SEGMENTATION BASED DENOISING USING MULTIPLE COMPACTION DOMAINS

Maneesh Singh, Prakash Ishwar, Krishna Ratakonda and Narendra Ahuja
University of Illinois at Urbana-Champaign
Beckman Institute and ECE Department
405 N. Mathews, Urbana, IL 61801.
{msingh,ishwar}@uiuc.edu

Abstract

In this paper, we propose a novel segmentation-based denoising algorithm. Segmentation yields intrinsically homogeneous and extrinsically heterogeneous regions. A denoising algorithm that uses Multiple Compaction Domains (MCD) is then applied on each of the resulting segments. Such a scheme retains important perceptual information in the segment boundaries while the denoising algorithm operates only on homogeneous segments. Further, the MCD algorithm is demonstrably superior to the classical denoising algorithms using transform domain thresholding [4]. Our algorithm yields better perceptual quality and superior PSNR as compared to MATLAB's adaptive Wiener filter.

1 Introduction

Image acquisition and transmission may yield images that are corrupted with additive noise. Such images can often be realistically modeled as a 2-dimensional signal corrupted with additive white gaussian noise (AWGN) with a known noise variance. If the image is modeled as a WSS signal, then the Weiner filter is provably optimal in minimizing the mean-square error. For most signals, such assumptions are valid locally and the adaptive realization of the Wiener filter [1] is often taken as a benchmark for all denoising algorithms.

Recently, some denoising algorithms have been proposed that use thresholding in (linear unitary) transform domains where the signal has a sparse representation [2]. The key idea here is that while the signal can be compacted into a few transform coefficients (in a suitably chosen transform domain), white noise cannot be thus compacted. Hence, small coefficients are more likely due to noise while large coefficients are due to the signal itself. Thus, thresholding in the transform domain is a reasonable approach.

Such techniques, however, are not tailored to explicitly preserve important structural information. It can be argued that the perceptual quality of the restored image depends significantly on the fidelity of the restoring algorithm in retaining the structure in the image. In this paper, the authors present two ideas for retaining such information. The first idea is segmentation based denoising. Segmentation explicitly represents structure in the shape of the segments. Thus, applying a denoising algorithms to segments preserves the structure represented by segment boundaries. Such ideas were explored in [3].

The second idea makes use of multiple compaction domains for denoising. While this approach does not explicitly make use of image structure but it overcomes some of the shortcomings of the classical denoising algorithms based on thresholding in the transform domain. The philosophy behind this approach is that while significant information can be captured in a transform domain, some information is lost due to thresholding. This information can be partially regained using another well-chosen transform domain which is *complementary* to the first transform domain. A number of such *complementary* transform domains where the signal is well compacted are chosen and a POCS-based technique is employed to iteratively restore the noisy signal from representations in these domains. The details of this Multiple Compaction Domain (henceforth called MCD) denoising algorithm can be found in [4].

It is important to note that linear transforms do not compact edges too well and as a result, like noise, edge energies are scattered among a large number of transform coefficients. Thus, threshold based denoising in the transform domain degrades edges, producing ringing artifacts. While MCD denoising method strives to reconstruct the information lost by thresholding in one domain by using a number of *complemen-*

0-7803-5467-2/99/ $ 10.00 © 1999 IEEE

tary compaction domains, segmentation based denoising preserves some significant structure explicitly by representing them as segment boundaries. Thus, while MCD denoising method implicitly tries to preserve intra-segment information, segmentation preserves the segment boundaries. Applying segmentation before the MCD algorithm has an additional benefit: since segments are intrinsically homogeneous, they are better compacted in the transform domain making the threshold-based denoising work better.

In Section 2, we present the outline of the segmentation based denoising algorithm. Subsection 2.1 describes the segmentation algorithm used. For the segments to be amenable to transform based denoising, it is necessary to extend them to rectangular supports. This is described in subsection 2.2. In subsection 2.3, the multiple compaction domain denoising algorithm is described. Results and conclusions are presented in Section 3.

2 The Denoising Algorithm

The denoising algorithm consists of three key steps. We discuss each of the steps in the following subsections.

2.1 Segmentation

We use the multiscale segmentation scheme given by Ahuja [5]. This algorithm does a multiscale tree-structured segmentation of a given image. The algorithm is fairly robust to Gaussian noise i.e. the presence of noise does not lead to much degradation in the segmentation scheme. This is because the segmentation scheme involves pixel population analysis in order to find regions and the noise being i.i.d., does not superimpose significant structure on the image.

In the final step of our algorithm, as described later, we perform thresholding in multiple transform domains. This thresholding operation is optimal only asymptotically under mild conditions on the signal statistics. In order to satisfy these conditions, we would like our regions to be have a minimum "statistically reliable" size. Let X_i be the pixels in the segment. Then, under the assumption that the noise is zero mean, the following criterion is specified for an accepted level of the validity of the asymptotic results $E(|\sum_{i=1}^{n}(X_i/n)|^2) < 1$ [3]. With Gaussian noise of standard deviation σ, we need the number of pixels in the region to be atleast σ^2.

In order to have segments of acceptable size, we start with the finest scale of segmentation. For any region with size below the threshold, we iteratively find nearest parent node along the path to the root of the tree which is larger than the specified threshold. If such a parent region exists, we use that and replace the subregions at all smaller scales. Else, we merge the region at the coarsest scale with the neighboring region which has minimum disparity in mean gray value from the current region.

Figure 1(b) shows the noisy Lena image and segmentation results before and after post-processing of the segments are shown in Figures 1(c) and 1(d) respectively.

2.2 Region Extension

Segmentation produces regions that may have arbitrary shapes. Linear Transforms, including wavelet transforms, are traditionally defined over square rectangular supports. The usual practice is to enclose the region in a rectangle and to apply zero padding. Defining transforms over regions extended by zero padding yields appreciable ringing artifacts because of the prominent edge at the region boundary. We propose to extend the region into the bounding rectangle so that most of the energy of the extended segment is compacted into *low frequency* transform coefficients.

Let C_a be the set of all signals defined in the bounding rectangle that have the same value as the original region over the region support. Let C_b be the set of all functions with transform coefficients at (some predefined set of) *high frequencies/fine scales* having value zero. In wavelet domain, for one level decomposition, it is equivalent to taking the low-low subimage while zeroing out the other subimages. Both these sets are closed and convex. We define our extended region as the one that lies in the intersection of these sets. We use POCS to find one such region in the intersection of these sets [6].

2.3 Multiple Domain Segment Denoising

Once we get the extended segments, we denoise each segment independently of the others. The denoising algorithm that we use is the Multiple Compaction Domain (MCD) algorithm [4]. We use two wavelet domains for complementary processing. In each domain we define a *confidence tube* of *radius* δ as the set of signals with transform coefficients s_i satisfying $|s_i - d_i| \leq \delta_i$, where d_i are the transform coefficients of the noisy signal and $\delta_i = \delta > 0$ if $|d_i| < \lambda = \sigma\sqrt{2\ln(N)}$ and $\delta_i = 0$ otherwise. By defining $\delta_i = 0$ for large coefficients, we force the estimates to agree with "reliable" component of the data. The confidence tube in each domain is both closed and convex. We use POCS to get the image that lies in the intersection of both convex sets.

POCS converges to a solution in the intersection of the two convex sets (if the intersection is non-empty) [6]. A reasonable choice for the initial point is the hard–thresholded signal estimate in any domain. An equivalent choice is the zero signal. The details of the MCD denoising algorithm can be found in [4].

3 Results

We carried out simulations on 512×512 Lena (Figure 1). The original Lena image (Figure 1(a)) is corrupted by AWGN with mean 0 and $\sigma = 15$. For region extension, we use a 2-level decomposition in DB3 where DBn represents Daubechies filter with n vanishing moments. DB3 and DB4 are the two wavelet domains used for segment denoising using MCD algorithm. Result for the MATLAB's 3×3 Adaptive Weiner Filter is shown in Figure 1(e). The result of our algorithm is given in Figure 1(f). One can see that our reconstruction preserves features better specially in Lena's hat and hair. The edges are better preserved and ringing artifacts are reduced along the region boundaries. Also, though not obvious in these images, perceptual quality of our reconstruction looks better on the computer monitors. Further, the Adaptive Weiner Filter gives a PSNR of 31.24dB while our method gives a PSNR of 31.3dB.

The results show that our algorithm is superior to MATLAB's 3×3 adaptive Wiener filter both in terms of PSNR as well as the perceptual quality.

Acknowledgments

The authors would like to thank Prof. Pierre Moulin[1] for many valuable discussions and suggestions.

References

[1] D. T. Kuan, A. A. Sawchuk, T. C. Strand and P. Chavel, "Adaptive Noise Smoothening Filter for Images with Signal-dependent Noise," *IEEE Trans. on Patt. Anal. and Mach. Intell.*, Vol. 2, pp 165–178, 1980.

[2] D. L. Donoho and I. Johnstone, "Ideal Spatial Adaptation by Wavelet Shrinkage," *Biometrika*, Vol. 81, pp. 425–455, 1994.

[3] K. Ratakonda and N. Ahuja, "Restoring Image Quality Through Structure Preserving De-noising," *ACCV*, Vol. 2, pp. 33-40, 1998.

[4] P. Ishwar, K. Ratakonda, P. Moulin and N. Ahuja, "Image Denoising Using Multiple Compaction Domains," *ICASSP*, Vol. 3, pp. 1889-92, 1998.

[5] N. Ahuja, "A Transform for the Detection of Multi-Scale Structure," *IEEE Transactions on PAMI*, December 1996.

[6] H. Stark and M. I. Sezan, "Image Processing Using Projection Methods," *Real-Time Optical Information Processing*, pp. 185–232, Academic Press, 1994.

[1] EE Dept., Univ. of Illinois at Urbana-Champaign

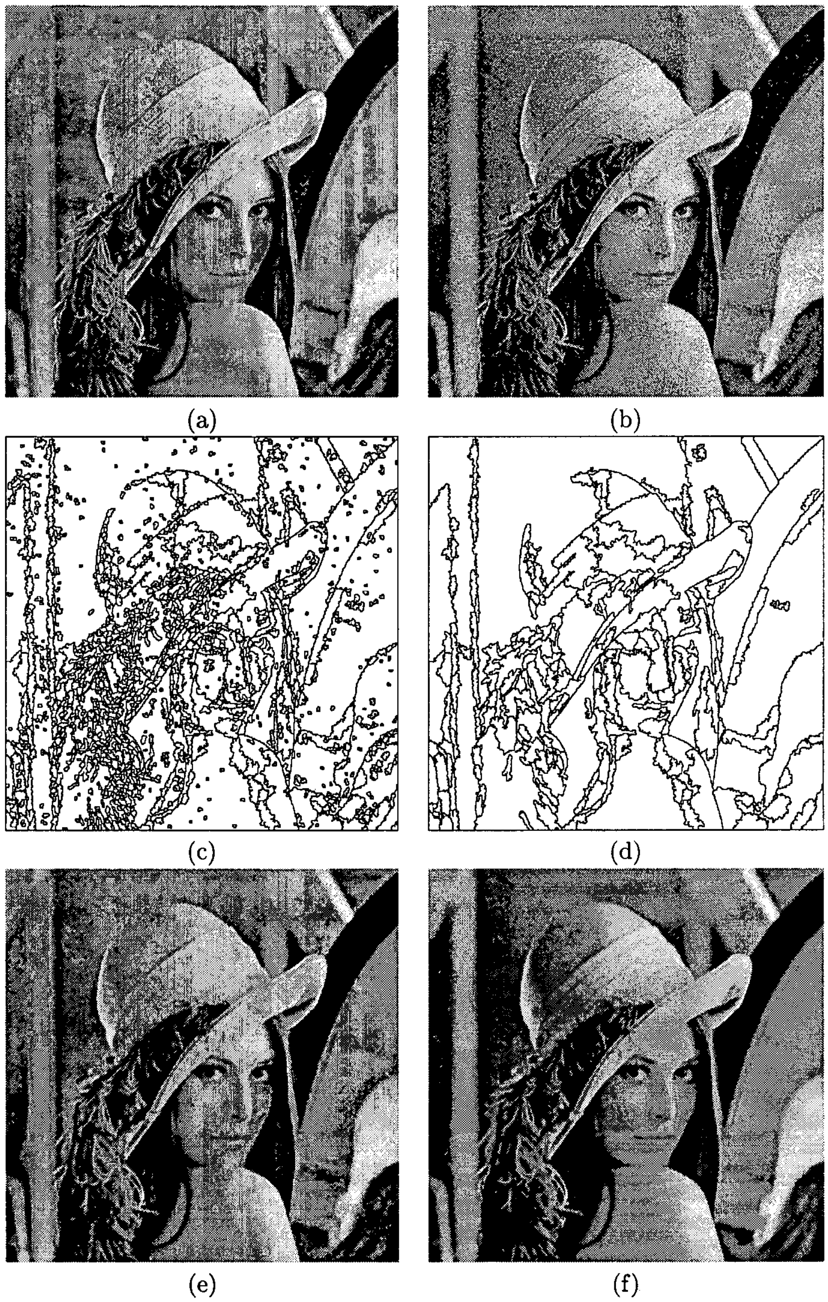

(a) (b) (c) (d) (e) (f)

Figure 1: (a) Original Lena image (b) Lena corrupted by AWGN, s.d. 15 (c) Segmentation at the finest scale (d) Segmentation after post processing using multiscale segmentation data (e) MATLAB's spatial Wiener filter (3×3 window) (f) Reconstructed Lena using segmentation based MCD algorithm