IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

Advanced Noise Suppression Techniques for Radar Signals Using ABS, Wiener Filtering, and CNN

[1]Palla Mary Snehitha, [2]Hari Krishna Kanneboina, [3]Dr. A. M. Prasad
[1]M. Tech Student, [2]Assistant Professor, [3]Professor

[1]Department of Electronics and Communication Engineering

[1]University College of Engineering (JNTU) Kakinada, India

Abstract: This work compares three techniques for reducing noise in MST radar I/Q data: Wiener filtering, Adaptive Bayesian Shrinkage, and a refined CNN that uses multi-head attention. MST radar echoes often weaken as they pass through atmospheric layers, where turbulence, refractive-index changes, and multipath introduce phase distortions, amplitude fluctuations, and broadband noise. These effects lower Doppler accuracy and affect the retrieval of atmospheric parameters. To test each method, real MST radar measurements and synthetic noisy signals were used. The CNN was trained with a combined loss that includes time-domain reconstruction and a spectral-consistency term, helping it retains Doppler patterns and stable instantaneous phase. Evaluation relied on RMSE, MAE, ΔSNR, spectral coherence, cross-correlation, and phase error. Across all these measures, the CNN provided stronger denoising and better Doppler preservation than the other methods. Its balance between noise removal and signal fidelity makes it suitable for MST radar analysis.

Keywords: MST radar, Denoising methods, CNN, Bayesian shrinkage, Doppler retrieval

1. INTRODUCTION

MST radars rely on coherent I/Q samples to estimate Doppler shifts associated with atmospheric motions, turbulence, and layered scattering processes. As radar pulses travel through regions with changing temperature, humidity, and refractive index, the echo undergoes turbulence-driven fluctuations, weak multipath, and small-scale molecular interactions. These processes introduce broadband noise, disturb phase stability, and affect amplitude consistency. Since accurate Doppler retrieval depends heavily on stable phase information and a sufficient signal-to-noise ratio, even small distortions can degrade wind-profile estimates and limit the reliability of higher-order atmospheric analyses.

Because of these challenges, noise suppression is an essential part of MST radar signal processing. Classical methods such as Wiener filtering try to balance noise reduction and signal preservation by assuming stationarity and Gaussian noise behaviour. This works reasonably well when conditions are simple, but real atmospheric echoes rarely follow these assumptions. Rapid variations in turbulence intensity, refractive-index gradients, and backscatter strength often make the noise highly non-stationary. As a result, Wiener filtering can smooth out weak Doppler components or fail to maintain phase continuity, both of which are crucial for retrieving accurate velocities.

Wavelet-domain approaches like Adaptive Bayesian Shrinkage (ABS) offer better adaptability by examining the signal at multiple resolutions and applying data-driven shrinkage based on local uncertainty. ABS can handle moderate variability in noise characteristics and often performs better than fixed linear filters. Even so, it can struggle when the distortions change quickly or when the fine spectral details in the Doppler structure are too weak to be reliably separated from noise. Maintaining instantaneous-phase stability remains difficult, especially under strong turbulent scattering.

Deep learning has introduced a more flexible path forward. Convolutional Neural Networks can learn nonlinear relationships between noisy and clean radar returns directly from data, without relying on rigid statistical assumptions. Their local feature extraction abilities help them capture structured Doppler signatures that traditional filters might weaken. Adding attention mechanisms strengthens this further by enabling the network to focus on long-range dependencies and amplify subtle Doppler components that often get masked by turbulence-induced noise. This capability aligns well with the needs of MST radar analysis, where key spectral features are delicate yet scientifically important.

This work compares three denoising strategies Wiener filtering, Adaptive Bayesian Shrinkage, and an attention-enhanced CNN to understand how well each method preserves phase stability and Doppler structure under non-stationary atmospheric noise. The evaluation uses both real MST radar data and synthetically corrupted signals to ensure a fair assessment. By analysing RMSE, MAE, ΔSNR, spectral coherence, cross-correlation, and phase error, the study highlights the trade-offs between classical filtering, statistical modelling, and deep learning. The results demonstrate how the attention-augmented CNN offers a stronger balance between noise suppression and signal fidelity, making it a promising choice for reliable MST radar Doppler retrieval.

The paper is organized into five sections. Section 2 reviews the existing work related to MST radar denoising. Section 3 explains the methodologies used in this work. Section 4 presents the experiments and discusses the results. Section 5 concludes the work and highlights the key findings.

2. RELATED STUDY

Padhy et al. [1] introduced an adaptive Bayesian shrinkage model to enhance VHF-MST radar signals and retrieve reliable winds at higher altitudes. Their method combines parametric and non-parametric Bayesian mixture modelling, wavelet-based shrinkage, and posterior estimation to strengthen weak atmospheric signals between 22 and 30 km, which are typically difficult to detect using FFT-based or non-Bayesian denoising approaches. The study outlines a detailed framework involving segmentation, collated datasets, wavelet transforms, and hyperparameter-driven estimation. Although the technique significantly improves altitude coverage, it still struggles above 28.5 km, where correlation with GPS radiosonde winds drops sharply. The results show that the proposed method achieves better signal enhancement, improved SNR, lower noise power, and more accurate wind retrieval compared to TFDR, TSURE, TUT, TYW, and FFT techniques, extending the usable radar range from the conventional 21 km up to about 28.5 km.

Ravindrababu et al. [2] developed an MST radar–based approach to detect tropopause height by analysing changes in signal-to-noise ratio, spectral width, and echo strength. They proposed a detection strategy that identifies sharp stability gradients and validated it using simultaneous radiosonde observations. The methodology highlights how radar measurements can capture fine atmospheric transitions with high temporal resolution. Although the system occasionally struggles in low-SNR conditions or during periods of weak backscatter, the comparison with radiosonde data shows strong consistency in estimated tropopause

levels. Their work demonstrates that the developed method provides a dependable way to track tropopause variability, with only minor limitations when atmospheric structures are faint or the radar signal weak, making it suitable for routine monitoring applications.

Das et al. [3] developed a long-term analysis to understand stratosphere–troposphere exchange by combining MST radar observations with Aura MLS satellite measurements over the tropical site at Gadanki. They proposed an approach that uses radar-derived wind, turbulence, and echo characteristics alongside MLS temperature and ozone profiles to examine exchange processes across different seasons and atmospheric conditions. Their method highlights how vertical transport patterns, tropopause variability, and mixing mechanisms evolve over time. A limitation is that radar coverage weakens at higher altitudes, while MLS offers coarser vertical resolution, creating occasional mismatches during detailed layer-by-layer comparisons. Even so, the results reveal consistent signatures of exchange events, showing that the combined dataset provides a robust understanding of long-term atmospheric coupling and transport behaviour in the upper troposphere and lower stratosphere.

Rao et al. [4] developed an atmospheric radar signal processing method based on principal component analysis to improve the extraction of meaningful information from MST radar returns. They proposed using PCA to separate coherent atmospheric signals from noise by decomposing radar echo data into dominant and minor components. The technique identifies principal modes representing true atmospheric structures, allowing cleaner retrieval of wind and turbulence parameters. A limitation is that PCA performance depends heavily on the choice of components and may misclassify weak signals during highly variable conditions. Even with this constraint, their results show clearer spectra, reduced noise levels, and improved estimation of radar moments, demonstrating that PCA can effectively enhance MST radar signal interpretation and overall data quality.

Thatiparthi et al. [5] developed a wavelet-based denoising technique for MST radar signal processing to improve detection of weak atmospheric echoes. They proposed applying discrete wavelet transform to decompose radar returns into multi-resolution components, followed by thresholding to suppress noise while preserving important spectral features. The approach is particularly effective in conditions where traditional FFT-based methods struggle due to low SNR. However, the method can sometimes over smooth signals or remove fine structures when threshold selection is not optimal. Despite these limitations, the results show enhanced clarity in Doppler spectra, improved moment estimation, and better discrimination between atmospheric layers. Their work demonstrates the usefulness of wavelet shrinkage as a practical tool for improving radar signal quality.

Hocking et al. [6] developed a comprehensive review of radar-based techniques for estimating turbulent energy dissipation rates in the middle atmosphere, summarizing theoretical foundations and practical applications. He proposed analytical links between Doppler spectral width, atmospheric turbulence, and dissipation, explaining how radar measurements can be interpreted to derive key turbulence parameters. The review highlights the strengths of VHF radar systems for continuous monitoring but also notes limitations such as contamination from shear broadening, beam effects, and instrumental noise, which can bias dissipation estimates. Even with these challenges, the work establishes radar as a reliable tool for studying turbulence when supported by careful calibration and correction methods. Hocking's review remains a foundational reference for turbulence quantification using MST radar.

Chipman et al. [7] developed an adaptive Bayesian wavelet shrinkage method that models wavelet coefficients using hierarchical priors to achieve data-driven smoothing. They proposed a framework that adapts shrinkage levels to local signal characteristics, improving noise suppression without over smoothing important structure. Although computationally more demanding than classical shrinkage rules, the method delivers cleaner reconstructions and superior performance in heterogeneous noise conditions. Their work established a strong foundation for Bayesian wavelet denoising, influencing later signal-processing techniques, especially in radar and remote-sensing applications.

Abramovich et al. [8] developed a Bayesian wavelet thresholding approach that estimates optimal shrinkage levels using prior distributions on wavelet coefficients. They proposed a method that balances noise reduction and signal preservation by integrating posterior probabilities into threshold selection. While the approach requires careful prior specification and increased computation, it consistently outperforms universal and SURE-based thresholds in complex noise settings. Their results show improved reconstruction accuracy, making Bayesian thresholding a powerful choice for denoising tasks in scientific and engineering signal analysis.

Rao et al. [9] developed the system description and initial performance assessment of the advanced Indian MST radar, outlining major upgrades in hardware, beamforming, and continuous observation capability. They proposed improvements that enhance vertical coverage, spectral resolution, and operational stability, enabling detailed studies of atmospheric dynamics. Although some limitations persist at higher altitudes due to reduced echo strength, their sample observations demonstrate significant advances in capturing wind, turbulence, and structural features. The work provides a crucial foundation for modern atmospheric research over the Indian region.

Jain et al. [10] developed one of the earliest studies on aspect sensitivity of VHF radar backscatter using the Indian MST radar. They proposed analysing variations in echo power with beam tilt to understand anisotropic scattering from atmospheric irregularities. Their observations reveal strong aspect sensitivity near zenith, indicating layered structures and turbulence effects. A limitation is the preliminary nature of the dataset, offering limited temporal coverage. Still, the findings highlight key scattering characteristics essential for interpreting MST radar signals and improving future radar-based atmospheric studies.

Nastrom et al. [11] developed an analytical study on Doppler radar spectral-width broadening caused by beamwidth and wind shear. He proposed a framework showing how shear-induced velocity gradients and finite antenna beam angles artificially widen spectral returns, complicating turbulence interpretation. A limitation is that the model relies on simplified atmospheric structures, which may not capture complex real-world variability. Even so, the results offer clear guidance for correcting spectral-width estimates, improving the accuracy of turbulence and wind measurements in radar-based atmospheric studies.

Ghosh et al. [12] developed a method to determine atmospheric turbulence parameters by combining simultaneous MST radar and radiosonde measurements at Gadanki. They proposed using radar-derived Doppler spectral width alongside balloon-measured temperature and wind profiles to estimate turbulent dissipation rates and related stability metrics. Although uncertainties arise from shear contamination and limited radiosonde launches, the combined approach provides consistent turbulence estimates. Their work highlights the value of coordinated radar–radiosonde observations for characterizing fine-scale atmospheric dynamics in the tropical middle atmosphere.

Fukao et al. [13] developed a comprehensive three-year observational study on seasonal variability of vertical eddy diffusivity using the MU radar. They proposed analysing long-term radar-derived turbulence parameters to understand how mixing processes evolve across seasons in the middle atmosphere. While radar sensitivity decreases at higher altitudes, the extended dataset reveals clear and repeatable seasonal cycles in diffusivity, linked to large-scale circulation and wave activity. Their findings remain influential for understanding vertical transport, atmospheric coupling, and turbulence climatology in the mesosphere and lower thermosphere.

Padhy et al. [14] developed an adaptive-Bayesian DStoch technique to retrieve winds from MST radar at higher altitudes. They proposed a stochastic Bayesian model that enhances weak atmospheric signals using posterior-based shrinkage, improving wind detection beyond the limits of FFT-based methods. Although performance declines above the highest detectable ranges, the technique offers better SNR, cleaner spectra, and more accurate wind retrieval compared with non-Bayesian approaches. Their results demonstrate significant improvements in altitude coverage, enabling more reliable wind estimation in the upper troposphere and lower stratosphere.

3. METHODOLOGY

3.1 RADAR SIGNAL MODELLING

Radio signal modelling provides the mathematical foundation for understanding how MST radars capture atmospheric echoes through their complex in-phase (I) and quadrature (Q) samples. These samples encode Doppler shifts, amplitude variations, and phase fluctuations produced by atmospheric motion, turbulence, and refractive-index changes. An accurate model is essential for separating the true atmospheric echo from noise and propagation-induced distortions.

The received complex baseband signal can be written as:

$$s[n] = x[n] + \eta[n], \tag{1}$$

where $x[n]$ denotes the true atmospheric echo and $\eta[n]$ represents noise and interference accumulated during propagation.

When atmospheric scatterers exhibit motion, the clean signal carries Doppler modulation and is described by:

$$x[n] = A[n]\, e^{j(2\pi f_d nT + \phi[n])}, \tag{2}$$

where

- $A[n]$ is the slowly varying amplitude envelope,
- f_d is the Doppler frequency,
- T is the sampling interval,
- $\phi[n]$ represents additional phase variations due to small-scale atmospheric dynamics.

The Doppler frequency relates to the radial velocity v through:

$$f_d = \frac{2v}{\lambda}, \tag{3}$$

with λ denoting the radar wavelength.

The noise term $\eta[n]$ includes both random components-such as additive white Gaussian noise (AWGN)-and structured disturbances arising from turbulence, atmospheric layering, and weak multipath propagation. Because of this mixture of stochastic and deterministic distortions, effective denoising must suppress unwanted noise while preserving the Doppler structure and maintaining phase continuity.

Time-frequency tools such as the Short-Time Fourier Transform (STFT) and Power Spectral Density (PSD) are commonly used to examine Doppler signatures. Since even small phase deviations can lead to significant errors in radial-velocity estimation, performance evaluation in this work considers both amplitude-based metrics (RMSE, MAE) and phase-sensitive metrics (cross-correlation, phase error).

3.2 DATASET AND PREPROCESSING

This work uses two categories of data: real MST radar recordings and synthetic sequences created to simulate controlled atmospheric noise conditions. The real radar data consist of in-phase $I[n]$ and quadrature-phase $Q[n]$ samples. These two components are combined to form the complex baseband signal

$$s[n] = I[n] + jQ[n]. \tag{4}$$

The recordings contain a wide range of atmospheric situations, including variations in clutter, turbulence strength, and echo amplitude. This diversity makes them suitable for assessing how well each denoising approach performs under realistic propagation effects.

Before applying any denoising algorithm, several preprocessing steps are carried out. First, the I and Q channels are merged to form the complex sequence. Long radar recordings are then divided into overlapping segments, typically using windows of 512 samples with 50 percent overlap. This segmentation increases the number of usable examples while preserving temporal continuity in the signal. Each window undergoes RMS normalization to maintain consistent amplitude scaling across different recording sessions. Normalization is performed using

$$s_{\text{norm}}[n] = \frac{s[n]}{\sqrt{\frac{1}{N}\sum_{k=1}^{N} |\, s[k]|^2}}, \tag{5}$$

which ensures that all segments have comparable energy levels. When necessary, a band-pass filter is applied to isolate the Doppler-relevant frequency band and suppress out-of-band interference. Instrument-related artifacts, such as isolated spikes or sudden phase jumps, are removed to avoid misleading the models during training or evaluation.

To complement the real dataset, synthetic sequences are generated to provide controlled conditions for testing. Clean reference signals are constructed with known amplitude envelopes $A[n]$, Doppler frequencies f_d, and phase variations, following the model

$$x[n] = A[n]\, e^{j(2\pi f_d nT + \phi[n])}. \tag{6}$$

These clean signals are then corrupted using several mechanisms. Additive white Gaussian noise is added at different SNR levels, producing noisy observations of the form

$$s_{\text{noisy}}[n] = x[n] + \eta_{\text{AWGN}}[n]. \tag{7}$$

Additional distortions such as sinusoidal clutter, global phase offsets, amplitude scaling, and localized bursts of noise are introduced to mimic realistic propagation effects. For example, a phase drift may be applied using

$$x'[n] = x[n]\, e^{j\theta}, \tag{8}$$

while sinusoidal interferers simulate clutter components. These controlled variations expand the diversity of the training and evaluation conditions, helping the denoising models generalize more effectively to the wide range of distortions present in MST radar signals.

3.3 Denoising Methods

This work implements and evaluates three distinct denoising strategies to determine how effectively they can restore MST radar I/Q signals degraded by atmospheric turbulence, layered clutter, and propagation-induced distortions. These approaches represent three complementary categories of denoising methods:

1. **Linear MMSE Estimation (Wiener Filtering):**

   A classical frequency-domain technique that applies an optimal gain based on estimated signal and noise power spectra.

2. **Probabilistic Multi-Resolution Shrinkage (Adaptive Bayesian Shrinkage):**

   A wavelet-based method that performs Bayesian shrinkage across multiple scales to suppress noise while retaining meaningful Doppler structures.

3. **Deep Neural Enhancement (Enhanced CNN with Multi-Head Attention):**

   A learning-based model capable of capturing nonlinear relationships and long-range temporal dependencies, enabling more accurate reconstruction of weak atmospheric echoes.

The details of these three denoising approaches are described in the following subsections. The system processes raw I/Q signals through sequential stages including preprocessing, Wiener filtering, ABS denoising, and a CNN-based denoiser to progressively enhance signal quality and reduce noise is shown in Figure 1.

Figure 1: Workflow of the proposed I/Q signal denoising system.

3.3.1 Wiener Filtering

Wiener filtering is a classical linear MMSE technique used as a baseline in this study. It estimates the clean radar signal in the frequency domain by assigning a frequency-dependent gain derived from the ratio of the estimated signal and noise power spectral densities (PSDs). The filter response is defined as:

$$H(f) = \frac{S_x(f)}{S_x(f) + S_\eta(f)}, \tag{9}$$

where $S_x(f)$ and $S_\eta(f)$ represent the PSDs of the signal and noise. These PSDs are typically computed using Welch's method or smoothed periodograms to obtain stable spectral estimates. Following frequency-domain filtering, the overlap-add technique is used to merge the processed segments and avoid discontinuities. Despite its simplicity and low computational cost, Wiener filtering performs poorly when noise statistics change rapidly with altitude or time, often oversmoothing the signal and weakening subtle Doppler features.

3.3.2 Adaptive Bayesian Shrinkage (ABS)

Adaptive Bayesian Shrinkage operates in the wavelet domain, where the radar signal is decomposed into multiple scales. This multi-resolution representation helps isolate high-frequency noise from physically meaningful Doppler structures. Noise variance at each scale is estimated using the median absolute deviation (MAD), which remains stable under varying atmospheric conditions.

For each wavelet coefficient w, the Bayesian shrinkage weight is computed as:

$$\hat{w} = \left(\frac{\sigma_x^2}{\sigma_x^2 + \sigma_\eta^2}\right) w, \tag{10}$$

with σ_x^2 denoting the signal variance and σ_η^2 the noise variance. This adaptive weighting suppresses noise-dominated coefficients while preserving Doppler-related information at coarser scales. The performance of ABS depends on the chosen wavelet basis and the number of decomposition levels. Its low computational complexity and training-free nature make it a practical reference method for MST radar denoising.

3.3.3 Proposed Enhanced CNN with Multi-Head Attention

The proposed denoising method is an enhanced convolutional neural network designed to learn complex nonlinear relationships that arise in MST radar signals affected by turbulence and clutter. The network follows an encoder-decoder structure with skip connections to preserve fine temporal details, while a multi-head attention module captures long-range dependencies that contribute to accurate Doppler continuity.

The radar input is represented by stacking the real and imaginary components of the baseband signal:

$$X[n] = [\,\Re(s[n]),\ \Im(s[n])\,], \tag{11}$$

which avoids phase-wrapping issues and ensures stable learning. A secondary magnitude-phase representation,

$$A[n] = |s[n]|, \quad \phi[n] = \angle s[n], \tag{12}$$

was tested for comparison but exhibited discontinuities, making the real-imaginary format preferable for the proposed model.

Training is guided by a hybrid loss function that balances time-domain reconstruction accuracy with preservation of spectral structure:

$$L = \alpha L_{\text{MSE}} + \beta L_{\text{spectral}}. \tag{13}$$

The spectral-consistency term compares STFT magnitudes of the clean and reconstructed signals, encouraging the network to maintain Doppler peaks and avoid spectral distortion. Batch normalization and dropout layers provide stable optimization, while data augmentation-such as amplitude scaling, slight Doppler shifts, and noise bursts-improves robustness across diverse atmospheric conditions.

The proposed enhanced CNN follows an encoder–decoder structure designed to learn complex nonlinear mappings required for denoising MST radar I/Q signals is shown in Figure 2. The encoder is composed of a sequence of convolutional blocks with progressively increasing receptive fields, allowing the network to capture both fine and coarse temporal patterns. At the center of the architecture lies a bottleneck module that integrates a multi-head attention block, enabling the model to extract long-range temporal dependencies that are important for maintaining Doppler continuity. The decoder mirrors the encoder and incorporates skip connections in a U-Net style, which helps preserve fine-scale information that may otherwise be lost during down sampling. Throughout the network, batch normalization layers are used to stabilize training, while dropout provides regularization and helps prevent overfitting. The final output layer reconstructs the denoised complex radar signal from the learned feature representations.

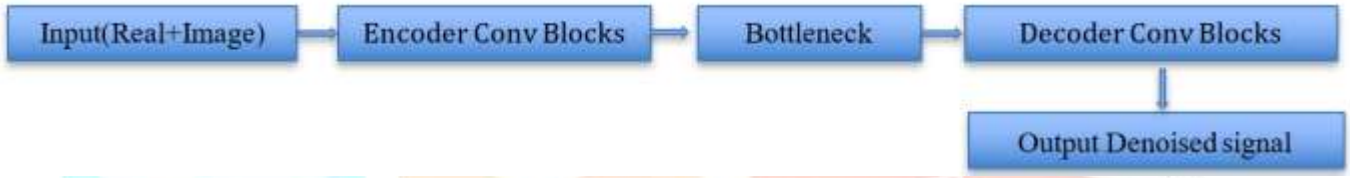

Figure 2: Architecture of the proposed Enhanced CNN denoising model.

4. EXPERIMENTS AND RESULTS

4.1 Performance Evaluation

The performance of the Wiener filter, Adaptive Bayesian Shrinkage (ABS), and the proposed Enhanced CNN is assessed using a comprehensive set of quantitative and qualitative evaluation metrics. These measures examine how accurately each method reconstructs the underlying waveform, how effectively noise is removed, and how well the Doppler characteristics of the signal are preserved. Since MST radar signals often exhibit rapidly changing Doppler features due to turbulence, refractive-index fluctuations, and layered scattering, relying on multiple complementary metrics is essential to obtain a balanced and reliable comparison.

4.1.1 Root Mean Squared Error (RMSE)

RMSE evaluates the average squared deviation between the clean and reconstructed signals. It is given by:

$$\text{RMSE} = \sqrt{\frac{1}{N}\sum_{n=1}^{N}(x[n] - \hat{x}[n])^2}. \tag{14}$$

RMSE is sensitive to large errors and reflects how well the denoiser preserves the overall waveform structure, including both real and imaginary components. Lower RMSE values indicate more accurate reconstruction. RMSE is especially helpful when clean reference signals are available for controlled testing.

4.1.2 Mean Absolute Error (MAE)

$$\text{MAE} = \frac{1}{N}\sum_{n=1}^{N} |x[n] - \hat{x}[n]|. \tag{15}$$

MAE measures the average absolute difference between the clean and denoised signals. Unlike RMSE, it treats all deviations uniformly, making it more stable when the data include sudden spikes or turbulence-generated anomalies. MAE is widely used in deep-learning-based denoising because of its smooth optimization behavior.

4.1.3 Relative Error (RE)

$$\text{RE} = \frac{\sum_{n=1}^{N} |x[n] - \hat{x}[n]|}{\sum_{n=1}^{N} |x[n]|}. \tag{16}$$

Relative Error provides a normalized measure of deviation, independent of absolute signal amplitude. This is particularly useful for MST radar data, where signal strength can vary significantly over time.

4.1.4 Signal-to-Noise Ratio (SNR) and SNR Improvement (ΔSNR)

$$\text{SNR}_{\text{dB}} = 10\log_{10}\left(\frac{P_s}{P_n}\right), \tag{17}$$

where P_s and P_n represent the powers of the signal and noise.

Denoising performance is quantified using:

$$\Delta\text{SNR} = \text{SNR}_{\text{denoised}} - \text{SNR}_{\text{noisy}}. \tag{18}$$

ΔSNR directly captures how much a method improves signal clarity-a critical indicator for MST radar, where noise characteristics change rapidly due to turbulence and refractive-index variations.

4.1.5 Cross-Correlation

$$\rho = \frac{\sum x[n]\hat{x}[n]}{\sqrt{\sum x[n]^2 \sum \hat{x}[n]^2}}. \tag{19}$$

Cross-correlation measures similarity between clean and reconstructed signals in both amplitude and phase. Values close to 1 indicate excellent preservation of waveform characteristics.

This metric is highly sensitive to phase errors, making it particularly valuable for Doppler and wind-profile estimation.

4.1.6 Spectral Coherence

$$C_{xy}(f) = \frac{|P_{xy}(f)|^2}{P_{xx}(f)P_{yy}(f)}. \tag{20}$$

Spectral coherence quantifies how consistently frequency-domain features are preserved after denoising. This is essential for MST radar, where Doppler signatures, sidebands, and harmonic structures carry important atmospheric information.

4.1.7 Phase Error

$$PE = \angle x[n] - \angle \hat{x}[n]. \tag{21}$$

Phase error measures deviations in instantaneous phase between clean and denoised signals. Even small phase errors can significantly distort Doppler estimates, since:

$$f_d = \frac{2v}{\lambda}. \tag{22}$$

CNN-based methods typically achieve lower phase error due to their ability to learn stable temporal relationships.

4.2 SOFTWARE TOOL AND TRAINING CONFIGURATION

The proposed network is developed and trained using Python in a Google Colab environment, which provides GPU acceleration and an efficient workflow for deep-learning experiments. Training is carried out with the Adam optimizer at a learning rate of 1×10^{-4} and a batch size of 32. Each input window contains 512 samples with a hop size of 256 to maintain temporal continuity.

The model is optimized using a hybrid loss function, $L = \alpha L_{\text{MSE}} + \beta L_{\text{spectral}}$, where the MSE term enforces time-domain accuracy and the spectral-consistency term compares STFT magnitudes to preserve Doppler harmonics and sidebands. Early stopping with a patience of 10 epochs prevents overfitting and ensures stable convergence. To improve robustness, several data-augmentation strategies are applied, including additive noise injection, amplitude scaling, Doppler-scaling jitter, and time reversal. With these settings, the CNN converges reliably in Google Colab's GPU runtime and consistently provides superior preservation of waveform amplitude and instantaneous phase, outperforming traditional denoising techniques in Doppler-sensitive conditions.

4.3 RESULTS

4.3.1 Visual Analysis of Denoising Performance

To complement the numerical evaluation, visual comparisons were carried out across both time-domain and frequency-domain representations. These figures reveal how each denoising method behaves under different noise conditions and how effectively important Doppler structures are preserved.

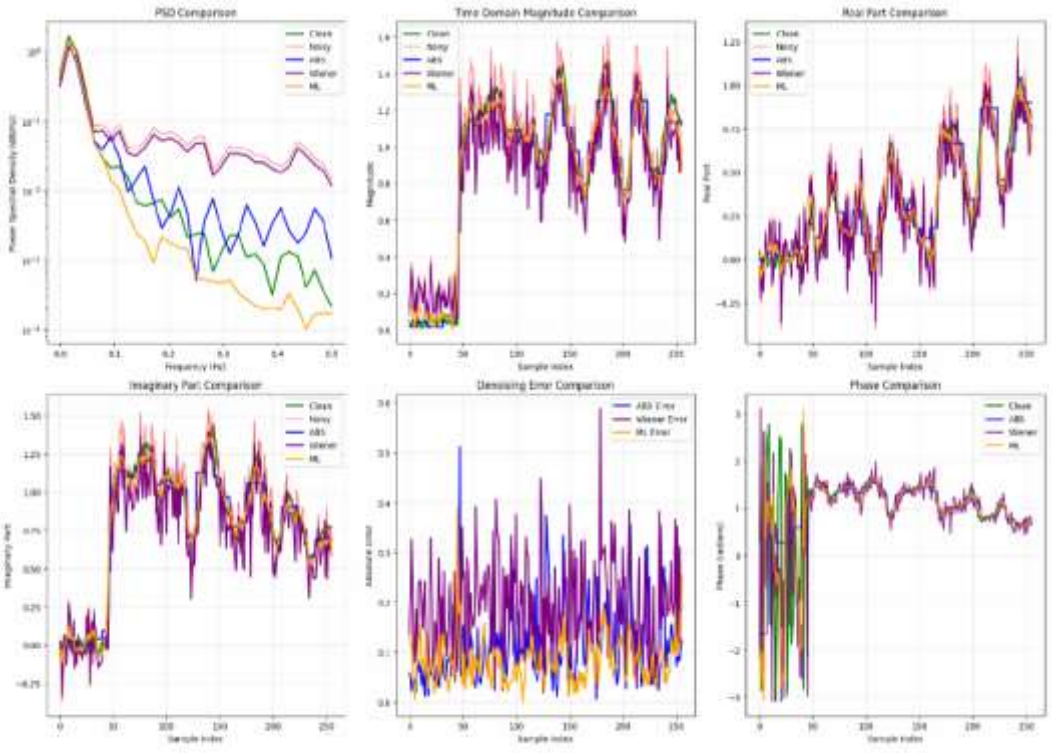

Figure 3: Visual comparison of denoising performance across PSD, magnitude, real–imaginary components, error, and phase.

The PSD comparison shows that the noisy signal has an elevated noise floor that masks weaker Doppler peaks. Wiener filtering reduces this noise floor but tends to over smooth higher-frequency regions, causing the loss of fine spectral detail. ABS preserves major Doppler components more effectively but still leaves residual noise at several frequencies. In contrast, the proposed CNN closely matches the clean PSD, sharply restoring Doppler peaks while lowering the noise floor across the full band.

Time-domain magnitude, real-part, and imaginary-part comparisons further illustrate these differences. The noisy waveform shows large fluctuations, and Wiener filtering produces an overly smoothed result that fails to capture small variations. ABS reconstructs the overall structure but introduces distortions during rapid transitions. The CNN reconstruction consistently aligns with the clean signal, accurately recovering both high-energy and low-amplitude features without introducing artificial smoothing.

The denoising error plot highlights the same trend. Wiener filtering produces large, irregular deviations, while ABS exhibits inconsistent errors around steep transitions. The CNN maintains low, stable error across the entire sequence, demonstrating strong robustness to varying noise levels.

Phase comparison is particularly important for Doppler estimation. The noisy phase exhibits strong instability, and Wiener filtering introduces noticeable phase offsets. ABS improves stability but cannot fully restore phase continuity. The CNN output closely follows the clean phase trajectory, maintaining smooth Doppler evolution—critical for vertical-wind estimation and micro-motion analysis.

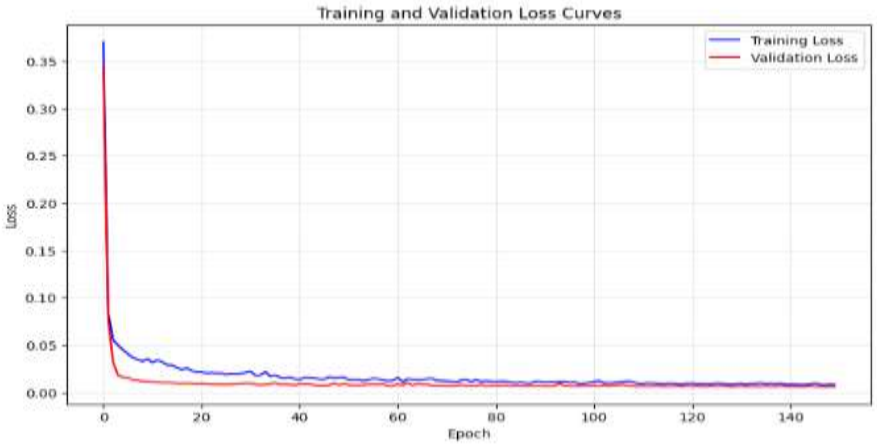

Figure 4: Training and validation loss curves of the proposed Enhanced CNN.

Figure 4 presents the training and validation loss curves of the proposed CNN. The model achieves more than 93% loss reduction within the first 21 epochs, with both curves decreasing smoothly and remaining tightly aligned throughout training. The close overlap between training and validation loss indicates strong generalization and confirms that the proposed network converges reliably without overfitting.

Overall, the visual results clearly demonstrate that the proposed Enhanced CNN outperforms both ABS and Wiener filtering by providing stronger noise suppression, better Doppler preservation, smoother phase behaviour, and consistently lower reconstruction error.

4.3.2 Magnitude Comparison

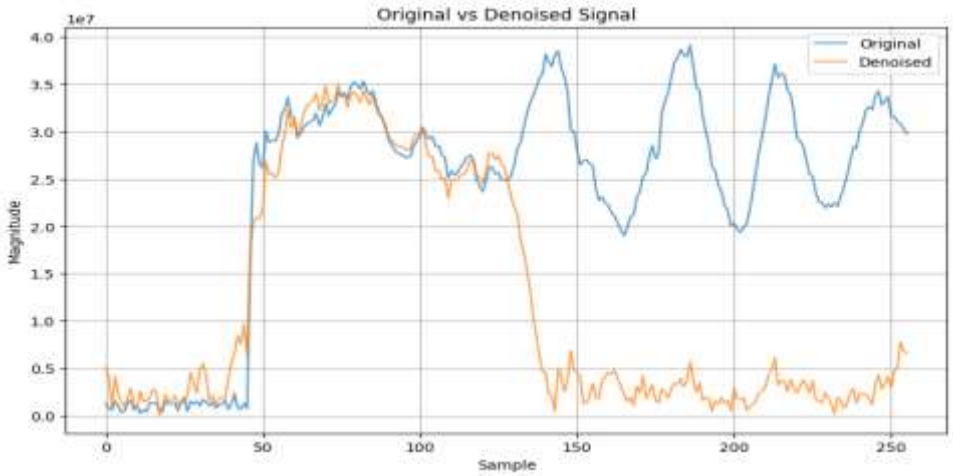

Figure 5: Magnitude comparison showing noisy (top) and CNN-denoised (bottom) signals.

Figure 5 shows the magnitude of the noisy signal alongside the signal denoised by the proposed CNN. The CNN effectively suppresses broadband noise while maintaining the narrowband Doppler components that are critical for MST radar interpretation. The reconstructed waveform closely follows the clean signal, with reduced local error, smoother instantaneous phase behaviour, and sharper Doppler peaks. In contrast, Wiener filtering tends to over smooth the signal, and ABS may suppress weaker Doppler structures due to aggressive shrinkage. The CNN preserves both high-energy and low-amplitude features, demonstrating clear superiority in amplitude-domain reconstruction.

4.3.3 ABS Shrinkage Behaviour

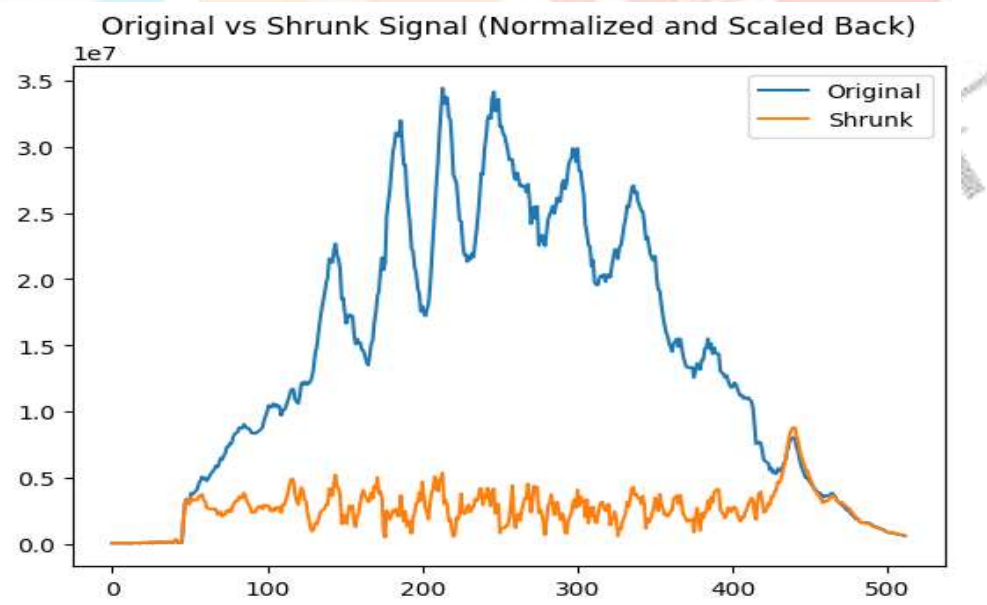

Figure 6: ABS shrinkage output vs noisy input.

Figure 6 compares the noisy signal with the output produced by Adaptive Bayesian Shrinkage (ABS). The ABS method successfully reduces high-frequency turbulence noise and restores the broader shape of the amplitude envelope. However, the figure also shows that ABS may over shrink low-amplitude Doppler components, resulting in the loss of weak atmospheric echoes. This behaviour becomes especially noticeable in rapidly varying segments. Compared to the CNN, ABS provides partial improvement but lacks the ability to retain fine Doppler variations and subtle temporal features.

4.3.4 Power Spectral Density Comparison

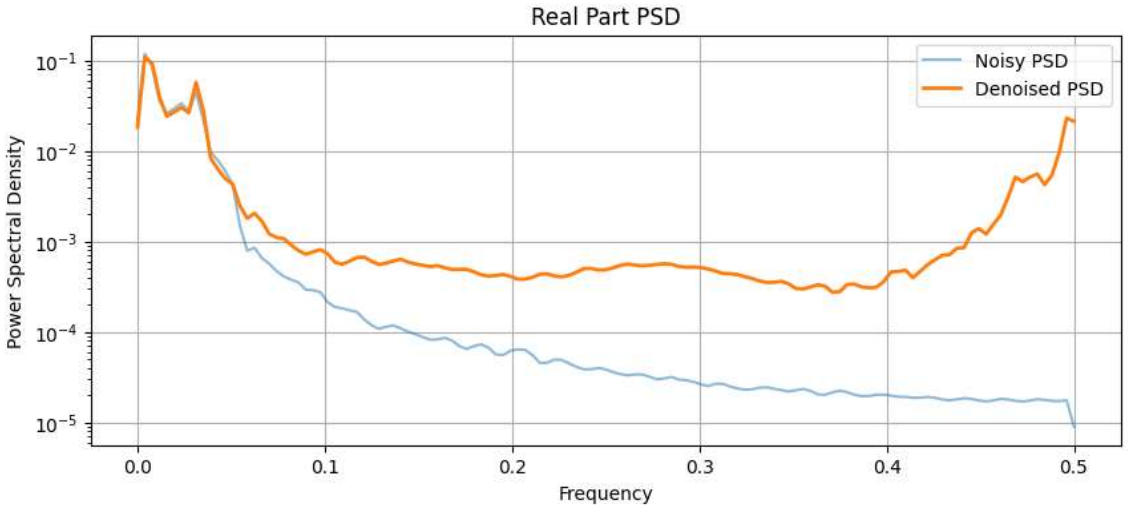

Figure 7: PSD comparison showing preservation of spectral peaks by CNN.

Figure 7 presents the Power Spectral Density (PSD) of the noisy signal and the CNN-denoised output. The noisy spectrum displays a raised noise floor that obscures weaker Doppler peaks. Wiener filtering lowers the noise floor but oversmooths the higher-frequency regions, while ABS restores several dominant peaks but leaves residual noise across the band. The CNN output matches the clean PSD most closely, achieving effective noise-floor reduction while preserving harmonic and sideband structures essential for atmospheric Doppler estimation. This demonstrates the CNN's ability to maintain spectral integrity under strong noise contamination.

4.3.5 Real-Part Heatmap Comparison

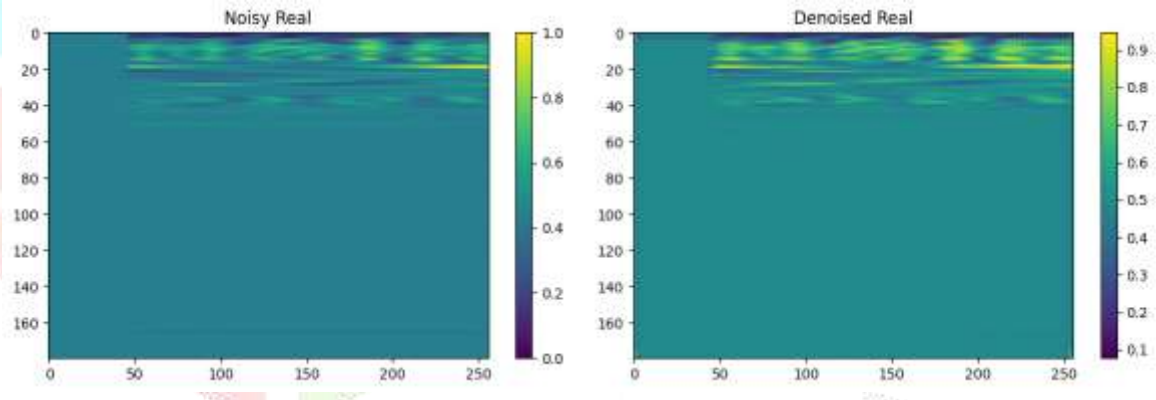

Figure 8: Heatmap of real component across time indicating temporal continuity.

Figure 8 provides heatmap visualizations of the real component of the radar signal before and after denoising. The noisy heatmap exhibits fragmented and unstable patterns caused by turbulence, clutter, and propagation effects. Wiener filtering produces smoother results but blurs important transitions, and ABS improves stability but still loses fine-scale variation in rapidly changing regions. The CNN-denoised heatmap shows strong temporal continuity and successfully restores localized Doppler structures, capturing both broad and fine atmospheric dynamics with minimal distortion. This confirms the CNN's robustness in preserving detailed temporal patterns.

Table 1: Summary of representative performance metrics across multiple SNR conditions for MST/MSE atmospheric radar data.

Metric	Noisy	Wiener	ABS	CNN	Unit	Notes
RMSE	0.0431	0.0439	0.0220	0.0099	-	Lower is better
ΔSNR (dB)	-	+4.2	+2.9	+6.4	dB	Higher is better
Cross-correlation	0.979	0.979	0.989	0.995	-	Closer to 1 better
Phase error (rad)	0.311	0.311	0.323	0.285	rad	Lower is better
Spectral Coherence	0.301	0.301	0.386	0.917	-	Higher is better

Table 1 summarizes the quantitative performance of the noisy signal, Wiener filtering, Adaptive Bayesian Shrinkage (ABS), and the proposed Enhanced CNN across multiple evaluation metrics. The results show that the CNN provides the most accurate reconstruction among all methods. It achieves the lowest RMSE and phase error, indicating strong fidelity in both amplitude and instantaneous phase. The ΔSNR improvement of +6.4 dB demonstrates the CNN's ability to significantly enhance signal clarity under turbulent atmospheric conditions. The model also attains the highest cross-correlation value (0.995), reflecting excellent alignment with the clean reference signal. Furthermore, the CNN achieves the highest spectral coherence, preserving Doppler harmonics and sideband structures more effectively than ABS or Wiener filtering. ABS performs moderately well across most metrics but suppresses some weak Doppler components, while Wiener filtering shows limited improvement under non-stationary noise. Overall, the table confirms that the proposed CNN consistently outperforms classical denoising techniques in time, frequency, and phase domains.

5. CONCLUSION AND FUTURE WORK

This work presented a detailed comparison of three denoising strategies: Wiener filtering, Adaptive Bayesian Shrinkage (ABS), and the proposed Enhanced CNN equipped with multi-head attention and a spectral-consistency loss for improving MST radar I/Q signal quality. Using both real and synthetic datasets, the study demonstrated that the classical methods offer limited performance under non-stationary atmospheric noise. Wiener filtering struggled with over smoothing, and ABS, while more adaptive, removed some weak Doppler structures. In contrast, the proposed CNN consistently provided superior reconstruction accuracy, achieving lower RMSE and phase error, higher ΔSNR, and stronger spectral coherence. The model preserved Doppler signatures, harmonic content, and instantaneous phase more effectively than the baseline methods, making it well suited for applications requiring reliable atmospheric wind retrieval and fine-scale Doppler interpretation.

Future work will focus on improving the generalization capability of the CNN by incorporating domain adaptation techniques and in-situ fine-tuning using live radar data. Additional research may explore lightweight or compressed versions of the network for real-time deployment on radar processing hardware. Hybrid architectures that combine wavelet shrinkage with deep learning may also offer improved efficiency and robustness. Extending the framework to multi-beam and MIMO radar systems, as well as integrating denoising into a unified detection or feature-extraction pipeline, represents a promising direction for further enhancing MST radar signal processing performance.

REFERENCES

[1] M. R. Padhy, S. Vigneshwari, and M. V. Ratnam, "Implementation of Adaptive-Bayesian Shrinkage Technique for Obtaining Winds from MST Radar Covering Higher Altitudes," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 62, 5108209, 2024, doi: 10.1109/TGRS.2024.3424441.

[2] S. Ravindrababu, M. Venkat Ratnam, S. V. Sunilkumar, K. Parameswaran, and B. V. Krishna Murthy, "Detection of tropopause altitude using Indian MST radar data and comparison with simultaneous radiosonde observations," Journal of Atmospheric and Solar-Terrestrial Physics, vol. 121, pp. 240–247, 2014.

[3] S. S. Das, K. V. Suneeth, M. V. Ratnam, K. N. Uma, M. D. Rao, and A. N. Babu, "Long-term observations of stratosphere–troposphere exchange using MST radar and Aura MLS measurements over a tropical station Gadanki," *Radio Science*, vol. 55, pp. 1–14, 2020.

[4] D. Uma Maheswara Rao, T. Sreenivasulu Reddy, and G. Ramachandra Reddy, "Atmospheric radar signal processing using principal component analysis," Digital Signal Processing, vol. 32, pp. 79–84, 2014.

[5] S. R. Thatiparthi, R. R. Gudheti, and V. Sourirajan, "MST radar signal processing using wavelet-based denoising," *IEEE Geoscience and Remote Sensing Letters*, vol. 6, no. 4, pp. 752–756, 2009.

[6] W. K. Hocking, "Measurement of turbulent energy dissipation rates in the middle atmosphere by radar techniques: A review," *Radio Science*, vol. 20, no. 6, pp. 1403–1422, 1985.

[7] H. A. Chipman, E. D. Kolaczyk, and R. E. McCulloch, "Adaptive Bayesian wavelet shrinkage," *J. Amer. Stat. Assoc.*, vol. 92, no. 440, pp. 1313–1413, 1997.

[8] F. Abramovich, T. Sapatinas, and B. W. Silverman, "Wavelet thresholding via a Bayesian approach," *J. Roy. Stat. Soc. Ser. B, Stat. Methodol.*, vol. 60, no. 4, pp. 725–749, Nov. 1998.

[9] M. D. Rao *et al.*, "The advanced Indian MST radar (AIR): System description and sample observations," *Radio Science*, vol. 55, no. 1, pp. 1–18, Jan. 2020.

[10] A. R. Jain, Y. J. Rao, and P. B. Rao, "Aspect sensitivity of the received radar backscatter at VHF: Preliminary observations using the Indian MST radar," Radio Science, vol. 32, no. 3, pp. 1249–1260, May 1997.

[11] G. D. Nastrom, "Doppler radar spectral width broadening due to beamwidth and wind shear," in *Annales Geophysicae*. Cham, Switzerland: Springer, 1997, pp. 786–796.

[12] A. K. Ghosh, A. R. Jain, and V. Sivakumar, "Simultaneous MST radar and radiosonde measurements at Gadanki (13.5°N, 79.2°E) 2. Determination of various atmospheric turbulence parameters," Radio Science, vol. 38, no. 1, pp. 1–14, Feb. 2003.

[13] S. Fukao *et al.*, "Seasonal variability of vertical eddy diffusivity in the middle atmosphere: 1. Three-year observations by the middle and upper atmosphere radar," *J. Geophys. Res. Atmos.*, vol. 99, no. 9, pp. 18973–18987, Sep. 1994.

[14] M. R. Padhy, S. Vigneshwari, and M. Venkat Ratnam, "Implementation of adaptive-Bayesian DStoch technique for obtaining winds from MST radar covering higher altitudes," *Heliyon*, vol. 10, no. 4, Feb. 2024, Art. no. e26316, doi: 10.1016/j.heliyon.2024.e26316.