Design of a Tight Frame of 2D Shearlets Based on a Fast Non-iterative Analysis and Synthesis Algorithm

Bart Goossens, Jan Aelterman, Hiêp Luong, Aleksandra Pižurica and Wilfried Philips

Ghent University, Dept. of Telecommunications and Information Processing, TELIN-IPI-IBBT, Sint-Pietersnieuwstraat 41, 9000 Gent, Belgium

ABSTRACT

The shearlet transform is a recent sibling in the family of geometric image representations that provides a traditional multiresolution analysis combined with a multidirectional analysis. In this paper, we present a fast DFT-based analysis and synthesis scheme for the 2D discrete shearlet transform. Our scheme conforms to the continuous shearlet theory to high extent, provides perfect numerical reconstruction (up to floating point rounding errors) in a non-iterative scheme and is highly suitable for parallel implementation (e.g. FPGA, GPU). We show that our discrete shearlet representation is also a tight frame and the redundancy factor of the transform is around 2.6, independent of the number of analysis directions. Experimental denoising results indicate that the transform performs the same or even better than several related multiresolution transforms, while having a significantly lower redundancy factor.

Keywords: multiresolution transforms, wavelets, shearlets

1. INTRODUCTION

In many applications, such as image restoration, image reconstruction, compressed sensing, it is often assumed that the "ideal" unknown image (i.e., the image that we want to recover) is sparse in a given basis (or frame). Sparseness means that within this basis (or frame), the image can be represented by a number of nonzero coefficients that is much smaller than the number of pixels in the image. Many techniques rely on the fact that such a "generic" sparsifying basis is already available, or at least will become available in the future.

Finding a good sparsifying representation for a general class of images, such as photographic images, is far from trivial. This is because images can be seen as a consequence of a complicated image formation process, consisting of physical processes (e.g., light reflection, absorption by materials, light scattering in fluids etc.) which can be characterized by *high-level* features, such as geometry, deterministic patterns (textures), material properties, lighting models, ... On the other hand, images are subject to camera distortions, such as out-of-focussedness (blur), lens distortion and/or other imperfections, ... Taking all these factors into account would yield overcomplicated and non-practical models.

Therefore, most authors focus the design of sparsifying transforms on exploiting *low-level* information, such as correlations between pixel intensities. One important class are the multiresolution transforms, which represent the image in a natural way by successively adding detail information in subsequent refinement steps. Classical tools such as the Fourier transform and the short-time Fourier transform do not allow the fine localization of image features in space, and it is not possible to determine the *exact* position of object edges. The discrete wavelet transform (DWT) offers a compromise between spatial and frequency localization, however, the transform is unable to optimally adapt to non-horizontal or non-vertical edge directions. For this reason, there has recently been a lot of interest in multiresolution representations that better adapt to the edge directions, i.e., transforms that also perform a *multidirectional* analysis. A few examples are: steerable pyramids,[1] dual-tree complex wavelets,[2] Marr-like wavelet pyramids,[3] 2-D (log) Gabor transforms,[4,5] contourlets,[6] ridgelets,[7,8] wedgelets,[9] bandelets,[10] brushlets,[11] curvelets,[12] phaselets,[13] directionlets[14] and surfacelets.[15]

The *shearlet* transform[16–18] is one of the most recent siblings in this family. This transform provides a traditional multiresolution analysis (such as the DWT) combined with a multidirectional analysis in arbitrary number of directions and is an optimally sparse representation for cartoon-like images[16] (more specifically, piecewise smooth functions with discontinuities along smooth curves). While most existing implementations of the discrete transform are either iterative or

Further author information: (Send correspondence to B. Goossens)
B. Goossens: E-mail: bart.goossens@telin.ugent.be, Telephone: +32 (0) 9 264 79 66

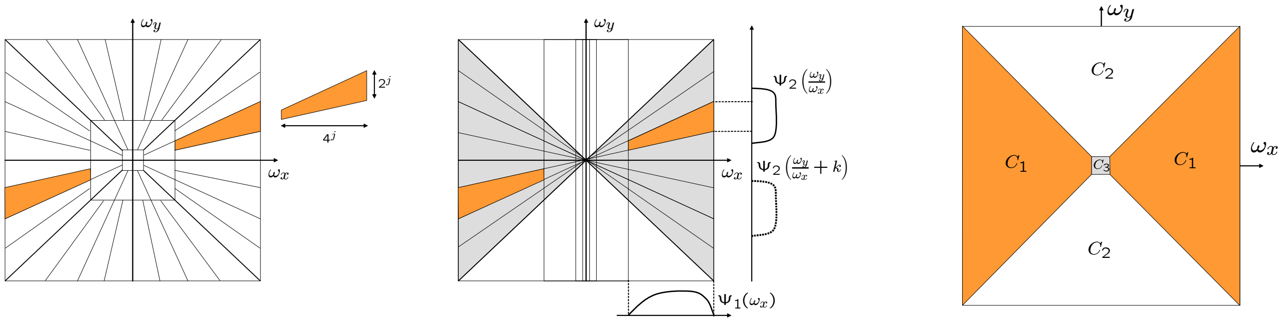

(a) (b) (c)

Figure 1. (a) Frequency tiling of the shearlet transform in trapezoidal shaped tiles (wedges).[16] (b) Individual components $\Psi_1(\omega_x)$ and $\Psi_2(\omega_y/\omega_x)$ of the Fourier transform of the shearlet mother function and the selection of orientations by the parameter k. (c) Partitioning of the 2-D frequency plane into two cones (C_1 and C_2) and a square (C_3) at the origin.

have a rather high redundancy factor, in this paper, we present a novel discrete transform that is self-invertible (which means that the backward transform is simply the adjoint transform), that has a non-iterative analysis and synthesis scheme, and a low redundancy factor. Our scheme also offers a number of other interesting features, such as alias-freeness, (approximate) shift-invariance and the ability to control the spatial support of the basis functions.

The remainder of this paper is organized as follows: in Section 2 we give a general overview of the shearlet transform. In Section 3, we present our DFT-based analysis and synthesis scheme for computing and inverting the shearlet transform. To illustrate the effectiveness of our approach, some results are given in Section 4. Finally, Section 5 concludes this paper.

2. BACKGROUND INFORMATION: THE SHEARLET TRANSFORM

The shearlet transform is a generalization of the wavelet transform with basis functions well localized in *space*, *frequency* and *orientation*. Let $\psi_{j,k,\mathbf{l}}(\mathbf{x})$ denote the shearlet basis functions (or in the remainder simply called shearlets), then the shearlet coefficients of an image $f(\mathbf{x}) \in L^2\left(\mathbb{R}^2\right)$ are given by:[19,20]

$$w_{j,k,\mathbf{l}} = \langle f, \psi_{j,k,\mathbf{l}} \rangle = \int_{\mathbb{R}^2} f(\mathbf{x}) \overline{\psi_{j,k,\mathbf{l}}(\mathbf{x})} \mathrm{d}\mathbf{x}, \tag{1}$$

where $j \in \mathbb{Z}$, $k \in \mathbb{Z}$ and $\mathbf{l} \in \mathbb{Z}^2$ denote the scale, orientation and the spatial location, respectively. The idea behind the shearlet transform is to combine geometry and multiscale analysis:[17] shearlets are formed by dilating, shearing and translating a mother shearlet function $\psi \in L^2\left(\mathbb{R}^2\right)$, as follows:

$$\psi_{j,k,\mathbf{l}}(\mathbf{x}) = |\det \mathbf{A}|^{j/2} \psi\left(\mathbf{B}^k \mathbf{A}^j \mathbf{x} - \mathbf{l}\right), \tag{2}$$

where $\mathbf{A}$ and $\mathbf{B}$ are invertible 2×2 matrices, with $\det \mathbf{B} = 1$ and where $|\det \mathbf{A}|^{j/2}$ is a normalization factor (such that $\left\|\psi_{j,k,\mathbf{l}}\right\| = \|\psi\|$). The shearlet functions are subject to a composite dilation $\mathbf{A}^j$ and geometrical transform $\mathbf{B}^k$. In this paper, we will use the following transform matrices:

$$\mathbf{A} = \begin{pmatrix} 4 & 0 \\ 0 & 2 \end{pmatrix} \quad \text{and} \quad \mathbf{B} = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}. \tag{3}$$

Here, $\mathbf{A}$ is an anisotropic scaling matrix (in the x-direction, the scaling is twice the scaling in the y-direction) and $\mathbf{B}$ is a geometric shear matrix. These transforms are illustrated in Figure 2.

The shearlet mother function is a composite wavelet, which is defined in the Fourier transform domain as:

$$\Psi(\omega) = \Psi_1(\omega_x) \Psi_2\left(\frac{\omega_y}{\omega_x}\right), \tag{4}$$

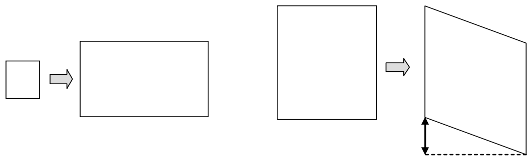

(a) (b)

Figure 2. Geometric transformations used by the shearlet transform (a) anisotropic dilation (matrix **A**). (b) shear (matrix **B**).

with $\omega = [\omega_x\ \omega_y]$, $\Psi_1(\omega_x)$ the Fourier transform of a wavelet function and $\Psi_2(\omega_y)$ a compactly supported bump function:

$$\Psi_2(\omega_y) \neq 0 \Leftrightarrow \omega_y \in [-1,1]. \tag{5}$$

By equation (5), the mother shearlet function is bandlimited in a hourglass-shaped region of the 2D frequency spectrum:

$$\Psi(\omega) \neq 0 \Leftrightarrow \left|\omega_y\right| < |\omega_x|. \tag{6}$$

Noting that the basis functions are obtained through shears and dilations of the mother shearlet function, this bandlimited property also directly controls the directional sensitivity of the basis functions: a shear operation on the mother shearlet function results in a shift in the argument of $\Psi_2(\omega_y/\omega_x)$:

$$\Psi\left(\left(\mathbf{B}^{-T}\right)^k \omega\right) = \Psi_1(\omega_x)\,\Psi_2\left(\frac{\omega_y}{\omega_x} - k\right), \tag{7}$$

and correspondingly (see Figure 1(b)): $\Psi\left(\left(\mathbf{B}^{-T}\right)^k \omega\right) \neq 0 \Leftrightarrow \omega_y \in [k\omega_x - |\omega_x|, k\omega_x + |\omega_x|]$. Hence the orientation of the basis function is controlled by the shear parameter k (see Figure 1(b)). Similarly, the anisotropic scaling leads to:

$$\Psi\left(\mathbf{A}^{-j}\omega\right) = \Psi_1\left(4^{-j}\omega_x\right)\Psi_2\left(2^j\frac{\omega_y}{\omega_x}\right). \tag{8}$$

We see that changing the scale parameter j results in a scaling in the argument of the wavelet Ψ_1, but it also affects the support of $\Psi\left(\mathbf{A}^{-j}\omega\right)$: $\Psi\left(\mathbf{A}^{-j}\omega\right) \neq 0 \Leftrightarrow 2^j\left|\omega_y\right| < |\omega_x|$. More concretely, when the scale parameter is increased by 1 (corresponding to a finer scale), the bandwidth of the shearlet is halved. If we further require that the set of shearlet functions cover the complete frequency spectrum, we can easily see that we will require *twice as many* shearlet functions $\psi_{j,k,\mathbf{l}}(\mathbf{x})$. Consequently, the number of analysis orientations *doubles* at every finer scale. Let us now consider composite dilation and shearing:

$$\Psi\left(\left(\mathbf{B}^{-T}\right)^k \mathbf{A}^{-j}\omega\right) \neq 0 \Leftrightarrow \omega_y \in \left[2^{-j}\left(k\omega_x - |\omega_x|\right), 2^{-j}\left(k\omega_x + |\omega_x|\right)\right]. \tag{9}$$

The last part of the equation corresponds to a wedge-shaped region in frequency space. Consequently, by changing the shear and scale parameters k and j, arbitrary wedges of the frequency plane can be selected.

So far, we considered vertical shearing and anisotropic dilation, with a larger scaling factor in the x-direction than in the y-direction. To obtain a more equal treatment of the horizontal and vertical directions, the frequency plane is usually split into two cones (for the high frequency band) and a square at the origin (for the low frequency band), as shown in Figure 1(c):

$$\begin{aligned}
C_1 &= \left\{(\omega_x,\omega_y) \in \mathbb{R}^2 \,|\, |\omega_x| \geq \omega_0, \left|\omega_y\right| \leq |\omega_x|\right\}, \\
C_2 &= \left\{(\omega_x,\omega_y) \in \mathbb{R}^2 \,|\, \left|\omega_y\right| \geq \omega_0, \left|\omega_y\right| > |\omega_x|\right\}, \\
C_3 &= \left\{(\omega_x,\omega_y) \in \mathbb{R}^2 \,|\, |\omega_x| < \omega_0, \left|\omega_y\right| < \omega_0\right\}.
\end{aligned}$$

with ω_0 the maximal frequency of the center square C_3. To treat horizontal and vertical frequencies equally, in cone C_2, the x- and y-components for $\mathbf{x}$ need to be switched before applying geometric transforms. This comes down to the following dilation and shear matrices in both cones:

$$\mathbf{A}_1 = \begin{pmatrix} 4 & 0 \\ 0 & 2 \end{pmatrix}, \mathbf{B}_1 = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \quad (\text{cone } C_1) \quad \text{and} \quad \mathbf{A}_2 = \begin{pmatrix} 2 & 0 \\ 0 & 4 \end{pmatrix}, \mathbf{B}_2 = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \quad (\text{cone } C_2). \tag{10}$$

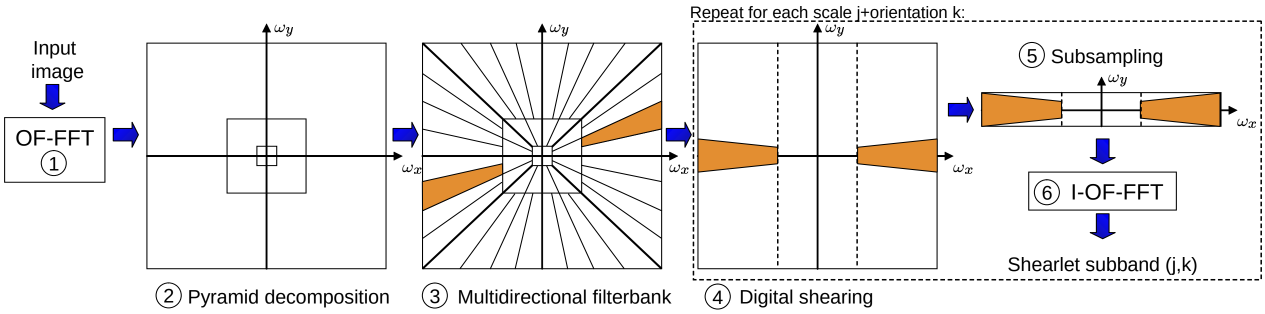

Figure 3. Overview of the analysis algorithm. OF-FFT is the odd-frequency DFT implemented using fast Fourier transforms (FFTs).

Consequently, the horizontal cone is dilated horizontally by factor 4 per scale, while the vertical cone is dilated vertically by factor 4. In the following, we make the distinction between both cones explicit by assigning different shearlet basis functions to each cone $d = 1,2$:

$$\psi_{j,k,\mathbf{l}}^{(d)}(\mathbf{x}) = |\det \mathbf{A}_d|^{j/2}\, \psi^{(d)}\left(\mathbf{B}_d^k \mathbf{A}_d^j \mathbf{x} - \mathbf{l}\right), \tag{11}$$

where $\psi^{(1)}(x,y) = \psi(x,y)$ and $\psi^{(2)}(x,y) = \psi(y,x)$. The resulting frequency tiling is illustrated in Figure 1(a).

3. THE NOVEL ANALYSIS AND SYNTHESIS ALGORITHM

3.1 Overview

Because the shearlet basis functions are bandlimited and directly defined in frequency domain (equation (4)), we will formulate both the analysis and synthesis algorithm in the DFT domain.* This implementation was first proposed in ref,[18] however in this paper we will present a number of refinements. Our discrete implementation is also different from other proposed implementations[17,20] in the sense that it is specifically designed as tight frame (which is important when performing multiscale thresholding) in a way that the transform is self-invertible, *without* requiring an iterative analysis/synthesis scheme. Our implementation also differs in the sense that it decimates not only between scales, but also between orientations, resulting in the possibility for a very low redundancy. The analysis algorithm consist of the following steps (see Figure 3):

1. Compute the odd-frequency discrete Fourier transform (OF-DFT) of the input image.
2. Analyze the OF-DFT of the input image using a Laplacian pyramid-like filter bank (with subsampling).
3. Apply a directional filter bank to the resulting subbands.
4. (*Optionally*) shear the subbands such that the spectral content lies in a central rectangle in frequency domain.
5. Perform a one-dimensional subsampling to get rid of the remaining zero-DFT coefficients.
6. Compute the inverse OF-DFT of every resulting subband.

The synthesis algorithm will simply reverse each step of the analysis algorithm, starting from step 6 to step 1. We will now discuss every step somewhat more in detail.

*Implementation in spatial domain is also possible, this will be the topic of a future paper.

3.2 Multiscale and multidirectional filter bank

The transform is conceived as a cascade of linear filters in the DFT domain. The concept of analysis and synthesis filter bank is shown in Figure 4. First, an isotropic wavelet filter $G_0(\omega_r)$ and complementary scaling filter $H_0(\omega_r)$ are used, where $\omega_r = \max(|\omega_x|, |\omega_y|)$. Here ω_x, ω_y denote 2-D continuous frequency coordinates. ω_r is the pseudo-radius in a pseudo-polar coordinate system (see ref[18]). For simplicity of the notations, we will stick to Discrete Time Fourier Transform (DTFT) definitions, transition to the equivalent formulas in the DFT domain can be obtained by substituting $\omega_y = 2\pi m/M$ and $\omega_x = 2\pi n/N$, where (m,n) are discrete frequency coordinates.

In essence, for this filter bank, any wavelet filter can be used, since the analysis scheme can be implemented as a Laplacian pyramid.[21] However, we advice to use orthogonal filters in order to have a tight frame.[22] On the other hand, the Laplacian pyramid can be made free of aliasing if the wavelet filters are bandlimited and if the decimation factors are adapted to the bandwidths of the filters. An example of filters obeying these conditions are Meyer wavelet filters:[23]

$$H_0(\omega_r) = \begin{cases} 1 & |\omega_r| < \frac{\pi}{4} \\ \cos\left(\frac{\pi}{2} v\left(\frac{4|\omega_r|}{\pi} - 1\right)\right) & \frac{\pi}{4} \le |\omega_r| \le \frac{\pi}{2} \\ 0 & \text{else} \end{cases} \quad \text{and} \quad G_0(\omega_r) = \begin{cases} 0 & |\omega_r| < \frac{\pi}{4} \\ \sin\left(\frac{\pi}{2} v\left(\frac{4|\omega_r|}{\pi} - 1\right)\right) & \frac{\pi}{4} \le |\omega_r| \le \frac{\pi}{2}, \\ 1 & \text{else} \end{cases} \tag{12}$$

with the corresponding synthesis filters equal to the analysis filters $\tilde{G}_0(\omega_r) = G_0(\omega_r)$, $\tilde{H}_0(\omega_r) = H_0(\omega_r)$ and with $v(x)$ defined by:

$$v(x) = \begin{cases} 3x^2 - 2x^3 & 0 \le x \le 1 \\ 0 & x < 0 \\ 1 & 1 < x \end{cases}$$

Filters for subsequent (coarser) scales are defined recursively based on the relation:

$$H_j(\omega_r) = H_{j-1}(\omega_r) H_0(4^j \omega_r) \quad \text{and} \quad G_j(\omega_r) = H_{j-1}(\omega_r) G_0(4^j \omega_r), \quad j = 1, 2, \ldots \tag{13}$$

Next, directional analysis is performed using a set of angular bump filters $R(\vartheta)$:

$$R(\vartheta) = \begin{cases} 0 & \vartheta < 0 \\ \sin\left(\frac{\pi}{2} v\left(\frac{\vartheta}{2\alpha}\right)\right) & 0 \le \vartheta \le 2\alpha \\ 1 & 2\alpha \le \vartheta < \pi \\ \cos\left(\frac{\pi}{2} v\left(\frac{\vartheta - \pi}{2\alpha}\right)\right) & \pi \le \vartheta \le \pi + 2\alpha \\ 0 & \text{else} \end{cases} \tag{14}$$

where $\alpha \in [0, \pi/2]$ is a parameter that controls the transition bandwidth of the shearlet filters in the angular direction, or more precisely, the overlap between the frequency supports of two neighboring directional shearlet filters. For $\alpha = 0$ we obtain "ideal" angular bump filters, which usually suffer from ringing artifacts (which is a limiting factor in many practical applications). The ringing can be reduced by making α somewhat larger. For $\alpha = \pi/2$, every shearlet filter shares half of its frequency support with a "neighboring" shearlet filter.

Based on these definitions, the resulting filters for direction $k = 1, \ldots, K_j$ and scale j are given by:

$$G_{j,k}^{\text{un}}(\omega_x, \omega_y) = \begin{cases} G_j(\omega_r) R\left(\alpha + k\pi - \left(1 + \frac{\omega_y}{\omega_x}\right) \frac{\pi K_j}{4}\right) & k = 1, \ldots, K_j/2 \\ G_j(\omega_r) R\left(\alpha + k\pi - \left(1 + \frac{\omega_x}{\omega_y}\right) \frac{\pi K_j}{4}\right) & k = K_j/2 + 1, \ldots, K_j \end{cases} \tag{15}$$

where the ranges $k = 1, \ldots, K_j/2$ and $k = K_j/2 + 1, \ldots, K_j$ respectively correspond to cones C_1 and C_2 in Figure 1(c). The superscript 'un' denotes the fact that the filters are unnormalized: the condition $\sum_{k=1}^{K_j} \left|G_{j,k}^{\text{un}}(\omega_r)\right|^2 = \left|G_j(\omega_r)\right|^2$ is not satisfied near the bisectors of the frequency plane (i.e. $\omega_x = \pm\omega_y$). Therefore, we normalize the filters $G_{j,k}^{\text{un}}(\omega_x, \omega_y)$ as follows:

$$G_{j,k}(\omega_x, \omega_y) = G_k^{\text{un}}(\omega_x, \omega_y) \left|G_j(\omega_r)\right| / \sqrt{\sum_{k=1}^{K_j} \left|G_{j,k}^{\text{un}}(\omega_r)\right|^2}. \tag{16}$$

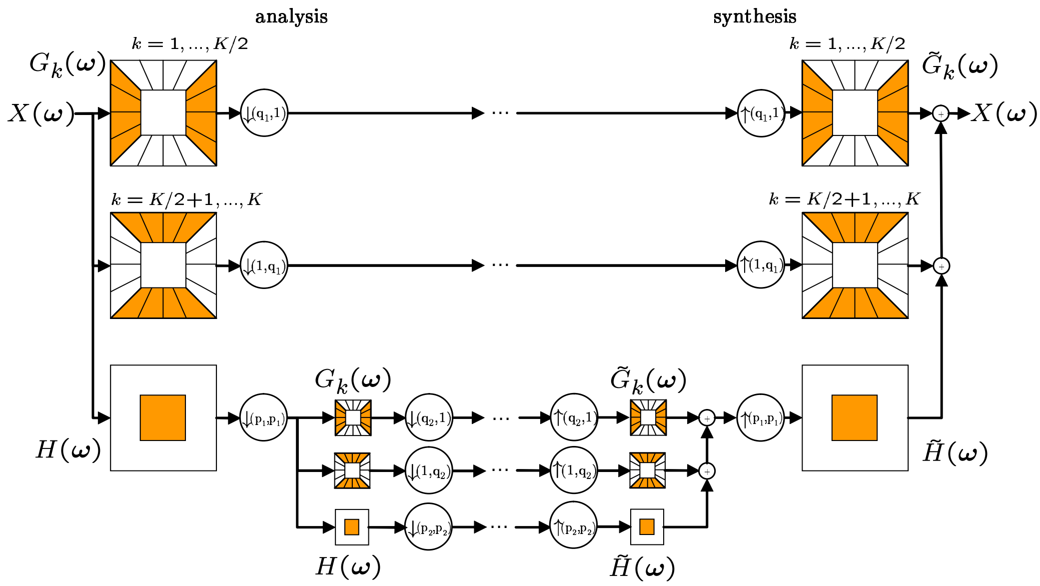

Figure 4. Shearlet analysis and synthesis filter bank.

Finally, note how the parameter K_j controls the number of directions in each scale. To conform to the continuous shearlet transform (Section 2), we choose $K_j = K_0 \cdot 2^{-j}$ ($j = 0,1,2,\ldots$), such that the number of orientations doubles at every finer scale.

3.3 Subsampling in frequency domain

Subsampling in frequency domain (or equivalently in spatial domain) - without special care - may destroy the perfect reconstruction (PR) and self-invertibility properties of the transform as we will show next. In the following we will discuss the subsampling of the *shearlet* coefficients in the DTFT domain. By changing the filters, the equations also apply to the scaling coefficients.

Recall that subsampling a real-valued filtered image with DTFT $X(\omega_x,\omega_y)G_{j,k}(\omega_x,\omega_y)$ by an integer factor D gives:

$$Y_{j,k}(\omega_x,\omega_y) = \frac{1}{\sqrt{D_x D_y}} \sum_{m=0}^{D_y-1} \sum_{n=0}^{D_x-1} X\left(\frac{2\pi m}{D_x}+\omega_x, \frac{2\pi n}{D_y}+\omega_y\right) G_{j,k}\left(\frac{2\pi m}{D_x}+\omega_x, \frac{2\pi n}{D_y}+\omega_y\right). \quad (17)$$

This formula basically expresses the creation of aliasing copies of the spectrum $X(\omega_x,\omega_y)$: the larger the subsampling factor, the more aliasing copies that are being created. Because of the real-valuedness of the input image, the following conjugate symmetry relationship holds:

$$X(\omega_x,\omega_y) = \overline{X(-\omega_x,-\omega_y)} \quad (18)$$

Things become slightly more complicated when considering the discretization of the frequency space, which is needed when computing the DFT. Let $\omega_y = 2\pi m/M$ and $\omega_x = 2\pi n/N$, then the conjugate symmetry property becomes:

$$X'(n,m) = X\left(\frac{2\pi n}{N}, \frac{2\pi m}{M}\right) = \overline{X'(N-n, M-m)}, \quad m = 1,\ldots,M-1, n = 1,\ldots,N-1. \quad (19)$$

Note that for the Nyquist frequency bins (corresponding to $m = M/2$ and/or $n = N/2$), the conjugate symmetry property takes the following form:

$$\Im\left[X'\left(\frac{N}{2},m\right)\right] = \Im\left[X'\left(n,\frac{M}{2}\right)\right] = 0. \quad (20)$$

$$\Re\left[X'\left(\frac{N}{2},m\right)\right] = \Re\left[X'\left(\frac{N}{2},M-m\right)\right] \quad \text{and} \quad \Re\left[X'\left(n,\frac{M}{2}\right)\right] = \Re\left[X'\left(N-n,\frac{M}{2}\right)\right] \quad (21)$$

where $\Re[\cdot]$ and $\Im[\cdot]$ are respectively the real and imaginary part of a complex number. One of the main issues is that the conjugate symmetry property of $Y(\omega_x, \omega_y)$ (21) is destroyed by applying the filter $G_k(\omega_x, \omega_y)$, because $G_k(\omega_x, \omega_y)$ does not satisfy (21). Consequently, by upsampling in the backward transform, the original image can not be recovered anymore! Fortunately, this issue only affects the Nyquist frequency coefficients, hence the influence on the reconstruction error is fairly limited. Nevertheless, we considered this problem as significant for a number of applications (e.g. compressed sensing), therefore we present three solutions:

- A first solution is to enforce filters satisfying (21) by partially sacrificing direction selectivity. This can be achieved by choosing $G'_{j,k}\left(\frac{N}{2}, m\right) = G'_{j,k}\left(n, \frac{M}{2}\right) = 1/\sqrt{K_j}$. This approach was taken in the implementation of the steerable pyramid transform,[1,24] however, because the bandlimitedness of the shearlet filters is lost, this solution makes it impossible to apply the angular subsampling in step 5 (otherwise PR would be lost).
- A second solution is to disregard the Nyquist frequency coefficients and to choose $G'_{j,k}\left(\frac{N}{2}, m\right) = G'_{j,k}\left(n, \frac{M}{2}\right) = 0$ after extracting and storing the Nyquist frequency coefficients separately. This results in a very small (but overall, negligible) amount of extra redundancy of the (final) shearlet transform coefficients. Because it is not so clear how the corresponding processing of these Nyquist coefficients should be done in practical applications (because the corresponding basis functions are not localized), this solution is considered to be rather dirty and tricky.
- Consider again the conjugate symmetry condition (19). If we write these conditions for all m and n as a linear system, we see that the problem arises when $n = N - n$ or $m = M - m$, resulting in a different equation (i.e. (21)) than for other combinations of (m, n). The main idea is now to shift the frequency grids such that all the equations have the same number of unknowns. Therefore, we use the following discretization of the frequency coordinates: $\omega_y = 2\pi(m + 0.5)/M$ and $\omega_x = 2\pi(n + 0.5)/N$. This corresponds to using an even-time odd-frequency OF-DFT, which is defined as follows:

$$\begin{aligned} X'(n,m) &= \frac{1}{\sqrt{MN}} \sum_{m'=0}^{M-1} \sum_{n'=0}^{N-1} x(n',m') \exp\left(-\frac{2\pi i n'}{N}(n+0.5) - \frac{2\pi i m'}{M}(m+0.5)\right) \\ &= \frac{1}{\sqrt{MN}} \sum_{m'=0}^{M-1} \sum_{n'=0}^{N-1} \left[\exp\left(-\frac{\pi i n'}{N} - \frac{\pi i m'}{M}\right) x(n',m')\right] \exp\left(-\frac{\pi i n' n}{N} - \frac{\pi i m' m}{M}\right), \end{aligned}$$

  which is nothing more than the DFT of the pre-modulated image $\exp\left(-\frac{\pi i n'}{N} - \frac{\pi i m'}{M}\right) x(n',m')$. Here, i is the imaginary unit. For the OF-DFT, the conjugate symmetry property becomes much simpler:

$$X'(n,m) = \overline{X'(N-1-n, M-1-m)}, \quad m = 0, \ldots, M-1, n = 0, \ldots, N-1. \tag{22}$$

  More importantly, there is no integer (n,m) for which $(n,m) = (N-1-n, M-1-m)$, hence conjugate symmetry does not impose real-valuedness on any of the OF-DFT coefficients. Consequently, it suffices to evaluate $G'_{j,k}(\omega_x, \omega_y)$ with the shifted frequency coordinates in order to ensure PR!

For this reason, we stick to the OF-DFT representation of our input image and the remaining operations (shearing) will be implemented in this domain as well.

3.4 Digital shearing in the OF-DFT domain

The next step in the transform is to shear the filtered subbands, such that the spectral content of each subband is contained in one central rectangle, as illustrated in step 5 of Figure 3. Because of the bandlimitedness of the shearlet filters, most of the OF-DFT coefficients will be zero and can be disregarded (by applying a proper decimation afterwards, see Subsection 3.5), without loss of information. The digital shearing step is optional and can be skipped if aliasing is not a major problem. First, we will explain how the digital shearing is being performed. Therefore, note that this operation can easily be expressed in the DTFT domain (omitting the scaling matrix $\mathbf{A}$ for simplicity):

$$Z_{j,k}(\omega) = Y_{j,k}\left(\left(\mathbf{B}^{-T}\right)^{s_{j,k}} \omega\right), \tag{23}$$

where $s_{j,k}$ is the shearing factor for the specific orientation k and scale j. For our frequency partitioning, we use the shearing factors $s_{j,k} = s_{j,k+K_j/2} = [-(K_j-2)+4(k-1)]/K_j$, with $k = 1,...,K_j/2$. Let us now consider shearing in cone C_1 (see Figure 1(c)). This shear matrix is $\left(\mathbf{B}_1^{-T}\right)^{s_{j,k}} = \begin{pmatrix} 1 & 0 \\ -s_{j,k} & 1 \end{pmatrix}$, such that $Z_{j,k}(\omega_x,\omega_y) = Y_{j,k}(\omega_x, \omega_y - s_{j,k}\omega_x)$. Applying the OF-DFT parametrization of frequency space gives:

$$Z'_{j,k}(n,m) = Y_{j,k}\left(2\pi\frac{n+0.5}{N}, 2\pi\frac{m+0.5}{M} - s_{j,k}2\pi\frac{n+0.5}{N}\right) = Y_{j,k}\left(2\pi\frac{n+0.5}{N}, 2\pi\frac{m+0.5-Ms_{j,k}(n+0.5)/N}{M}\right). \quad (24)$$

Because $Ms_{j,k}(n+0.5)/N$ is generally not an integer, $Z'_{j,k}(n,m)$ can not directly be related to the OF-DFT $Y'_{j,k}(n,m) = Y_{j,k}(2\pi(n+0.5)/N, 2\pi(m+0.5)/M)$, hence a fractional delay is required. A simple workaround would be to round the fractional delay $Mk(n+0.5)/N$ to the nearest integer, however we found that this deteriorates the spatial localization of the shearlet basis functions. Instead, several methods are available for performing a fractional delay.[25,26] In this paper, we will use bandlimited interpolation (also used in, e.g., ref[27]) because the bandlimited interpolation filter has a perfect linear phase response. The formula can be given in terms of the Dirichlet kernel $\rho_{m,m'}(n)$:

$$Z'_{j,k}(n,m) = \sum_{m'=0}^{M-1} Y'_{j,k}(n,m')\rho_{m,m'}(n) = \sum_{m'=0}^{M-1} Y'_{j,k}(n,m')\frac{\sin\left(\pi\left(m\text{-}m'\text{-}Ms_{j,k}(n+0.5)/N\right)\right)}{M\sin\left(\pi\left((m\text{-}m')/M\text{-}s_{j,k}(n+0.5)/N\right)\right)}, \quad (25)$$

and the practical implementation can be done using a DFT along the ω_y-axis (or similarly along the ω_x-axis, for the cone C_2). Let us denote by $Z''_{j,k}(n,l)$ and $Y''_{j,k}(n,l)$ the DFT transforms of respectively $Y'_{j,k}(n,m)$ and $Z'_{j,k}(n,m)$ along the second dimension (ω_y), then the interpolation simply amounts to a modulation of the phase of the complex coefficients:

$$Z''_{j,k}(n,l) = Y''_{j,k}(n,l)\exp\left(\frac{2\pi i}{N}s_{j,k}(n+0.5)r_l\right) \quad \text{with} \quad r_l = \begin{cases} l & 0 \le l < \frac{M}{2} \\ M-1-l & \frac{M}{2} \le l < M \end{cases} \quad (26)$$

Afterwards, an inverse DFT along the same dimension is used to obtain $Z'_{j,k}(n,m)$ again.

3.5 Fractional one-dimensional subsampling

The angular filtering step in the filter bank from Subsection 3.2 actually *increases* the redundancy factor of the transform by a factor K_j per scale for square images. Luckily, by the bandlimitedness of the filters $G_{j,k}(\omega_x,\omega_y)$ the redundancy factor can be made approximately† independent of K_j. This allows analyzing images with a large number of directional shearlet filters without demanding extra redundancy. Note that by definition, the filters $G_{j,k}(2\pi(n+0.5)/N, 2\pi(m+0.5)/M)$ are zero outside their frequency support. The maximal size of the frequency support in the vertical direction is $q = 2\lceil M(1+\alpha/\pi)/K\rceil$.‡ In the next step, we will subsample the resulting subbands by a factor $M/q = M/\left(2\lceil M(1+\alpha/\pi)/K\rceil\right)$ in the vertical direction. Remark that this subsampling factor is not necessarily *integer*! Hence, an adaptation of (17) is needed. The fractional subsampling operation can be defined in the OF-DFT domain as follows:

$$U'_{j,k}(n,m) = \sqrt{\frac{q}{M}}\sum_{m''=-\infty}^{+\infty}\tilde{Z}''_{j,k}(n,m+m''q) \quad \text{with} \quad \tilde{Z}''_{j,k}(n,m) = \begin{cases} Z''_{j,k}(n,m) & 0 \le m < M \\ 0 & \text{else} \end{cases}, \quad (27)$$

which is equivalent to (17) for integer M/q. Here, the factor $\sqrt{\frac{q}{M}}$ is an energy normalization constant, this constant is required in order to have a tight frame.

It is easy to check that $U'_{j,k}(n,m)$ satisfies the conjugate symmetry property (22):

$$U'_{j,k}(N-1-n, M-1-m) = \sqrt{\frac{q}{M}}\sum_{m''=-\infty}^{+\infty}\tilde{Z}''_{j,k}(N-1-n, M-1-m-m''q) = \sqrt{\frac{q}{M}}\sum_{m''=-\infty}^{+\infty}\overline{\tilde{Z}''_{j,k}(n,m+m''q)} = \overline{U'_{j,k}(n,m)}.$$

†in the sense that a close upper bound for the redundancy factor can be written that is independent of K_j.

‡Here we rounded q upward to be a multiple of two in order to have even DFT dimensions.

Finally, we need to check if PR can be achieved with this subsampling and digital shearing scheme. Therefore, we express that the shearing and subsampling does not affect the PR (which already held for the filter bank in Subsection 3.2). To undo the subsampling of (27), it is sufficient to divide by the normalization constant $\sqrt{q/M}$ and to multiply the obtained frequency response with the synthesis filter $\overline{G'_{j,k}(n,m)}$. Hence, for obtaining PR, this result should be equivalent to applying both the analysis filter $G'_{j,k}(n,m)$ and synthesis filter $\overline{G'_{j,k}(n,m)}$ to the input image. This directly leads to the following equation:

$$\sqrt{\frac{M}{q}}\overline{G'_{j,k}(n,m)}U'_{j,k}(n,m) = \left|G'_{j,k}(n,m)\right|^2 X'(n,m). \tag{28}$$

By substituting (27) and (25), the left hand side of (28) becomes:

$$\begin{aligned}\sqrt{\frac{M}{q}}\overline{G'_{j,k}(n,m)}U'_{j,k}(n,m) &= \overline{G'_{j,k}(n,m)}\sum_{m''=-\infty}^{+\infty}\sum_{m'=0}^{M-1} G'_{j,k}(n,m')X'(n,m')\rho_{m+m''q,m'}(n) \\ &= \left|G'_{j,k}(n,m)\right|^2 X'(n,m) + \sum_{m''=-\infty}^{+\infty}\sum_{\substack{m'=0\\ m-m'+qm''\neq 0}}^{M-1}\left[\overline{G'_{j,k}(n,m)}G'_{j,k}(n,m')\rho_{m+m''q,m'}(n)\right]X'(n,m').\end{aligned} \tag{29}$$

Because of the bandlimitedness of the filters, we have that $\overline{G'_{j,k}(n,m)}G'_{j,k}(n,m') = 0$ if $|m-m'| > q/2$. Consequently, the second term in (29) becomes zero if $\rho_{m,m_0}(n) \neq 0 \Leftrightarrow -q/2 \leq m < q/2$ with m_0 fixed. Hence PR imposes a *finite support* to $\rho_{m,m'}(n)$ (in terms of m and m'). We will now consider the following scenarios:

- *No digital shearing* (Subsection 3.4) is used. This is formally equivalent to using the Dirac-kernel $\rho_{m,m'}(n) = \delta(m-m')$. Since this kernel has a finite support, PR is guaranteed.
- Digital shearing with *rounding-to-nearest integer* of the shifts. This can be expressed using the kernel $\rho_{m,m'}(n) = \delta(m-m'-\text{round}(Ms_{j,k}(n+0.5)/N))$. Again, PR is guaranteed. As mentioned before, this is not a good option since the spatial localization of the basis functions is partially destroyed.
- Digital shearing with *fractional shifts*. The Dirichlet kernel $\rho_{m,m'}(n) = \frac{\sin\left(\pi\left(m-m'-Ms_{j,k}(n+0.5)/N\right)\right)}{M\sin\left(\pi\left((m-m')/M-s_{j,k}(n+0.5)/N\right)\right)}$ has a support of length $M > q$. However, this means that PR is not possible! To work around this problem, we use a trick: we split the fractional shifts into an integer part and a fractional part:

$$Ms_{j,k}(n+0.5)/N = \lfloor Ms_{j,k}(n+0.5)/N\rfloor + \text{frac}\left(Ms_{j,k}(n+0.5)/N\right). \tag{30}$$

  As shown above, PR holds for shifting with the integer part. To have overall PR, we perform the fractional shift $\text{frac}\left(Ms_{j,k}(n+0.5)/N\right)$ using a *"modified"* Dirichlet kernel designed for a support of size q. This gives:

$$\rho_{m,m'}(n) = \sum_{m''=-\infty}^{+\infty}\delta(m\text{-}m'\text{-}\lfloor Ms_{j,k}(n+0.5)/N\rfloor)\frac{\sin\left(\pi\left(m''\text{-}m\text{-frac}\left(Ms_{j,k}(n+0.5)/N\right)q/M\right)\right)}{q\sin\left(\pi\left(m''\text{-}m\text{-frac}\left(Ms_{j,k}(n+0.5)/N\right)q/M\right)/q\right)}I\left(\left|m''\text{-}m\right| \leq \frac{q}{2}\right) \tag{31}$$

  where $I(x)$ is the indicator function. As $q \to M$, this modified Dirichlet kernel approaches the original Dirichlet kernel from (25), see Figure 5. The novel digital shearing operation can efficiently be implemented as follows: 1) shear the input DFT coefficients $Y'_{j,k}(n,m')$ using an integer shift (which is very fast). 2) The obtained DFT coefficients will be supported in a rectangle of size $N \times q$ in frequency domain (or of size $q \times M$ for cone C_2). Perform a DFT of length q along the second dimension (or first dimension for cone C_2). 3) Apply the phase modulation step from (26). 4) Perform the inverse DFT of length q along the second dimension.

So we have shown that PR is possible in all three of the above scenarios. The redundancy factor of the resulting transform can then easily be checked: for a fixed scale j, the redundancy is:

$$R_j = K_j\left(\lceil M(1+\alpha/\pi)/K_j\rceil/M + \lceil N(1+\alpha/\pi)/K_j\rceil/N\right) \approx 2\left(1+\alpha/\pi\right),$$

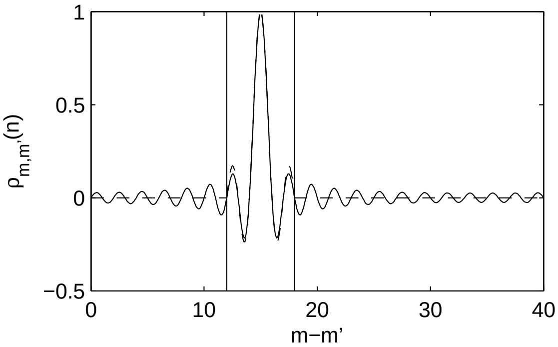

Figure 5. Illustration of the Dirichlet kernel (solid line) and its modified version according to (31) for $q = 6$ (dashed line). The larger the interval, the better the modified version will approximate the original kernel.

which is (approximately, due to the ceiling operation) independent of the number of orientations K_j, and proportional to the angular transition bandwidth of the shearlet filters α. The overall redundancy factor the transform is equal to the redundancy factor of a Laplacian pyramid with extra redundancy R_j applied to the highpass outputs of the pyramid:

$$R = R_1 + \frac{R_2}{4} + \frac{R_3}{16} + ... \approx \frac{8}{3}\left(1 + \frac{\alpha}{\pi}\right)$$

Since $0 \leq \alpha \leq \pi/2$, we find that the redundancy factor of transform is in the range $8/3 = 2.66... \leq R \leq 16/3 = 5.33...$.

3.6 Properties of the transform

It can be shown that the transform, implemented using the above algorithm, has the following properties:

- The transform is completely s*hift invariant* for displacements perpendicular to the shearing direction. This means that a shift in the input image in that particular direction results in the same shift in the subbands of the transform. On the other hand, shifts in other directions depend on the *shearing matrix* for the considered subband. In any case, shifts can be compensated if necessary. E.g., some applications exploit inter-dependencies both spatially and between multiresolution coefficients at different scales/orientations. By knowing how shifts in the input image affect the transform subbands, local analysis is greatly facilitated. For digital curvelet transforms,[12] this is not trivial because the translation lattices are not fixed. The shearlet transform offers an improvement here. If the input image is shifted by $(\Delta n, \Delta m)$, then the subband at scale j and orientation k is shifted by:

$$\begin{pmatrix} \Delta n' \\ \Delta m' \end{pmatrix} = \mathbf{A}^{-j} \begin{pmatrix} 1 & -s_{j,k} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \Delta n \\ \Delta m \end{pmatrix} = \mathbf{A}^{-j}\mathbf{B}^{-s_{j,k}} \begin{pmatrix} \Delta n \\ \Delta m \end{pmatrix}$$

  We remark that, in case of using the modified Dirichlet kernel, equation (31), strictly speaking, the shift-invariance property is lost. This can be alleviated by slightly increasing q (e.g. by 10%) to compensate the differences between the modified Dirichlet kernel and the true Dirichlet kernel.

- It can be shown that the basis elements constitute a *tight (Parseval) frame*. The shearlet basis is called a tight frame, if for all $f \in \ell^2(\mathbb{R}^2)$ the following equality holds:

$$\|f\|^2 = a_0 \left(\sum_{j,k,\mathbf{l}} \left| w_{j,k,\mathbf{l}} \right|^2 + \sum_{\mathbf{l}} |v_{\mathbf{l}}|^2 \right), \tag{32}$$

  where a_0 is a constant and where $v_{\mathbf{l}}$ is the scaling coefficient at position $\mathbf{l}$. The tight frame condition states that the image energy in spatial domain and in transform domain is equal (up to a known constant factor).

- The size of the spatial support of the shearlet basis functions obeys the *parabolic scaling law*. Because the shearlet filters are bandlimited, it is obvious that the basis functions are infinitely supported. Nevertheless, because of the relative fast decay of the functions, it is useful to put some estimate on the size of the functions (for example the size

(a) Original image | (b) noisy ($\sigma = 25$) PSNR=20.16dB | (c) wavelets PSNR=30.73dB RED=1 | (d) FSP PSNR=31.60dB RED=18.66 | (e) shearlets ($\alpha = \pi/2$) PSNR=31.50dB RED=5.66

Figure 6. BLS-GSM denoising results for several multiresolution transforms.

of the support that covers a certain percentage of the total energy of the basis function). To do so, we will consider a shearlet supported in the central rectangle in cone C_1 at scale j. The sizes of the supports of other shearlets can then be obtained by scaling, shearing and/or rotating the size of the support of this single basis function. First, we note that the height of the support of the function is proportional to 4^j, due to anisotropic scaling. On the other hand, we compute the vertical bandwidth of the shearlet function as the distance between the center transition frequencies. This vertical bandwidth is given by: $b_j = \frac{2^{-j}}{K_j}(1+\frac{\alpha}{\pi})$ (see (8) and (15)). The width of the support of the function is then inversely proportional to b_j:

$$\text{width} \propto 2^j K_0 \left(1+\alpha/\pi\right)^{-1} \quad \text{and} \quad \text{height} \propto 4^j$$

Consequently, the parabolic scaling law is found as follows: height $\approx$ width$^2 \left(1+\frac{\alpha}{\pi}\right)/K_0$. By this property, the basis functions are elongated, which is a useful property for representing object edges in images. The directional selectivity of the transform can be controlled by specifying K_0, the number of orientations for the finest scale of the transform. Furthermore, α, which controls the bandwidth of the shearlet filters, both influences the redundancy factor of the transform and the support size of the basis functions.

These properties are particularly useful for applications that make use of this transform. For example, in the context of image denoising, it has been found that shift-invariant transforms consistently yield better performance (e.g., in terms of image artifacts) than transforms that are not shift-invariant. Also, a Parseval relationship between the spatial domain and shearlet domain coefficients ensures that a “small” energy correction on the shearlet coefficients (e.g., due to shrinkage) results in a correspondingly small correction of pixel intensities. Finally, the freedom in choosing the number of orientations K_0 for the finest scale and the parameter α allows to trade-off the directional adaptivity properties of the transform versus the redundancy factor.

4. RESULTS

In Figure 6, we show denoising results for the removal of stationary white Gaussian noise from images, using three different multiresolution transforms: 1) the decimated DWT with 5 scales and using the Daubechies’ wavelet with two vanishing moments, 2) the full steerable pyramid transform (FSTP)[28] with 5 scales and 8 orientations and 3) the shearlet transform with 3 scales, 16 orientations for the finest scale ($K_0 = 16$) and $\alpha = \pi/2$. In each transform domain, we used the BLS-GSM estimator with the same parameters as in ref[28] (without inclusion of a parent coefficient in the local neighborhood vector), and finally we obtained the denoised image by applying the backward transform. It can be seen that the shearlet domain denoising method better reconstructs the edges and line-like structures in the images, leading to a very high visual quality with a relatively low redundancy factor (compared to the FSTP). On an Intel Core 2 Quad Q9550 processor at 2.83 GHz, the (single threaded) Matlab implementation of the forward and backward shearlet transform each take about 4.1 sec. for processing a 512×512 grayscale image. We also partially implemented our approach on a GPU (i.e., skipping steps 4-5 of Figure 3) using CUDA and the cuFFT library. Processing times for a NVidia GTX560 Ti GPU are, on average, 32 msec for the forward transform and 33 msec for the backward transform.

(a) Original image | (b) noisy ($\sigma = 50$) PSNR=14.14dB | (c) wavelets PSNR=24.79dB RED=1 | (d) FSP PSNR=25.62dB RED=18.66 | (e) shearlets ($\alpha = \pi/2$) PSNR=26.36dB RED=5.66

Figure 7. BLS-GSM denoising results for several multiresolution transforms.

5. CONCLUSION

In this paper, we presented a fast DFT-based analysis and synthesis scheme for the 2D discrete shearlet transform. This scheme implements the shearlet transform in such a way that it is consistent to the continuous shearlet theory and offers a number of mathematical properties, such as supporting shift invariance, being a Parseval frame and parabolic scaling of the basis functions. The transform has a low redundancy factor (2.6 to 5.2, fully controlled by the parameter α), independent of the number of analysis directions.

Acknowledgments

The authors would like to thank Prof. Ivan Selesnick for the nice discussions that contributed to this work.

REFERENCES

1. Simoncelli, E., Freeman, W. T., Adelson, E. H., and Heeger, D. J., "Shiftable Multi-scale Transforms," *IEEE Trans. Information Theory* **38**(2), 587–607 (1992).
2. Selesnick, I. W., Baraniuk, R. G., and Kingsbury, N. G., "The Dual-Tree Complex Wavelet Transform," *IEEE Signal Processing Magazine* **22**, 123–151 (nov 2005).
3. Van De Ville, D. and Unser, M., "Complex Wavelet Bases, Steerability, and the Marr-like pyramid," *IEEE Trans. Image Processing* **17**, 2063–2080 (nov 2008).
4. Lee, T., "Image Representation Using 2D Gabor Wavelets," *IEEE Trans. Pattern Analysis and Machine Intelligence* **18**, 1 (oct 1996).
5. Fischer, S., Šroubek, F., Perrinet, L., Redondo, R., and Cristobal, G., "Self-Invertible 2D Log-Gabor Wavelets," *International Journal of Computer Vision* **75**, 231–246 (nov 2007).
6. Do, M. N. and Vetterli, M., "The contourlet transform: An efficient directional multiresolution image representation," *IEEE Trans. Image Process.* **14**(12), 2091–2106 (2005).
7. Candès, E., *Ridgelets: Theory and Applications*, PhD thesis, Departement of Statistics, Stanford University (1998).
8. Do, M. N. and Vetterli, M., "The finite ridgelet transform for image representation," *IEEE Trans. Image Processing* **12**, 16–28 (jan 2003).
9. Romberg, J., Wakin, M., and Baraniuk, R., "Multiscale wedgelet image analysis: fast decompositions and modeling," in [*Proc. International Conference on Image Processing 2002*], **3**, 585–588 (24–28 June 2002).
10. LePennec, E. and Mallat, S., "Sparse geometric image representation with bandelets," *IEEE Trans. Image Processing* **14**, 423–438 (april 2005).
11. Coifman, R. R. and Meyer, F. G., "Brushlets: a tool for directional image analysis and image compression," *Appl. Comp. Harmonic Anal.* **5**, 147–187 (1997).
12. Candès, E., Demanet, L., Donoho, D., and Ying, L., "Fast Discrete Curvelet Transforms," *Multiscale modeling and simulation* **5**(3), 861–899 (2006).
13. Gopinath, R. A., "The Phaselet Transform - An Integral Redundancy Nearly Shift-Invariant Wavelet Transform," *IEEE Trans. Signal Processing* **51**, 1792–1805 (july 2003).

14. Velisavljevic, V., Beferull-Lozano, B., Vetterli, M., and Dragotti, P., "Directionlets: anisotropic multidirectional representation with separable filtering," *IEEE Trans. Image Processing* **15**, 1916–1933 (july 2006).
15. Lu, Y. M. and Do, M. N., "Multidimensional Directional Filter Banks and Surfacelets," *IEEE Trans. Image Processing* **16**, 918–931 (april 2007).
16. Guo, K. and Labate, D., "Optimally Sparse Multidimensional Representation using Shearlets," *SIAM J Math. Anal.* **39**, 298–318 (2007).
17. Easley, G., Labate, D., and Lim, W., "Sparse Directional Image Representations using the Discrete Shearlet Transform," *Applied and Computational Harmonic Analysis* **25**, 25–46 (2008).
18. Goossens, B., Aelterman, J., Luong, H. Q., Pižurica, A., and Philips, W., "Efficient Design of a Low Redundant Discrete Shearlet Transform," in [*2009 International Workshop on Local and Non-Local Approximation in Image Processing (LNLA2009)*], 112–124, IEEE, Tuusula, Finland (August 19-21 2009).
19. Guo, K., Labate, D., and Lim, W., "Edge Analysis and Identification using the Continuous Shearlet Transform," *Applied and Computational Harmonic Analysis* **27**, 24–46 (jul 2009).
20. Yi, S., Labate, D., Easley, G., and Krim, H., "A Shearlet Approach to Edge Analysis and Detection," *IEEE Trans. Image Processing* **18**, 929–941 (may 2009).
21. Burt, P. J. and Adelson, E. H., "The Laplacian pyramid as a compact image code," *IEEE Trans. Commun.* **COM-31**, 532–540 (april 1983).
22. Do, M. N. and Vetterli, M., "Frame reconstruction of the Laplacian pyramid," in [*IEEE Conf. on Acoustics, Speech and Signal Processing*], **6**, 3641–3644 (2001).
23. Daubechies, I., [*Ten Lectures on Wavelets*], Society for Industrial and Applied Mathematics, Philadelphia (1992).
24. Simoncelli, E. P. and Freeman, W. T., "The Steerable Pyramid: A flexible architecture for Multi-scale Derivative Computation," in [*Proc IEEE Int.Conf. Image Processing*], 444–447 (oct 1995).
25. Thiran, J. P., "Recursive digital filters with maximally flat group delay," *IEEE Trans. Circ. Theory* **18**, 659–664 (Nov. 1971).
26. Laakso, T. I., Välimäki, V., Karjalainen, M., and Laine, U. K., "Splitting the unit delay – tools for fractional delay filter design," *IEEE Sig. Process. Mag.* **13**, 30–60 (1996).
27. Condat, L., Van De Ville, D., and Forster-Heinlein, B., "Reversible, Fast, and High-Quality Grid Conversions," *IEEE Trans. Image Processing* **17**, 679–693 (may 2008).
28. Portilla, J., Strela, V., Wainwright, M., and Simoncelli, E., "Image denoising using scale mixtures of gaussians in the wavelet domain," *IEEE Transactions on image processing* **12**, 1338–1351 (November 2003).