

SURE-LET Interscale-Intercolor Wavelet Thresholding for Color Image Denoising

Florian Luisier^a and Thierry Blu^a

^aBiomedical Imaging Group (BIG), Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland

ABSTRACT

We propose a new orthonormal wavelet thresholding algorithm for denoising color images that are assumed to be corrupted by additive Gaussian white noise of known intercolor covariance matrix. The proposed wavelet denoiser consists of a linear expansion of thresholding (LET) functions, integrating both the interscale and intercolor dependencies. The linear parameters of the combination are then solved for by minimizing Stein’s unbiased risk estimate (SURE), which is nothing but a robust unbiased estimate of the mean squared error (MSE) between the (unknown) noise-free data and the denoised one. Thanks to the quadratic form of this MSE estimate, the parameters optimization simply amounts to solve a linear system of equations.

The experimentations we made over a wide range of noise levels and for a representative set of standard color images have shown that our algorithm yields even slightly better peak signal-to-noise ratios than most state-of-the-art wavelet thresholding procedures, even when the latter are executed in an undecimated wavelet representation.

Keywords: Color image denoising, orthonormal wavelet transform, SURE minimization, interscale consistencies, intercolor similarities

1. INTRODUCTION

Many efficient denoisers make use of multiresolution tools such as the wavelet transform.^{1,2} Its appealing property is its ability to concentrate most of the signal information into few large coefficients, while the noise—assumed as additive, Gaussian and white—is uniformly spread throughout all coefficients. In the transformed domain, a simple processing such as *thresholding* becomes then very efficient.³⁻⁷

The most straightforward strategy to remove additive Gaussian white noise from color images is simply to apply an existing grayscale denoiser separately in each color channel of the standard Red-Green-Blue (RGB) representation. However, this approach is far from being optimal, due to the potentially strong common information shared by the color channels. To improve this basic solution, the images are usually denoised in a more “decorrelated” color space, such as the luminance-chrominance (YUV) representation.⁸ An alternative or complement to this approach consists in devising specific non-separable color image denoising algorithms.^{6,9-11}

In this paper, we propose to extend our previous grayscale denoiser⁷ to color image denoising (see¹² for a more complete theory on multichannel image denoising) in an *orthonormal* wavelet representation. Following the general approach introduced in,^{7,13} we build our wavelet estimator as a *linear* expansion of thresholding (LET) functions taking into account both the interscale consistencies and the intercolor similarities. The unknown weights of the linear combination are then solved for by minimizing Stein’s unbiased risk estimate (SURE),¹⁴ which is a robust estimate of the mean squared error (MSE) between the —unknown— noise-free image and the denoised one, that does not require any prior on the original image. In the SURE-LET framework, the noise-free wavelet coefficients are thus not even considered as random; the only randomness is induced by the noise. Our linear parametrization conjugated with the quadratic form of SURE reduces the parameters optimization to a simple resolution of a linear system of equations.

Further author information:

florian.luisier@epfl.ch; thierry.blu@epfl.ch

2. SURE-LET FOR COLOR IMAGES

In this paper, we consider N -pixel images with 3 color channels. We denote these color images by a $3 \times N$ matrix whose columns are the color component values of each pixel:

$$\mathbf{x} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N] \quad \text{where } \mathbf{x}_n = \begin{bmatrix} x_{n,1} \\ x_{n,2} \\ x_{n,3} \end{bmatrix}.$$

We assume that these images are corrupted by an additive color component-wise Gaussian white noise $\mathbf{b} = [\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_N]$ of known intercolor covariance matrix \mathbf{R} , i.e.

$$\mathcal{E} \{ \mathbf{b}_n \mathbf{b}_{n'}^T \} = \mathbf{R} \delta_{n-n'}$$

We denote the resulting noisy color image by $\mathbf{y} = [\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_N]$ and we thus have:

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

Note that, in our work, the original image \mathbf{x} is *not* considered as a realization of some random process; the only randomness comes from the noise \mathbf{b} , and therefore, the noisy image \mathbf{y} is also random.

The quality of the denoised image $\hat{\mathbf{x}}$ will be evaluated by the widely used *Peak Signal-to-Noise Ratio* (PSNR)

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

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

$$\begin{aligned} \text{MSE} &= \frac{1}{3N} \text{Tr} \{ (\hat{\mathbf{x}} - \mathbf{x})(\hat{\mathbf{x}} - \mathbf{x})^T \} \\ &= \frac{1}{3N} \sum_{n=1}^N \|\hat{\mathbf{x}}_n - \mathbf{x}_n\|^2 \end{aligned}$$

The denoising process will contain a *critically sampled Orthonormal Wavelet Transform* (OWT) applied to each color channel (see Figure 1). Denoting the resulting wavelet subbands at scale $j \in [1, J]$ and orientation $o \in [1, 3]$ by superscripts, we thus have:

$$\mathbf{y}^{j,o} = \mathbf{x}^{j,o} + \mathbf{b}^{j,o} \tag{3}$$

The linearity and orthonormality of the OWT has two important consequences:

1. The noise statistics are preserved in the transform domain, i.e. its wavelet coefficients are Gaussian and independent within and between the subbands, with the same intercolor covariance matrix \mathbf{R} .
2. The MSE in the image domain is simply a weighted sum of the subband MSE's.

Thanks to these two properties, an independent processing $\theta^{j,o}$ of each individual noisy wavelet subbands $\mathbf{y}^{j,o}$ becomes particularly suitable. The thresholding function $\theta^{j,o}$ will also involve an interscale predictor $\mathbf{p}^{j,o}$ obtained by appropriately filtering the lowpass subband at the same scale j (see Figure 2), as detailed in.⁷ Notice that, due to the statistical independence between subbands of different scale or orientation, $\theta^{j,o}$ and $\mathbf{p}^{j,o}$ are statistically independent.

From now on, we will drop the subband superscripts j and o , and consider the denoising of a color wavelet subband $\mathbf{y} = \mathbf{x} + \mathbf{b}$, given its independent interscale predictor \mathbf{p} , by using a $\mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{R}^3$ function θ relating the coefficients of \mathbf{y} and \mathbf{p} to the coefficients of the estimate $\hat{\mathbf{x}}$ through

$$\hat{\mathbf{x}}_n = \theta(\mathbf{y}_n, \mathbf{p}_n), \quad \text{for } n = 1, 2, \dots, N. \tag{4}$$

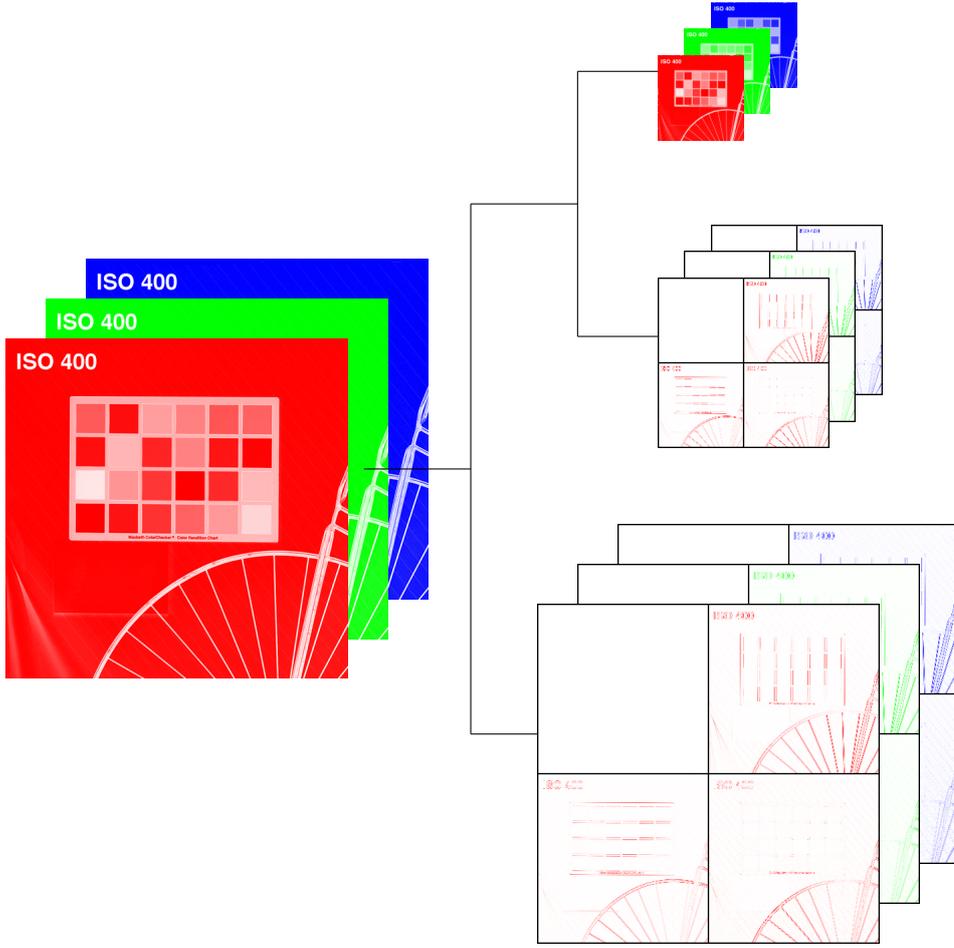


Figure 1. Two iterations of a 2-D orthonormal wavelet transform applied to a RGB image.

Our aim will be then to find the θ that will minimize the MSE defined in (3). Since we do not have access to the noise-free data \mathbf{x} , we will instead rely on an adapted version of Stein's unbiased risk estimate (SURE)¹⁴ to accurately estimate the actual MSE, as shown below:

THEOREM 2.1. *Assume that $\theta(\cdot, \cdot)$ is (weakly) differentiable with respect to its first variable. Then, if the estimate $\hat{\mathbf{x}}$ is built according to (4), the following random variable*

$$\epsilon = \frac{1}{3N} \sum_{n=1}^N \|\theta(\mathbf{y}_n, \mathbf{p}_n) - \mathbf{y}_n\|^2 + \frac{2}{3N} \sum_{n=1}^N \text{Tr} \{ \mathbf{R}^T \nabla_1 \theta(\mathbf{y}_n, \mathbf{p}_n) \} - \frac{1}{3} \text{Tr} \{ \mathbf{R} \} \quad (5)$$

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

$$\mathcal{E} \{ \epsilon \} = \mathcal{E} \left\{ \frac{1}{3N} \sum_{n=1}^N \|\hat{\mathbf{x}}_n - \mathbf{x}_n\|^2 \right\}.$$

Here, we have denoted by $\nabla_1 \theta$ the matrix containing the partial derivatives of the color components of

In this framework, the unbiased MSE estimate defined in (5) is *quadratic* in the unknown parameters \mathbf{A} , and its minimization boils down to the following linear system of equations:

$$\mathbf{A} = \mathbf{M}^{-1}\mathbf{C} \quad (7)$$

where

$$\begin{aligned} \mathbf{M} &= \sum_{n=1}^N \Theta(\mathbf{y}_n, \mathbf{p}_n) \Theta(\mathbf{y}_n, \mathbf{p}_n)^T \\ \mathbf{C} &= \sum_{n=1}^N \left(\Theta(\mathbf{y}_n, \mathbf{p}_n) \mathbf{y}_n^T - (\nabla_1 \Theta(\mathbf{y}_n, \mathbf{p}_n))^T \mathbf{R} \right) \end{aligned}$$

Notice that if \mathbf{M} is not a full rank matrix, we can simply take its pseudo-inverse to choose among the admissible solutions.

3. A NEW INTERSCALE-INTERCOLOR THRESHOLDING FUNCTION

We now propose a natural extension of the denoising function presented in⁷ by taking into account the potentially strong similarities between the various color channels. We will use $K = 4$ thresholding θ_k 's, one for each particular class of coefficients. These classes are determined by a smooth decision function $\gamma(x) = e^{-\frac{|x|}{3-|x|}}$ applied to a weighted combination of the predictors/wavelet coefficients, leading to the following interscale-intercolor thresholding function:

$$\begin{aligned} \theta(\mathbf{y}_n, \mathbf{p}_n) &= \underbrace{\gamma(\mathbf{p}_n^T \mathbf{R}^{-1} \mathbf{p}_n)}_{\text{small predictors/small coefficients}} \underbrace{\gamma(\mathbf{y}_n^T \mathbf{R}^{-1} \mathbf{y}_n)}_{\text{small coefficients}} \mathbf{a}_1^T \mathbf{y}_n \\ &+ \underbrace{(1 - \gamma(\mathbf{p}_n^T \mathbf{R}^{-1} \mathbf{p}_n)) \gamma(\mathbf{y}_n^T \mathbf{R}^{-1} \mathbf{y}_n)}_{\text{large predictors/small coefficients}} \mathbf{a}_2^T \mathbf{y}_n \\ &+ \underbrace{\gamma(\mathbf{p}_n^T \mathbf{R}^{-1} \mathbf{p}_n) (1 - \gamma(\mathbf{y}_n^T \mathbf{R}^{-1} \mathbf{y}_n))}_{\text{small predictors/large coefficients}} \mathbf{a}_3^T \mathbf{y}_n \\ &+ \underbrace{(1 - \gamma(\mathbf{p}_n^T \mathbf{R}^{-1} \mathbf{p}_n)) (1 - \gamma(\mathbf{y}_n^T \mathbf{R}^{-1} \mathbf{y}_n))}_{\text{large predictors/large coefficients}} \mathbf{a}_4^T \mathbf{y}_n \end{aligned} \quad (8)$$

4. EXPERIMENTS

4.1 Context

All the experiments of this section have been carried out on $N = 512 \times 512$ RGB images from the set presented in Figure 3. We have applied our interscale-interchannel thresholding algorithm after 5 decomposition levels of an *orthonormal* wavelet transform (OWT) using the standard Daubechies symlets² with eight vanishing moments (*sym8* in Matlab). We have assumed that in the red-green-blue (RGB) color representation, the intercolor noise covariance matrix is given by:

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

This assumption implies that, in other color spaces, there will usually be noise correlations between the color channels. However, contrary to the other algorithms that have been previously published (see Section 4.2), ours is quite insensitive to the color representation (variations of ± 0.1 dB), thanks to our linear parametrization which acts as an "optimal"—in the minimum SURE sense—color space projector. From now on, we will thus apply our intercolor algorithm in the RGB representation only.



Figure 3. Test images used in the experiments, referred to as *Image 1* to *Image 8* (numbered from left-to-right and top-to-bottom).

4.2 Comparisons

We have chosen to compare our method with two state-of-the-art multiresolution-based denoising algorithms:

- Pižurica’s *et al.* *ProbShrink-YUV*,⁸ which is an application of their original grayscale denoiser in the luminance-chrominance color space in the undecimated wavelet transform.
- Pižurica’s *et al.* *ProbShrink-MB*,⁶ which is a multiband extension of their original grayscale denoiser. For color image denoising, it has to be applied in the standard RGB representation, and for equal noise variance in each channels.

We have applied these two algorithms with the—highly redundant—*undecimated* wavelet transform (UWT) using the code of the authors* with their suggested parameters; we have considered the same number of decomposition levels and the same wavelet (*sym8*) as with our method. Since these algorithms have been shown in^{6,8} to favorably compare with the multiband wavelet thresholding described in,⁹ as well as with the vector-based linear minimum mean squared error estimator proposed in,¹⁰ they constitute a good reference for evaluating our solution.

In a first experiment, we have corrupted the test images with additive Gaussian white noise having the same variance in each RGB channel. The PSNR results are displayed in Table 1. Despite being performed in a *non-redundant* wavelet representation, our solution gives even slightly better (+0.3 dB on average) output PSNRs than both variants of the *ProbShrink* applied in the *undecimated* wavelet representation. From a visual point of view, our algorithm holds its own against these state-of-the-art multiresolution-based redundant approaches (see Figure 4).

In a second experiment, the test images have been corrupted with additive Gaussian white noise having a different power in each RGB channel ($\sigma_R = 38.25$, $\sigma_G = 25.50$, $\sigma_B = 12.75$). The PSNR results are displayed in Figure 5 and a visual example is shown in Figure 6. As can be observed, our *SURE-LET* approach outperforms the *ProbShrink-YUV* in terms of PSNR (around +1 dB).

*Available at: <http://telin.rug.ac.be/~sanja/>

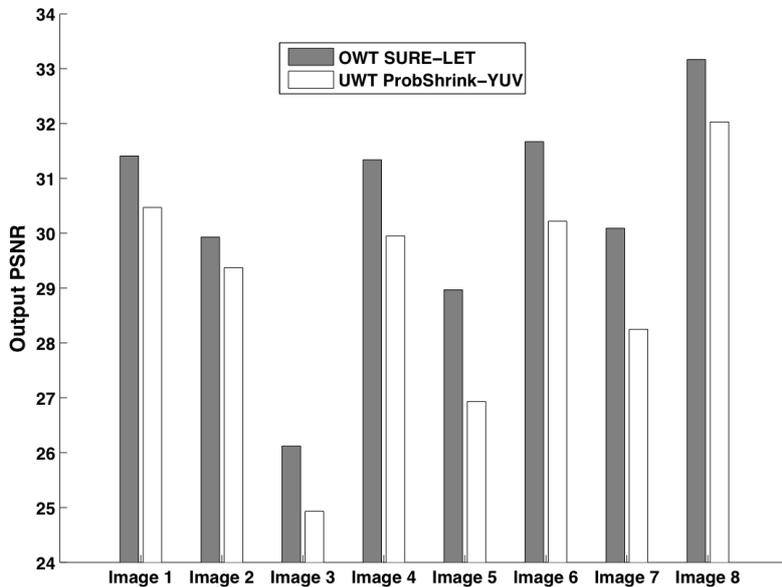


Figure 4. (A) Part of the noise-free *Image 5*. (B) Part of the noisy *Image 5*: PSNR = 18.59 dB. (C) Result of our *SURE-LET* (OWT *sym8*): PSNR = 26.79 dB. (D) Result of the *ProbShrink-MB* (UWT *sym8*): PSNR = 25.44 dB. (E) Result of the *ProbShrink-YUV* (UWT *sym8*): PSNR = 26.59 dB.

Table 1. Comparison of Color Denoising Algorithms (Same Noise Level in Each RGB Channel)

$\sigma_R = \sigma_G = \sigma_B$	5	10	20	30	50	100	5	10	20	30	50	100
Input PSNR [dB]	34.15	28.13	22.11	18.59	14.15	8.13	34.15	28.13	22.11	18.59	14.15	8.13
Method	<i>Image 1</i> 512×512						<i>Image 2</i> 512×512					
<i>OWT SURE-LET</i>	37.80	34.64	31.90	30.33	28.35	25.66	36.62	33.35	30.72	29.19	27.16	24.48
<i>UWT ProbShrink-YUV</i> ^s	36.52	34.22	31.63	30.08	28.03	25.23	35.33	33.02	30.59	28.98	26.83	24.11
<i>UWT ProbShrink-MB</i> ^b	37.46	34.42	31.47	29.83	27.76	25.03	36.33	33.35	30.74	29.20	26.85	24.28
Method	<i>Image 3</i> 512×512						<i>Image 4</i> 512×512					
<i>OWT SURE-LET</i>	35.12	30.49	26.64	24.71	22.59	20.37	39.11	35.70	32.29	30.29	27.77	24.77
<i>UWT ProbShrink-YUV</i> ^s	33.98	29.59	26.12	24.27	22.27	20.07	38.37	35.41	31.53	29.41	27.09	24.15
<i>UWT ProbShrink-MB</i> ^b	34.83	30.15	26.17	24.16	21.98	19.81	38.78	35.23	31.80	29.81	26.87	23.97
Method	<i>Image 5</i> 512×512						<i>Image 6</i> 512×512					
<i>OWT SURE-LET</i>	38.35	33.67	29.21	26.80	24.03	20.88	39.87	35.88	32.01	29.88	27.45	24.64
<i>UWT ProbShrink-YUV</i> ^s	37.47	33.24	28.96	26.56	23.86	20.68	39.21	35.78	31.92	29.77	27.27	24.25
<i>UWT ProbShrink-MB</i> ^b	36.71	31.99	27.69	25.44	22.84	19.93	38.69	34.66	30.87	28.82	26.52	23.78
Method	<i>Image 7</i> 512×512						<i>Image 8</i> 512×512					
<i>OWT SURE-LET</i>	38.69	34.24	30.29	28.15	25.63	22.72	41.05	37.56	34.00	31.88	29.26	26.11
<i>UWT ProbShrink-YUV</i> ^s	37.87	33.71	29.91	27.90	25.47	22.58	40.27	37.13	33.56	31.58	29.08	25.77
<i>UWT ProbShrink-MB</i> ^b	37.05	32.64	28.88	26.87	24.58	21.91	40.49	36.92	33.36	31.24	28.61	25.55

- Notes: 1. *UWT* stands for the redundant Undecimated Wavelet Transform, whereas *OWT* stands for the non-redundant Orthonormal Wavelet Transform.
2. Output PSNRs have been averaged over ten noise realizations.

Figure 5. PSNRs comparison between our algorithm (*OWT SURE-LET*) and a state-of-the-art multiresolution-based method (*UWT ProbShrink-YUV*) for a different noise power inside each RGB color channel (input PSNR: 19.33 dB).

Finally, we must emphasize that the execution of the un-optimized Matlab implementation of our algorithm only lasts around 6s for 512×512 color images on a Power Mac G5 workstation with 1.8GHz CPU[†]. To compare with, the *ProbShrink* requires approximately 19s under the same conditions.

5. CONCLUSION

The natural extension of our previous grayscale denoiser has been shown to be particularly efficient for denoising color images. Thanks to an appropriate integration of both the interscale and intercolor similarities inside the SURE-LET framework, our orthonormal wavelet estimator is even competitive —both qualitatively and computationally— with the best state-of-the-art multiresolution algorithms which use highly redundant transforms.

[†]The interested reader can check our online demo.¹⁵

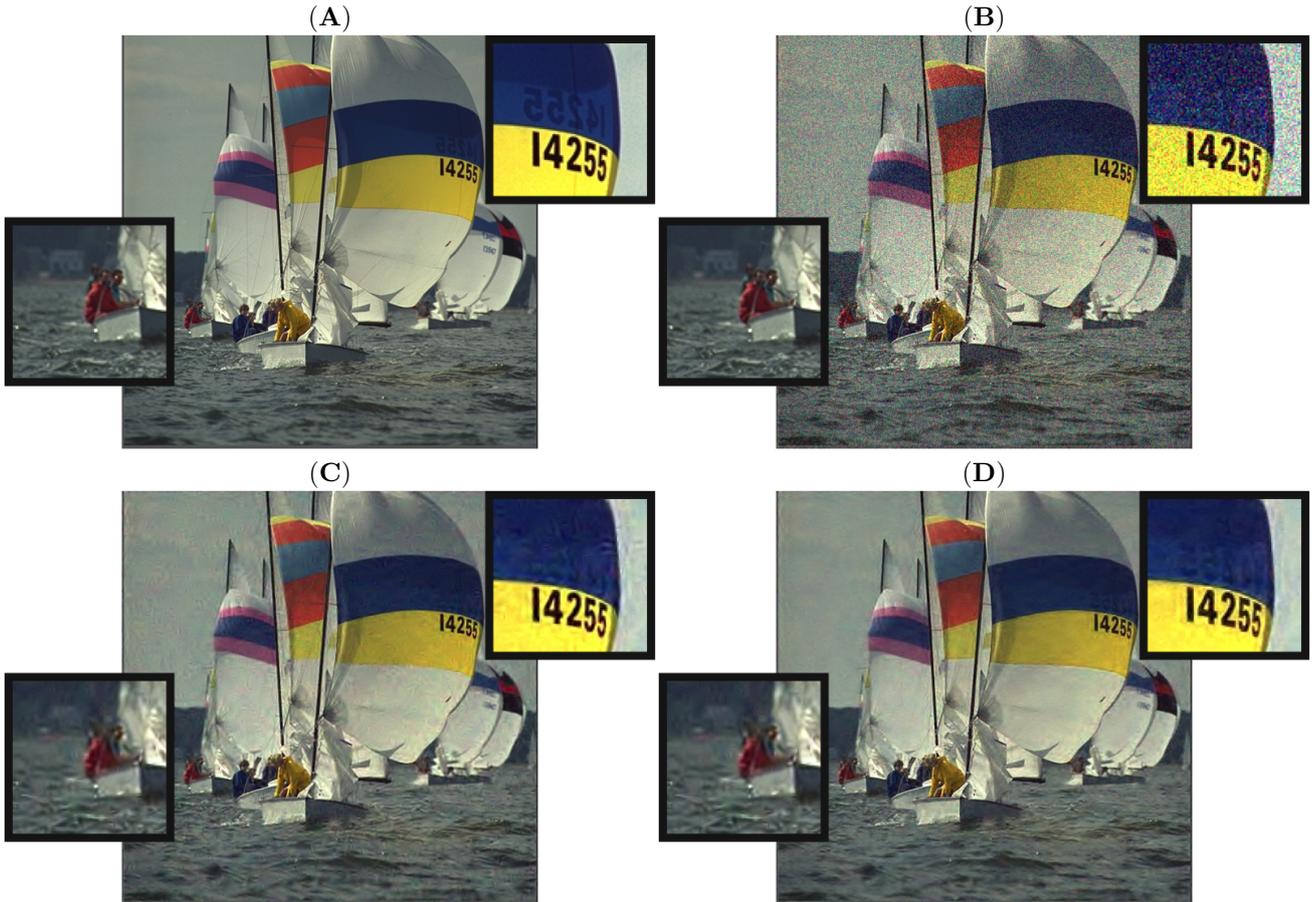


Figure 6. (A) Noise-free *Image 6*. (B) Noisy *Image 6*: PSNR = 19.33 dB. (C) Result of the *ProbShrink-YUV* (UWT *sym8*): PSNR = 30.25 dB. (D) Result of our *SURE-LET* (OWT *sym8*): PSNR = 31.67 dB.

REFERENCES

1. S. Mallat, "A Theory for Multiresolution Signal Decomposition: The Wavelet Representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **11**, pp. 674–693, July 1989.
2. I. Daubechies, *Ten lectures on wavelets*, vol. 61 of *CBMS-NSF Regional Conference Series in Applied Mathematics*, SIAM, Philadelphia, PA, 1992.
3. D. L. Donoho and I. M. Johnstone, "Adapting to Unknown Smoothness via Wavelet Shrinkage," *Journal of the American Statistical Association* **90**, pp. 1200–1224, December 1995.
4. F. Abramovitch, T. Sapatinas, and B. W. Silverman, "Wavelet Thresholding via a Bayesian Approach," *Journal of the Royal Statistical Society. Series B* **60**(4), pp. 725–749, 1998.
5. J. Portilla, V. Strela, M. J. Wainwright, and E. P. Simoncelli, "Image Denoising using Scale Mixtures of Gaussians in the Wavelet Domain," *IEEE Transactions on Image Processing* **12**, pp. 1338–1351, November 2003.
6. A. Pizurica and W. Philips, "Estimating the Probability of the Presence of a Signal of Interest in Multiresolution Single- and Multiband Image Denoising," *IEEE Transactions on Image Processing* **15**, pp. 654–665, March 2006.
7. F. Luisier, T. Blu, and M. Unser, "A New SURE Approach to Image Denoising: Interscale Orthonormal Wavelet Thresholding," *IEEE Transactions on Image Processing* **16**, pp. 593–606, March 2007.
8. A. Pizurica, W. Philips, and P. Scheunders, "Wavelet domain denoising of single-band and multiband images adapted to the probability of the presence of features of interest," *Proc. SPIE* **5914**, September 2005.

9. P. Scheunders, "Wavelet Thresholding of Multivalued Images," *IEEE Transactions on Image Processing* **13**, pp. 475–483, April 2004.
10. P. Scheunders and J. Driesen, "Least-Squares Interband Denoising of Color and Multispectral Images," *IEEE International Conference on Image Processing* **2**, pp. 985–988, October 2004.
11. A. Benazza-Benyahia and J.-C. Pesquet, "Building Robust Wavelet Estimators for Multicomponent Images Using Stein's Principle," *IEEE Transactions on Image Processing* **14**, pp. 1814–1830, November 2005.
12. F. Luisier and T. Blu, "SURE-LET Multichannel Image Denoising: Interscale Orthonormal Wavelet Thresholding," *Submitted to IEEE Transactions on Image Processing*.
13. T. Blu and F. Luisier, "The SURE-LET Approach to Image Denoising," *IEEE Transactions on Image Processing (in press)*.
14. C. Stein, "Estimation of the Mean of a Multivariate Normal Distribution," *The Annals of Statistics* **9**, pp. 1135–1151, 1981.
15. SURE-LET Color Denoising Demo: <http://bigwww.epfl.ch/demo/suredenoising-color/index.html>.