IMPROVED BM3D FOR CORRELATED NOISE REMOVAL

Marcella Matrecano, Giovanni Poggi and Luisa Verdoliva
Department of Biomedical, Electronic, and Telecommunication Engineering, University of Naples Federico II
Via Claudio 21, 80121 Naples, Italy

Keywords: Denoising, Correlated Noise, Nonlocal Filtering, BM3D.

Abstract: Most of the literature on denoising focuses on the additive-white-gaussian-noise (AWGN) model. However, in many important applicative fields, images are typically affected by non-Gaussian and/or colored noise, in which cases AWGN-based techniques fall much short of their promises. In this paper, we propose a new denoising technique for correlated noise based on the non-local approach. We start from the well-known BM3D algorithm, which can be considered to be the state of the art in AWGN denoising, and modify it in various critical steps in order to take into account the non-whiteness of noise. Experimental results on several test images corrupted by correlated noise confirm the potential of the proposed technique.

1 INTRODUCTION

The noise is an unpredictable perturbation which disturbs a signal causing random fluctuations of the observed variables. Generally speaking, it is an issue of considerable importance in any acquiring and processing data system, especially imaging techniques. The scientific literature offers a plethora of denoising functions often included in commercial software as tools to support and simplify the extraction of significant information from noisy images.

In fact, image denoising has been the object of intense research from the very beginning of the digital era, with tremendous performance improvements over the years, especially with the advent of wavelet-domain shrinkage techniques originally proposed by Donoho and Johnstone (1994). Nonetheless, this “mature” field has seen yet another big leap forward in recent years with the introduction of the non-local filtering concept, first formalized in the non-local means (NLM) algorithm proposed by Buades, Coll and Morel (2005).

The major breakthrough of NLM consists in selecting in a very sensible way the set of pixels used to estimate the true value of a target pixel, that is, not the pixels closest to the target, but those deemed to be the most similar to it. In practice, for each target pixel z^T of the noisy image $z(n)$, the surrounding patch P^T is extracted and compared with all patches P_i in a given neighborhood of the target. The patches that are more similar to the target patch, according to a suitable “distance” measure, are associated with the most relevant predictor pixels, that is

$$z^T = \sum w_i z_i \tag{1}$$

where the weight w_i is a decreasing function of the distance $d_i = d(P^T, P_i)$.

The nonlocal approach turns out to be especially effective in the presence of edges and textures, when patches are well characterized and bring valuable information on the pixel context, providing a significant performance improvement w.r.t. conventional techniques both in the spatial and transform domain. Following the success of NLM many variations have been proposed, among which the block-matching 3d (BM3D) algorithm (Dabov, Foi, Katkovnik and Egiazarian, 2007), which appears to be the current state of the art in image denoising and restoration in general.

Although a thorough description of BM3D is out of the scope of this work, we need to recall here its basic steps. The first action taken by BM3D, just like in NLM, is to locate similar patches by means of block-matching. Unlike in NLM, however, all such patches are then collected in a 3D structure which undergoes a decorrelating transform (typically wavelet) so as to exploit both spatial and contextual dependencies. Once a sparse representation is obtained, some form of shrinkage is used to remove

Matrecano M., Poggi G. and Verdoliva L..
IMPROVED BM3D FOR CORRELATED NOISE REMOVAL.
DOI: 10.5220/0003854101290134
In *Proceedings of the International Conference on Computer Vision Theory and Applications* (VISAPP-2012), pages 129-134
ISBN: 978-989-8565-03-7

noise components, before going back in the image domain. Since filtered patches can overlap, several estimates of the same pixel are typically obtained, which are weighted to compute a “basic estimate” $\hat{y}_{basic}$ of the denoised image. At this point, the noisy image *z(n)* undergoes the denoising process anew, with the difference that block-matching takes place on the basic estimate $\hat{y}_{basic}$ of the clean image so as to obtain more reliable matches, and wavelet shrinkage is replaced by empirical Wiener filtering, with statistics computed again on $\hat{y}_{basic}$.

Both NLM and BM3D have been proposed in the context of AWGN image denoising and, therefore, work poorly in all situations where the noise cannot be considered Gaussian nor white. Nonetheless, the nonlocal approach keeps making full sense, and hence there is much interest in adapting the basic algorithms to such new conditions. For example, with reference to synthetic aperture radar (SAR) images, where the speckle is clearly non-Gaussian (actually, not even additive) suitable ad hoc versions of NLM and BM3D have been proposed by Deledalle, Denis and Tupin (2009) and by Parrilli, Poderico, Angelino, and Verdoliva (2011) respectively, with very good results.

The problem of nonlocal image denoising in the presence of colored noise has been already addressed as well. A version of NLM for colored noise (NLM-C) is proposed by Goossens, Luong, Pizurica and Philips (2008) where the noise is assumed to come from the linear filtering of white Gaussian noise. Given the impulse-response of the filter, and hence all noise statistics, the Authors replace the Euclidean distance, used originally to compute the similarity among patches, with the Mahalanobis distance which takes the noise covariance matrix into account. Alternatively, to reduce the computational load, they apply a prewhitening linear filter to the noisy image and use the resulting image to compute the weights by the Euclidean distance. Numerical experiments show NLM-C to provide significant improvements, both visually and in terms of PSNR, not only over basic NLM (called NLM-W in this context) but also w.r.t. to some recent wavelet-based denoising techniques for colored noise: BLS-GSM (Portilla, Strela, Wainwright & Simoncelli, 2003) and MP-GSM (Goossens, Pizurica, 2009).

Also BM3D has been already adapted, by Dabov, Foi, Katkovnik and Egiazarian (2008), to the case of correlated noise. The Authors observe that the decorrelating transform, used before shrinkage, outputs coefficients which, due to noise nonwhiteness, have variances $\sigma^2_{2D}(i)$ that do depend on the coefficient index *i*. This fact is taken into account in various phases of the algorithm: by using a weighted block distance computed in the transform domain; by using a different shrinkage threshold for each coefficient; and by aggregating filtered blocks based on their expected noise level.

In this work, based on the above ideas, we propose a new version of BM3D for correlated noise. We use the basic strategy of the original BM3D algorithm because of its strong rationale, but modify it in several steps to keep into account the actual noise statistics. In particular, we improve the block matching by resorting to image prewhitening, and the shrinkage (hard thresholding in the first step, and Wiener filtering in the second step) by taking into account the different noise variances of coefficients and improving their estimate.

2 PROPOSED ALGORITHM

Since the proposed algorithm is a modification of BM3D, we describe and discuss here only the differences w.r.t. the original algorithm (Dabov et al., 2007). The first and probably most important improvement concerns the block matching, based on straight Euclidean distance in the original algorithm.

The ultimate goal of block matching is to find out the *signal* patches that most resemble the *signal* target patch. However, since the clean image is not available, at least in the first step, one can only work on *signal+noise* patches. Therefore, it can happen that some patches happen to be close to the target not because of an actual similarity of signal but as the effect of the random patterns of noise. This event, relatively uncommon in the AWGN case, can become a serious problem in the presence of strong correlated noise, when independent noise samples are reduced. If the noise is very structured and comparable in intensity with the signal it can dominate the block matching phase, leading to the selection of patches loosely related (in terms of signal) with one another and, eventually, to a poor performance. Therefore, in nonlocal approaches it is very important to counter this problem. To this end, we carry out a prewhitening of the noisy image. Let *z* be the observed noisy image, related to the noise-free image *y* by

$$z(n) = y(n) + h(n) * u(n) \qquad (2)$$

where *u(n)* is stationary white noise independent of *y(n)*, and *h(n)* is a linear filter. The prewhitened image $z_{pw}(n)$ is then computed as the inverse transform of

$$Z_{pw}(\omega) = Z(\omega)\frac{1}{\max(\epsilon, |H(\omega)|)} \qquad (3)$$

where $X(\omega)$ indicated discrete Fourier transform of $x(n)$, and ε is a small positive constant added to ensure stability. The prewhitened image is then used to locate the best matching patches, while all other processing steps take place on the original image. It is worth underlying that this approach is quite different from that of Dabov et al. (2008), let us call it BM3D-C, where no prewhitening is carried out but block similarity is computed in the transform domain with a weighted Euclidean distance, with smaller weights associated to noisier coefficients to reduce their detrimental effects.

On the contrary, for our second modification, concerning coefficient shrinkage, we follow closely the approach proposed by Dabov et al. (2008). In particular, focusing on the first step of BM3D, we carry out hard thresholding using a different threshold for each coefficient, proportional to the expected variance of noise. In formulas

$$\lambda(i,j) = \lambda\, \sigma_{2D}(i) \qquad (4)$$

where i and j are indexes associated with the 2D spatial and 1D transforms, respectively, and λ is a constant. By so doing, we increase our chances to suppress large coefficients originated exclusively by noise and, at the same time, to keep small coefficients with significant signal contribution. As for the variances, following the model in (2), they can be computed as:

$$\sigma_{2D}^2(i) = \frac{\sigma^2}{N_1^2}\left\| |H(\omega)| F\{\psi_{T_{2D}}^{(i)}\} \right\|_2^2 \qquad (5)$$

where $\psi_{T_{2D}}^{(i)}$ is the i-th basis element of T_{2D}, $F\{\cdot\}$ is the discrete Fourier transform operator, σ^2 is the overall noise variance and N_1 is the block dimension.

Our last improvement concerns the empirical Wiener filtering in the second step which, just like hard thresholding in the first step, can be adapted to take into account the actual noise variances of coefficients according to

$$W(i,j) = \frac{\left|T_{3D}^{wie}(Y_{basic}^{z^T})(i,j)\right|^2}{\left|T_{3D}^{wie}(Y_{basic}^{z^T})(i,j)\right|^2 + \sigma_{2D}^2(i)} \qquad (6)$$

where $\left|T_{3D}^{wie}(Y_{basic}^{z^T})\right|^2 (i,j)$ is the energy of the 3D transform coefficients of the basic estimate group where patches more similar to the target one P^T of z^T are collected.

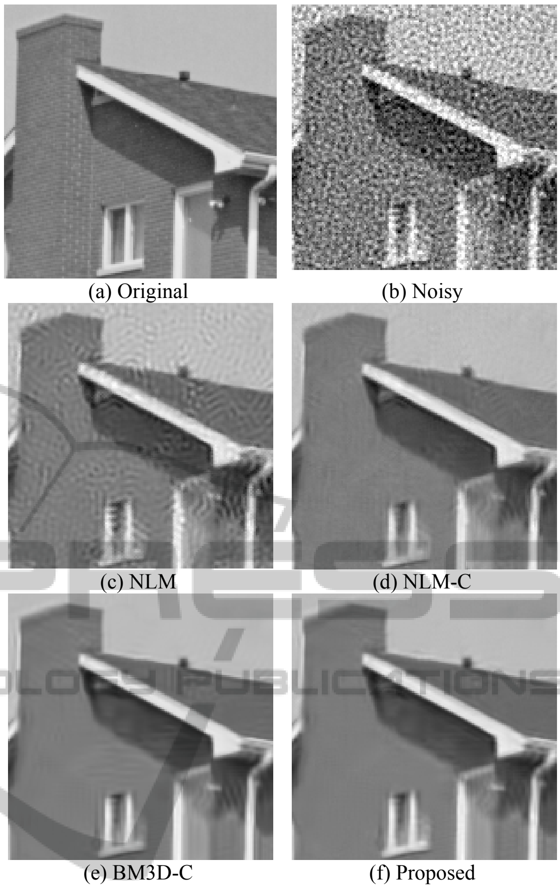

Figure 1: Visual results for a crop-outs of House, corrupted with correlated noise (σ=30); (a) original image, (b) noisy image, (c) NLM filtered image, (d) NLM-C filtered image, (e) BM3D-C filtered image, (f) proposed technique filtered image.

Here, we propose to estimate the variances $\sigma_{2D}^2(i)$ starting from the basic estimate of the clean image, $\hat{y}_{basic}$, provided by the first step. As a matter of facts, BM3D exploits the knowledge of $\hat{y}_{basic}$ in the second step to accomplish several tasks: $\hat{y}_{basic}$ is used to carry out the block matching process, which is why prewhitening is not required anymore, and also to obtain reliable estimates of signal statistics for the Wiener filtering. In the same manner, assuming that $\hat{y}_{basic}$ is a reliable estimate of the clean image y, the difference between the noisy image and the denoised one:

$$\hat{v}(n) = z(n) - \hat{y}_{basic}(n) \approx h(n) * u(n) \qquad (7)$$

can be assumed to be a good estimate of the actual correlated noise. Therefore, by working on this noise image we can actually *measure* the coefficient noise variances, in each single group of blocks, rather than estimating them according to (5).

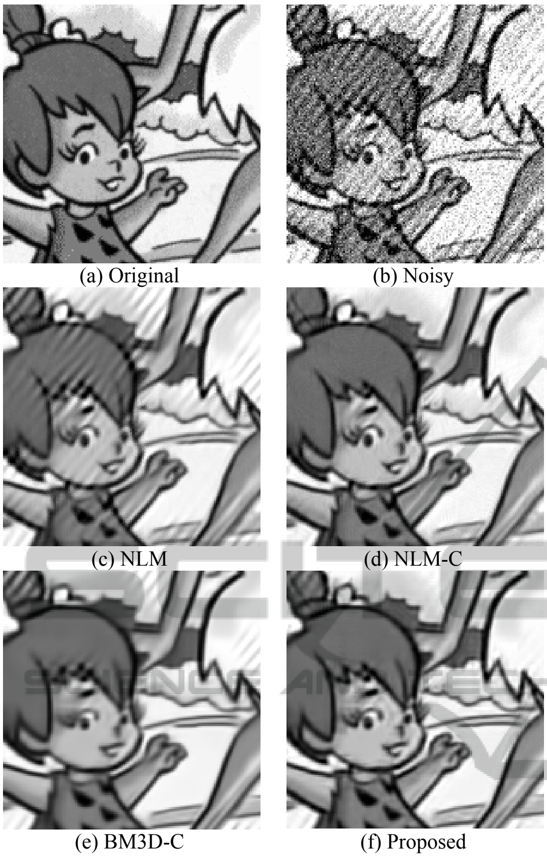

Figure 2: Visual results for a crop-outs of Flinstone, corrupted with correlated noise (σ=40); (a) original image, (b) noisy image, (c) NLM filtered image, (d) NLM-C filtered image, (e) BM3D-C filtered image, (f) proposed technique filtered image.

3 EXPERIMENTAL RESULTS

In this Section we describe the results of a limited set of experiments chosen to allow a comparison with reference techniques and highlight the major phenomena of interest.

In particular, we have simulated four type of colored noise: a bandpass noise with σ=30 on "House" image (Figure 1), and a line pattern noise like that found on analogue video with σ=40 on "Flinstones" image (Figure 2), exactly as they appear in the work of Goossens et al. (2008). In addition, we simulated two other types of colored noise, one characterized by a power spectral density with circulsar symmetry, as usually found in digital pictures, with σ=35 on "House" image (Figure 3) and a double line pattern noise with σ=30 again on "House" image (Figure 4), as a further test of the efficacy of the proposed method.

Focusing on the numerical results reported in Table 1, we note first of all that both NLM and BM3D degrade their performance as the noise becomes more structured, very likely because of the detrimental effects of noise on block matching.

In almost all cases, BM3D keeps a 2dB edge w.r.t. NLM, confirming that the BM3D two-steps structure leads to better results.

Obviously, the colored-noise versions of the algorithms, NLM-C and BM3D-C, provide a significant gain w.r.t. the basic versions, especially for the case of more structured noise. Note that the gain of NLM-C over NLM is much stronger than that of BM3D-C over BM3D, especially for structured noise.

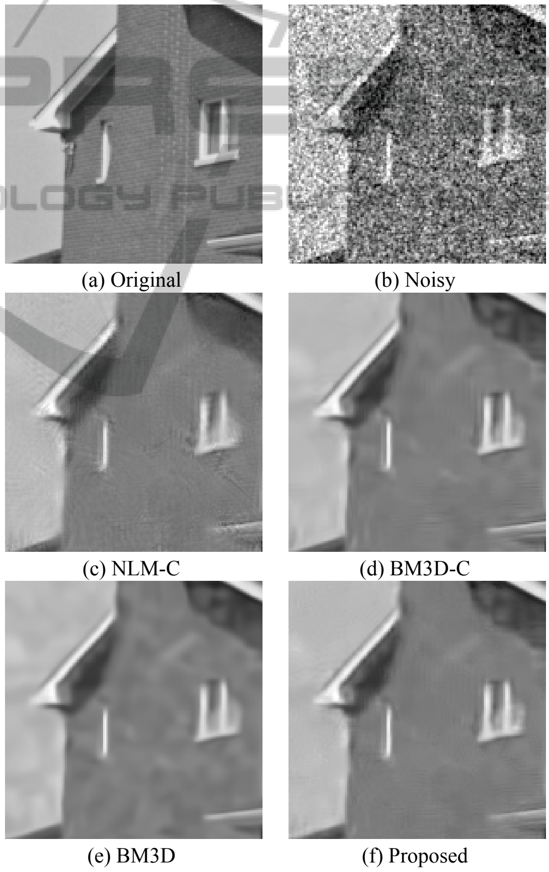

Figure 3: Visual results for a crop-outs of House, corrupted with correlated noise (σ=35); (a) original image, (b) noisy image, (c) NLM-C filtered image, (d) BM3D-C filtered image, (e) BM3D filtered image, (f) proposed technique filtered image.

This backs up our conjecture that strong structured noise very much impairs block matching, and that the solution proposed by Dabov et al. (2008) for BM3D does not really solve the problem.

On the contrary, using the pre-whitened image for the block matching, the risk of patch mis-classifications is reduced. In fact, plain BM3D with prewhitening provides already a performance comparable or superior to that of BM3D-C, see Table 1.

Table 1: PSNR results.

	House (σ=35)	House (σ=30)		Flinst. (σ=40)
	circular psd	band pass	stripes	stripes
Noisy	17.25	18.59	18.59	16.09
NLM	26.68	28.51	26.67	22.51
BM3D	28.12	30.94	28.91	22.01
NLM-C	28.11	30.74	32.01	25.44
BM3D-C	28.77	31.44	30.78	23.63
BM3D/p.w.	28.50	31.38	31.04	24.23
Proposed	**28.77**	**31.62**	**32.26**	**25.54**

Our proposed version of BM3D, which includes also improvements in the shrinkage phase and in the variance estimation, turns out to work well with all types of noise (unlike the reference techniques which fail in one or another situations) and provides consistently the best performance.

For example, the performance of NLM-C is good on streaked noise, but not so much for granular noise, even worse than BM3D developed in AWGN hypotheses. This behaviour is probably due to noise spectral characteristics. The whitening operation before filtering is probably more “invasive” for band pass noise than for other types, altering the underlying signal characteristics, thereby reducing the denoising effectiveness.

On the contrary, BM3D-C returns better results for the band pass noise and only modest results for streaked noise. Although suitable estimates of the transform coefficient variances are used in block-matching to reduce the influence of noisier coefficients, the particular and repetitive pattern of streaked noise strongly influences the block distance measure. If the block matching is compromised, the remaining filtering part is ineffective.

Cleverly combining the two approaches and exploiting the two-step BM3D structure, the method proposed in this work leads to robust results on all types of correlated noise simulated.

The visual inspection of the filtered images from Figure 1 to Figure 4 further reinforces the positive judgement on the proposed algorithm. In addition to the reference image and the noisy one, in Figure 1 and Figure 2 we show a visual comparison between our technique and NLM-C, BM3D-C and NLM developed for white noise, as suggested by Goossens et al. (2008) in their work. In Figure 3 and Figure 4, instead, we have chosen BM3D as a comparison technique developed for white noise, because it is the basis of the technique for removing correlated noise that we propose.

In any case, the proposed method removes most of the noise, like the other colored-noise algorithms, but presents a smaller number of ghost artifacts and of generally lower intensity.

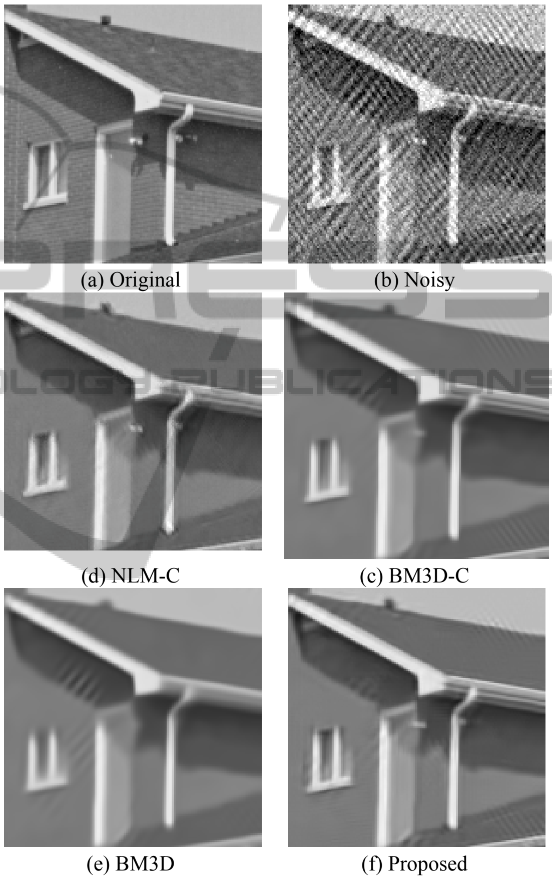

Figure 4: Visual results for a crop-outs of House, corrupted with correlated noise (σ=30); (a) original image, (b) noisy image, (c) NLM-C filtered image, (d) BM3D-C filtered image, (e) BM3D filtered image, (f) proposed technique filtered image.

4 CONCLUSIONS

Since noise is a primary cause of reduced image analysis capability in many application fields, in this work we focus on the problem of correlated noise removal. To this regard, we introduce a modified version of BM3D for correlated noise reduction, in which three significant changes are introduced with respect to the original algorithm. In particular, we improve the block matching by resorting to image prewhitening, and the shrinkage (hard thresholding in the first step, and Wiener filtering in the second step) by taking into account the different noise variances of coefficients and by improving their estimate.

This method turns out to be competitive both with state of art nonlocal algorithm for colored noise and with the original BM3D.

REFERENCES

Buades, A., Coll, B. and Morel, J. M., (2005). A Review of Image Denoising Algorithms, with a New One. *Multiscale Modeling & Simulation*, 4(2), 490-530.

Dabov, K., Foi, A., Katkovnik, V. and Egiazarian, K., (2007). Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Transactions on Image Processing,* 16(8), 2080-95.

Dabov, K., Foi, A., Katkovnik, V. and Egiazarian, K., (2008). Image restoration by sparse 3D transform-domain collaborative filtering. *Proceedings of SPIE*, 213462(213462), 681207-12.

Deledalle, C., Denis, L. and Tupin, F., (2009). Iterative weighted maximum likelihood denoising with probabilistic patch-based weights. *IEEE Transactions on Image Processing,* 18(12), 2661-72.

Donoho, D. and Johnstone, J. M. (1994). Ideal spatial adaptation by wavelet shrinkage. *Biometrika*, 81(3), 425-455.

Goossens, B. and Pizurica, A., (2009). Image denoising using mixtures of projected Gaussian scale mixtures. *Image Processing, IEEE*.

Goossens, B., Luong, Q., Pizurica, A. and Philips, W., (2008). An improved non-local denoising algorithm. In *2008 Proc. Int. Workshop on Local and Non-Local Approximation in Image Processing*. 1026–3.

Parrilli, S., Poderico, M., Angelino, C. V. and Verdoliva, L., (2012). A Nonlocal S AR Image Denoising Algorithm Based on LLMMSE Wavelet Shrinkage. *IEEE Transactions on Geoscience and Remote Sensing* (99), 1–11.

Portilla, J., Strela, V., Wainwright, M. J. and Simoncelli, E. P., (2003). Image denoising using scale mixtures of Gaussians in the wavelet domain. *IEEE transactions on image processing*, 12(11), 1338-51.