SURE-LET MULTICHANNEL IMAGE DENOISING: UNDECIMATED WAVELET THRESHOLDING

*Florian Luisier**

Biomedical Imaging Group
Ecole Polytechnique Fédérale de Lausanne

Thierry Blu

Dept of Electronic Engineering
The Chinese University of Hong Kong

ABSTRACT

We propose an extension of the recently devised SURE-LET grayscale denoising approach for multichannel images. Assuming additive Gaussian white noise, the unknown *linear* parameters of a transform-domain *pointwise multichannel* thresholding are *globally* optimized by minimizing Stein's unbiased MSE estimate (SURE) in the image-domain.

Using the undecimated wavelet transform, we demonstrate the efficiency of this approach for denoising color images by comparing our results with two other state-of-the-art denoising algorithms.

Index Terms— Multichannel image denoising, multichannel wavelet thresholding, multichannel SURE minimization

1. INTRODUCTION

Transform-domain processing has become a standard procedure for efficient image denoising, where the distorting noise is often assumed as additive and Gaussian. In particular, multiresolution tools such as the wavelet transform [1] often provide a *sparse* signal representation that makes relatively simple operations —such as thresholding— very efficient [2, 3, 4, 5].

Recently, we developed a general procedure to optimize a large class of processing applied in a possibly redundant *linear* transform-domain [5]. It combines the *image-domain* minimization of Stein's unbiased risk estimate [6] (SURE) with a linear expansion of thresholds (LET). Thus, we called this approach SURE-LET. The very competitive results obtained for grayscale images have encouraged us to extend the SURE-LET strategy to multichannel image denoising. For such signals, the application of existing grayscale denoisers on each separate channel is often sub-optimal due to the potentially strong interchannel similarities. Significantly better results are obtained by devising specific multichannel algorithms such as [3, 4] or "decorrelating" the data. The latter solution has been widely exploited in color image denoising, where the data are first projected on a more suitable color space (e.g. an orthogonal luminance-chrominance (YUV) space in [7]) than the standard red-green-blue (RGB) representation.

In this paper, we first recall the two key ingredients, SURE and linear parametrization, of the SURE-LET approach and give an explicit expression of a Stein-like unbiased estimate of the mean-squared error (MSE) between a denoised multichannel image and its underlying unknown noise-free version. Secondly, we propose a simple pointwise multichannel thresholding function, which exploits the interchannel similarities to efficiently reduce the noise while preserving most of the detail. Finally, we show promising results in color image denoising, using the undecimated wavelet transform (UWT).

2. CONTEXT

We denote a N-pixel image with C channels, typically $C = 3$ for RGB images, as a $C \times N$ matrix whose columns are the channel values of each pixel:

$$\mathbf{x} = [\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N], \text{ where } \mathbf{x}_n = [x_{n,1}, x_{n,2}, \ldots, x_{n,C}]^{\mathrm{T}}$$

This image is assumed to be corrupted by an additive "channelwise" Gaussian white noise $\mathbf{b} = [\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_N]$ with known interchannel covariance matrix $\mathbf{\Gamma}$, i.e.

$$\mathscr{E}\left\{\mathbf{b}_n \mathbf{b}_{n'}^{\mathrm{T}}\right\} = \mathbf{\Gamma}\, \delta_{n-n'}$$

where $\mathscr{E}\{\cdot\}$ stands for the mathematical expectation.

The resulting noisy image is denoted by $\mathbf{y} = [\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_N]$. Thus we can write the following:

$$\mathbf{y} = \mathbf{x} + \mathbf{b}. \tag{1}$$

We assume that the only source of randomness is the noise $\mathbf{b}$. As a consequence, the noisy image $\mathbf{y}$ is also random, whereas the original noise-free image $\mathbf{x}$ is assumed *not* to be drawn from some random process.

Hence, denoising the image $\mathbf{y}$ consists in finding an estimate $\hat{\mathbf{x}}$ of $\mathbf{x}$ as a function of only the noisy observation $\mathbf{y}$.

*This work was supported by the Center for Biomedical Imaging (CIBM) of the Geneva–Lausanne Universities and the EPFL, the foundations Leenaards and Louis-Jeantet, as well as by the Swiss National Science Foundation under grant 200020-109415.

The quality of the denoising will be evaluated using the classic *Peak Signal-to-Noise Ratio* (PSNR)

$$\mathrm{PSNR} = 10\log_{10}\left(\frac{255^2}{\mathrm{MSE}}\right)\ \mathrm{dB} \tag{2}$$

which involves the standard *Mean-Squared Error* (MSE) criterion:

$$\mathrm{MSE} = \frac{1}{CN}\,\mathrm{Tr}\left\{(\hat{\mathbf{x}}-\mathbf{x})(\hat{\mathbf{x}}-\mathbf{x})^{\mathrm{T}}\right\} \tag{3}$$

3. MULTICHANNEL SURE-LET

3.1. Principle

In this paper, we will focus on linear transform-domain denoising, in particular:

$$\hat{\mathbf{x}} = \mathbf{F}(\mathbf{y}) = \big(\mathbf{R}\boldsymbol{\Theta}(\underbrace{\mathbf{D}\mathbf{y}^{\mathrm{T}}}_{\mathbf{w}})\big)^{\mathrm{T}} \tag{4}$$

where:

- $\mathbf{D}$ and $\mathbf{R}$ are respectively the $L\times N$ decomposition and $N\times L$ reconstruction matrix associated with the considered linear transformation such that $\mathbf{RD} = \mathrm{Identity}$. For redundant transforms (e.g. the undecimated wavelet transform), $L > N$.
- $\boldsymbol{\Theta}(\mathbf{w}) = \left[\boldsymbol{\theta}_l(\mathbf{w}_l)^{\mathrm{T}}\right]_{1\le l\le L}$ with $\boldsymbol{\theta}_l(\mathbf{w}_l) = [\theta_{l,c}(\mathbf{w}_l)]_{1\le c\le C}$ is a (spatially) pointwise multichannel thresholding function that will be defined in Section 3.2.

Our denoising process $\mathbf{F}(\mathbf{y})$ follows an extended version of the SURE-LET approach proposed in [5], i.e.

1. $\mathbf{F}$ is built as a *linear* expansion of possibly nonlinear functions $\mathbf{F}_k$:

$$\begin{aligned}\mathbf{F}(\mathbf{y}) &= \sum_{k=1}^{K}\mathbf{a}_k^{\mathrm{T}}\underbrace{\left(\mathbf{R}\boldsymbol{\Theta}_k(\mathbf{D}\mathbf{y}^{\mathrm{T}})\right)^{\mathrm{T}}}_{\mathbf{F}_k(\mathbf{y})}\\ &= \underbrace{[\mathbf{a}_1^{\mathrm{T}},\mathbf{a}_2^{\mathrm{T}},\ldots,\mathbf{a}_K^{\mathrm{T}}]}_{\mathbf{A}^{\mathrm{T}}}\begin{bmatrix}\mathbf{F}_1(\mathbf{y})\\ \mathbf{F}_2(\mathbf{y})\\ \vdots\\ \mathbf{F}_K(\mathbf{y})\end{bmatrix}\end{aligned} \tag{5}$$

   where $\mathbf{A}$ is the $KC\times C$ matrix of unknown parameters.

2. In general (i.e. non-orthogonal transforms), the search for the optimal, in the minimum MSE sense, parameters $\mathbf{A}$ has to be performed *globally* in the *image-domain* [5]. In practice, we cannot compute the actual MSE defined in (3), since the original image is obviously not available. Instead, we determine these unknown parameters $\mathbf{A}$ by minimizing *Stein's Unbiased Risk Estimate* [6], which is an *unbiased* estimate of the actual MSE, based on the noisy image $\mathbf{y}$ only. For multichannel signals, this MSE estimate is given by:

Theorem 1 *The following random variable*

$$\begin{aligned}\epsilon &= \frac{1}{NC}\,\mathrm{Tr}\left\{(\mathbf{F}(\mathbf{y})-\mathbf{y})(\mathbf{F}(\mathbf{y})-\mathbf{y})^{\mathrm{T}}\right\} +\\ &\quad\frac{2}{NC}\,\mathrm{div}\left\{\boldsymbol{\Gamma}\mathbf{F}(\mathbf{y})\right\} - \frac{1}{C}\,\mathrm{Tr}\left\{\boldsymbol{\Gamma}\right\}\end{aligned} \tag{6}$$

is an unbiased estimate of the MSE, i.e.

$$\mathscr{E}\left\{\epsilon\right\} = \frac{1}{NC}\mathscr{E}\left\{\mathrm{Tr}\left\{(\hat{\mathbf{x}}-\mathbf{x})(\hat{\mathbf{x}}-\mathbf{x})^{\mathrm{T}}\right\}\right\},$$

where

$$\mathrm{div}\left\{\mathbf{F}(\mathbf{y})\right\} = \sum_{c=1}^{C}\sum_{n=1}^{N}\frac{\partial \mathrm{F}_{n,c}(\mathbf{y})}{\partial y_{n,c}}$$

is a generalized divergence operator.

In the LET framework (5), and after a few formal algebraic manipulations, the expression (6) can be rewritten as follows:

$$\begin{aligned}\epsilon &= \frac{1}{NC}\,\mathrm{Tr}\left\{\mathbf{A}^{\mathrm{T}}\mathbf{M}\mathbf{A} - 2\mathbf{P}^{\mathrm{T}}\mathbf{A}\right\} +\\ &\quad\frac{1}{NC}\,\mathrm{Tr}\left\{\mathbf{y}\mathbf{y}^{\mathrm{T}}\right\} - \frac{1}{C}\,\mathrm{Tr}\left\{\boldsymbol{\Gamma}\right\}\end{aligned} \tag{7}$$

where we have defined:

$$\mathbf{M} = [\mathbf{F}_k(\mathbf{y})\mathbf{F}_l^{\mathrm{T}}(\mathbf{y})]_{1\le k,l\le K} \text{ and } \mathbf{P} = [\mathbf{p}_k]_{1\le k\le K}$$

where

$$\begin{aligned}\mathbf{p}_k &= \mathbf{F}_k(\mathbf{y})\mathbf{y}^{\mathrm{T}} -\\ &\quad\boldsymbol{\Gamma}\begin{bmatrix}\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{\Theta}'_{k;1,1}(\mathbf{w}) & \cdots & \boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{\Theta}'_{k;1,C}(\mathbf{w})\\ \vdots & \ddots & \vdots\\ \boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{\Theta}'_{k;C,1}(\mathbf{w}) & \cdots & \boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{\Theta}'_{k;C,C}(\mathbf{w})\end{bmatrix}\end{aligned}$$

with $\boldsymbol{\alpha} = \mathrm{diag}\{\mathbf{DR}\}$ and $\boldsymbol{\Theta}'_{k;i,j}(\mathbf{w}) = \left[\frac{\partial\theta_{k;l,i}(\mathbf{w}_l)}{\partial w_{l,j}}\right]_{1\le l\le L}$.

The computation of the diagonal of $\mathbf{DR}$ can be performed offline using the numerical algorithm described in [5]. Therefore, we do not need the explicit form of these decomposition and reconstruction matrices.

The minimization of (7) with respect to the unknown parameters $\mathbf{A}$ boils down to the following linear system of equations:

$$\mathbf{A} = \mathbf{M}^{-1}\mathbf{P} \tag{8}$$

which makes our algorithm fast and simple to implement.

3.2. Multichannel Thresholding

Now, we propose to extend the monochannel denoiser, described in [5], by taking into account the potentially strong interchannel similarities as follows:

$$\boldsymbol{\theta}_l(\mathbf{w}_l) = \mathbf{a}_1^{\mathrm{T}}\mathbf{w}_l + \gamma\left(\mathbf{w}_l^{\mathrm{T}}\boldsymbol{\Gamma}^{-1}\mathbf{w}_l\right)\mathbf{a}_2^{\mathrm{T}}\mathbf{w}_l \tag{9}$$

where we have experimentally chosen $\gamma(x) = \mathrm{e}^{-\frac{1}{C}\left(\frac{x}{9}\right)^4}$ as discriminator between large/small coefficients. If $C = 1$, we obviously obtain the grayscale thresholding function of [5].

The potential PSNR disparities between the channels are explicitly taken into account inside the argument of the zone selector γ, as well as implicitly through the $C \times C$ parameter matrices $\mathbf{a}_1$ and $\mathbf{a}_2$.

4. EXPERIMENTS

We have applied our multichannel SURE-LET thresholding for denoising color images using J iterations[1] of the undecimated wavelet transform (UWT)[2] with symmetric boundaries handling. A new thresholding function (9) was applied inside each of the $3J$ wavelet highpass subbands, while keeping the lowpass residual unchanged. All the parameters[3] have been optimized in the *image-domain* due to the non-orthogonality of the UWT. The experiments have been executed on the set of standard RGB images shown in Figure 1.

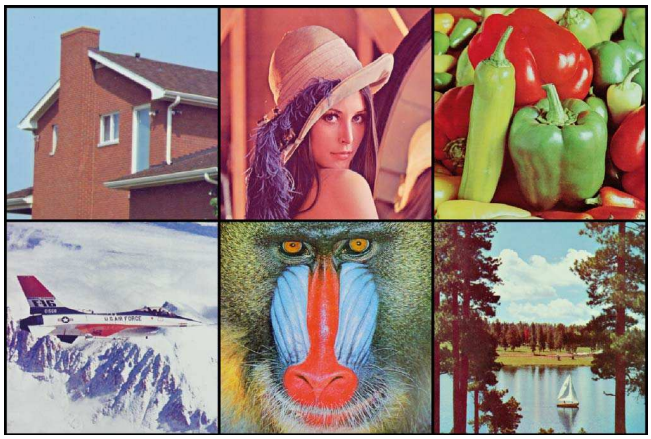

Fig. 1. Test images used in the experiments, referred to as *Image 1* to *Image 6* (numbered from left to right, top-down).

We have corrupted these test images with an additive colorwise Gaussian white noise of known intercolor covariance matrix:

$$\mathbf{\Gamma} = \begin{bmatrix} \sigma_R^2 & 0 & 0 \\ 0 & \sigma_G^2 & 0 \\ 0 & 0 & \sigma_B^2 \end{bmatrix}$$

In Table 1 we compare our PSNR results with those obtained by running two other state-of-the-art color denoising algorithms:

- Pižurica *et al.* ProbShrink-MB[4] [4], which is a multiband extension of their original wavelet-based grayscale denoiser. We have applied it with the same transform (UWT Haar) as ours.
- Foi *et al.* pointwise SA-DCT[5] [7], which is the application of their grayscale shape-adaptive DCT denoiser in the opponent color space, but using the adaptive neighborhoods defined in the luminance channel for all channels.

We notice that we obtain a significant gain (about $+1$ dB) over the ProbShrink-MB, and similar results to the pointwise SA-DCT. Moreover, our denoised images contain very few color artifacts, and represent a good trade-off between noise removal and preservation of small details (see Figure 2).

From a computational point of view, the execution of our current unoptimized Matlab implementation lasts around 12s for 256×256 RGB images, which is slightly faster than the two other algorithms that make use of pre-compiled codes.

An interesting fallout of our linear parametrization is that the resulting thresholding function (9) is very robust to any linear color space transformation (e.g. similar results are obtained in the RGB and in the more "decorrelated" luminance-chrominance representation), which is not the case with other algorithms.

Table 1. Comparison of some state-of-the-art color image denoising methods.

$\sigma_R = \sigma_G = \sigma_B$	5	10	20	30	50	100
Input PSNR	34.15	28.13	22.11	18.59	14.15	8.13
Method	**Image 1** 256×256					
ProbShrink-MB[4]	37.99	34.68	31.84	30.19	27.94	24.73
SA-DCT[7]	**38.61**	**35.64**	**32.96**	**31.38**	**29.15**	**25.90**
Our method	38.49	35.38	32.73	31.18	29.09	25.81
Method	**Image 2** 512×512					
ProbShrink-MB[4]	37.43	34.35	31.61	30.03	28.04	25.45
SA-DCT[7]	37.69	34.97	**32.62**	**31.21**	**29.25**	26.47
Our method	**37.98**	**35.02**	32.51	31.05	29.16	**26.56**
Method	**Image 3** 512×512					
ProbShrink-MB[4]	36.47	33.46	31.10	29.73	27.82	25.04
SA-DCT[7]	36.79	33.68	**31.55**	**30.31**	**28.51**	**25.80**
Our method	**36.82**	**33.76**	31.45	30.16	28.43	**25.80**
Method	**Image 4** 512×512					
ProbShrink-MB[4]	38.82	35.31	32.11	30.31	27.86	24.59
SA-DCT[7]	39.45	**36.36**	**33.38**	**31.59**	29.22	26.05
Our method	**39.49**	36.35	33.31	31.52	**29.23**	**26.10**
Method	**Image 5** 512×512					
ProbShrink-MB[4]	34.85	30.06	26.12	24.18	22.01	19.95
SA-DCT[7]	**35.22**	30.61	26.88	**25.06**	**23.02**	20.52
Our method	35.18	**30.65**	**26.93**	25.03	22.91	**20.69**
Method	**Image 6** 512×512					
ProbShrink-MB[4]	35.84	31.98	28.97	27.41	25.48	22.88
SA-DCT[7]	**36.34**	32.33	29.39	27.92	26.10	23.62
Our method	36.21	**32.44**	**29.57**	**28.07**	**26.24**	**23.76**

Note: output PSNRs have been averaged over eight noise realizations.

[1] $J = 4$ (resp. $J = 5$) for 256×256 (resp. 512×512) color images.

[2] Similarly to what was observed in [5], Haar filters turn out to give the best results.

[3] Globally $K = 3J \times 3^2 \times 2$ linear parameters

[4] Matlab code available online at `http://telin.rug.ac.be/~sanja/`

[5] Matlab code available online at `http://www.cs.tut.fi/~foi/SA-DCT/#ref_software`

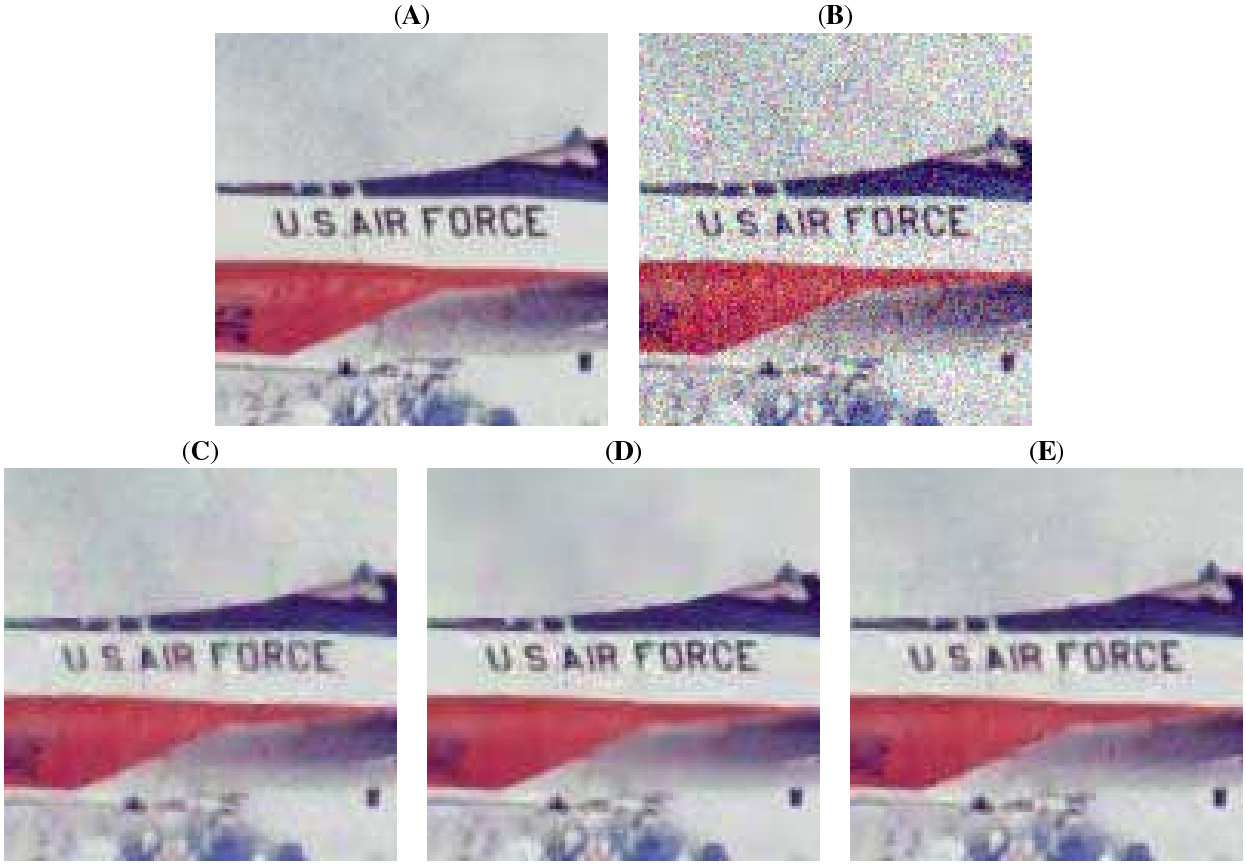

Fig. 2. (**A**) Part of the noise-free *Image 4*. (**B**) Part of the noisy *Image 4*: PSNR = 18.59 dB. (**C**) Result of the *ProbShrink-MB*: PSNR = 30.31 dB. (**D**) Result of the pointwise SA-DCT: PSNR = 31.59 dB. (**E**) Result of our method: PSNR = 31.52 dB.

5. CONCLUSION

We presented an extension of the SURE-LET denoising approach introduced in [5] to properly handle multichannel images. For color image denoising, the resulting pointwise multichannel denoiser applied within the undecimated wavelet transform outperforms other state-of-the-art wavelet-based algorithms. Despite its lower complexity, our strategy is even competitive with respect to a recent shape-adaptive method.

6. REFERENCES

[1] Stéphane Mallat, "A Theory for Multiresolution Signal Decomposition: The Wavelet Representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 7, pp. 674–693, July 1989.

[2] David L. Donoho and Iain M. Johnstone, "Adapting to Unknown Smoothness via Wavelet Shrinkage," *Journal of the American Statistical Association*, vol. 90, no. 432, pp. 1200–1224, December 1995.

[3] A. Benazza-Benyahia and J.-C. Pesquet, "Building Robust Wavelet Estimators for Multicomponent Images Using Stein's Principle," *IEEE Transactions on Image Processing*, vol. 14, no. 11, pp. 1814–1830, November 2005.

[4] Aleksandra Pižurica and Wilfried Philips, "Estimating the Probability of the Presence of a Signal of Interest in Multiresolution Single- and Multiband Image Denoising," *IEEE Transactions on Image Processing*, vol. 15, no. 3, pp. 654–665, March 2006.

[5] Thierry Blu and Florian Luisier, "The SURE-LET Approach to Image Denoising," *IEEE Transactions on Image Processing*, vol. 16, no. 11, pp. 2778–2786, November 2007.

[6] C. Stein, "Estimation of the Mean of a Multivariate Normal Distribution," *The Annals of Statistics*, vol. 9, pp. 1135–1151, 1981.

[7] Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian, "Pointwise Shape-Adaptive DCT for High-Quality Denoising and Deblocking of Grayscale and Color Images," *IEEE Transactions on Image Processing*, vol. 16, no. 5, pp. 1395–1411, May 2007.