

A split-augmented Lagrangian Algorithm for Spectral Factorization of a set of 2D Directional Filters and Application to the Design of Compact Shearlet Frames

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

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

ABSTRACT

In this paper, we first briefly review the directional properties of the Dual-Tree complex wavelet transform and we investigate how the directional selectivity of the transform can be increased (i.e., to obtain more than 6 orientations per scale). To this end, we describe a new augmented Lagrangian optimization algorithm to jointly perform the 2D spectral factorization of a set of 2D directional filters, with a high numerical accuracy. We demonstrate how this approach can be used to design compactly supported shearlet frames that are tight. Finally, a number of experimental results are given to show the merits of the resulting shearlet frames.

Keywords: multiresolution analysis, complex wavelets, shearlets

1. INTRODUCTION

The success of recent image restoration and reconstruction techniques highly relies on the fact that images are sparse in a certain basis. This basis is then closely related to features observed in these images (e.g., points, lines, curves, textures). Finding good representations (bases or frames) for general classes of images (e.g., photographic images) is however a difficult problem. Sparsifying transforms focus on exploiting low-level information in the images (e.g. correlation between intensities of neighboring pixels). An important class of sparsifying transforms are the multiresolution transforms, which represent the image in a natural way by successively adding detail information in subsequent refinement steps. The quest for good sparsifying transforms has led to the development of steerable pyramids,¹ dual-tree complex wavelets,² 2-D (log) Gabor transforms,^{3,4} contourlets,⁵ ridgelets,^{6,7} wedgelets,⁸ bandelets,⁹ brushlets,¹⁰ curvelets,¹¹ phaselets,¹² directionlets¹³ and surfacelets.¹⁴

Despite the large number of available x-let transforms, there is still room left for further improvement: today's applications require processing of 3D (or even 4D) datasets in a reasonable amount of time. Several transforms that have been developed in 2D do not scale well to higher dimensions because of the large computational requirements or the excessive amount of memory needed to store each individual directional subband.

Shearlets^{15,16} are a relatively novel sibling in the family of multidirectional and multiresolution representations and have the main advantage of allowing a very fine directional analysis with an arbitrary number of directions per scale. Furthermore, shearlets are well suited for representing data defined on a Cartesian grid. In particular, the latter opens a number of possibilities to reduce the redundancy of the transform.^{17,18} Although "traditional" shearlets are bandlimited and hence most efficiently implemented using fast Fourier transforms (FFTs), recently compactly supported shearlet systems have been proposed^{19,20} as an alternative. These compactly supported systems can then efficiently be implemented using several finite impulse response filters. Despite the resulting shearlets having nice directional properties and efficient implementations are possible using digital shearing and the DWT, the resulting frame bounds were estimated to be rather high.¹⁹ Frame bounds are important indicators of numerical stability of the transform inversion: when a transform coefficient is slightly modified, they determine the corresponding reconstruction error. Tight frames (i.e., with equal frame bounds) or snug frames (with frame bound ratio close to 1) are highly desired in practical applications such as image reconstruction and image restoration.

Bart Goossens is a post-doctoral researcher of FWO Flanders. Send correspondence to Bart Goossens (e-mail: bart.goossens@telin.ugent.be)

In this paper, we analyze the properties of the dual-tree complex wavelet transform (DT-CWT) in Section 2, which has equal frame bounds and provides a directional analysis in 6 orientations (in 2D) and is computationally efficient due to the use of separable filters. We investigate whether it is possible to design a compactly supported tight shearlet frame based on insights brought by the DT-CWT. Based on a novel augmented Lagrangian filter optimization method (see Section 3), we find that it is possible to design tight frames of compactly supported shearlets. In Section 4, we illustrate the properties of the transform in image approximation and denoising applications. Finally, Section 5 concludes this paper.

2. OVERVIEW OF THE DT-CWT AND IMPROVING ITS DIRECTIONAL SELECTIVITY

The dual-tree complex wavelet transform (DT-CWT)²¹ is directly related to the DWT and also provides a multiresolution analysis (MRA).²² In contrast to the DWT, the DT-CWT uses complex-valued wavelets instead of real-valued wavelets. When additionally the wavelets fulfill the so-called Hilbert-transform pair property (see further), the transform can also be designed in such a way that it performs a *multi-directional* analysis (which practically means that features that have a dominant direction, such as edges in images, can be more compactly represented). This is beneficial for analyzing higher-dimensional data. In this Section, we give a brief explanation of how the DT-CWT achieves its directionality properties and how the directionality can be further improved (e.g. by increasing the number of orientations).

2.1 One-dimensional complex wavelets

The complex wavelets of the DT-CWT obey a special property:²³ the imaginary part of the wavelet is the Hilbert transform of the real-part.*

$$\psi_c(t) = \psi(t) + i\mathcal{H}\{\psi\}(t). \quad (1)$$

where i is the imaginary unit. Here, the Hilbert transform of $\psi(t)$ is defined as $\mathcal{H}\{\psi\}(t) = \frac{1}{\pi} \int_{-\infty}^{+\infty} \frac{\psi(\tau)}{t-\tau} d\tau$ or, in the Fourier domain $\widehat{\mathcal{H}\{\psi\}}(\omega) = H(\omega)\widehat{\psi}(\omega)$ with:

$$H(\omega) = \begin{cases} -i, & \omega > 0 \\ 0, & \omega = 0 \\ i, & \omega < 0 \end{cases}$$

Based on (1), it can easily be shown that $\psi_c(t)$ is analytic ($\psi_c(t)$ is only supported on positive frequencies ($\omega > 0$)), such that the magnitude of the resulting complex wavelet coefficients is shift-invariant. In Figure 1, an example of a complex wavelet is shown. It can be seen that the complex wavelet is well localized in time and that its magnitude is similar to a bump function.

By projecting a real-valued function $f(t)$ onto the complex-valued wavelet functions:

$$\psi_{a,b}(t) = \frac{1}{\sqrt{a}} \psi_c\left(\frac{t-b}{a}\right), \quad (2)$$

complex wavelet coefficients $\langle f, \psi_{a,b} \rangle$ are obtained, from which the magnitude and phase can be easily computed: a large magnitude then reveals the presence of a singularity (e.g. edge, texture, point, ...), while the phase indicates the position of the singularity within the support of the complex wavelet.²³

In practice, computing the inner products one by one is rather slow, therefore discrete implementations of the DT-CWT transform make use of the fast DWT decomposition scheme (see Ref.²³): two fast DWTs are applied in parallel, one DWT for the real part of the complex wavelet, a second DWT for the imaginary part. Because there are two wavelet decomposition trees, the transform is called the *dual-tree* complex wavelet transform.

*We remark that (1) is not the only way to construct complex wavelets. More general complex wavelets $\psi_c(t)$ are investigated e.g. by Belzer et al.²⁴ and Lina and Mayrand²⁵, however, these wavelets are not analytic.

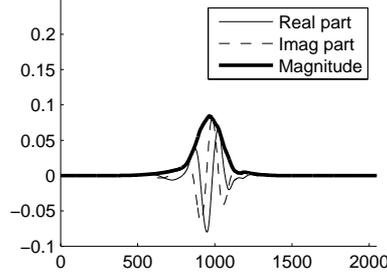


Figure 1. A complex wavelet $\psi_c(t)$, together with its real and imaginary parts.

2.2 Two-dimensional complex wavelets

As in the real-valued DWT, higher-dimensional wavelets are constructed as a tensor-product of one-dimensional wavelets.²³ Let us consider a pair of complex wavelets $\psi_{c,x}(t) = \psi_x(t) + i\mathcal{H}\{\psi_x\}(t)$ and $\psi_{c,y}(t) = \psi_y(t) + i\mathcal{H}\{\psi_y\}(t)$ with Fourier transforms $\widehat{\psi}_{c,x}(\omega) = (1 + iH(\omega))\widehat{\psi}_x(\omega)$ and $\widehat{\psi}_{c,y}(\omega) = (1 + iH(\omega))\widehat{\psi}_y(\omega)$, respectively. The Fourier transform of the corresponding complex wavelet is:

$$\begin{aligned}\widehat{\psi}_{c,1}(\omega_x, \omega_y) &= \widehat{\psi}_{c,x}(\omega_x)\widehat{\psi}_{c,y}(\omega_y) \\ &= \widehat{\psi}_x(\omega_x)\widehat{\psi}_y(\omega_y)(1 - H(\omega_x)H(\omega_y) + i(H(\omega_x) + H(\omega_y))) \\ &= 4\widehat{\psi}_x(\omega_x)\widehat{\psi}_y(\omega_y)Y(\omega_x)Y(\omega_y),\end{aligned}\quad (3)$$

with $Y(\omega)$ the Heaviside step function. From (3) it can be seen that $\widehat{\psi}_{c,1}(\omega_x, \omega_y)$ only passes positive horizontal and vertical frequencies, hence $\widehat{\psi}_{c,1}(\omega_x, \omega_y)$ is supported on the first quadrant of the 2D frequency plane. The same way, it is possible to design complex wavelets that pass frequencies in other quadrants of the frequency plane, by using conjugates of the one-dimensional complex wavelets, e.g.:

$$\widehat{\psi}_{c,2}(\omega_x, \omega_y) = \overline{\widehat{\psi}_{c,x}(\omega_x)}\widehat{\psi}_{c,y}(\omega_y). \quad (4)$$

It can be shown that the frequency response of $\psi_{c,2}(x, y)$ is given by $\widehat{\psi}_{c,2}(\omega_x, \omega_y) = 4\widehat{\psi}_x(\omega_x)\widehat{\psi}_y(\omega_y)Y(-\omega_x)Y(\omega_y)$, which is supported in the second quadrant of the frequency plane (i.e. $\omega_x < 0$ and $\omega_y > 0$). A frequency domain illustration is given in Figure 2.

Now, if we look at the frequency response of the real part and imaginary part of the 2D complex wavelets, i.e.:

$$\widehat{\psi}_{c,k}(\omega_x, \omega_y) = \widehat{\psi}_{re,k}(\omega_x, \omega_y) + i\widehat{\psi}_{im,k}(\omega_x, \omega_y), \quad k = 1, 2 \quad (5)$$

then we find that, by the conjugate symmetry property of the Fourier transform, both parts have the same magnitude response:

$$\begin{aligned}|\widehat{\psi}_{re,1}(\omega_x, \omega_y)| &= |\widehat{\psi}_{im,1}(\omega_x, \omega_y)| = 2|\widehat{\psi}_x(\omega_x)||\widehat{\psi}_y(\omega_y)|(Y(\omega_x)Y(\omega_y) + Y(-\omega_x)Y(-\omega_y)) \\ |\widehat{\psi}_{re,2}(\omega_x, \omega_y)| &= |\widehat{\psi}_{im,2}(\omega_x, \omega_y)| = 2|\widehat{\psi}_x(\omega_x)||\widehat{\psi}_y(\omega_y)|(Y(-\omega_x)Y(\omega_y) + Y(\omega_x)Y(-\omega_y))\end{aligned}\quad (6)$$

where the last factor consisting of the Heaviside functions selects two quadrants of the Fourier space (see Figure 2(d)-(e)). Hence, by using a pair of complex wavelets: $\Psi_{c,x}(\omega_x)$, $\Psi_{c,y}(\omega_y)$ and their conjugates $\overline{\Psi_{c,x}(\omega_x)}$, $\overline{\Psi_{c,y}(\omega_y)}$, followed by taking the real parts of the resulting complex wavelet coefficients, we obtain two orientation bands with angles $\pm 45^\circ$. This elegantly solves the directionality problem of the separable DWT, that cannot distinguish between orientation angles $+45^\circ$ and -45° . When wavelet filters are used together with scaling filters in a multiresolution approach, an analysis is possible in 6 orientation angles (in 2D) instead of 3. Also remarkable is that the real or imaginary parts of the complex wavelets are *non-separable*. However, due to the tensor-product, the DT-CWT can still be implemented efficiently based on *separable* wavelet filters. The next question is whether the directional selectivity of the complex wavelets can be further increased while retaining a tight frame representation. A trivial solution is by making use of complex wavelet packet based-transforms, which can have arbitrary numbers of orientations. However, due to the separability of the complex wavelets (3), the equi-energy contours of the complex wavelet frequency responses have a rectangular shape rather than a wedge-like shape. Wedge-like shapes are desired in practice, because they permit parabolic scaling²⁶ for arbitrary orientations. A design of compactly supported shearlets can thus offer a solution here.

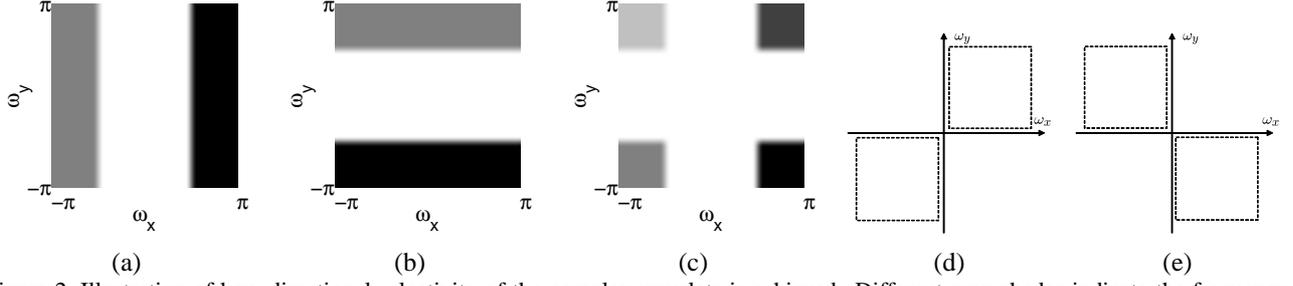


Figure 2. Illustration of how directional selectivity of the complex wavelets is achieved. Different gray-shades indicate the frequency spectra of different complex wavelets: (a) $\widehat{\psi}_{c,x}(\omega_x)$ and $\widehat{\psi}_{c,x}(-\omega_x)$, (b) $\widehat{\psi}_{c,y}(\omega_y)$ and $\widehat{\psi}_{c,y}(-\omega_y)$, (c) The resulting 2D complex wavelets $\widehat{\psi}_{c,k}(\omega_x, \omega_y)$ obtained as tensor products of 1D complex wavelets, (d) Frequency support of $\psi_{re,1}(\omega_x, \omega_y)$ (or $\psi_{im,1}(\omega_x, \omega_y)$), (e) Frequency support of $\psi_{re,2}(\omega_x, \omega_y)$ (or $\psi_{im,2}(\omega_x, \omega_y)$).

3. DESIGN OF A PARSEVAL FRAME OF COMPACTLY SUPPORTED SHEARLETS

3.1 Overview of 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 $f(\mathbf{x}) \in L^2(\mathbb{R}^2)$ are given by:^{27,28}

$$w_{j,k,\mathbf{l}} = \langle f, \psi_{j,k,\mathbf{l}} \rangle = \int_{\mathbb{R}^2} f(\mathbf{x}) \overline{\psi_{j,k,\mathbf{l}}(\mathbf{x})} d\mathbf{x}, \quad (7)$$

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 shearlet transform combines geometry and multiscale analysis:¹⁶ shearlets are formed by dilating, shearing and translating a mother shearlet function $\psi \in L^2(\mathbb{R}^2)$:

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

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 $\|\psi_{j,k,\mathbf{l}}\| = \|\psi\|$). The shearlet functions are subjected to a composite dilation \mathbf{A}^j and geometric transformation \mathbf{B}^k , e.g.,

$$\mathbf{A} = \begin{pmatrix} 4 & 0 \\ 0 & 2 \end{pmatrix} \quad \text{and} \quad \mathbf{B} = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}. \quad (9)$$

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. The shearlet mother function is defined in the Fourier transform domain as the product of the Fourier transform of a wavelet ($\Psi_1(\omega_x)$) and a compactly supported bump function ($\Psi_2(\omega_y)$) satisfying $\Psi_2(\omega_y) \neq 0 \Leftrightarrow \omega_y \in [-1, 1]$:

$$\Psi(\omega_x, \omega_y) = \Psi_1(\omega_x) \Psi_2\left(\frac{\omega_y}{\omega_x}\right), \quad (10)$$

Consequently, the mother shearlet function is bandlimited in a hourglass-shaped region of the 2D frequency spectrum:

$$\Psi(\omega) \neq 0 \Leftrightarrow |\omega_y| < |\omega_x|. \quad (11)$$

In practical implementations, 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):

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

with ω_0 the maximal frequency of the center square C_3 . This enables treating the horizontal and vertical frequencies in images equally, e.g. using the following dilation and shear matrices in both cones:

$$\mathbf{A}_1 = \begin{pmatrix} 4 & 0 \\ 0 & 2 \end{pmatrix}, \mathbf{B}_1 = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \quad (\text{cone } C_1) \quad \text{and} \quad \mathbf{A}_2 = \begin{pmatrix} 2 & 0 \\ 0 & 4 \end{pmatrix}, \mathbf{B}_2 = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \quad (\text{cone } C_2). \quad (12)$$

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 $d = 1, 2$:

$$\psi_{j,k,1}^{(d)}(\mathbf{x}) = |\det \mathbf{A}_d|^{j/2} \psi^{(d)}(\mathbf{B}_d^k \mathbf{A}_d^j \mathbf{x} - \mathbf{1}), \quad (13)$$

where $\psi^{(1)}(x, y) = \psi(x, y)$ and $\psi^{(2)}(x, y) = \psi(y, x)$.

The shearlet transform based on the band-limited Meyer wavelet can efficiently be implemented in the DFT domain^{16,17} and with a limited redundancy factor ranging from 2.66 to 5.32 in 2D and 3.5 to 7 in 3D, depending on the overlap of the bump functions $\Psi_2(\omega_y + k)$, $k \in \mathbb{Z}$. In contrast to (9), an isotropic scaling matrix, e.g., $\mathbf{A} = \begin{pmatrix} 2 & 0 \\ 0 & 2 \end{pmatrix}$, can also be used, this leads to directional shearlet transforms with the same number of orientations per scale.

3.2 A split-augmented Lagrangian algorithm

In Section 3.1, shearlets are defined in Fourier space as the product of the Fourier transform of a 1D wavelet ($\Psi_1(\omega_x)$) and a compactly supported bump function ($\Psi_2(\omega_y/\omega_x)$). Instead of adopting a frequency domain-based design for the shearlet basis functions, we will define the shearlet filters directly in the spatial domain. The support of $\mathcal{F}^{-1}\Psi_1(\omega_x)$ can be made compact, simply by selecting a compactly supported wavelet, such as the Daubechies' wavelets. The remaining problem is controlling the spatial support of $\mathcal{F}^{-1}\Psi_2(\omega_y/\omega_x)$. Our goal is therefore designing directional finite impulse response (FIR) filters for $\mathcal{F}^{-1}\Psi_2((\omega_y + k\omega_x)/\omega_x)$, $k \in \mathbb{Z}$.

As a starting point, we simply truncate the impulse response $\mathcal{F}^{-1}\Psi_2((\omega_y + k\omega_x)/\omega_x)$ to a support that is smaller than some given upper bound L , where $L > 1$ (i.e. we set filter coefficients outside this support to zero). Remark that simply truncating the impulse response does not guarantee perfect reconstruction of the corresponding transform, so extra steps are needed.

In the following, we present a numerical optimization technique that can be used to recover the perfect reconstruction and that simultaneously leads to a synthesis which is the adjoint of the analysis (giving us a Parseval frame). Let \mathbf{P}_k denote a vector of length L^2 , that is obtained by column-stacking the truncated impulse response of the analysis filter $\mathcal{F}^{-1}\Psi_2((\omega_y + k\omega_x)/\omega_x)$. Furthermore, let \mathbf{F}_1 denote a discrete 1D DFT matrix of size $L \times L$. Then the 2D Fourier transform matrix of size $L^2 \times L^2$ can be expressed as the Kronecker product of two 1D DFT matrices ($\mathbf{F}_2 = \mathbf{F}_1 \otimes \mathbf{F}_1$). Similarly, $\tilde{\mathbf{P}}_k$ is the synthesis filter corresponding to $\mathcal{F}^{-1}\Psi_2((\omega_y + k\omega_x)/\omega_x)$. Using this notation, 2D operations on the impulse responses are translated to matrix operations on 1D vectors. In order to build a tight frame of shearlets, we require that $\mathbf{F}_2 \tilde{\mathbf{P}}_k$ is the complex conjugate of $\mathbf{F}_2 \mathbf{P}_k$, i.e. $\tilde{\mathbf{P}}_k = \mathbf{F}_2^H \overline{\mathbf{F}_2 \mathbf{P}_k}$. Due to the properties of the DFT and the column-stacked order for the elements of \mathbf{P}_k , this is equivalent to stating that $\tilde{\mathbf{P}}_k$ is simply the reverse of the vector \mathbf{P}_k .

To obtain perfect reconstruction, the sum of the squared magnitudes of the filter responses in the DFT domain should have sum 1:

$$\sum_{k=1}^K \overline{\mathbf{F}_2 \tilde{\mathbf{P}}_k} \cdot \mathbf{F}_2 \mathbf{P}_k = \sum_{k=1}^K |\mathbf{F}_2 \mathbf{P}_k|^2 = \mathbf{1}_{L^2 \times 1} \quad (\text{perfect reconstruction}), \quad (14)$$

where $\mathbf{1}_{L^2 \times 1}$ is a vector of length L^2 that consists of ones and where \cdot is the Schur product (or element-wise product) of two vectors. Here, $|\cdot|$ denotes the element-wise modulus of a complex-valued vector.

Let $\mathbf{P}^{(0)}$ be the truncated impulse response of $\mathcal{F}^{-1}\Psi_2((\omega_y + k\omega_x)/\omega_x)$ (which we will call prototype filters in the remainder). Again, the L^2 values of the impulse response are re-ordered into a vector of length L^2 . Γ is a diagonal projection matrix with diagonal elements 0 and 1 that projects \mathbf{P}_k (or $\tilde{\mathbf{P}}_k$) onto a filter with the desired truncated support that is at least twice[†] as small as $L \times L$. Then our goal is to find filters \mathbf{P}_k such that 1) perfect reconstruction (equation (14))

[†]This is to avoid periodization effects (i.e., the circular convolution) caused by DFT-based filtering.

holds and 2) the following conditions are fulfilled:

$$\begin{aligned}\mathbf{P}_k &\approx \mathbf{P}^{(0)} && \text{(approximation for the truncated prototype filters)} \\ (\mathbf{I} - \Gamma)\mathbf{P}_k &= \mathbf{0} && \text{(compact support)}\end{aligned}$$

This leads us to the following constrained optimization problem:

$$\min_{\mathbf{P}_k, k=1, \dots, K} \sum_{k=1}^K \left\| \Gamma \mathbf{P}_k^{(0)} - \mathbf{P}_k \right\|^2 \quad \text{s.t.} \quad \sum_{k=1}^K |\mathbf{F}_2 \mathbf{P}_k|^2 = \mathbf{1}_{L^2 \times 1} \text{ and } (\mathbf{I} - \Gamma)\mathbf{P}_k = \mathbf{0} \quad (15)$$

Based on the fact that the squared magnitude responses of the prototype filters $|\mathbf{F}_2 \mathbf{P}_k^{(0)}|^2$ are known, the above problem could be solved by 1) designing the squared magnitude responses $|\mathbf{F}_2 \mathbf{P}_k|^2$ and 2) obtaining $\mathbf{F}_2 \mathbf{P}_k$ using a 2D spectral factorization method from $|\mathbf{F}_2 \mathbf{P}_k|^2$. Although several 2D spectral factorization methods exist^{29,30} based on 1D spectral factorization combined with helix filtering, we have noted that this approaches usually work well when the support of the filters \mathbf{P}_k is sufficiently small (e.g. smaller than 5×5). When calculating solutions for larger support size, one has to deal with numerical stability issues due to the limited floating point precision.

Instead, we solve the above joint spectral factorization problem using direct numerical methods. First, we note that the condition $(\mathbf{I} - \Gamma)\mathbf{P}_k$ is redundant as long as the initial solution \mathbf{P}_k is truncated (e.g. $\mathbf{P}_k = \Gamma \mathbf{P}_k$). This is due to the fact that $\left\| \Gamma \mathbf{P}_k^{(0)} - \mathbf{P}_k \right\|^2 = \left\| \Gamma \mathbf{P}_k^{(0)} - \Gamma \mathbf{P}_k \right\|^2$, from which it can be noted that non-zero values $(\mathbf{I} - \Gamma)\mathbf{P}_k$ do not contribute to the objective function. We can simplify the above problem to:

$$\min_{\mathbf{P}_k, k=1, \dots, K} \sum_{k=1}^K \left\| \Gamma \mathbf{P}_k^{(0)} - \mathbf{P}_k \right\|^2 \quad \text{s.t.} \quad \sum_{k=1}^K |\mathbf{F}_2 \mathbf{P}_k|^2 = \mathbf{1}_{L^2 \times 1} \quad (16)$$

Because the cost function is in a quadratic form, minimization amounts to a linear system of equations that can be solved using sparse solvers. This leads to the tedious computation of a large sparse system matrix and depending on the solver being used, numerical stability issues may arise. A simpler numerical algorithm to solve (16) can be obtained using the split-augmented Lagrangian method. This involves introducing splitting variables $\mathbf{R}_k = \mathbf{F}_2 \mathbf{P}_k$, leading to the following Lagrange function:

$$\sum_{k=1}^K \left\| \Gamma \mathbf{P}_k^{(0)} - \mathbf{P}_k \right\|^2 + \frac{\mu}{2} \sum_{k=1}^K \left\| \mathbf{F}_2^H \mathbf{R}_k - \mathbf{P}_k \right\|^2 + \sum_{k=1}^K \boldsymbol{\lambda}_k^T (\mathbf{F}_2^H \mathbf{R}_k - \mathbf{P}_k) + \frac{\mathbf{v}^T}{2} \left(\sum_{k=1}^K |\mathbf{R}_k|^2 - \mathbf{1}_{L^2 \times 1} \right) \quad (17)$$

Here, the second term is a penalization term with weight $\mu > 0$. $\boldsymbol{\lambda}_k$ and \mathbf{v} are Lagrange multipliers for $\mathbf{F}_2^H \mathbf{R}_k = \mathbf{P}_k$ and $\sum_{k=1}^K |\mathbf{R}_k|^2 = \mathbf{1}_{L^2 \times 1}$, respectively. In (17), we divided both μ and \mathbf{v} by 2 because this factor is then eliminated when computing the derivative of the Lagrange function. Obviously, this does not affect the solution of the problem in any sense. By differentiating and equating to zero, it can be verified that the solution of this Lagrangian optimization problem is found by Algorithm 1.

Algorithm 1 Split-augmented Lagrangian method for the joint spectral factorization of a set of squared frequency responses $|\mathbf{F}_2 \mathbf{P}_k|^2$, $k = 1, \dots, K$. ε is a sufficiently small number, e.g. $\varepsilon = 10^{-10}$.

Initialization $\mathbf{P}_k = \Gamma \mathbf{P}_k^{(0)}$, $\mathbf{R}_k = \mathbf{F}_2 \mathbf{P}_k$, $\boldsymbol{\lambda}_k = \mathbf{0}$, $\mu = 1$, $\varepsilon = 10^{-10}$

do

$$\mathbf{P}_k \leftarrow \frac{\Gamma \mathbf{P}_k^{(0)} + \mu \Gamma \mathbf{F}_2^H \mathbf{R}_k + \boldsymbol{\lambda}_k}{1 + \mu} \quad k = 1, \dots, K \quad \text{(update step 1)} \quad (18)$$

$$(\mathbf{R}_k)_m \leftarrow \frac{|(\mathbf{F}_2 \mathbf{P}_k)_m|}{\sqrt{\sum_{k'=1}^K |(\mathbf{F}_2 \mathbf{P}_{k'})_m|^2}} \quad k = 1, \dots, K, m = 1, \dots, L^2 \quad \text{(update step 2)} \quad (19)$$

$$\boldsymbol{\lambda}_k \leftarrow \boldsymbol{\lambda}_k + (\mathbf{F}_2^H \mathbf{R}_k - \mathbf{P}_k) \quad k = 1, \dots, K \quad \text{(update step 3)} \quad (20)$$

repeat until $\sum_{k=1}^K \left\| \mathbf{F}_2^H \mathbf{R}_k - \mathbf{P}_k \right\|^2 < \varepsilon$

This algorithm allows for an intuitive interpretation:

1. Update step 1: we start from \mathbf{R}_k (DFT domain) and go to the spatial domain ($\mathbf{F}_2^H \mathbf{R}_k$). We truncate the support of the filter, and average the result with the truncated impulse response of the prototype filter $\Gamma \mathbf{P}_k^{(0)}$. Next, we add a small correction $\lambda_k/(1 + \mu)$. By applying update step iteratively, we can interpret the loop as a recursive filter (that lets the result converge to a result that well approximates the prototype filters $\Gamma \mathbf{P}_k^{(0)}$).
2. Update step 2: we transform the result from update step 1 back to the DFT domain, and we discard the phase information of the DFT coefficients (by calculating the magnitude). Next, we normalize the resulting frequency responses such that they sum up to 1 and such that PR is ensured. By discarding the phase information, we force the filter impulse responses to be localized around the origin, which subsequently minimizes the potential energy lost by the projection $\Gamma \mathbf{P}_k$.
3. Update step 3: by update steps 1 and 2, \mathbf{R}_k is no longer equal to $\mathbf{F}_2 \mathbf{P}_k$. The variable λ_k keeps track of the current residual error between $\mathbf{F}_2^H \mathbf{R}_k$ and \mathbf{P}_k (update step 3). This step generally improves the convergence of the method.
4. The whole procedure is iterated until the error between the current solution \mathbf{P}_k and $\mathbf{F}_2^H \mathbf{R}_k$ is sufficiently small.

The compactly supported shearlet transform can then be implemented as in Figure 3:

1. Apply the horizontal directional filters $\mathbf{P}_k, k = 1, \dots, K/2$ to the image.
2. Apply the vertical directional filters $\mathbf{P}_k, k = K/2, \dots, K$ to the image.
3. Apply a 1D horizontal DT-CWT of J scales to the horizontally filtered image.
4. Apply a 1D vertical DT-CWT of J scales to the vertically filtered image.

The inverse transform can be computed by reversing the different steps. In this scheme, the resulting shearlet basis functions are complex-valued. The representation is a Parseval frame, because each step of the analysis algorithm preserves energy.

4. EXPERIMENTAL RESULTS AND DISCUSSION

Algorithm 1 allows to compute many directional filters with relatively large support (e.g. 16×16) based on some computationally simple update steps: the Fourier matrix multiplications \mathbf{F}_2^H and \mathbf{F}_2 can all be implemented efficiently using FFTs. Although the computation of the filters only has to be performed once, we found it interesting to investigate the number of iterations required to reach convergence. In Figure 4, we show the absolute reconstruction error (measured as

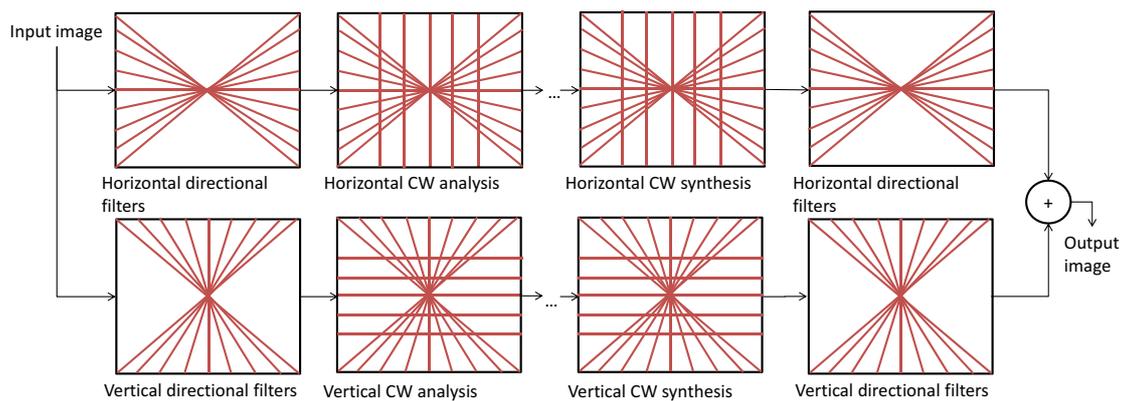


Figure 3. Implementation of our proposed compactly supported shearlet transform (CW=complex wavelet).

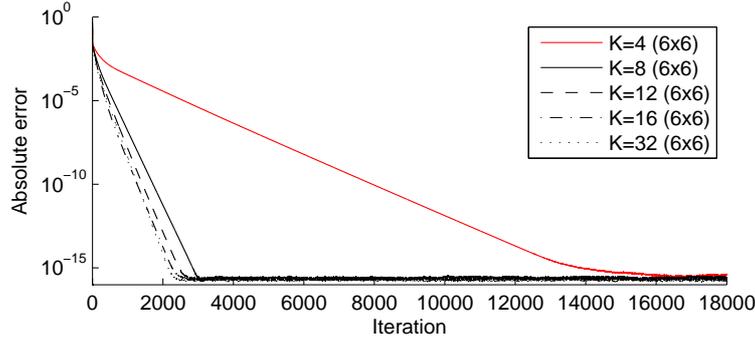


Figure 4. Convergence of the augmented Lagrangian algorithm for different number of orientations (K).

$\sum_{m=1}^{L^2} \left| \sum_{k=1}^K |\mathbf{R}_{k,m}|^2 - 1 \right|$) as function of the iteration index. Here, we use compact filters of support size 6×6 and $L = 24$. When considering eight or more orientations, we find that convergence is reached in about 3000 iterations. For $K = 4$, the convergence takes significantly longer (about 17000 iterations). Nevertheless, in Matlab this procedure takes no longer than a few seconds (up to one minute for $K = 4$). The absolute errors (around 10^{-16}) correspond to a reconstruction PSNR of around 300dB, which is due to the limited floating point precision. Also for other support sizes and other values of L and K we obtained perfect reconstruction (up to floating point precision).

In Figure 5, the resulting shearlet basis functions are depicted for $K = 12$ and $K = 24$. It can be noted that the compact shearlets resemble tiny edges and that they lack the long tail of the bandlimited shearlets or curvelets. It is apparent that the basis functions are sheared versions of each other, which is according to the shearlet definition (Section 3.1). The basis functions have an excellent spatial localization and directional selectivity, which is beneficial for practical applications.

To investigate the advantages of the compact support of the shearlets, we first perform a simple experiment using a grayscale line drawing of a rose (Figure 6(a)). Due to the curve-like discontinuities in this image, traditional transforms based on wavelets often have problems approximating these kind of images: typically many non-zero transform coefficients are needed to approximate the edge. In this experiment, we analyze the line drawing using different multiresolution transforms. We use only one single scale of the transform, and we select the 1% most significant coefficients (other non-significant coefficients as well as the scaling coefficients are set to zero). The 1% most significant coefficients are then further quantized to -1 and +1. Next, the image is reconstructed using the inverse transform and the intensity range is

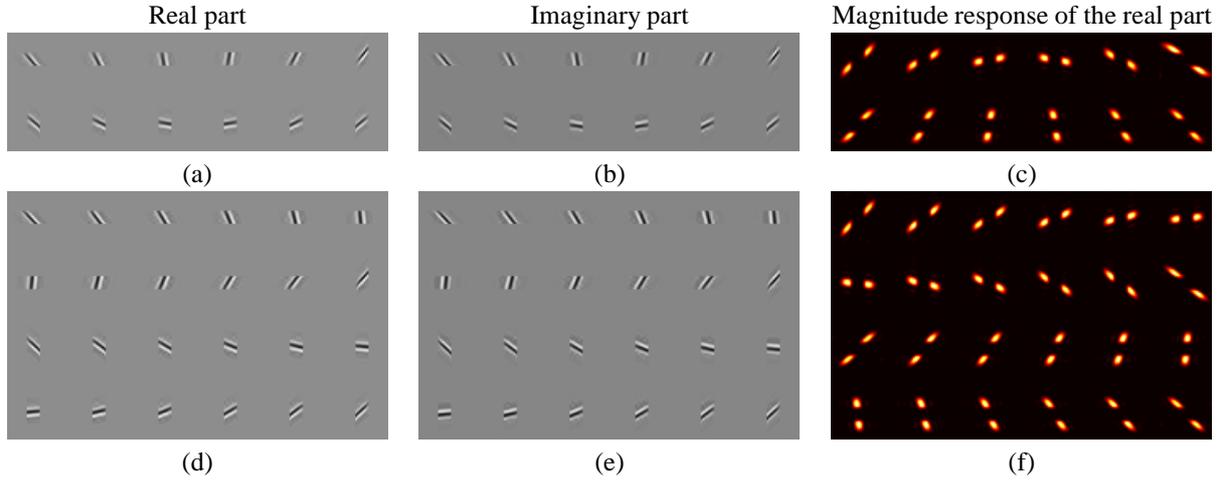


Figure 5. Compactly supported shearlet basis functions and their corresponding magnitude responses: (a)-(c) 12 orientations, (d)-(f) 24 orientations. For (c) and (f), the origin of frequency space is in the center of the image (in between the two blobs). The color conventions are as follows: (top row) gray corresponds to zero, white to positive values and gray to negative values. (bottom row) yellow: large positive magnitude, red: small positive magnitude, black: zero.

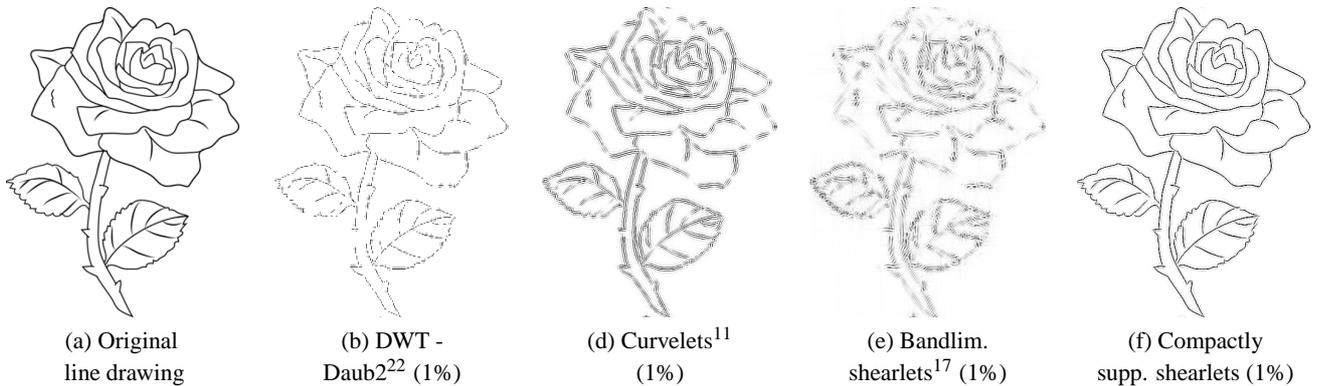


Figure 6. Approximation of the line drawing of a rose, using x-let coefficients with 3 possible values (-1, 0 and 1). Enclosed in parentheses is the relative number of x-let coefficients that are non-zero after processing.

rescaled to the range 0-255. The results are shown in Figure 6(b)-(f). It is clear that the DWT (Figure 6(b)) is not well able to approximate edges that are not horizontal or vertical. The curvelets[‡] (Figure 6(c)) and bandlimited shearlets introduce ringing in the reconstructed image, which is here amplified by the quantization to -1 and +1. It is remarkable that the compact shearlet transform (with $K = 8$ orientations) is well able to reconstruct the curves of the original drawing. One can argue that this is due to the high redundancy of our scheme (a factor 8), while other transforms such as curvelets and shearlets have a slightly lower redundancy (about 7 and 5.6 respectively). Nevertheless, the number of retained coefficients in Figure 6(f) was 20972, which is smaller than the original number of non-zero pixels (26502).

As this experiment suggests that our compactly supported shearlets are well able to approximate and localize small curve-discontinuities in images, we apply the different transforms to a charcoal drawing (see Figure 7), and we use the same procedure as above. Here, both the compactly supported shearlets from Kittipoom et al.¹⁹ as well as our method give a good result. Using our compact shearlet design (see Figure 7(d)), the delineation of the contours appears to be slightly better.

Next, we compare the different transforms in an image denoising application. We added artificial white Gaussian noise with standard deviation 25 to the line drawing from Figure 7(a). Then, we use the BLS-GSM method from Portilla et al.³¹ to remove noise from the images. We then compare the denoising performance for different multiresolution transforms, each using the same number of scales. Visual results are given in Figure 8. First, we remark that in this type of comparisons, the PSNR metric is actually not very useful, because this metric is not well able to capture ringing artifacts in images. Here, we provide PSNR values to illustrate that the transforms are within the same range of performance. Nevertheless, visually it can be noted that the curvelets and bandlimited shearlets cause a lot of ringing and cross-over artifacts. The cross-over artifacts appear due to the fact that long-tailed basis functions can not approximate well corner structures in images: other transform coefficients are needed to compensate this, which are suppressed to much by the denoising. As a result, line-structures appear that extend existing structures in the image. Using the compactly supported shearlets, the cross-over artifacts as well as the Gibbs-type ringing can be avoided. Compared to the compactly supported separable shearlets from Kittipoom et al.¹⁹ (Figure 8(e)), our denoised image in (Figure 8(f)) appears slightly sharper. This might be caused by the improved frame bounds. For other test images (e.g. lena, barbara, ...), the denoising performance is similar.

Finally, we remark that there are many extensions possible to the shearlet design procedure from Section 3. First, it is easy to impose symmetry relationships (such as the fact that the ‘horizontal’ shearlets are 90°-rotated versions of the ‘vertical’ shearlets). Next, due to the flexibility in defining the cost function (17), it is possible to include decimation operations in the optimization: in the current design, the redundancy factor of the transform is proportional to the number of analysis orientations K . By including decimations, the redundancy can further be reduced, in a similar way as we did

[‡]Here, we used the wrapping based curvelet implementation from Candès et al.¹¹

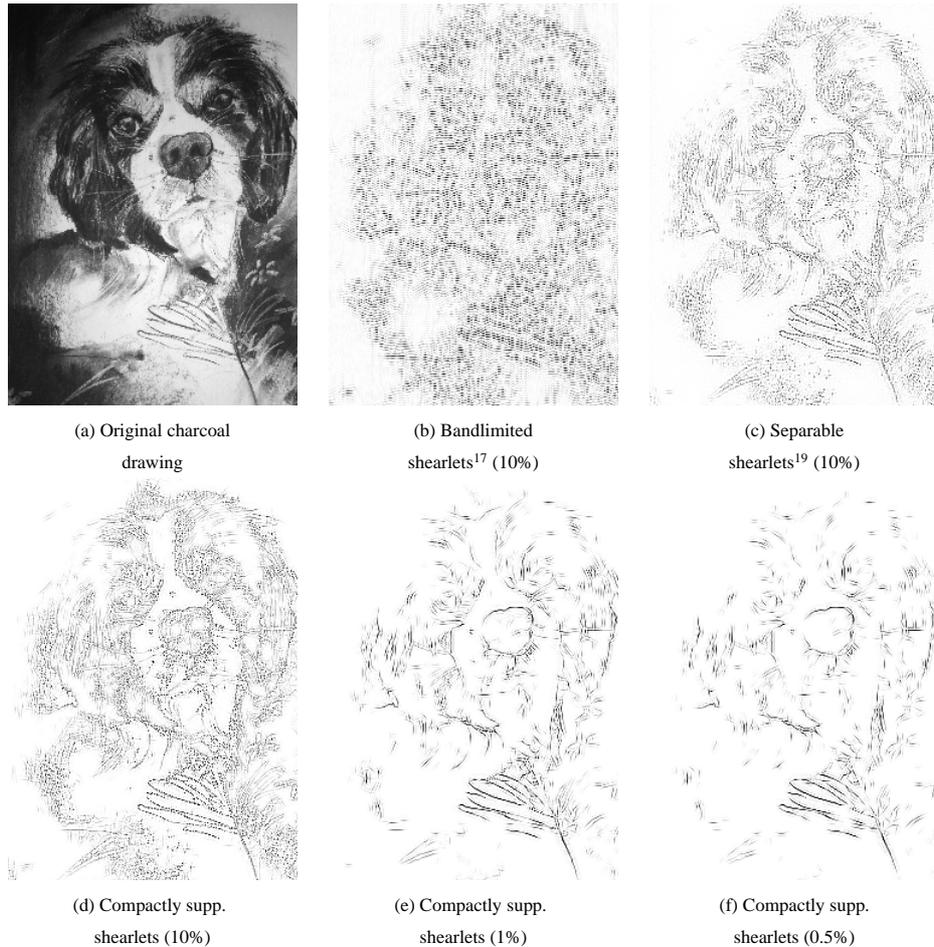


Figure 7. Obtaining a sketch from a charcoal drawing: comparison of different multiresolution transforms. Enclosed in parentheses is the relative number of x-let coefficients that are non-zero after processing. Original charcoal drawing: Liliane Goossens, *Dolly*, 2012.

for bandlimited shearlets in Ref.¹⁷ Our current results are encouraging: we are already able to construct a near-Parseval frame, with basis functions with good directional properties (exact details will be published elsewhere). Another extension is to define the shearlets in a separable way, similar to Kittipoom et al.,¹⁹ but in such a way that the separable supports are sheared versions of each other. This would be very useful for the processing of higher dimensional data (3D, 4D, ...), because the resulting algorithms are then based on separable filtering, yielding a computational performance that is similar to the DWT or the DT-CWT.

5. CONCLUSION

To tackle the design problem of tight shearlet frames with compactly supported basis functions, we first reviewed how directional selectivity is obtained using the Dual-Tree complex wavelet transform (DT-CWT), which possesses many of the desired properties (tight frame, basis functions with compact support). One disadvantage is that the DT-CWT does not enable wedge-like analysis filters in the frequency domain with arbitrary number of orientations. Therefore, we introduced an augmented split-Lagrangian algorithm, that allows us to compute compactly supported shearlet filters with an iterative algorithm requiring relatively simple computational steps. Experimental results reveal that the resulting shearlet basis functions resemble tiny edges and that these functions are particularly interesting for analyzing and synthesizing line drawings (even though they can still be used for the wide class of photographic images). We showed an example where we successfully extracted a pencil sketch from charcoal drawing. Finally, due to the flexibility of the design method, many future extensions become possible, such as decimation schemes for reducing the redundancy of the transform, as filtering schemes based on separable filters to decrease the computation time.

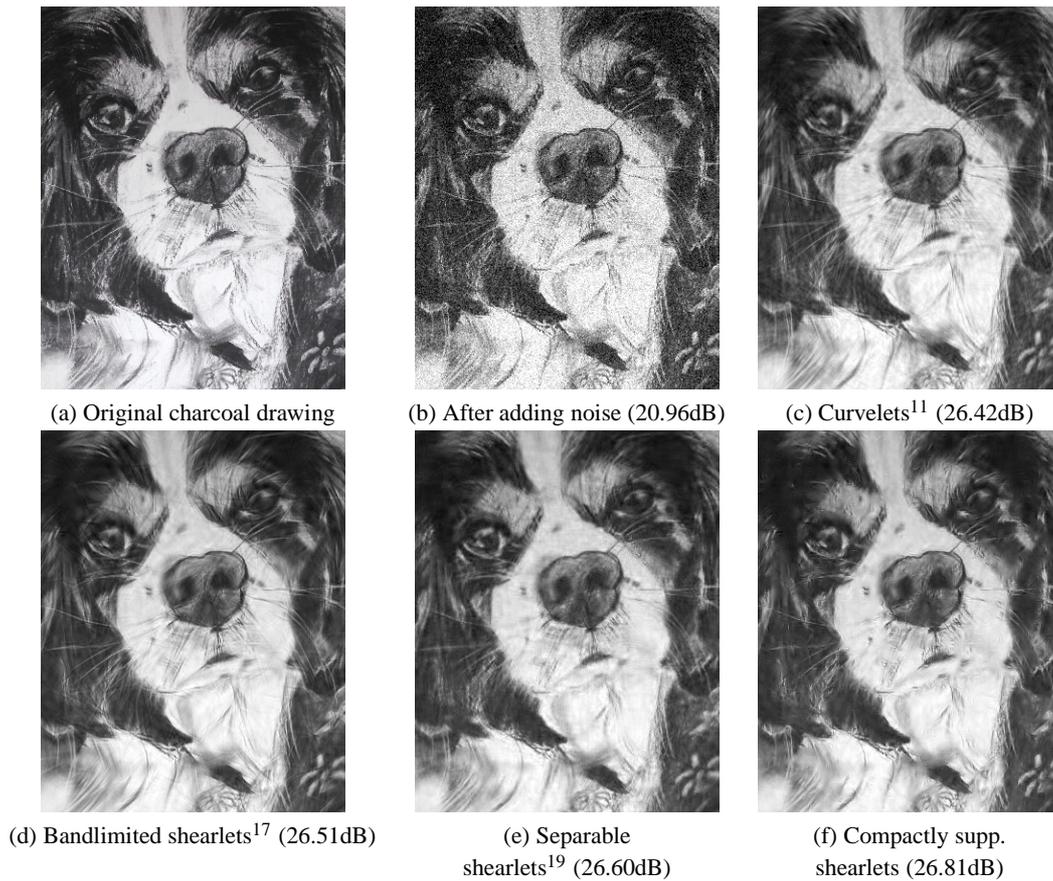


Figure 8. Image denoising experiment using BLS-GSM and different multiresolution transforms. Images are cropped for visualization purposes. Original charcoal drawing: Liliane Goossens, *Dolly*, 2012.

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. Lee, T., "Image Representation Using 2D Gabor Wavelets," *IEEE Trans. Pattern Analysis and Machine Intelligence* **18**, 1 (oct 1996).
4. 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).
5. Do, M. N. and Vetterli, M., "The contourlet transform: An efficient directional multiresolution image representation," *IEEE Trans. Image Process.* **14**(12), 2091–2106 (2005).
6. Candès, E., *Ridgelets: Theory and Applications*, PhD thesis, Departement of Statistics, Stanford University (1998).
7. Do, M. N. and Vetterli, M., "The finite ridgelet transform for image representation," *IEEE Trans. Image Processing* **12**, 16–28 (jan 2003).
8. 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).
9. LePennec, E. and Mallat, S., "Sparse geometric image representation with bandelets," *IEEE Trans. Image Processing* **14**, 423–438 (april 2005).
10. 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).

11. Candès, E., Demanet, L., Donoho, D., and Ying, L., “Fast Discrete Curvelet Transforms,” *Multiscale modeling and simulation* **5**(3), 861–899 (2006).
12. Gopinath, R. A., “The Phaselet Transform - An Integral Redundancy Nearly Shift-Invariant Wavelet Transform,” *IEEE Trans. Signal Processing* **51**, 1792–1805 (july 2003).
13. 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).
14. Lu, Y. M. and Do, M. N., “Multidimensional Directional Filter Banks and Surfacelets,” *IEEE Trans. Image Processing* **16**, 918–931 (april 2007).
15. Kutyniok, G. and Labate, D., “Construction of Regular and Irregular Shearlets,” *J. Wavelet Theory and Appl.* **1**, 1–10 (2007).
16. 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).
17. Goossens, B., Aelterman, J., Luong, H., Pižurica, A., and Philips, W., “Design of a tight frame of 2d shearlets based on a fast non-iterative analysis and synthesis algorithm,” in [*SPIE Optics and Photonics*], p. 81381Q (Aug. 2011).
18. Goossens, B., Luong, H., Aelterman, J., Pižurica, A., and Philips, W., “Efficient Multiscale and Multidirectional Representation of 3D Data using the 3D Discrete Shearlet Transform,” in [*IEEE Optics & Photonics 2011, Wavelets and Sparsity XIV*], p. 81381Z (Aug. 21-25 2011). Invited paper.
19. Kittipoom, P., Kutyniok, G., and Lim, W.-Q., “Construction of Compactly Supported Shearlet Frames,” *Constr. Approx.* **35**, 12–72 (2012).
20. Kutyniok, G., Lemvig, J., and Lim, W.-Q., “Optimally Sparse Approximation of 3D Functions by Compactly Supported Shearlet Frames,” *SIAM J Math. Anal.* **44**, 2962–3017 (2012).
21. Kingsbury, N. G., “Complex wavelets for shift invariant analysis and filtering of signals,” *Journal of Applied and Computational Harmonic Analysis* **10**, 234–253 (may 2001).
22. Daubechies, I., [*Ten Lectures on Wavelets*], Society for Industrial and Applied Mathematics, Philadelphia (1992).
23. 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).
24. Belzer, B., Lina, J.-M., and Villasenor, J., “Complex, linear-phase filters for efficient image coding,” *IEEE Trans. Signal Process.* **43**, 2425–2427 (Oct. 1995).
25. Lina, J.-M. and Mayrand, M., “Complex Daubechies wavelets,” *Appl. Comp. Harm. Analysis* **2**(3), 219–229 (1995).
26. Guo, K. and Labate, D., “Optimally Sparse Multidimensional Representation using Shearlets,” *SIAM J Math. Anal.* **39**, 298–318 (2007).
27. 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).
28. 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).
29. Fomel, S., Sava, P., Rickett, J., and Claerbout, J., “The wilson-burg method of spectral factorization with application to helical filtering,” *Geophysical Prospecting* **51**, 409–420 (2003).
30. Norell, B., Burkadov, O., Andersson, M., and Knutsson, H., “Approximate spectral factorization for design of efficient sub-filter sequences,” tech. rep., Technical Report LiTH-MAT-R-2011/14-SE, Department of Mathematics, Linköping University (2011).
31. 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).