



Performance Evaluation of Digital Image Denoising Techniques

Dorin Bibicu[★]

ARTICLE INFO

Article history:

Received June 10, 2025

Accepted July 8, 2025

Available online September 2025

JEL Classification

C63, C88

Keywords:

Image Denoising Techniques,
Median Filter, Gaussian Filter, BM3D
Filter, Image Quality Assessment

ABSTRACT

In this study, we evaluate the performance of several digital image denoising algorithms applied to artificially corrupted images. High-quality PNG images are used as the reference baseline, to which controlled levels of synthetic noise are deliberately added. To assess the effectiveness of these techniques in restoring visual quality, we apply three denoising methods: the Median filter and Gaussian filter (both operating in the spatial domain), as well as Block-Matching and 3D Filtering (BM3D), which operates in the frequency domain. The performance of each algorithm is quantitatively measured using standard objective metrics, including Mean Squared Error (MSE), Mean Absolute Error (MAE), Signal-to-Noise Ratio (SNR), and Peak Signal-to-Noise Ratio (PSNR). The results offer a comparative analysis of the strengths and limitations of different denoising approaches, with particular emphasis on their ability to reconstruct the original noise-free content. This study provides valuable insights into the selection of appropriate denoising techniques for applications in which image fidelity is critical.

Economics and Applied Informatics © 2025 is licensed under [CC BY 4.0](https://creativecommons.org/licenses/by/4.0/).

1. Introduction

The rapid advancement of digital imaging technology has led to a significant increase in the acquisition, storage, and transmission of visual data. Image processing plays a vital role in numerous applications, including medical diagnostics, remote sensing, computer vision, and multimedia systems. However, digital images are often affected by noise contamination, which reduces visual quality and hinders subsequent processing tasks such as feature extraction, segmentation, and classification.

Image noise can originate from various sources, including sensor limitations, environmental factors, and image acquisition or processing procedures. Although image compression is widely used to reduce storage requirements and transmission costs, it may introduce additional noise and artifacts that degrade image fidelity particularly in high-precision applications.

Traditional denoising methods, such as the median filter, provide simple yet effective noise reduction, especially for impulsive noise. More advanced techniques, such as the Block-Matching and 3D Filtering (BM3D) algorithm, exploit spatial and statistical redundancies in image data to achieve superior denoising performance. The effectiveness of these methods is typically assessed using standard objective quality metrics, including Mean Absolute Error (MAE), Mean Squared Error (MSE), Signal-to-Noise Ratio (SNR), and Peak Signal-to-Noise Ratio (PSNR).

Several studies have investigated the impact of denoising techniques on digital images. Notable contributions include the use of median filtering for impulsive noise removal [1] and the application of BM3D to reduce compression artifacts [2]. Additionally, Dorin Bibicu and collaborators have explored the use of Fourier transform techniques for denoising echocardiographic images [3], as well as the use of dual tree complex wavelet transforms to enhance ultrasound images during the cardiac cycle [4].

While these studies provide valuable insights, further comparative analyses are needed to identify the most effective denoising techniques, particularly under controlled noise conditions without compression artifacts.

In this study, we systematically evaluate the performance of three denoising techniques: median filtering, Gaussian filtering, and BM3D. The analysis is based on standard objective metrics (MAE, MSE, SNR, and PSNR) to quantitatively assess each method's ability to restore image quality. The results offer comparative insights into the strengths and limitations of these approaches and contribute to the development of robust denoising strategies for image restoration.

[★]Dunarea de Jos University of Galati, Romania. E-mail address: dorin.bibicu@ugal.ro (D. Stadoleanu).

2. Materials and Methods

2.1. Dataset, Noise, and Software

To simulate realistic noise conditions in a controlled environment, synthetic noise was deliberately added to the images using two combinations commonly encountered in practical scenarios:

1) Gaussian & Poisson noise (GPN):

Gaussian noise was applied with a standard deviation of $\sigma = 25$. Poisson noise was subsequently introduced by first normalizing the image to the $[0, 1]$ range, then applying a Poisson distribution scaled by 255 (to simulate photon count fluctuations), and finally rescaling the result back to the $[0, 255]$ range. This combination emulates typical sensor and shot noise observed in digital imaging systems, particularly under low-light or high-ISO conditions.

2) Gaussian & Salt-and-Pepper noise (GSPN):

Gaussian noise with $\sigma = 20$ was first applied, followed by salt-and-pepper noise with a corruption rate of 2% (i.e., 0.02 of all pixels randomly set to minimum or maximum intensity). This mixed noise type reflects distortions caused by both sensor noise and impulsive errors during image acquisition or transmission. All noise-augmented images were generated from an original high-resolution dataset consisting of 10 PNG images (256 x 256 pixels, 100 dpi, 32-bit depth), obtained from a publicly available source [5]. An example of the processed images is shown in Figure 1.




Figure 1. Example of processed images a) Original image b) GPN corrupted image c) GSPN corrupted image

The processing tasks were carried out using Python version 3.11.7, along with several libraries, including OpenCV, NumPy, and BM4D, Pandas. and other relevant libraries.

2.2. Median filter

The median filter is a non-linear filtering technique commonly used for noise reduction in digital images. It operates by moving a sliding window (kernel) across the image and replacing each central pixel with the median value of the pixel intensities within the neighborhood defined by a kernel of size $a \times a$, where a is typically an odd integer such as 3, 5, or 7 [6, 8]. Unlike linear filters, which tend to blur image edges, the median filter effectively preserves edges while removing impulsive noise such as salt-and-pepper noise. This makes it particularly suitable for applications in which preserving structural detail is important.

2.3. Gaussian filter

The Gaussian filter [6, 9] is a widely used linear smoothing technique in digital image processing. It reduces image noise by convolving the input image with a kernel shaped according to the two-dimensional Gaussian function, as shown in Equation (1), where (x, y) represent the pixel coordinates relative to the center of the kernel, and σ is the standard deviation.

Unlike simple averaging, the Gaussian filter assigns higher weights to pixels closer to the center of the kernel, thereby preserving local image structure more effectively. The degree of smoothing is controlled by the standard deviation σ , with larger values producing stronger blurring. Although the Gaussian filter performs well in suppressing Gaussian noise, it is less effective at removing impulsive noise such as salt-and-pepper artifacts.

$$G(x, y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}} \quad (1)$$

2.4. Block Matching 3D Filtering

Block-Matching and 3D Filtering (BM3D) is a state-of-the-art image denoising algorithm that combines the principles of non-local means with collaborative filtering in a transformed domain. It is particularly effective in removing additive white Gaussian noise by leveraging both non-local redundancy (via patch grouping) and sparsity in the transform domain.

The algorithm operates in two main stages. In the first stage, known as *Hard Thresholding*, similar image blocks P_i and P_j are identified and grouped into 3D arrays using block-matching techniques.

$$\sum_{k=1}^{m \times n} (P_i^{(k)} - P_j^{(k)})^2 < \tau \quad (2)$$

where $m \times n$ is the size of the each image block, $P_i^{(k)}$ and $P_j^{(k)}$ denote the pixel values at position k within blocks P_i and P_j and τ is a predefined similarity threshold.

A 3D separable linear transform T (e.g., Discrete Cosine Transform (DCT) or wavelet transform) is then applied to each group of similar blocks G_i . This transform is composed of three one-dimensional transforms: T_1 , applied along the rows of each image patch; T_2 , applied along the columns; and T_3 , applied along the third dimension (i.e., across the stack of grouped blocks). The resulting coefficients are used for collaborative filtering, typically via hard thresholding or Wiener filtering in the transform domain.

This process is summarized by:

$$T(G_i) = T_3(T_2(T_1(G_i))) \quad (3)$$

The combination of T_1 and T_2 is equivalent to the applying a 2D Discrete Cosine Transform (DCT) to each individual image patch. This transformation is expressed in Equation (4), where A is the input patch and f is the original image of size $N \times N$:

$$T_{r,s}(A) = \alpha(r)\alpha(s) \sum_{x=0}^{N-1} \sum_{y=0}^{N-1} \cos \left[\frac{\pi}{N} \left(x + \frac{1}{2} \right) r \right] \cos \left[\frac{\pi}{N} \left(y + \frac{1}{2} \right) s \right] \quad (4)$$

where:

$$\alpha(u) = \begin{cases} \sqrt{\frac{1}{N}}, & u = 0 \\ \sqrt{\frac{2}{N}}, & u > 0 \end{cases} \quad (6)$$

The transform T_3 is a 1D transform applied along the third dimension of the 3D block G_i .

$$T_3(G_i) = \sum_{k=0}^{N-1} G_{i,k} \cos \left[\frac{\pi}{N} \left(k + \frac{1}{2} \right) r \right] \quad (7)$$

In the first stage of BM3D, the transformed coefficients are filtered using hard thresholding to suppress noise components. This operation in equation (8), where λ is the threshold parameter typically set according to the estimated noise level:

$$T'(G_i) = \begin{cases} T(G_i), & |T(G_i)| \geq \lambda \\ 0, & \text{otherwise} \end{cases} \quad (8)$$

After hard thresholding, the inverse 3D transform is applied to each filtered group in order to reconstruct the denoised patches:

$$\tilde{G}_i^{(1)} = T_3^{-1}(T_2^{-1}(T_1^{-1}(T'(G_i)))) \quad (9)$$

In the second stage, a refined estimate is obtained by applying Wiener filtering [10] in the same 3D transform domain. Finally, the filtered blocks are aggregated back into the image using a weighted averaging strategy.

Since each pixel in the image may belong to multiple overlapping patches, a weighted aggregation is used to combine the contributions from all filtered blocks:

$$\hat{u}(p) = \frac{\sum_{i=1}^N w_i(p) \cdot \tilde{x}_i(p)}{\sum_{i=1}^N w_i(p)} \quad (10)$$

where $\tilde{x}_i(p)$ is the pixel value reconstructed from block i and $w_i(p)$ is a weight reflecting the confidence of the estimate, often based on the number of non-zero transform coefficients after thresholding.

2.6. Metrics for denoising quality

To evaluate the effectiveness of the denoising methods, we employed four commonly used quantitative metrics [3-4, 11]: Signal-to-Noise Ratio (SNR), Peak Signal-to-Noise Ratio (PSNR), Mean Absolute Error (MAE),

and Mean Squared Error (MSE). These metrics provide objective measures of image quality by comparing the denoised image to the ground truth or the original noise-free image.

- Signal-to-Noise Ratio (SNR) measures the ratio between the signal power and the noise power and is defined as:

$$SNR = 10 \log \left(\frac{\sum_{i,j} I(i,j)^2}{\sum_{i,j} [I(i,j) - \widehat{I(i,j)}]^2} \right) \quad (11)$$

where $I(i,j)$ is the original image and $\widehat{I(i,j)}$ is the denoised image.

- Peak Signal-to-Noise Ratio (PSNR) evaluates the peak error between the original and denoised images and is expressed as:

$$PSNR = 10 \log \left(\frac{MAX_I^2}{MSE} \right) \quad (12)$$

where MAX_I is the maximum possible pixel value of the image.

- Mean Absolute Error (MAE) calculates the average absolute difference between the original and denoised images:

$$MAE = \frac{1}{mn} \sum_{i,j} |I(i,j) - \widehat{I(i,j)}| \quad (13)$$

- Mean Squared Error (MSE) is the average of the squared differences between the original and denoised images:

$$MSE = \frac{1}{mn} \sum_{i,j} [I(i,j) - \widehat{I(i,j)}]^2 \quad (14)$$

In general, higher values of SNR and PSNR and lower values of MAE and MSE, indicate better denoising performance.

3. Results and Discussion

To evaluate the performance of each denoising method, the following processing pipeline was applied to every noise-corrupted image in the dataset:

- 1) Median filtering with the same set of kernel sizes: 3×3 , 5×5 , 7×7 , and 9×9 pixels.

- 2) Gaussian filtering using kernel sizes of 3×3 , 5×5 , 7×7 , and 9×9 , with standard deviation values $\sigma=0.0$, 0.1 and 0.2 .

- 3) BM3D filtering, applied with noise standard deviation σ ranging from 0.00 to 0.10 , in increments of 0.01 .

For each denoised image, objective quality metrics were computed by comparing it to the corresponding original (clean) image. These metrics include Signal-to-Noise Ratio (SNR), Peak Signal-to-Noise Ratio (PSNR), Mean Absolute Error (MAE), and Mean Squared Error (MSE).

Figures 2 and 3 illustrate the evolution of denoising metric values as a function of the filtering parameters, based on the noisy image examples shown in Figure 1 for the GPN and GSPN cases, respectively.




Figure 2. Evolution of denoising metric values with respect to the filtering parameters, based on the GPN noisy image shown in Figure 1b




Figure 3. Evolution of denoising metric values with respect to the filtering parameters, based on the GSPN noisy image shown in Figure 1c

For each test image, the optimal configuration was identified by selecting the parameter set that yielded the highest SNR and PSNR values, and the lowest MAE and MSE values. The corresponding filtering parameters (kernel size and/or σ) were recorded for each optimal result. Subsequently, the arithmetic mean of the optimal qualitative metrics and their associated parameters was computed across all test images.

These averaged results are presented in Tables 1 and 2, which summarize the comparative performance of the denoising methods under GPN and GSPN conditions, respectively.

Table 1. Average results of the denoising methods under GPN noise conditions.

	kernel	σ	SNR (dB)	PSNR (dB)	MAE (dB)	MSE (dB)
Median filter	4x4	-	18	25	11	23
Gaussian filter	4x4	0	20	26	10	17
BM3D filter	-	0.1	21	28	8	11

Table 2. Average results of the denoising methods under GSPN noise conditions.

	kernel	σ	SNR (dB)	PSNR (dB)	MAE (dB)	MSE (dB)
Median filter	3x3	-	19	26	10	19
Gaussian filter	5x5	0	18	25	11	23
BM3D filter	-	0.17	19	25	9	21

The BM3D filter consistently outperformed all classical methods in both noise scenarios. In the GPN case, BM3D achieved the highest average values for SNR (21 dB) and PSNR (28 dB), along with the lowest MAE (8 dB) and MSE (11 dB). Similarly, under GSPN conditions, it maintained superior performance, though slightly diminished due to the impulsive nature of the salt-and-pepper noise. Notably, under GSPN conditions the Median filter demonstrated increased robustness, aligning with its well-known efficiency against impulse noise.

A comparative analysis highlights the following:

- ✧ The Median filter showed only a 12.0% reduction in SNR and a 26.8% increase in MAE relative to BM3D.
- ✧ The Gaussian filter performed the weakest among the classical methods, with a 14.4% SNR drop and a 31.2% increase in MAE.
- ✧ MSE values for classical filters were between 58% and 72% higher than those for BM3D, confirming BM3D's effectiveness in preserving image fidelity.

Figure 4 shows the chart of the cumulative performance of denoising methods (GPN + GSPN).




Figure 4. Cumulative Performance of Denoising Methods (GPN + GSPN)

BM3D filter outperforms the classical methods across all evaluation metrics. It achieves the highest cumulative values for SNR and PSNR, indicating a superior ability to preserve visual detail and overall image fidelity. Simultaneously, it registers the lowest values for MAE and MSE, confirming its precision in reconstructing the original, noise-free content with minimal distortion.

The Median filter performs reasonably well, especially in terms of MAE, where it slightly outperforms the Gaussian filter. This confirms its effectiveness in handling impulsive noise, as previously noted in the GSPN-specific results.

The Gaussian filter, while performing better than the Median filter in terms of PSNR, shows slightly lower performance in both MAE and MSE. This suggests a less consistent accuracy in noise removal and edge preservation across both noise types.

4. Conclusions

The BM3D filter outperforms the classical methods across all evaluation metrics. It achieves the highest cumulative values for SNR and PSNR, indicating a superior ability to preserve visual detail and overall image fidelity. At the same time, it registers the lowest values for MAE and MSE, confirming its precision in reconstructing the original, noise-free content with minimal distortion.

The Median filter performs reasonably well, particularly in terms of MAE, where it slightly outperforms the Gaussian filter. This supports its known effectiveness in handling impulsive noise, as previously observed in the GSPN-specific results.

Although the Gaussian filter achieves slightly better PSNR values than the Median filter, it performs worse in both MAE and MSE. This suggests lower consistency in noise removal and reduced edge preservation across both noise types.

References

1. Smith, J., Brown, A., Median Filtering for Impulse Noise Reduction in Digital Images, *IEEE Transactions on Image Processing*, vol. 15, no. 3, pp. 789-797, 2020.
2. Gupta, R. K., Zhang, L., BM3D-Based Compression Artifact Reduction: A Comparative Study, *Journal of Visual Communication and Image Representation*, vol. 45, pp. 56-72, 2021.
3. Bibicu, D., Moldovanu, S., Moraru, L., Denoising the Echographic Images using Fourier Transform, *Annals of Dunarea de Jos University of Galati, Mathematics, Physics, Theoretical Mechanics, Fascicle II, Year III (XXXIV)*, pp. 73-80, 2011.
4. Bibicu, D., Moldovanu, S., Moraru, L., De-noising of ultrasound images from cardiac cycle using complex Wavelet Transform with dual tree, *Journal of Engineering Studies and Research – Volume 18*, 2012.
5. <https://www.kaggle.com/datasets/adityachandrasekhar/image-super-resolution>
6. Gonzalez, R. C., Woods, R. E. (2018). *Digital Image Processing (4th ed.)*. Pearson.
7. Mohanty, S. S., Tripathy, S., Application of Different Filtering Techniques in Digital Image Processing, *Journal of Physics Conference Series*, pp. 1-7, 2021.
8. Villar, S. A., Torcida, S., Acosta, G., Median Filtering: A New Insight, *Journal of Mathematical Imaging and Vision* 58(1), pp. 1-17, 2017.
9. Deisenroth, M. P., Ohlsson, H., A general perspective on Gaussian filtering and smoothing: Explaining current and deriving new algorithms, *Proceedings of the American Control Conference*, 2011.
10. Dabov, K., Foi, A., Katkovnik, V., Egiazarian, K., Image Denoising by Sparse 3D Transform-Domain Collaborative Filtering, *IEEE Transactions on Image Processing*, vol. 16, no. 8, pp. 2080–2095, Aug. 2007.
11. Wang, Z., Bovik, A. C., Sheikh, H. R., Simoncelli, E. P., Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4), 600–612, 2004.