IMAGE DENOISING USING MULTISCALE DIRECTIONAL COSINE BASES

Zhen Yao   Nasir Rajpoot

Dept. of Computer Science, University of Warwick
Coventry, CV4 7AL, United Kingdom
Email: {yao, nasir}@dcs.warwick.ac.uk

ABSTRACT

Motivated by the fact that in images, there is usually a presence of local strongly oriented harmonics, a representation which is both well-localised in frequency and orientation is desirable to efficiently describe such oriented harmonic features. Here we introduce a family of multiscale trigonometric bases for a bi-variate function called the Multiscale Directional Cosine bases for image denoising tasks. Our results show the promise of the new bases which almost consistently outperform other image representation bases on natural images.

1. 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 2D basis functions are simply tensor products of the corresponding 1D 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 [1] 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.

The ability to efficiently analyse and describe directional patterns is of fundamental importance for image analysis and compression. In order to represent higher dimensional image features, Candés and Donoho [2] proposed a directional multiresolution representation named *ridgelets* which deal effectively with linear singularities in 2D. The idea of *curvelets* [3] is to represent a curve as a superposition of functions of various lengths and widths obeying the scaling law. Starck et al. [4] employed the curvelet transform in the denoising task. The curvelet coefficients are more than $16J + 1$ times redundant, with J scales of decompositions. Not only does the curvelet denoising outperform wavelet-based methods in terms of PSNR, the visual appearance of the resulting images is better as well. Do and Vetterli [5]

used similar ideas to construct curvelets that can be computed with a perfect reconstruction filter bank, called the *contourlet transform*, which combines the Laplacian pyramid transform with directional filter banks.

However, the assumption that natural images are characterised solely by linear edges is not true. In images, there is usually a presence of local strongly oriented harmonics (textures) separated by curvilinear edges. A sparse representation which is both well-localised in frequency and orientation is desirable to efficiently describe such oriented harmonic features. Here we describe a set of bases called *Multiscale Directional Cosine Bases* which can efficiently represent local oriented harmonics, and with a local directional cosine packet analysis, we can accommodate the "texture + edge" model by representing both directional periodic ridges and ridglets in a unified framework.

For the rest of this paper, we introduce a family of bases called the "Multiscale Directional Cosine" (MDC) Bases in section 2. Next we show some results from our denoising experiments using the MDC bases in section 3. The paper concludes with a summary in section 4.

2. MULTISCALE DIRECTIONAL COSINE BASES

First, we define the directional cosine basis in 2D as

$$\mathcal{C}_{k,\theta} = \lambda_k \cos(\pi k(x \cos \theta + y \sin \theta)) \tag{1}$$

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

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

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. In fact, if we substitute the cosine basis function with $e^{ix\xi}$, we obtain the polar Fourier basis representation. However,
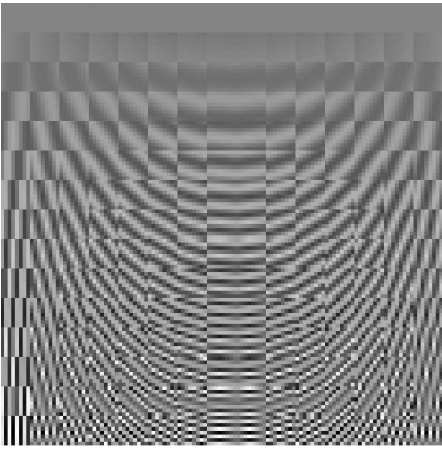
Fig. 1. The 8×8 directional cosine basis vectors

the proposed transform has several advantages when compared with traditional Fourier basis. First, the transform is real-to-real instead of real-to-complex. Second, since the cosine basis assumes a symmetric extension, its approximation error decays more rapidly than Fourier series. Lastly, unlike in the Fourier domain, the Cartesian-to-polar coordinate conversion has already been done in the Radon transform to get the sense of directionality.

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 multiresolutional. A prototypical MDC function has the form

$$\psi_{k,\theta,s,\mathbf{t}}(\mathbf{x}) = b\left(\frac{\mathbf{x} - \mathbf{t}}{s}\right) \mathcal{C}_{k,\theta}\left(\frac{\mathbf{x} - \mathbf{t}}{s}\right). \qquad (2)$$

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 MDC transform is similar to the digital curvelet construction. While the discrete cosine transform and discrete Radon transform [6] are well studied in the literature, a combination of these two transforms gives us the discrete directional cosine operator. Unlike the digital curvelet implementation in [4], which is very redundant, the multiresolution property of the MDC transform is given by the well-known Laplacian pyramid [7]. Therefore, the transform is over-complete by some 33% in 2D if non-overlapping windows are used. The discrete MDC of a 2D vector $\mathbf{x}$, at scale s is given by

$$\mathbf{X}_s = \mathcal{C}_n(\mathbf{I} - \mathbf{G}_{s,s+1}\mathbf{G}_{s+1,s})\mathbf{x}_s. \qquad (3)$$

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

$$\mathbf{x}_s = \prod_{l=0}^{s-1} \mathbf{G}_{l+1,l}\mathbf{x}. \qquad (4)$$

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. 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 cosine basis. 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 analysis, 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.

3. IMAGE DENOISING EXPERIMENTS

A good basis for reprensenting images should be able to capture important features of interest. In this case, the MDC bases should allow us to map directional periodic pattern patches into coefficients of significant magnitude. In order to demonstrate the efficiency of the proposed representation, we conducted image denoising experiment by simply thresholding the coefficients in the transform domain with certain threshold. The denoising experiments are performed in such settings:

1. The Laplacian pyramid is decomposed at 5 levels of subbands.

2. The window size n is chosen at 16×16, modulated with a squared cosine with 50% overlapping.

3. We used both directional cosine and directional cosine packet as the operator $\mathcal{C}_n$ on the windowed block.

The thresholding we use is a form of the universal thresholding proposed in [8], 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, the value of $a = 0.08$ was found to give satisfactory result. For the directional local cosine packets, $a = 0.062$ was used. The lowpass subband is left intact.

The results are compared with two algorithms. The first one is a wavelet-packet based wavelet shrinkage algorithm which is described in [9], called $\mathcal{S}$-Bayes, with the thresholding function being a modified version of the *BayesShrink* [10]. The best wavelet packet basis is sought by using the Shannon entropy function and cycle-spinning [11] 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 some periodic signals. The second algorithm is a modified version

Image	Noise (dB)	TI $\mathcal{S}$-Bayes	Curvelet	Directional Cosine	Directional Cosine Packets
barbara	0	14.75	14.59	14.89	15.19
	5	16.08	15.93	16.42	16.85
	10	18.00	17.64	18.44	18.85
	15	21.12	19.64	20.34	20.71
	20	24.92	21.41	21.88	22.10
lena	0	17.06	17.10	17.24	17.75
	5	18.96	18.82	18.99	19.63
	10	21.08	20.83	21.05	21.89
	15	23.49	23.10	23.35	24.26
	20	26.15	25.35	25.60	26.31
grain	0	12.93	13.06	13.10	13.22
	5	13.81	13.66	13.84	14.29
	10	15.87	15.01	15.51	16.18
	15	18.72	17.28	18.00	18.84
	20	22.02	20.05	20.56	21.38

Table 1. The comparative image denoising results in SNR

of curvelet. The curvelet implementation in [4] which uses a much more redundant overcomplete wavelet frame than our MDC, is a "specialised" transform to perform denoising task instead of general-purpose image processing. In order to carry out a fair comparison, here the local ridgelets are placed on the Laplacian pyramid as well in our setting, in order to carry out a fair comparison.

Here we present results from several typical natural images: 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 singularities at different scales; and grain image is a texture image which is considered to be very difficult to compress.

Table 1 gives denoising results in SNR using four algorithms. It seems that directional local cosine transform is the overall winner, although the TI-$\mathcal{S}$-BayesShrink sometimes outperforms in low noise levels, since the BayesShrink tends to optimise the MSE output. However, the results on other three images are more visually pleasing, preserving important directional features on the image. A detailed head-to-head comparison is presented in Figure 2 on barbara. It is obvious that from the TI-$\mathcal{S}$-BayesShrink denoised image, the diagonal stripes 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 curvelets and directional cosine bases, it is not suprising to see it gives better results than its two counterparts. However, it involves considerable amount of extra computations, since the best basis has to be sought for each
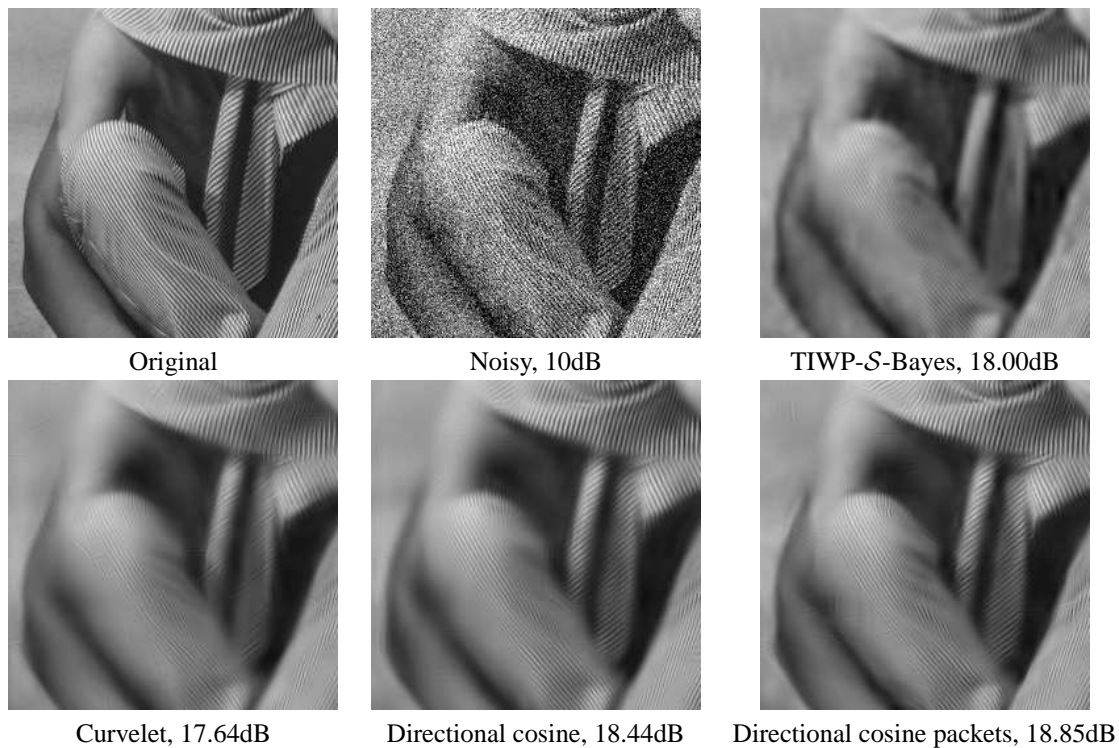
of the Radon slices. It seems that such heavy computation burden is not a fair tradeoff, at least at present, since the directional cosine transform denoised images are visually similar to directional local cosine's results.

4. CONCLUSION

We have presented a set of new bases which have localisation in space and frequency, orientation selectivity and employ a multiresolution pyramidal framework allowing analyse of the image at different scales. In a sense the MDC bases qualify as geometric wavelets and share a lot of similarities with other directional wavelet bases proposed previously, but the construction allows us to capture local directional texture patches and linear features at ease. Unlike the traditional 2D local DCT constructed by tensor-product, the directional cosine basis provides orientation selectivity like the Fourier basis but avoids the unpleasant convergence rate of the Fourier series on intervals.

The effectiveness of the MDC bases was tested against the state-of-the-art translation-invariant wavelet packet based shrinkage method and a directional basis, the so-called *curvelet*. The new bases demonstrated a strong potential in the experiments, often outperforming the other methods. While producing much visually pleasant output than the wavelet packets with optimal threshold, the MDC seems to be able to capture a wider range of directional features than the curvelet, even without the local cosine treatment on Radon slices.

Of course the transform presented here with the Laplacian pyramid for the denoising experiments is only a taste for conducting multiresolution analysis with the directional cosine (packet) bases. A wide variation on the theme is pos-

| Original | Noisy, 10dB | TIWP-$\mathcal{S}$-Bayes, 18.00dB |
| Curvelet, 17.64dB | Directional cosine, 18.44dB | Directional cosine packets, 18.85dB |

Fig. 2. Detailed comparative denoising results on `barbara`

sible, 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. It is our intention to put forward this basis in a general way in this work to popularise its usage in various kinds of image processing tasks.

5. REFERENCES

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

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

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

[4] J. Starck, E.J. Candés, and D.L. Donoho, "The curvelet transform for image denoising," *IEEE Transactions on Image Processing*, vol. 11, no. 6, pp. 670–684, June 2002.

[5] M.N. Do and M. Vetterli, "Contourlet," in *Beyond Wavelets*, G.V. Welland, Ed. Academic Press, 2003.

[6] 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 Scientific Computing.

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

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

[9] N. Rajpoot, Z. Yao, and R. Wilson, "Adaptive wavelet restoration of noisy video sequences," in *Proceedings IEEE ICIP 2004*, Singapore, 2004.

[10] S.G. Chang, B. Yu, and M. Vetterli, "Adaptive wavelet thresholding for image denoising and compression," *IEEE Transaction on Image Processing*, vol. 9, no. 9, pp. 1532–1546, September 2000.

[11] R.R. Coifman and D.L. Donoho, "Translation invariant denoising," in *Wavelets in Statistics*, A. Antoine and G. Oppenheim, Eds., pp. 125–150. Springer, New York, 1995.