Unitary Transforms Using Time-Frequency Warping for Digital Holograms of Deep Scenes

David Blinder ⓘ, *Student Member, IEEE*, Colas Schretter, *Member, IEEE*, Heidi Ottevaere, *Member, IEEE*, Adrian Munteanu ⓘ, *Member, IEEE*, and Peter Schelkens, *Member, IEEE*

Abstract—**With the advent of ultrahigh-resolution holographic displays, viewing macroscopic deep scenes with large viewing angles becomes a possibility. These deep holograms possess different signal properties in contrast with common applications where the scene content is assumed to lie around a planar slice. Therefore, the conventional approach of refocusing at a fixed depth is ineffective. There is a need for an efficient invertible transformation that is able to account for the wide depth range of macroscopic three-dimensional scenes. To this end, we derive necessary invertibility conditions for the diffraction from nonplanar surfaces for symmetric light propagation kernels, such as Fresnel diffraction. We construct a unitary transform for modeling deep holographic scenes using a generalization of linear canonical transforms. From the symplectic properties of the time–frequency domain, we obtain invertibility conditions of the transforms depending on surface shape, hologram bandwidth, and wavelength. These transforms can be subsequently combined with other sparsifying transforms for compression. Experiments demonstrate one application in lossy coding of holograms by implementing a computationally efficient subset of the transforms for piecewise depth profiles that is combined with the JPEG 2000 codec. Results show improved reconstruction quality. A significant visual gain is observed as the depth information is well preserved under identical encoding rates in contrast to using Fresnel propagation at a fixed depth. This paper shows that it is possible to effectively represent holograms of variable-depth scenes and our local adaptive transform leads to a practical holographic compression framework.**

Index Terms—**Holography, Transforms, Transform Coding, Computer Graphics.**

I. INTRODUCTION

HOLOGRAPHY is a technique for capturing and reconstructing the wave field of light, encoding a three-dimensional representation of a scene. In digital holography, numerical processing is involved either after recording or for displaying the wave field. For hologram recording, the sample is placed into an interferometric setup, where typically a CCD-camera will record the interference pattern encoding the three-dimensional structure of the object. This has many metrological applications, such as cell imaging, lens characterization, deformation analysis and particle velocimetry [1]. Moreover, this approach can be used to capture a scene for viewing in a holographic display setup. Alternatively, holograms can be computer-generated to be used in display systems. Computer-generated holograms are computationally very expensive (especially for the resolutions required in display systems), requiring more advanced techniques to have tractable computing requirements [2], [3]. But even then, the computations are too costly for real-time viewing, making efficient representations an attractive choice for display systems.

Holography can be considered as the ultimate display system, since it can capture all aspects of the human visual system: continuous parallax, no accommodation-vergence conflict and it can potentially represent features at resolutions matching the diffraction limit of light. One of the main bottlenecks for current holographic displays is the bandwidth: high quality holograms for display at appreciable viewing angles need resolutions requiring hundreds of megapixels up to gigapixels. Current state-of-the-art holographic display systems often combine many high resolution spatial-light modulators to approach these resolutions [4], [5]. Efficient representation of these large data volumes is therefore an important task.

Holograms possess signal properties that differ significantly from photographic images: because no imaging lens is used, localized information will spread out over the entire hologram. Moreover, diffuse object surfaces will give rise to speckle patterns containing pervasive high-frequency information. These properties will cause standard wavelets, favoring localized features and low frequencies, to be inappropriate for efficiently representing holograms. Therefore, alternative transforms are required for effectively representing digital holograms.

A. Related Work

Some transforms address the aforementioned problem by modeling the general properties of holographic signals, e.g., addressing the prevalence of directionality (due to the interference fringes) and high-frequencies. In [6], the standard JPEG 2000 framework [7] was extended with the directional-adaptive discrete wavelet transform [8] and packet decompositions for

D. Blinder, C. Schretter, A. Munteanu, and P. Schelkens are with the Department of Electronics and Informatics (ETRO), Vrije Universiteit Brussel, Brussels B-1050, Belgium, and also with the imec, Leuven B-3001, Belgium (e-mail: dblinder@etrovub.be; cschrett@etrovub.be; acmuntea@etrovub.be; pschelke@etrovub.be).

H. Ottevaere is with the Department of Applied Physics and Photonics, Brussels Photonics Team (B-PHOT), Vrije Universiteit Brussel, Brussels B-1050, Belgium (e-mail: hottevae@b-phot.org).

improved compression performance. In [9], a non-separable adaptive vector lifting scheme was used to better encode holographic data. Wavelet-bandelets have also been used to better compress holograms [10]. The drawback of these type of approaches is the lack of "scene-awareness" (the actual objects and their properties are indirectly modeled), and the resulting nonlinear relationships between the objective and perceived distortions of the hologram after lossy compression.

A second approach is to mimic light-field representations, by cutting the hologram into small pieces and using every reconstructed hologram piece as a small aperture for a view, which can e.g., be subsequently coded with light field coders or video codecs [11]. The main drawback of this approach is the strong speckle noise present due to the limited aperture size, making intra-view and inter-view modeling difficult. An approach using compressed sensing aiming to address this problem was recently proposed [12].

Other approaches will model the diffraction of light, such as Fresnelets [13], which combine Fresnel diffraction with B-spline wavelets. This approach is limited by the fact that unitary Fresnel diffraction is confined to a fixed depth. This is sufficient for objects with shallow depth-of-focus, parallel to the camera plane (as is often the case in holographic microscopy), but causes problems for deep scenes where it is impossible to simultaneously refocus the whole scene using a single Fresnel transform, causing strong degradations after compression at out-of-focus regions [14].

B. Problem Statement

Holograms with content spanning deep volumes are suboptimally represented by methods currently found in literature (for visible light, we are typically considering scenes of at least 1 cm deep). This is mainly due to its signal properties, differing substantially from natural imagery, as shown in the example hologram in Fig. 1. We want to generalize the Fresnelet approach [13], which models diffraction for objects at a specific depth to content placed at variable depths. The goal is to describe an invertible transform closely tied to the diffraction of arbitrary-shaped surfaces that is efficiently computable and to derive the bounds for which it can operate.

C. Contributions

In this paper, we generalize the latter approach by deriving the criteria for unitary transforms modeling the diffraction of non-constant surface depths, which can be combined with wavelets to model deep scenes more accurately. The proposed transform is guaranteed to be invertible when the first derivative of the surfaces' depth map is below a specific bound (see (20)). This approach can be combined with conventional transforms and codecs to improve compressibility. Our particular contributions are as follows:

- We make the connection between unitary transforms represented by warpings of symplectic manifolds [15] and Linear Canonical Transforms (LCTs) and their generalizations; we subsequently apply this framework to express our proposed transform and derive its properties.
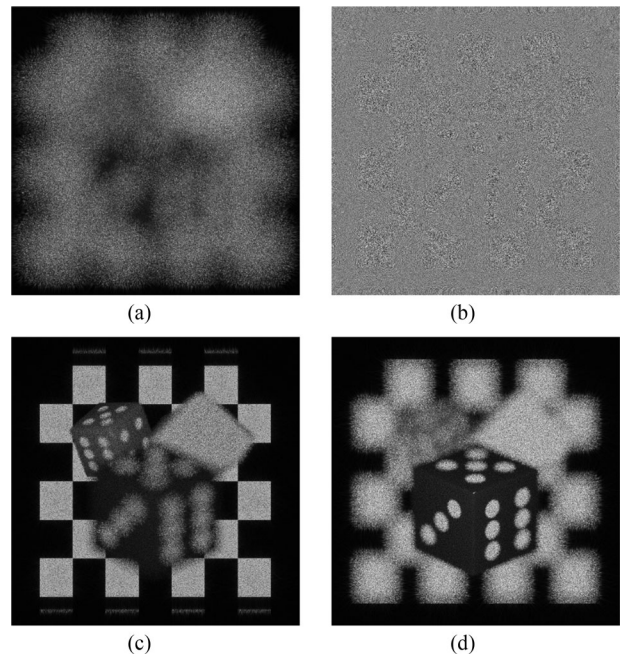

Fig. 1. Uncompressed dice hologram data: (a) absolute value, (b) phase, (c) focused at back plane, and (d) focused at front plane. Source: Advanced Media Coding Lab at IRT b-com.

- We generalize the approach of combining wavelets with convolutional Fresnel diffraction to an efficiently computable approximation for non-planar surfaces.
- We derive the bounds required for invertibility of the proposed transform.
- We implement an approximation of the transform and combine it with the JPEG 2000 codec, evaluating the increase in compression performance and visual quality.

D. Outline

In Section II, we describe the signal properties of holograms and motivate the choice of a depth-adaptive transform. In Section III, we demonstrate the unitarity of arbitrary-surface diffraction and derive the invertibility conditions for diffraction kernels in general and Fresnel diffraction in particular. Then, in Section IV, we explore potential application of the method and propose a compression framework combining our proposed transform with the JPEG 2000 codec. In Section V, we compare different versions of the transform and report significant improvements in the visual quality of the coded holograms over existing methods. We briefly discuss some of the limitations in Section VI, and conclude in Section VII.

II. SIGNAL PROPERTIES OF HOLOGRAMS

A. Diffraction of Light

The propagation of electromagnetic waves is governed by Maxwell's equations, describing a vectorial theory of light. When the medium is linear, isotropic, homogeneous, non-dispersive and nonmagnetic, light diffraction reduces to a

complex-valued scalar model [16]. If we consider a monochromatic coherent source, diffraction can be modelled with the Rayleigh-Sommerfield Diffraction formula [16]

$$U(\mathbf{p}) = \frac{1}{i\lambda} \iint_S U(\mathbf{x}) \frac{e^{ik\|\mathbf{p}-\mathbf{x}\|}}{\|\mathbf{p}-\mathbf{x}\|} \cos(\mathbf{n}, \mathbf{p}-\mathbf{x}) d\mathbf{x} \qquad (1)$$

evaluating the scalar field U at a point $\mathbf{p}$, where λ is the wavelength of the light, $k = \frac{2\pi}{\lambda}$ is the wavenumber, $\mathbf{x} \in S$ are points on a surface S over which is integrated and $\mathbf{n}$ is the surface normal of S in $\mathbf{x}$. The aforementioned diffraction integral is generally not efficiently computable: the integral has to be evaluated over all points $\mathbf{x}$ for every point $\mathbf{p}$. Nevertheless, if the surface S is planar, diffraction to a parallel target plane can be written as a convolution, which is much more efficient. Such a propagation can be modelled precisely using the Angular Spectrum Method [16]:

$$\mathcal{A}(\omega_x, \omega_y, z) = \mathcal{A}(\omega_x, \omega_y, 0) e^{i\pi z \sqrt{\lambda^{-2} - \omega_x^2 - \omega_y^2}} \qquad (2)$$

where $\mathcal{A}(\omega_x, \omega_y, z)$ is the 2D Fourier transform of the wavefield in the xy plane at depth z.

For most applications, we can approximate the spherical point spread functions further with parabolic wavefronts by using a quadratic Taylor approximation of the diffraction kernel; this reduces to Fresnel diffraction:

$$U(x, y, z) = \frac{e^{ikz}}{i\lambda z} \iint_{\mathbb{R}^2} U(x', y', 0) e^{\frac{ik}{2z}[(x-x')^2 + (y-y')^2]} dx' dy' \qquad (3)$$

This integral can be efficiently computed using convolutions with chirp-functions. By rewriting the (unitary) Fresnel transform as a convolution, we get:

$$U(x, y, z) = U(x, y, 0) * K_z(x, y)$$

$$\text{with } K_z(x, y) = \frac{1}{\sqrt{\lambda z}} e^{\frac{i\pi}{\lambda z}(x^2 + y^2)} \qquad (4)$$

This means we can describe the propagation of light by convolving with a quadratic chirp function. Due to the separability of the Fresnel transform, we will mostly describe holographic signals as one-dimensional, confined to the xz-plane for notational simplicity.

With increasing z, the light will be increasingly spread over all pixels of the hologram. This is the main reason why holograms exhibit signal characteristics which differ significantly from natural images. Methods for effectively addressing this (and other) issues will be discussed in the next section.

B. Representing Holograms Efficiently

Holography is often used for metrology in small-scale/microscopic applications. Because of the small distances, the depth of the object will be encoded as a phase profile, since these small depth variations will induce varying phase delays. This allows digital holographic microscopes to acquire the phase shift image of a sample, giving quantifiable information about the optical distance besides the conventional bright field image. In this type of setup, the object generally resides at a fixed depth-plane parallel to the camera. This allows for refocusing

the object using an (inverse) Fresnel transform with the right z, resulting in a wavefield looking like a natural image. This is why e.g., Fresnelets are very effective for representing these types of holograms: they essentially consist of B-spline wavelets combined with a Fresnel transform. Unfortunately, in the macroscopic case, involving large viewing angles, this approach often does not work anymore. There are two main reasons:

1) the scene is deeper, consisting of objects at multiple depths. Due to the large hologram aperture size, it is not possible to simultaneously have the whole scene in focus using a single Fresnel transform. Fresnelets will be ineffective at representing depths too far from the chosen z.

2) diffuse surfaces and large viewing angles. In holography, the diffraction angle of light θ is described by the following relationship: $\lambda f = \sin(\theta)$, where f stands for the spatial frequency of the signal.

Diffuse surfaces, emitting light in all directions, will therefore contain frequencies covering the whole frequency spectrum, unlike natural imagery which tend to exhibit a power spectrum proportional to $1/f^2$ [17]. The diffusiveness, combined with the large hologram aperture, will severely reduce the depth of focus, making Fresnelet-like approaches less effective. Moreover, wavelets will strongly penalize high frequencies after compression, filtering out off-axis viewing angles and leading to a "keyhole-effect": objects will appear to be in a tunnel and depth perception will be impaired because of the blurring of sharp details.

C. No Optimal Transform for Hologram Compression

Suppose we have an optimal discrete transform $A \in \mathrm{GL}(n, \mathbb{C})$ (general linear group) for a specific holographic signal emanating from (multiple) objects x. Then $y = Ax$ will be highly suitable for subsequent coding. However, the objects can in principle be present at many different depths. A (thin) object displaced at some z will result in a diffraction pattern on the hologram, which can e.g., be modeled by the Fresnel transform of x at distance z, written as $x_z = K_z x$. The signal on the hologram will thus be efficiently represented by $y = \tilde{A}x_z = AK_{-z}x_z$.

However, we want to assess whether we can effectively represent scenes containing multiple depths. For $\lim_{z \to 0} K_z$, we get $K_0 = I$, the identity matrix. For large propagation distances, we can use the Fraunhofer diffraction model, which is the Fourier transform of x (up to a constant): $\lim_{z \to \infty} K_z = \mathcal{F}$. This means that (in the limit) we would need a transform A which can be simultaneously sparse in both the spatial and the frequency domain, as well as for representing all depths in between. This is impossible due to the Heisenberg uncertainty principle, meaning that any transform A will by definition be suboptimal, resulting in some sparsity concessions for holograms in general. To address this limitation one could use overcomplete transforms or adaptive transforms. Our goal is to address the latter in the next section: how can we construct transforms for modeling objects with variable depth content?
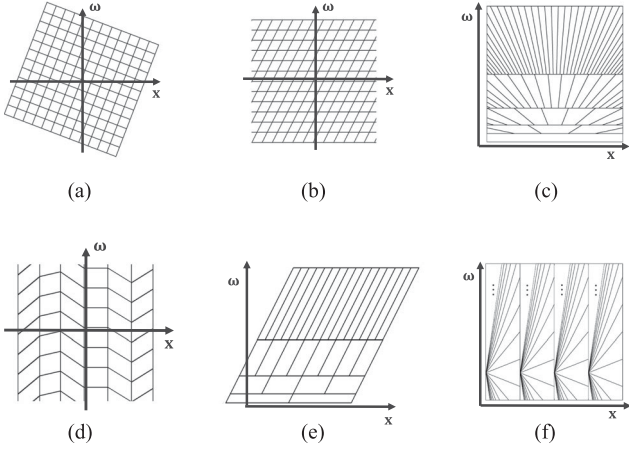
Fig. 2. Representations of the Heisenberg boxes in Wigner space corresponding to various (composite) unitary transforms: (a) fractional Fourier transform, (b) Fresnel transform, (c) fan-chirp transform with wavelets, (d) modulated Malvar-Wilson basis, (e) fresnelets, and (f) resetting chevron basis. Source (c) and (f): [21].

Fig. 3. Time-frequency warpings associated to the Fresnel diffraction of surfaces with bounded first derivatives. The top row (C_∞) is the original continuous surface, the second row (C_1) is the piecewise linear approximation used as a proposed transform in this paper; the last row (C_0) corresponds to a regular Fresnel transform, which is inadequate to efficiently represent surfaces with strongly varying depth.

III. INVERTIBILITY CONDITIONS FOR THE PROPAGATION OF NON-PLANAR SURFACES

A. Unitary Transform as Tilings of the Time-Frequency Domain

Classical linear analysis of signals using Fourier transforms, conventional wavelets or Gabor atoms are often suboptimal for the analysis and representation of many types of highly non-stationary signals, such as speech [18], radar [19] and holograms among many others. It is often useful to describe the properties of such signals indirectly using time-frequency (TF) analysis [20], which is a $2n$-dimensional representation of a n-dimensional signal depicting time and frequency domains simultaneously. We will refer to this representation as TF space, which is also known in optics as "Phase space" or "Wigner space".

Specifically, several unitary transforms have been constructed with non-axis aligned Heisenberg boxes [21], [22] (Fig. 2). Some of these can be computed using operations such as axis resamplings (e.g., Baraniuk's Fan-chirp and Resetting Chevron bases [21]); others use multiple Fourier transforms and all-pass filters, such as the Linear Canonical Transform (LCT) [23]–[25], on which we will elaborate.

LCTs are linear transforms of the TF domain described by the symplectic matrix group $\mathrm{Sp}(2n, \mathbb{R})$, defined by

$$S \in \mathrm{Sp}(2n, \mathbb{R}) \subseteq \mathrm{SL}(2n, \mathbb{R}) \quad \Longleftrightarrow \quad S^\mathsf{T} J S = J \quad (5)$$

where $\mathrm{SL}(2n, \mathbb{R})$ is the special linear group. J is given by

$$J = \begin{pmatrix} 0_n & I_n \\ -I_n & 0_n \end{pmatrix} \quad (6)$$

where 0_n is a $n \times n$ matrix of zeros and I_n is a $n \times n$ identity matrix. For every element $S \in \mathrm{Sp}(2n, \mathbb{R})$ given by

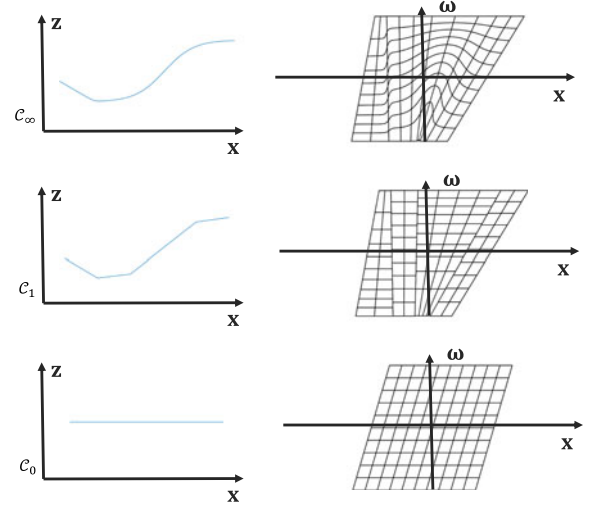$$S = \begin{pmatrix} A & B \\ C & D \end{pmatrix} \quad (7)$$

where A to D are $n \times n$ matrices, there is a bijective mapping between the double covering of $\mathrm{Sp}(2n, \mathbb{R})$ and the metaplectic group $\mathrm{Mp}(n, \mathbb{R})$, which are composed of unitary operators; namely, every symplectic matrix is mapped to a pair of unitary operators differing only by a sign. The mapping operator $\mathcal{M}_S$ for the typical use case (when $\det B \neq 0$) is defined as [25]

$$\mathcal{M}_S[f](\tilde{\mathbf{x}}) := (\det iB)^{-1/2}$$
$$\cdot \int_{\mathbb{R}^n} e^{i\pi(\tilde{\mathbf{x}}^\mathsf{T} D B^{-1} \tilde{\mathbf{x}} - 2\mathbf{x}^\mathsf{T} B^{-1} \tilde{\mathbf{x}} + \mathbf{x}^\mathsf{T} B^{-1} A \mathbf{x})} f(\mathbf{x}) d\mathbf{x}. \quad (8)$$

The LCTs comprise among others the Fourier transform (a 90° rotation matrix), fractional fourier transforms (arbitrary rotation matrices) and the Fresnel transform (shearing matrix); see Fig. 2. By combining the unitary Fresnel Transform with B-spline wavelets, we get the Fresnelets [13]. Unfortunately, LCTs are insufficient for accurately describing diffraction from a surface with non-constant depth: we would need a transform with a varying amount of shearing in TF-space, which cannot be expressed using a LCT; this is illustrated in Fig. 3.

The question is: can we construct unitary transforms with these desired properties? To maximize the potential sparsity, we would like design reversible transforms which accurately model the diffraction of surfaces with a non-constant depth map. In the following subsection, we will derive the surface constraints to guarantee invertibility, using a generalization of the LCT.

B. Non-Planar Diffraction as a Spatially-Varying Convolution

Diffraction $u'(x)$ of a signal $u(x)$ on a planar surface in a perpendicular direction with a distance z can be described as a convolution with some kernel $h(x)$:

$$u'(x) = u(x) *^z h(x) \quad (9)$$

The power convolution $*^z$ with a real valued exponent z can be described conventionally in Fourier space as

$$U'(\omega) = U(\omega) \cdot H(\omega)^z \tag{10}$$

Suppose now we have a surface with a variable depth profile described by $z(x)$ emitting light with complex amplitude $u(x)$. The measured signal on the hologram is then described by the spatially-varying convolution:

$$u'(x) = u(x) *^{z(x)} h(x) \tag{11}$$

When $z(x)$ is constant, this reduces to the traditional case where we get a unitary transform using e.g., the Fresnel transform or the Angular Spectrum Method. Otherwise, we cannot directly express the transform in Fourier space because of the space-varying property. We can use a different approach to determine for what $z(x)$ and $h(x)$ the expression is unitary.

For determining these relationships, we would like to be able to generalize the mapping $\mathrm{Sp}(2n, \mathbb{R}) \leftrightarrow \mathrm{Mp}(2n, \mathbb{R})$, i.e., describing nonlinear canonical transforms. For this purpose we will be using mathematical tools from classical quantum mechanics. There is a close relationship between the phase space of particles (describing position and velocity) and the time-frequency representation of signals; e.g., Heisenberg uncertainty principle originating from quantum mechanics is used in signal processing to express the trade-off between spatial and frequency localization of transform elements.

C. The Heisenberg Uncertainty Principle Expressed in Symplectic Geometry

We will use one of the main results from [15], where it has been proven that there exists a bijective correspondence between "Hamiltonian symplectomorphisms" and unitary transforms:

$$\mathcal{P} \, \mathrm{Ham}(2n, \mathbb{R}) \leftrightarrow \mathcal{P} \, \mathrm{U}(L^2(\mathbb{R}^n)) \tag{12}$$

- We denote by $\mathcal{P} \, \mathrm{Ham}(2n, \mathbb{R})$ the set of all one-parameter families $f_t \in \mathrm{Ham}(2n, \mathbb{R})$ depending smoothly on t and passing through the identity at $t = 0$.
- $\mathcal{P} \, \mathrm{U}(L^2(\mathbb{R}^n))$ is the set of all strongly continuous one-parameter families of unitary operators F_t on $L^2(\mathbb{R}^n)$ depending smoothly on t such that F_0 is the identity operator.

$\mathrm{Ham}(2n, \mathbb{R})$ is a particular group of diffeomorphisms, describing operations on the $2n$-dimensional TF space representing n-dimensional signals. A diffeomorphism is an isomorphism between two smooth manifolds $\mathcal{M}$ and $\mathcal{N}$. When the source and target manifold are the same $\mathcal{M} = \mathcal{N}$, i.e., an automorphism, it can be seen intuitively as a smooth invertible warping of the manifold (in this case the TF space).

A diffeomorphism is also a symplectomorphism $\mathrm{Symp}(2n, \mathbb{R})$ when it preserves the symplectic structure, effectively satisfying the Heisenberg uncertainty of a transform. For this to hold, the diffeomorphism must be volume-preserving, but this is a necessary, not a sufficient condition. An intuitive understanding of symplectomorphisms is the reformulation of Gromov's non-squeezing theorem [26], stating that any $2n$-dimensional phase space ball cannot be deformed by a canonical transform so that it can be squeezed through a hole in any plane of
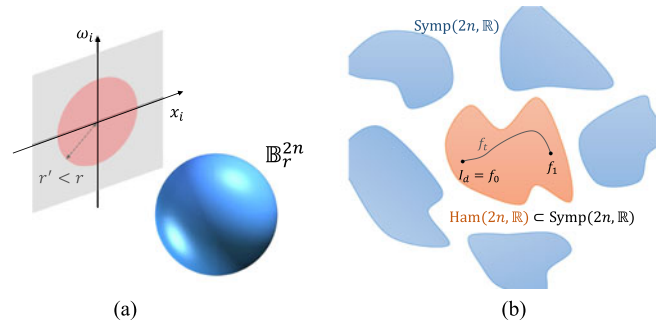

Fig. 4. Illustrations of the properties of symplectic geometry. (a) A symplectomorphism cannot squeeze a phase space ball $\mathbb{B}_r^{2n}$ with radius r into a hole in any plane (x_i, ω_i) with radius $r' < r$. (b) Visualization of $\mathrm{Ham}(2n, \mathbb{R}) \subset \mathrm{Symp}(2n, \mathbb{R})$, which is the path-component containing the identity transform.

conjugate coordinates (x_i, ω_i) (i.e., pair of joint time and frequency axes), whose hole is smaller than the cross-section of the ball (see Fig. 4(a)).

This means that the Heisenberg uncertainty principle of a n-dimensional signal shall not be violated when looking at that signal along any axis. In the 1D case, the symplectic condition reduces to area-preservation of any neighborhood in the 2D TF space. For a more rigorous description of symplectic manifolds, we refer to the appendix.

A symplectomorphism f_1 is also Hamiltonian when it is isotopic with the identity mapping. In topological terms, $\mathrm{Ham}(2n, \mathbb{R})$ is the path-component of $\mathrm{Symp}(2n, \mathbb{R})$ containing the identity. Concretely, there should be a continuous path in the space of symplectomorphisms $f_t \in \mathrm{Symp}(2n)$ parametrized by $t \in [0, 1]$ connecting the identity mapping $f_0 = I_d$ with the target symplectomorphism f_1 (Fig. 4(b)).

Using these properties, we can express desired properties of a transform directly in TF domain, while guaranteeing that the corresponding transform is unitary (and therefore invertible). This approach allows us to tie in to and generalize the well-known LCT and can simplify calculations that might otherwise involve lengthy integral equations.

D. Non-Planar Surface Diffraction as a Unitary Transform

Without loss of generality, we can state that any unitary propagation kernel h expressed in Fourier space as H must be an all-pass filter of the form $H(\omega) = e^{i\pi G(\omega)}$, where $G : \mathbb{R} \to \mathbb{R}$; hence, the kernel can only cause phase delays, no amplitude changes; otherwise, the transform will not preserve the energy of the signal for all possible u, violating the conditions for a unitary transform. Furthermore, we will assume the propagation kernel $h(x)$ to be symmetric; this will cause G to be symmetric as well, i.e., $G(\omega) = G(-\omega)$.

For the transform, we consider the class of signals with a finite bandwidth 2η. In the time-frequency representation, this is described by a symplectic manifold $\mathcal{M} = \mathbb{R} \times \,]-\eta, +\eta[$. This choice can not only be justified by the limited bandwidth of practical measurement systems, but is an inherent property of holography, as the bandwidth is bounded by the diffraction limit determined by the wavelength of the coherent light source.
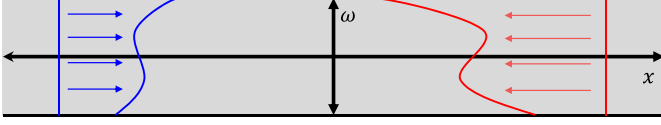
Fig. 5. Visual representation of a symplectomorphism operating on the TF domain for signal points corresponding to a space-varying convolution. Every signal point of $u(x)$ corresponds to a vertical line (much like a Dirac pulse) in the TF domain; every such line will be warped by a convolution with a different kernel with differing phase shifts. The warping of the points within a line is given by the sought $\xi(x, \omega)$.
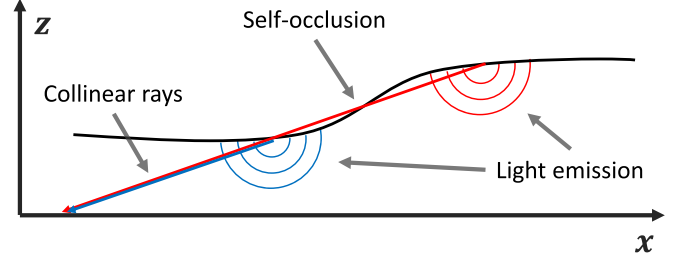

Fig. 6. The angular spectrum method models the diffraction of spherical waves, meaning that any non-planarity will ensue in some self-occlusion (and thus prevent reversibility). Intuitively, this can be seen in the figure, where two collinear rays map onto the same point on the detector at $z = 0$, making them indistinguishable.

Therefore, every point of the signal, corresponding to a vertical bounded line in $\mathcal{M}$, similar to a Dirac pulse, will experience a frequency-dependent phase delay described by the diffeomorphism $\Phi : (x, \omega) \mapsto (x + z_t(x)g(\omega), \omega)$, with the instantaneous frequency $g(\omega) = \frac{\partial}{\partial \omega} G(\omega)$ as illustrated in Fig. 5.

This still gives us a degree of freedom, as we still can choose how to redistribute the points within each line. This is determined by the diffeomorphism $\mathcal{W} : (x, \omega) \mapsto (x, \xi_t(x, \omega))$, warping the points using (an unknown) $\xi_t(x, \omega)$.

We therefore have to find out when the following composite diffeomorphism is a symplectomorphism:

$$\mathcal{T} = \Phi \circ \mathcal{W} = \begin{cases} \hat{x} = x + z_t(x)g(\xi_t(x, \omega)) \\ \hat{\omega} = \xi_t(x, \omega) \end{cases} \quad (13)$$

where we choose $z_t(x) = t \cdot z(x)$, with $t \in [0, 1]$. $z_0(x) = 0$ will thus lead to the identity mapping and $z_1(x) = z(x)$ will be the target depth profile. From the appendix, we can infer that the condition for symplecticness for a 1D signal is given by

$$\frac{\partial \hat{x}}{\partial x} \cdot \frac{\partial \hat{\omega}}{\partial \omega} - \frac{\partial \hat{\omega}}{\partial x} \cdot \frac{\partial \hat{x}}{\partial \omega} = 1 \quad (14)$$

If this holds for all $t \in [0, 1]$, we have a Hamiltonian symplectomorphism. We get the following differential equation:

$$\left(\frac{\partial}{\partial \omega} \xi_t(x, \omega) \right) \left(\frac{\partial}{\partial x} z_t(x) \cdot g(\xi_t(x, \omega)) + 1 \right) = 1 \quad (15)$$

After solving for $\xi_t(x, \omega)$, we get:

$$\frac{\partial}{\partial x} z_t(x) \cdot G(\xi_t(x, \omega)) + \xi_t(x, \omega) = \omega + \mathcal{C} \quad (16)$$

The constant $\mathcal{C}$ is yet to be determined. Because the diffeomorphism must also be Hamiltonian, $\xi_t(x, \omega)$ must be an isotopy parametrized by t for every chosen x. Isotopies on a line interval must preserve point ordering, from which it follows that the limiting extremities $\xi_t(x, \pm\eta) = \pm\eta$ are fixpoints, giving us $\mathcal{C} = \partial_x z_t G(\pm\eta)$:

$$\frac{\partial}{\partial x} z_t(x) \Big(G(\xi_t(x, \omega)) - G(\pm\eta) \Big) + \xi_t(x, \omega) = \omega \quad (17)$$

Symplectomorphisms are invertible, so $\xi_t(x, \omega)$ must be invertible in ω as well, for every x, t. We can find the inverse mapping

$$\xi_t^{-1}(x, \omega) = \frac{\partial}{\partial x} z_t(x) \Big(G(\omega) - G(\pm\eta) \Big) + \omega \quad (18)$$

Because of their invertibility property and because of $\xi_0(x, \omega) = \omega, \forall x \in \mathbb{R}$ both $\xi_t(x, \omega)$ and $\xi_t^{-1}(x, \omega)$ have to be strictly increasing, i.e.,

$$\frac{\partial}{\partial \omega} \xi_t^{-1}(x, \omega) > 0 \quad (19)$$

From (18) and (19), we get the inequality $\partial_x z_t g(\omega) > -1$. Given the symmetry of G, g must be antisymmetric. We thus get the equivalent condition $\partial_x z_t g(\omega) < 1$ to hold as well. Combining these properties, we get the condition for symplecticness:

$$\left| \frac{\partial}{\partial x} z(x) \right| < \inf_{]-\eta, +\eta[} \left(\left| \frac{\partial}{\partial \omega} G(\omega) \right|^{-1} \right) \quad (20)$$

giving us one of the main results of this paper: for every symmetric unitary propagation kernel, any surface with a sufficiently bounded first derivative has an associated unitary transform.

It is interesting to note that this relationship is precisely inverted w.r.t. the Shannon sampling condition for phase aliasing: more bandwidth means we can unambiguously retrieve the unwrapped phase of steeper surfaces. Conversely, less bandwidth means we can invertibly transform surfaces with steeper surface features. This means we can still transform deep scenes which are aliased in their phase information, which is often the case for holograms of macroscopic objects. (21) shown at the bottom of the page.

Using the Angular Spectrum Method for the diffraction spherical waves (2), we can determine that for the one-dimensional case, the propagation kernel produces a phase delay $G(\omega) = \sqrt{\lambda^{-2} - \omega^2}$ defined for a holographic signal bandlimited by the diffraction limit $\eta = \lambda^{-1}$. When computing the infimum for

$$\begin{cases} \left(\delta_{ix} + \frac{\partial}{\partial i} z_x(x, y) \cdot g_x(\xi_x, \xi_y) \right) \left(\frac{\partial}{\partial \omega_j} \xi_x \right) + \left(\delta_{iy} + \frac{\partial}{\partial i} z_y(x, y) \cdot g_y(\xi_x, \xi_y) \right) \left(\frac{\partial}{\partial \omega_j} \xi_y \right) = \delta_{ij} \quad \text{for} \quad i, j \in \{x, y\} \\ \left(\frac{\partial}{\partial \xi_x} g_y(\xi_x, \xi_y) \cdot z_y(x, y) - \frac{\partial}{\partial \xi_y} g_x(\xi_x, \xi_y) \cdot z_x(x, y) \right) \cdot \left(\frac{\partial \xi_y}{\partial \omega_x} \frac{\partial \xi_x}{\partial \omega_y} - \frac{\partial \xi_x}{\partial \omega_x} \frac{\partial \xi_y}{\partial \omega_y} \right) = 0 \end{cases} \quad (21)$$

(20), we get:

$$\inf_{]-\lambda^{-1},\lambda^{-1}[} \left| \frac{\sqrt{\lambda^{-2} - \omega^2}}{\omega} \right| = 0 \qquad (22)$$

Therefore enforcing that the first derivative of a surface must be identically zero. This means that the angular spectrum method can only be used for parallel plane-to-plane diffraction if invertibility is required. This property can also be understood intuitively: in Fig. 6, we can observe that for any non-planar shape, the surface will necessarily self-occlude some rays and cause ambiguity when attempting to retrieve the signal.

E. Non-planar Fresnel Diffraction

We can use an approximation to the diffraction model instead by using the Taylor expansion of $G(\omega)$:

$$\sqrt{\lambda^{-2} - \omega^2} = \frac{1}{\lambda} - \frac{\lambda}{2}\omega^2 - \frac{\lambda^3}{8}\omega^4 - \frac{\lambda^5}{16}\omega^6 - \frac{5\lambda^7}{128}\omega^8 + \dots \qquad (23)$$

Using any Taylor approximation, we get a polynomial which is guaranteed to be bounded, resulting in a strictly positive infimum. Depending on the order of the chosen polynomial, we get a trade-off between precision and maximum surface derivative. The second-order approximation corresponds to Fresnel diffraction (amounting to parabolic wavefronts), which is sufficiently precise for most applications involving holography. Using this approximation, the symplectomorphism becomes:

$$\begin{cases} \hat{x} = x - \lambda z(x)\xi(x,\omega) \\ \hat{\omega} = \xi(x,\omega) \end{cases} \qquad (24)$$

We now get a spatially varying adaptive shear of the TF domain. Note that if we pick $z(x) = z$ to be constant, it reduces to a regular shear, resulting in a conventional LCT. Furthermore, $\xi(x,\omega)$ and its inverse become:

$$\xi(x,\omega) = \begin{cases} \frac{1 - \sqrt{\lambda^2 \eta^2 (\partial_x z(x))^2 - 2\lambda \partial_x z(x)\omega + 1}}{\lambda \partial_x z(x)} & \text{if } \partial_x z(x) \neq 0 \\ \omega & \text{if } \partial_x z(x) = 0 \end{cases} \qquad (25)$$

$$\xi^{-1}(x,\omega) = -\frac{\lambda}{2}\frac{\partial}{\partial x}z(x)\omega^2 + \omega + \frac{\lambda}{2}\frac{\partial}{\partial x}z(x)\eta^2 \qquad (26)$$

If we assume that the surface $z(x)$ continuous and differentiable everywhere, $z(x)$ can be locally approximated by a linear function αx with arbitrary precision by choosing a small enough Δx. This gives us (locally) that $\partial_x z(x) = \alpha$. Using (26), we get:

$$\xi^{-1}(\omega) = -\frac{1}{2}\lambda\alpha\omega^2 + \omega + \frac{1}{2}\lambda\alpha\eta^2 \qquad (27)$$

Note that the diffeomorphism is not a function of x anymore (this is true for all valid $G(\omega)$). This means that the symplectomorphism can be computed directly in the Fourier domain, in this case using a quadratic distortion. Note that this particular transform is very similar to the Fan-chirp transform found in [21], but modulated as to have the origin of the fan to lie outside of the manifold $\mathcal{M}$. This transform effectively computes a tilted surface diffraction using Fresnel diffraction. This bears similarities to [27] where tilted diffraction is used with angular

spectrum diffraction, but the latter lacks invertibility because of (22).

How close is this (local) quadratic warping to Fresnel diffraction of a tilted plane? The coordinate warping of a source wavefield u to a target wavefield $\hat{u}$ is described via the Fourier transform $\mathcal{F}$:

$$\mathcal{F}\{\hat{u}(x)\} = \hat{U}(\omega) = U(\xi^{-1}(\omega)) = \int_{\mathbb{R}} u(x)e^{2\pi i x \xi^{-1}(\omega)}dx \qquad (28)$$

Using the Fourier inversion theorem, we get:

$$\hat{u}(\hat{x}) = \iint_{\mathbb{R}^2} u(x)e^{2\pi i(x\xi^{-1}(\omega) - \hat{x}\omega)}d\omega dx$$

$$= \int_{\mathbb{R}} \frac{u(x)}{\sqrt{i\lambda\alpha x}} e^{\frac{i\pi}{\lambda\alpha x}((x-\hat{x})^2 + (\lambda\alpha\eta x)^2)}dx \qquad (29)$$

We get the same expression as for Fresnel diffraction (in (3)), where z is substituted by αx, up to a differing attenuation in the denominator. Fortunately, this discrepancy will be reduced when the average distance of the surface to the hologram is large enough.

F. Multidimensional Non-planar Fresnel Diffraction

This approach can also be generalized to higher-dimensional signals. For a n-dimensional signal, we define an associated $2n$-dimensional symplectic manifold $\mathcal{M} = \mathbb{R}^n \times \mathbb{B}_\eta^n$, where $\mathbb{B}_r^n$ stands for the n-dimensional open ball with radius r, formally: $\mathbb{B}_r^n = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\| < r\}$. We will choose $\lambda = 1$ to simplify the subsequent expressions.

In the 2D case, the 4D symplectomorphism becomes:

$$\mathcal{T} = \begin{cases} \hat{x} = x + z_x(x,y)\frac{\partial}{\partial \omega_x}G(\hat{\omega}_x, \hat{\omega}_y) \\ \hat{y} = y + z_y(x,y)\frac{\partial}{\partial \omega_y}G(\hat{\omega}_x, \hat{\omega}_y) \\ \hat{\omega}_x = \xi_x(\omega_x, \omega_y) \\ \hat{\omega}_y = \xi_y(\omega_x, \omega_y) \end{cases} \qquad (30)$$

We drop the subscript t for notational clarity. The conditions for symplecticness are now given by a system of five partial differential equations (21), where δ is the Kronecker delta function ($\delta_{ij} = 1$ if $i = j$, $\delta_{ij} = 0$ if $i \neq j$). Note that the general expression has two depth maps z_x and z_y: this allows for the modeling of some astigmatic phenomena, where the rays in the x-plane and y-plane have different foci, which can be useful for some applications. However, for most intents and purposes, we can assume that $z = z_x = z_y$.

It is difficult to find a closed-form solution with this general description alone. We therefore make the additional assumption that points on the manifold only undergo motion in the direction of the surface gradient: otherwise, it would mean that light rays (associated to points in $\mathcal{M}$) would change direction and thus bend in mid-air, which is impossible for free space propagation. Working with these assumptions, we get the following
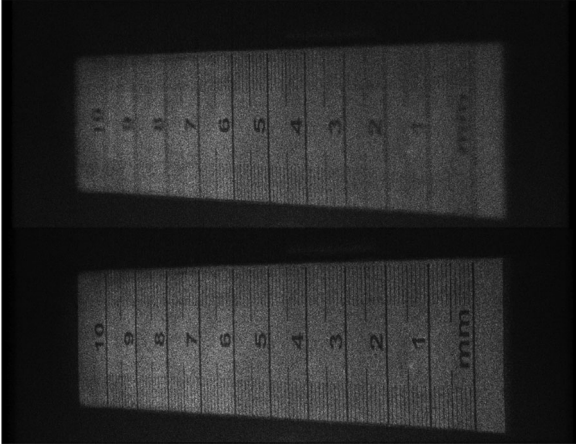
Fig. 7. Backpropagated hologram of a tilted ruler. The upper image is reconstructed by repropagating the hologram to the depth corresponding to the middle of the ruler (at 4.1 mm), the lower image uses the quadratic warping to correct for the tilt w.r.t. the hologram plane. Notice how the defocus at the ruler edges is corrected. Source: ZIF at the Warsaw University of Technology.

expressions for 2D arbitrary surface Fresnel diffraction:

$$\xi_x^{-1}(x, y, \omega_x, \omega_y) = \omega_x + \frac{\partial_x z(x,y)}{2}(\omega_x^2 + \omega_y^2 - 1) \quad (31)$$

$$\xi_y^{-1}(x, y, \omega_x, \omega_y) = \omega_y + \frac{\partial_y z(x,y)}{2}(\omega_x^2 + \omega_y^2 - 1) \quad (32)$$

Here too, we get a pure Fourier domain warping function when $z(x, y)$ is a plane.

We evaluate this warping of the Fourier space on the experimentally acquired hologram of a tilted ruler: this hologram was recorded using a camera with a pixel pitch $p = 3.105\ \mu$m and with a laser of wavelength of $\lambda = 642.1$ nm. The ruler was placed at a distance of 4.1 mm and tilted along the x-axis with an angle of approximately $12°$. The resulting reconstructions with and without frequency rewarping are shown in Fig. 7.

IV. APPLICATIONS

Because the transform is unitary, it can substitute the (unitary) Fresnel transform currently used in many existing systems to extend their application domain to deep scenes.

For the sensing of holograms with sparse/incomplete data, it could benefit applications such as viewpoint inference and despeckling of holograms [12], scene reconstruction from subsampled holograms using compressed sensing [28], or for compressive holographic tomography [29]. The transform bounds also indicate what types of object surface shapes are possible to (fully) recover from a hologram.

The proposed transform could also have applications for e.g., iterative algorithms, where the source and destination wave fields must be transformed into each other. This could benefit (non)convex optimization algorithms for resolution enhancement [30] or ping-pong algorithms for computer-generated holography [31], extended to deep scenes.

Other application of the transform could benefit depth-adaptive hologram transformations, for subsequent filtering
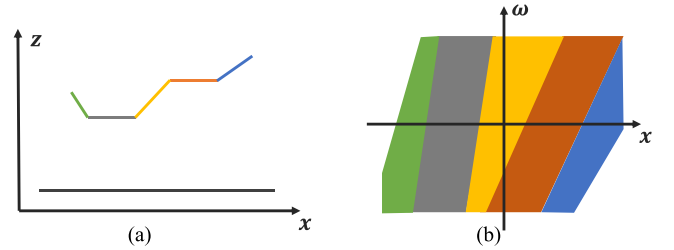

Fig. 8. The piecewise constant emissive surface approximation in (a) will map to wedge-shaped footprints in the TF domain in (b) (see matching colors). Together, they will seamlessly fold the TF domain.

and/or coding. In particular, we will focus on this paper on the lossy compression of digital holograms.

A. Lossy Compression of Deep Holograms

Although we can analytically model diffraction between arbitrary surfaces, it is not practically desirable to model the general case, especially for compression purposes, because of the following reasons:

- evaluating the integral in transform for an arbitrary $z(x, y)$ amounts to computing a different point spread function for every sample each affecting all hologram pixels, resulting in $O(n^2)$ computations for n samples.
- The associated depth map has to be estimated precisely (or given) for the corresponding hologram.
- This depth map will cause significant overhead when coding the hologram.

Instead, we propose to approximate the diffraction for coding using a piecewise polygonal approximation with only a few elements. The rationale is that small deviations from the actual depth map will only result in limited defocus and this will not affect the compressibility significantly. For planar surface pieces, we have an analytical expression in the Fourier domain (31), (32). Moreover, this approach will only require the storage of a small amount of vertices, minimizing overhead. Also, when the depth map is not available, only a coarse approximation estimate is required. Actually, this approach can be viewed as a folding of the time-frequency domain using neatly fitting tiles; this is illustrated in Fig. 8. This transform is then intended to be subsequently compatible with conventional transforms and codecs.

The suitability of this proposed transform will be evaluated for coding holograms of deep scenes. We divide each hologram in a regular square grid with a limited number of segments, which are each diagonally divided into two right triangles. We thus only need to store the depth of the corners of all grid elements. Note that we can potentially improve performance further by using meshes with adaptive triangulations, but this is currently considered out of scope.

In the forward transform, each square is propagated to its designated center using the unitary Fresnel propagation kernel, then subdivided into two right triangles which are each separately transformed in their respective Fourier domains, then
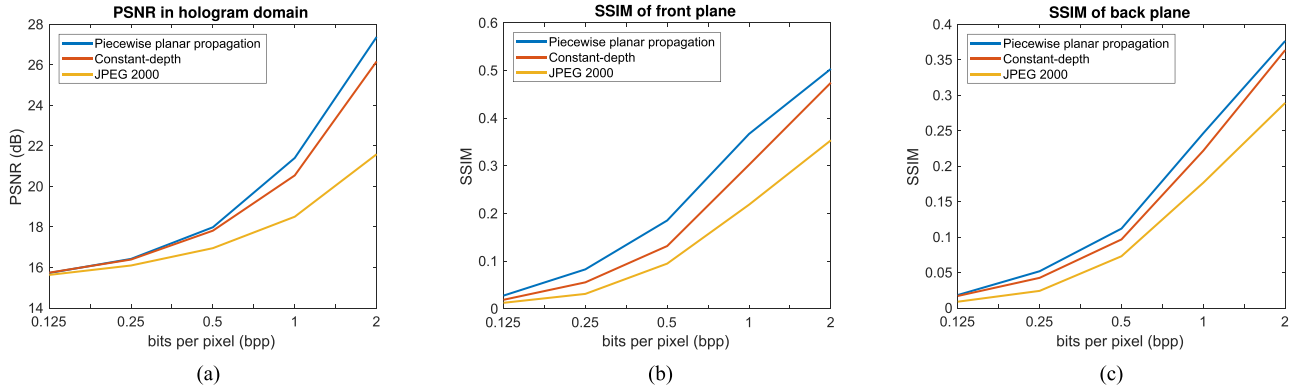
Fig. 9. Compression results using the three tested codec configurations. The PSNR (expressed in dB) is shown in (a), the SSIM of the magnitudes of the backpropagated hologram for the front and back of the scene are shown respectively in (b) and (c).
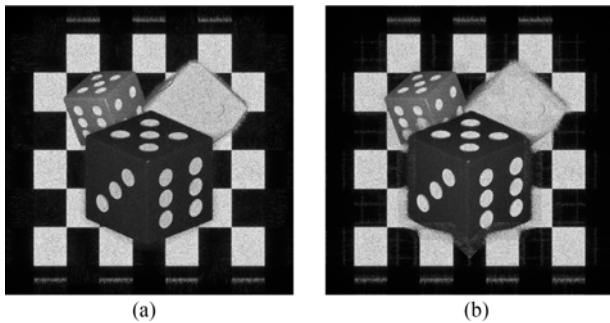

Fig. 10. Backpropagated holograms using non-planar Fresnel diffraction. (a) The non-approximated model. (b) The piecewise planar model.

resampled using the following expression:

$$U'(\omega_x, \omega_y) = \sqrt{\frac{\partial(\xi_x, \xi_y)}{\partial(\omega_x, \omega_y)}} U(\xi_x(\omega_x, \omega_y), \xi_y(\omega_x, \omega_y)) \quad (33)$$

using the Jacobian to preserve the unitarity of the axis rewarping. The resampling is done using 2D bi-cubic spline interpolation. Note that this approximation can impact the distortion at high coding rates; more accurate resampling techniques exist, but often come at increased computational costs [32].

To code the hologram, we need a (coarse) representation of the depth map. When the depth is not known, it has to be estimated. Depth estimation for shallow and smooth surfaces is fairly straightforward in holography: it mainly involves the unwrapping of the phase. This approach is unfortunately mostly infeasible for our use case, because the surfaces are rough, severely aliased in the phase information, contain speckle noise and have multiple discontinuities. For these types of holograms, depth maps could be extracted by e.g., maximizing some sharpness metric in a focal stack of refocused versions of the hologram [33].

V. Experiments

For the experiments, we use a computer-generated hologram containing multiple dices generated using the technique described in [3], see Fig. 1. The hologram is complex-valued, with a resolution of 1280×1280 pixels, with a pixel pitch of

$p = 6.4\,\mu\text{m}$. We use the red channel, corresponding to a wavelength of $\lambda = 640\,\text{nm}$. The signal has a maximal bandwidth of $2\eta = 1/p$. We use a 16×16 square grid for the piecewise planar approximation.

The experiments section is divided in two parts: first, we will compare the difference in quality and calculation time between the proposed piecewise planar approximation and the non-approximated spatially varying Fresnel transform; in the second part we will evaluate the suitability of the proposed transform for compression and compare it with other conventional transforms.

A. Comparison With the Non-Approximated Model

The non-approximated model consists of evaluating the integral (or its inverse) from (3). However, we cannot straightforwardly discretize this expression, because this would cause aliasing: the frequencies of the PSF K_z (4) for points too far from its center will surpass the Nyquist rate. To address this problem, we will use a solution inspired on the Wavefront Recording Plane method [34], but adapted to Fresnel diffraction: we only evaluate the discretized PSF at points that do not violate the Nyquist rate. Formally, the discretized integral of the propagated hologram U' for an surface z with complex amplitude U is found by:

$$U'(x, y) = \sum_{x', y'} U(x', y') \cdot K_{z(x', y')}(x - x', y - y') \cdot m_{(x, x', y, y')} \quad (34)$$

where the binary masking function m is given by

$$m_{(x, x', y, y')} = \begin{cases} 1 & \text{if} \quad \sqrt{(x - x')^2 + (y - y')^2} \leq \frac{\lambda |z(x', y')|}{2} \\ 0 & \text{otherwise} \end{cases} \quad (35)$$

Aside from avoiding aliasing, this approach will markedly improve calculation times over an integration over the complete hologram, since we only need to evaluate the expression for points within a limited radius. For the inverse transform, the
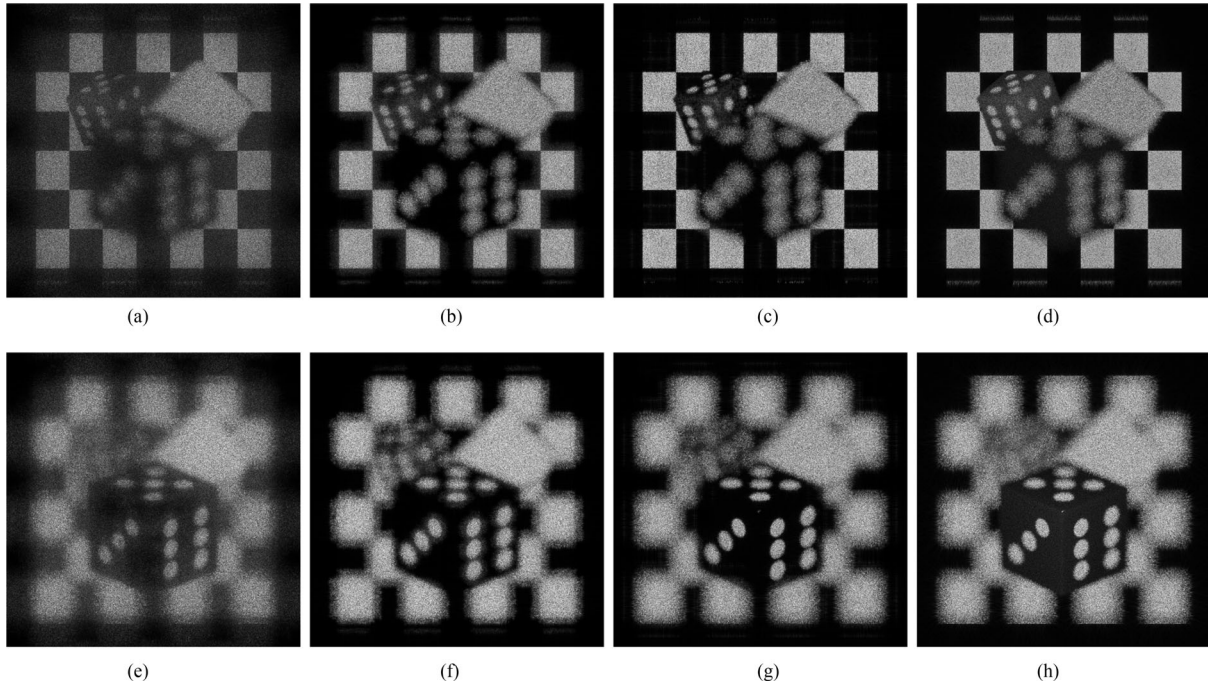
Fig. 11. The compressed dice hologram using 1 bpp. The upper row of subfigures is focused at the back plane, the lower row is focused at the front plane. (a) and (e) default JPEG 2000. (b) and (f) constant-depth transform with JPEG 2000. (c) and (g) Piecewise planar transform with JPEG 2000. (d) and (h) are the originals.

expression becomes:

$$U(x, y) = \sum_{x', y'} U'(x', y') \cdot K_{z(x,y)}(x - x', y - y') \cdot m_{(x', x, y', y)}$$

(36)

We implemented the code in MATLAB R2017b on a machine with an Intel Core i7 6700K CPU and 32GB RAM, running on the OS Windows 10. The proposed piecewise transform took 2.1 s, while the discrete integration took about 46 minutes, making it impractical for real-time applications. The transformed holograms are shown on Fig. 10. Note that the non-approximated model will not be exactly reversible, because the depth map does not satisfy (20) in every point (e.g., at the edges of the dices, where some light leaking is visible). The proposed transform causes some blurring at the dice edges because of the gradual z-change of the piecewise planar depth map, but the overall scene is sharply in focus. This property will benefit compression performance, as shown in the next subsection.

B. Evaluating Compression Performance

We compare the proposed transform with two other transforms: standard JPEG 2000, and a Fresnelet-inspired transform, which we will call the "constant-depth" transform, combining the unitary Fresnel transform propagated to the middle of the scene with the CDF97 wavelets instead of B-Spline wavelets for better compatibility with JPEG 2000. Although B-spline wavelets have superior theoretical properties in the context of diffraction [13] and have shown to be effective for autofocus applications [35], it has been shown that Fresnel transformed CDF97 wavelets are more effective at sparsely representing holograms [36]. Both the real and imaginary components of the signal are coded as separate 8-bit quantized channels, us-

ing a 4-level Mallat CDF97 wavelet decomposition and 32×32-sized codeblocks. The depth values are not encoded in the bitstream.

The compression results are shown in Fig. 9. Our proposed transform consistently outperforms the constant-depth and classic JPEG 2000 transforms. Although the gain in PSNR is rather modest, the visual difference between the reconstructed holograms is quite significant; this is shown in Figs. 11 and 12, showing reconstructed holograms compressed using various methods at 1 bpp and 0.25 bpp respectively. We also report the SSIM of the reconstructed amplitude images at the front and back planes to better quantify the improvement brought by our proposed algorithm.

Even at higher rates, JPEG 2000 introduces a lot of noise; at lower rates, the scene is barely recognizable. The constant-depth transform performs better, but still exhibits noise. Particularly, the depth information is strongly impacted: objects at their respective depths appear more blurry and the defocus is insufficient. At lower rates, the difference between the front and back refocusings is almost imperceptible. The proposed transform preserves the depth information at both rates, and has sharper details.

VI. LIMITATIONS

In this section, we briefly discuss some of the limitations of the method in its current form, and some potential solutions. Some type of holograms or scenes that will be suboptimally representable using the proposed transform are the following:

- Holograms with too high surface gradients, surpassing the bound of (20). Using the piecewise linear approximation will partially remedy this shortcoming by regularizing the depth map, as is the case for the used dice hologram. But
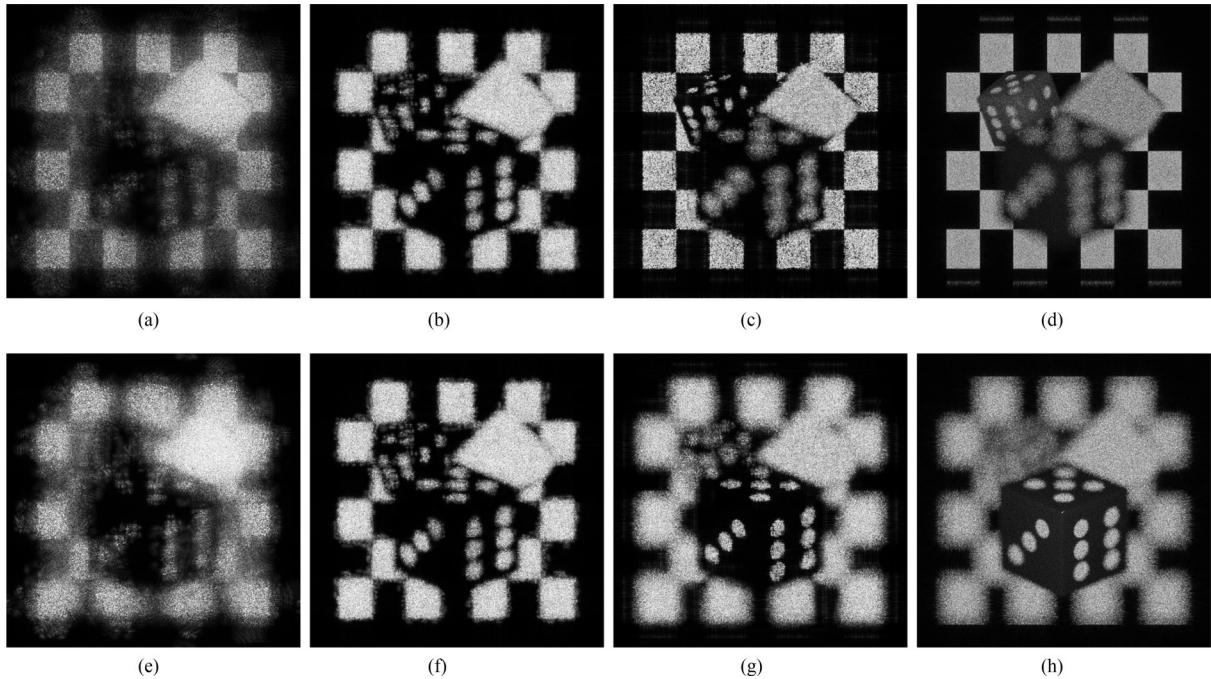
Fig. 12. The compressed dice hologram using 0.25 bpp. The upper row of subfigures is focused at the back plane, the lower row is focused at the front plane. (a) and (e) default JPEG 2000. (b) and (f) constant-depth transform with JPEG 2000. (c) and (g) Piecewise planar transform with JPEG 2000. (d) and (h) are the originals.
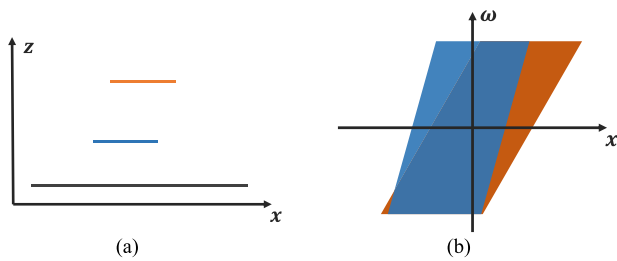

Fig. 13. Example of a failure case. The scene in (a) cannot be accurately represented using a (single) depth map, and has a discontinuity which is too large. The corresponding TF segments in (b) will overlap because of occlusion, violating the condition for a unitary transform.

this will not work for the general case, where objects are very steep or have many large discontinuities. A potential solution would involve segmenting the TF domain into separate pieces, who each individually satisfy (20).
- TF manifolds with anisotropic bandwidths. More general expressions for TF warping would be required for such content.
- Objects placed directly behind each other in wide viewing angle holograms. Such scenes cannot be faithfully represented by a single depth map (see example in Fig. 13). Aside from TF segmentation, the hologram could be divided into different views, each with a corresponding depth map.

In future work, we intend to further investigate other applications of the transform, more efficient and accurate techniques for computing the TF warping operator, generalizing the transform to broaden the application domain and overcome some of the previously mentioned limitations, and integrating the transform in a video codec.

VII. CONCLUSION

We propose a construction of unitary transforms modeling the diffraction of non-planar surfaces for deep holography. We motivate the reason and choices made for the operator, derive its properties and constraints for validity, and apply a computationally efficient approximation on a deep hologram. We also show its close ties to the Linear Canonical Transform, by using its generalization found in the domain of symplectic differential geometry. The application of data compression was explored, showing improved compressibility over conventional models, both objectively and subjectively, has been demonstrated. With this work, we hope to advance the understanding of transforms tailored for diffraction and to have an positive impact on the compression performance of digital holograms.

APPENDIX
SYMPLECTIC MANIFOLDS

Differentiable manifolds are topological manifolds with a globally defined differential structure, thereby behaving locally as a linear space, enabling one to do calculus on such a manifold. These manifolds can be equipped with a differential form, which is a coordinate-independent approach to multi-variable calculus on differential manifolds. This subject is too large to fully address in an appendix, so we refer to [37] for more information.

A symplectic manifold is an even-dimensional smooth manifold equipped with a symplectic form ω, which is a closed

non-degenerate differential 2-form. A diffeomorphism f which preserves this symplectic form is called a symplectomorphism. Formally, a mapping $f : (\mathcal{M}, \boldsymbol{\omega}) \rightarrow (\mathcal{N}, \boldsymbol{\omega}')$ defined on manifolds $\mathcal{M}$ and $\mathcal{N}$ with associated differential forms $\boldsymbol{\omega}$ and $\boldsymbol{\omega}'$ respectively is a symplectomorphism when:

$$f^* \boldsymbol{\omega}' = \boldsymbol{\omega} \qquad (37)$$

where f^* is the pull-back of f [37]. In our case, we consider f to be automorphic (i.e., $\mathcal{M} = \mathcal{N}$), meaning that the symplectic form should be invariant under f mapping $\mathcal{M}$ to itself. Writing this out, gives us:

$$f^* \left(\sum \boldsymbol{\omega}'_{kl} \mathrm{d}\mathbf{x}'_k \wedge \mathrm{d}\mathbf{x}'_l \right) = \sum \boldsymbol{\omega}_{ij} \mathrm{d}\mathbf{x}_i \wedge \mathrm{d}\mathbf{x}_j \qquad (38)$$

using the wedge product $\wedge$; equivalently,

$$\sum_{k,l} (\boldsymbol{\omega}'_{kl} \circ f) \frac{\partial}{\partial \mathbf{x}_i} f_k \cdot \frac{\partial}{\partial \mathbf{x}_j} f_l = \boldsymbol{\omega}_{ij} \qquad (39)$$

When using the symplectic normal form $\boldsymbol{\omega} = J$ (from (6)), (39) reduces to a set of partial differential equations, namely:

$$\forall i, j \in \{1, 2n\} : \sum_{k=1}^{n} \frac{\partial f_k}{\partial \mathbf{x}_i} \cdot \frac{\partial f_{k+n}}{\partial \mathbf{x}_j} - \frac{\partial f_{k+n}}{\partial \mathbf{x}_i} \cdot \frac{\partial f_k}{\partial \mathbf{x}_j} = J_{ij} \qquad (40)$$

These equations are equivalent to the Poisson bracket operator [37], [38].

As mentioned previously (12), there is a bijective correspondence between the space of Hamiltonian symplectomorphisms and (a subset of) unitary transforms. This mapping can be computed in general using the following steps [39]. First, the associated Hamiltonian H is computed using the following equation:

$$H(\mathbf{x}, t) = - \int_0^1 \mathbf{x}^T \boldsymbol{\omega} \left(\frac{\partial}{\partial t} f_t \circ f_t^{-1} \right) (\beta \mathbf{x}) d\beta \qquad (41)$$

where f_t is a Hamiltonian symplectomorphism parametrized by t (so that f_0 corresponds to the identity mapping), and f_t^{-1} is its inverse. Then, using Weyl quantization, the associated Hamiltonian operator $\hat{H}$ operating on some function $u(\mathbf{x})$ is given by:

$$(\hat{H}u)(\mathbf{x}) = \iint e^{2\pi i (\mathbf{x} - \mathbf{y})t} H \left(\frac{\mathbf{x} + \mathbf{y}}{2}, t \right) u(\mathbf{x}) d\mathbf{y} dt \qquad (42)$$

Finally, the corresponding unitary transform $\hat{U}$ can be found by solving a Schrödinger equation:

$$i\hbar \frac{\partial}{\partial t} \hat{U}_t = \hat{H} \hat{U}_t \qquad (43)$$

where typically $\hbar = \pi/2$ is chosen; $\hat{U}_0$ is the identity operator. Unfortunately, finding closed-form analytical solutions using this generic procedure is difficult (and perhaps impossible) for many practical cases. Instead, we use a different approach in the paper, but we've added this information nonetheless for completeness.

REFERENCES

[1] T. Kreis, "Application of digital holography for nondestructive testing and metrology: A review," *IEEE Trans. Ind. Informat.*, vol. 12, no. 1, pp. 240–247, Feb. 2016.

[2] A. Symeonidou, D. Blinder, A. Munteanu, and P. Schelkens, "Computer-generated holograms by multiple wavefront recording plane method with occlusion culling," *Opt. Express*, vol. 23, no. 17, pp. 22149–22161, 2015.

[3] A. Gilles, P. Gioia, R. Cozot, and L. Morin, "Hybrid approach for fast occlusion processing in computer-generated hologram calculation," *Appl. Opt.*, vol. 55, pp. 5459–5470, Jul. 2016.

[4] T. Kozacki, G. Finke, P. Garbat, W. Zaperty, and M. Kujawińska, "Wide angle holographic display system with spatiotemporal multiplexing," *Opt. Express*, vol. 20, pp. 27473–27481, Dec. 2012.

[5] H. Sasaki, K. Yamamoto, Y. Ichihashi, and T. Senoh, "Image size scalable full-parallax coloured three-dimensional video by electronic holography," *Sci. Rep.*, vol. 4, 2014, Art. no. 4000.

[6] D. Blinder, T. Bruylants, H. Ottevaere, A. Munteanu, and P. Schelkens, "JPEG 2000-based compression of fringe patterns for digital holographic microscopy," *Opt. Eng.*, vol. 53, no. 12, pp. 123102–123102, 2014.

[7] P. Schelkens, A. Skodras, and T. Ebrahimi, *The JPEG 2000 Suite*. Hoboken, NJ, USA: Wiley, 2009.

[8] C.-L. Chang and B. Girod, "Direction-adaptive discrete wavelet transform for image compression," *IEEE Trans. Image Process.*, vol. 16, no. 5, pp. 1289–1302, May 2007.

[9] Y. Xing, M. Kaaniche, B. Pesquet-Popescu, and F. Dufaux, "Vector lifting scheme for phase-shifting holographic data compression," *Opt. Eng.*, vol. 53, no. 11, pp. 112312–112312, 2014.

[10] L. T. Bang, Z. Ali, P. D. Quang, J.-H. Park, and N. Kim, "Compression of digital hologram for three-dimensional object using wavelet-bandelets transform," *Opt. Express*, vol. 19, no. 9, pp. 8019–8031, 2011.

[11] Y.-H. Seo, H.-J. Choi, and D.-W. Kim, "3D scanning-based compression technique for digital hologram video," *Signal Process., Image Commun.*, vol. 22, no. 2, pp. 144–156, 2007. (Special issue on three-dimensional video and television).

[12] Y. Rivenson, M. A. Shalev, and Z. Zalevsky, "Compressive fresnel holography approach for high-resolution viewpoint inference," *Opt. Lett.*, vol. 40, pp. 5606–5609, Dec. 2015.

[13] M. Liebling, T. Blu, and M. Unser, "Fresnelets: New multiresolution wavelet bases for digital holography," *IEEE Trans. Imag. Process.*, vol. 12, no. 1, pp. 29–43, Jan. 2003.

[14] E. Darakis, T. J. Naughton, and J. J. Soraghan, "Compression defects in different reconstructions from phase-shifting digital holographic data," *Appl. Opt.*, vol. 46, pp. 4579–4586, Jul. 2007.

[15] M. A. de Gosson and B. J. Hiley, "Imprints of the quantum world in classical mechanics," *Found. Phys.*, vol. 41, no. 9, pp. 1415–1436, 2011.

[16] J. W. Goodman, *Introduction to Fourier Optics*, 3rd ed. Englewood, CO, USA: Roberts Company Publ., Dec. 2004.

[17] A. van der Schaaf and J. van Hateren, "Modelling the power spectra of natural images: Statistics and information," *Vis. Res.*, vol. 36, no. 17, pp. 2759–2770, 1996.

[18] M. Képesi and L. Weruaga, "Adaptive chirp-based time–frequency analysis of speech signals," *Speech Commun.*, vol. 48, no. 5, pp. 474–492, 2006.

[19] F. Millioz and M. Davies, "Sparse detection in the chirplet transform: Application to FMCW radar signals," *IEEE Trans. Signal Process.*, vol. 60, no. 6, pp. 2800–2813, Jun. 2012.

[20] L. Cohen, *Time-Frequency Analysis* (Electrical Engineering Signal Processing). Englewood Cliffs, NJ, USA: Prentice-Hall, 1995.

[21] R. G. Baraniuk and D. L. Jones, "Shear madness: New orthonormal bases and frames using chirp functions," *IEEE Trans. Signal Process.*, vol. 41, no. 12, pp. 3543–3549, Dec. 1993.

[22] R. Coifman, G. Matviyenko, and Y. Meyer, "Modulated Malvar–Wilson bases," *Appl. Comput. Harmon. Anal.*, vol. 4, no. 1, pp. 58–61, 1997.

[23] M. Moshinsky and C. Quesne, "Linear canonical transformations and their unitary representations," *J. Math. Phys.*, vol. 12, no. 8, pp. 1772–1780, 1971.

[24] M. de Gosson, "Pseudo-differential operators, " in, *Symplectic Methods in Harmonic Analysis and in Mathematical Physics*. Basel, Switzerland: Springer, 2011.

[25] J. J. Healy, M. Alper Kutay, H. M. Ozaktas, and J. T. Heridan, *Linear Canonical Transforms: Theory and Applications*. New York, NY, USA: Springer, 2016.

[26] M. A. de Gosson, "The symplectic camel and the uncertainty principle: The tip of an iceberg?" *Found. Phys.*, vol. 39, pp. 194–214, Feb. 2009.

[27] K. Matsushima, H. Schimmel, and F. Wyrowski, "Fast calculation method for optical diffraction on tilted planes by use of the angular spectrum of plane waves," *J. Opt. Soc. Amer. A*, vol. 20, no. 9, pp. 1755–1762, 2003.

[28] S. Bettens, C. Schretter, N. Deligiannis, and P. Schelkens, "Bounds and conditions for compressive digital holography using wavelet sparsifying bases," *IEEE Trans. Comput. Imag.*, vol. 3, no. 4, pp. 592–604, Dec. 2017.

[29] D. J. Brady, K. Choi, D. L. Marks, R. Horisaki, and S. Lim, "Compressive holography," *Opt. Express*, vol. 17, pp. 13040–13049, Jul. 2009.

[30] C. Schretter, D. Blinder, S. Bettens, H. Ottevaere, and P. Schelkens, "Regularized non-convex image reconstruction in digital holographic microscopy," *Opt. Express*, vol. 25, pp. 16491–16508, Jul. 2017.

[31] R. G. Dorsch, A. W. Lohmann, and S. Sinzinger, "Fresnel ping-pong algorithm for two-plane computer-generated hologram display," *Appl. Opt.*, vol. 33, pp. 869–875, Feb. 1994.

[32] A. Jarrot, C. Ioana, and A. Quinquis, "Toward the use of the time-warping principle with discrete-time sequences," *J. Comput.*, vol. 2, no. 6, pp. 49–55, 2007.

[33] Z. Ren, N. Chen, and E. Y. Lam, "Extended focused imaging and depth map reconstruction in optical scanning holography," *Appl. Opt.*, vol. 55, pp. 1040–1047, Feb. 2016.

[34] T. Shimobaba, N. Masuda, and T. Ito, "Simple and fast calculation algorithm for computer-generated hologram with wavefront recording plane," *Opt. Lett.*, vol. 34, pp. 3133–3135, Oct. 2009.

[35] M. Liebling and M. Unser, "Autofocus for digital Fresnel holograms by use of a Fresnelet-sparsity criterion," *J. Opt. Soc. Amer. A*, vol. 21, pp. 2424–2430, Dec. 2004.

[36] S. Bettens, H. Yan, D. Blinder, H. Ottevaere, C. Schretter, and P. Schelkens, "Studies on the sparsifying operator in compressive digital holography," *Opt. Express*, vol. 25, pp. 18656–18676, Aug. 2017.

[37] J. Lee, *Introduction to Smooth Manifolds* (Graduate Texts in Mathematics), New York, NY, USA: Springer, 2012.

[38] K. Wolf, *Geometric Optics on Phase Space* (Theoretical and Mathematical Physics). Berlin, Germany: Springer, 2004.

[39] M. A. de Gosson, "Hamiltonian deformations of Gabor frames: First steps," *Appl. Comput. Harmon. Anal.*, vol. 38, no. 2, pp. 196–221, 2015.

Heidi Ottevaere (M'03) received the Graduate degree in electrotechnical engineering with majors in photonics from the Vrije Universiteit Brussel, Brussel, Belgium, in 1997, and the Ph.D. degree in applied sciences in 2003 from the Vrije Universiteit Brussel for her work entitled: Refractive microlenses and micro-optical structures for multiparameter sensing: a touch of micro-photonics. She has been a Full Professor with the Faculty of Engineering, Vrije Universiteit Brussel, since October 1, 2009. She is responsible for the Instrumentation and Metrology Platform, Photonics Innovation Center, Gooik, Belgium, and for the research unit Biophotonics, Brussels Photonics Team B-PHOT chaired by Prof. H. Thienpont. 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. He made research efforts in fundamental, applied, and industrial photonics. She authored 172 Web of Science Core Collection cited papers of which 76 high-impact journal papers and 96 papers in international conference proceedings. She was a promoter of 10 Ph.D.s, edited 2 conference proceedings, authored 7 chapters in books, was invited speaker at more than 25 international conferences, and is a coinventor of 2 patents.

David Blinder (S'13) was born in Brussels, Belgium, in 1990. He received the M.Sc. degree (*summa cum laude*) in applied sciences and engineering from the Vrije Universiteit Brussel, Brussel, Belgium, and the École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland, in 2013. He is with the Department of Electronics and Informatics (ETRO), VUB, where he is currently working toward the Ph.D. degree in signal processing in digital holography. His research interests include computer-generated holography, inverse problems, sparse representations, and image, video, and hologram coding.

Colas Schretter (M'07) received the M.Sc. degree (with great honors) in computer science in 2004, and the additional Master of Advanced Studies (DEA) degree from the Université libre de Bruxelles, Brusseles, Belgium, in 2006. In February 2007, he moved to Germany for a scientist position with Philips Research, Aachen, Germany. Meanwhile, he received a Marie Curie EST fellowship for the doctoral thesis, awarded with *magna cum laude* in April 2010, with the Otto-von-Guericke University. Thereafter, he was the recipient of a personal Marie Curie ERG grant for a postdoctoral project with the RWTH Aachen University and a visiting researcher with the Politecnico di Milano, Milano, Italy. He was selected for an M+Vision 2012 fellowship at MIT, Cambridge, USA, and has been in Belgium since October 2013, as a Senior Research Scientist with the ETRO Department, Vrije Universiteit Brussel. His research interests include problems in high-dimensional image formation using regularized inverse methods, multiscale statistical density estimation and quasi-Monte Carlo numerical integration.

Adrian Munteanu (M'07) received the M.Sc. degree in electronics and telecommunications from the Politehnica University of Bucharest, Bucharest, Romania, in 1994, the M.Sc. degree in biomedical engineering from the University of Patras, Patras, Greece, in 1996, and the Doctorate degree (*magna cum laude*) in applied sciences from Vrije Universiteit Brussel, Brussels, Belgium, in 2003. He is a Professor with the Department of Electronics and Informatics (ETRO), Vrije Universiteit Brussel, Brussel, Belgium. During 2004–2010, he was a Postdoctoral Fellow with the Fund for Scientific Research Flanders (FWO), Brussel, Belgium, and since 2007, he has been a Professor with VUB. His research interests include image, video and 3-D graphics coding, distributed visual processing, 3-D graphics, error-resilient coding, multimedia transmission over networks, and statistical modeling. He is the author of more than 300 journal and conference publications, book chapters, and contributions to standards and holds 7 patents in image and video coding. He was an Associate Editor for the IEEE TRANSACTIONS ON MULTIMEDIA. He was the recipient of the 2004 BARCO-FWO prize for his Ph.D. work, and of several prizes and scientific awards in international journals and conferences.

Peter Schelkens (M'98) holds a Professorship with the Department of Electronics and Informatics, Vrije Universiteit Brussel, Brussels, Belgium, and is a Research Group Leader with the imec, Leuven, Belgium. In 2013, he received an EU ERC Consolidator Grant focusing on digital holography. He authored or coauthored more than 250 journal and conference publications. He is the Coeditor of the books *The JPEG 2000 Suite* (Wiley, 2009) and *Optical and Digital Image Processing* (Wiley, 2011). His current research focuses on the multidimensional signal processing while especially focusing on cross-disciplinary research. He is a member of the ISO/IEC JTC1/SC29/WG1 (JPEG) standardization committee. He is an Associate Editor for the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY and *Signal Processing: Image Communications*.