

EFFICIENT DESIGN OF A LOW REDUNDANT DISCRETE SHEARLET TRANSFORM

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

Ghent University, Department of Telecommunications and Information Processing (TELIN-IPI-IBBT)
Sint-Pietersnieuwstraat 41, B-9000 Ghent, Belgium
bart.goossens@telin.ugent.be

ABSTRACT

Recently, there has been a huge interest in multiresolution representations that also perform a multidirectional analysis. The Shearlet transform provides both a multiresolution analysis (such as the wavelet transform), and at the same time an optimally sparse image-independent representation for images containing edges. Existing discrete implementations of the Shearlet transform have mainly focused on specific applications, such as edge detection or denoising, and were not designed with a low redundancy in mind (the redundancy factor is typically larger than the number of orientation subbands in the finest scale). In this paper, we present a novel design of a Discrete Shearlet Transform, that can have a redundancy factor of 2.6, independent of the number of orientation subbands, and that has many interesting properties, such as shift-invariance and self-invertability. This transform can be used in a wide range of applications. Experiments are provided to show the improved characteristics of the transform.

1. INTRODUCTION

It is well known that while traditional separable multidimensional wavelets are efficient for approximating images with point-wise singularities, they are not very effective for line singularities. For this reason, there has recently been a lot of interest in multiresolution representations that have basis elements with a much better directional selectivity, i.e. that also perform a multidirectional analysis. To name a few: steerable pyramids [1, 2], dual-tree complex wavelets [3–7], steerable complex wavelets [8], Marr-like wavelet pyramids [9], 2-D (log) Gabor transforms [10, 11], contourlets [12–14], ridgelets [15, 16], wedgelets [17, 18], bandelets [19], brushlets [20], curvelets [21, 22], phaselets [23], directionlets [24] and surfacelets [25].

One of the most recent siblings in this family of representations is the Shearlet transform [26–30], that provides the mathematical rigidness of a traditional multiresolution analysis (such as the wavelet transform), and that is at the same time an optimally sparse, image-independent representation for images containing edges. In this work, we further investigate the implementation of the Discrete

Shearlet Transform (DST), and we show that a design is possible with a redundancy factor as low as 2.6, while the number of orientation subbands can be arbitrarily high and while offering highly desirable properties such as shift invariance and self-invertability of the transform. This combination of properties is very difficult to achieve with existing x-let transforms. Further, we put special attention on design choices and we point out possibilities for further improvement. Results are given to demonstrate the improved characteristics of this transform.

The remainder of this paper is as follows: in Section 2 we give an overview the Continuous Shearlet Transform (CST) and its properties that are of importance to our work. In Section 3 we explain a number of existing DST implementations and their advantages or drawbacks. The novel DST implementation is presented in Section 4. Results and a discussion are given in Section 5. Finally, Section 6 concludes this paper.

2. THE CONTINUOUS SHEARLET TRANSFORM

2.1. Shearlet basis functions

The CST is a multiresolution 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 CST of an image $f(\mathbf{x}) \in L_2(\mathbb{R}^2)$ is defined by [29, 30]:

$$[S\mathcal{H}_\psi f](j, k, \mathbf{l}) = \int_{\mathbb{R}^2} f(\mathbf{x}) \psi_{j,k,\mathbf{l}}(\mathbf{l} - \mathbf{x}) d\mathbf{x} \quad (1)$$

where $j \in \mathbb{R}$, $k \in \mathbb{R}$ and $\mathbf{l} \in \mathbb{R}^2$ denote the scale, orientation and the spatial location, respectively. The idea behind the Continuous Shearlet Transform (CST) is to combine geometry and multiscale analysis [27]. Shearlets are formed by dilating, shearing and translating a mother shearlet function $\psi \in L_2(\mathbb{R}^2)$, as follows:

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

where \mathbf{A} and \mathbf{B} are invertible 2×2 matrices, with $\det \mathbf{B} = 1$. The normalization factor $|\det \mathbf{A}|^{j/2}$ has been chosen such that the norm $\|\psi\|_2 = \|\psi_{j,k,\mathbf{l}}\|_2$ for all j, k, \mathbf{l} . The basis functions are subject to a composite dilation \mathbf{A}^j and geometrical transform \mathbf{B}^k . For the shearlet analysis, we will use the following transform matrices:

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

B. Goossens is supported in part by Special Research Fund (BOF) of Ghent University, Belgium. A. Pižurica is a postdoctoral researcher of the Fund for the Scientific Research in Flanders (FWO) Belgium.




Figure 1. Geometric transformations used by the Shearlet transform (a) anisotropic dilation (matrix \mathbf{A}). (b) shear (matrix \mathbf{B}).

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 1.

The shearlet mother function is a composite wavelet that satisfies appropriate admissibility conditions [29], and that is defined in the Fourier transform domain as:

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

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

$$\Psi_2(\omega_y) = 0 \Leftrightarrow \omega_y \notin [-1, 1]. \quad (5)$$

Note that by this condition, the mother shearlet function is bandlimited in a diagonal band of the 2-D frequency spectrum. Because 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. To see this, let us investigate the effect of a shear operation on the mother shearlet function. For the shear transform in (3), we have:¹

$$\Psi(\mathbf{B}^k \boldsymbol{\omega}) = \Psi_1(\omega_x) \Psi_2\left(k - \frac{\omega_y}{\omega_x}\right) \quad (6)$$

which means that a shear operation results in a shift in the argument of $\Psi_2(\omega_y/\omega_x)$, hence the orientation of the basis function is controlled by the parameter k (see Figure 3b). Similarly, the anisotropic scaling leads to:

$$\Psi(\mathbf{A}^j \boldsymbol{\omega}) = \Psi_1(4^{-j} \omega_x) \Psi_2\left(2^{-j} \frac{\omega_y}{\omega_x}\right). \quad (7)$$

Here 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 the bump function Ψ_2 . More concretely, when the scale parameter is increased by 1, the bandwidth of the shearlet is halved (hence the shearlet has a finer directional selectivity).

2.2. Shearlets on the cone

So far, we considered shear operations in the vertical direction and anisotropic dilation, with a larger scaling factor in the x-direction than in the y-direction. To obtain a

¹Here, we rely on the fact that the Fourier transform of a geometrically transformed function $f(\mathbf{A}\mathbf{x})$ is given by $|\det \mathbf{A}|^{-1} \mathcal{F}\{f\}(\mathbf{A}^{-T}\boldsymbol{\omega})$, with $\mathcal{F}\{f\}$ the Fourier transform of f .




Figure 2. Partitioning of the 2-D frequency plane into two cones (C_1 and C_2) and a square (C_3) at the origin.

more equal treatment of the horizontal and vertical directions, the frequency plane is split into two cones (for the high frequency band) and a square at the origin (for the low frequency band), as shown in Figure 2 [26]:

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

with ω_0 the maximal frequency of the the center square C_3 . This square is added to be able to construct a shearlet tight frame [26, 27]. 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 using the following dilation and shear matrices in both cones:

$$\begin{aligned} \mathbf{A}_1 &= \begin{pmatrix} 4 & 0 \\ 0 & 2 \end{pmatrix}, \mathbf{B}_1 = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \\ \mathbf{A}_2 &= \begin{pmatrix} 2 & 0 \\ 0 & 4 \end{pmatrix}, \mathbf{B}_2 = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}. \end{aligned}$$

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 $d = 1, 2$:

$$\psi_{j,k,\mathbf{l}}^{(d)}(\mathbf{x}) = |\det \mathbf{A}_d|^{j/2} \psi(\mathbf{B}_d^k \mathbf{A}_d^j \mathbf{x} - \mathbf{l}) \quad (8)$$

Analogously to the wavelet transform [31], it is natural to discretize the scale, orientation and position indices. In the remainder, we will therefore restrict j, k, \mathbf{l} to discrete (integer) values. The resulting frequency tiling is illustrated in Figure 3a.

2.3. Tight frames of shearlets

Next, we want to represent an arbitrary function $f \in L_2(\mathbb{R}^2)$ by a set of projections of this function onto the shearlet basis elements, $\langle f, \psi_{j,k,\mathbf{l}}^{(d)} \rangle$. The family of functions

$$\left\{ \psi_{j,k,\mathbf{l}}^{(1)}(\mathbf{x}), \psi_{j,k,\mathbf{l}}^{(2)}(\mathbf{x}) \mid j \in \mathbb{Z}, k \in \mathbb{Z}, \mathbf{l} \in \mathbb{Z}^2, j \geq 0 \right\}$$




Figure 3. (a) Frequency tiling of the Shearlet transform in trapezoidal shaped tiles (wedges) [26]. (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 .

forms a tight frame (or Parseval frame) if there exist a positive constant $A > 0$ so that for all $f \in L_2(\mathbb{R}^2)$ the Parseval relationship holds:

$$\sum_{j,k,1,d} \left| \langle f, \psi_{j,k,1}^{(d)} \rangle \right|^2 = A \|f\|_2^2 \quad (9)$$

This relationship implies that any function in $L_2(\mathbb{R}^2)$ can be expanded into a set of functions [31]:

$$f = A^{-1} \sum_{j,k,1,d} \langle f, \psi_{j,k,1}^{(d)} \rangle \psi_{j,k,1}^{(d)} \quad (10)$$

For the shearlet functions, equation (9) amounts to a simple design constraint for the functions $\Psi_1(\omega)$ and $\Psi_2(\omega)$:

$$\sum_{j \geq 0} |\Psi_1(4^{-j}\omega)|^2 = 1 \quad \text{for } |\omega| \geq \omega_0 \quad (11)$$

$$\sum_{k=-2^j}^{2^j-1} |\Psi_2(2^j\omega - k)|^2 = 1 \quad \text{for } |\omega| \leq \pi \quad (12)$$

which practically means that the sum of the energies of $\Psi_1(\omega)$ and $\Psi_2(\omega)$ for respectively scaled frequencies and shifted frequencies must be one. The Parseval relationship holds for the part of the frequency plane that excludes the center square (see Figure 2), although this can be trivially extended to the complete frequency plane by adding a bandlimited scaling function [27]:

$$\Phi(\omega_x, \omega_y) = \begin{cases} \tilde{\Phi}(\omega_x) & |\omega_y| \leq |\omega_x| < \omega_0 \\ \tilde{\Phi}(\omega_y) & |\omega_x| \leq |\omega_y| < \omega_0 \\ 0 & \text{else} \end{cases} \quad (13)$$

with $\tilde{\Phi}(\omega)$ a 1-D scaling function that satisfies:

$$\sum_{j \geq 0} |\Psi_1(4^{-j}\omega)|^2 + |\tilde{\Phi}(\omega)|^2 = 1 \quad \text{for } |\omega| < \omega_0 \quad (14)$$

By comparing equation (2) to equation (13), it can be noted that the scaling function is more or less isotropic. This behavior resembles the isotropy of the scaling functions in the 2-D discrete wavelet transform (DWT), with the only difference that in the DWT, 2-D scaling functions are formed by a tensor products of one-dimensional scaling functions, instead of being defined per cone.

2.4. Shearlets or curvelets?

Shearlets are very similar to curvelets in the sense that both perform a multiscale and multidirectional analysis. Each basis element has a frequency support that is contained in a rectangle of size proportional to $2^j \times 4^j$ (or $4^j \times 2^j$) in both transforms, which means that the length of the frequency support is approximately the squared width of the frequency support. This property is called *parabolic scaling*, hence the frequency supports become increasingly thin as j decreases [22, 27]. Both transforms have very similar asymptotic approximation properties: for images $f(\mathbf{x})$ that are C^2 everywhere except near edges, where $f(\mathbf{x})$ is piecewise C^2 , the approximation error of a reconstruction with the N -largest coefficients ($f_N(\mathbf{x})$) in the shearlet/curvelet expansion is given by [22, 26]:

$$\|f - f_N\|_2^2 \leq B \cdot N^{-2} (\log N)^3, \quad N \rightarrow \infty$$

with B a constant. Because this is the optimal approximation rate for this type of functions [26], this property is often referred to as *optimal sparsity*. Still, there are a number of differences between shearlets and curvelets [27]:

- Shearlets are generated by applying a family of operators to a single function, while curvelet basis elements are not in the form of equation (2).
- Shearlets are normally associated to a fixed translation lattice, while curvelets are not. This is of importance for applications: when combining information from multiple scales and orientations (e.g.

to model inter- or intrascale dependencies), curvelet techniques need to take into account that the translation lattice is not fixed.

- In the construction of the shearlet tight frame above, the number of orientations doubles at every scale, while in the curvelet frame, this number doubles at every other scale.
- Shearlets are associated to a multiresolution analysis, while curvelets are not.

Perhaps the most primary advantage, that we want to point out in this work, is that shearlets allow for a much less redundant sparse tight frame representation, while offering shift invariance.

3. EXISTING DISCRETE SHEARLET TRANSFORMS

3.1. Direct discretization

The most straightforward way to design the DST is to apply a direct discretization to the shearlet filters $\psi_{j,k,1}^{(d)}(\mathbf{x})$. This approach is applied for example in [28]. First, the functions $\Psi_1(\omega)$, $\Psi_2(\omega)$ and $\tilde{\Phi}(\omega)$ are designed according to the constraints in equations (11)-(13) (see [27] for the design details). Next, the shearlet filters can be expressed in the discrete time Fourier transform domain as:

$$\begin{aligned}\Psi_{j,k,\mathbf{0}}^{(1)}(\omega) &= 2^{3j/2}\Psi_1(4^{-j}\omega_x)\Psi_2\left(2^{-j}\left(\frac{\omega_y}{\omega_x} - k\right)\right) \\ \Psi_{j,k,\mathbf{0}}^{(2)}(\omega) &= 2^{3j/2}\Psi_1(4^{-j}\omega_y)\Psi_2\left(2^{-j}\left(\frac{\omega_x}{\omega_y} - k\right)\right)\end{aligned}$$

A frequency domain implementation for an image of size $N \times M$ follows by sampling $\Psi_{j,k,1}^{(1)}(\omega)$ and $\Psi_{j,k,1}^{(2)}(\omega)$ in the points $\omega = \left(\frac{2\pi n}{N}, \frac{2\pi m}{M}\right)^T$, for $m = 0, \dots, M-1$ and $n = 0, \dots, N-1$. Let $B(m, n)$ denote the Discrete Fourier Transform (DFT) of the input image, then the DFTs of the shearlet coefficient subbands can be computed as:

$$W_{j,k}^{(d)}(m, n) = B(m, n)\Psi_{j,k,\mathbf{0}}^{(d)}\left(\frac{2\pi n}{N}, \frac{2\pi m}{M}\right),$$

with $j = 1, \dots, J$, $d = 1, 2$ and $k = -2^j, \dots, 2^j - 1$.

The DFT of the scaling coefficients is given by:

$$W^s(m, n) = B(m, n)\Phi\left(4^{-J}\frac{2\pi n}{N}, 4^{-J}\frac{2\pi m}{M}\right),$$

where J denotes the number of scales. The DFT of the reconstructed image $\tilde{B}(m, n)$ is obtained by applying the appropriate reconstruction filters to the shearlet subbands and by summing the results:

$$\begin{aligned}\tilde{B}(m, n) &= W^s(m, n)\Phi\left(4^{-J}\frac{2\pi n}{N}, 4^{-J}\frac{2\pi m}{M}\right) + \\ &\quad \sum_{j,k,d} W_{j,k}^{(d)}(m, n)\Psi_{j,k,\mathbf{0}}^{(d)}\left(\frac{2\pi n}{N}, \frac{2\pi m}{M}\right)\end{aligned}$$

where \bar{x} denotes the complex conjugate of x .

Even though the scheme is computationally simple and shift-invariant, the redundancy factor is high due to the lack of downsampling operations in the decomposition. More specifically, the redundancy factor is:

$$1 + \sum_{j=1}^J 2^{j+1} = 2^{J+2} - 3,$$

when choosing 2^{J+1} orientations for the first scale. For example, using 3 scales gives redundancy factor 29! In analogy to the undecimated DWT, we will call this transform the *undecimated DST* because of the lack of decimations.

3.2. Related existing implementations

Recently, a number of related DST implementations have been proposed. Easley et al. [27] propose a discrete implementation with one of the main applications in image denoising. In their work, a Laplacian pyramid is followed by windowing filters in the Pseudo-Polar DFT domain. By including decimations in the Laplacian pyramid, the redundancy of the transform is reduced. Because the redundancy factor per scale of the transform increases linearly with the number of orientations for that scale, the overall redundancy factor is still high. Finally, we remark that the Laplacian pyramid representation of [32] that is used in [27] is not a tight frame in its standard form, however a tight frame can be constructed by using orthogonal pyramid filters [33].

Yi et al. [30] outline a different implementation for edge detection and analysis. In their implementation, there is an explicit distinction between horizontal and vertical shearlets. However, the authors do not take further steps to reduce the redundancy, as they choose to stay faithful to the CST in terms of edge analysis. Further, it is not clear which cascade algorithm would be the best to do the inverse transform of this scheme, as no reconstruction algorithm is proposed.

These implementations are designed for specific applications in mind and have a redundancy factor that is lower (or equal) than the redundancy factor of the undecimated DST. The following questions arise:

1. Can this DST be designed to use in a wide range of applications?
2. Is it possible to devise a DST that is maximally decimated, i.e. with redundancy factor as low as possible? And what can be said about the shift-variance of such schemes?

In the next subsection, we will explain how do design such a transform.

4. NEW DESIGN OF THE DISCRETE SHEARLET TRANSFORM

As we explained in Section 2.2, for the CST there is an explicit separation of the horizontal cone C_1 and the vertical

cone C_2 . An obvious discrete realization would be to use hourglass-shaped filters. We prefer not to do this, as this either increases the redundancy factor by 2, or causes angular aliasing when including decimations in the angular filtering.² The presence of angular aliasing is very cumbersome in practical applications as it severely degrades the directional selectivity of the basis functions. Instead, we apply only one angular filtering stage at each scale to directly split up all orientation subbands, which also has the advantage that the corresponding filterbank is conceptually more clean.

To proceed, we will define shearlet filters in pseudo-polar frequency coordinates (FC). Every shearlet filter will extract a wedge-shaped region of the 2-D frequency plane; these wedge filters can be easily described in pseudo-polar FC.

4.1. Pseudo-polar coordinate system

We use FC (ω_r, ϑ) in a pseudo-polar grid [34], that is consistent to a polar grid, in the sense that the pseudo-angle is in the range $\vartheta \in [-\pi, \pi]$. The corresponding conversion from Cartesian coordinates (ω_x, ω_y) to pseudo-polar FC (ω_r, ϑ) is given by:

$$\omega_r(\omega_x, \omega_y) = \sqrt{\frac{1 + \max(|\omega_x|^2, |\omega_y|^2)}{1 + \pi^{-2}}} \quad (15)$$

$$\vartheta(\omega_x, \omega_y) = \begin{cases} \frac{\pi}{4} \left(\frac{\omega_y}{\omega_x} \right) & |\omega_x| > |\omega_y| \text{ and } \omega_x \geq 0 \\ \frac{\pi}{4} \left(2 - \frac{\omega_x}{\omega_y} \right) & |\omega_x| < |\omega_y| \text{ and } \omega_y \geq 0 \\ \frac{\pi}{4} \left(4 + \frac{\omega_y}{\omega_x} \right) & |\omega_x| > |\omega_y| \geq 0 \text{ and } \omega_x < 0 \\ \frac{\pi}{4} \left(-4 + \frac{\omega_y}{\omega_x} \right) & |\omega_x| > |\omega_y| < 0 \text{ and } \omega_x < 0 \\ \frac{\pi}{4} \left(-2 - \frac{\omega_x}{\omega_y} \right) & |\omega_x| < |\omega_y| \text{ and } \omega_y < 0 \end{cases} \quad (16)$$

where we replace the fractions ω_x/ω_y and ω_y/ω_x by 0 whenever the denominator becomes 0. The denominator in (15) has been chosen such that $\omega_r(\pm\pi, \pm\pi) = \pi$. Important to note is that adding a constant to the pseudo-angle ϑ corresponds to a vertical shear transform if both $(\omega_x, \omega_y) \in C_1$ and the transformed point also belong to C_1 . Equivalently, if $(\omega_x, \omega_y) \in C_2$, adding a constant to ϑ corresponds to a horizontal shear transform if the transformed point also belongs to C_2 . Transiting from C_1 to C_2 (or vice versa) can be done using a cascade of a horizontal and vertical shear transform. The pseudopolar grid defined above is illustrated in Figure 4. It can be noted that contours of equal pseudo-radial frequencies ω_r define concentric squares around the origin, instead of circles as is the case with a polar grid.

4.2. Filter bank

The filter bank design that we propose is mostly related to the design of the Steerable Pyramid transform [2], in the

²Such an angular filterbank with decimation is used in e.g. the contourlet transform [13].




Figure 4. Pseudo-polar coordinate system.

sense that we use a frequency partitioning into $K_j \geq 2$ orientation subbands and a lowpass subband at each scale. In our scheme, the specific way of decimating the horizontal and vertical orientation subbands, is different, as well as the pseudo-polar grid for defining the filters and the filters being used. This will allow us to further subsample the orientation subbands.

At each scale of our DST, we use a $(K_j + 1)$ -band recursive decomposition into K_j (bandpass or highpass) orientation bands and a lowpass band. Let us denote the scaling analysis filters as $H(\omega)$ and the shearlet analysis filters as $G_k(\omega)$, with $k = 1, \dots, K_j$ the index of the orientation band. The scaling synthesis filters and shearlet synthesis filters are $\tilde{H}(\omega)$ and $\tilde{G}_k(\omega)$, respectively. The analysis and synthesis filter bank is shown in Figure 5. In our filter bank, the image is first filtered using the oriented shearlet filters $G_k(\omega)$, subsequently the result is decimated with a scale-dependent factor q_j , in the direction orthogonal to the main filter orientation (i.e. horizontal or vertical). Next, the filter bank is iterated on the decimated output of the scaling filters, where the decimation factor for the scaling step j is denoted by p_j . The synthesis filter bank is entirely analogous, hence when designing the filters appropriately, the filter bank can be made to be self-inverting.

To design the filters, we express the perfect reconstruction equations for Figure 5, and try to find filters that satisfy these equations. The first perfect reconstruction (PR) condition for this filter bank is given by:

$$H(\omega)\tilde{H}(\omega) + \sum_{k=1}^{K_j} G_k(\omega)\tilde{G}_k(\omega) = 1, \omega \in \Omega \quad (17)$$

with $\Omega = [-\pi, \pi] \times [-\pi, \pi]$. For a decimated transform, other PR conditions are needed, to state that the aliasing caused by the downsampling operations should cancel itself. We will call these conditions the aliasing canceling conditions. Note that for some combinations of (p, q) PR is not even possible (e.g. $(p, q) = (4, 2)$), in that case, the PR conditions are conflicting.

We investigate $p = 4$ and $q = 1$ with anisotropic di-

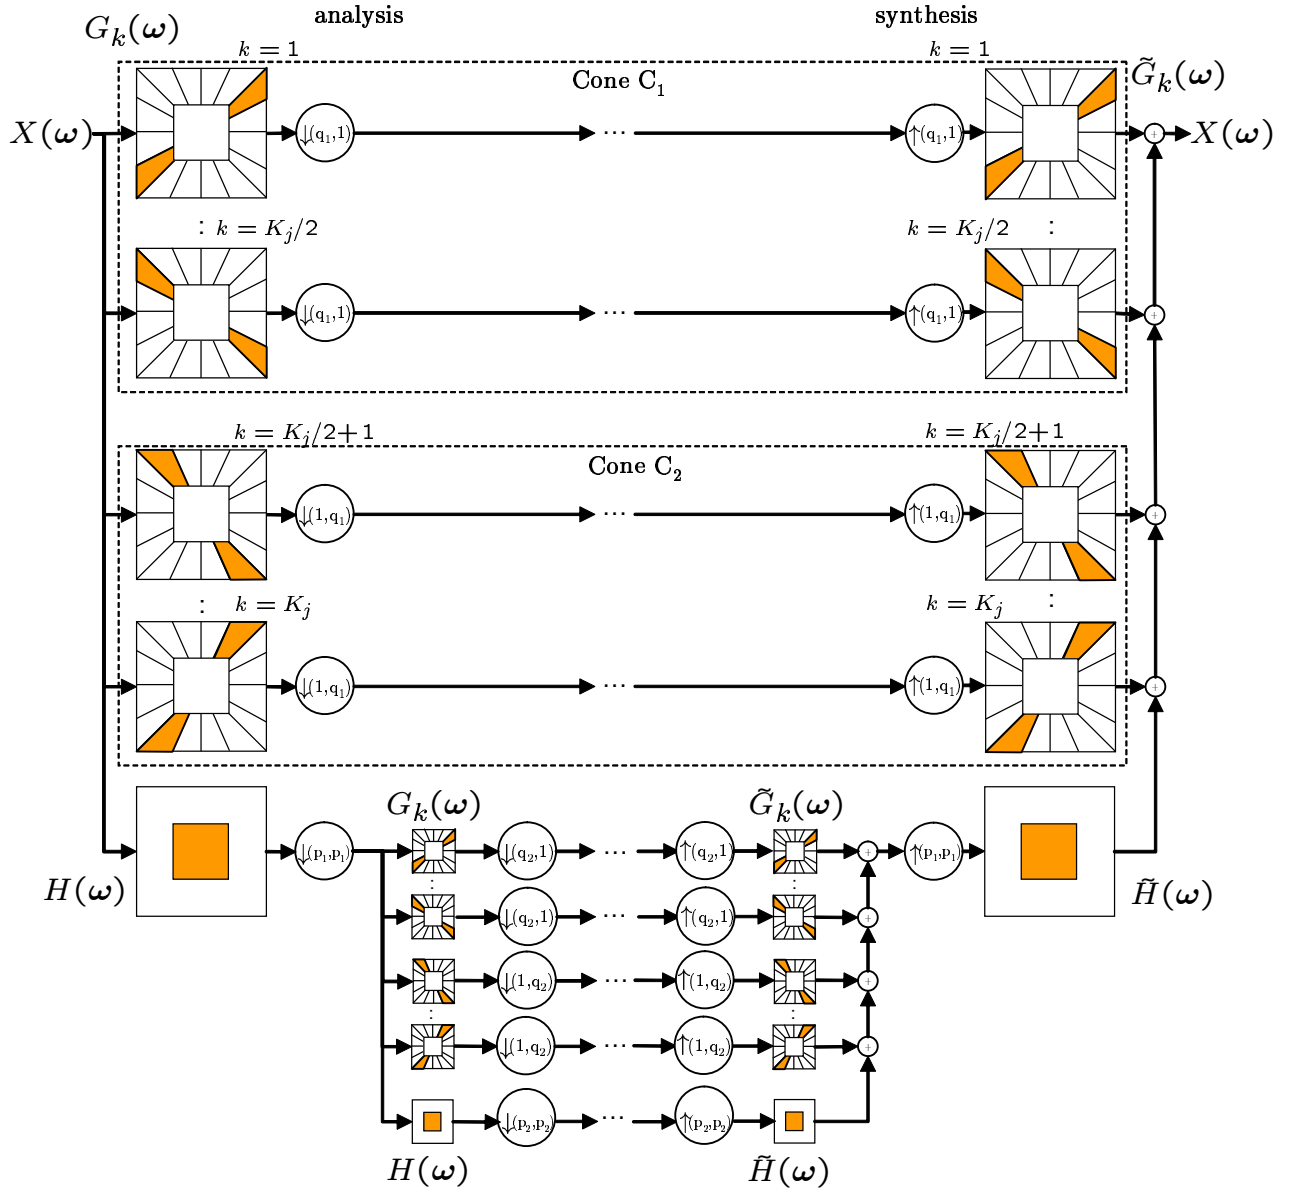


Figure 5. Shearlet analysis and synthesis filterbank.

lation matrices as in Section 2. The aliasing canceling PR conditions are:

$$H(\omega_x, \omega_y) \tilde{H}\left(\omega_x + \frac{m\pi}{2}, \omega_y + \frac{n\pi}{2}\right) = 0, \quad (18)$$

with $m = 0, \dots, 3$, $n = 0, \dots, 3$ and $(m, n) \neq 0$. Notably, equation (18) only affects the scaling filter and not the shearlet filters: the scaling filter $H(\omega)$ must have frequency support $[-\frac{\pi}{4}, \frac{\pi}{4}] \times [-\frac{\pi}{4}, \frac{\pi}{4}]$. Consequently, $H(\omega)$ cannot have compact support in spatial domain (see [31]). Nevertheless, because of the lack of aliasing, the transform can be made to be *shift-invariant*, in a similar way as done for the steerable pyramid transform [1, 35]. To do so, we define the filters in separable pseudo-polar FC (see

Section 4.1):

$$\begin{aligned} H(\omega_r, \vartheta) &= H_0(\omega_r), \\ G_k(\omega_r, \vartheta) &= G_0(\omega_r) \sum_{i=-\infty}^{+\infty} R\left(\frac{(\vartheta + i\pi)K_j}{\pi} - k + 1\right), \\ \tilde{H}(\omega_r, \vartheta) &= \tilde{H}_0(\omega_r), \\ \tilde{G}_k(\omega_r, \vartheta) &= \tilde{G}_0(\omega_r) \sum_{i=-\infty}^{\infty} \tilde{R}\left(\frac{(\vartheta + i\pi)K_j}{\pi} - k + 1\right) \end{aligned} \quad (19)$$

with $H_0(\omega_r)$ the frequency response of a 1-D scaling filter, $G_0(\omega_r)$ the frequency response of a 1-D wavelet filter and $R(\vartheta)$ a real-valued compactly supported bump function. In equation (19), the bump function is periodized in ϑ with period π to construct filters with real-valued

impulse responses. In practice we can assume that $\vartheta \in [-\pi, \pi]$, by the construction of the pseudo-polar grid (see Section 4.1). Consequently, the summation in (19) only needs to iterate over a finite number of values for i .

Using the above filters, the PR conditions come down to:

$$H_0(\omega_r)\tilde{H}_0(\omega_r) + G_0(\omega_r)\tilde{G}_0(\omega_r) = 1, \quad \omega_r \in [-\pi, \pi]$$

$$\sum_{k=1}^K \sum_{i=-\infty}^{+\infty} R\left(\frac{(\vartheta + i\pi)K_j}{2\pi} - k + 1\right) \tilde{R}\left(\frac{(\vartheta + i\pi)K_j}{2\pi} - k + 1\right) = 1, \quad \vartheta \in [-\pi, \pi] \quad (20)$$

$$H_0(\omega_r) = \tilde{H}_0(\omega_r) = 0, \quad |\omega_r| > \frac{\pi}{4} \quad (21)$$

In Figure 7a, an example of radial filters satisfying these equations is shown. It can be seen that the scaling filter has a band center frequency $\sim \pi/8$, as a result the frequency resolution of this DST may be rather poor (for the second scale, this frequency becomes $\sim \pi/32$; hence much of the frequency content of the image is contained in the first scale). Therefore, we propose to replace (21) by a less strong condition:

$$H_0(\omega_r) = \tilde{H}_0(\omega_r) = 0, \quad |\omega_r| > \frac{\pi}{2}$$

and modify the decimation operations appropriately, such that there is no information loss and hence PR is still possible. For the first scale, we set $p_1 = 2$, $q_1 = 1$ and starting from the second scale ($j > 1$), we use $p_j = 4$ and $q_j = 2$. In Table 1, the decimation factors p_j and q_j are listed per scale. The modified Meyer wavelet and scaling filters with adjusted frequency scaling is shown in Figure 7b. The use of the Meyer wavelet here is an appealing choice due to its excellent localization properties in both time and frequency and also because the filters are defined directly in frequency domain [31]:

$$H_0(\omega_r) = \begin{cases} 1 & |\omega_r| < \frac{\pi}{4} \\ \cos\left(\frac{\pi}{2}v\left(\frac{4|\omega_r|}{\pi} - 1\right)\right) & \frac{\pi}{4} \leq |\omega_r| \leq \frac{\pi}{2}, \\ 0 & \text{else} \end{cases}$$

$$G_0(\omega_r) = \begin{cases} 0 & |\omega_r| < \frac{\pi}{4} \\ \sin\left(\frac{\pi}{2}v\left(\frac{4|\omega_r|}{\pi} - 1\right)\right) & \frac{\pi}{4} \leq |\omega_r| \leq \frac{\pi}{2}, \\ 1 & \text{else} \end{cases}$$

$$\tilde{H}_0(\omega_r) = H_0(\omega_r),$$

$$\tilde{G}_0(\omega_r) = G_0(\omega_r)$$

where we use the following interpolation function (Figure 6):

$$v(x) = \begin{cases} 3x^2 - 2x^3 & 0 \leq x \leq 1 \\ 0 & x < 0 \\ 1 & 1 < x \end{cases} \quad (22)$$

for $x \in [0, 1]$. This function is chosen such that it satisfies $v(x) = 1 - v(1 - x)$ (see [31]), while being a C^3 function.




Figure 6. The interpolation function $v(x)$.

scale j	p_j	q_j
1	2	1
2	4	1
3	4	2
4	4	4
5	4	4

Table 1. Proposed decimation factors for the DST with anisotropic dilation (also see Figure 5).

Similarly, angular filters satisfying (21) are given by:

$$R(x) = \tilde{R}(x) = \begin{cases} 0 & x < -\frac{1+\alpha}{2} \\ \sin\left(\frac{\pi}{2}v\left(\frac{\alpha+2x+1}{2\alpha}\right)\right) & |x + \frac{1}{2}| \leq \frac{\alpha}{2} \\ 1 & |x| < \frac{1-\alpha}{2} \\ \cos\left(\frac{\pi}{2}v\left(\frac{\alpha+2x-1}{2\alpha}\right)\right) & |x - \frac{1}{2}| \leq \frac{\alpha}{2} \\ 0 & \text{else} \end{cases}$$

with $\alpha \in [0, \frac{1}{2}]$ a constant parameter that determines the bandwidth of the angular filters. In Figure 7c, $R(x)$ is depicted for different values of α . Higher values of α correspond to a slower decay of the transition bandwidth. The corresponding filters $\tilde{G}_k(\omega_r, \vartheta)$ for $\omega_r = \frac{\pi}{2}$ and $\alpha = \frac{1}{2}$ are shown in Figure 7d. The choice of α has an influence on the redundancy factor of the DST. We will go deeper into this in Section 4.4.

4.3. Folding and angular decimation

By the compact support of $R(x)$, the filters $G_k(\omega_r, \vartheta)$ are supported on trapezoidal wedges in the frequency plane. Outside these wedges, the filtered DFT coefficients are 0. In case of more than two orientations ($K > 2$), we can partially get rid of the extra redundancy in two different ways (see Figure 8):

- (*Folding*) Shear the filtered subbands such that the frequency support is fully contained in the central rectangles as shown in Figure 8a. Subsequently, a vertical decimation can be applied to the subbands in cone C_1 and a horizontal decimation to the subbands in cone C_2 . Note that a suitable (possibly non-integer) decimation factor needs to be chosen, we will explain this further on. For the shear transform, we rely on bandlimited interpolation (most efficiently implemented in the DFT domain). For the exact details of the shear transform implementation, we refer to [36].
- (*Non-folding*) Perfect reconstruction is possible even without folding. Therefore we need to make sure

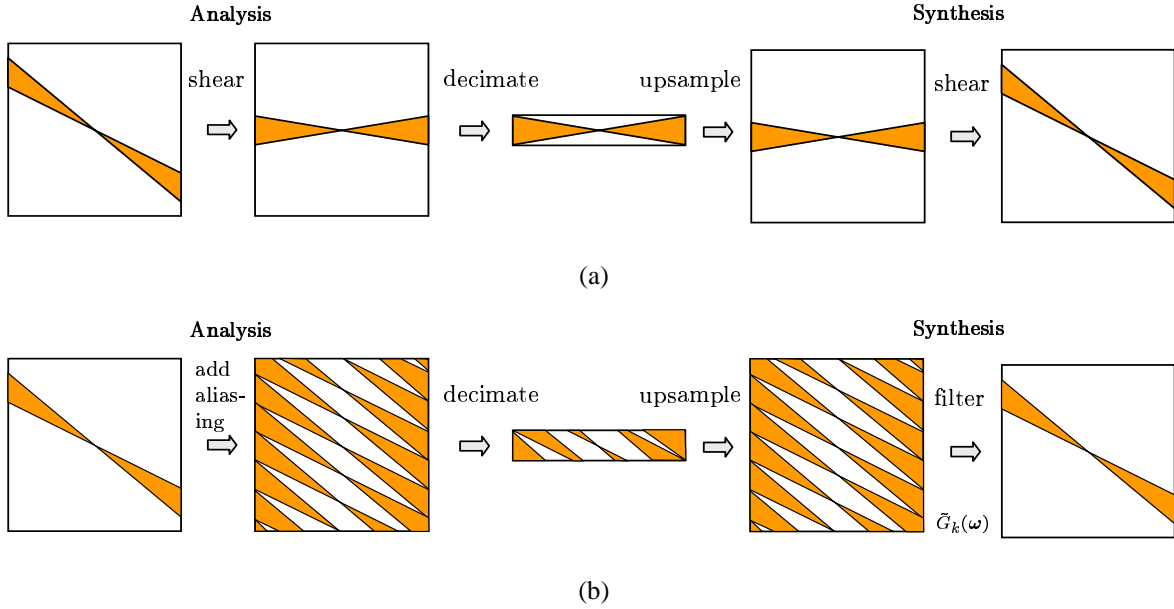


Figure 8. Strategies to reduce the redundancy of the DST. (a) Perfect reconstruction by shear operations and decimating (*Folding*), (b) Perfect reconstruction by decimating without shearing (*Non-folding*). See text also.

that the aliasing caused by the decimations does not contaminate the content of the wedges of interest. This can again be done by choosing the decimation factor suitably (actually the same as in the folding strategy). This approach is illustrated in Figure 8b. Even though many aliasing copies are produced during decimation, the original wedges can be perfectly reconstructed after applying the reconstruction filter $\tilde{G}_k(\omega_r, \vartheta)$.

Both schemes have the same redundancy factor. The difference is that in the *folding* strategy, the translation lattice is sheared, while without folding, the translation lattice remains Cartesian, which can be an advantage in certain applications. Additionally, the *non-folding* strategy heavily relies on aliasing and it is easy to show that the non-folding strategy is *not* shift-invariant, whereas the folding strategy *is* shift-invariant.

For a squared subband at scale j of size N_j , we compute the decimation factor from Figure 8 as follows:

$$d_j = \max \left(1, \frac{N_j}{\lceil (1 + 2\alpha)N_j / (K_j/2) \rceil} \right) \quad (23)$$

In our implementation, the folding is performed in the DFT domain; N_j/d_j then determines the integer number of DFT coefficients to keep per row or column. For this reason a ceiling function is present in equation (23).

More importantly, because we have K_j orientation subbands per scale, we see that the redundancy for scale j of the transform, that is proportional to $K_j/d_j \approx 2(1 + 2\alpha)$, becomes independent of K_j !³

Because the filters are defined in frequency domain, our current implementation of this filterbank makes use of

³Up to small deviations caused by the ceiling operation, but this is usually neglectable.

number of scales J	α		
	$\frac{1}{32}$	$\frac{1}{8}$	$\frac{1}{2}$
1	2.19	2.56	4.06
2	2.66	3.13	5.00
3	2.67	3.14	5.03
4	2.67	3.15	5.03
5	2.67	3.15	5.03

Table 2. Redundancy factors for J scales, computed using equation (24).

FFTs. Because the filters $H_0(\omega_r)$, $G_0(\omega_r)$, $\tilde{H}_0(\omega_r)$ and $\tilde{G}_0(\omega_r)$ are bandlimited, the filters do not have compact support in spatial domain. Nevertheless, it is possible to approximate the impulse response by truncation, as proposed e.g. in [37] for the steerable pyramid filters.

4.4. Computation of the redundancy factor

The redundancy factor for our scheme is given by the recursive formula:

$$R \approx \left[\frac{2}{q_1} + \frac{1}{p_1^2} \left(\frac{2}{q_2} + \frac{1}{p_2^2} \left(\frac{2}{q_3} + \frac{1}{p_3^2} \left(\frac{2}{q_4} + \dots \right) \right) \right) \right] \cdot (1 + 2\alpha) + 2^{-2 \sum_{j=1}^J p_j} \quad (24)$$

with p_j and q_j as listed in Table 1 and where the approximation sign is due to neglecting the ceiling operation in equation (23). Each fraction $2/q_j$ corresponds to shearlet subbands for scale j , while the fractions $1/p_j^2$ are related to the decimations during the scaling steps. In Table 2, redundancy factors of the transform are given with respect to the number of scales J and the parameter α .

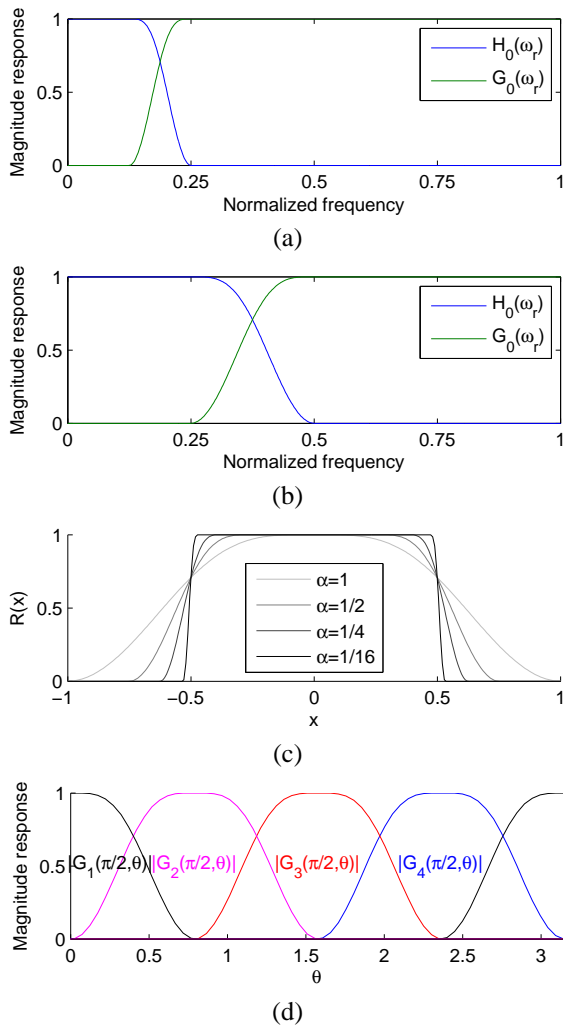


Figure 7. (a) Shearlet radial magnitude responses for dilation factor 4 (using the Meyer wavelet), (b) Shearlet filter radial magnitude responses with proposed adjustment to increase the lowpass center band frequency (using the Meyer wavelet), (c) Angular response $R(x)$, (d) Shearlet filter magnitude responses for the constant radial frequency $\omega_r = \pi/2$.

5. RESULTS AND DISCUSSION

First, we will look closer at some of the properties of our DST implementation. In Figure 9, shearlet basis functions are shown for different scales and orientations. Even though the size of the support of these basis functions is not finite, these functions have a fast decay and are well localized in space, frequency and orientation. In Figure 10, the frequency response of one shearlet function is depicted. This function has a very compact support in frequency domain and a clear orientation. In Figure 11, the influence of the parameter α on the spatial localization of the shearlet is illustrated: the smaller α , the larger the side-lobes of the basis function, hence the worse the spatial localization. As mentioned before, the redundancy of our DST is directly related to α (values are given in Table 2), so we can conclude that α allows us to trade off

redundancy to spatial localization.

In Figure 12, the DST subband decomposition of the “blackman” test image is shown. Here, we used the folding strategy to reduce the redundancy of the transform. In the magnified part of one of the subbands, it can be noted that the DST coefficients are not subject to spurious oscillations (aliasing) as for example shown in Figure 12d. These spurious oscillations are often present in subbands of shift-variant transforms and are very disturbing as the local energy signature of edges depends on the exact edge position.

Another interesting experiment is the trade-off between spatial localization and the number of orientation subbands while keeping $\alpha = 1/2$. For our DST implementation, this also means that the redundancy factor of the transform remains constant. In Figure 13, we reconstruct the zone plate image from the DST coefficients from one orientation subband from the first (finest) scale of the transform, while increasing the number of orientations K_1 for the first scale. We observe that for $K_1 \geq 16$, edges with orientation around 135° are still well detected, however, their position becomes less certain, as the responses are more spread over the entire image. This shows that the spatial localization properties of the shearlet basis elements are heavily reduced in this case.

As a final example, we investigate the approximation quality of different multiresolution transforms. We start from a test image, apply a given wavelet or shearlet transform to this image and we reconstruct the image from the 2.5% largest wavelet or shearlet coefficients (in magnitude). In Figure 14, the results are given for the zone plate image and the barbara image, for the decimated DWT, the undecimated DWT, the dual-tree complex wavelet transform (DT-CWT) and the DST. We also list the redundancy factor of each transform in the figure, because this factor plays a big role here. The DWT is a shift variant transform, and the aliasing creates disturbing artifacts in the end result. The undecimated DWT has the largest number of coefficients retained in absolute terms, however the basis functions of this transform corresponding to the HH_j subbands have a poor directional selectivity, which causes here the blurring of some of the edges. The DT-CWT basis functions have an excellent spatial localization, but are only able to distinguish 6 orientations, also causing a fair amount of blurring here (see Figure 14c/g). The DST gives here the best visual result, mainly because of its excellent directional selectivity and shift-invariance.

6. CONCLUSION

In this paper, we described a novel design of the Discrete Shearlet Transform with a redundancy factor that is very low and independent of the number of orientation subbands, while offering shift-invariance. The filters of this transform are designed in pseudo-polar frequency coordinates, based on the Meyer wavelet. A special decimation scheme is applied to the filtered subbands, to further reduce the redundancy. This results in a multiresolution transform, that allows to easily make a trade-off in spatial

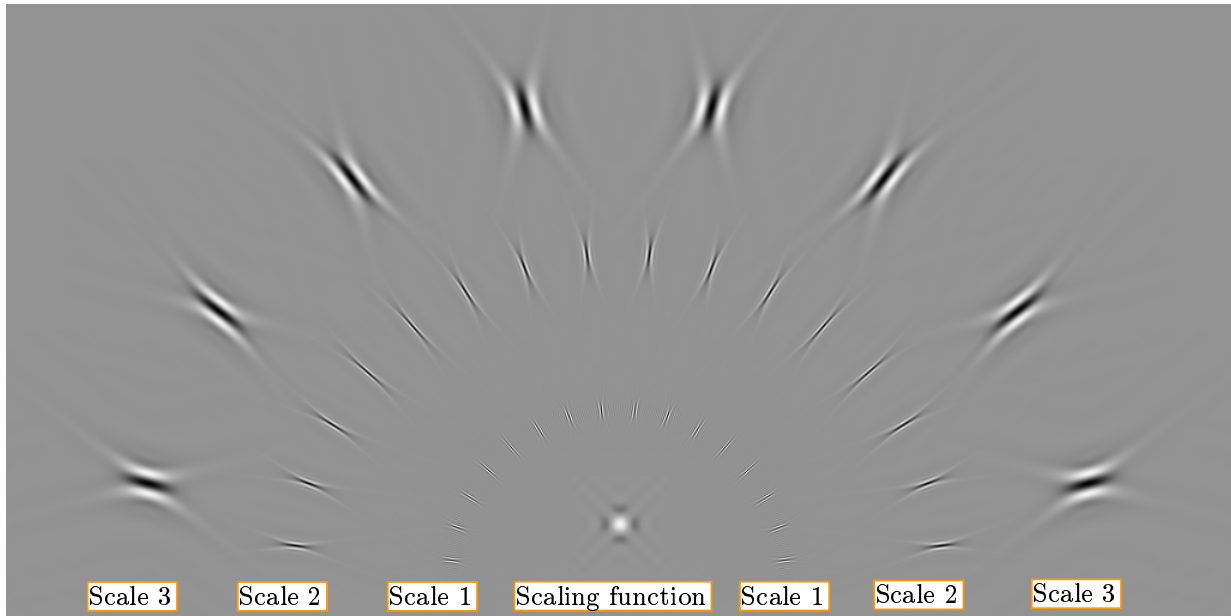


Figure 9. Shearlets basis elements for $\alpha = \frac{\pi}{2}$. For illustration purposes, we used $K_1 = 16$ (instead of 32) orientations for the first scale.




Figure 12. (a) The “blackman” image. (b) DST decomposition (with folding) of the “blackman” image into three scales. Eight orientations are used for the first scale. Note that the subbands for the finest scale of cone C_2 have been rotated 90° . (c) Crop out of one subband, to illustrate the lack of aliasing. (d) Aliasing in a subband of the DST (implemented without folding) of the image in (a).




Figure 13. Illustration of the directional selectivity of the DST (a) the zone plate image (b)-(f) reconstruction from one orientation subband of a DST with one scale ($J = 1$), and for an increasing number of orientations (K_1).




Figure 14. Reconstruction from 2.5% of the x -let coefficients, for (a),(e) DWT with 2 scales, (b),(f) Undecimated DWT with 2 scales, (c),(g) dual-tree complex wavelet transform [6] with 2 scales, (d),(h) DST with 1 scale (because of anisotropic dilation with factor 4). *Top row*: crop out of the Barbara image, *bottom row*: crop out of the zone plate image (see Figure 13a). Between parentheses is the redundancy factor of each transform in this experiment.




Figure 10. Magnitude response of a shearlet basis element.




Figure 11. Shearlets for different values of the parameter α , showing the impact of the angular bandwidth on the spatial support of the shearlets.

localization, frequency localization and directional selectivity, while keeping the redundancy factor more or less constant. There are a number of interesting open issues remaining. First is how the idea of scale and orientation shiftability from [1] can be applied to this transform, without sacrificing redundancy. A second issue is that the redundancy of the transform can possibly be decreased further, while allowing certain forms of aliasing (hence being not completely shift-invariant). Finally, another open issue is the design of non-separable or separable shearlet filters with compact support in spatial domain, while staying faithful to the CST.

Acknowledgements: The authors would like to thank Prof. Ingrid Daubechies for the fruitful and inspiring discussions that initiated this work.

7. REFERENCES

- [1] W. Freeman and E. Adelson, "Design and use of steerable filters," *IEEE Transactions on Pattern analysis and Machine Intelligence*, vol. 13, no. 9, pp. 891–906, 1991.
- [2] E. Simoncelli, W. T. Freeman, E. H. Adelson, and D. J. Heeger, "Shiftable Multi-scale Transforms," *IEEE Trans. Information Theory*, vol. 38, no. 2, pp. 587–607, 1992.
- [3] N. G. Kingsbury, "Image processing with complex wavelets," *Philos. Trans. R. Soc. London A., Math Phys. Sci.*, vol. 357, no. 1760, pp. 2543–2560, 1999.
- [4] N. G. Kingsbury, "Complex Wavelets for Shift Invariant Analysis and Filtering of Signals," *Journal of Applied and Computational Harmonic Analysis*, vol. 10, pp. 234–253, May 2001.
- [5] F. C. A. Fernandes, R. L. C. Van Spaendonck, and C. S. Burrus, "A new framework for complex wavelet transforms," *IEEE Trans. Signal Processing*, vol. 51, no. 7, pp. 1825–1837, 2003.
- [6] I. W. Selesnick, R. G. Baraniuk, and N. G. Kingsbury, "The Dual-Tree Complex Wavelet Transform," *IEEE Signal Processing Magazine*, vol. 22, pp. 123–151, Nov. 2005.
- [7] C. Chau, L. Duval, and J.-C. Pesquet, "Image analysis using a dual-tree m-band wavelet transform," *IEEE Trans. Image Processing*, vol. 15, pp. 2397–2412, Aug. 2006.
- [8] A. A. Barath and J. Ng, "A steerable complex wavelet construction and its application to image denoising," *IEEE Trans. Image Processing*, vol. 14, no. 7, pp. 948–959, 2005.
- [9] D. Van De Ville and M. Unser, "Complex Wavelet Bases, Steerability, and the Marr-like pyramid," *IEEE Trans. Image Processing*, vol. 17, pp. 2063–2080, Nov. 2008.
- [10] T. Lee, "Image Representation Using 2D Gabor Wavelets," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 18, p. 1, Oct. 1996.
- [11] S. Fischer, F. Šroubek, L. Perrinet, R. Redondo, and G. Cristobal, "Self-Invertible 2D Log-Gabor Wavelets," *International Journal of Computer Vision*, vol. 75, pp. 231–246, Nov. 2007.
- [12] M. N. Do and M. Vetterli, *Beyond Wavelets, chapter Contourlets*. Academic Press, 2003.
- [13] M. N. Do and M. Vetterli, "The contourlet transform: An efficient directional multiresolution image representation," *IEEE Trans. Image Process.*, vol. 14, no. 12, pp. 2091–2106, 2005.
- [14] A. da Cunha, J. Zhou, and M. Do, "The nonsubsampled contourlet transform: theory, design, and applications," *IEEE Trans. Image Processing*, vol. 15, pp. 3089–3101, Oct 2006.
- [15] E. Candès, *Ridgelets: Theory and Applications*. PhD thesis, Department of Statistics, Stanford University, 1998.
- [16] M. N. Do and M. Vetterli, "The finite ridgelet transform for image representation," *IEEE Trans. Image Processing*, vol. 12, pp. 16–28, Jan. 2003.

- [17] D. L. Donoho, "Wedgelets: Nearly minimax estimation of edges," *Annals of Statistics*, vol. 27, no. 3, pp. 859–897, 1999.
- [18] J. Romberg, M. Wakin, and R. Baraniuk, "Multi-scale wedgelet image analysis: fast decompositions and modeling," in *Proc. International Conference on Image Processing 2002*, vol. 3, pp. 585–588, 24–28 June 2002.
- [19] E. LePennec and S. Mallat, "Sparse geometric image representation with bandelets," *IEEE Trans. Image Processing*, vol. 14, no. 4, pp. 423–438, 2005.
- [20] R. R. Coifman and F. G. Meyer, "Brushlets: a tool for directional image analysis and image compression," *Appl. Comp. Harmonic Anal.*, vol. 5, pp. 147–187, 1997.
- [21] J. L. Starck, E. J. Candès, and D. L. Donoho, "The curvelet transform for image denoising," *IEEE Trans. on Image Process.*, vol. 11, pp. 670–684, 2000.
- [22] E. Candès, L. Demanet, D. Donoho, and L. Ying, "Fast Discrete Curvelet Transforms," *Multiscale modeling and simulation*, vol. 5, no. 3, pp. 861–899, 2006.
- [23] R. A. Gopinath, "The Phaselet Transform - An Integral Redundancy Nearly Shift-Invariant Wavelet Transform," *IEEE Trans. Signal Processing*, vol. 51, no. 7, pp. 1792–1805, 2003.
- [24] V. Velisavljevic, B. Beferull-Lozano, M. Vetterli, and P. Dragotti, "Directionlets: anisotropic multi-directional representation with separable filtering," *IEEE Trans. Image Processing*, vol. 15, no. 7, pp. 1916–1933, 2006.
- [25] Y. M. Lu and M. N. Do, "Multidimensional Directional Filter Banks and Surfacelets," *IEEE Trans. Image Processing*, vol. 16, no. 4, pp. 918–931, 2007.
- [26] K. Guo and D. Labate, "Optimally Sparse Multidimensional Representation using Shearlets," *SIAM J Math. Anal.*, vol. 39, pp. 298–318, 2007.
- [27] G. Easley, D. Labate, and W. Lim, "Sparse Directional Image Representations using the Discrete Shearlet Transform," *Applied and Computational Harmonic Analysis*, vol. 25, pp. 25–46, 2008.
- [28] G. Easley, D. Labate, and F. Colonna, "Shearlet Based Total Variation for Denoising," *IEEE Trans. Image Process.*, vol. 18, no. 2, pp. 260–268, 2009.
- [29] K. Guo, D. Labate, and W. Lim, "Edge Analysis and Identification using the Continuous Shearlet Transform," *Applied and Computational Harmonic Analysis*, vol. 27, pp. 24–46, July 2009.
- [30] S. Yi, D. Labate, G. Easley, and H. Krim, "A Shearlet Approach to Edge Analysis and Detection," *IEEE Trans. Image Processing*, vol. 18, pp. 929–941, May 2009.
- [31] I. Daubechies, *Ten Lectures on Wavelets*. Philadelphia: Society for Industrial and Applied Mathematics, 1992.
- [32] P. J. Burt and E. H. Adelson, "The Laplacian pyramid as a compact image code," *IEEE Trans. Commun.*, vol. COM-31, no. 4, pp. 532–540, 1983.
- [33] M. N. Do and M. Vetterli, "Frame reconstruction of the Laplacian pyramid," in *IEEE Conf. on Acoustics, Speech and Signal Processing*, vol. 6, pp. 3641–3644, 2001.
- [34] A. Averbuch, R. Coifman, D. Donoho, M. Elad, and M. Israeli, "Fast and Accurate Polar Fourier Transform," *Journal on Appl. and Comp. Harm. Analysis*, vol. 21, pp. 145–167, 2006.
- [35] E. P. Simoncelli and W. T. Freeman, "The Steerable Pyramid: A flexible architecture for Multi-scale Derivative Computation," in *Proc IEEE Int. Conf. Image Processing*, (Washington, DC.), Oct. 1995.
- [36] L. Condat, D. Van De Ville, and B. Forster-Heinlein, "Reversible, Fast, and High-Quality Grid Conversions," *IEEE Trans. Image Processing*, vol. 17, pp. 679–693, May 2008.
- [37] K. Castleman, M. Schulze, and Q. Wu, "Simplified design of steerable pyramid filters," in *Proceedings of the 1998 IEEE Int. Symposium on Circuits and Systems (ISCAS)*, vol. 5, pp. 329–332, May 1998.