JOINT IMAGE SHARPENING AND DENOISING BY 3D TRANSFORM-DOMAIN COLLABORATIVE FILTERING

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

Institute of Signal Processing, Tampere University of Technology,
P.O. Box 553, Tampere 33101, Finland
firstname.lastname@tut.fi

ABSTRACT

In order to simultaneously sharpen image details and attenuate noise, we propose to combine the recent block-matching and 3D filtering (BM3D) denoising approach, based on 3D transform-domain collaborative filtering, with alpha-rooting, a transform-domain sharpening technique. The BM3D exploits grouping of similar image blocks into 3D arrays (groups) on which collaborative filtering (by hard-thresholding) is applied. We propose two approaches of sharpening by alpha-rooting; the first applies alpha-rooting individually on the 2D transform spectra of each grouped block; the second applies alpha-rooting on the 3D-transform spectra of each 3D array in order to sharpen fine image details shared by all grouped blocks and further enhance the interblock differences. The conducted experiments with the proposed method show that it can preserve and sharpen fine image details and effectively attenuate noise.

1. INTRODUCTION

The problem of sharpening image details while attenuating noise arises in practice as a preprocessing step prior to segmentation, binarization (by thresholding) in the context of image analysis, classification, pattern recognition. The procedure is also very useful in consumer applications, where it is utilized to obtain more visually appealing digital photographs. Various methods for image sharpening exist. Traditional ones rely on linear filtering that boosts higher frequencies or on elementwise transformations such as the histogram-based methods (equalization, matching, and shaping). Current advances [1, 2, 3] in the field exploit transforms such as wavelet decompositions and trigonometric transforms (DCT, DFT, etc.) to improve the efficiency of the sharpening. The sharpening is typically achieved by amplifying certain parts of the transform spectrum. Example of a well established such technique is the alpha-rooting [4, 5], where the amplification is achieved by taking the α-root of the magnitude of the original coefficients, for some constant $\alpha > 1$. A recent work [3] showed that histograms of the logarithm of a transform spectrum can also be utilized for sharpening using methods such as histogram matching and shaping.

This work was supported by the Academy of Finland, project No. 213462 (Finnish Centre of Excellence program [2006 - 2011]).

An inherent drawback of most sharpening methods is the amplification of the noise component that is inevitably present when dealing with practical applications.

Recently, we proposed a highly effective image denoising method [6], namely the block-matching and 3D filtering (BM3D), based on filtering of enhanced sparse non-local image representations. Using block-wise processing, we group blocks similar to the currently processed one into 3D arrays, termed "groups". Subsequently, we apply what we call "collaborative filtering" on each of these groups. In [6] we showed that this filtering can be efficiently realized by shrinkage (e.g., hard-thresholding) in 3D-transform domain. It is exactly in the 3D-transform domain where the enhanced sparse representation of the true signal is attained. Hence, the shrinkage of the transform coefficients allows for both good noise attenuation and faithful detail preservation. The result of the collaborative filtering of a group is a set of grouped block estimates, which are then returned to their original location in the image domain. Since the block estimates from a given group can mutually overlap and also overlap with ones from different groups, we aggregate them by a weighted averaging in order to obtain a single estimate of each image pixel.

In this paper we propose to combine the BM3D denoising approach with alpha-rooting, in order to simultaneously sharpen image details and attenuate noise. Hard-thresholding is applied on the 3D-transform spectrum of each group to attenuate the noise. Subsequently, sharpening is realized by modifying the thresholded transform coefficients of the grouped blocks. We show experimental results corresponding to two approaches, one with alpha-rooting performed on the 3D transform spectrum and the other with alpha-rooting applied separately on the 2D transform spectrum of each grouped block. The joint application of the collaborative filtering and alpha-rooting allows for both good noise suppression and effective preservation and sharpening of even very fine image details. It is worth noting that the adopted 3D transform is a separable composition of a 2D transform (on each block) and a 1D transform (in the "temporal" dimension, along which blocks are stacked), which makes possible the application of alpha-rooting on 2D transform spectra. In the following sections we present the developed method and show that it can be very effective when applied on both artificially degraded images and real retinal medical images.

2. PROPOSED METHOD

We consider a noisy image $z = y+\eta$, where y is the noise-free image (with poor contrast) and η is i.i.d. Gaussian noise with zero mean and variance σ^2. Following are the two variations of the proposed method, denominated BM3D-SH3D and BM3D-SH2D. The former performs alpha-rooting on 3D transform spectra and the latter on 2D transform spectra. Their corresponding flowcharts are shown in Figure 1.

1. Process overlapping blocks in a raster scan. For each such block, do the following:
   (a) Use block-matching to find the locations of the blocks in z that are similar to the currently processed one. Form a 3D array (group) by stacking the blocks located at the obtained locations.
   (b) Apply a 3D transform on the formed group.
   (c) Attenuate the noise by hard-thresholding the 3D transform spectrum.
   (d) Sharpening. We propose the following two alternatives.
      - BM3D-SH3D. Apply alpha-rooting on the hard-thresholded 3D transform spectrum and invert the 3D transform to produce filtered grouped blocks.
      - BM3D-SH2D. Invert the 1D transforms along the temporal dimension of the formed 3D array, then perform alpha-rooting separately on the 2D transform spectra of each grouped block and subsequently invert the 2D transforms to produce filtered grouped blocks.
2. Return the filtered blocks to their original locations in the image domain and compute the resultant filtered image by a weighted average of these filtered blocks.

Except for the alpha-rooting and a modification of the aggregation weights, both described below, the rest of the steps of the algorithm are taken without modification from the first step of BM3D [6] (the one that uses hard-thresholding). We refer the reader to [6] for details.

Given a transform spectrum t of a signal, which contains a DC coefficient termed $t(0)$, the alpha-rooting is performed as

$$t_{sh}(i) = \begin{cases} \operatorname{sign}[t(i)]\,|t(0)|\left|\frac{t(i)}{t(0)}\right|^{\frac{1}{\alpha}}, & \text{if } t(0) \neq 0 \\ t(i), & \text{otherwise,} \end{cases} \tag{1}$$

where t_{sh} is the spectrum of the resultant signal and where an exponent $\alpha > 1$ results in sharpening of image details [5]. This simple technique is applied for both the BM3D-SH3D and the BM3D-SH2D on 3D and 2D transform spectra, respectively. The alpha-rooting is motivated by fact that, except for the DC, the transform basis elements extract differential information of the signal. Its effectiveness is based on amplifying these coefficients according to Equation (1). Performing hard-thresholding prior to alpha-rooting resolves problems [7] of this sharpening technique regarding noise amplification (which causes spike artifacts).

As shown in [6], the overlapping filtered blocks can be effectively aggregated using weights that are inversely proportional to the total variance of the filtered groups to which the respective blocks belong. While for the denoising method described there, this total variance can be approximated as σ^2 times the number of retained (i.e. non-zero) coefficients after thresholding, in the case of joint denoising and sharpening, due to the amplification caused by alpha-rooting, the variance of each sharpened coefficient in the group could be different from the original σ^2.

We derive a simple estimator for the total residual variance in each group. Let t be the hard-thresholded 3D transform spectrum of a group and t_{sh} be the spectrum of the sharpened signal obtained by Equation 1. If $t(0) \neq 0$, we can rewrite $t_{sh}(i)$ as the product of two independent terms $\varrho_0 = |t(0)|^{1-\frac{1}{\alpha}}$ and $\rho_i = \operatorname{sign}[t(i)]\,|t(i)|^{\frac{1}{\alpha}}$. Assuming that the moduli of $t(0)$ and $t(i)$ are much larger than their standard deviation σ (which can be a reasonable assumption, since $t(0)$ is a DC coefficient and $t(i)$ survived thresholding), the variances of the product of these two terms can be roughly approximated (using the first partial derivatives) as follows

$$\begin{aligned} \operatorname{var}\{t_{sh}(i)\} &\simeq \left(\varrho_0'(E\{t(0)\})\,E\{\rho_i\}\right)^2 \operatorname{var}\{t(0)\} + \\ &\quad + \left(\rho_i'(E\{t(i)\})\,E\{\varrho_0\}\right)^2 \operatorname{var}\{t(i)\} \\ &\simeq \left(1-\frac{1}{\alpha}\right)^2 |E\{t(0)\}|^{-\frac{2}{\alpha}} E^2\{\rho_i\}\,\sigma^2 + \\ &\quad + \frac{1}{\alpha^2} |E\{t(i)\}|^{\frac{2}{\alpha}-2} E^2\{\varrho_0\}\,\sigma^2 \\ &\simeq \left(1-\frac{1}{\alpha}\right)^2 \left|\frac{E\{t(i)\}}{E\{t(0)\}}\right|^{\frac{2}{\alpha}} \sigma^2 + \\ &\quad + \frac{1}{\alpha^2} |E\{t(i)\}|^{\frac{2}{\alpha}-2} |E\{t(0)\}|^{2-\frac{2}{\alpha}} \sigma^2 \\ &\simeq \left(1-\frac{1}{\alpha}\right)^2 |t(0)|^{-\frac{2}{\alpha}} |t(i)|^{\frac{2}{\alpha}} \sigma^2 + \\ &\quad + \frac{1}{\alpha^2} |t(i)|^{\frac{2}{\alpha}-2} |t(0)|^{2-\frac{2}{\alpha}} \sigma^2 \\ &= \omega_i \sigma^2 \end{aligned}$$

Thus, the total variance of the filtered (i.e., thresholded and sharpened) group is approximated as

$$v = \sigma^2 + \sum_{t(i)\neq 0, i>0} \omega_i \sigma^2. \tag{2}$$

Consequently, the weights used for the aggregation are defined as the reciprocal of the above v. Note that these

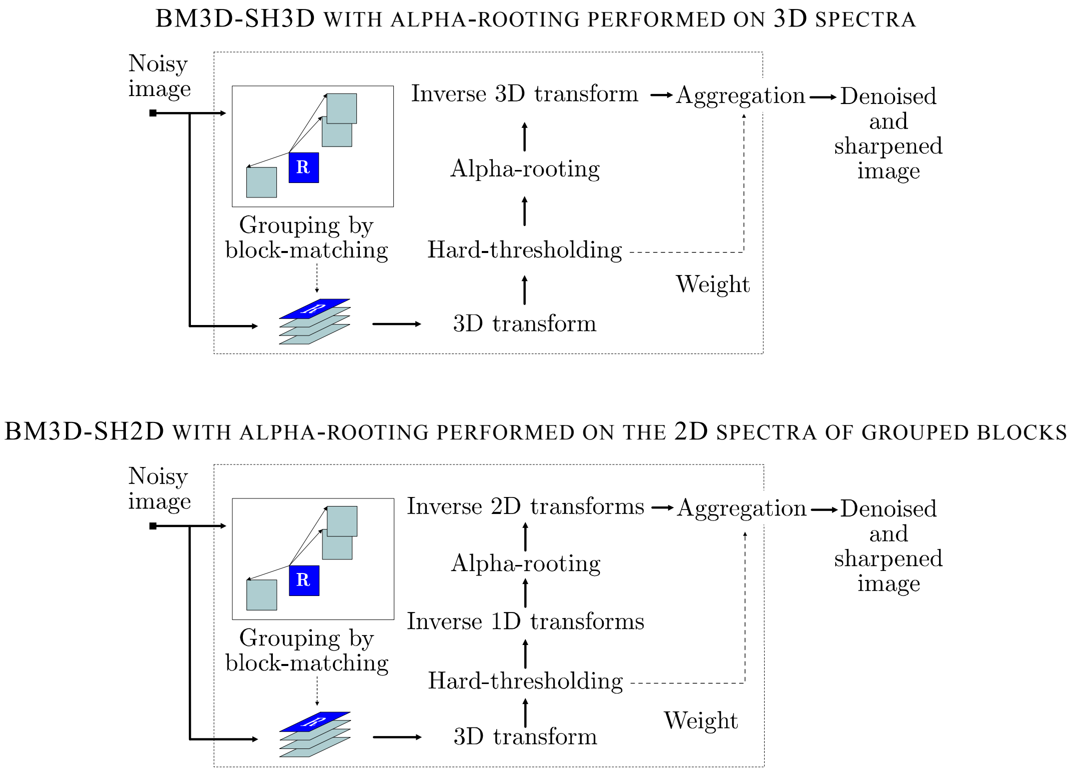

Figure 1. Flowcharts of the two variations of the proposed method. Top: BM3D-SH3D, which applies alpha-rooting on the 3D spectrum of each group; bottom: BM3D-SH2D, which applies alpha-rooting on the 2D spectrum of each grouped block.

weights are used only for BM3D-SH3D and for the BM3D-SH2D, for simplicity, we resort to the same weights used in BM3D [6]. (i.e. $\omega_i = 1$). Further, one can observe that if $\alpha = 1$, we obtain $\omega_i = 1$. Thus, (2) can be interpreted as a generalization of the weights used in [6].

3. RESULTS

We present experimental results from the two variations of the proposed method. Their implementation along with the original noise-free and the filtered images are available online[1]. We used the same algorithm parameters as in BM3D [6]; one can refer to the provided Matlab codes for additional details. The exponent α sets the desired level of sharpening. An illustration of applying the proposed method for a few values of α is given in Figure 2. If not specified otherwise, in our experiments we used $\alpha = 1.5$.

In Figures 4 and 5 one can see the results of applying the proposed method on images degraded by noise with $\sigma = 10, 20$. It can be observed that fine details are well preserved and sharpened while the noise is suppressed. It is worth recalling that the BM3D is particularly effective for objects that are formed by blocks for which there can be found plenty of similar blocks at different spatial locations. Such objects are edges, repeating patterns, and textures. Hence, these (e.g., the repeating circular objects in *Harbour* and the edges of the building in *Pentagon*) are well preserved and subsequently sharpened.

[1]Matlab code available at `www.cs.tut.fi/~foi/GCF-BM3D`.

Figure 3 shows an example of denoising and sharpening a fragment from a real retinal image with added noise of $\sigma = 10, 20$. The resultant filtered images preserve and sharpen most of the image details while efficiently suppress noise. Moreover, the filtered images clearly reveal details which are hard to see even in the original image.

4. DISCUSSION AND CONCLUSIONS

We developed a method for joint denoising and sharpening of grayscale images. It inherits the outstanding denoising potential of the BM3D [6] and in combination with alpha-rooting, it achieves perceptually appealing filtered images where fine details are effectively sharpened.

Let us discuss the two variations of the proposed method. The BM3D-SH3D differs from BM3D-SH2D in that alpha-rooting is applied on the 3D transform spectrum which conveys information of the temporal differences in a group in addition to the spatial one (we call "temporal" the dimension along which blocks are stacked together in a group). In this 3D transform domain, the application of Equation (1) can amplify also differential information along this temporal dimension (and not along the spatial ones). It is noteworthy that in the case of ideal grouping, i.e. having identical blocks in a group, the two approaches produce identical filtered blocks (and differ only in the weighted aggregation of these blocks). The visual results of both approaches (Figures 4, 5, 3, 7) are quite similar by subjective evaluation. Therefore, we abstain from concluding which of them is better in practice. Instead, we

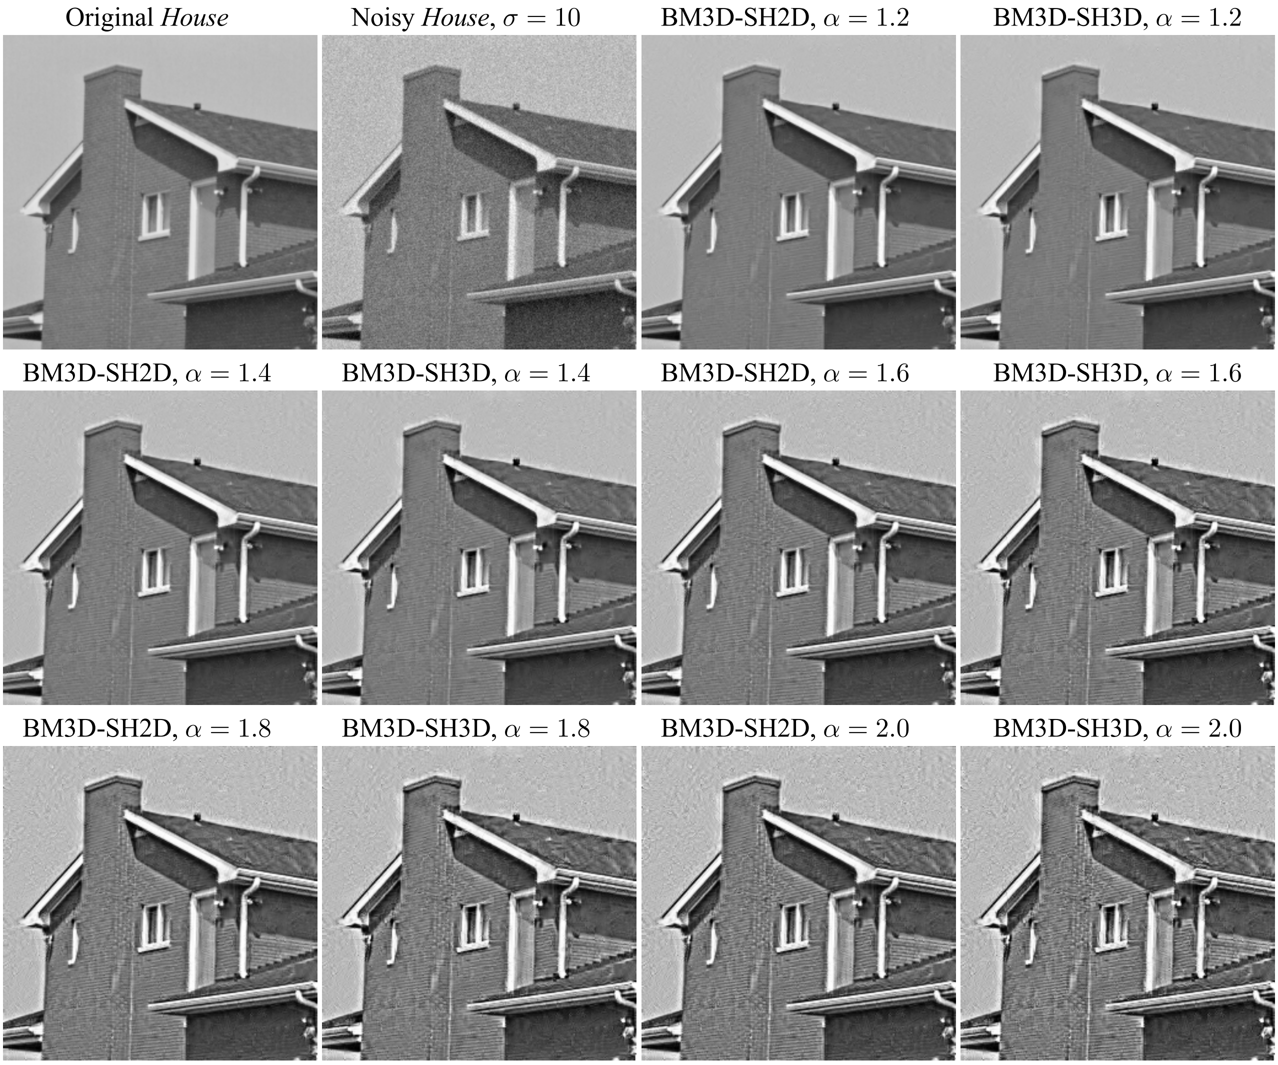

Figure 2. Result of filtering with BM3D-SH2D and BM3D-SH3D for $\alpha = 1.2, 1.4, 1.6, 1.8, 2.0$. Higher values of α correspond to stronger sharpening of the details. The processed image is *House* with noise of standard deviation $\sigma = 10$.

only point to some of the noticeable differences. In the enlarged fragment of *Pentagon* in Figure 5, one can observe that for $\sigma = 10$ the BM3D-SH3D suppresses slightly better the noise in the smooth areas, mainly due to the aggregation weights defined from the total variance estimate in Equation (2). In Figure 7, we see that the BM3D-SH2D can preserve more effectively the oscillating curve. In addition, Figure 6 presents a plot of line 365 of *Pentagon*, which shows that BM3D-SH3D attains slightly better sharpening results than the BM3D-SH2D.

The proposed method can be extended for color data, where the denoising is performed on all three luminance-chrominance color channels as in [8] and the alpha-rooting is performed on all channels or only on the luminance one, depending on the application.

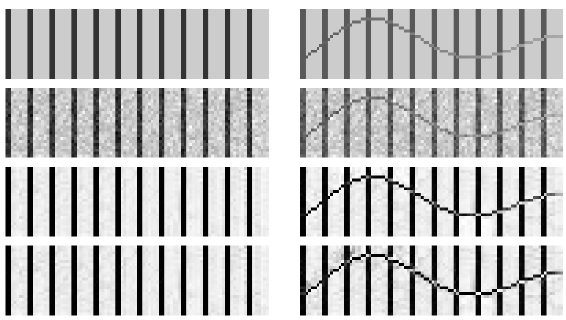

Figure 7. First row: noise-free artificial test images 24×96 pixels; second row: corresponding noisy images with $\sigma = 15$; third row: results of applying BM3D-SH2D; fourth row: results of applying BM3D-SH3D. For the experiments in this figure, both algorithms used $\alpha = 2$ and the sliding step of the BM3D was set to unity.

5. REFERENCES

[1] D.-Y. Tsai, Y. L. M. Sekiya, S. Sakaguchi, and I. Yamada, "A method of medical image enhancement using wavelet analysis," in *Proc. IEEE Int. Conf. Signal Process.*, vol. 1, August 2002, pp. 723–726.

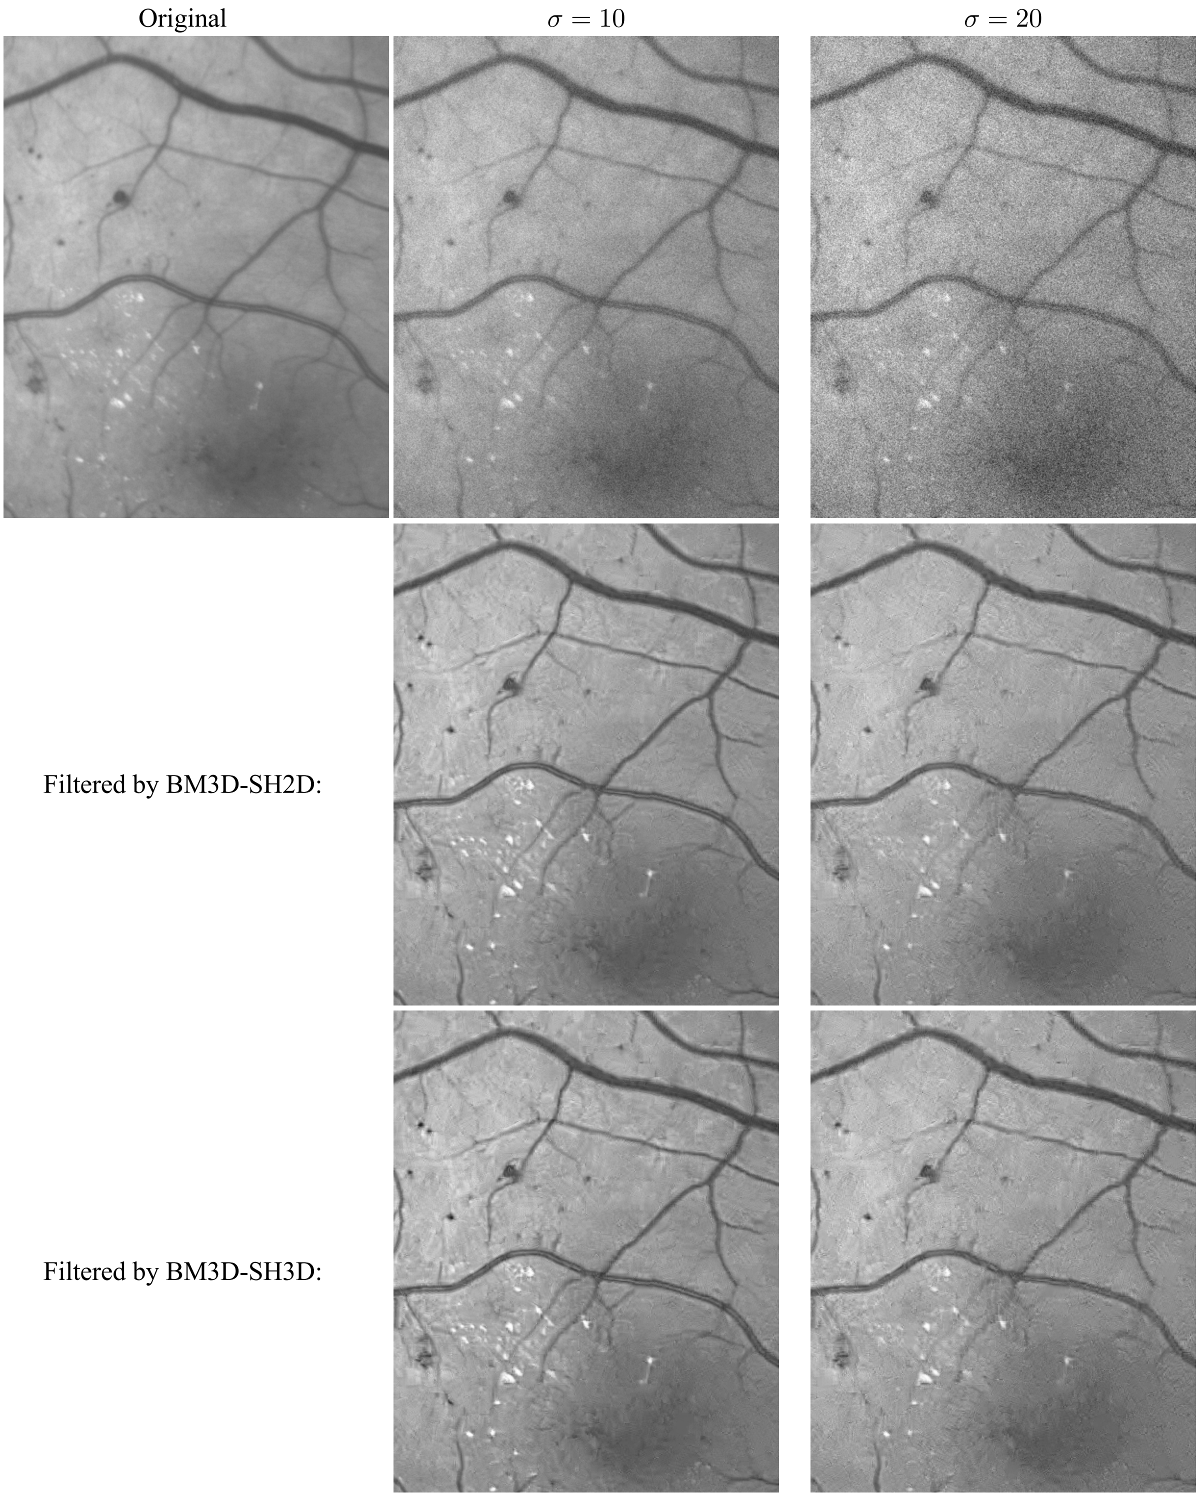

Figure 3. Results of applying the proposed methods on a fragment of a retinal fundus image. The second and third colums correspond to noise with $\sigma = 10$ and $\sigma = 20$, respectively.

[2] S. Hatami, R. Hosseini, M. Kamarei, and H. Ahmadi, "Wavelet based fingerprint image enhancement," in *Proc. IEEE Int. Symp. Circ. Syst. (ISCAS)*, Kobe, Japan, May 2005.

[3] S. Agaian, B. Silver, and K. A. Panetta, "Transform coefficient histogram-based image enhancement algorithms using contrast entropy," *IEEE Trans. Image Process.*, vol. 16, no. 3, pp. 741–758, March 2007.

[4] W. K. Pratt, *Digital Image Processing*. John Wiley and Sons, New York, 1978.

[5] S. Aghagolzadeh and O. K. Ersoy, "Transform image

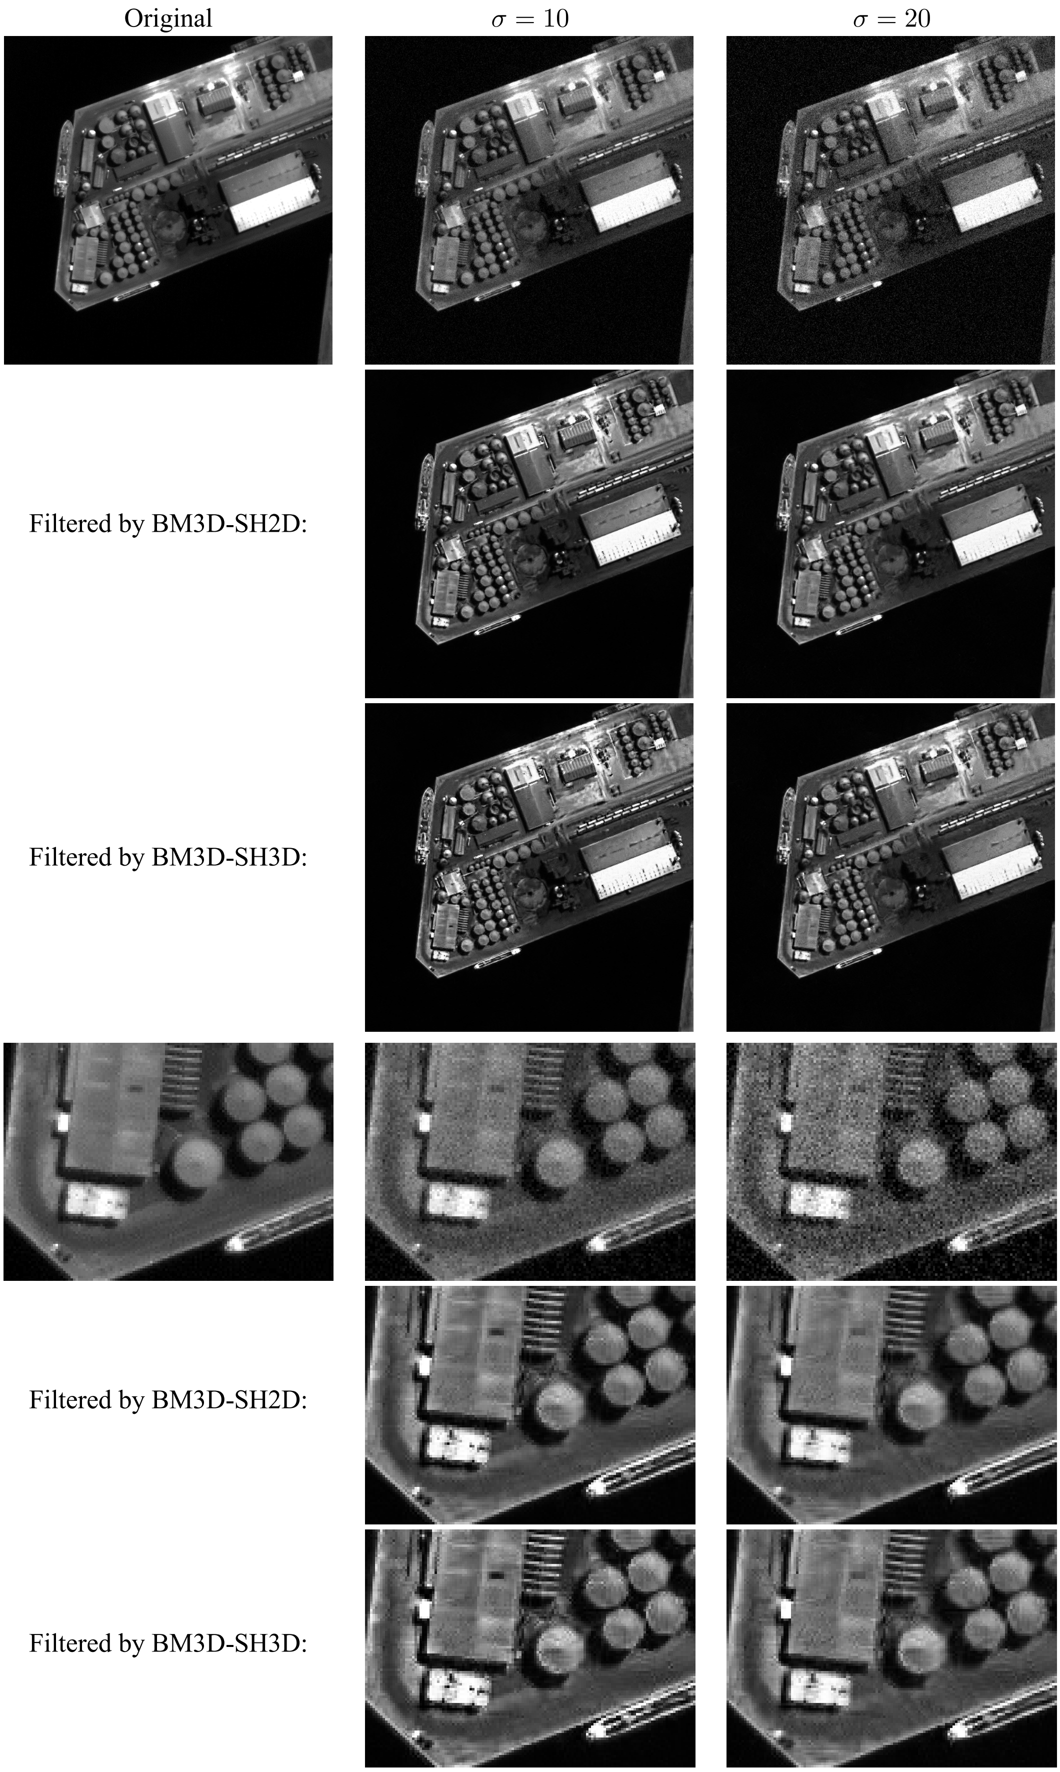

Figure 4. Results of applying the proposed methods on the noisy *Harbour* image (upper half) and a fragment of it (lower half). The second and third colums correspond to noise with $\sigma = 10$ and $\sigma = 20$, respectively.

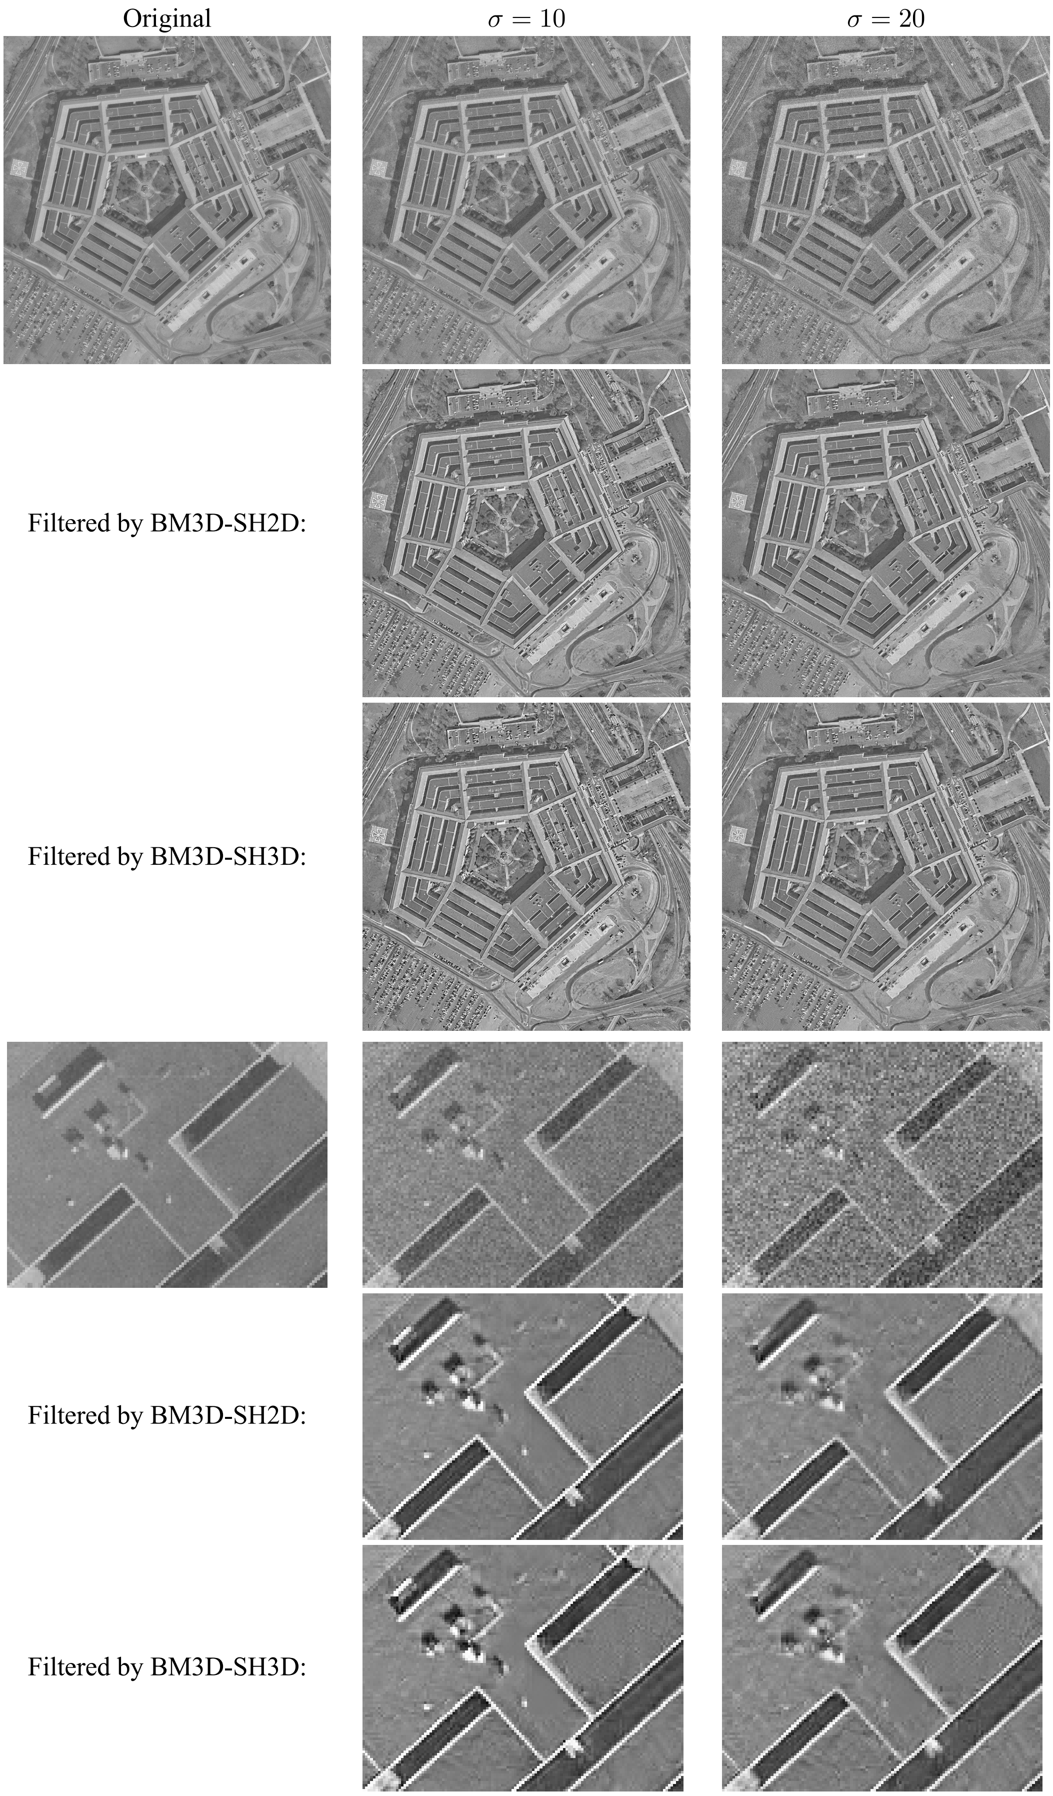

Figure 5. Results of applying the proposed methods on the noisy *Pentagon* image (upper half) and a fragment of it (lower half). The second and third colums correspond to noise with $\sigma = 10$ and $\sigma = 20$, respectively.

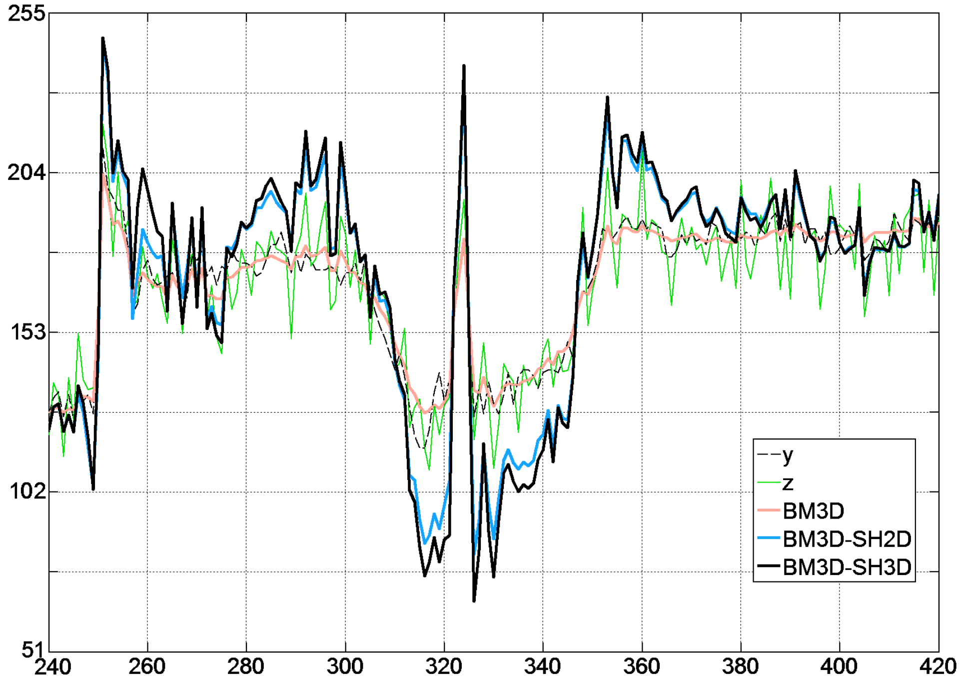

Figure 6. Line 365 of *Pentagon* image for: the noise-free image y, a noisy ($\sigma = 10$) image z, the result of BM3D denoising with no sharpening, the result of joint denoising and sharpening with BM3D-SH2D, and the result of joint denoising and sharpening with BM3D-SH3D.

enhancement," *Opt. Eng.*, vol. 31, no. 3, pp. 614–626, March 1992.

[6] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3D transform-domain collaborative filtering," *IEEE Trans. Image Process.*, vol. 16, no. 8, August 2007.

[7] J. McClellan, "Artifacts in alpha-rooting of images," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process. (ICASSP)*, vol. 5, April 1980, pp. 449–452.

[8] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Color image denoising via sparse 3D collaborative filtering with grouping constraint in luminance-chrominance space," in *Proc. IEEE Int. Conf. Image Process.*, San Antonio, Texas, September 2007.