

Image restoration by sparse 3D transform-domain collaborative filtering

Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian

Department of Signal Processing, Tampere University of Technology

P.O. Box 553, 33101 Tampere, Finland

firstname.lastname@tut.fi

ABSTRACT

We propose an image restoration technique exploiting regularized inversion and the recent block-matching and 3D filtering (BM3D) denoising filter. The BM3D employs a non-local modeling of images by collecting similar image patches in 3D arrays. The so-called collaborative filtering applied on such a 3D array is realized by transform-domain shrinkage. In this work, we propose an extension of the BM3D filter for colored noise, which we use in a two-step deblurring algorithm to improve the regularization after inversion in discrete Fourier domain. The first step of the algorithm is a regularized inversion using BM3D with collaborative hard-thresholding and the second step is a regularized Wiener inversion using BM3D with collaborative Wiener filtering. The experimental results show that the proposed technique is competitive with and in most cases outperforms the current best image restoration methods in terms of improvement in signal-to-noise ratio.

Keywords: image restoration, deconvolution, deblurring, block-matching, collaborative filtering

1. INTRODUCTION

Image blurring is a common degradation in imaging. In many cases, the blurring can be assumed space-invariant and thus modeled as a convolution of the true image with a fixed point-spread function (PSF). Such a model is given by

$$z(x) = (y \otimes v)(x) + \eta(x), \quad (1)$$

where y is the true (non-degraded) image, v is a blur PSF, η is i.i.d. Gaussian noise with zero mean and variance σ^2 , and $x \in X$ is a 2D coordinate in the image domain X . The inversion of the blurring is in general an ill-posed problem; thus, even noise with very small magnitude, such as truncation noise due to limited-precision arithmetic, can cause extreme degradations after naive inversion. Regularization is a well known and extensively studied approach to alleviate this problem. It imposes some regularity conditions (e.g., smoothness) on the obtained image estimate and/or on its derivatives. Numerous approaches that employ regularization have been proposed; an introduction can be found for example in the books.^{1,2} In particular, an image restoration scheme that comprises of regularized inversion followed by denoising has been a basis of the current best-performing restoration methods.^{3,4} Such denoising after the inversion can be considered as part of the regularization since it attenuates the noise in the obtained solution (i.e. the solution is smoothed).

Various denoising methods can be employed to suppress the noise after the inversion. Filtering in multiresolution transform domain (e.g., overcomplete wavelet and pyramid transforms) was shown⁴⁻⁶ to be effective for this purpose. In particular, the SV-GSM,⁴ which employs Gaussian scale mixtures in overcomplete directional and multiresolution pyramids, is among the current best image deblurring methods. Another denoising technique used after regularized inversion^{3,7,8} is the LPA-ICI⁹ which exploits a non-parametric local polynomial fit in anisotropic estimation neighborhoods. The best results of the methods based on LPA-ICI were achieved by the shape-adaptive discrete cosine transform (SA-DCT) deblurring³ where the denoising is realized by shrinkage of the SA-DCT applied on local neighborhoods whose arbitrary shapes are defined by the LPA-ICI.

This work was partly supported by the Academy of Finland, project No. 213462 (Finnish Centre of Excellence program [2006 - 2011]); the work of K. Dabov was supported by the Tampere Graduate School in Information Science and Engineering (TISE).

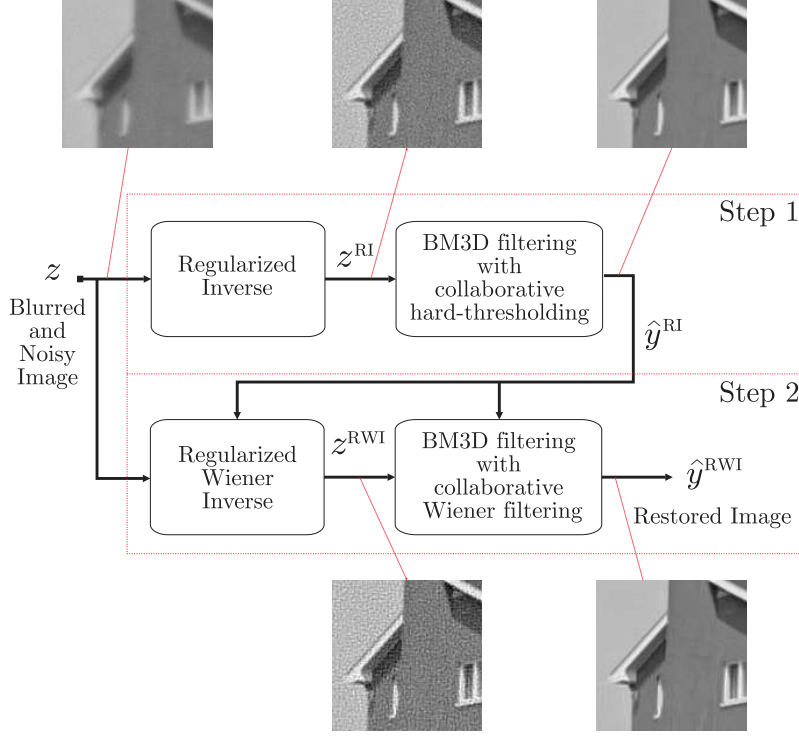


Figure 1. Flowchart of the proposed deconvolution algorithm. A fragment of *House* illustrates the images after each operation.

In this work we follow the above restoration scheme (regularized inversion followed by denoising) exploiting an extension of the block-matching and 3D filtering¹⁰ (BM3D). This filter is based on the assumption that there exist mutually similar patches within a natural image — the same assumption used in other non-local image filters such as.^{11,12} The BM3D processes a noisy image in a sliding-window (block) manner, where block-matching is performed to find blocks similar to the currently processed one. The blocks are then stacked together to form a 3D array and the noise is attenuated by shrinkage in a 3D-transform domain. This results in a 3D array of filtered blocks. A denoised image is produced by aggregating the filtered blocks to their original locations using weighted averaging. This filter was shown¹⁰ to be highly effective for attenuation of additive i.i.d. Gaussian (white) noise. The contribution of this work includes

- extension of the BM3D filter for additive colored noise, and
- image deblurring method that exploits the extended BM3D filter for improving the regularization after regularized inversion in Fourier transform domain.

The paper is organized as follows. The developed image restoration method and the extension of the BM3D filter are presented in Sections 2. Simulation results and a brief discussion are given in Section 3 and relevant conclusions are made in Section 4.

2. IMAGE RESTORATION WITH REGULARIZATION BY BM3D FILTERING

The observation model given in Equation (1) can be expressed in discrete Fourier transform (DFT) domain as

$$Z = YV + \tilde{\eta}, \tag{2}$$

where Y , V , and $\tilde{\eta}$ are the DFT spectra of y , v , and η , respectively. Capital letters denote DFT of a signal; e.g. $Z = \mathcal{F}\{z\}$, $V = \mathcal{F}\{v\}$; the only exception in that notation is for $\tilde{\eta} = \mathcal{F}\{\eta\}$. Due to the normalization of the

Postprint note: in Eq. (3), as in the rest of the paper, the straight vertical brackets are used for two different purposes:
1. to denote the modulus (i.e., magnitude) of a transform spectrum, as in $|V|$,
2. to denote the cardinality of a set, as in $|X|$.

forward DFT, the variance of $\tilde{\eta}$ is $|X|\sigma^2$, where $|X|$ is the cardinality of the set X (i.e., $|X|$ is the number of pixels in the input image).

Given the input blurred and noisy image z , the blur PSF v , and the noise variance σ^2 , we apply the following two-step image deblurring algorithm, which is illustrated in Figure 1.

Proposed two-step image deblurring algorithm

Step 1. Regularized Inversion (RI) using BM3D with collaborative hard-thresholding.

1.1. The regularized inverse z^{RI} is computed in DFT domain as

$$\begin{aligned} T^{\text{RI}} &= \frac{\bar{V}}{|V|^2 + \alpha_{\text{RI}} |X| \sigma^2} \\ z^{\text{RI}} &= \mathcal{F}^{-1} \{ T^{\text{RI}} Z \} \\ &= \mathcal{F}^{-1} \left\{ Y \frac{|V|^2}{|V|^2 + \alpha_{\text{RI}} |X| \sigma^2} \right\} + \mathcal{F}^{-1} \{ \tilde{\eta} T^{\text{RI}} \}, \end{aligned} \quad (3)$$

where α_{RI} is a regularization parameter determined empirically. Note that the obtained inverse z^{RI} is the sum of $\mathcal{F}^{-1} \left\{ Y \frac{|V|^2}{|V|^2 + \alpha_{\text{RI}} |X| \sigma^2} \right\}$, a biased estimate of y , and the colored noise $\mathcal{F}^{-1} \{ \tilde{\eta} T^{\text{RI}} \}$.

1.2. Attenuate the colored noise in z^{RI} given by Eq. (3) using BM3D with collaborative hard-thresholding (see Section 2.1); the denoised image is denoted \hat{y}^{RI} .

Step 2. Regularized Wiener inversion (RWI) using BM3D with collaborative Wiener filtering.

2.1. Using \hat{y}^{RI} as a reference estimate, compute the regularized Wiener inverse z^{RWI} as

$$\begin{aligned} T^{\text{RWI}} &= \frac{\bar{V} |\hat{Y}^{\text{RI}}|^2}{|V \hat{Y}^{\text{RI}}|^2 + \alpha_{\text{RWI}} |X| \sigma^2}, \\ z^{\text{RWI}} &= \mathcal{F}^{-1} \{ T^{\text{RWI}} Z \}, \\ z^{\text{RWI}} &= \mathcal{F}^{-1} \left\{ Y \frac{|V \hat{Y}^{\text{RI}}|^2}{|V \hat{Y}^{\text{RI}}|^2 + \alpha_{\text{RWI}} |X| \sigma^2} \right\} + \mathcal{F}^{-1} \{ \tilde{\eta} T^{\text{RWI}} \} \end{aligned} \quad (4)$$

where, analogously to Eq. (3), α_{RWI} is a regularization parameter and z^{RWI} is the sum of a biased estimate of y and colored noise.

2.2. Attenuate the colored noise in z^{RWI} using BM3D with collaborative Wiener filtering (see Section 2.2) which also uses \hat{y}^{RI} as a pilot estimate. The result \hat{y}^{RWI} of this denoising is the final restored image.

The BM3D filtering of the colored noise (Steps 1.2 and 2.2) plays the role of a further regularization of the sought solution. It allows the use of relatively small regularization parameters in the Fourier-domain inverses, hence reducing the bias in the estimates z^{RI} and z^{RWI} , which are instead essentially noisy. The BM3D denoising filter¹⁰ is originally developed for additive white Gaussian noise. Thus, to enable the attenuation of colored noise, we propose some modifications to the original filter.

Before we present the extensions that enable attenuation of colored noise, we recall how the BM3D filter works; for details of the original method one can refer to.¹⁰ The BM3D processes an input image in a sliding-window manner, where the window (block) has a fixed size $N_1 \times N_1$. For each processed block a 3D array is

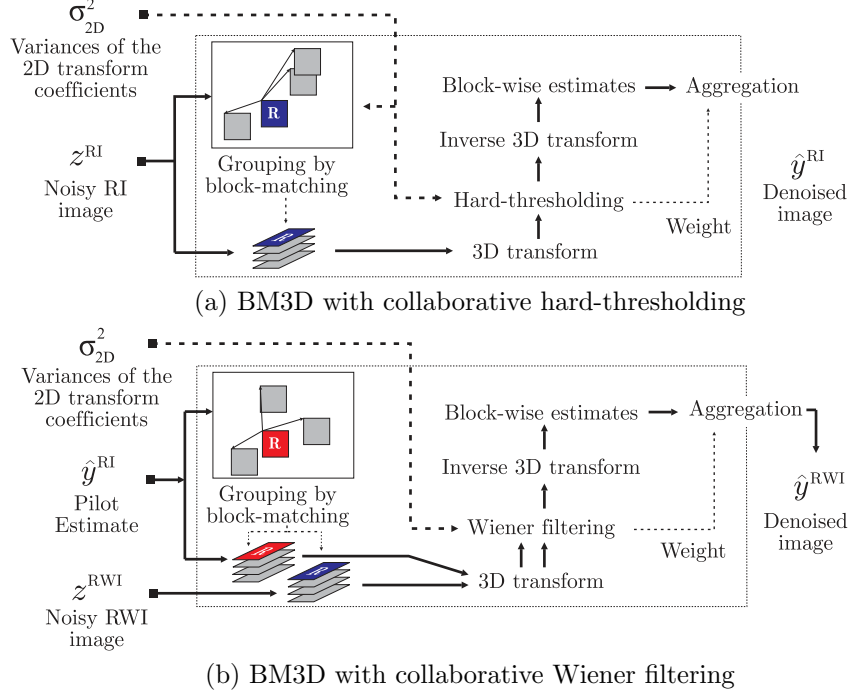


Figure 2. Flowcharts of the BM3D filter extensions for colored-noise removal.

formed by stacking together blocks (from various image locations) which are similar to the current one. This process is called “grouping” and is realized by block-matching. Consequently, a separable 3D transform \mathcal{T}_{3D} is applied on the 3D array in such a manner that first a 2D transform, \mathcal{T}_{2D} , is applied on each block in the group and then a 1D transform, \mathcal{T}_{1D} , is applied in the third dimension. The noise is attenuated by shrinkage (e.g. hard-thresholding or empirical Wiener filtering) of the \mathcal{T}_{3D} -transform spectrum. Subsequently, the transform \mathcal{T}_{3D} is inverted and each of the filtered blocks in the group is returned to its original location. After processing the whole image, since the filtered blocks can (and usually do) mutually overlap, they are aggregated by weighted averaging to form a final denoised image.

If the transforms \mathcal{T}_{2D} and \mathcal{T}_{1D} are orthonormal, the grouped blocks are non-overlapping, and the noise in the input image is i.i.d. Gaussian, then the noise in the \mathcal{T}_{3D} -transform domain is also i.i.d. Gaussian with the same constant variance. However, if the noise is colored as in the case of Eq. (3), then the variances $\sigma_{2D}^2(i)$, for $i = 1, \dots, N_1^2$, of the \mathcal{T}_{2D} -transform coefficients are in general non constant. In the following subsections, we extend the BM3D filter to attenuate such colored noise. We note that the developed extensions are not necessarily restricted to the considered image restoration scheme but are applicable to filtering of colored noise in general.

Let us introduce the notation used in what follows. With z_x^{RI} we denote a 2D block of fixed size $N_1 \times N_1$ extracted from z^{RI} , where $x \in X$ is the coordinate of the top-left corner of the block. Let us note that this block notation is different from the one (capital letter with subscript) used in¹⁰ since the capital letter in this paper is reserved for the DFT of an image. A group of collected 2D blocks is denoted by a bold-face letter with a subscript indicating the set of its grouped blocks’ coordinates: e.g., \mathbf{z}_S^{RI} is a 3D array composed of the blocks $z_x^{\text{RI}}, \forall x \in S \subseteq X$.

2.1. BM3D with collaborative hard-thresholding (Step 1.2)

This filtering is applied on the noisy z^{RI} given by Eq. (3). The variances of the coefficients of a \mathcal{T}_{2D} -transform (applied to an arbitrary image block) are computed as

$$\sigma_{2D}^2(i) = \frac{\sigma^2}{|X|} \left\| T^{\text{RI}} \mathcal{F} \left\{ \psi_{\mathcal{T}_{2D}}^{(i)} \right\} \right\|_2^2, \quad \forall i = 1, \dots, N_1^2, \quad (5)$$

where $\psi_{\mathcal{T}_{2D}}^{(i)}$ is the i -th basis element of \mathcal{T}_{2D} . The flowchart of the BM3D with collaborative hard-thresholding extended for color-noise attenuation is given in Figure 2(a).

The variances σ_{2D}^2 are used in the block-matching to reduce the influence of noisier transform coefficients when determining the block-distance. To accomplish this, the block-distance is computed as the ℓ^2 -norm of the difference between the two \mathcal{T}_{2D} -transformed blocks scaled by the corresponding standard deviations of the \mathcal{T}_{2D} -transform coefficients. Thus, the distance is given by

$$d(z_{x_R}^{\text{RI}}, z_x^{\text{RI}}) = N_1^{-2} \left\| \frac{\mathcal{T}_{2D} \{z_{x_R}^{\text{RI}}\} - \mathcal{T}_{2D} \{z_x^{\text{RI}}\}}{\sigma_{2D}} \right\|_2^2, \quad (6)$$

where $z_{x_R}^{\text{RI}}$ is the current reference block, z_x^{RI} is an arbitrary block in the search neighborhood, and the operations between the three $N_1 \times N_1$ arrays $\mathcal{T}_{2D} \{z_{x_R}^{\text{RI}}\}$, $\mathcal{T}_{2D} \{z_x^{\text{RI}}\}$, and σ_{2D} are elementwise. After the best-matching blocks are found (their coordinates are saved as the elements of the set S_{x_R}) and grouped together in a 3D array, collaborative hard-thresholding is applied. It consists of applying the 3D transform \mathcal{T}_{3D} on the 3D group, hard-thresholding its spectrum, and then inverting the \mathcal{T}_{3D} . To attenuate the colored noise, the hard-threshold is made dependent on the variance of each \mathcal{T}_{3D} -transform coefficient. Due to the separability of \mathcal{T}_{3D} , this variance depends only on the corresponding 2D coordinate within the \mathcal{T}_{3D} -spectrum; thus, along the third dimension of a group the variance and hence the threshold are the same. The hard-thresholding is performed by an elementwise multiplication of the \mathcal{T}_{3D} -spectrum $\mathcal{T}_{3D} \{z_{S_{x_R}}^{\text{RI}}\}$ with the 3D array \mathbf{h}_{x_R} defined as

$$\mathbf{h}_{x_R}(i, j) = \begin{cases} 1, & \text{if } \left| \mathcal{T}_{3D} \{z_{S_{x_R}}^{\text{RI}}\}(i, j) \right| > \lambda_{3D} \sigma_{2D}(i) \\ 0, & \text{otherwise,} \end{cases} \quad \forall i = 1, \dots, N_1^2, \forall j = 1, \dots, |S_{x_R}|,$$

where i is a spatial-coordinate index and j is an index of the coefficients in the third dimension, λ_{3D} is a fixed threshold coefficient and $|S_{x_R}|$ denotes the cardinality of the set S_{x_R} .

After all reference blocks are processed, the filtered blocks are aggregated by a weighted averaging, producing the denoised image \hat{y}^{RI} . The weight for all filtered blocks in an arbitrary 3D group is the inverse of the sum of the variances of the non-zero transform coefficients after hard-thresholding; for a 3D group using $x_R \in X$ as reference, the weight is

$$w_{x_R}^{\text{ht}} = \frac{1}{\sum_{\substack{i=1, \dots, N_1^2 \\ j=1, \dots, |S_{x_R}|}} \mathbf{h}_{x_R}(i, j) \sigma_{2D}^2(i)}.$$

2.2. BM3D with collaborative Wiener filtering (Step 2.2)

The BM3D with collaborative empirical Wiener filtering uses \hat{y}^{RI} as a reference estimate of the true image y . Since the grouping by block-matching is performed on this estimate and not on the noisy image, there is no need to modify the distance calculation as in Eq. (6). The only modification from Step 2 of the original BM3D filter concerns the different variances of the \mathcal{T}_{3D} -transform coefficients in the empirical Wiener filtering. This filtering is performed by an elementwise multiplication of the \mathcal{T}_{3D} -spectrum $\mathcal{T}_{3D} \{z_{S_{x_R}}^{\text{RWI}}\}$ with the Wiener attenuation coefficients \mathbf{w}_{x_R} defined as

$$\mathbf{w}_{x_R}(i, j) = \frac{\left| \mathcal{T}_{3D} \{\hat{y}_{S_{x_R}}^{\text{RI}}\}(i, j) \right|^2}{\left| \mathcal{T}_{3D} \{\hat{y}_{S_{x_R}}^{\text{RI}}\}(i, j) \right|^2 + \sigma_{2D}^2(i)}, \quad \forall i = 1, \dots, N_1^2, \forall j = 1, \dots, |S_{x_R}|,$$

where, similarly to Eq. (5), the variances σ_{2D}^2 of the \mathcal{T}_{2D} -transform coefficients are computed as

$$\sigma_{2D}^2(i) = \frac{\sigma^2}{|X|} \left\| T^{\text{RWI}} \mathcal{F} \left\{ \psi_{\mathcal{T}_{2D}}^{(i)} \right\} \right\|_2^2, \quad \forall i = 1, \dots, N_1^2. \quad (7)$$

For an arbitrary $x_R \in X$, the aggregation weight for its corresponding filtered 3D group is

$$w_{x_R}^{\text{wie}} = \frac{1}{\sum_{\substack{i=1, \dots, N_T^2 \\ j=1, \dots, |S_{x_R}|}} \mathbf{w}_{x_R}^2(i, j) \sigma_{2D}^2(i)}.$$

The flowchart of the BM3D with collaborative Wiener filtering extended for color-noise attenuation is given in Figure 2(b).

3. RESULTS AND DISCUSSION

We present simulation results of the proposed algorithm, whose Matlab implementation is available online.¹³ All parameters, obtained after a rough empirical optimization, are fixed in all experiments (invariant of the noise variance σ^2 , blur PSF v , and image y) and can be inspected from the provided implementation. In our experiments, we used the same blur PSFs and noise combinations as in.⁴ In particular, these PSFs are:

- PSF 1: $v(x_1, x_2) = 1 / (1 + x_1^2 + x_2^2)$, $x_1, x_2 = -7, \dots, 7$,
- PSF 2: v is a 9×9 uniform kernel (boxcar),
- PSF 3: $v = [1 \ 4 \ 6 \ 4 \ 1]^T [1 \ 4 \ 6 \ 4 \ 1] / 256$,
- PSF 4: v is a Gaussian PSF with standard deviation 1.6,
- PSF 5: v is a Gaussian PSF with standard deviation 0.4.

All PSFs are normalized so that $\sum v = 1$.

Table 1 presents a comparison of the improvement in signal-to-noise ratio (ISNR) for a few methods^{3, 4, 6, 14–16} among which are the current best.^{3, 4} The results of ForWaRD⁶ were obtained with the Matlab codes¹⁷ made available by its authors, for which we used automatically estimated regularization parameters. The results of the SA-DCT deblurring³ were produced with the Matlab implementation,¹⁸ where however we used fixed regularization parameters in all experiment in order to have fair comparison (rather than using regularization parameters dependent on the PSF and noise). The results of the GSM method¹⁹ and the SV-GSM⁴ are taken from.⁴ In most of the experiments, the proposed method outperforms the other techniques in terms of ISNR. We note that the results of the four standard experiments used in the literature (e.g.,^{3, 6, 7, 20}) on image restoration are included in Table 1 as follows.

- Experiment 1: PSF2, $\sigma^2 = 0.308$, and *Cameraman* image.
- Experiment 2: PSF1, $\sigma^2 = 2$, and *Cameraman* image.
- Experiment 3: PSF1, $\sigma^2 = 8$, and *Cameraman* image.
- Experiment 4: PSF3, $\sigma^2 = 49$, and *Lena* image.

The visual quality of some of the restored images can be evaluated from Figures 4, 5, and 6. One can see that fine details are well preserved and there are few artifacts in the deblurred images. In particular, ringing can be seen in some images such as the ones shown in Figure 3, where a comparison with the SA-DCT deblurring³ is made. The ringing is stronger (and ISNR is lower) in the estimate obtained by the proposed technique. We explain this as follows; let us recall that each of the noisy images z^{RI} and z^{RWI} (input to the extended BM3D filter) is sum of a bias and additive colored noise; the exact models of z^{RI} and z^{RWI} are given by Eq. (3) and (4), respectively. The ringing is part of the bias and thus it is not modeled as additive colored noise. Hence, if the ringing magnitude is relatively high, the BM3D fails to attenuate it and it is preserved in the final estimate, as in Figure 3.

By comparing the results corresponding to \hat{y}^{RI} and \hat{y}^{RWI} in Table 2, one can see the improvement in ISNR after applying the second step (RWI using BM3D with collaborative Wiener filtering) of our two-step restoration scheme. This improvement is significant and can be explained as follows. First, the regularized Wiener inverse is more effective than the regularized inverse because it uses the estimated power spectrum for the inversion given

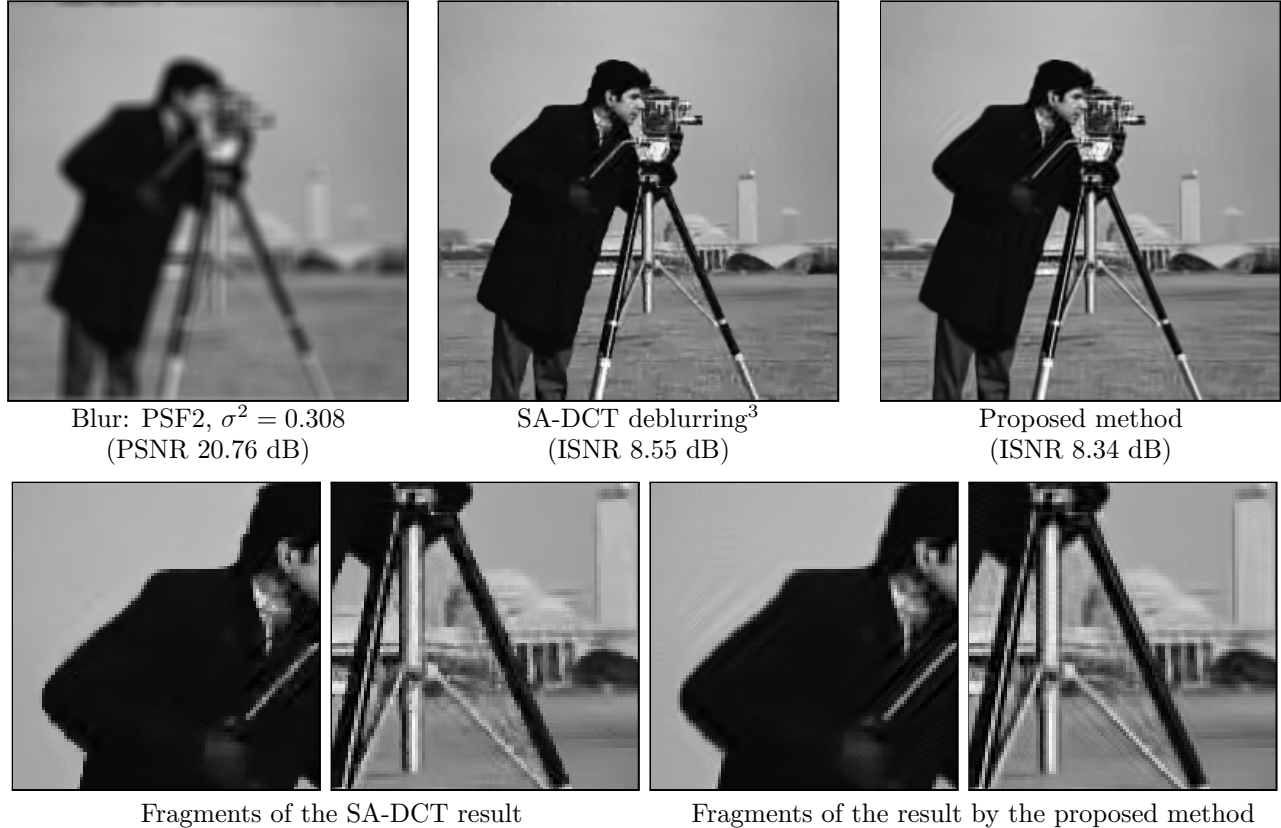


Figure 3. Comparison of the proposed method with the SA-DCT deconvolution method for *Cameraman* and PSF 2 blur kernel.

in Eq. (4). Second, the block-matching in the BM3D filtering is more accurate because it is performed within the available estimate \hat{y}^{RI} rather than within the input noisy image z^{RWI} . Third, the empirical Wiener filtering used by the BM3D in that step is more effective than the simple hard-thresholding used in the first step. In fact, the first step can be considered as an adaptation step that significantly improves the actual restoration performed by the second step.

In Table 2, we also provide (in the row corresponding to $\hat{y}_{\text{naive}}^{\text{RWI}}$) the results of the naive approach of using the original BM3D filter¹⁰ rather than the one extended for colored noise. This filter was applied on z^{RI} and z^{RWI} by assuming additive i.i.d. Gaussian noise, whose variance was computed as $\sigma_{\text{WGN}}^2 = N_1^{-2} \sum_{i=1}^{N_1^2} \sigma_{2\text{D}}^2(i)$, where $\sigma_{2\text{D}}^2(\cdot)$ is defined in Eq. (5) and (7) for z^{RI} and z^{RWI} , respectively. This variance calculation was empirically found to be better (in terms of ISNR) than estimating a noise variance from the noisy images z^{RI} and z^{RWI} . The benefit of using the BM3D for colored noise reaches 1 dB; in particular, the benefit is substantial for those experiments where the noise in z^{RI} and z^{RWI} is highly colored.

4. CONCLUSIONS

The developed image deblurring method outperforms the current best techniques in most of the experiments. This performance is in line with the BM3D denoising filter¹⁰ which is among the current best denoising filters. The proposed colored-noise extension of the BM3D is not restricted to the developed deblurring method and it can in general be applied to filter colored noise.

Future developments might target attenuation of ringing artifacts by exploiting the SA-DCT transform²¹ which, as shown in Figure 3, is effective in suppressing them.

Blur	PSF 1		PSF 2	PSF 3	PSF 4	PSF 5
σ^2	2	8	0.308	49	4	64
<i>Cameraman</i>						
Input PSNR	22.23	22.16	20.76	24.62	23.36	29.82
ForWaRD ⁶	6.76	5.08	7.34	2.40	3.14	3.92
GSM ¹⁹	6.84	5.29	-1.61	2.56	2.83	3.81
EM ¹⁶	6.93	4.88	7.59	-	-	-
Segm.-based Reg. ¹⁵	7.23	-	8.04	-	-	-
GEM ¹⁴	7.47	5.17	8.10	-	-	-
BOA ²⁰	7.46	5.24	8.16	-	-	-
Anis. LPA-ICI ⁷	7.82	5.98	8.29	-	-	-
SV-GSM ⁴	7.45	5.55	7.33	2.73	3.25	4.19
SA-DCT ³	8.11	6.33	8.55	3.37	3.72	4.71
Proposed	8.19	6.40	8.34	3.34	3.73	4.70
<i>Lena</i>						
Input PSNR	27.25	27.04	25.84	28.81	29.16	30.03
Segm.-based Reg. ¹⁵	-	-	-	1.34	-	-
GEM ¹⁴	-	-	-	2.73	-	-
BOA ²⁰	-	-	-	2.84	-	-
ForWaRD ⁶	6.05	4.90	6.97	2.93	3.50	5.42
EM ¹⁶	-	-	-	2.94	-	-
Anis. LPA-ICI ⁷	-	-	-	3.90	-	-
SA-DCT ³	7.55	6.10	7.79	4.49	4.08	5.84
Proposed	7.95	6.53	7.97	4.81	4.37	6.40
<i>House</i>						
Input PSNR	25.61	25.46	24.11	28.06	27.81	29.98
ForWaRD ⁶	7.35	6.03	9.56	3.19	3.85	5.52
GSM ¹⁹	8.46	6.93	-0.44	4.37	4.34	5.98
SV-GSM ⁴	8.64	7.03	9.04	4.30	4.11	6.02
SA-DCT ³	9.02	7.74	10.50	4.99	4.65	5.96
Proposed	9.32	8.14	10.85	5.13	4.56	7.21
<i>Barbara</i>						
Input PSNR	23.34	23.25	22.49	24.22	23.77	29.78
ForWaRD ⁶	3.69	1.87	4.02	0.94	0.98	3.15
GSM ¹⁹	5.70	3.28	-0.27	1.44	0.95	4.91
SV-GSM ⁴	6.85	3.80	5.07	1.94	1.36	5.27
SA-DCT ³	5.45	2.54	4.79	1.31	1.02	3.83
Proposed	7.80	3.94	5.86	1.90	1.28	5.80

Table 1. Comparison of the output ISNR [dB] of a few deconvolution methods (only the rows corresponding to “Input PSNR” contain PSNR [dB] of the input blurry images).



Blur: PSF1, $\sigma^2 = 2$



Output ISNR 9.32 dB



Blur: PSF1, $\sigma^2 = 8$



Output ISNR 8.14 dB



Blur: PSF2, $\sigma^2 = 0.308$



Output ISNR 10.85 dB

Figure 4. Deblurring results of the proposed method for *House*.



Blur: PSF3, $\sigma^2 = 49$



Output ISNR 4.81 dB




Blur: PSF4, $\sigma^2 = 4$



Output ISNR 4.37 dB



Blur: PSF5, $\sigma^2 = 64$



Output ISNR 6.40 dB

Figure 5. Deblurring results of the proposed method for *Lena*.



Blur: PSF1, $\sigma^2 = 8$



Output ISNR 3.94 dB



Blur: PSF2, $\sigma^2 = 0.308$



Output ISNR 5.86 dB



Blur: PSF3, $\sigma^2 = 49$



Output ISNR 1.90 dB

Figure 6. Deblurring results of the proposed method for *Barbara*.

Blur \rightarrow	PSF 1		PSF 2	PSF 3	PSF 4	PSF 5
$\sigma^2 \rightarrow$	2	8	0.308	49	4	64
\hat{y}^{RI}	7.13	5.16	7.52	2.31	3.23	2.46
\hat{y}^{RWI}	8.19	6.40	8.34	3.34	3.73	4.70
$\hat{y}_{\text{naive}}^{\text{RWI}}$	7.17	6.25	8.14	2.57	2.71	4.63

Table 2. ISNR comparison for: the basic estimate \hat{y}^{RI} ; the final estimate \hat{y}^{RWI} ; the final estimate $\hat{y}_{\text{naive}}^{\text{RWI}}$ obtained using the original BM3D filter instead of the one extended for colored noise. The test image was *Cameraman*.

REFERENCES

1. C. Vogel, *Computational Methods for Inverse Problems*, SIAM, 2002.
2. P. C. Hansen, *Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion*, SIAM, Philadelphia, 1997.
3. A. Foi, K. Dabov, V. Katkovnik, and K. Egiazarian, "Shape-Adaptive DCT for denoising and image reconstruction," in *Proc. SPIE Electronic Imaging: Algorithms and Systems V*, **6064A-18**, (San Jose, CA, USA), January 2006.
4. J. A. Guerrero-Colon, L. Mancera, and J. Portilla, "Image restoration using space-variant Gaussian scale mixtures in overcomplete pyramids," *IEEE Trans. Image Process.* **17**, pp. 27–41, January 2007.
5. R. Neelamani, H. Choi, and R. G. Baraniuk, "Wavelet-domain regularized deconvolution for ill-conditioned systems," in *Proc. IEEE Int. Conf. Image Process.*, **1**, pp. 204–208, (Kobe, Japan), October 1999.
6. R. Neelamani, H. Choi, and R. G. Baraniuk, "Forward: Fourier-wavelet regularized deconvolution for ill-conditioned systems," *IEEE Trans. Signal Process.* **52**, pp. 418–433, February 2004.
7. V. Katkovnik, A. Foi, K. Egiazarian, and J. Astola, "Directional varying scale approximations for anisotropic signal processing," in *Proc. European Signal Process. Conf.*, pp. 101–104, (Vienna, Austria), September 2004.
8. V. Katkovnik, K. Egiazarian, and J. Astola, "A spatially adaptive nonparametric regression image deblurring," *IEEE Trans. Image Process.* **14**, pp. 1469–1478, October 2005.
9. V. Katkovnik, K. Egiazarian, and J. Astola, *Local Approximation Techniques in Signal and Image Process.*, vol. PM157, SPIE Press, 2006.
10. K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3D transform-domain collaborative filtering," *IEEE Trans. Image Process.* **16**, pp. 2080–2095, August 2007.
11. A. Buades, B. Coll, and J. M. Morel, "A review of image denoising algorithms, with a new one," *Multiscale Modeling and Simulation* **4**(2), pp. 490–530, 2005.
12. C. Kervrann and J. Boulanger, "Optimal spatial adaptation for patch-based image denoising," *IEEE Trans. Image Process.* **15**, pp. 2866–2878, October 2006.
13. K. Dabov and A. Foi, "BM3D Filter Matlab Demo Codes." <http://www.cs.tut.fi/~foi/GCF-BM3D>.
14. J. Dias, "Fast GEM wavelet-based image deconvolution algorithm," in *Proc. IEEE Int. Conf. Image Process.*, **3**, pp. 961–964, (Barcelona, Spain), September 2003.
15. M. Mignotte, "An adaptive segmentation-based regularization term for image restoration," in *Proc. IEEE Int. Conf. Image Process.*, **1**, pp. 901–904, (Genova, Italy), September 2005.
16. M. A. T. Figueiredo and R. D. Nowak, "An EM algorithm for wavelet-based image restoration," *IEEE Trans. Image Process.* **12**, pp. 906–916, August 2003.
17. R. Neelamani, "Fourier-Wavelet Regularized Deconvolution (ForWaRD) Software for 1-D Signals and Images." <http://www.dsp.rice.edu/software/ward.shtml>.
18. A. Foi and K. Dabov, "Pointwise Shape-Adaptive DCT Demobox." <http://www.cs.tut.fi/~foi/SA-DCT>.
19. J. Portilla and E. P. Simoncelli, "Image restoration using Gaussian scale mixtures in the wavelet domain," in *Proc. IEEE Int. Conf. Image Process.*, **2**, pp. 965–968, (Barcelona, Spain), September 2003.
20. M. A. T. Figueiredo and R. Nowak, "A bound optimization approach to wavelet-based image deconvolution," in *Proc. IEEE Int. Conf. Image Process.*, **2**, pp. 782–785, (Genova, Italy), September 2005.
21. A. Foi, V. Katkovnik, and K. Egiazarian, "Pointwise Shape-Adaptive DCT for high-quality denoising and deblocking of grayscale and color images," *IEEE Trans. Image Process.* **16**, pp. 1395–1411, May 2007.