

A NEW CONTOURLET TRANSFORM WITH SHARP FREQUENCY LOCALIZATION

Yue Lu and Minh N. Do

Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA
Email: {yuelu, minhdo}@uiuc.edu; Web: www.ifp.uiuc.edu/~{yuelu, minhdo}

ABSTRACT

The contourlet transform was proposed as a directional multiresolution image representation that can efficiently capture and represent singularities along smooth object boundaries in natural images. Its efficient filter bank construction as well as low redundancy make it an attractive computational framework for various image processing applications. However, a major drawback of the original contourlet construction is that its basis images are not localized in the frequency domain. In this paper, we analyze the cause of this problem, and propose a new contourlet construction as a solution. Instead of using the Laplacian pyramid, we employ a new multiscale decomposition defined in the frequency domain. The resulting basis images are sharply localized in the frequency domain and exhibit smoothness along their main ridges in the spatial domain. Numerical experiments on image denoising show that the proposed new contourlet transform can significantly outperform the original transform both in terms of PSNR (by several dB's) and in visual quality, while with similar computational complexity.

Index Terms— Contourlet transform, multiscale pyramid, directional filter banks, image denoising

1. INTRODUCTION

Studying and exploiting the special properties of natural images has been one of the most important tasks in image processing. One key distinguishing feature of natural images is that they have intrinsic geometrical structures, for example, along object boundaries.

Recently, Do and Vetterli [1] proposed the contourlet transform as a directional multiresolution image representation that can efficiently capture and represent smooth object boundaries in natural images. The contourlet transform is constructed as a combination of the Laplacian pyramid [2] and the directional filter banks (DFB) [3]. Conceptually, the flow of operation can be illustrated by Figure 1(a), where the Laplacian pyramid iteratively decomposes a 2-D image into lowpass and highpass subbands, and the DFB are applied to the highpass sub-




Fig. 1. The original contourlet transform. (a) Block diagram. (b) Resulting frequency division.

bands to further decompose the frequency spectrum. Using ideal filters, the contourlet transform will decompose the 2-D frequency spectrum into trapezoid-shaped regions as shown in Figure 1(b).

In practice, when non-ideal filters are used, the resulting contourlets do not have the desired sharp frequency domain localization. Although the majority of the energy in each subband is still concentrated on the ideal support regions, there are also significant amount of aliasing components showing up at locations far away from the desired support. This kind of frequency aliasing is undesirable, since the resulting contourlets in the spatial domain are not smooth along their main ridges and exhibit some fuzzy artifacts. Consequently, this jeopardizes the efficiency of contourlets in representing smooth boundaries in natural images. This phenomenon has also been independently observed by other researchers, e.g., in [4], and through various personal communications to us.

In Section 2, we provide a simple pictorial explanation of the cause of this frequency non-localization problem. As a solution, we propose in Section 3 a new construction of the contourlet transform, in which the non-localization problem is greatly alleviated. Numerical experiments are presented in Section 4 to confirm the superiority of the proposed new construction over the original transform. We conclude the paper in Section 5.

This work was supported by the National Science Foundation under Grant CCR-0237633 (CAREER).




Fig. 2. Block diagram of the contourlet transform with two levels of multiscale decomposition. Gray regions represent the ideal passband support of the component filters. Left: The iterated form. Right: The equivalent parallel form.

2. A PICTORIAL ANALYSIS OF THE FREQUENCY NON-LOCALIZATION PROBLEM

In the left part of Figure 2, we show the block diagram of the contourlet transform with two levels of multiscale decomposition, followed by angular decomposition. Note that the Laplacian pyramid shown in the diagram is a simplified version of its actual implementation [2, 1]. Nevertheless, this simplification serves our explanation purposes satisfactorily. By using the multirate identities, we can rewrite the filter bank into its equivalent parallel form, as shown in the right part of Figure 2. In the following discussions, we will concentrate on channel 2 of the filter bank.

In Figure 3(a), we show a more realistic illustration of one of the directional filters from the DFB, when we use non-ideal filters. Gray regions in the figure represent the ideal passband, and patterned regions represent the aliasing areas concentrated along two parallel lines ($\omega_2 = \pm\pi$).

Two reasons contribute to this aliasing effect. The first one is due to the periodicity of 2-D frequency spectrums of discrete signals. In Figure 3(a), the patterned regions marked by p are actually the transition bands of the wedge-shaped filters, folded back through 2π periodization. The other reason is intrinsic to the frequency partitioning of the DFB. Using the argument of permissible passband supports proposed by Chen and Vaidyanathan [5], we can show that perfect reconstruction and frequency domain localization cannot be achieved simultaneously by a critically-sampled filter bank with the frequency partitioning of the DFB. In other words, since the DFB are critically-sampled and have perfect reconstruction, their component filters must exhibit aliasing components outside of the desired passband regions.

When the DFB is combined with a multiscale decomposition as in the contourlet transform, the aliasing problem becomes a serious issue. For instance, to calculate the equivalent filter of the second channel in Figure 2, the directional filter must first be upsampled by 2 along each dimensions, as shown in Figure 3(b). As a result of the upsampling, the aliasing components (again represented by the patterned-regions in the figure) are folded towards the lowpass regions and concentrated mostly along two lines ($\omega_2 = \pm\frac{\pi}{2}$). Combining the upsampled DFB with the bandpass filter shown in Figure 3(c), we can get the contourlet subband filter for channel 2 in Figure 3(d). We

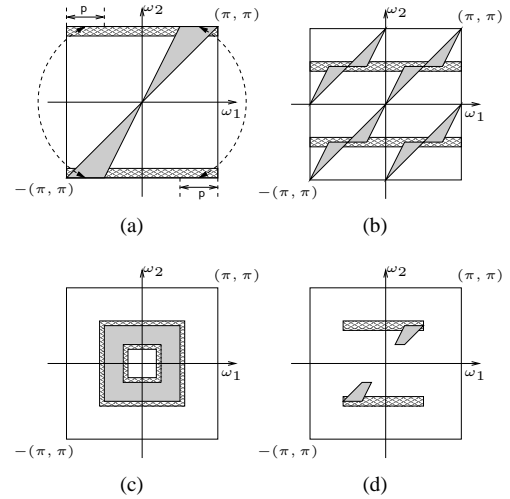


Fig. 3. Illustration of the frequency domain aliasing problem of the contourlet transform. Gray regions represent the ideal passband support. Patterned regions represent the aliasing components or transition bands. (a) One directional filter. (b) The directional filter after being upsampled by 2 along each dimension. (c) A bandpass filter from the Laplacian pyramid. (d) The resulting contourlet subband.

can see that contourlets are not localized in frequency, with substantial amount of aliasing components outside of the desired trapezoid-shaped support.

3. A NEW CONTOURLET CONSTRUCTION WITH SHARP FREQUENCY LOCALIZATION

3.1. Construction

In Figure 4, we show a new construction of the contourlet transform. We still use the DFB for directional decomposition. However, an important distinction from the original contourlet transform is that, instead of using the Laplacian pyramid, we employ a new pyramid structure for the multiscale decompo-




Fig. 4. The block diagram of the new contourlet transform. Only the analysis part is shown, while the synthesis part is exactly symmetric.

sition, which is conceptually similar to the one used in the steerable pyramid [6].

In the diagram, we use $L_i(\omega)$ ($i = 0, 1$) to represent the lowpass filters and $D_i(\omega)$ ($i = 0, 1$) to represent the highpass filters in the multiscale decomposition, with $\omega \stackrel{\text{def}}{=} (\omega_1, \omega_2)$. The DFB is attached to the highpass branch at the finest scale and bandpass branches at all coarser scales. The lowpass filter $L_0(\omega)$ in the first level is downsampled by d along each dimension, with d being a number to be determined shortly, and the lowpass filter $L_1(\omega)$ in the second level is downsampled by $(2, 2)$. To have more level of decomposition, we can recursively insert at point a_{n+1} a copy of the diagram contents enclosed by the dashed rectangle.

As an important difference from the Laplacian pyramid shown in Figure 2, the new multiscale pyramid can employ a different set of lowpass and highpass filters for the first level and all other levels. As will be seen shortly, this is a crucial step in reducing the frequency-domain aliasing of the DFB.

In our current implementation, we specify the lowpass filters $L_i(\omega)$ ($i = 0, 1$) in the frequency domain as $L_i(\omega) = L_i^{1d}(\omega_1) \cdot L_i^{1d}(\omega_2)$, and $L_i^{1d}(\omega)$ is a 1-D lowpass filter with passband frequency $\omega_{p,i}$ and stopband frequency $\omega_{s,i}$ and a smooth transition band, defined as

$$L_i^{1d}(\omega) = \begin{cases} 1 & \text{for } |\omega| \leq \omega_{p,i}, \\ \frac{1}{2} + \frac{1}{2} \cos \frac{(|\omega| - \omega_{p,i})\pi}{\omega_{s,i} - \omega_{p,i}} & \text{for } \omega_{p,i} < |\omega| < \omega_{s,i}, \\ 0 & \text{for } \omega_{s,i} \leq |\omega| \leq \pi, \end{cases} \quad (1)$$

for $|\omega| \leq \pi$ and $i = 0, 1$.

Assuming the aliasing introduced by the downsampling operations can be completely cancelled, we can simplify the perfect reconstruction condition for the multiscale pyramid as

$$|L_i(\omega)|^2 + |D_i(\omega)|^2 \equiv 1, \quad \text{for } i = 0, 1. \quad (2)$$

Once we have specified the lowpass filters, the highpass filters $D_i(\omega)$ can be obtained from (2) to ensure perfect reconstruction.

3.2. Parameter Selection

According to different choices of d , we can have three different variants for the new contourlet construction.

When $d = 1$: In this case, the lowpass filter in the first level is not downsampled. The constraints on $\omega_{p,i}$ and $\omega_{s,i}$ are as follows. First, to achieve an approximate octave-band decomposition as the Laplacian pyramid does, we need

$$\frac{\omega_{p,0} + \omega_{s,0}}{2} = \frac{\pi}{2} \quad \text{and} \quad \frac{\omega_{p,1} + \omega_{s,1}}{2} = \frac{\pi}{4}.$$

Secondly, to simplify the perfect reconstruction condition, we want the aliasing introduced by the downsampling operations in the second and all subsequent levels can be completely cancelled, and hence $\omega_{s,1} < \frac{\pi}{2}$.

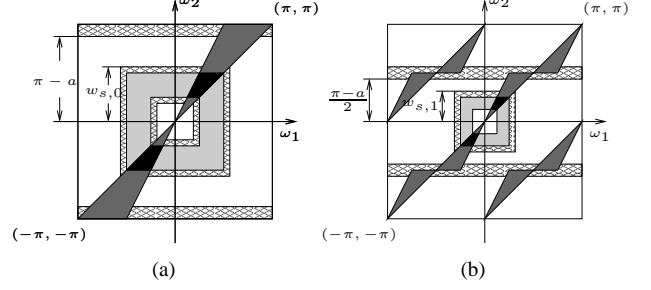


Fig. 5. Illustration of how aliasing components are cancelled in the new construction with $d = 1$. (a) The case for the second level. (b) The case for the third and all subsequent levels.

Most importantly, we want the new multiscale pyramid to cancel the aliasing components of the DFB, and this can be illustrated pictorially. In Figure 5, we show the bandpass filter (in light-gray) from the multiscale decomposition, together with a directional filter (in dark-gray). The aliasing and transition regions are shown as patterned regions. Let a denote the maximum width of the aliasing component of the DFB filter. As we can see from the figure, as long as $\omega_{s,0} \leq \pi - a$ and $\omega_{s,1} \leq \frac{\pi - a}{2}$, the aliasing component of DFB can be completely cancelled out by the bandpass filter. The resulting contourlet filter (shown in black) will be localized in the frequency domain, corresponding to a single direction.

When $d = \frac{3}{2}$, the lowpass filter $L_0(\omega)$ in the first level is downsampled by a non-integer factor of 1.5 along each dimension. This is implemented as upsampling by 2 followed an anti-aliasing filter and then downsampling by 3.

When $d = 2$, the lowpass filters are always downsampled by $(2, 2)$, just like the Laplacian pyramid. Consequently, this version of the new contourlet transform has the same redundancy ratio of around 1.33 as the original one.

For the latter two cases, the conditions on $\omega_{p,i}$ and $\omega_{s,i}$ can be inferred similarly as in the case for $d = 1$, and the details are omitted here. To summarize, we list the parameters used by the three different configurations in Table 1, together with the resulting redundancy ratio.

Table 1. Parameters and redundancy ratios of different configurations of the new contourlet transform

d	$\omega_{p,0}$	$\omega_{s,0}$	$\omega_{p,1}$	$\omega_{s,1}$	Redundancy
1	$\pi/3$	$2\pi/3$	$\pi/6$	$\pi/3$	≈ 2.33
1.5	$5\pi/14$	$9\pi/14$	$19\pi/72$	$35\pi/72$	≈ 1.60
2	$4\pi/21$	$10\pi/21$	$4\pi/21$	$10\pi/21$	≈ 1.33

4. NUMERICAL EXPERIMENT

In this section, we will use ‘‘Contourlet’’ to denote the original transform described in [1]. We use ‘‘Contourlet-2.3’’, ‘‘Contourlet-1.6’’, ‘‘Contourlet-1.3’’ to denote the three different variants of the proposed new contourlet transform, with the numbers cor-




Fig. 6. Comparison of basis images. Top row: frequency domain. Bottom row: spatial domain. From left to right: Contourlet, Contourlet-2.3, Contourlet-1.6, Contourlet-1.3.

Table 2. PSNR values of the denoised images.

σ	Lena			Peppers		
	30	40	50	30	40	50
Contourlet	26.82	25.51	24.51	26.59	25.27	24.15
Contourlet-1.3	28.79	27.66	26.80	28.63	27.48	26.60
Contourlet-2.3	29.44	28.08	27.03	29.13	27.73	26.71

responding to their respective redundancy ratios.

4.1. Basis Images

In Figure 6, we compare the frequency and spatial domain basis images of the original contourlet transform with those from the new contourlet transform proposed in this paper. As we can see from Figure 6(a), the original contourlet transform suffers from the frequency non-localization problem. In sharp contrast, all three variants of the new construction produce basis images that are well-localized in the frequency domain, as shown by Figure 6(b)-(d). The improvement in the frequency localization is also reflected in the spatial domain. As shown in Figure 6(e) - (h), the spatial regularity of contourlets can be greatly improved by using the new construction.

4.2. Denoising

In this experiment, we compare the denoising performance of the proposed new contourlet transform with that of the original transform, by using the standard hard thresholding denoising method.

Table 2 shows the PSNR (in dB) of the denoised images by using different transforms. Although Contourlet-1.3 has the same redundancy ratio and similar computational cost as the original contourlet transform, it outperforms the latter by more than 2 dBs. Even greater improvement is achieved by Contourlet-2.3, due to its increased redundancy.

Figure 7 shows a “zoom-in” comparison of the denoised “Peppers” images. Due to the improved regularity of the basis elements, the reconstructed images from Contourlet-1.3 is visually much better than the one from Contourlet.




Fig. 7. Comparison of denoised “Peppers” images by using Contourlet (left) and Contourlet-1.3 (right).

Recently, the proposed new contourlet transform has been employed in a more advanced denoising algorithm [7], and achieves promising results.

5. CONCLUSION

In this paper, we proposed a new construction for the contourlet transform. Compared with the old version, the new construction produces basis images with much better localization in the frequency domain and regularity in the spatial domain. In applications such as image denoising, we show the proposed new contourlet construction significantly outperforms the original transform.

6. REFERENCES

- [1] M. N. Do and M. Vetterli, “The contourlet transform: an efficient directional multiresolution image representation,” *IEEE Trans. Image Proc.*, vol. 14, no. 12, December 2005.
- [2] P. J. Burt and E. H. Adelson, “The Laplacian pyramid as a compact image code,” *IEEE Trans. Comm.*, vol. COM-31, pp. 532–540, April 1983.
- [3] R. H. Bamberg and M. J. T. Smith, “A filter bank for the directional decomposition of images: theory and design,” *IEEE Trans. Signal Proc.*, vol. 40, no. 4, pp. 882–893, April 1992.
- [4] E. J. Candès, L. Demanet, D. L. Donoho, and L. Ying, “Fast discrete curvelet transforms,” Tech. Rep., Applied and Computational Mathematics, California Institute of Technology, 2005.
- [5] T. Chen and P. P. Vaidyanathan, “Considerations in multidimensional filter bank design,” in *Proc. IEEE Int. Symp. Circ. and Syst.*, Chicago, Illinois, USA, May 1993, pp. 643–646.
- [6] E. P. Simoncelli, W. T. Freeman, E. H. Adelson, and D. J. Heeger, “Shiftable multiscale transforms,” *IEEE Trans. Inform. Th., Special Issue on Wavelet Transforms and Multiresolution Signal Analysis*, vol. 38, no. 2, pp. 587–607, March 1992.
- [7] M. Elad, B. Matalon, and M. Zibulevsky, “Image denoising with shrinkage and redundant representations,” in *Proc. IEEE Int. Conf. on Computer Vision and Pattern Recognition*, New York, USA, June 2006.