

Discrete directional wavelet bases and frames for image compression and denoising

Pier Luigi Dragotti¹, Vladan Velisavljević², Martin Vetterli^{2,3} and Baltasar Beferull-Lozano²

¹ Electrical and Electronic Engineering
Imperial College, Exhibition Road,
London SW7 2BT, United Kingdom

² École Polytechnique Fédérale de Lausanne (EPFL)
CH - 1015 Lausanne, Switzerland.

³ Electrical Engineering and Computer Science Department,
University of California at Berkeley, Berkeley CA 94720

ABSTRACT

The application of the wavelet transform in image processing is most frequently based on a separable construction. Lines and columns in an image are treated independently and the basis functions are simply products of the corresponding one dimensional functions. Such method keeps simplicity in design and computation, but is not capable of capturing properly all the properties of an image. In this paper, a new truly separable discrete multi-directional transform is proposed with a subsampling method based on lattice theory. Alternatively, the subsampling can be omitted and this leads to a multi-directional frame. This transform can be applied in many areas like denoising, non-linear approximation and compression. The results on non-linear approximation and denoising show interesting gains compared to the standard two-dimensional analysis.

Keywords: Wavelets, denoising, Non-linear approximation

1. INTRODUCTION

At the heart of many image processing tasks, there is an efficient image representation that can capture significant information of an image in a compact form. Over the last decade or so, the wavelet transform has emerged as the dominating tool in image processing. The success of the wavelet transform is mainly due to its ability to characterize certain classes of signals with few transform coefficients.¹¹ In particular, wavelets as time-frequency localized bases are particularly suited to represent one dimensional (1-D) piecewise smooth functions.¹⁴ Piecewise smooth signals are often encountered in practice (for example, one scan-line of an image).

In two dimensions (2-D), the situation is more open. Despite their widespread success in image processing, wavelets are not adequate at representing images. In images, discontinuities are straight line or smooth curves. The 2-D wavelet transform is a separable transform given by the tensor-product of two 1-D wavelets along the horizontal and vertical directions. For this reason, this separable transform is good at isolating horizontal and vertical edges, but it is not adequate at treating more complex discontinuities. This indicates that more powerful image representations are needed.

A key insight on the construction of more efficient bases for image representations was provided by Candés and Donoho.^{3,4} They introduced the *ridgelet* and the *curvelet* transforms. The ridgelet transform is efficient at representing 2-D functions made of two smooth regions divided by a straight line. The curvelet transform can deal with more complex discontinuities. The curvelet transform is obtained by filtering the image and then by applying a block ridgelet transform on each bandpass image. Both these transforms show the potential of non-separable methods, but come at a price in terms of design and computational complexity. Indeed, ridgelets and curvelets are defined in continuous spaces (e.g., in \mathbb{R}^2) and the construction of equivalent transforms in discrete spaces is not trivial. For instance, the implementations of the curvelet transform proposed in^{9,16} are very redundant, therefore, of scarce interest in applications such as compression. Recall that in a redundant transform, the number of transform coefficients is larger than the number of samples in the original image.

Other non-separable approaches, such as directional filter banks^{1,8} or steerable pyramid¹⁵ suffer from similar drawbacks. Some separable approaches have been made in¹⁹ but not on discrete spaces and the directional transform proposed in¹² is again overcomplete therefore of less interest for compression. Finally, the bandelet representation proposed by Mallat and LePennec¹³ and the edge adapted multiscale transform of Cohen and Matei⁵ require an edge detection step followed by an adaptive transform, for this reason these methods tend to be complex.

In our approach, we wish to retain the simplicity of the separable wavelet transform, while realizing some of the potential of non-separable schemes. The focus is on discrete framework that can lead to algorithmic implementations. We introduce a discrete directional wavelet transform that acts much like a standard separable transform but allows more directionalities. This can be done by introducing digital directions in discrete space. Along these directions, it is then possible to apply the wavelet transform or the wavelet frame decomposition.

While the design of multi-directional frames is pretty straightforward,² the construction of multi-directional bases having good representation properties presents a more difficult challenge due to the subsampling.

In this paper, we use a lattice theory based method to define a new convenient directional subsampling method. The proposed method considers each intersection of the digital lines and cosets produced by shifted versions of a lattice as independent transform lines (called *co-lines*). The transform is still separable and it allows many (much more than only two) transform directions. Moreover, the subsampling issue is solved simply and clearly for a general combination of different angles.

The outline of the paper is as follows: In the next section, we introduce the notion of directional and multi-directional transforms and presents some examples of multi-directional wavelet frames. Section 3 focuses on the construction of multi-directional wavelet bases. In this section we show how the subsampling problem is solved using lattice theory. In Section 4, we present some possible applications of this new transform. In particular, we concentrate on non-linear approximation and denoising. Finally, we conclude in Section 5.

2. DIRECTIONAL TRANSFORMS AND FRAMES

The standard separable wavelet transform simply consists of two 1-D wavelet transforms along the horizontal and vertical directions. Subsampling is first performed on columns and then on rows. The process is then iterated on the transformed image. Notice that one can shift out the subsamplers in the decomposition obtaining in this way a redundant transform. The directional transform we propose is still separable but we allow it to be taken along more directions.

As a first step, we need to define directions in discrete spaces. We follow the definition of a discrete line given in.² The line is determined by its slope and shift by the following equation:

$$y[n] = \lfloor kx[n] \rfloor + \lfloor B \rfloor \quad (1)$$

where $k = \tan(\theta)$ represents the slope and belongs to the range $0 \leq k \leq 1$, and B represents the real-valued shift parameter. This definition of a discrete line insures that each pixel in an image belongs exactly to one line for a chosen slope. Lines with slopes out of the range may be obtained by symmetry or by rotating the space. This gives access to a wealth of directions in \mathbb{Z}^2 .

In our approach, one is free to choose any of the digital directions given in Eq. (1) and to apply the 1-D wavelet transform (or any other one dimensional transform) along that direction. Moreover, the process can be iterated and at each iteration one can choose a new direction. In particular, a one-directional wavelet transform produces two subbands, where the high-pass subband does not contain any smooth object along the transform direction. Then, one can iterate the process on one or on both subbands and can choose a different direction. The process can be iterated several times and this leads to a multi-resolution decompositions along multiple directions. Notice that according to the chosen 1-D transform and to the kind of iteration, one can obtain a wide range of directional transforms that span from orthogonal/ biorthogonal transforms to frames. In the case of multi-directional wavelet bases, the challenging part is in the subsampling step and this will be discussed in more details in the next section.

Alternatively, one can avoid the subsampling by applying a 1-D wavelet frame on each chosen direction. This leads to a multi-directional frame that can be used for instance as an efficient direction discriminator. This is conceptually shown in Figure 1. Also, one can apply the 1-D wavelet transform along one direction

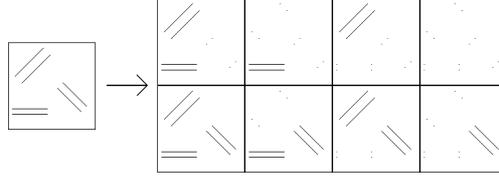


Figure 1. An example of an iterated directional transform with three directions equal to 0° , 45° and 135° . Some of the transformed images are omitted.

and then, rather than iterate the process on the transformed image, one can chose a new direction and apply the corresponding directional wavelet transform on the original image again. The block diagram in Figure 2 illustrates this second kind of directional redundant transform. In Section 4.2, we show that this transform is very useful in image denoising.

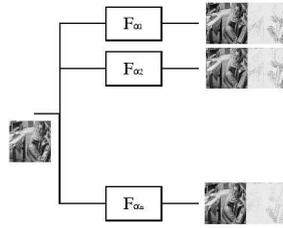


Figure 2. An example of a multi-directional wavelet frames. Each box represents a wavelet transform along the direction α_i , the union of all these transforms gives the directional wavelet frame.

An interesting property of these transforms is the following^{17,18}:

PROPOSITION 1. *Directional or multi-directional transforms when built on top of one-dimensional orthogonal transforms or tight frames, again lead to orthogonal transforms or frames.*

3. DIRECTIONAL AND MULTI-DIRECTIONAL BASES

To design the multi-directional wavelet basis, we need to use the concept of integer lattices. The subsampling procedure will be explained in terms of these lattices.

A full rank integer lattice Λ consists of the set of points obtained by taking linear combinations of two linearly independent vectors where both the components of the vectors and the coefficients are integers. Any integer lattice Λ is always a sublattice of the ordinary cubic integer lattice \mathbb{Z}^2 , that is $\Lambda \subset \mathbb{Z}^2$, and can be represented by a (non-unique) generator matrix:

$$M_\Lambda = \begin{pmatrix} s_{11} & s_{12} \\ s_{21} & s_{22} \end{pmatrix} = \begin{pmatrix} \vec{v}_1 \\ \vec{v}_2 \end{pmatrix}. \quad (2)$$

The corresponding lattice Λ specified by the generator matrix M_Λ is given by

$$\Lambda = \{x : x = u_1\vec{v}_1 + u_2\vec{v}_2, u_i \in \mathbb{Z}, i = 1, 2\}.$$

An example with $M_\Lambda = \begin{pmatrix} 2 & 1 \\ 0 & 1 \end{pmatrix}$ is shown in Figure 3(a).

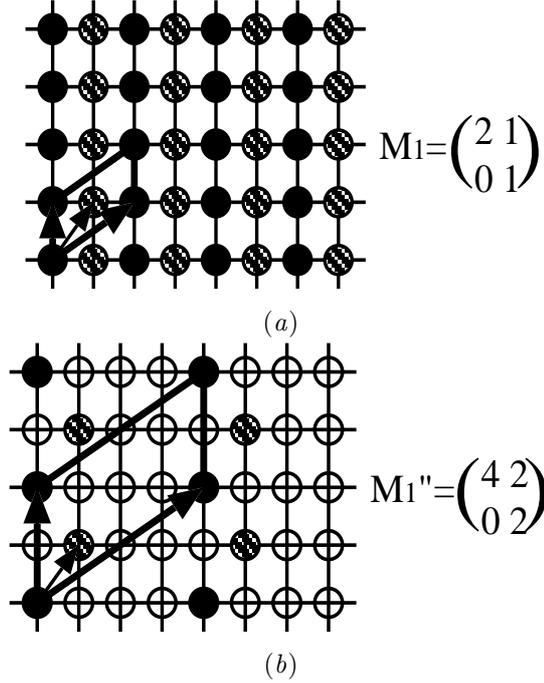


Figure 3. (a) The initial cubic lattice is partitioned in two cosets of the lattice defined by M_1 corresponding to shift vectors $\vec{c}_0 = (0,0)$ and $\vec{c}_1 = (1,1)$, (b) the subsampled version is given by two cosets of the lattice defined by M_1'' (sublattice of M_1) with the same shift vectors.

It can be shown from lattice theory⁷ that given an integer Λ with a generator matrix having determinant $\det(M_\Lambda)$, the cubic lattice \mathbb{Z}^2 is partitioned into $\det(M_\Lambda)$ cosets $B_{\Lambda,c_j} = \Lambda + c_j, j = 0, \dots, |\det(M_\Lambda)| - 1$, which are shifted versions of the sublattice Λ .

In addition, call a set of pixels that belong to a digital line described by (1) as $L_{r,b}$. A co-line is defined as the intersection of a coset and a digital line:

$$\tilde{L}_{\Lambda,c_j,r,b} = L_{r,b} \cap B_{\Lambda,c_j}, r \in \mathbb{Q} \quad (3)$$

If M_Λ is a generator matrix for Λ , then all the possible generator matrices for Λ are given by UM_Λ where U is a unimodular matrix, that is, a matrix with integer components and $|\det(U)| = 1$.

Although in terms of lattice representation all generator matrices are equivalent, in terms of the way the sampling procedure is performed, the row vectors of each particular generator matrix can be associated with the digital directions along which the 1D filterings and subsampling will take place. Consider a lattice Λ_1 generated by a matrix $M_{\Lambda_1} = \begin{pmatrix} s_{11}^{(1)} & s_{12}^{(1)} \\ s_{21}^{(1)} & s_{22}^{(1)} \end{pmatrix}$. The slopes of the transform directions are defined as $r_1 = s_{12}^{(1)}/s_{11}^{(1)}$ and $r_2 = s_{22}^{(1)}/s_{21}^{(1)}$.

First we apply the one-dimensional wavelet transform along co-lines with slope r_1 in all cosets, $\{\tilde{L}_{\Lambda_1,c_j,r_1,b}\}$, where $j = 0, \dots, |\det M_{\Lambda_1}| - 1, b \in \mathbb{Z}$. Subsampling along that direction is done independently in each coset and the set of kept points belongs to a sublattice defined by a new generator matrix $M'_{\Lambda_1} = \begin{pmatrix} 2\vec{v}_1 \\ \vec{v}_2 \end{pmatrix}$. Discarding each second sample along each transform co-line secures a valid subsampling in the sense that perfect reconstruction condition is satisfied. The process is continued in the same way along the second slope r_2 . The generator matrix that describes a sublattice of pixels that survived the two sampling procedures is given by $M''_{\Lambda_1} = \begin{pmatrix} 2\vec{v}_1 \\ 2\vec{v}_2 \end{pmatrix} =$

$2M_{\Lambda_1}$. The corresponding lattice Λ_1'' is clearly a sublattice of the initial one containing a quarter of the samples. The corresponding example is shown in Figure 3(b).

The matrix $2M_{\Lambda_1}$ produces four times more cosets. However, only the initial cosets should be kept. The other three-quarters of the cosets represent the discarded samples.

The next iteration is obtained by using a lattice Λ_2 which is in general a sublattice of Λ_1'' , that is $\Lambda_2 \subseteq \Lambda_1''$. This means that any generator matrix for Λ_2 is given by:

$$M_{\Lambda_2} = T_1 \cdot M_{\Lambda_1}'', T_1 = \begin{pmatrix} t_{11}^{(1)} & t_{12}^{(1)} \\ t_{21}^{(1)} & t_{22}^{(1)} \end{pmatrix}, t_{ij}^{(1)} \in \mathbb{Z} \quad (4)$$

The new pair of tangent coefficients is:

$$r_3 = \frac{t_{11}^{(1)} s_{12}^{(1)} + t_{12}^{(1)} s_{22}^{(1)}}{t_{11}^{(1)} s_{11}^{(1)} + t_{12}^{(1)} s_{21}^{(1)}} \text{ and } r_4 = \frac{t_{21}^{(1)} s_{12}^{(1)} + t_{22}^{(1)} s_{22}^{(1)}}{t_{21}^{(1)} s_{11}^{(1)} + t_{22}^{(1)} s_{21}^{(1)}}.$$

In the case $|\det(T_1)| = 1$, then $\Lambda_2 = \Lambda_1''$ and sampling along the digital lines with slopes r_3 and r_4 results simply in a different resampling of Λ_1'' , which we call redirection step. After the subsampling, we will obtain simply $\Lambda_2'' = 2\Lambda_1''$.

If the resampling is not unimodular, (4) yields more cosets. Then, each coset that survived the previous subsampling step is divided into $|\det T_1|$ new cosets.

In total there are $|\det M_{\Lambda_2}|/4$ cosets that are to be processed independently in the second step. The step involves two 1D filterings and subsamplings equivalently as in the first step, but along the new angles.

The iteration can be continued in a similar way as many times as desired. The redirection step can always be applied on a subsampled version of the initial generator matrix M_{Λ_1} . Indeed, consider the change from M_{Λ_2}'' to M_{Λ_3} . We have that:

$$M_{\Lambda_3} = C \cdot M_{\Lambda_2}'' = 2C \cdot M_{\Lambda_2} = 4C \cdot T_1 \cdot M_{\Lambda_1} = T_2 \cdot 4M_{\Lambda_1}$$

where C is an integer matrix and $T_2 = C \cdot T_1$. In general the following holds:

$$M_{\Lambda_{i+1}} = T_i \cdot 2^i M_{\Lambda_1}. \quad (5)$$

The tangent coefficients after each change are given by:

$$r_{2i-1} = \frac{t_{11}^{(i)} s_{12}^{(1)} + t_{12}^{(i)} s_{22}^{(1)}}{t_{11}^{(i)} s_{11}^{(1)} + t_{12}^{(i)} s_{21}^{(1)}}, r_{2i} = \frac{t_{21}^{(i)} s_{12}^{(1)} + t_{22}^{(i)} s_{22}^{(1)}}{t_{21}^{(i)} s_{11}^{(1)} + t_{22}^{(i)} s_{21}^{(1)}} \quad (6)$$

Using a concatenation of unimodular redirection steps is desirable in non-linear approximation of images, but this constrains the choice of directions because we need to have that:

$$|\det T_i| = |t_{11}^{(i)} t_{22}^{(i)} - t_{12}^{(i)} t_{21}^{(i)}| = 1 \quad (7)$$

However (6) and (7) together still allow a number of directions to be used. The following proposition gives the good approximation property of our construction.

PROPOSITION 2. *Let $\{M_{\Lambda_i}\}_{i=1}^K$ be generator matrices of a set of lattices $\{\Lambda_i\}_{i=1}^K$ which satisfy the nesting property explained above. If 1D filterings and subsamplings are applied along the directions contained in the matrices $\{M_{\Lambda_i}\}_{i=1}^K$, then: a) perfect reconstruction can be achieved and b) concatenated operations of filtering and subsampling along the different directions do not create directional interactions, that is, the samples that are kept at each iteration are aligned along the next filtering and subsampling directions.*

This avoidance of directional interaction is the crucial property to achieve good non-linear approximation.

4. APPLICATIONS

4.1. Non-linear approximation

It is widely accepted that a basis with good approximation properties is necessary to develop efficient compression algorithms. If the dominant directions in an image are the horizontal and vertical directions, then the traditional wavelet transform performs a good non-linear approximation of the image. However, it fails when other directions are present in the image. The multi-directional basis we have presented in the previous section can efficiently approximate a larger class of images. This is stated more precisely in the following proposition¹⁷:

PROPOSITION 3. *Consider a bilevel polygonal image of size $N \times N$ containing information along D different directions. The best wavelet decomposition leads to a number of non-zero coefficients of the order $O(kN)$, where k is the number of lines in the image. The best decomposition with the multi-directional wavelet basis leads to a number of non-zero coefficients of the order $O(lk \log_2^D N)$, where l is the total number of cosets in the decomposition.*

Consider, for example, the two synthetic images given in Figure 4. The traditional wavelet transform performs well in the first case and it is easy to see that the number of non-zero coefficients using a Haar wavelet and a correct decomposition of the image is smaller than $4 \log_2^2 N$. In the second case, however, the standard wavelet transform fails to represent the discontinuity at 45° and it attains a poor approximation ($O(N)$).

The multi-directional wavelet performs well in both cases. In particular in the second case, one can use a lattice generated by the matrix $M_\Lambda = \begin{pmatrix} 2 & 0 \\ 1 & 1 \end{pmatrix}$, this gives two co-lines and an approximation of the order of $8 \log_2^2 N$.

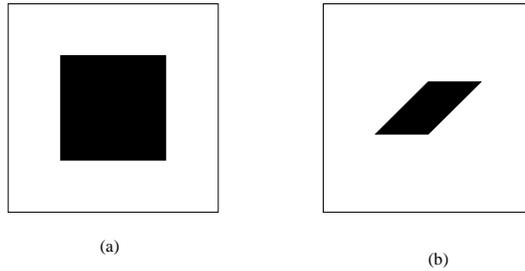


Figure 4. Two simple synthetic images.

A numerical comparison between the multi-directional approximation and the standard one is made on a synthetic image shown in Figure 6(a). The proposed directional transform performs largely better than the standard one as shown in Figure 6(b). The multi-directional method uses a wavelet decomposition along 0° , 45° , 90° and 135° and some elements of the corresponding directional basis are shown in Figure 5.



Figure 5. Some multi-directional scaling functions for an iteration of 45° and 135° . (a,b) An equal number of 45° and 135° steps, (c) more steps along 45° , (d) more steps along 135° .



Figure 6. (a) The original synthetic image, (b) The results of non-linear approximation by the standard (the dotted line) and the multi-directional method (the full line).

4.2. Denoising

In images, noise is usually removed by thresholding the coefficients obtained by the standard wavelet transform.¹⁰ The main advantage of this approach is its simplicity. In our method, we threshold the coefficients obtained by a directional wavelet frame. Our algorithm is in spirit similar to cycle-spinning.⁶

We consider several different directions, for each direction we apply the corresponding directional wavelet transform, hard-threshold the wavelet coefficients and invert the wavelet transform to obtain a denoised version of the image. Finally, the different denoised versions of the image are averaged to obtain the final image. Figure 7 shows with a block diagram the proposed algorithm.

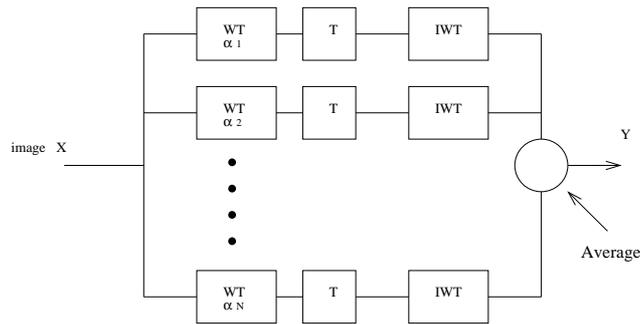


Figure 7. Denoising with directional wavelet frames. The coefficients of each directional transform are thresholded, the transform is inverted and the resulting denoised images are averaged.

Figure 8 shows an example of denoising of the image 'Cameraman'. The original image is affected by additive Gaussian noise and both methods are applied. We can see that the new method which uses in this case forty different directions largely outperform the classical algorithm.

5. CONCLUSIONS

Traditional image decomposition with wavelets involves only two directions. The method we have proposed in this paper goes beyond this limitation and calculates separable wavelet transform along a set of different directions. Simplicity and low computational complexity are still maintained. This new transform can be used in many image processing applications. Very promising results have been shown in image denoising and non-linear approximation.

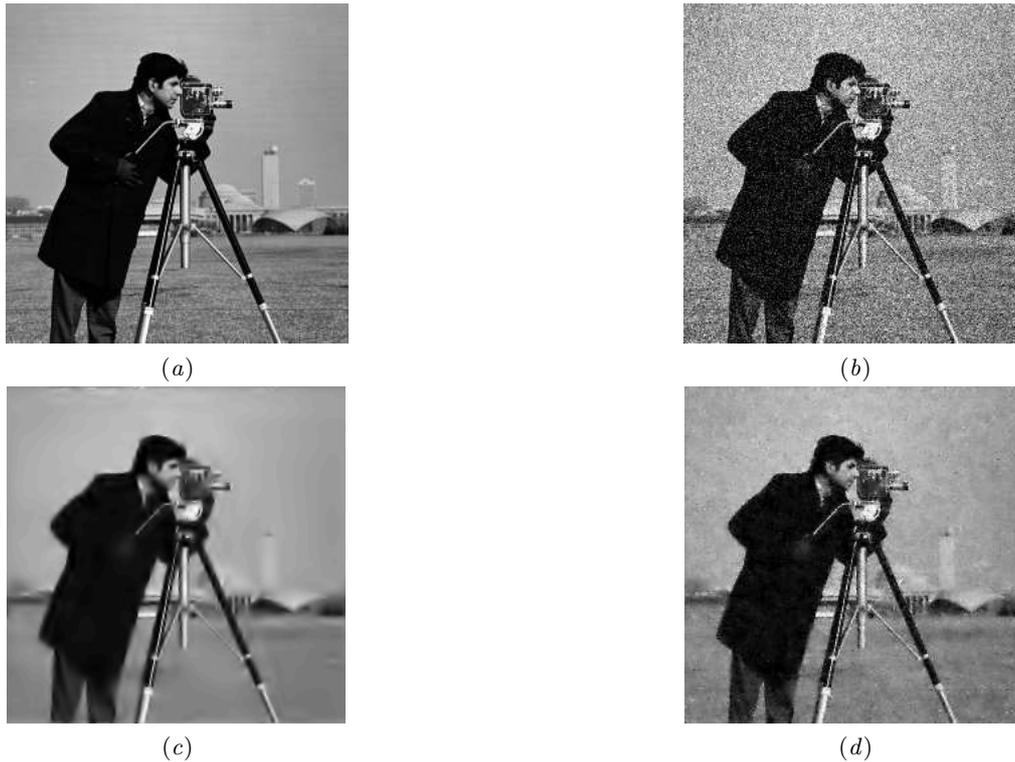


Figure 8. Denoising of the image 'Cameraman'. (a) Original image, (b) noisy version (PSNR=19dB), (c) the denoised image with standard undecimated wavelet transform (PSNR=22.46dB), (d) denoise with the directional wavelet frame (PSNR=25.15dB).

REFERENCES

1. R.H. Bamberg and M.J.T. Smith. A filter bank for the directional decomposition of images: Theory and design. *IEEE Trans. Signal Processing*, pages 882–893, April 1992.
2. J.E. Bresenham. Algorithms for computer control of a digital plotter. *IBM Systems Journals*, 4(1):25–30, 1965.
3. E. J. Candès and D. L. Donoho. Curvelets – a suprisingly effective nonadaptive representation for objects with edges. In A. Cohen, C. Rabut, and L. L. Schumaker, editors, *Curve and Surface Fitting*, Saint-Malo, 1999. Vanderbilt University Press.
4. E. J. Candès and D. L. Donoho. Ridgelets: a key to higher-dimensional intermittency? In *Phil. Trans. Royal Society London*, 1999.
5. A. Cohen and B. Matei. Compact representation of images by edge adapted multiscale transforms. In *IEEE Int. Conf. Image Processing*, Vancouver, Canada, September 2000.
6. R. R. Coifman and D.L. Donoho. Translation invariant denoising. *Technical Report 475, Dept. of Statistics, Stanford University*, May 1995.
7. J.H. Conway and N.J.A. Sloane. *Sphere packing, lattices and groups*. Springer-Verlag, 1998.
8. Minh Do. *Directional Multiresolution Image Representations*. PhD, dissertation, Swiss Federal Institute of Technology, November 2001.
9. D.L. Donoho and M.R. Duncan. Digital curvelet transform: strategy, implementation and experiments. In *Proc. Aerosense 2000, Wavelet Applications VII*, pages 12–29, 2000.
10. D.L. Donoho and I.M. Johnstone. Ideal spatial adaptation via wavelet shrinkage,. *Biometrika*, 81:425–455, December 1994.
11. D.L. Donoho, M. Vetterli, R.A. DeVore, and I. Daubechies. Data compression and harmonic analysis. *IEEE Trans. on Information Theory*, 44(6):2435–2476, October 1998.

12. N. Kingsbury. Complex wavelets for shift invariant analysis and filtering of signals. *Journal of Applied and Computational Harmonic Analysis*, 10:234–253, May 2001.
13. E. LePennec and S. Mallat. Image compression with geometrical wavelets. In *IEEE Int. Conf. Image Processing*, Vancouver, Canada, September 2000.
14. S. Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, 1998.
15. E.P. Simoncelli, W.T. Freeman, E.H. Adelson, and D.J. Heeger. Shiftable multiscale transforms. *IEEE Trans. on Information Theory*, 38(2):587–607, March 1992.
16. J.L. Starck, E.J.. Candés, and D.L. Donoho. The curvelet transform for image denoising. *IEEE Trans. Image Processing*, pages 670–684, June 2002.
17. V. Velisavljević, B. Beferull-Lozano, M. Vetterli, and P.L. Dragotti. Discrete multi-directional wavelet bases. *Technical Report*. In preparation.
18. V. Velisavljević, P.L. Dragotti, and M. Vetterli. Directional wavelet transforms and frames. In *IEEE Int. Conf. on Image Processing*, Rochester, NY, September 2002.
19. R.A. Zuidwijk. Directional and time-scale wavelet analysis. *SIAM Journal on Mathematical Analysis*, 31(2):416–430, 2000.