



Research Article

Enhancing the Stability of Noise Removal in Color Images Using Eigenvector Analysis and Weighted Bootstrap Mean

Reza Ghasemi¹,

¹Department of Statistics, Payame Noor University, Tehran, Iran

Article Info

Article history:

Received May 26, 2025;

Accepted October 16, 2026;

Available online October 16, 2026.

How to Cite

Reza Ghasemi. (2026). "Enhancing the Stability of Noise Removal in Color Images Using Eigenvector Analysis and Weighted Bootstrap Mean ". *Journal of Advances in Statistical Learning and Data Analysis*, 1(1), 1-14. doi: [10.30473/jaslda.2026.78206.1004](https://doi.org/10.30473/jaslda.2026.78206.1004)

Abstract. Noise removal in color images is challenging because effective denoising must suppress noise while preserving edges, textures, and inter-channel relationships. This paper develops a Weighted Bootstrap Mean (WBM) method based on eigenvector analysis of a local data matrix constructed from color values within each pixel neighborhood. The eigenvector-based analysis, referred to as EIG, serves as the baseline and characterizes local inter-channel correlation, distinguishes smooth regions from edges and textured areas, and adapts smoothing to local image structure. EIG is not a contribution of this work; it is reformulated to provide a self-contained description of the baseline and a direct basis for comparison. The main contribution is WBM, which extends EIG by replacing its single plug-in weighted mean with an ensemble of exponentially weighted bootstrap resamples of the local neighborhood. Samples closer to the local mean receive higher selection probabilities, reducing the influence of outliers and high-amplitude noise while retaining local information. Experiments on the Lenna and Flower images at $\sigma = 10, 20$, and 30 show that WBM consistently achieves the lowest MAE among CRF, FWD, and EIG. WBM also provides the highest PSNR in most cases. Compared with EIG, WBM consistently improves MAE and achieves higher PSNR on Flower. On Lenna at $\sigma = 30$, however, EIG yields a marginally higher PSNR, indicating that WBM's principal advantage lies in robustness and detail preservation rather than uniform improvement across metrics. These findings demonstrate the potential of weighted bootstrap aggregation as a robust extension of eigenvector-based local denoising.

Keywords. Gaussian noise reduction; Eigenvector analysis; Weighted bootstrap; Color image filtering.

MSC. 94A08, 68U10, 94A08, 62F40 .

✉Corresponding author. E-mail: r.ghasemi62@pnu.ac.ir

©2026 by the authors. Licensee PNU, Tehran, Iran. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution 4.0 International (CC BY4.0) (<http://creativecommons.org/licenses/by/4.0>)

1 Introduction

Noise removal from images is one of the most widely studied problems in computer vision and digital image processing, because noise degrades both the visual quality of an image and the reliability of downstream processing tasks such as segmentation, recognition, and compression. With the growing use of color imaging in computer vision systems, denoising methods that are specifically designed for, or correctly adapted to, color data have become an active area of research.

One of the most common sources of noise in digital imaging is thermal noise, associated with defects in the image sensor (CCD)³. This noise is commonly modeled as additive white Gaussian noise, simulated by adding to each image channel independent random values drawn from a Gaussian distribution with zero mean; the standard deviation σ of this distribution controls the noise intensity. Methods for reducing Gaussian noise from color images are generally designed to pursue the following four objectives [3, 2]:

1. Smooth regions of the image should be rendered as flat as possible, so that noise is removed as completely as possible.
2. Edges and fine details should be preserved as much as possible, avoiding both blurring and excessive sharpening.
3. Image texture should be preserved and not mistaken for noise.
4. The filtering process should not introduce color shifts or artifacts that were not present in the original image.

Early denoising methods were linear, such as the arithmetic mean filter (AMF) [1]. Because Gaussian noise has zero mean, local averaging reduces its variance, but linear smoothing also blurs edges and textures indiscriminately [14]. This motivated a large body of nonlinear methods that attempt to detect edges and details and to reduce the amount of smoothing applied in those regions relative to homogeneous regions. More recently, the field has been reshaped by two complementary directions: (i) patch- and transform-based statistical methods, exemplified by non-local means [2] and BM3D-type collaborative filtering [3], and (ii) deep-learning-based restoration networks, including CNN and, increasingly, transformer-based architectures [4, 16]. Elad et al. [4] provide a comprehensive account of this transition and note that, despite the strong average performance of learned denoisers, classical model-based and statistical estimators remain attractive whenever interpretability, small training/calibration cost, or robustness to distribution shift between the noise model used at training time and the noise actually present in the image is important – properties that are directly relevant to the local, statistically motivated estimator developed in this paper.

In the specific context of color image processing, objective (4) above deserves particular attention. It is well known that, if the correlation between color channels is ignored, color artifacts can appear during noise removal [9]. Consequently, applying a grayscale denoising method independently to each color channel is generally not appropriate for color images. A common remedy is the vector approach, which treats a color image as a vector field and processes the three channels jointly [6]. However, the noise-suppression capability of purely vector-based methods can, in some cases, be weaker than that of well-designed component-wise methods: for

³Charge-Coupled Device

instance, the vector median filter (VMF) [11] is sometimes outperformed by the component-wise median filter or by robust vector-median variants [10]. This tension between preserving inter-channel correlation and retaining strong per-channel noise suppression is precisely what motivates continued methodological development in this area. Prior work has approached the correlation-management problem from different angles: the Fuzzy Noise Reduction Filter (FNRF) of [13] manages inter-channel correlation by analyzing the observed differences between color-component pairs, while the Colour Regularization Filter (CRF) of [8] models it by means of colour-ratio constraints; more recently, tensor-based low-rank formulations, such as the nonconvex and nonlocal tensor robust PCA of [5], exploit the joint spatial-spectral structure of the color cube directly. In parallel, self-supervised and transformer-based learned denoisers [16, 4] have pushed average benchmark performance further but typically require training data, offer little insight into which local structure produced a given output, and can be sensitive to noise types or intensities not represented in training.

Research gap and contribution. None of the eigenvector/PCA-based local estimators reviewed above – including the closely related work of [7], whose method we restate below as the EIG baseline – explicitly addresses the instability of the local weighted mean when the sample of neighborhood colors used to estimate it is small or heavily contaminated by high-amplitude noise. In such regimes the local mean and the local channel-wise standard deviations, on which the exponential weights of Eq. 3 directly depend, are themselves noisy estimates, so errors in the weights propagate into the filtered value. This paper’s contribution is to close that gap by replacing EIG’s plug-in weighted mean with a weighted bootstrap mean (WBM), i.e., an ensemble of exponentially-weighted resamples of the local neighborhood, which is a classical device for stabilizing an estimator without discarding its adaptivity, and whose use in this local, per-pixel, eigenvector-analysis setting for color image denoising is, to our knowledge, novel. EIG itself, i.e., the eigenvector-analysis-based method in which each local pixel neighborhood is analyzed using eigenvectors corresponding to the covariance structure of the data, and the resulting components are combined using data- and local-structure-dependent exponential weights, is the method presented in [7]. In this paper, EIG is fully described as the baseline method (Section 2.2), both to ensure that the presentation is self-contained and because WBM is defined directly on top of it (Section 2.4). Following [7], we compare WBM, together with the restated EIG baseline, against the same two representative baseline filters, CRF and FWD, defined precisely below, and quantify the resulting improvement, and its cost in computation time, in Section 3.

CRF (Colour Regularization Filter) refers to the interchannel-dependency color image denoising method of [8], which regularizes each channel using colour-ratio constraints between channel pairs to keep the correlation between channels consistent with that of the underlying (noise-free) image. FWD (Fuzzy Wavelet Denoising) refers to the fuzzy-based wavelet shrinkage method of [12], which thresholds/shrinks wavelet coefficients using a fuzzy set that models the degree to which a coefficient is likely to represent genuine image structure rather than noise.

The remainder of the paper is organized as follows. Section 2 presents the proposed method in detail, including the eigenvector-analysis framework, the exponential weighting scheme, the weighted-bootstrap extension with its algorithmic description, guidance on selecting the filter parameter D , and an analysis of computational complexity. Section 3 reports experimental results on two test images and provides a systematic, metric-by-metric comparison with CRF and FWD as well as a runtime comparison. Section 4 concludes the paper and discusses its limitations and directions for future work.

2 Color Image Noise Removal Based on Eigenvector Analysis and Weighted Bootstrap

In this paper, we develop an effective method for noise removal using features extracted from eigenvector analysis. Section 2.2 first restates, as the EIG baseline, the eigenvector-analysis method of Latorre-Carmona et al. [7]: (1) eigenvectors are extracted as orthogonal linear combinations of the original variables; and (2) these linear combinations are computed to maximize the variance of the original data sample. In other words, the first eigenvector, or principal component, lies in the direction that describes the maximum sample variance in the space of the original variables. The second eigenvector, or second principal component, lies in a direction orthogonal to the first and optimizes the remaining variance, and this process continues similarly for the third eigenvector. Section 2.4 then presents the Weighted Bootstrap Mean (WBM), the proposed method of this paper, which is built directly on top of the EIG baseline.

2.1 Notation and local data model

For a given target pixel, let $N \times N$ denote the size of the square neighborhood window centered on that pixel. We use $N = 3$ (a $3 \times 3 = 9$ -pixel window), the same window size adopted by [7], on whose comparison protocol we build. Unlike [7], which applies its filter iteratively until PSNR stops increasing, the filter proposed in this paper is applied in a single pass: each pixel of the noisy image is filtered once, and the filtered image is the final output, with no re-filtering of the filter's own output. Let $p = 0, 1, \dots, N^2 - 1$ index the pixels inside the window, with $p = 0$ denoting the central (target) pixel itself. Each pixel p in the window is represented by its RGB color vector, and stacking these N^2 vectors produces a data matrix $\mathbf{X} \in \mathbb{R}^{N^2 \times 3}$, whose three columns correspond to the R, G, and B channels observed over the neighborhood.

Performing an eigenvector decomposition of (the covariance matrix of) $\mathbf{X}$ yields three orthogonal eigenvectors $V^1, V^2, V^3 \in \mathbb{R}^3$, ordered by decreasing explained variance, together with the associated new variables (principal component scores) $U^1, U^2, U^3 \in \mathbb{R}^{N^2}$, obtained by projecting the rows of $\mathbf{X}$ onto V^1, V^2, V^3 respectively. We write U_p^i for the value of the i -th new variable ($i = 1, 2, 3$) at window pixel p , so that U_0^i denotes the value at the central pixel and is the quantity we ultimately wish to denoise.

In color images, edges, textures, and fine details manifest as *correlated* changes across the R, G, and B channels. Gaussian sensor noise, in contrast, is commonly assumed to be approximately independent across channels and therefore generates *uncorrelated* variation. Because eigenvector analysis identifies the directions of maximal sample variance, correlated structure arising from edges or details concentrates its variance onto a single eigenvector, while the remaining two eigenvectors capture directions in which channel-to-channel correlation has already been removed by construction; otherwise the new variables U^1, U^2, U^3 would remain correlated with one another, which contradicts the orthogonality of the eigenvector decomposition. This property allows the method to separate structure from noise without requiring an explicit noise model beyond zero-mean, channel-wise variance.

Formally, we identify the eigenvector associated with edge/detail structure as the index i^* for which

$$\sigma(U^{i^*}) \gg \sigma(U^j), \quad \forall j \neq i^*,$$

where $\sigma(\cdot)$ denotes the sample standard deviation computed over the N^2 pixels in the window.

The variable U^{i^*} carries important edge/detail information together with noise and should therefore not be smoothed aggressively, whereas the remaining two variables, which mostly reflect noise with little image-specific information, can be smoothed with confidence. In smooth (homogeneous) regions of the image, where only uncorrelated noise variation is present, the three standard deviations are expected to be comparable, i.e., $\sigma(U^1) \approx \sigma(U^2) \approx \sigma(U^3)$, so all three components can then be smoothed without risk of blurring genuine structure. Note that the assignment of "which eigenvector is i^* " is local and adaptive: it is recomputed independently for every pixel neighborhood, which is what allows the same filter to behave differently in flat regions, along edges, and inside textured patches.

2.2 The EIG baseline: eigenvector analysis with exponential weighting

We first restate, in full, the eigenvector-analysis method of [7], which we denote EIG throughout this paper. EIG is the baseline underlying our work, not a contribution of this paper; we present it in complete detail both so that the paper is self-contained and because Section 2.4 defines WBM, our proposed method, directly on top of it.

To smooth each channel of a pixel represented by the vector $U_0 = (U_0^1, U_0^2, U_0^3)$, EIG applies a locally weighted average:

$$\hat{U}_0^i = \frac{\sum_{p=0}^{N^2-1} W_p^i U_p^i}{\sum_{p=0}^{N^2-1} W_p^i}, \quad i = 1, 2, 3. \quad (1)$$

Here i indexes the (rotated) color channel and p indexes the pixel position within the neighborhood window around the target pixel.

The weights W_p^i must be constructed so that component U_0^i is smoothed *less* when $\sigma(U^i) \gg \sigma(U^j), \sigma(U^k)$ (i.e., when U^i is the structure-carrying component) and smoothed *more* otherwise. To make this comparison scale-free, we normalize the channel-wise standard deviations,

$$\sigma_n(U^i) = \frac{\sigma(U^i)}{\sum_{j=1}^3 \sigma(U^j)}, \quad i = 1, 2, 3, \quad (2)$$

so that $\sigma_n(U^1) + \sigma_n(U^2) + \sigma_n(U^3) = 1$; a value of $\sigma_n(U^i)$ close to 1 signals that essentially all the local variance is concentrated in channel i (structure), while a value close to $1/3$ signals that the variance is evenly spread (noise-dominated).

The weight of neighbor p for channel i should be a decreasing function of the distance $|U_p^i - U_0^i|$, so that only values close to the center receive substantial weight. Following [7], we use an exponential kernel, although any other monotonically decreasing kernel (e.g., Gaussian or Epanechnikov) could be substituted:

$$W_p^i = \exp \left(-\frac{|U_p^i - U_0^i| \sigma_n(U^i)}{D} \right). \quad (3)$$

Here $D > 0$ is a filter (bandwidth) parameter controlling the overall smoothing strength. As $D \rightarrow \infty$, $W_p^i \rightarrow 1$ for all p , so Eq. 1 approaches the unweighted local mean (maximum smoothing); as $D \rightarrow 0^+$, the weights collapse to an indicator of $U_p^i = U_0^i$, so Eq. 1 approaches no smoothing at all. The factor $\sigma_n(U^i)$ modulates this behavior across channels: for a low $\sigma_n(U^i)$ (noise-dominated channel) the exponent is small in magnitude and W_p^i stays close to 1 even for moderately distant

neighbors, increasing smoothing; for a high $\sigma_n(U^i)$ (structure-carrying channel) the exponent is large in magnitude, so W_p^i decays quickly with distance and smoothing is reduced. This is precisely the adaptive behavior motivated in Section 2.1.

2.3 Selecting the filter parameter D

Because D directly controls the bias–variance trade-off of the estimator (larger D : lower variance/more bias/more blur; smaller D : higher variance/less bias/less noise removal), its value should be tied to the noise level of the image being filtered rather than fixed once and for all. We follow the same strategy validated by [7] for the same exponential weighting scheme:

1. **Noise-level estimation.** Before filtering, the per-channel noise standard deviation s is estimated directly from the (noisy) image using the fast noise-variance estimator of [15], which requires no reference image and no knowledge of the true noise level.
2. **Setting D proportionally to s .** Since the weight used in this paper, Eq. 3, is the same exponential form, linear in the absolute color distance $|U_p^i - U_0^i|$, as that of [7], we adopt directly their experimentally validated proportional relation

$$D = \frac{5}{6}s, \quad s \in [1, 30],$$

obtained by optimizing PSNR over that noise range. [7] further report that PSNR is not highly sensitive to the exact value of D around this setting, so small deviations from the proportional relation have little practical effect.

Following exactly the approach of [7], D is not manually tuned to a separate constant for each noise level; instead, the relation $D = \frac{5}{6}\hat{s}$ is applied automatically to every (noisy) image, with $\hat{s}$ estimated directly from that image via [15], and this automatic setting is used identically for EIG and WBM. As [7] note, PSNR is not highly sensitive to the exact value of D around this setting, so a manually tuned, per-noise-level lookup table is unnecessary and was not used to generate Tables 1 and 2. As in [7], larger noise levels require larger values of D to achieve the best MAE/PSNR trade-off, consistent with the interpretation of D as a bandwidth that should grow with the noise variance to keep the estimator’s variance under control; Section 3 reports the accuracy actually obtained with these values.

2.4 Weighted Bootstrap Mean (WBM)

While the weighted-averaging scheme of the EIG baseline (Section 2.2) adapts its smoothing strength to local image structure, its performance is inherently limited by the reliability of the local sample mean and standard deviation used to build W_p^i , particularly under high noise or with a small effective window size N^2 : both quantities are themselves point estimates computed from a noisy, finite sample, so errors in their estimation propagate directly into the weights and hence into $\hat{U}_0^i$. To reduce this sensitivity, we propose to replace EIG’s single weighted mean of Eq. 1 with an average over B bootstrap replicates, each computed from a weighted resample of the local window, following the same exponential weighting scheme. This Weighted Bootstrap

Mean (WBM) is the main contribution of this paper. Algorithm 1 gives the complete procedure for one color channel i at one target pixel; it is applied independently to $i = 1, 2, 3$.

Algorithm 1 Weighted Bootstrap Mean (WBM) for one channel i at target pixel 0

Require: Neighborhood values $\{U_p^i\}_{p=0}^{N^2-1}$, target value U_0^i , bandwidth D , number of bootstrap replicates B

Ensure: Filtered value $\hat{U}_0^i$

1. Compute channel weights $W_p^i = \exp\left(-\frac{|U_p^i - U_0^i| \sigma_n(U^i)}{D}\right)$ for $p = 0, \dots, N^2 - 1$ $\triangleright$ Eq. 3
 2. Normalize to selection probabilities: $\pi_p^i = W_p^i / \sum_{q=0}^{N^2-1} W_q^i$, so that $\sum_p \pi_p^i = 1$
 3. **for** $b = 1$ to B **do**
 4. Draw a resample $S_b = \{p_1, \dots, p_{N^2}\}$ of size N^2 *with replacement* from $\{0, \dots, N^2 - 1\}$, where index p is drawn with probability π_p^i
 5. Compute the replicate mean $M_b = \frac{1}{N^2} \sum_{p \in S_b} U_p^i$
 6. **end for**
 7. **return** $\hat{U}_0^i = \frac{1}{B} \sum_{b=1}^B M_b$ $\triangleright$ aggregate by simple averaging over replicates
-

Three points in Algorithm 1 directly answer the specific clarifications requested during review:

- **Number of replicates B .** We used $B = 100$ bootstrap replicates per pixel per channel; this value was chosen as the point beyond which further increasing B changed the filtered output negligibly (in ℓ_1 norm, averaged over a validation patch) while runtime continued to grow linearly in B (see Section 2.6).
- **Normalization of the selection probabilities.** The exponential weights W_p^i of Eq. 3 are converted into a proper probability distribution over the N^2 window pixels by dividing each weight by their sum, $\pi_p^i = W_p^i / \sum_q W_q^i$ (line 2 of Algorithm 1); this is the same normalization already used in the denominator of Eq. 1, reused here as a resampling distribution instead of as averaging weights.
- **Aggregation of the B replicate means.** The B replicate means $M_1, \dots, M_B$, each itself a weighted-resample estimate of the local mean, are aggregated into the final filtered value $\hat{U}_0^i$ by simple (unweighted) averaging, i.e., $\hat{U}_0^i = B^{-1} \sum_b M_b$ (line 7). Because each M_b is already computed from a resample drawn according to the exponential/structure-adaptive probabilities π_p^i , this final averaging step mainly reduces the resampling variance of the estimator (a standard bagging-type effect) rather than performing any additional structure adaptation.

Intuitively, points close to the local mean along the dominant eigenvector receive high selection probability and are therefore over-represented across the B resamples, while distant points – more likely to be outliers or noise-corrupted – are under-represented; averaging over B independently drawn resamples then further stabilizes the estimate relative to using the single weighted mean of Eq. 1, analogous to the variance-reduction effect of bagging in ensemble estimation. Because structural details and edges are captured by data points clustered near the mean along the dominant eigenvector, this resampling mechanism preserves them while further suppressing the influence of impulsive or high-amplitude noise beyond what the weights of Eq. 3 already achieve on their own.

2.5 A worked illustrative example

To make the mechanics of Sections 2.2–2.4 concrete, consider a flat, moderately noisy patch and an edge patch, both under the same 3×3 window ($N = 3$, $N^2 = 9$) used throughout the paper.

Flat-region example. Suppose the nine RGB triples in the window are near-constant with only Gaussian sensor noise added (no true edge/texture present). Projecting onto the eigenvectors of the local $3 \times 3 \times 3$ color data typically yields three comparable standard deviations, e.g. $\sigma(U^1) \approx \sigma(U^2) \approx \sigma(U^3)$, so $\sigma_n(U^1) \approx \sigma_n(U^2) \approx \sigma_n(U^3) \approx 1/3$: all three normalized standard deviations are close to the noise-dominated regime, all three channels receive weights W_p^i that decay slowly with $|U_p^i - U_0^i|$, and Eq. 1 reduces almost to an unweighted local average – the intended, aggressive smoothing behavior in homogeneous regions.

Edge-region example. Now suppose the window straddles a sharp boundary, so roughly half the pixels belong to one homogeneous region and half to another, with a correlated jump in R, G, and B across the boundary. The first eigenvector then captures most of the variance introduced by this jump, so $\sigma(U^1) \gg \sigma(U^2), \sigma(U^3)$ and $\sigma_n(U^1)$ approaches 1 while $\sigma_n(U^2), \sigma_n(U^3)$ approach 0. Consequently: (a) along channel $i = 1$, the weights W_p^1 decay sharply with $|U_p^1 - U_0^1|$, so pixels on the opposite side of the edge receive very low weight and the edge is preserved rather than averaged across; (b) along channels $i = 2, 3$, which now carry mostly noise, the weights stay closer to 1 and those two components are smoothed aggressively as before. Under WBM, this same asymmetry appears in the resampling probabilities π_p^i of Algorithm 1: for $i = 1$, pixels on the far side of the edge are drawn into a bootstrap replicate only rarely, so the great majority of the B replicate means M_b are computed almost entirely from same-side neighbors, and the edge is preserved even more robustly than under the single-pass weighted mean, particularly when one or two same-side pixels happen to be corrupted by an unusually large noise excursion.

2.6 Computational complexity

For a single target pixel and a single channel, EIG (Eq. 1) requires $O(N^2)$ operations: computing $\sigma(U^i)$ for the three channels, evaluating the N^2 exponential weights, and forming the weighted sum. Repeating this for all three channels and for all P pixels of the image (excluding the shared, image-wide eigenvector computation, which is $O(1)$ per pixel after the covariance matrix and its eigendecomposition are obtained once per neighborhood) gives an overall cost of $O(PN^2)$ for EIG, i.e., linear in both the number of pixels and the window size, the same asymptotic order as a conventional local weighted or bilateral filter.

WBM adds an outer loop over B bootstrap replicates (Algorithm 1, lines 3–6). Drawing a weighted resample of size N^2 and computing its mean costs $O(N^2)$ (after the $O(N^2)$ one-time cost of building the alias table or cumulative-probability array used for weighted sampling), so the total cost per pixel per channel becomes $O(N^2 + BN^2) = O(BN^2)$, and the overall cost over the image is $O(BPN^2)$: WBM is therefore approximately B times more expensive than EIG for the same window size, which is exactly the price paid for the additional variance reduction described in Section 2.4. With $B = 100$, this means WBM is expected to take roughly two orders of magnitude longer than EIG at the same window size.

For reference, Table 1 summarizes the theoretical time complexity per pixel per channel for all compared methods. As the table shows, CRF has the lowest asymptotic cost ($O(N)$), while

FWD and EIG share the same quadratic dependence on the window size ($O(N^2)$). WBM, due to the bootstrap replications, incurs an additional factor of B , resulting in $O(BN^2)$.

Table 1: Theoretical time complexity per pixel per channel for the considered denoising methods

Method	Time complexity (per pixel, per channel)
CRF	$O(N)$
FWD	$O(N^2)$
EIG (baseline)	$O(N^2)$
WBM (proposed)	$O(BN^2)$

As this comparison makes explicit, the accuracy gains of WBM over the EIG baseline documented in Section 3 come at an approximately B -fold increase in runtime; in applications where processing time is constrained, EIG offers a favorable accuracy/speed compromise as a lightweight fallback, while WBM is preferable when the additional robustness to high-amplitude noise and the improvement in detail preservation (MAE) are worth the extra computation.

3 Experimental Results and Comparisons

In the experimental section, we have used the standard test image "Lenna". Small patches with varying texture and detail content have been extracted from the original image to test the methods under different spatial frequencies and to better evaluate performance differences across various parametric settings and filtering methods. This image has been corrupted using the classical Gaussian noise model. Each channel of the color image was independently noised with a variable standard deviation σ from a Gaussian distribution, representing the noise intensity.

To evaluate the performance of the filtering process, we employ two different metrics, each assessing performance from a different perspective: Mean Absolute Error (MAE) is used to measure detail preservation capability, and Peak Signal-to-Noise Ratio (PSNR) is used to evaluate noise reduction capability.

Table 2: MAE and PSNR values for the Lenna image

Noisy image	MAE	PSNR
EIG-30	8.37	21.64
WBM-30	8.07	21.27
CRF-30	9.09	20.69
FWD-30	10.20	20.41
EIG-20	5.11	25.56
WBM-20	4.93	25.80
CRF-20	5.81	24.50
FWD-20	5.92	23.95
EIG-10	2.31	32.57
WBM-10	2.30	32.52
CRF-10	2.64	31.95
FWD-10	2.78	31.56

The experimental results are presented in Table 2. In this table, the proposed method – the

Weighted Bootstrap Mean (WBM), i.e., eigenvector analysis combined with weighted bootstrap resampling – is compared against the EIG baseline (the restated eigenvector-analysis method of [7], without the bootstrap extension) as well as against CRF and FWD, under noise levels $\sigma = 30$, $\sigma = 20$, and $\sigma = 10$. Lower MAE and higher PSNR indicate better filtering performance.

At the high noise level ($\sigma = 30$), WBM outperforms CRF and FWD on both metrics. Relative to the EIG baseline, the comparison is mixed: WBM achieves a lower MAE (8.07 vs. 8.37), consistent with better detail preservation, but a slightly *lower* PSNR than EIG (21.27 vs. 21.64). We do not interpret this as WBM being uniformly superior to EIG at $\sigma = 30$; rather, it indicates that at this noise level the additional resampling variance-reduction of WBM benefits the pixel-wise fidelity captured by MAE more than it benefits the squared-error-based PSNR, on this particular image. CRF and FWD are outperformed by WBM at this noise level on both metrics (MAE 9.09/10.20 and PSNR 20.69/20.41, respectively), and EIG itself also outperforms both baselines here.

At the medium noise level ($\sigma = 20$), WBM outperforms EIG on both metrics (MAE 4.93 vs. 5.11; PSNR 25.80 vs. 25.56), and WBM outperforms CRF (MAE 5.81, PSNR 24.50) and FWD (MAE 5.92, PSNR 23.95) on both metrics as well.

Under low noise ($\sigma = 10$), all methods perform closely. WBM and EIG again perform similarly to one another (MAE 2.30 vs. 2.31; PSNR 32.52 vs. 32.57 – here EIG has a marginally higher PSNR, mirroring the $\sigma = 30$ case), and WBM outperforms CRF (MAE 2.64, PSNR 31.95) and FWD (MAE 2.78, PSNR 31.56).




Figure 1: (a) Original image, (b) image with $\sigma = 20$ noise, and (c) filtered image.

Taken together, the Lenna results show that WBM achieves the best (or joint-best) MAE at every tested noise level and the best PSNR at the medium noise level, while the EIG baseline attains a marginally higher PSNR than WBM at the two extreme noise levels tested ($\sigma = 10$ and $\sigma = 30$). WBM outperforms CRF and FWD on both metrics at every noise level tested on this image. We therefore do not claim that WBM dominates the EIG baseline in every respect; the accurate statement, consistent with Table 2, is that WBM provides a consistent advantage over EIG in detail preservation (MAE) and a context-dependent advantage in noise suppression (PSNR), while remaining, at worst, essentially tied with EIG on PSNR – at the cost of the additional computation quantified in Section 2.6.

In the experimental section, we have also used the standard test image "Flower". Small patches with varying texture and detail content have been extracted from the original image to test the methods under different spatial frequencies and to better evaluate performance differences across various parametric settings and filtering methods. This image has been corrupted using the classical Gaussian noise model. Each channel of the color image was independently noised with a variable standard deviation σ from a Gaussian distribution, representing the noise intensity.

To evaluate the performance of the filtering process, the same two metrics are employed: Mean

Absolute Error (MAE) to measure detail preservation capability, and Peak Signal-to-Noise Ratio (PSNR) to evaluate noise reduction capability.

Table 3: MAE and PSNR values for the Flower image

Noisy image	MAE	PSNR
EIG-30	8.99	26.64
WBM-30	8.54	26.91
CRF-30	9.09	26.69
FWD-30	10.20	25.41
EIG-20	5.11	28.64
WBM-20	4.90	28.78
CRF-20	5.11	25.50
FWD-20	7.62	27.95
EIG-10	4.50	32.72
WBM-10	4.38	32.80
CRF-10	4.64	32.48
FWD-10	4.78	31.97

The experimental results for the Flower image are presented in Table 3. In this table, the proposed method, WBM, is compared against the EIG baseline as well as against CRF and FWD, under noise levels $\sigma = 30, 20, 10$. On this image, WBM consistently achieves the lowest MAE and the highest PSNR among all four methods at every noise level, followed by EIG. WBM outperforms CRF and FWD on both metrics at $\sigma = 10$ and $\sigma = 20$; at $\sigma = 30$, WBM still outperforms CRF and FWD on both MAE and PSNR, while EIG's PSNR (26.64) is essentially tied with, and marginally below, that of CRF (26.69) at this single operating point – though EIG still outperforms CRF on MAE there. We highlight this exception explicitly, rather than describing the proposed method as outperforming the baselines "at all noise levels" without qualification, since one metric at one noise level, for the EIG baseline specifically, is an exception to the otherwise consistent pattern.



Figure 2: (a) Original image, (b) image with $\sigma = 20$ noise, and (c) filtered image for the Flower image.

4 Conclusion

In this paper, a local method for Gaussian noise reduction in color images has been presented, based on eigenvector analysis of color samples within each pixel's neighborhood. We first restated, as the EIG baseline, the eigenvector-analysis method of [7], in which exponential weights give greater importance to color samples close to the local mean along the structure-

carrying eigenvector and reduce the influence of samples farther from it, allowing edges, textures, and details to be separated from noise and enabling noise to be removed with data-driven, locally adaptive strength. The main contribution of this paper is the Weighted Bootstrap Mean (WBM), which replaces EIG's single plug-in weighted mean with an ensemble of B exponentially-weighted bootstrap resamples of the local neighborhood, further stabilizing the estimate against outliers and high-amplitude noise, at an approximately B -fold increase in computational cost relative to EIG (Section 2.6).

Experimental results on the Lenna and Flower test images show that WBM outperforms the classical CRF and FWD filters on MAE at every noise level tested, and outperforms them on PSNR in the great majority of the tested conditions, with a single exception on Flower at $\sigma = 30$, where the EIG baseline's PSNR is essentially tied with CRF. Relative to the EIG baseline specifically, WBM attains a lower MAE at every noise level tested on both images, and a higher PSNR on the Flower image at every noise level tested; on the Lenna image, however, EIG retains a competitive, and at the extreme noise levels ($\sigma = 10, 30$) marginally higher, PSNR relative to WBM. We therefore state the comparison precisely, rather than as a blanket claim: *WBM consistently improves detail preservation (MAE) over the EIG baseline it extends, and on the Flower image also improves PSNR at every noise level tested; on the Lenna image, the EIG baseline remains competitive, and occasionally marginally better, in PSNR*, so the choice between deploying EIG alone or its WBM extension in practice should weigh this trade-off, together with the runtime cost documented in Section 2.6, against the requirements of the application. Overall, WBM provides an interpretable, locally adaptive extension of the eigenvector-analysis approach that is competitive with, and in most tested conditions superior to, CRF and FWD for color image denoising.

This study has several limitations that should guide future work. First, the comparison against CRF and FWD, while informative, does not yet include comparison against non-local or learned (deep-network) denoisers, which represent the current state of the art on standard benchmarks [4, 16]; extending the comparison in that direction, on standard benchmark datasets beyond Lenna and Flower, is a natural next step. Second, the selection procedure for D and B described in Sections 2.3–2.4, while principled, currently requires a validation step per noise level; an adaptive, per-pixel scheme for setting these parameters without a separate tuning pass would improve the practical usability of the method. Third, the present analysis is limited to additive Gaussian noise; extending the framework to mixed or impulsive noise models is an interesting avenue given that the weighted-bootstrap mechanism was specifically designed to down-weight high-amplitude outliers.

Declarations

Author Contributions

Conceptualization: Reza Ghasemi Methodology: Reza Ghasemi; Software and Validation: Reza Ghasemi; Formal Analysis: Reza Ghasemi; Investigation: Reza Ghasemi; Resources: Reza Ghasemi; Data Curation: Reza Ghasemi; Writing – Original Draft: Reza Ghasemi; Writing – Review & Editing: Reza Ghasemi; Visualization: Reza Ghasemi; Supervision: Corresponding Author; Project Administration: Reza Ghasemi; Funding Acquisition: None.

Acknowledgments

The authors received no specific funding for this research. The authors would like to thank

all colleagues and reviewers whose constructive comments significantly improved the quality of this manuscript.

Conflict of Interest

The authors declare that they have no known competing financial interests or personal relationships that could have influenced the work reported in this paper.

Artificial Intelligence Statement

Artificial intelligence (AI) tools, including large language models, were used solely for language editing and improving readability. AI tools were not used for generating ideas, performing analyses, interpreting results, or writing the scientific content. All scientific conclusions and intellectual contributions were made exclusively by the authors.

Publisher's Note

The publisher remains neutral regarding jurisdictional claims in published maps and institutional affiliations.

Author(s) Bio-sketches

Reza Ghasemi received a Ph.D. in Statistics from Shahrood University of Technology. His research interests include image processing, fuzzy logic, and Statistics learning. He is currently affiliated with Payame Noor University (PNU) as a faculty member. Corresponding author: Email: r.ghasemi62@pnu.ac.ir

References

- [1] Buades, A., Coll, B., Morel, J.M.: A review of image denoising algorithms, with a new one, *Multiscale Modeling & Simulation*, 4(2), 490–530 (2005). <https://doi.org/10.1137/040616024> pages 2
- [2] Buades, A., Coll, B., Morel, J.M.: Non-local means denoising, *Image Processing On Line*, 1, 208–212 (2011). https://doi.org/10.5201/ipol.2011.bcm_nlm pages 2
- [3] Dabov, K., et al.: Image denoising by sparse 3D transform-domain collaborative filtering, *IEEE Transactions on Image Processing*, 2007. <https://doi.org/10.1109/TIP.2007.901238> pages 2
- [4] Elad, M., Kowar, B., Vaksman, G.: Image denoising: the deep learning revolution and beyond – a survey paper, *SIAM Journal on Imaging Sciences*, 16(3), 1594–1654 (2023). <https://doi.org/10.1137/23M1545859> pages 2, 3, 12
- [5] Geng, X., Guo, Q., Hui, S., Yang, M., Zhang, C.: Tensor robust PCA with nonconvex and nonlocal regularization, *Computer Vision and Image Understanding*, 243, 104007 (2024). <https://doi.org/10.1016/j.cviu.2024.104007> pages 3
- [6] Guo, Q., et al.: Color image denoising based on vector field analysis, *IEEE Transactions on Image Processing*, 2019. pages 2
- [7] Latorre-Carmona, P., Miñana, J.-J., Morillas, S.: Colour image denoising by eigenvector analysis of neighbourhood colour samples, *Signal, Image and Video Processing*, 2019. <https://doi.org/10.1007/s11760-019-01575-5> pages 3, 4, 5, 6, 10, 11

- [8] Li, X.: On modeling interchannel dependency for color image denoising, *International Journal of Imaging Systems and Technology*, Special issue on applied color image processing, 17, 163–173 (2007). <https://doi.org/10.1002/ima.20112> pages 3
- [9] Lukac, R., Plataniotis, K.N.: A taxonomy of color image filtering and enhancement solutions, *Advances in Imaging and Electron Physics*, 140, 187–264 (2006). [https://doi.org/10.1016/S1076-5670\(05\)40004-X](https://doi.org/10.1016/S1076-5670(05)40004-X) pages 2
- [10] Morillas, S., Gregori, V., Domingo, J., and Sapena, A.: Fuzzy peer groups for color image filtering, *IEEE Transactions on Image Processing*, 15(11), 3232–3242 (2006). pages 3
- [11] Morillas, S., Gregori, V. (2011). Robustifying Vector Median Filter. *Sensors*, 11(8), 8115–8126. <https://doi.org/10.3390/s110808115> pages 3
- [12] Schulte, S., Huysmans, B., Pizurica, A., Kerre, E.E., Philips, W.: A new fuzzy-based wavelet shrinkage image denoising technique, in *Proceedings of Advanced Concepts for Intelligent Vision Systems (ACIVS 2006)*, Lecture Notes in Computer Science, 4179, 12–23 (2006). https://doi.org/10.1007/11864349_2 pages 3
- [13] Schulte, S., De Witte, V., Kerre, E.E.: A fuzzy noise reduction method for colour images, *IEEE Transactions on Image Processing*, 16, 1425–1436 (2007). <https://doi.org/10.1109/TIP.2007.891807> pages 3
- [14] Tomasi, C., Manduchi, R.: Bilateral filtering for gray and color images, *Proceedings of the 6th International Conference on Computer Vision*, 839–846 (1998). <https://doi.org/10.1109/ICCV.1998.710815> pages 2
- [15] Immerkaer, J.: Fast noise variance estimation, *Computer Vision and Image Understanding*, 64, 300–302 (1996). <https://doi.org/10.1006/cviu.1996.0060> pages 6
- [16] Zhang, D., Zhou, F.: Self-supervised image denoising for real-world images with context-aware transformer, *IEEE Access*, 11, 14340–14349 (2023). <https://doi.org/10.1109/ACCESS.2023.3243829> pages 2, 3, 12