4.92	0.42	0.32	0.38	0.40	0.00	0.23	0.55	0.61	0.58	0.54	0.50	0.25
Microlenses	4.12	-0.05	0.03	0.02	-0.04	0.00	-0.25	0.08	0.17	0.17	0.07	0.00	-0.15
Ball	5.67	-0.11	0.25	0.24	0.21	0.00	0.18	0.50	0.51	0.51	0.49	0.45	0.40
Scratch	4.45	-0.02	0.11	0.12	0.10	0.00	0.06	0.10	0.11	0.12	0.11	0.08	0.04
Seaweed 1	3.69	0.55	0.26	0.28	0.25	0.00	0.19	0.16	0.00	0.01	0.19	0.15	0.10
Seaweed 2	4.16	-2.37	0.21	0.23	0.21	0.00	0.14	0.07	-0.07	-0.07	0.09	0.06	0.00
Seaweed 3	3.97	-2.07	0.16	0.17	0.14	0.00	0.07	0.02	-0.10	-0.09	0.04	0.01	-0.06
Coin	4.52	0.13	-0.09	-0.09	-0.13	0.00	-0.39	0.06	0.15	0.15	0.05	0.02	-0.23
Mirror	6.14	0.00	0.78	0.78	0.77	0.00	0.56	1.55	1.65	1.65	1.54	1.48	1.22
Pollen	5.26	0.61	0.31	0.31	0.28	0.00	0.12	0.98	1.06	1.06	0.97	0.93	0.70
Sine	5.44	0.46	0.34	0.38	0.41	0.00	0.36	0.97	0.93	0.93	0.98	1.01	0.82
Average	4.75	-0.19	0.25	0.26	0.25	0.00	0.13	0.45	0.46	0.45	0.46	0.42	0.29

Note: The bold values indicate which compression parameters give the highest gain for a given JPEG 2000 configuration.

packet decompositions alone. This is caused by the recording setup in which the fringes align with the horizontal and vertical axes. Such axis alignment during the image acquisition is, in fact, sub-optimal as it minimizes the available bandwidth for the spectral separation of the real and conjugate image parts. The results also show that, even without the DA-DWT, most packet decomposition structures already significantly improve the compression efficiency for most holograms.

Table 7 shows the BD-PSNR results relative to the 4-level Mallat configuration using the 9×7 wavelet kernel. These results indicate that the 9×7 kernel with lossy coding provides the largest compression performance gain when applying the 4-level partial packet decomposition, in combination with the DA-DWT transform. Again, similar to lossless coding, the Seaweed images benefit even more from using the packet decompositions alone.

The results from both Tables 6 and 7 show that even in a JPEG 2000 constraint application, the compression efficiency for off-axis holographic images can seriously benefit from the use of a limited packet decomposition, such as the 3-level full packet or 4-level partial packet structures. However, our proposed extensions enhance the compression efficiency even more drastically. The DA-DWT proves to be

a very powerful tool, significantly increasing the compression performance on off-axis microscopic holography data.

It should be noted that the measured distortion introduced by the lossy compression of the recorded hologram is not necessarily directly proportional to the actual perceived distortion of the reconstructed object. This depends entirely on the nature of the introduced distortion. Improvement of the reconstruction quality can be achieved by modifying the used quality metric (which is the conventional MSE employed by the JPEG-2000 standard), so that it better models the relation between objective and subjective distortions of the hologram. However, as noted in Sec. 2 and due to the large number of possible requirements, it is unlikely that creating a universal quality metric directly applicable for every possible type of measurement would be desirable or possible. Still, it is possible to further improve upon the default MSE-based distortion metric (e.g., a weighted MSE metric would allow one to assign lower weights to code-blocks representing frequencies lying far from the carrier frequency, which generally contain less important information). This is similar to the visual frequency weighting used in JPEG 2000 compression for improving the perceived quality of regular imagery.^{36,37} However, this subject is beyond the scope of our paper.

Table 7 Results for lossy compression, with the values representing the BD-PSNR improvements (in dB) w.r.t. to the 4-level Mallat decomposition, in the range of 0.25 to 2.00 bpp. The second column shows the results obtained with JPEG, while the other columns report the results obtained with lossy JPEG 2000, using the 9×7 wavelet kernel. Column-notations use abbreviations for Mallat (M), partial packet (PP), full packet (FP), and DA-DWT-enabled (+DA), all preceded by the number of decomposition levels. Only the columns marked with an asterisk are JPEG 2000 Part 1/Part 2 compliant. The last row shows the averages for the holographic images.

Lossy 9×7	JPEG	3FP*	4PP*	4FP	5M*	5FP	3FP + DA	4M + DA	5M + DA	4PP + DA	4FP + DA	5FP + DA
Barbara	-4.36	0.45	0.19	-0.24	-0.01	-2.42	0.45	0.64	0.64	0.02	-0.32	-2.27
Lena	-2.72	-0.36	-0.70	-0.91	0.00	-2.61	-0.41	0.00	0.00	-0.77	-0.97	-2.67
Mandrill	-5.12	-0.22	-0.35	-0.54	0.00	-1.54	-0.18	0.04	0.05	-0.33	-0.50	-1.48
Neuron	-4.16	4.44	5.15	5.77	0.06	5.86	5.66	5.72	5.69	5.96	5.96	5.91
Erythrocyte	-5.20	5.45	6.35	6.78	0.07	6.34	7.03	7.18	6.69	7.19	7.05	5.97
Microlenses	-3.18	2.77	3.19	3.20	0.04	1.97	3.13	2.86	2.94	3.22	3.11	2.25
Ball	-8.65	4.36	4.50	4.86	0.02	5.03	5.72	5.09	5.09	5.73	5.72	5.63
Scratch	-2.39	2.36	2.53	2.65	0.03	2.59	2.27	1.99	2.02	2.37	2.35	2.26
Seaweed 1	-1.30	5.15	5.99	6.46	0.09	6.68	4.56	0.14	0.22	5.30	5.72	5.86
Seaweed 2	-1.27	3.95	4.41	4.53	0.05	4.58	3.22	-0.32	-0.27	3.71	3.79	3.79
Seaweed 3	-2.78	3.64	4.02	4.07	0.04	4.06	2.89	-0.49	-0.45	3.25	3.31	3.25
Coin	-1.84	0.75	0.79	0.83	0.01	-0.08	1.25	1.00	1.01	1.30	1.29	0.40
Mirror	-2.57	8.70	9.13	9.52	0.03	8.80	11.61	11.38	11.40	11.77	11.66	10.36
Pollen	-8.90	3.49	4.24	5.25	0.01	4.65	7.02	6.99	7.00	7.45	7.42	6.47
Sine	-10.65	4.43	5.40	6.37	0.03	6.66	7.13	5.92	5.94	7.69	8.24	7.97
Average	-4.41	4.12	4.64	5.02	0.04	4.76	5.12	3.96	3.94	5.41	5.47	5.01

Note: The bold values indicate which compression parameters give the highest gain for a given JPEG 2000 configuration.

6 Conclusions

We demonstrate how JPEG 2000 can be efficiently used to compress microscopic off-axis holograms by proposing two extensions to the standard:

1. We replace the existing ADSs feature such that any decomposition structure becomes available. Along with this extension, we provide the means to efficiently signal these AD structures in the code-stream. This means that our proposed code-stream syntax for the XAD marker requires up to 10 times less bits in the header than that of JPEG 2000's AD syntax (ADS and DFS markers) for equal decomposition styles, and it has a lower implementation complexity.
2. We introduce a practical implementation of a block-based DA-DWT for JPEG 2000 (DA-DWT). From the results, it is clear that the compression performance for off-axis microscopic holography data benefits from employing the DA-DWT, even with the overhead of signaling the direction vectors in the code-stream.

In doing so, we realized a framework that is specific enough to compress DHM data with significantly improved compression performance, and yet general enough to leave room for subsequent filtering and postprocessing of the hologram data, depending on the use case. Additionally, we postulate that this framework can be extended to other imaging

technologies based on fringe pattern data, as they largely share the frequency and directionality properties of DHM data. The encoding framework also allows for using other basis functions as well.

Using the proposed techniques, we report significant compression performance gains of 1.3 up to 11.6 dB (BD-PSNR) for lossy compression and bit-rate reductions of over 1.6 bpp for lossless compression of off-axis holographic images.

Acknowledgments

We would like to thank Nicolas Pavillon, the Lyncée Tec SA (Lausanne, Switzerland) and Ahmed El Mallahi (Microgravity Research Center, ULB, Brussels) for providing the digital holographic recordings used in these experiments. The research leading to these results has received funding from the Research Foundation Flanders (FWO) with project no. G014610N and the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC Grant Agreement n. 617779 (INTERFERE).

References

1. M. A. Kronrod, N. Merzlyako, and N. P. Yaroslavski, "Reconstruction of holograms with a computer," *Sov. Phys. Tech. Phys.* **17**(2), 333–334 (1972).
2. E. Cuque, P. Marquet, and C. Depeursinge, "Simultaneous amplitude-contrast and quantitative phase-contrast microscopy by numerical

- reconstruction of Fresnel off-axis holograms," *Appl. Opt.* **38**, 6994–7001 (1999).
3. F. Charrière et al., "Characterization of microlenses by digital holographic microscopy," *Appl. Opt.* **45**, 829–835 (2006).
 4. A. E. Shortt, T. J. Naughton, and B. Javidi, "Histogram approaches for lossy compression of digital holograms of three-dimensional objects," *IEEE Trans. Image Process.* **16**, 1548–1556 (2007).
 5. G. A. Mills and I. Yamaguchi, "Effects of quantization in phase-shifting digital holography," *Appl. Opt.* **44**, 1216–1225 (2005).
 6. Y. Xing, B. Pesquet-Popescu, and F. Dufaux, "Compression of computer generated phase-shifting hologram sequence using AVC and HEVC," *Proc. SPIE* **8856**, 88561M (2013).
 7. E. Darakis and J. J. Soraghan, "Compression of interference patterns with application to phase-shifting digital holography," *Appl. Opt.* **45**, 2437–2443 (2006).
 8. T. Senoh et al., "Multiview image and depth map coding for holographic tv system," *Opt. Eng.* **53**(11), 112302 (2014).
 9. Y. Xing et al., "Vector lifting scheme for phase-shifting holographic data compression," *Opt. Eng.* **53**(11), 112312 (2014).
 10. M. Liebling, T. Blu, and M. Unser, "Fresnelets: new multiresolution wavelet bases for digital holography," *IEEE Trans. Image Process.* **12**, 29–43 (2003).
 11. M. Unser, A. Aldroubi, and M. Eden, "On the asymptotic convergence of b-spline wavelets to gabor functions," *IEEE Trans. Inf. Theory* **38**, 864–872 (1992).
 12. K. Viswanathan, P. Gioia, and L. Morin, "Wavelet compression of digital holograms: towards a view-dependent framework," *Proc. SPIE* **8856**, 88561N (2013).
 13. D. Blinder et al., "Wavelet coding of off-axis holographic images," *Proc. SPIE* **8856**, 88561L (2013).
 14. T. Bruylants et al., "Microscopic off-axis holographic image compression with JPEG 2000," *Proc. SPIE* **9138**, 91380F (2014).
 15. M. K. Kim, "Principles and techniques of digital holographic microscopy," *SPIE Rev.* **1**(1), 018005 (2010).
 16. E. N. Leith and J. Upatnieks, "Reconstructed wavefronts and communication theory," *J. Opt. Soc. Am.* **52**, 1123–1128 (1962).
 17. J. W. Goodman, *Introduction to Fourier Optics*, 3rd ed., Roberts and Company Publishers, Greenwood Village, Colorado (2004).
 18. M. Liebling and M. Unser, "Comparing algorithms for reconstructing digital off-axis fresnel holograms," *Proc. SPIE* **6016**, 60160M (2005).
 19. E. Cuhe, P. Marquet, and C. Depeursinge, "Spatial filtering for zero-order and twin-image elimination in digital off-axis holography," *Appl. Opt.* **39**, 4070–4075 (2000).
 20. C. Liu et al., "Elimination of zero-order diffraction in digital holography," *Opt. Eng.* **41**(10), 2434–2437 (2002).
 21. N. Pavillon et al., "Artifact-free reconstruction from off-axis digital holograms through nonlinear filtering," *Proc. SPIE* **7723**, 77231U (2010).
 22. H. Xia, M. Li, and M. Tang, "Contrast between the wavelet transform with coefficients selection method and the traditional frequency domain filtering method for digital hologram reconstruction," *Proc. SPIE* **7848**, 78481I (2010).
 23. N. Pavillon et al., "Suppression of the zero-order term in off-axis digital holography through nonlinear filtering," *Appl. Opt.* **48**, H186–H195 (2009).
 24. Z. Ma et al., "Numerical iterative approach for zero-order term elimination in off-axis digital holography," *Opt. Express* **21**, 28314–28324 (2013).
 25. P. Ferraro et al., "Extended focused image in microscopy by digital holography," *Opt. Express* **13**, 6738–6749 (2005).
 26. T. Colomb et al., "Numerical parametric lens for shifting, magnification, and complete aberration compensation in digital holographic microscopy," *J. Opt. Soc. Am. A* **23**, 3177–3190 (2006).
 27. L. T. Bang et al., "Compression of digital hologram for three-dimensional object using wavelet-bandelets transform," *Opt. Express* **19**, 8019–8031 (2011).
 28. D. Taubman, "EBCOT: embedded block coding with optimized truncation," *IEEE Trans. Image Process.* **9**(7), 1158–1170 (1998).
 29. D. Taubman, "High performance scalable image compression with EBCOT," *Int. Conf. Image Process.* **9**(7), 1158–1170 (2000).
 30. S. G. Mallat, "A theory for multiresolution signal decomposition: the wavelet representation," *IEEE Trans. Pattern Anal. Mach. Intell.* **11**, 674–693 (1989).
 31. M. D. Adams and F. Kossentini, "Reversible integer-to-integer wavelet transforms for image compression: performance evaluation and analysis," *IEEE Trans. Image Process.* **9**(6), 1010–1024 (2000).
 32. C.-L. Chang and B. Girod, "Direction-adaptive discrete wavelet transform for image compression," *IEEE Trans. Image Process.* **16**, 1289–1302 (2007).
 33. T. Bruylants et al., "On the use of directional transforms for still image coding," *Proc. SPIE* **8135**, 81350L (2011).
 34. T. Bruylants, A. Munteanu, and P. Schelkens, "Wavelet based volumetric medical image compression," *Elsevier Signal Process.: Image Commun.*, accepted (2014).
 35. G. Bjontegaard, "Calculation of Average PSNR Differences Between RD-Curves," in *VCEG M33*, ITU-T, Austin, Texas (2001).
 36. P. Schelkens, A. Skodras, and T. Ebrahimi, *The JPEG 2000 Suite*, Wiley Publishing, Chichester, West Sussex, United Kingdom (2009).
 37. Z. Liu, L. Karam, and A. Watson, "JPEG2000 encoding with perceptual distortion control," *IEEE Trans. Image Process.* **15**, 1763–1778 (2006).

David Blinder is a PhD student working at the Vrije Universiteit Brussel (VUB). He received his BSc degree in electronics and information technology engineering from the VUB and graduated with the MSc degree in applied sciences and engineering at the VUB and the École Polytechnique Fédérale de Lausanne (EPFL) in 2013. His research focuses on the efficient representation and compression of static and dynamic holograms.

Tim Bruylants graduated with an MSc degree in 2001 at the University of Antwerp. In 2005, he participated as a member of the Forms Working Group (W3C). In 2006, he became a PhD student at the VUB. His main research topic is the compression of medical volumetric datasets. He is an active member of the ISO/IEC JTC1/SC29/WG1 (JPEG) and WG11 (MPEG) standardization committees. He is coeditor of the JPEG 2000 Part 10 (JP3D) specification.

Heidi Ottevaere has been a full professor at Vrije Universiteit Brussel (VUB), since 2009. She is responsible for the instrumentation and metrology platform at the Photonics Innovation Center and for the biophotonics research unit of the Brussels Photonics Team (B-PHOT). She is coordinating and working on multiple research and industrial projects focusing on the design, fabrication, and characterization of different types of photonic components and systems in the field of biophotonics, interferometry, holography, and imaging.

Adrian Munteanu has been a professor at VUB since 2007 and a research leader of the 4Media group at the iMinds Institute in Belgium. He is the author or coauthor of more than 200 journal and conference publications, book chapters, patent applications, and contributions to standards. He is the recipient of the 2004 BARCO-FWO prize for his PhD work. He currently serves as an associate editor for *IEEE Transactions on Multimedia*.

Peter Schelkens currently holds a professorship at the VUB and is a research director at the iMinds Research Institute. In 2013, he obtained an EU ERC Consolidator Grant focusing on digital holography. He (co-)authored over 200 journal and conference publications and books. He is an associate editor of the *IEEE Transactions on Circuits and Systems for Video Technology*. He is also participating in the ISO/IEC JTC1/SC29/WG1 (JPEG) and WG11 (MPEG) standardization activities.