

Image Interpolation Using Multiscale Geometric Representations

Nickolaus Mueller, Yue Lu and Minh N. Do

Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign

ABSTRACT

With the ever increasing computational power of modern day processors, it has become feasible to use more robust and computationally complex algorithms that increase the resolution of images without distorting edges and contours. We present a novel image interpolation algorithm that uses the new contourlet transform to improve the regularity of object boundaries in the generated images. By using a simple wavelet-based linear interpolation scheme as our initial estimate, we use an iterative projection process based on two constraints to drive our solution towards an improved high-resolution image. Our experimental results show that our new algorithm significantly outperforms linear interpolation in subjective quality, and in most cases, in terms of PSNR as well.

Keywords: Image interpolation, contourlet transform, geometric regularity, directional multiresolution image representation

1. INTRODUCTION

Image interpolation is the process of increasing the resolution of a given image. One such application to image interpolation can be found in streaming video websites such as YouTube, which often store video at low resolutions (e.g. 352×288 pixels CIF format) for various reasons. The problem presents itself in that users often wish to expand the size of the video to watch at full screen with resolutions of 1024×768 or higher, and this process requires that the images be interpolated to the higher resolution. Another application comes from the emergence of HDTV displays. To better utilize the display prowess of these state-of-the-art viewing devices, input signals coming from a low-resolution source (e.g. legacy standard definition video) must first be converted to higher resolutions through interpolation. Although the temporal correlation in video signals allows the interpolation to be done using multiframe super-resolution techniques, for our purposes we will examine interpolation algorithms on a still frame image in this paper.

The most common methods used in practice, such as bilinear and bicubic interpolation, require only a small amount of computation. However, because they are based on an oversimplified slow varying image model, these simple methods often produce images with various problems along object boundaries, including aliasing, blurring, and zigzagging edges. Various interpolation algorithms have been proposed to reduce edge artifacts, aiming at obtaining images with regularity (i.e. smoothness) along edges. In one of the earliest papers on the subject, Jensen and Anastassiou¹ propose to estimate the orientation of each edge in the image by using projections onto an orthonormal basis and the interpolation process is modified to avoid interpolating across the edge. Allebach and Wong² propose to use an estimate of the high-resolution edge map to iteratively correct the interpolated pixels. Another low-cost, yet quite effective method called data dependent triangulation³ performs interpolation by dividing every four-pixel square into two triangles and restricts interpolation to within the triangles.

Instead of explicitly estimating edges, Li and Orchard⁴ propose a statistical estimation method that tunes interpolation coefficients according to local edge structures. While this method produces superior results, its computational complexity is prohibitive due to the large window size used to estimate local covariances. Improved interpolation methods have also been cast as a non-linear PDE problem,⁵ where the algorithm attempts to satisfy

Send correspondence to N. Mueller (nmueller@uiuc.edu).

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

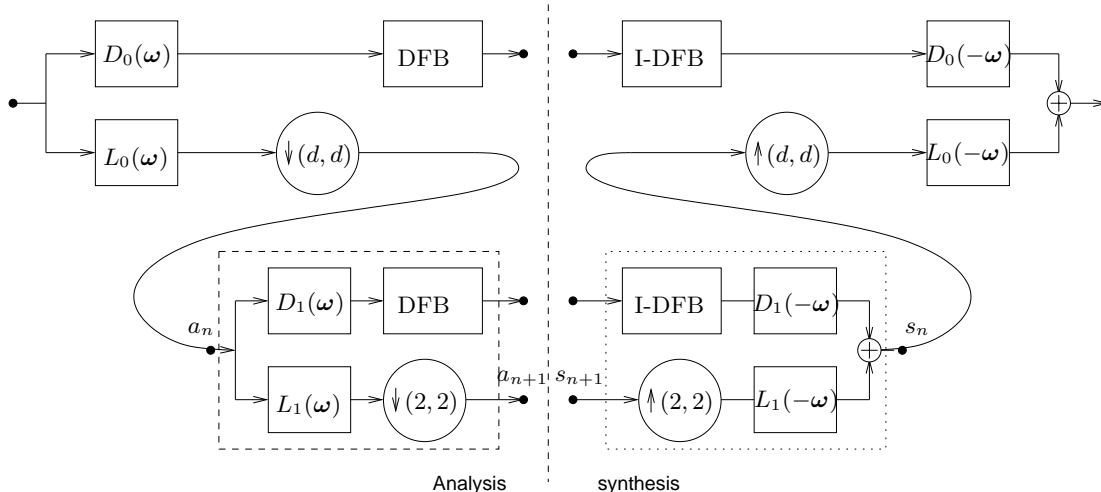


Figure 1. The block diagram of the new contourlet transform. The directional filter banks (DFB) are attached to the highpass subbands of the multiscale pyramid at each level.

smoothness and orientation constraints. Recently, methods have been proposed which perform interpolation in a transform (e.g. wavelet) domain.^{6,7} These algorithms assume the low-resolution image to be the lowpass output of the wavelet transform and utilize dependence across wavelet scales to predict the “missing” coefficients in the more detailed scales.

In this paper, we propose a novel image interpolation algorithm based on the new contourlet transform,⁸ which provides an efficient multiscale geometric representation for natural images. By enforcing a sparsity constraint on the coefficients of the contourlet transform, we provide an efficient mechanism that helps improve the regularity along edges in the resulting images. Furthermore, to satisfy the observation constraint provided by the given low-resolution image, we employ an iterative projection procedure in the wavelet domain, originally proposed by Guleryuz.⁹

The rest of the paper is organized as follows. Section 2 discusses the new contourlet transform and provides motivation for its use in our algorithm. In Section 3, we discuss motivation for our two constraints and construct an algorithm to improve linear interpolation. We report the results of our experiments in Section 4 and conclude the paper in Section 5.

2. THE NEW CONTOURLET TRANSFORM

The contourlet transform was proposed by Do and Vetterli¹⁰ as a directional multiresolution transform that can efficiently capture and represent smooth object boundaries in natural images. In this paper, we employ a new version of the contourlet transform, recently proposed by Lu and Do.⁸ As an improvement over the original version, the basis images from the new transform are sharply localized in the frequency domain and exhibit smoothness along their main ridges in the spatial domain. In practice, this improved regularity of basis images helps reduce the artifacts in processed images, consequently making the new contourlet transform a more favorable choice for various image processing applications.

We show in Figure 1 the construction of the new contourlet transform. We use the directional filter banks¹¹ (denoted by DFB and I-DFB for the analysis and synthesis parts, respectively) to achieve directional decomposition. Instead of using the Laplacian pyramid¹² as in the original contourlet transform, we employ a new pyramid structure for the multiscale decomposition, which is conceptually similar to the one used in the steerable pyramid.¹³

In Figure 1, 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




Figure 2. (a) One basis image of the new contourlet transform in the frequency domain. (b) The same basis image, but in the spatial domain.

branch at the finest scale and bandpass branches at all coarser scales. The lowpass filter $L_0(\boldsymbol{\omega})$ in the first level is downsampled by d along each dimension, with d being a number chosen from 1 and 2, and the lowpass filter $L_1(\boldsymbol{\omega})$ in the second level is downsampled by $(2, 2)$. To have more levels of decomposition, we can recursively insert at point a_{n+1} a copy of the diagram contents enclosed by the dashed rectangle in the analysis part, and at point s_{n+1} a copy of the diagram contents enclosed by the dotted rectangle in the synthesis part.

As an important difference from the Laplacian pyramid used in the original contourlet transform, the new multiscale pyramid shown in Figure 1 can employ a different set of lowpass and highpass filters for the first level and all other levels. This turns out to be a crucial step in reducing the frequency-domain aliasing of the DFB. We leave the detailed explanation for this issue as well as the specification of the filters $L_i(\boldsymbol{\omega})$ and $D_i(\boldsymbol{\omega})$ to reference.⁸

Figure 2 shows one contourlet basis image in the frequency and spatial domains. In general, contourlet basis images in space are localized contour-like functions (hence the name of the transform), which are smooth along their main ridges but oscillatory across the normal directions. The entire collection of basis images from the transform spans a large range of directions, scales, and spatial locations. We can envision that these basis elements can be used to efficiently capture and represent different object boundaries in natural images.

3. INTERPOLATION BY ITERATING CONSTRAINTS

3.1. Main Ideas

We look to design an algorithm that will upsample a given low-resolution image (by factor of 2^n with $n = 1, 2, \dots$, for the sake of simplicity) that will improve the quality of interpolation in edge regions of the image. Like several papers that perform interpolation in the transform domain, we assume our low-resolution image to be the lowpass output of an n -level wavelet transform applied to the ground truth high-resolution image (to which we do not have access).

The main idea of our algorithm is to alternately enforce two constraints, which is similar to the technique proposed in reference.⁹ The first constraint comes from the assumption that the given low-resolution image is the downsampled output of the lowpass anti-aliasing filter in a wavelet transform. This observation constraint can be enforced by requiring the high-resolution image to have our given low-resolution image as its lowpass output of the wavelet transform. The second constraint is based on a model for natural images. Since the contourlet transform described in Section 2 provides us with a sparse image representation well-suited to preserving contours and edges, we assume that the unknown high-resolution image should be sparse in the contourlet transform domain. As shown in Section 3.3, we enforce this constraint through a hard-thresholding of the contourlet coefficients.

3.2. Observation Constraint

Figure 3 illustrates the observation constraint, in which we assume the given low-resolution image, denoted by \boldsymbol{x}_L , is the lowpass subband of an n -level wavelet transform of the *unknown* high-resolution image \boldsymbol{x} , while all the coefficients in the highpass subbands have been discarded. As a simple way to get an estimate $\hat{\boldsymbol{x}}_0$ of the high-resolution image, we can take the inverse wavelet transform by keeping \boldsymbol{x}_L as the lowpass band and zeropadding




Figure 3. The observation model. We assume the given low-resolution image x_L is the lowpass subband of a wavelet transform applying on the unknown high-resolution image x . A simple estimation of the high-resolution image, denoted by \hat{x}_0 , can be obtained by an inverse wavelet transform, which uses x_L as the lowpass and zero coefficients for all highpass subbands.

all highpass subbands. It has been observed by several authors^{14,15} that this simple scheme often outperforms other more sophisticated interpolation methods, which do not take into account the anti-aliasing operation in the image generation process.

Note that the set of all images whose lowpass wavelet subbands are equal to the given low-resolution image x_L forms a linear variety. Consequently, for any given image y , we can calculate the best approximation (in L^2 norm) to y , subject to our observation constraint, through orthogonal projection. Let \mathbf{W} and \mathbf{W}^{-1} represent the forward and inverse wavelet transforms, respectively; denote \mathbf{P} as the diagonal projection matrix of 1's and 0's that keeps the lowpass wavelet coefficients and zeros out the high frequency subband coefficients, and let $\mathbf{P}^\perp = \mathbf{I} - \mathbf{P}$. If we use orthonormal wavelet transforms, then the projection of any image y onto the constraint set can be calculated by

$$\hat{y} = \mathbf{W}^{-1}(\mathbf{P}^\perp \mathbf{W} y + \mathbf{P} \mathbf{W} \hat{x}_0), \quad (1)$$

where \hat{x}_0 is the estimation of the high-resolution image obtained as in Figure 3.

3.3. Sparsity Constraint

The contourlet transform provides us with a sparse image representation (i.e. there are many more insignificant coefficients than significant ones), and these coefficients well represent the edge regions of the image. Note that there can be several different ways to enforce the sparseness constraint (e.g. see the pursuit schemes described in reference¹⁶). For the sake of simplicity, we choose to use a direct hard-thresholding scheme in our proposed algorithm.

Intuitively, we view our estimate to the high-resolution image as a noisy version of the true image. Enforcing our sparsity constraint works to denoise the estimation of the interpolated signal while retaining the important coefficients near edges. Denote \mathbf{C} and \mathbf{C}^{-1} as the forward and inverse contourlet transforms, respectively; \mathbf{D}_T as the diagonal matrix that, given some threshold value T , zeros out insignificant coefficients in the coefficient vector whose absolute values are smaller than T ; \hat{x} the noisy high-resolution image; and \tilde{x} the denoised high-resolution image. The sparseness constraint by hardthresholding can be written as

$$\tilde{x} = \mathbf{C}^{-1} \mathbf{D}_T \mathbf{C} \hat{x}. \quad (2)$$

3.4. Iterative Contourlet-Based Image Interpolation

We show in Figure 4 the block diagram of the proposed iterative contourlet-based image interpolation scheme, which can be summarized as follows:

1. We start our algorithm by taking x_0 , obtained by the simple wavelet interpolation shown in Figure 3, as the initial estimate of the high-resolution image.

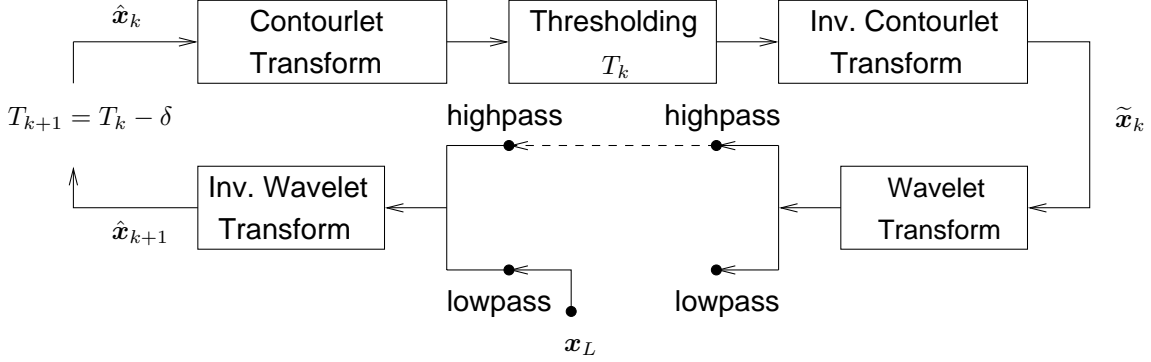


Figure 4. Block diagram of the proposed contourlet-based interpolation algorithm.

2. We then attempt to improve the quality of interpolation, particularly in regions containing edges and contours, by iteratively enforcing the observation constraint as well as the sparseness constraint. Let $\hat{\mathbf{x}}_k$ represent the estimate at the k th step. By combining (1) and (2), the new estimate $\hat{\mathbf{x}}_{k+1}$ can then be obtained by

$$\hat{\mathbf{x}}_{k+1} = \mathbf{W}^{-1}(\mathbf{P}^\perp \mathbf{W} \mathbf{C}^{-1} \mathbf{D}_{T_k} \mathbf{C} \hat{\mathbf{x}}_k + \mathbf{P} \mathbf{W} \hat{\mathbf{x}}_0).$$

3. Following the same principle of the sparseness-based image recovery algorithm proposed by Guleryuz,⁹ we gradually decrease the threshold value T_k by a small amount δ in each iteration, i.e., $T_{k+1} = T_k - \delta$. This has been shown to be effective in circumventing the nonconvexity of the sparseness constraint.
4. Return to step 2 and keep iterating, until the generated images converge or a predetermined maximum iteration number has been reached.

4. EXPERIMENTAL RESULTS

4.1. Implementation

In this section, we report several experiment results of the proposed image interpolation algorithm. For comparison, we also show the results obtained by four other interpolation methods, including the bilinear interpolation, the simple wavelet-based linear interpolation implemented by zero-padding highpass subbands, the data-dependent triangulation method (DDT),³ and the new edge-directed interpolation (NEDI) proposed by Li and Orchard.⁴ Each algorithm has been run in MATLAB, with the NEDI code provided to us by the authors. In all experiments, we use the “symlet” of length 16 for the wavelet transform and the filters specified in reference⁸ for the contourlet transform, whose redundancy ratio is around 2.33 (i.e. $d = 1$ in Figure 1). Meanwhile, the initial threshold in our algorithm has been chosen to be $T_0 = 10$ and is decreased by $\delta = 0.2$ in each iteration, with a maximum of 10 iterations.

We use several standard test images of size 512×512 , including *Lena*, *Gaussian Disc*, *Peppers*, and *Barbara*. To show the true power of the interpolation algorithms, we first downsampled each image by a factor of four and then interpolated the result back to its original size. This provides a better comparison than factor of two interpolation, and is well justified if one compares the ratio of NTSC scan lines (240 per frame) to state-of-the-art HDTV (1080 per frame), which is a factor of 4.5. For each algorithm, we have assumed the presence of an anti-aliasing filter in the downsampling process. Although we can achieve higher PSNR if we do not use such a filter, object geometries become distorted and images have a noticeably worse subjective quality to them.

4.2. Quality Assessment

We show in Figure 5 to Figure 8 the zoom-in comparisons of different algorithms on the test images. We can see that our algorithm outperforms the simple wavelet-based linear interpolation scheme in every test image. Significant improvement is present in the synthetic *Gaussian Disc* image when comparing edge smoothness and




Figure 5. The zoom-in comparison of the *Lena* image. Shown in parentheses are the corresponding PSNR values for the region shown.

regularity of the two algorithms. Our algorithm is also clearly superior to both bilinear and DDT interpolation, as it is able to better reconstruct the high-frequency areas of the image without distorting the smooth regions.

Due to the length of the filters used in the contourlet transform, the proposed algorithm can create some annoying ringing artifacts in some images (e.g. Figure 6(e)). The ringing artifacts are no worse than in linear interpolation, and are at least smoothed out by our algorithm in comparison. Furthermore, ringing artifacts tend to be much more noticeable in synthetic images compared to natural images.

The NEDI algorithm often outperforms our algorithm in edge regions because it does not create the same ringing artifacts. It has been observed, however, that despite NEDI’s thresholding used to apply bilinear interpolation to smooth regions of the image, the algorithm tends to have a smearing effect on smooth and textured regions as a result of its large window size. This effect is very noticeable in Figure 5(e) and Figure 8(e), and our algorithm seems to better reconstruct many regions of those two images.

In Table 1, we provide PSNR values of each algorithm for different images. It is noted that there is not necessarily a correlation between PSNR and visual appearance. In particular, the relatively lower PSNR values of DDT and NEDI are largely due to the fact that these two algorithms do not make use of the anti-aliasing filter in the image generation process. In many cases, our algorithm achieves a gain in PSNR over linear interpolation, similar to the gains provided by DDT and NEDI over bilinear interpolation.

4.3. Computational Complexity

In Figure 9, we provide a plot of the gain in PSNR at each iteration of the algorithm for images that show large overall PSNR gain. For the synthetic *Gaussian Disc* image, we show substantial gains through ten iterations

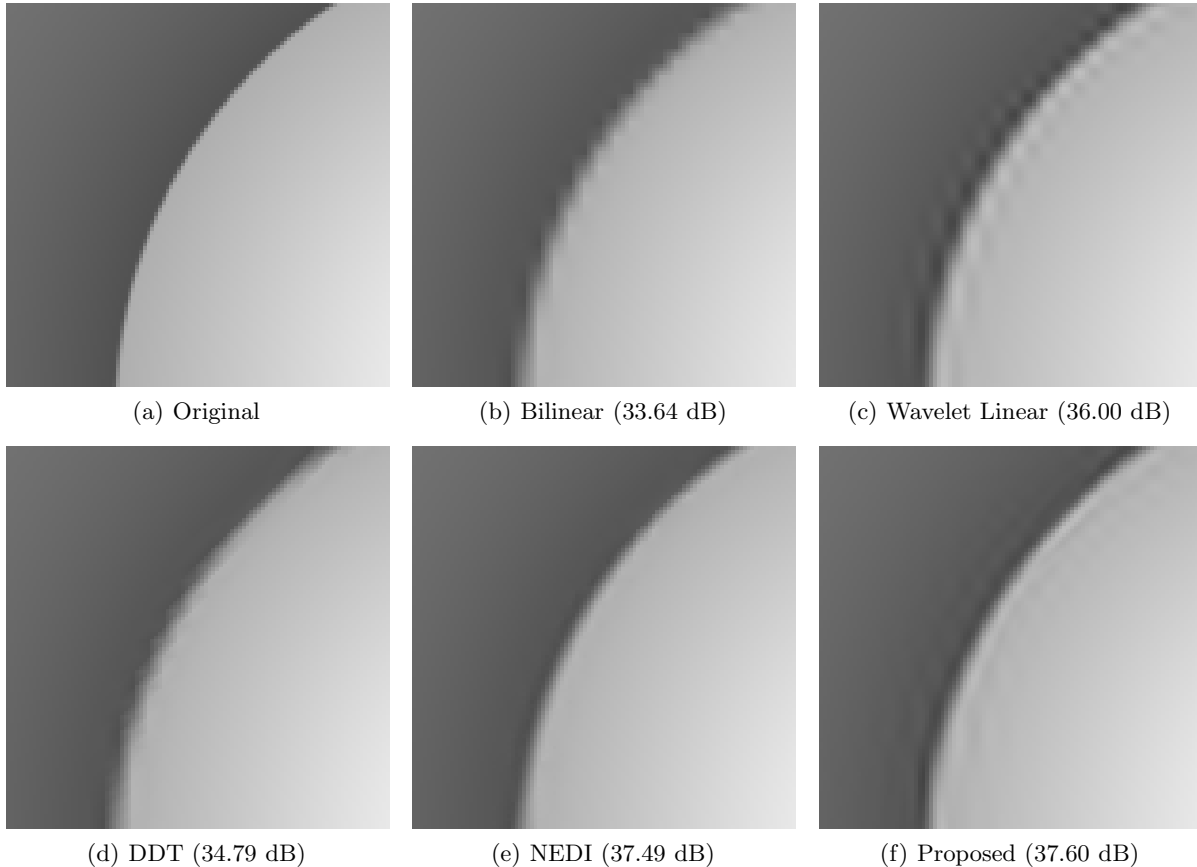


Figure 6. The zoom-in comparison of the *Gaussian Disc* image. Shown in parentheses are the corresponding PSNR values for the region shown.

	PSNR				
	Proposed	Wavelet Linear	DDT	NEDI	Bilinear
Lena	30.29	29.98	29.11	28.97	28.68
Gaussian Disc	42.92	41.38	40.26	42.70	39.41
Peppers	29.39	28.89	28.03	25.92	27.76
Barbara	24.06	24.04	23.82	23.54	23.77

Table 1. PSNR values (in dB) of several images and interpolation algorithms.

of the algorithm. However, for the natural *Lena* and *Peppers* images, we see the algorithm begins to provide diminishing returns after four or five iterations. We offer that although our algorithms were run for 10 iterations to provide our best PSNR values, substantial gains can be made with fewer iterations when minimizing computation time is an issue. We verify this claim in Figure 10 by noting that we still obtain a much improved version of the image by the fourth and fifth iteration. Our algorithm still runs slightly faster than the NEDI algorithm when run for ten iterations, and is considerably faster when run for only five iterations.

5. CONCLUSION

In this paper, we have presented a novel image interpolation algorithm using the new contourlet transform. Our algorithm improves the regularity along edges in the interpolated image by treating each successive approximation to the high-resolution image as a noisy approximation and subsequently enforcing a sparsity constraint on the contourlet coefficients. Based on the given low-resolution image, we also enforce an observation constraint in




Figure 7. The zoom-in comparison of the *Peppers* image. Shown in parentheses are the corresponding PSNR values for the region shown.

the wavelet domain. By iterating this procedure and alternating between constraints, our algorithm converges towards an improved high-resolution image.

We have shown improvement over simple wavelet-based interpolation on a subjective scale, and in many cases with an improvement in PSNR. Our algorithm is competitive to the NEDI algorithm and runs faster. We expect to further reduce the computational complexity of our algorithm with future work, as we try to improve its rate of convergence.

In a future work, we intend to rigorously examine convergence of our algorithm as well as examine the effect of varying threshold values on the rate of convergence and quality of the final image. We also will be searching for ways to minimize the ringing artifacts that are a result of the long filters used in our transforms.

REFERENCES

1. K. Jensen and D. Anastassiou, "Subpixel edge localization and the interpolation of still images," *IEEE Trans. Image Proc.* **4**, pp. 285–295, March 1995.
2. J. Allebach and P. W. Wong, "Edge-Directed interpolation," in *Proc. IEEE Int. Conf. on Image Proc.*, 1996.
3. D. Su and P. Willis, "Image interpolation by pixel level data-dependent triangulation," *Computer Graphics Forum* **23**, pp. 189–201, June 2004.
4. X. Li and M. T. Orchard, "New edge-directed interpolation," *IEEE Trans. Image Proc.* **10**, pp. 1521–1527, October 2001.



Figure 8. The zoom-in comparison of the *Barbara* image. Shown in parentheses are the corresponding PSNR values for the region shown.




Figure 9. PSNR gain vs. number of iterations of the Lena, Gaussian Disc, and Peppers images

5. H. Jiang and C. Moloney, "A new direction adaptive scheme for image interpolation," in *Proc. IEEE Int. Conf. on Image Proc.*, pp. 369–372, 2002.
6. W. K. Carey, D. B. Chang, and S. S. Hermami, "Regularity-preserving image interpolation," *IEEE Trans. Image Proc.* **8**, pp. 1293–1297, September 1999.
7. D. D. Muresan and T. W. Parks, "Prediction of image detail," in *Proc. IEEE Int. Conf. on Image Proc.*,




Figure 10. *Lena* image at each iteration of the algorithm. The generated images begin to converge after 4 or 5 iterations.

pp. 323–326, 2000.

8. Y. Lu and M. N. Do, “A new contourlet transform with sharp frequency localization,” in *Proc. IEEE Int. Conf. on Image Proc.*, (Atlanta, USA), October 2006.
9. O. G. Guleryuz, “Nonlinear approximation based image recovery using adaptive sparse reconstructions and iterated denoising: Part I - theory,” *IEEE Trans. Image Proc.* **15**, pp. 539–554, March 2006.
10. M. N. Do and M. Vetterli, “The contourlet transform: an efficient directional multiresolution image representation,” *IEEE Trans. Image Proc.* **14**, December 2005.

11. R. H. Bamberger and M. J. T. Smith, "A filter bank for the directional decomposition of images: theory and design," *IEEE Trans. Signal Proc.* **40**, pp. 882–893, April 1992.
12. P. J. Burt and E. H. Adelson, "The Laplacian pyramid as a compact image code," *IEEE Trans. Commun.* **COM-31**, pp. 532–540, April 1983.
13. 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* **38**, pp. 587–607, March 1992.
14. O. G. Guleryuz, "Predicting wavelet coefficients over edges using estimates based on nonlinear approximations," in *Proc. IEEE Data Compression Conference*, April 2004.
15. X. Li, "Image resolution enhancement via data-driven parametric models in the wavelet space." under review, 2007.
16. J.-L. Starck, M. Elad, and D. Donoho, "Redundant multiscale transforms and their application for morphological component analysis," *Journal of Advances in Imaging and Electron Physics* **132**, pp. 287–348, 2004.