Directional Wavelet Analysis with Fourier-type Bases for Image Processing

Zhen Yao, *Student Member, IEEE* , Nasir Rajpoot, *Member, IEEE* and Roland Wilson

***Abstract*— Motivated by the fact that in natural images, there is usually a presence of local strongly oriented features such as directional textures and linear discontinuities, a representation which is both well-localised in frequency and orientation is desirable to efficiently describe those oriented features. Here we introduce a family of multiscale trigonometric bases for image processing using Fourier-type constructions, namely, the multiscale directional cosine transform and the multiscale Fourier transform. We also show that by seeking an adaptive basis locally, the proposed bases are able to capture both oriented harmonics as well as discontinuities, although the complexity of such adaptiveness varies significantly. We conducted denoising experiments with the proposed bases and the results show great promise of the proposed directional wavelet bases.**

***Index Terms*— Directional wavelets, curvelets, Fourier transform, cosine transform, denoising, restoration.**

I. Introduction

THE application of transforms in image processing is often based on a separable construction. Rows and columns in an image are treated independently and the two-dimensional basis functions are simply tensor products of the corresponding one-dimensional functions. Such method keeps simplicity in terms of design and computation, but is not capable of capturing properly all the interesting features of an image. For example, the orthonormal separable wavelet transform [27] in higher dimensions is seriously limited in its ability to efficiently represent higher dimensional features such as lines. Furthermore, the lack of frequency selectivity remains an elusive problem with most techniques operating in the wavelet domain.

Edges and textures in an image can exist at all possible locations, orientations, and scales. The ability to efficiently analyse and describe directional patterns is thus of fundamental importance for image analysis and image compression. The idea that biological visual systems might analyse image along dimensions such as orientation, scale and frequency (ie. bandpass) dates back to the work by Hubel and Wiesel [25] in the 1960's. In the computational vision literature, the idea of analysing images along multiple orientations appears at the beginning of the seventies with the Binford-Horn's line-finder [2], [23] and later work by Granlund [22]. Many edge-based image representations have then been elaborated [20], [42] with different edge detection procedures and image approximations using jump models along these edges. To refine these models, multiscale edge representations using wavelet maxima [29] or an edge-adapted multiresolution representation [10] have also been studied. Edge based image representations with complete orthonormal families of foveal wavelets in [32] and footprints [19] have been introduced and studied to reconstruct the main image edge structures. To stabilize the edge detection, global optimization procedures have also been elaborated by Donoho [17], Shukla et al. [36] and Wakin et al. [40]. The optimal configuration of edges is then calculated with an image segmentation over dyadic squares using fast dynamic programming algorithms over quadtrees. Instead of describing the image geometry through edges, which are most often ill-defined, Le Pennec and Mallat later proposed a basis named *"bandelets"* [33] which characterises the image geometry with a geometric flow of vectors. Recently, Peyré and Mallat presented the *second generation bandelets* [34]. The decomposition is computed first by the standard wavelet transform, followed by adaptive geometric orthogonal filters. The compression results are significantly better than wavelet-based coders.

All the approaches previously discussed are *adaptive* representations, in the sense that the bases are adapted to the signal/image contents. Meanwhile, from a different heuristic principle, a number of researchers have been working on developing *fixed* directional representation bases for natural images. The idea of curvelets [7] is to represent a curve as a superposition of functions of various lengths and widths obeying the scaling law $width \approx length^2$. Several different methods were proposed to construct the curvelets. A digital implementation for the curvelet transform, more commonly referred as the *curvelet-99* was used in [37] for noise removal by Starck et al. The transform first decomposes the image into subbands, i.e., separating the object into a series of disjoint scales, using the *algorithme á trous* wavelet transform. Each scale is then analysed by means of a local windowed ridgelet [6] transform. The proposed transform is $16J+1$ times redundant, with J being the number of scales for decomposition. The same authors later proposed a combined approach with curvelets and wavelets in denoising [38]. Such joint sparse representation idea is related to the idea of *Matching Pursuit* (MP) and *Basis Pursuit* (BP) [28], and another application in image deconvolution was presented in [39].

While the redundancy certainly is a advantage in the area of image restoration, it is by no means an ideal transform for compression and other tasks. In order to construct a form of discrete curvelet frame with less redundancy, Do and

Z. Yao is with the Computer Science Department, University of Warwick, CV4 7AL UK. E-mail: yaozhen@ieee.org

N. Rajpoot is with the Computer Science Department, University of Warwick, CV4 7AL UK. E-mail: nasir@dcs.warwick.ac.uk

R. Wilson is with the Computer Science Department, University of Warwick, CV4 7AL UK. E-mail: rgw@dcs.warwick.ac.uk

Vetterli [15], [16] pioneered the *"contourlet"* transform by marrying the Laplacian pyramid and a directional filter bank. Such approach is called *"double filter bank"* structure. The Laplacian pyramid mainly is used for separating isotropic features into different resolutions, then the directional filter links the point discontinuities into linear structures. This allows contourlets to efficiently approximate a smooth contour at different scales. The double filter bank design certainly allows the contourlet to be flexibly constructed. In [26], Lu and Do developed a critically sampled contourlet transform called *"CRISP-contourlet* using a combined iterated non-separable filter bank for both multiscale and directional decomposition. A non-subsampled contourlet transform was recently proposed [14]. The Laplacian pyramid was substituted with a 2-channel non-subsampled 2D filter bank which is similar to the *á trous* wavelet expansion. However, with J levels of decomposition, it has $J+1$ redundancy. By contrast, the 2-D *á trous* algorithm by tensor product has $3J+1$ redundancy. The whole transform has $1+\sum_{j=1}^{J} 2^{l_j}$, where l_j denotes the number of levels in the transform at the j-th scale. Experimental results suggest that the transform compares favourably to other existing denoising and enhancement methodds reported in literature.

The curvelets can also be conveniently constructed from a frequency tiling approach. Such idea later adopted by Candés and Donoho [8] in constructing *second generation curvelets* which do not require ridgelets. Such tight frame can be computed more efficiently than the previous curvelet-99 implementation. A recent report [5] details its implementation using unequally-spaced fast Fourier transforms (USFFT) and the wrapping of specially selected Fourier samples. Both implmentations are improved in the sense that they are conceptually simpler, faster and far less redundant. The same strategy was used in constructing a 3D curvelet transform [43] whose basis functions are planar patches. These digital implementations can be found in the `CurveLab` distribution.

However, the assumption that natural images are characterised solely by linear edges is not true. Evidently we have seen attempts to separate the image into additive ingredients [31] - usually one is textural and the other is piecewisely smooth. This suggests that there is usually a presence of local strongly oriented harmonics (textures) separated by curvilinear edges. Sparse representations which are both well-localised in frequency and orientation is desirable to efficiently describe such oriented harmonic features. Also, it would be ideal to accommodate both directional linear features as well as directional periodic textures in a unified manner according to the *"image=texture+edge"* model. In this paper, we show that directional wavelet analysis can be performed with directional trigonometric transforms localised in a multiscale framework. We introduce the *Multiscale Directional Cosine Bases* in section II which can efficiently represent local oriented harmonics, and with a local directional cosine packet analysis, we can accommodate both directional periodic ridges and ridgelets which is a dual basis to the ridgelet packets. In section III, we show that directional singularities and harmonics can also be captured by the *Multiresolution Fourier transform* using a Gaussian model of its magnitude spectrum with less computational burden. Next we show some results from our denoising experiments with both transforms and compare them with other wavelet transforms in section IV. The paper concludes with a summary in section V.

II. The Multiscale Directional Cosine Transform

Like the 2D orthonormal wavelet transform, the discrete cosine transform (DCT) in 2D is also formed by tensor product, resulting in basis functions which look like "chessboard" patterns. Therefore we need to define a directional cosine operator in order to bring the orientation parameter into the transform. Also, we will need to localise the basis spatially in order to capture local features. This section describes the construction of the Multiscale Directional Cosine Transform (MDCT).

A. The Directional Cosine Basis

First we define the parametric space $\Gamma = \{\gamma = (k, \vec{\theta})\}$ where $k \in [0, 2\pi), \vec{\theta} \in \mathbf{S}^{d-1}$, $\vec{\theta}$ is on the unit sphere $\mathbf{S}^{d-1}$ in dimension d which indicates orientation and k indicates the frequency. Consider a family of orthonormal trigonometric basis for $L^2([0,1])$, derived from Fourier transform $\hat{f}(\xi) = \int e^{-ix\xi} f(x)dx$.

1) $\{\sqrt{2}\sin(\frac{2k+1}{2}\pi x)\}$, $k = 0, 1, 2, 3, ...$
2) $\{\sqrt{2}\sin(k\pi x)\}$, $k = 1, 2, 3, ...$
3) $\{\sqrt{2}\cos(\frac{2k+1}{2}\pi x)\}$, $k = 0, 1, 2, 3, ...$
4) $\{1, \sqrt{2}\cos(k\pi x)\}$, $k = 1, 2, 3, ...$

We denote such a trigonometric basis as $c_k(x)$, and the corresponding transform can be written as $\langle f, c_k\rangle$. Now we define the continuous directional trigonometric transform on a multi-variate function $f(\mathbf{x}), \mathbf{x} \in \mathbb{R}^d$:

$$\mathcal{C}_\gamma(\mathbf{x}) = c_k(\vec{\theta} \cdot \mathbf{x}) \tag{1}$$

Since $\langle f, c_k\rangle$ is essentially a Fourier transform, we have the admissibility condition

$$K_{\mathcal{C}} = \int \frac{|\hat{\mathcal{C}}(\xi)|^2}{|\xi|^d} d\xi < \infty \tag{2}$$

and the reconstruction is

$$f = \int \langle f, \mathcal{C}_\gamma\rangle \mathcal{C}_\gamma \mu(d\gamma) \tag{3}$$

The Parseval relation holds

$$\|f\|_2^2 = \int |\langle f, \mathcal{C}_\gamma\rangle|^2 \mu(d\gamma) \tag{4}$$

For a general image representation, we choose the $c_k = \{1, \sqrt{2}\cos(k\pi x)\}$, known as the cosine II basis, which has faster decay on interval $[0, 1]$ than the Fourier transform. We then have the directional 2D basis

$$\mathcal{C}_{k,\theta} = \lambda_k \cos(\pi k(x\cos\theta + y\sin\theta)) \tag{5}$$

where $\lambda_k = \begin{cases} 1 & \text{if } k = 0 \\ \sqrt{2} & \text{if } x \neq 0. \end{cases}$

The directional 2D continuous cosine transform is defined as

$$\begin{aligned} \mathcal{C}f(k,\theta) &= \langle f, \mathcal{C}_{k,\theta}\rangle \\ &= \int_{\mathbb{R}^2} \lambda_k f(x,y) \cos(\pi k(x\cos\theta + y\sin\theta)) dxdy \end{aligned}$$

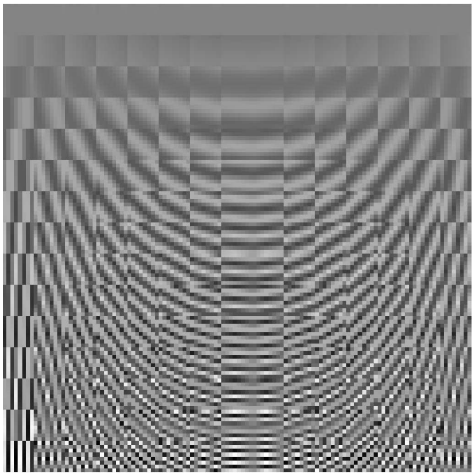

Fig. 1. The 8×8 directional cosine basis vectors

The directional cosine basis vectors are indexed by frequency k and direction θ, as can be seen in Figure 1. It is obvious that the basis vectors look similar to the Fourier basis despite the fact that the directional cosine transform is real-to-real instead of real-to-complex and its approximation error decays more rapidly than the Fourier counterpart, due to its symmetrical boundary extension.

B. Directional Cosine Packets

Smooth local trigonometric bases proposed by Coifman and Meyer [11] and by Malvar [30] use smooth window functions to split the signal and to fold overlapping parts back into the pieces so that the orthogonality is preserved. Therefore, the folded signal is suited for representation by a trigonometric basis. At its simplest, a typical local cosine basis function has the form :

$$\phi_{n,k}(x) = b_n(x)\cos(k+\frac{1}{2})\pi x. \tag{6}$$

where b_n is a smooth window or a *bell function*.

The usefulness of applying smooth local trigonometric bases to focus on local interesting properties of a signal is well studied, and their applications such as MP3 audio compression have been demonstrated to be successful. As we have discussed, image usually consists regions of homogenous textures separated by linear edges and contours. It is then natural to consider applying local cosine basis on Radon projected slices in order to represent both periodic patterns and some linear singularities. In this work, we use the Coifman-Wickerhauser's entropy-based best basis algorithm [13] to look for the best local cosine basis with dyadic interval. The resulting adaptive basis is similar to one possible "*ridgelet packets*" construction mentioned in [21]. We will thereafter refer to such dictionary of bases in the Radon domain as the *"directional cosine packets"*.

C. Multiscale Digital Implementation

For an image representation basis to be useful, the basis vectors should be localised both in space and frequency, and they should have certain orientation selectivity. More importantly, to capture patterns of interest at different scales, the basis need to be multiresolution. A prototypical MDC function has the form

$$\psi_{k,\theta,s,\mathbf{t}}(\mathbf{x}) = b\left(\frac{\mathbf{x}-\mathbf{t}}{s}\right)\mathcal{C}_{k,\theta}\left(\frac{\mathbf{x}-\mathbf{t}}{s}\right). \tag{7}$$

where k, θ, $\mathbf{t}$ and s denotes the frequency, orientation, location and scale parameters of the function respectively and $b(\cdot)$ is the smooth bell function chosen along with the sampling interval to ensure invertibility of the discrete form of the transform.

The discrete implementation of MDCT is similar to the digital curvelet-99 construction. While the discrete cosine transform and discrete Radon transform [1] are well studied in the literature, a combination of these two transforms gives us the discrete directional cosine operator. Unlike curvelet-99 which is very redundant, the multiresolution property of the MDC transform is given by the well-known decimated Laplacian pyramid [3]. The discrete MDC of a 2D vector $\mathbf{x}$, at scale s is given by

$$\mathbf{X}_s = \mathcal{C}_n(\mathbf{I} - \mathbf{G}_{s,s+1}\mathbf{G}_{s+1,s})\mathbf{x}_s. \tag{8}$$

where $\mathbf{X}_s$ denotes the transform at scale s, $\mathcal{C}_n$ is the discrete directional cosine transform operator with window size $n \times n$, $\mathbf{I}$ is the identity operator, $\mathbf{x}_s$ is the Gaussian pyramid representation of $\mathbf{x}$ at scale s

$$\mathbf{x}_s = \prod_{l=0}^{s-1}\mathbf{G}_{l+1,l}\mathbf{x}. \tag{9}$$

and $\mathbf{G}_{s,s+1}, \mathbf{G}_{s+1,s}$ are the raising and lowering operators associated with transitions between levels in the Gaussian pyramid. We certainly have the choice of using the directional cosine packets as the transform operator $\mathcal{C}_n$ by substituting the cosine transform by a cosine packet operator ψ_n, forming a semi-adaptive basis. In this way, the MDC packet basis is able to capture a wide range of directional features at different resolutions.

III. The Multiresolution Fourier Transform

The MDCT is similar to the curvelet-99 transform, only the wavelet ridge function is replaced by the cosine basis. With the local cosine analysis on Radon slices, the MDC packet bases fits well to the *" image = edges + textures"* model.

However, the proposed bases have two limitations. The first problem is that best basis for local cosine packets has to be sought on every Radon slice, making the computation extremely expensive. Secondly, the Radon transform is a redundant transform and its inverse introduces some numerical errors. One might note that the Radon transform is directly related to the Fourier transform by the Fourier Slice Theorem, briefly stated as below:

Theorem 1: (Fourier Slice Theorem). The 1D Fourier transform with respect to t of the projection $Rf(t,\theta)$ is equal to a central slice, at angle θ, of the 2D Fourier transform of the function $f(x,y)$, that is,

$$\hat{Rf}(t,\theta) = \hat{f}(\xi\cos\theta, \xi\sin\theta). \tag{10}$$

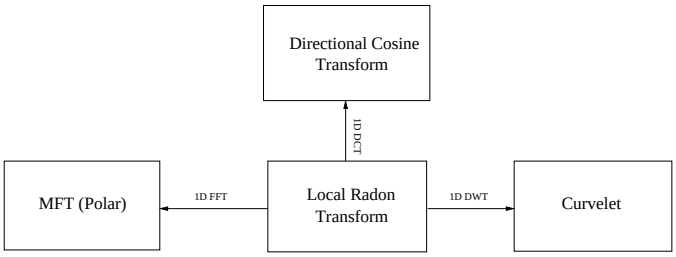

Fig. 2. The relationships between MFT, Radon, Curvelet and directional cosine basis

where

$$\hat{f}(\xi_1, \xi_2) = \int\int f(x, y) e^{-2\pi i(x\xi_1 + y\xi_2)} dxdy.$$

is the 2D Fourier transform of $f(x, y)$.

Since many discrete Radon transform are implemented via this theorem, it suggests that we may be able to perform such directional analysis using the standard Fourier transform.

A. The MFT Implementation

The *Multiresolution Fourier Transform* (MFT) [4], [41] has been proposed as a combination of STFT and wavelet methods, which inherits many of the desired features of both. With the windowing function $g(t)$, the transform of a function $f \in L^2(\mathbb{R})$ at position u frequency ξ and scale s is defined as:

$$Mf(u, \xi, s) = \frac{1}{\sqrt{s}} \int_{-\infty}^{+\infty} f(t) g(s(t-u)) e^{-i\xi t} dt. \quad (11)$$

In effect, it is simply a stack of windowed Fourier transforms, in which the scale of the analysis window is varied systematically with the stack index. As a general image analysis tool, it has been applied in feature extraction and segmentation with music and image analysis, such as music note segmentation and extracting boundary curves in a multiresolution fashion [41]. It has also been used in texture synthesis and analysis [24] and many other areas.

The discrete implementation of MFT can take many forms. Similar to the construction of the digital MDCT described before. We build the MFT on top of the Laplacian pyramid, then on each level of the pyramid, windowed Fourier transform is performed with the same window size regardless of the scale. The discrete MFT of a 2D vector $\mathbf{x}$, at scale s is given by

$$\hat{\mathbf{x}}_s = \mathcal{F}_n(\mathbf{I} - \mathbf{G}_{s,s+1}\mathbf{G}_{s+1,s})\mathbf{x}_s. \quad (12)$$

where $\mathcal{F}_n$ is the discrete Fourier transform operator with window size $n \times n$. The closeness of the Burt and Adelson filter to a Gaussian function gives the pyramid virtually isotropic behavior, which can be well exploited by the high frequency resolution of the Fourier basis. The whole transform is some 5.33 times redundant if overlapping window is used.

We can see that the only difference between MFT and the MDCT is that the operator used here is just a Fourier transform. In fact, the polar separablity of the Fourier transform suggests that it is also a directional trigonometric transform and Radon transform was implemented via the Fourier-slice theorem by inverse Fourier transform on Fourier polar slices. The relations between MFT, MDCT, Radon and curvelet transform can be illustrated in Figure 2. It is obvious that it requires an inverse Fourier transform and a cosine transform to convert the Fourier domain into the directional cosine domain. Although it has some advantage in approximation convergence, the extra computation is 2 times more than the conventional Fourier transform.

B. Gaussian Modelling of Fourier Spectrum

The Fourier basis is a natural representation for directional periodic patterns, although it decays slower than a cosine basis in terms of approximation. In order to perform some sort of *"curvelet"* analysis, we need a model for linear features in the Fourier domain. Fortunately, such model is not difficult to derive, since a line in the spatial domain will be transformed into another line in its Fourier domain perpendicular to its direction.

The magnitude intensity of the local Fourier spectrum can be modelled as a single 2D Gaussian function with its centroid fixed at the origin, which means the Gaussian is zero-mean :

$$G(\mathbf{x}) = \frac{1}{2\pi} \exp\left(\frac{-\mathbf{x}^T C^{-1} \mathbf{x}}{2}\right). \quad (13)$$

The covariance matrix C of the Gaussian $G(\cdot)$ can be obtained from the inertia tensor of the spectrum.

$$C = \sum_{\vec{\omega}} |\hat{f}(\vec{\omega})|^2 \vec{\omega}\vec{\omega}^T. \quad (14)$$

where the $\hat{f}(\vec{\omega})$ denotes the Fourier coefficients. With the covariance matrix C, the shape and the orientation of the Gaussian is determined. The centroid of the feature $\mathbf{x}_0$ can then be estimated by taking the pairwise average correlations between neighbouring coefficients in each of the horizontal and vertical directions :

$$\mathbf{x}_0 = \frac{B}{2\pi} \arg\left(\sum_{\vec{\omega}} \hat{f}(\vec{\omega} - 1)\hat{f}(\vec{\omega})^*\right). \quad (15)$$

where B is the windowing size.

The choice of using the Gaussian function is due to several reasons. First, the uncertainty principle states that the Gaussian function can achieve optimal spread in space and frequency, and it is smooth in both domains. Secondly, the shape of the 2D Gaussian function can be both isotropic and anisotropic. When the covariance matrix gives an anisotropic Gaussian distribution, this suggests that the spatial feature is a linear shape. In other cases, the Gaussian blobs will tend to be isotropic. A simple measure of the anisotropy can be obtained by performing the Principal Components Analysis (PCA) on the covariance matrix C, yielding two eignvalues λ_1 and λ_2, where $\lambda_1 \leq \lambda_2$. The measure is simply:

$$\mathcal{A} = \left|\frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2}\right| \quad (16)$$

To test the effectiveness of our model, we use a simple linear feature with Gaussian white noise. The noisy line (see Figure

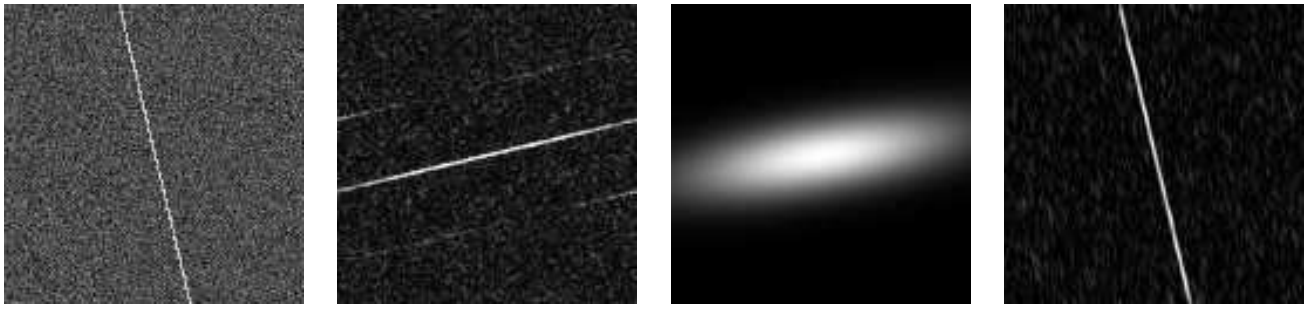

(a) (b) (c) (d)

Fig. 3. (a) a noisy line, (b) the Fourier transform of the noisy line, (c) The Gaussian filter estimated from inertia tensor, (d) the denoised line.

3.(a)) is transformed into its Fourier domain (see Figure 3.(b)). In order to suppress noise, we can use the $G(\mathbf{x})$ as a frequency filter in the Fourier domain. However, the problem is that the inertia tensor itself is easily affected by noise. Our solution is to apply thresholding on the noisy transformed data as a stage of pre-processing, in order to suppress the noise energy. The inertia tensor C then can be more reliably estimated from the thresholded data. Therefore, the resulting Gaussian frequency filter is estimated from a thresholded version of Figure 3.(b). The inversion is a clean line image without most of the noise energy. In this way, we have achieved a directional *ridgelet*-like analysis with the Fourier basis, based on the single-feature hypothesis. While the assumption is not realistic for a natural image, such library of wave packets will work well locally in a multiresolution setting. A combination of this model to the MFT allows us to analyse the signal adaptively, so that many features including contours and textures can be captured effectively.

IV. Image Denoising Experiments

Good bases for reprensenting images should be able to capture important features of interest, so that the reconstruction requires as few basis functions as possible. The bases' effectiveness can be tested by performing denoising experiments by simple thresholding in the transformed domain. For the MDCT and MDC packet transform, the denoising experiments are performed in such settings:

- The Laplacian pyramid is decomposed at 5 levels of subbands.
- The window size n is chosen at 16×16, modulated with a squared cosine.
- The windows are overlapped by 50%.

The thresholding we use is a form of the universal thresholding proposed in [18], multiplied by an extra constant a, $\Theta = a\sqrt{2 \log N}\sigma / 1.23^L$, where $N = n^2 = 256$ here and L denotes the level of decomposition, while $L = 0$ corresponds to the highest frequency subband. For directional cosine denoising, $a = 0.08$ was found to give satisfactory result. For the directional local cosine packets, $a = 0.062$ was used. The lowpass subband is left intact.

The settings for the MFT are generally the same as for the MDCT, only with a little sophistication on estimating the filter and $a = 0.8$

1) Within each Fourier transformed block $\hat{B}$, we apply the threshold $8\sqrt{2 \log 256}\sigma$ on that to obtain $\hat{B}_T$, from which the inertia tensor C will be estimated.
2) If $\mathcal{A}_C > 0.43$, which means there is a strong directional feature present, the Gaussian filter generated from C will apply on the original noisy spectrum $\hat{B}$ to obtain $\hat{B}_G$, the inverse Fourier transform is taken on $(\hat{B}_G + \hat{B}_T)/2$ as the denoised block.
3) Otherwise, we take the $\hat{B}_T$ as the denoised block.

The Fourier-type transforms are compared with two algorithms. The first is a wavelet-packet based wavelet shrinkage algorithm which is described in [35], called $\mathcal{S}$-Bayes in which the thresholding function is a modified version of the *BayesShrink* [9]. The best wavelet packet basis is sought by using the Shannon entropy function and cycle-spinning [12] is used to suppress the pseudo-Gibbs artifact. Essentially such treatment gives the translation invariance to the wavelet packet basis, which is known to be good in representing periodic signals as well as discontinuities.

The second is a modified version of curvelet. The curvelet-99 implementation reported in [37], which uses a much more redundant overcomplete wavelet frame than our MDCT and MFT, is a "specialised" transform to perform denoising task instead of general-purpose image processing. Therefore, here the local ridgelets are placed on the Laplacian pyramid as in our setting, in order to carry out a fair comparison.

We have conducted experiments on a wide range of natural images. Three of them present some typical characteristics: `barbara` contains some directional and non-directional periodic textures; `lena`, which can be regarded as one of the "curvelet-friendly" image, since it mainly consists of linear discontinuities at different scales; the `grain` image is a texture image which was considered to be very difficult to compress. It contains many directional components, however very irregular.

Table I gives denoising results in SNR by those five bases, where the best numbers are stressed in bold. We see that the best results are always among the MDC packet and the MFT, while the MDC packet seems to be more effective in more noisy situations. This is due to the fact that the Gaussian filters are estimated from the noise-sensitive inertia tensor. When the spectrum are dominated by the noise energy, a simple thresholding would fail to preserve the signal information. However, the MFT can be considered as the overall winner: it compares

TABLE I
THE COMPARATIVE IMAGE DENOISING RESULTS IN SNR

Image	Noise (dB)	TIWP	Curvelet	MDCT	MDC Packet	MFT
barbara	0	14.75	14.59	14.89	**15.19**	14.80
	5	16.08	15.93	16.42	16.85	**17.10**
	10	18.00	17.64	18.44	18.85	**19.68**
	15	21.12	19.64	20.34	20.71	**22.46**
	20	24.92	21.41	21.88	22.10	**24.65**
lena	0	17.06	17.10	17.24	**17.75**	16.63
	5	18.96	18.82	18.99	**19.63**	19.02
	10	21.08	20.83	21.05	**21.89**	21.60
	15	23.49	23.10	23.35	24.26	**24.43**
	20	26.15	25.35	25.60	26.31	**26.94**
grain	0	12.93	13.06	13.10	**13.22**	12.85
	5	13.81	13.66	13.84	**14.29**	14.01
	10	15.87	15.01	15.51	16.18	**16.30**
	15	18.72	17.28	18.00	18.84	**19.66**
	20	22.02	20.05	20.56	21.38	**22.82**

well with other denoising methods and its computational cost is much lower than the rest of the methods. We also notice that TIWP-S-BayesShrink sometimes outperforms at some low noise levels, since the BayesShrink tends to optimise the MSE output. However, the visual qualities of other candidates are more pleasing, preserving important directional features on the image.

A detailed head-to-head comparison is presented in Figure 4 on `barbara`. It is obvious that from the TIWP-$\mathcal{S}$-BayesShrink thresholded image, the diagonal strips are absent on the cloth in the middle, although a few of such patterns can be seen on the trousers. The curvelet is able to recover some of those directional patterns, but incomplete nonetheless. These features are restored almost completely by our proposed methods.

Since the directional cosine packets can be regarded as a generalisation of the curvelets and directional cosine bases, it is not suprising to see it gives better results than these two counterparts. However, it introduces considerable amount of extra computations, since the best basis has to be sought for each of the Radon slices. Also, the inverse Radon transform from incomplete data introduces numerical errors in the reconstruction. Although visually MFT and MDC packet both captured linear and oscillating patterns, MFT's reconstruction is much sharper and outperforms by quite a margin in SNR. On the other hand, the best-basis search ensures that the denoised image from MDC packet is smoother and almost "artifact-free". Considering the visual/statistical performance and the complexity, the MFT with Gaussian filter is the overall winner.

V. CONCLUSION

In this paper, we have reviewed a growing literature body on directional wavelets' construction, analysis and their applications. The contribution of this paper is to introduce sets of Fourier-type bases which have localisation in space and frequency, orientation selectivity, and employing a multiresolution pyramidal framework allowing analyses of images at different scales.

In a sense, the Fourier-type bases qualify as geometrical wavelets and share a lot of similarities with other directional wavelet bases proposed previously. But the semi-adaptiveness allows us to capture local directional texture patches and linear features at ease and we have shown that by a simple Gaussian frequency filter model of magnitude spectrum intensity, analysis of directional harmonics and linear features can be carried out much more efficiently than any other directional wavelet bases proposed to date. It can be considered as a parametric curvelet representation, or a *generalised directional wavelet packets* and the simple inertia tensor method has demonstrated to be a good substitute for the Cartesian-polar conversion. More importantly, its computational cost is a big advantage, since it does not involve a notion of Radon transform, nor the best-basis search as in some adaptive representations.

The effectiveness of the proposed bases was tested against the state-of-the-art translation-invariant wavelet packet based shrinkage method and the curvelets. The new bases demonstrated a strong potential in the experiments, outperforming the opponents by quite a margin. While producing much visually pleasant output than the wavlet packets with optimal threshold, the MDC bases seems to be able to capture a wider range of directional features than the curvelet, even without the local cosine treatment.

The denoising experiments show the effectiveness of conducting multiresolution analysis with these bases. A wide variation on the theme is possible, for example using variable sized windows on the original image might be another possibility, or to use the *algorithme á trous* subband decomposition for better denoising results. The usage of the Gaussian frequency filter and its parameter estimation in noisy environments are still under investigation. It is our intention to put forward these bases in a general way in this work to popularise their usage in various kinds of image processing tasks.

REFERENCES

[1] A. Averbuch, R. Coifman, D.L. Donoho, and M. Israeli. Fast slant stack: A notion of radon transform for data in a cartesian grid which is rapidly computible, algebraically exact, geometrically faithful and invertible. to appear in SIAM Scientifi c Computing.

[2] T. Binford. Inferring surfaces from images. *Artifi cial Intelligence*, 17:205–244, 1981.

[3] P.J. Burt and E.H. Adelson. The Laplacian pyramid as a compact image code. *IEEE Transactions on Communications*, 31:532–540, 1983.

[4] A. Calway. *The Multiresolution Fourier Transform: A General Purpose Tool for Image Analysis*. PhD thesis, University of Warwick, 1989.

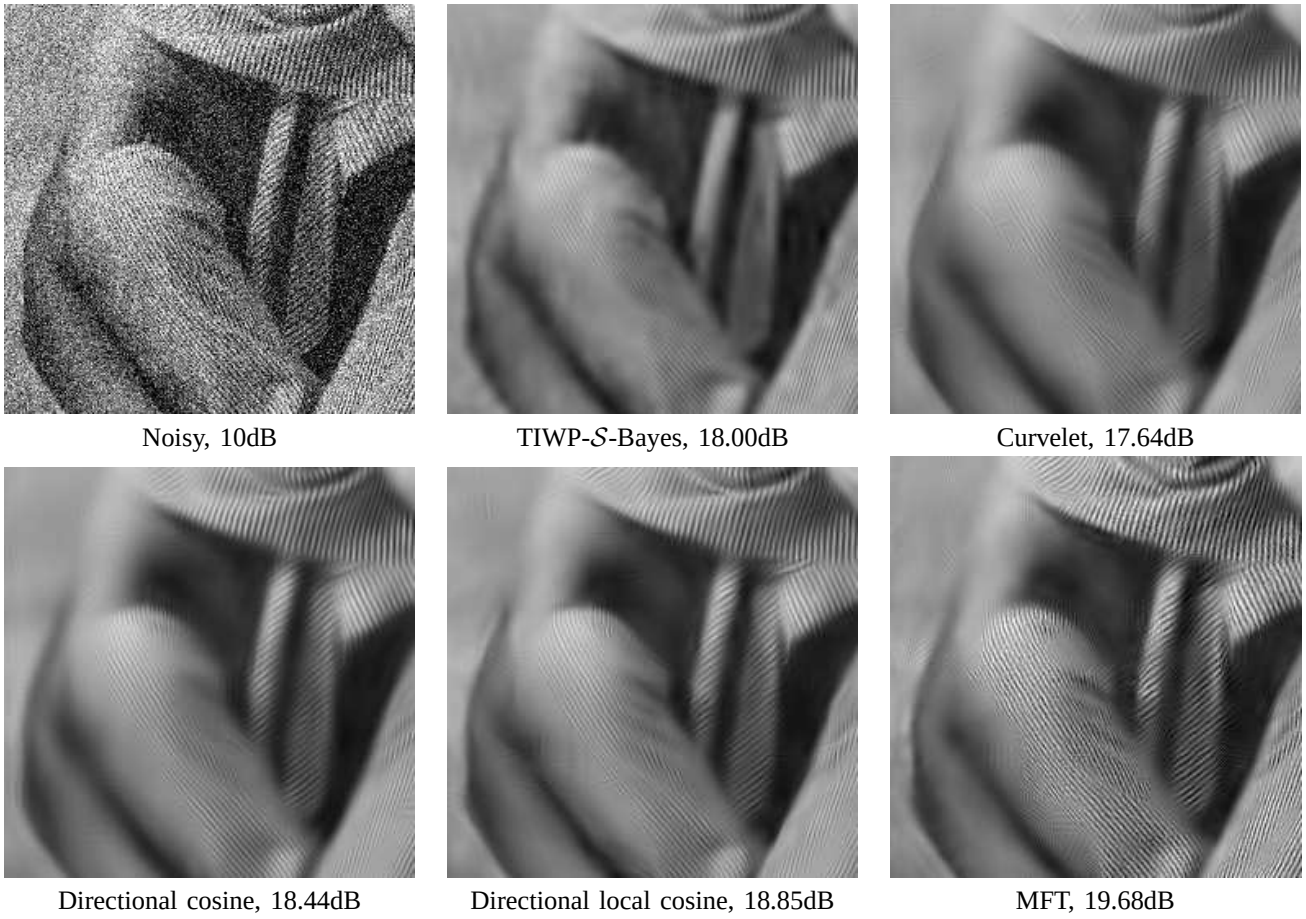

Fig. 4. Detailed comparative denoising results on `barbara`

[5] E. Candés, L. Demanet, D. Donoho, and L. Ying. Fast discrete curvelet transforms. Technical report, California Institute of Technology, July 2005.

[6] E.J. Candés. *Ridgelets: Theory and applications*. PhD thesis, Dept. of Stats, Stanford Univ., Stanford, CA, 1998.

[7] E.J. Candés and D.L. Donoho. Curvelets - a suprisingly effective nonadaptive representation for objects with edges. In C. Rabut, A. Cohen, and L.L. Schumaker, editors, *Curves and Surfaces*, pages 105–120. Vanderbilt University Press, Nashville, TN, 2000.

[8] E.J. Candés and D.L. Donoho. New tight frames of curvelets and optimal representations of objects with C^2 smooth singularities. Technical report, Department of Statistics, Stanford University, 2002.

[9] S.G. Chang, B. Yu, and M. Vetterli. Adaptive wavelet thresholding for image denoising and compression. *IEEE Transaction on Image Processing*, 9(9):1532–1546, September 2000.

[10] A. Cohen and B. Matei. Nonlinear subdivisions schemes: Applications to image processing. In A. Iske, E. Quack, and M. Floater, editors, *Tutorial on Multiresolution in Geometric Modelling*. Springer, New York, 2002.

[11] R. Coifman and Y. Meyer. Remarques sur l'analyse de fourier à fenêtre. *C.R. Acad. Sci. Paris Sér. I Math.*, I(312):259–261, 1991.

[12] R.R. Coifman and D.L. Donoho. Translation invariant denoising. In A. Antoine and G. Oppenheim, editors, *Wavelets in Statistics*, pages 125–150. Springer, New York, 1995.

[13] R.R. Coifman and M.V. Wickerhauser. Entropy-based algorithms for best basis selection. *IEEE Trans. on Information Theory*, 38(2):713–718, 1992.

[14] A.L. Cunha, J. Zhou, and M.N. Do. The nonsubsampled contourlet transform:Theory, Design and Applications. *IEEE Transactions on Image Processing*, 2005. submitted.

[15] M.N. Do and M. Vetterli. Contourlet. In G.V. Welland, editor, *Beyond Wavelets*. Academic Press, 2003.

[16] M.N. Do and M. Vetterli. The contourlet transform: an efficient directional multiresolution image representation. Submitted to IEEE Transactions Image Processing, 2003.

[17] D. Donoho. Wedgelets: Nearly-minimax estimation of edges. *Ann. Stat.*, 27:353–382, 1999.

[18] D.L. Donoho and I.M. Johnstone. Ideal spatial adaptation via wavelet shrinkage. *Biometrika*, 81:425–455, 1994.

[19] M. Dragotti and M. Vetterli. Wavelet footprints: Theory, algorithm and applications. *IEEE Transactions on Signal Processing*, 51(5):1306–1323, May 2003.

[20] J. Elder. Are edges imcomplete? *International Journal Computer Vision*, 34(2/3):97–122, 1999.

[21] A.G. Flesia, H. Hel-Or, A. Averbuch, E.J. Candés, R.R. Coifman, and D.L. Donoho. Digital implementation of ridgelet packets. In G. Welland, editor, *Beyond Wavelets*, pages 31–60. Academic Press, Sep. 2003.

[22] G.H. Granlund. In search of a general picture processing operator. *Computer Graphics and Image Processing*, 8:155–173, 1978.

[23] B. Horn. The Binford-Horn linefinder. Technical Report Memo 285, MIT AI Lab, 1971.

[24] T.I. Hsu and R.G. Wilson. A two-component model of texture for analysis and synthesis. *IEEE Transactions on Image Processing*, 7:1466–1476, 1998.

[25] D.H. Hubel and T.N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology*, pages 106–154, 1962.

[26] Y. Lu and M.N. Do. CRISP-contourlet: a critically sampled directional multiresolution image representation. In *Proceedings SPIE conference on Wavelet Applications in Signal and Image Processing*, San Diego, Aug. 2003.

[27] S. Mallat. A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 11:674–693, July 1989.

[28] S. Mallat and Z. Zhang. Atomic decomposition by basis pursuit. *IEEE Transactions on Signal Processing*, 41(12):3397–3415, 1993.

[29] S. Mallat and S.S. Zhong. Wavelet transform maxima and multiscale edges. In M.B. Rskai et al., editor, *Wavelets and their applications*. Jones and Bartlett, 1992.

[30] H.S. Malvar. Lapped transforms for effi cient transform/subband coding. *IEEE Trans. Acoust. Speech Signal Processing*, 38:969–978, 1990.

[31] F. Meyer, A. Averbuch, and R. Coifman. Multilayered image representation: Application to image compr ession. *IEEE Transactions on Image Processing*, 11:1072–1080, 2002.

[32] E. Le Pennec and S. Mallat. Image compression with geometrical wavelets. In *presented at IEEE ICIP*, Vancouver, BC, Canada, Sep. 2000.

[33] E. Le Pennec and S. Mallat. Sparse geometric image representations with bandelets. *IEEE Transactions on Image Processing*, 14(4):423–438, April 2005.

[34] G. Peyré and S. Mallat. Discrete bandelets with geometric orthogonal fi lters. In *Proceedings IEEE ICIP*, 2005. submitted.

[35] N. Rajpoot, Z. Yao, and R. Wilson. Adaptive wavelet restoration of noisy video sequences. In *Proceedings IEEE ICIP 2004*, Singapore, 2004. To Appear.

[36] R. Shukla, P.L. Dragotti, M.N. Do, and M. Vetterli. Rate-distortion optimized tree structured compression algorithms. *IEEE Transactions on Image Processing*, to be published.

[37] J. Starck, E.J. Candés, and D.L. Donoho. The curvelet transform for image denoising. *IEEE Transactions on Image Processing*, 11(6):670–684, June 2002.

[38] J. Starck, D.L. Donoho, and E.J. Very high quality image restoration by combining wavelets and curvelets. In *Proceedings of SPIE Wavelets: Applications in Signal and Image Processing IX*, volume 4478, pages 9–19, 2001.

[39] J. Starck, M.K. Nguyen, and F. Murtagh. Wavelets and curvelets for image deconvolution: a combined approach. *Signal Processing*, 83:2279–2283, 2003.

[40] M. Wakin, J. Romberg, H. Choi, and R. Baraniuk. Rate-distortion optimized image compression using wedgelets. In *Proceedings IEEE ICIP*, volume 3, pages 237–240, 2002.

[41] R. Wilson, A.D. Calway, and E.R.S. Pearson. A generalized wavelet transform for fourier analysis: the multiresolution fourier transform and its application to image and audio signal analysis. *IEEE Trans. on Information Theory*, 38:674–690, 1992.

[42] X. Xue and X. Wu. Image compression based on multi-scale edge compensation. In *Proceedings IEEE ICIP*, volume 3, pages 560–564, 1999.

[43] L. Ying, L. Demanet, and E.J. Candés. 3d discrete curvelet transform. In *Proceedings SPIE Wavelet XI*, volume 5914, 2005.