Effective Speckle Noise Reduction Using Transformed Bayesian Likelihood with Wiener-Based and Sketch-Based Geometric Priors

Ming-Hsun Mo [1], Pin-Wen Huang [2], and Jian-Jiun Ding [3*]
Graduate Institute of Communication Engineering, National Taiwan University, Taipei, Taiwan
E-mail: r12942060@ntu.edu.tw [1], r10942146@ntu.edu.tw [2], and jjding@ntu.edu.tw [3*]

***Abstract*— An efficient filter for speckle noise reduction in Synthetic Aperture Radar (SAR) images is proposed. Traditional SAR despeckling techniques often suffer from the trade-offs between computational cost, edge preservation, and adaptability to complex scene structures. To address these limitations, the proposed method introduces three improvements: (1) A log-domain distance metric derived from a transformed Bayesian likelihood is proposed. It linearizes the multiplicative noise model and allows fast and stable patch similarity computation. (2) A Wiener-based prior estimation that leverages local statistics to provide a more accurate reflectance prior is adopted. It reduces the estimation bias and enhancing structure preservation. (3) A sketch-based anisotropic weighting mechanism that integrates geometric priors based on gradient orientation and edge strength is proposed. It enables edge-aware aggregation and superior texture retention. Extensive experiments demonstrate that the proposed algorithm consistently outperforms existing approaches in terms of both the peak signal-to-noise ratio (PSNR) and the structural similarity index (SSIM), while also achieving substantial computational efficiency and improving the visual quality.**

I. Introduction

Synthetic Aperture Radar (SAR) images are inherently degraded by multiplicative speckle noise [1], a signal-dependent disturbance that severely impairs visual quality and complicates tasks such as segmentation, classification, and change detection.

The Lee filter [2, 3], employs adaptive local averaging based on neighborhood mean and variance, offering computational simplicity but over-smoothing edges and fine textures due to fixed-window limitations. Subsequently, the Frost filter [4, 5], derived from a minimum mean square error (MMSE) criterion under autoregressive and chi-squared noise assumptions, introduces an exponential adaptive kernel that adjusts smoothing according to local statistics. Despite improved edge preservation, the Frost filter remains limited by isotropic windowing, stationary statistical assumptions, and intensive per-pixel computations.

Addressing limitations of local approaches, BM3D (Block Matching and 3D Filtering) [6] emerged as a state-of-the-art method for additive Gaussian noise, leveraging non-local patch redundancy through collaborative transforms. Its SAR-specific variant, SAR-BM3D [7, 8], modifies patch similarity metrics using speckle-specific likelihood measures and employs wavelet-domain transforms to accommodate multiplicative noise statistics. Although effective, SAR-BM3D still faces considerable computational demands, single-scale operation limitations, and sensitivity to parameter tuning.

The Bayesian Non-Local Means (BNLM) method further integrates statistical speckle priors, applying Bayesian likelihood weighting within non-local neighborhoods. Zhong *et al.* enhanced BNLM [9] through sigma-preselection to reduce bias from noisy priors, but introduced multiple hyperparameters and exhaustive search complexity. BNLM thus remains computationally intensive and sensitive to parameter choices, especially when distinguishing genuine image textures from speckle noise.

A more recent trend explores the use of denoising diffusion probabilistic models (DDPMs) [10]. The DDPM-based despeckling reconstructs clean SAR images by reversing a multi-step Gaussian noise diffusion process conditioned on the speckled input. To mitigate stochastic artifacts, the authors incorporated a cycle-spinning strategy during inference. While the method theoretically leverages the advantages of diffusion models in image restoration, its practical performance is questionable. Despite the complexity of the model, their reported PSNR improvement over baseline deep learning models was minimal, only +0.0014 dB on average, while the per-image inference time reached 60 seconds, significantly limiting its real-time applicability. This performance bottleneck may be partly attributed to the small training dataset (208 synthetic images), which is insufficient for training high-capacity generative models like DDPMs.

Collectively, these existing methods reveal ongoing challenges in balancing denoising performance, adaptivity, and computational efficiency. Our proposed method directly addresses these limitations and apply the following techniques:

979-8-3315-7206-8/25/$31.00

1. Log-domain distance metric based on transformed Bayesian likelihood:

   This transformation not only accelerates patch similarity computation but also retains the statistical integrity of the original multiplicative model, yielding both efficiency and accuracy improvements.
2. Wiener-based prior estimation:

   The Wiener-based prior dynamically adapts to local variance, effectively reducing the estimation bias while better preserving fine structures such as edges and textures.
3. Structure-aware anisotropic weighting:

   By aligning kernel anisotropy with local edge orientation, this mechanism enhances edge fidelity, suppresses cross-structure blurring, and better preserves geometrical features critical in SAR images.

Thereby offering improved despeckling performance suitable for practical SAR applications.

II. PROPOSED METHOD

A. A New Form of BNL Filter

Let $v(x)$ represent the noisy observation at pixel and let $u(x)$ denote the noise-free reflectance (intensity or amplitude) at the corresponding pixel in a SAR image. Additionally, the notations $v(x)$ and $u(x)$ are also used to represent the vectorized patches centered at pixel x with dimensions. The multiplicative speckle noise observed in SAR images can thus be mathematically formulated as: $v(x) = u(x) \cdot s(x)$, where $s(x)$ represents the multiplicative speckle component, typically modeled by a Gamma distribution governed by the number of looks L.

The refined BNLM estimate $\hat{u}(x)$ is computed pixel-wise as the weighted average of all values u(y) in the neighborhood $\Delta(x)$ around pixel x. This is formally given by

$$\hat{u}(x) = \frac{\sum_{y\in\Delta(x)} p\big(v(x)\,|\,u(y)\big)\,p\big(u(y)\big)u(y)}{\sum_{y\in\Delta(x)} p\big(v(x)\,|\,u(y)\big)\,p\big(u(y)\big)}, \quad (1)$$

where $\hat{u}(x)$ is obtain pixelwise as the weighted average of all gray values $u(y)$ in the neighborhood $\Delta(x)$ of x. The term $p(v(x)|u(y))p(u(y))$ acts as the similarity measure between $v(x)$ and $u(y)$.

We illustrate the formulation of the modified BNL filter using an intensity image as our example, noting that the amplitude scenario yields a comparable formulation. Under the assumption of fully developed, statistically independent speckle, the conditional distribution $p(v(x)|u(y))$ introduced in (1) can be rewritten as

$$p\big(v(x)\,|\,u(y)\big) = \prod_{n=1}^{N\times N} p\big(v_n(x)\,|\,u_n(y)\big), \quad (2)$$

where $u_n(y)$ and $v_n(x)$ noting the n^{th} pixel in the corresponding patches, respectively. Assuming that $u_n(y)$ belongs to the set of potential reflectance values corresponding to $v_n(x)$, the conditional probability density function $p(v_n(x)|u_n(y))$ for an intensity SAR image with L-looks can be formulated as shown in

$$p\big(v_n(x)\,|\,u_n(y)\big) = \frac{v_n(x)^{L-1}}{\Gamma(L)}\left(\frac{L}{u_n(y)}\right)^{L}\exp\left(-\frac{Lv_n(x)}{u_n(y)}\right) \quad (3)$$

where $\Gamma(\cdot)$ is the gamma function. Based on (2), (3) can be rewritten as

$$p\big(v(x)\,|\,u(y)\big) \propto \exp\left(\frac{-1}{h^2}\sum_{n=1}^{N\times N}\left(\frac{v_n(x)}{u_n(y)}\right) + \ln\big(u_n(y)\big) - \frac{L-1}{L}\ln\big(v_n(x)\big)\right), \quad (4)$$

and h is the smoothing parameter.

While the original BNLM formulation employs a similarity metric defined in the spatial domain, this formulation incurs substantial computational overhead due to the use of explicit division and logarithmic operations on raw pixel values. In contrast, by transitioning to the logarithmic domain via variable substitution and carefully deriving the transformed likelihood, we obtain a new distance metric that retains the statistical properties of the original model but enables much more efficient computation.

Taking the logarithm and summing across all patch elements yields the linear-domain distance measure based on (4):

$$d_{lin}(x,y) = \sum_{n\in patch}\left[\frac{v_n(x)}{u_n(y)} + \ln u_n(y) - \frac{L-1}{L}\ln v_n(x)\right]. \quad (5)$$

Let $s_n = \log v_n$, $t_n = \log u_n$. Then $v_n = e^{s_n}$, $u_n = e^{t_n}$. The transformation of variables yields:

$$p\big(s_n\,|\,t_n\big) = p\big(v_n = e^{s_n}\,|\,u_n = e^{t_n}\big)\left|\frac{dv_n}{ds_n}\right| = p\big(v_n\,|\,u_n\big)e^{s_n}. \quad (6)$$

Therefore, (4) can be written as:

$$\begin{aligned} p\big(s_n\,|\,t_n\big) &= \frac{L^L}{\Gamma(L)}\frac{e^{s_n(L-1)}}{e^{t_nL}}\exp\left(-L\frac{e^{s_n}}{e^{t_n}}\right)e^{s_n} \\ &\propto \exp\Big[-L\big(e^{s_n-t_n} - (s_n - t_n)\big)\Big]. \end{aligned} \quad (7)$$

Assuming independence across patch pixels, the joint conditional density leads to a log-domain distance based on (7):

$$d_{\log}(x,y) = -\ln \prod_{n\in patch} p\left(s_n(x) \mid t_n(y)\right) = L \sum_{n\in patch} \left[e^{s_n(x)-t_n(y)} - \left(s_n(x) - t_n(y)\right)\right]. \quad (8)$$

This becomes the dedicated distance metric for SAR despeckling under the log-domain formulation.

The result is constructed identically to our BNLM framework:

$$p_{\log}\left(v(x) \mid u(y)\right) \propto \exp\left[-\frac{1}{h^2} L \sum_{n\in patch} \left[e^{s_n(x)-t_n(y)} - \left(s_n(x) - t_n(y)\right)\right]\right]. \quad (9)$$

Let σ^2 denote the variance of speckle noise. As commonly adopted in related work such as that in [11], the smoothing parameter h is modeled to scale linearly with the noise level, i.e., $h = k\sigma$, where $k \approx 2$ is found to be effective for intensity SAR imagery. At the same time, the speckle standard deviation is estimated by our previous noise estimation algorithm.

Previous works such as that in [12] commonly assume a uniform prior distribution over $u(y)$ to simplify the Bayesian estimator, i.e., $p(u(y)) = 1/|\Delta(x)|$, where $|\Delta(x)|$ denotes the number of candidates in the search region. Additionally, earlier BNLM implementations tend to approximate both the unknown reflectance $u(y)$ and the target $u(x)$ by their noisy counterparts $v(y)$ and $v(x)$, respectively. While this substitution simplifies computation, it inevitably introduces estimation bias, especially in high-variance or heterogeneous regions.

In contrast, our method proposes a refined estimation strategy where the prior estimate $u'(y)$ is not derived from local mean reflectance, but rather from the output of a preliminary Wiener filter applied to the speckled image. This choice provides a more stable and informative priori estimation that leverages local variance characteristics to enhance noise suppression while retaining signal structures.

Further, we integrate a sketch-based structure-aware selection mechanism to guide the computation of similarity weights. By analyzing gradient orientation and edge strength, this mechanism designates whether a pixel belongs to a structural region. If so, anisotropic weighting based on geometric alignment is applied, enabling better preservation of directional features such as edges and textures. Critically, rather than assuming a uniform prior $p(u'(y))$, our approach treats $p(u'(y)) \propto G(x,y)$, where $G(x,y)$ is the sketch-based anisotropic Gaussian kernel, thus embedding structural information directly into the prior. The aggregation subset $N(x) \subset \Delta(x)$ is thereby refined not only by intensity similarity but also by structural conformity.

Together, these two contributions (1) replacing the prior mean with Wiener-based prior estimation, and (2) applying sketch-based structural referencing, yield the improved estimator:

$$\hat{u}(x) = \frac{\sum_{y\in N(x)} p\left(v(x) \mid u'(y)\right) G(x,y)\, u'(y)}{\sum_{y\in N(x)} p\left(v(x) \mid u'(y)\right) G(x,y)}, \quad (10)$$

where $p(v(x)|u'(y))$ is the usual log-domain likelihood term, $G(x,y)$ is the sketch-based anisotropic Gaussian kernel (so that structural patches are weighted more heavily), and $N(x)$ is the preselected subset of $\Delta(x)$.

B. A Wiener-based Prior Estimation

In the original BNLM framework, the a priori mean $u'(y)$ is computed via a local σ-filter that aggregates neighboring pixel intensities to approximate the true reflectance before speckle corruption. This local mean plays a critical role in reducing estimation bias by supplying a more accurate representation of $u(y)$ when evaluating the likelihood $p(v(x)|u(y))$. However, σ-filter–based means are themselves susceptible to residual speckle and may not optimally balance noise suppression with detail preservation in heterogeneous SAR regions. In contrast, our proposed method replaces this σ-filter mean with the output of a Wiener filter applied to the raw speckled image. Specifically, by first executing a spatially adaptive Wiener filter, parameterized by local variance estimates, we obtain a smoothed image $u_w(y)$ that better approximates the underlying noise-free reflectance. This Wiener-filtered image serves as the new prior estimate:

$$u'(y) = u_w(y) = Wiener\left(v(y)\right). \quad (11)$$

Because the Wiener filter incorporates both local mean and variance information, its output suppresses speckle more effectively than a simple mean filter while retaining fine structural details. Consequently, when computing nonlocal weights, we substitute $u'(y) = u_w(y)$ into the likelihood term $p(v(x)|u'(y))$, thereby reducing bias and improving denoising accuracy.

C. A Sketch-Based Structural Prior

In the original BNLM framework, the prior distribution over candidate patches $u(y)$ is assumed uniform

$$p(u(y)) = \frac{1}{|\Delta(x)|}, \quad (12)$$

where $|\Delta(x)|$ denotes the total number of pixels in the search window. Under this assumption, all pixels in $\Delta(x)$ are treated equally.

While this simplification reduces the computational burden, it completely ignores any geometric or structural information, leading to over-smoothing along edges or thin linear features. To address these shortcomings, we introduce a sketch-based structural prior that replaces the uniform assumption. Let $G(x)$ denote an anisotropic Gaussian kernel constructed from local gradient orientation $\theta(x)$ and edge strength. We now set $p(u'(y)) \propto G(x,y)$, which biases the prior in favor of pixels y aligned with the same local structure as x. Hence, the combined nonlocal weight becomes

$$w(x,y) = p\big(v(x) \,|\, u'(y)\big) G(x,y), \tag{13}$$

where $u'(y)$ is the Wiener-filtered prior estimate of the unknown reflectance at y.

First, we extract sketch map by computing Sobel gradient magnitude $G_{mag}(x)$ and orientation $\theta(x)$ at each pixel. Apply Otsu's method to threshold $G_{mag}(x)$, yielding a binary edge mask $E(x)$.

Then we define the sketch map $S(x)$ as

$$S(x) = \big[G_{mag}(x) > \tau \big] \vee E(x), \tag{14}$$

where τ is a chosen gradient threshold, ensuring strong structural cues are captured, and v is the pixel-wise product operator.

Secondly, we want one axis of the Gaussian to lie along the structure (the "major" axis, with larger variance), and the other axis across the structure (the "minor" axis, with smaller variance). To do this: Let $\Delta x = x_x - y_x$ and $\Delta y = x_y - y_y$ be the relative offsets from pixel x to candidate y. We can rotate these offsets by the local orientation $\theta = \theta(x)$:

$$f_1 = -\Delta y \sin\theta + \Delta x \cos\theta, \quad f_2 = \Delta y \cos\theta + \Delta x \sin\theta. \tag{15}$$

Choose a minor-axis standard deviation σ_s (e.g., 1–2 pixels) and an elongation factor $\lambda > 1$. Then we define

$$G(x,y) = \exp\left(-\left(\frac{f_1^2}{\sigma_s^{\,2}} + \frac{f_2^{\,2}}{(\lambda\sigma_s)^2} \right) \right). \tag{16}$$

Here, the major axis (lengthened by λ) aligns with the local structure, allowing smoothing along edges, while the minor axis remains narrow to avoid averaging across the edge.

In other words, across-structure weight (term with f_1) penalizes large offsets perpendicular to the edge and along-structure weight (term with f_2) decays more slowly along the edge direction encouraging stretching of the kernel along the line or edge. If candidate y lies along the same structure direction as x, then $f_1 \approx 0$ and $|f_2|$ may be moderate, yielding a relatively large $G(x,y)$. Such pixels receive increased weight in the BNLM average.

Conversely, if y rosses a structural boundary, f_1 becomes large, causing $G(x,y)$ to become exponentially small. These cross-boundary pixels thus contribute negligibly, preserving edge sharpness. By embedding this sketch-based prior into the BNLM algorithm, we ensure that $w(x,y)$ strongly favors pixels that both statistically align in intensity (via the likelihood) and geometrically align in structure (via $G(x,y)$). The combined estimator $u'(x)$ thus achieves a superior balance between despeckling and edge/line retention, outperforming the original uniform-prior BNLM in preserving fine geometric details.

III. Experimental Results

We evaluate our proposed despeckling method on a comprehensive dataset of 208 synthetic SAR airport images sourced from the publicly available RadarS SAR Dataset [13]. Each image measures 512×512 pixels and has been synthetically corrupted with multiplicative speckle noise with $std = 100/255$.

For quantitative assessment, two standard image-quality metrics are employed: Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM). PSNR measures pixel-level differences between denoised and original noise-free images, while SSIM assesses perceived structural fidelity, capturing luminance, contrast, and structural details. Together, these metrics comprehensively characterize despeckling performance in terms of both global and local structural accuracy.

Table I summarizes the averaged PSNR and SSIM scores obtained by different despeckling methods across the entire dataset. Compared to classical local filters (Lee and Frost), transform-domain methods (SAR-BM3D), and original BNLM, our proposed filter consistently achieves the highest PSNR (23.1713 dB) and SSIM (0.7797), indicating significant improvement in both noise suppression and structure preservation.

Moreover, Table II presents the average PSNR improvement ($\Delta PSNR$) relative to the noisy baseline. Our method attains an 8.1183 dB gain, substantially surpassing the gains achieved by traditional Lee (5.78 dB), Frost (5.996 dB), SAR-BM3D (5.964 dB), and original BNLM (7.1756 dB). This clearly demonstrates the effectiveness of our enhanced prior and anisotropic weighting strategies.

Visual comparisons are presented in Figure 1. The Lee filter exhibits noticeable blurring, especially around edges and linear features, while the Frost filter preserves edges slightly better but still lacks fine textural accuracy. SAR-BM3D introduces block-like artifacts, particularly evident along straight lines. Original BNLM maintains edge sharpness but tends to blur delicate structural details due to isotropic weighting.

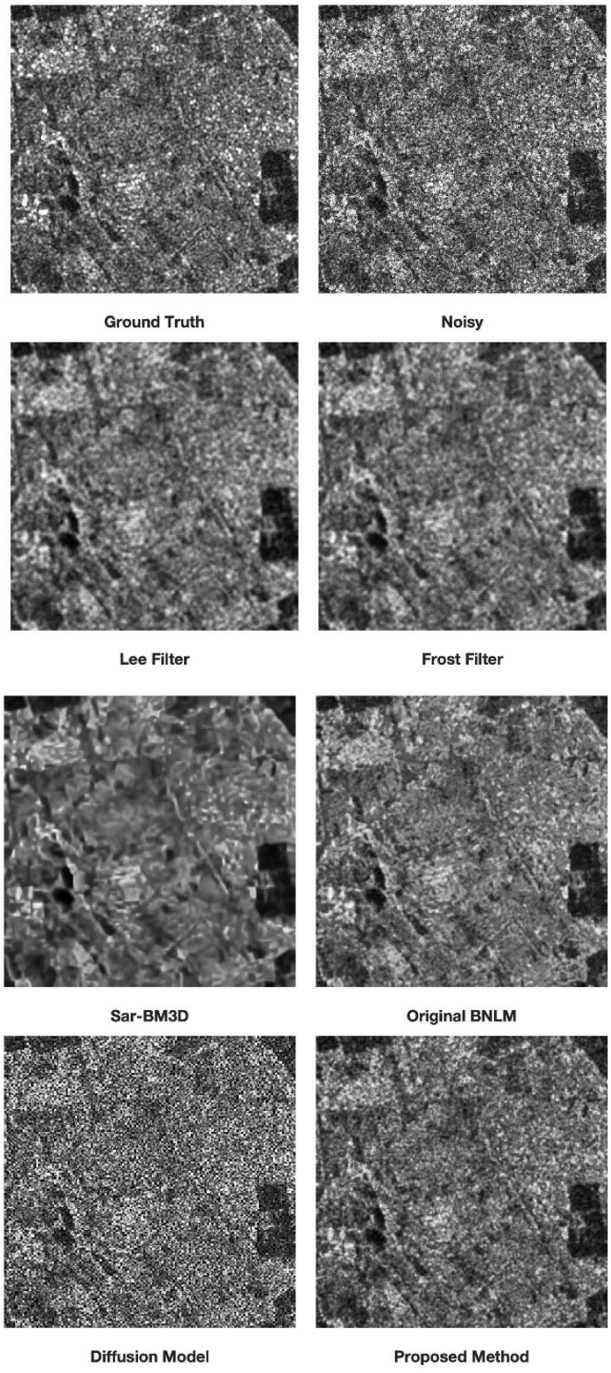

Figure 1 Denoising results with "BEIJING_shahejichang_2_11".

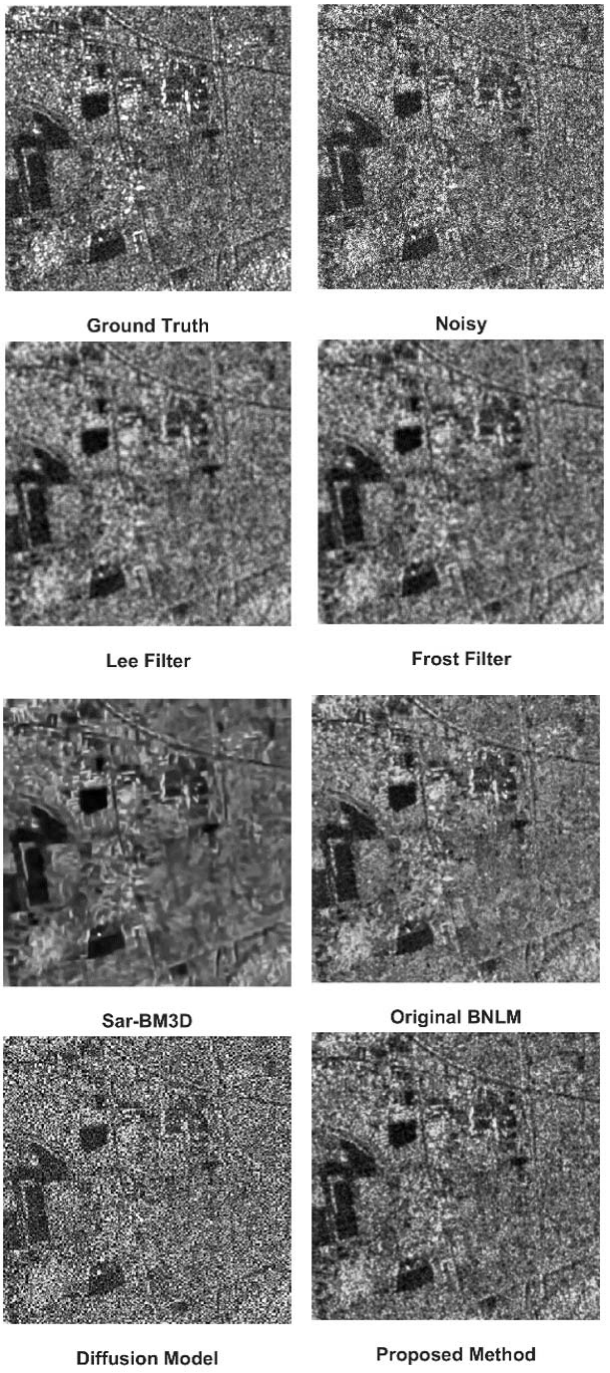

Figure 2 Denoising results of Image "BEIJING_shahejichang_1_4".

Table I Average PSNR and SSIM across All Test Images

Method	PSNR (dB)	SSIM
Noisy Input	15.053	0.3437
Lee Filter	20.833	0.5366
Frost Filter	21.049	0.5593
SAR-BM3D	21.017	0.5297
Original BNLM	22.2286	0.6987
Diffusion Model	15.0546	0.3511
Proposed	**23.1713**	**0.7797**

Table II Average PSNR Gain over Noisy Inputs

Method	Δ_{PSNR} (dB)
Lee Filter	+5.78
Frost Filter	+5.996
SAR-BM3D	+5.964
Original BNLM	+7.1756
Diffusion Model	+0.0016
Proposed	**+8.1183**

In contrast, the proposed method significantly improves visual quality, retaining sharp edges and subtle textures. The sketch-based anisotropic kernel effectively guides the nonlocal weighting to align with structural geometry, visibly enhancing the clarity of runways and building contours. Annotated SSIM values in the visual examples reinforce our method's superior ability to preserve structural coherence.

IV. Conclusions

In this study, we presented an enhanced Bayesian Nonlocal Means (BNLM) filter specifically designed for effective speckle noise reduction in SAR imagery. The proposed method integrates three main innovations: (1) a log-domain distance metric. In addition to computational efficiency, the log-domain metric linearizes the multiplicative noise model, ensuring numerical stability, enhancing robustness to extreme intensity variations, and preserving the statistical characteristics of the original SAR image data. (2) a robust Wiener-based prior estimation, leveraging local mean and variance information, the Wiener-based prior effectively reduces estimation bias and significantly enhances structural fidelity. This prior accurately reflects local reflectivity variations, thereby preserving critical edge and texture details. and (3) a structure-aware anisotropic weighting mechanism guided by geometric sketch maps. This anisotropic weighting mechanism facilitates edge-aware smoothing. It preserves structural integrity, prevents blurring across edges, and effectively maintains geometric features, ensuring high-quality image restoration results. Our extensive experiments demonstrate that the method significantly surpasses conventional despeckling algorithms in terms of the PSNR, the SSIM, and visual quality. Additionally, our approach achieves considerable computational efficiency, making it suitable for practical large-scale and near-real-time SAR image processing applications. Future research directions include further computational optimizations, GPU-based implementations, and validations across more diverse SAR imaging scenarios.

References

[1] H. Maître, *Processing of Synthetic Aperture Radar (SAR) Images*, John Wiley & Sons, 2013.

[2] J. Jaybhay and R. Shastri. "A study of speckle noise reduction filters," *Signal & Image Processing: An International Journal*, vol. 6, issue 3, pp. 71-80, 2015.

[3] A. S. Yommy, R. Liu, and S. Wu, "SAR image despeckling using refined Lee filter," in *IEEE Int. Conf. Intelligent Human-Machine Systems and Cybernetics*, vol. 2, pp. 260-265, 2015.

[4] J. Li, W. Yu, Y. Wang, Z. Wang, J. Xiao, Z. Yu, and D. Zhang, "Guidance-aided triple-adaptive frost filter for speckle suppression in the synthetic aperture radar image," *Remote Sensing*, vol. 15, issue 3, article 551, 2023.

[5] S. Banerjee, S. S. Chaudhuri, R. Mehra, and A. Misra, "A comprehensive survey on frost filter and its proposed variants," in *Int. Conf. Communication and Electronics Systems*, pp. 109-114, 2020.

[6] K. Dabov, A. Foi, V. Katkovnik and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE Trans. Image Processing*, vol. 16, issue 8, pp. 2080-2095, Aug. 2007.

[7] D. Cozzolino, S. Parrilli, G. Scarpa, G. Poggi, and L. Verdoliva, "Fast adaptive nonlocal SAR despeckling," *IEEE Geoscience and Remote Sensing Letters*, vol. 11, issue 2, pp. 524-528, 2014.

[8] F. Sica, D. Cozzolino, X. X. Zhu, L. Verdoliva, and G. Poggi, "InSAR-BM3D: A nonlocal filter for SAR interferometric phase restoration," *IEEE Trans, Geoscience and Remote Sensing*, vol. 56, issue 6, pp. 3456-3467, 2018.

[9] H. Zhong, Y. Li, and L. Jiao, "SAR image despeckling using Bayesian nonlocal means filter with sigma preselection," *IEEE Geoscience and Remote Sensing Letters*, vol. 8, issue 4, pp. 809-813, July 2011.

[10] M. V. Perera, N. G. Nair, W. G. C. Bandara, and V. M. Patel, "SAR despeckling using a denoising diffusion probabilistic model," *IEEE Geoscience and Remote Sensing Letters*, vol. 20, pp. 1-5, 2023.

[11] J. A. Buades, B. Coll, and J. M. Morel, "A review of image denoising algorithms, with a new one," *Multiscale Modeling & Simulation*, vol. 4, issue 2, pp. 490–530, 2005.

[12] C. Kervrann, J. Boulanger, and P. Coupe, "Bayesian non-local means filter, image redundancy and adaptive dictionaries for noise removal," in *Int. Conf. Scale Space and Variational Methods in Computer Vision*, pp. 520–532, 2007.

[13] D. Wang, F. Zhang, F. Ma, W. Hu, Y. Tang, and Y. Zhou, "A benchmark sentinel-1 SAR dataset for airport detection," *IEEE J. Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 15, pp. 6671-6686, 2022.