Scientific Research Publishing

Neuroscience and Medicine, 2025, 16(3), 114-141
https://www.scirp.org/journal/nm
ISSN Online: 2158-2947
ISSN Print: 2158-2912

Enhancing MRI Image Quality through Deep CNN-Augmented Denoising: A Comparative Study of Standard and Hybrid Filters

Samuel Ocen[1,2], Lawrence Muchemi[1], Michealina Almaz Yohannis[1]

[1]Department of Computer Science and Informatics, University of Nairobi, Nairobi, Kenya
[2]Department of Computer Science, Mountains of the Moon University, Fort Portal, Uganda
Email: samocenuel@gmail.com

How to cite this paper: Ocen, S., Muchemi, L. and Yohannis, M.A. (2025) Enhancing MRI Image Quality through Deep CNN-Augmented Denoising: A Comparative Study of Standard and Hybrid Filters. *Neuroscience and Medicine*, **16**, 114-141.
https://doi.org/10.4236/nm.2025.163013

Received: June 5, 2025
Accepted: August 23, 2025
Published: August 26, 2025

Abstract

Magnetic Resonance Imaging (MRI) is commonly applied to clinical diagnostics owing to its high soft-tissue contrast and lack of invasiveness. However, its sensitivity to noise, attributable to hardware limitations, patient motion, and acquisition parameters, remains a long-term source of concern that normally leads to impaired diagnostic accuracy. Traditional denoising filters such as Gaussian, Wavelet, Anisotropic Diffusion, and Non-Local Means (NLM) have previously been employed to reduce these issues, but at the cost of typically sacrificing noise removal against structural detail. More recent advances in deep learning, in particular Convolutional Neural Networks (CNNs), have indicated excellent promise in overcoming these limitations in being capable of learning data-driven, highly robust feature representations. This study proposes and compares an extensive denoising pipeline that unites CNNs with both classical and hybrid filters to improve the quality of MRI images. The pipeline encompasses not only independent classical filters but also new hybrid pipelines like NLM-Gaussian Fusion, NLM-Gaussian Sequential, and Wavelet-Anisotropic Diffusion (WAD), all enriched with deep CNNs. Experiments were performed using a publicly shared MRI dataset of 5141 training images and 1279 test images. Performance was gauged using Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM), and Mean Squared Error (MSE). Results show that while the Wavelet filter performed best as a standalone denoising filter, when CNNs were integrated with hybrid filters, there were monumental improvements in all metrics tried, with the pipeline of NLM-Gaussian Fusion + CNN achieving a PSNR of 30.4 and SSIM of 0.65. Moreover, visualizations such as SSIM heatmaps and loss-epoch convergence plots supported the efficacy of the proposed models in preserving structural

information. The study reveals that the use of CNNs with conventional and hybrid filters offers synergistic benefits, especially for low-resource clinical setups where denoising quality is paramount. The novelty in this research comes from its systematic benchmarking of traditional as well as hybrid filtering pipelines supplemented with CNNs, providing empirical evidence for their utility in real-world medical imaging scenarios. These findings not only contribute to the widening ambit of AI-assisted image reconstruction but also provide functional avenues towards enhancing the dependability of diagnosis in resource-scarce healthcare environments.

Keywords

MRI Denoising, Convolutional Neural Networks (CNNs), Hybrid Filters, Image Quality Enhancement, Medical Image Processing

1. Introduction

Magnetic Resonance Imaging (MRI) is a mainstay of current diagnostic medicine due to its non-invasiveness and superior soft tissue contrast [1]. However, it is also highly susceptible to noise artifacts caused by hardware limitations, patient motion, and long acquisition times, which can significantly degrade image quality [2] [3]. High noise levels not only reduce diagnostic clarity for radiologists but also impair downstream tasks such as segmentation and classification [4].

To mitigate these challenges, several denoising methods have been explored. Classical filters like Gaussian, Anisotropic Diffusion, Non-Local Means (NLM), and Wavelet-based approaches are widely used due to their simplicity and interpretability [5]-[8]. However, these methods often struggle to balance noise suppression with preservation of structural details. Over-smoothing effects may obscure critical anatomical boundaries [9].

Deep learning, particularly Convolutional Neural Networks (CNNs), has transformed image denoising by learning complex mappings between noisy and clean images [10] [11]. More recently, hybrid pipelines that combine classical filters with CNNs have gained attention for their potential to merge handcrafted smoothing with learned representations [12].

Despite this progress, comparative research examining the effectiveness of CNNs when integrated with both classical and hybrid denoising filters remains limited. Prior studies often focus on a single method or dataset, lacking extensive benchmarking across diverse filtering strategies. In particular, the combination of CNNs with emerging hybrid filters, such as NLM-Gaussian fusion and Wavelet-Anisotropic Diffusion, has not been rigorously evaluated.

To address this gap, this study proposes a systematic investigation of standard, hybrid, and CNN-augmented denoising pipelines using a large publicly available MRI dataset. The goal is to assess their comparative effectiveness using standard

evaluation metrics and practical visualizations in contexts relevant to low-resource clinical environments.

1.1. Contributions and Objectives of the Study

This study contributes significantly to medical image processing by addressing the long-standing problem of removing noise from Magnetic Resonance Imaging (MRI) scans without compromising essential anatomical information [13]. While traditional filtering methods such as Gaussian, Wavelet, Anisotropic Diffusion, and Non-Local Means (NLM) are largely used in MRI denoising, they rarely provide the optimum balance between noise reduction and preservation of fine image details [8] [6]. In addition, the integration of classical denoising with deep learning techniques, specifically Convolutional Neural Networks (CNNs), has yet to be exhaustively explored in normal and hybrid filtering pipelines [14]. This paper fills this gap by proposing and evaluating an integrated denoising strategy that combines CNNs with several classical and hybrid filters and therefore leverages both learned and hand-designed representations. The key contributions of this paper are fourfold. First, it thoroughly tests and compares the performance of four standard filters (Wavelet, Gaussian, Anisotropic Diffusion, and NLM) and three hybrid filtering methods (Wavelet-Anisotropic Diffusion, NLM-Gaussian Fusion, and NLM-Gaussian Sequential) in combination with a deep CNN model. Second, it introduces a novel experimental setup where CNN-augmented filtering pipelines are compared on a big MRI dataset under the same settings. Thirdly, this paper provides thorough quantitative evaluations using industry-standard metrics, Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM), and Mean Squared Error (MSE), and qualitative visualizations such as SSIM heatmaps and loss convergence plots. Finally, this piece of research is set against the backdrop of low-resource clinical settings where computationally efficient algorithms for high-quality image reconstruction are of paramount importance [10].

To guide the research, the following objectives were established:

1) To implement and evaluate the performance of the classical denoising filters, Wavelet, Gaussian, NLM, and Anisotropic Diffusion, on noisy MRI scans.
2) To establish and evaluate hybrid denoising pipelines through the integration of traditional filters, viz., Wavelet-Anisotropic Diffusion, NLM-Gaussian Fusion, and NLM-Gaussian Sequential.
3) To integrate a Convolutional Neural Network (CNN) with each filtering pipeline and evaluate the improvement in denoising accuracy.
4) Comparing all denoising methods, baseline, hybrid, and CNN-augmented, based on PSNR, SSIM, and MSE values.
5) Visualizing and explaining the effect of various denoising strategies via SSIM heatmaps, reconstruction figures, and training loss plots for comparison and understanding.

1.2. Organization of the Paper

This paper is divided into six major sections to give a comprehensive discussion

of CNN-aided denoising techniques for MRI image enhancement. Following the Introduction in Section 1, Section 2 gives a comparative review of pertinent literature, stating current conventional, hybrid, and deep learning-based denoising techniques. Section 3 outlines the materials and methods used in this work, including dataset description, preprocessing methods, suggested model architecture, evaluation metrics, and experimental procedure. Section 4 provides an extensive analysis of quantitative and qualitative outcomes backed by PSNR, SSIM, and MSE measures, and visualizations including heatmaps and training loss plots. Section 5 situates the results within the larger body of earlier research, clarifies the clinical implications, and outlines the applicability of the proposed model in low-resource settings. Finally, Section 6 concludes the key findings, model performance observations, and gives future work directions.

2. Related Literature

The study [15] revealed that medical imaging plays a vital role in the diagnosis and treatment of disease, but it can be hard to interpret with noise and artifacts. Denoising techniques such as deep learning can improve image quality by removing noise and artifacts. This review article provides an overview of currently used deep learning methods for denoising medical images, out of which 40% use deep convolutional neural networks, 18% use encoder-decoder, 15% use artificial intelligence-based approaches, and 2% use multilayer perceptron, 35% of the images have Gaussian noise, and much effort must be put into developing standardized denoising models that can work optimally on a wide variety of medical images.

The study [16] revealed that image processing of digital images is a very crucial part of building image-based applications. In medical imaging, the image processing step is one of the most crucial steps that needs maximum accuracy to detect and determine the nature of the disease. Its objective is to surmount the noise problem of medical images and preserve information and edges in the image. Medical images can be enhanced by removing noise with traditional and Deep Learning (DL) methods. CNN-based DL methods have produced outstanding results in the processing stage for reducing noise in medical images. DL is a strong and efficient technique for real noise estimation and feature representation extraction from images. This paper is a survey of image denoising methods for medical images, considering sources of noise and types of noise. Concepts of noise reduction (denoising) for various methods are discussed. Additionally, there is a comparative study that will elucidate the advantages and disadvantages of each method. Finally, some potential future work trends are outlined.

The study [17] introduces a teacher-student network model that uses the Noise-ContextNet Block to detect and mitigate noise in medical imaging. The model is refined for computational efficiency without compromising denoising fidelity. Experiments show significant qualitative enhancements across various imaging modalities, setting a new benchmark in medical image denoising.

The study [18] reviews noise reduction techniques in medical imaging and addresses issues like quantum, electronic, and radiation interference. Spatial, frequency domain, statistical, probability-based, and adaptive filtering methods are contrasted based on their capability to suppress noise while losing no diagnostic information. The study also identifies future research directions, like the implementation of advanced Machine Learning techniques and multimodal data fusion.

The study [19] presented an unsupervised medical image denoising technique that learns the noise patterns from the images and constructs denoised images. It uses patch-based dictionaries to learn the noise indirectly and residue learning to learn it directly. It is extended to 2D and 3D images and covers different medical imaging instruments. The K-SVD algorithm is utilized for sparse representation, and a deep residue network for residue training. Tests on MRI/CT images illustrate that the algorithm preserves significant information and improves visual quality.

The study [20] revealed that advancement in medical imaging technology enables the obtaining of valuable medical information for disease diagnosis, treatment, and research with accuracy. The technologies image the interior of the human body vivo, enabling the direct examination of organ function and identification of pathogens. The growing imaging data, however, needs further processing by computers for denoising and feature preservation. As big volumes and charts are complicated, machine learning is typically utilized for automation and image classification, where it surpasses conventional techniques in most applications.

The study [21] proposed a novel denoising method for medical CT images based on the U-network and multi-attention. The method utilizes three attention modules of local, multi-feature channel, and hierarchical. The local module localizes feature information, the multi-feature channel module learns and extracts features, and the hierarchical module extracts rich feature information. The enhanced learning module with enhanced learning capacity learns and remembers minute details using the integration of multi-layer convolution, batch normalization, and activation function layers. Experimental studies validate that the method achieved 34.7329 of PSNR and 0.9293 of SSIM for $\sigma = 10$ in the QIN_LUNG_CT dataset, and achieved 28.9163 of PSNR and 0.8602 of SSIM in the Mayo Clinic LDCT Grand Challenge dataset.

The study [22] introduced a new attention-guided denoising convolutional neural network (ADNet), comprising a sparse block, feature enhancement block, attention block, and reconstruction block. SB realizes a balance between performance and efficiency, FEB enhances expressive power, AB extracts noise information, and RB reconstructs clean images. ADNet records excellent performance in synthetic and real noisy images and blind denoising.

The study [23] revealed that image denoising faces challenges from noise sources like Gaussian, impulse, salt, pepper, and speckle. Convolutional neural networks (CNNs) have gained attention for image denoising. This paper explores various CNN techniques, categorized and analyzed, and examines popular datasets. The study provided a review of CNN image denoising, outlines motivations

and principles, and highlights potential challenges and directions for future research.

The study [24] proposed a progressive residual and attention mechanism fusion with a light network to address this issue. This architecture reduces network parameters and extracts local features, combining shallow convolutional features with deeper ones. The network performs better than more than 20 current methods on six datasets, showcasing outstanding image denoising performance while preserving necessary features such as edges and textures. The model can be used in many image-focused applications.

Synthesis of the Literature

The reviewed studies affirm CNNs as the leading approach for medical image denoising, often outperforming traditional and hybrid methods. Emerging attention-based, unsupervised, and GAN-driven models show promise in preserving structural integrity, as shown in **Table 1**. However, standardization across datasets and noise types remains a key challenge for real-world deployment.

Table 1. Synthesis of literature on medical image denoising.

Study	Objective	Method/Architecture	Noise Type/Source	Key Contributions
[15]	Survey of deep learning-based denoising in medical imaging	40% CNNs, 18% encoder-decoder, 15% AI-based, 2% MLP	35% involve Gaussian noise	Highlights the need for standardized DL denoising models
[16]	Survey of denoising methods and their impact on diagnostic accuracy	CNN-based deep learning, traditional filters	General noise from acquisition	DL surpasses traditional methods; key for noise estimation and edge preservation
[17]	Enhance denoising with efficiency	Teacher-Student Network (NoiseContextNet Block)	Multi-modal noise	Efficient denoising with fidelity across modalities
[18]	Comparative analysis of standard denoising techniques	Spatial, frequency, statistical, and adaptive filtering	Quantum, electronic, radiation noise	Highlights ML and multimodal fusion as future directions
[19]	Unsupervised denoising through noise learning	Patch-based dictionary + K-SVD + deep residual learning	Unknown/synthetic MRI and CT noise	Preserves key features; applicable to 2D and 3D imaging
[20]	Address the growing complexity of imaging data	Machine learning and automation	High-volume imaging datasets	ML surpasses conventional methods in denoising and classification
[21]	Denoising CT images with enhanced attention	U-Net + multi-attention (local, channel, hierarchical)	Gaussian noise (σ = 10)	Achieved PSNR 34.73 / SSIM 0.9293 on QIN_LUNG_CT dataset
[22]	CT denoising for low-dose applications	PWGAN-WSHL with hybrid loss function	LDCT noise + artifacts	Outperforms SoTA methods; effective artifact minimization
[23]	Design an attention-based blind denoising model	ADNet (sparse, feature enhancement, attention, reconstruction blocks)	Synthetic + real-world noise	High performance in both synthetic and blind denoising tasks
[24]	Denoising with minimal parameters and high fidelity	Progressive residual + attention + lightweight CNN	Gaussian, impulse, salt & pepper, speckle	Top performance on 6 datasets; excellent edge preservation

3. Materials and Methods

3.1. Dataset Description

This study utilized 6420 T1-weighted brain MRI images sourced from the Alzheimer's Disease Neuroimaging Initiative (ADNI) [25]. The dataset was stratified across four clinical classes: Mild Demented, Moderate Demented, Non-Demented, and Very Mild Demented. A total of 5141 images were used for training, and 1279 for testing. Each image was resized to 128 × 128 pixels and normalized to the [0, 1] intensity range to standardize the inputs across all experiments.

3.2. Data Preprocessing

To support consistent evaluation across all denoising pipelines, a standardized preprocessing procedure was applied. The brain MRI dataset was organized into four diagnostic categories: Mild Demented, Moderate Demented, Non-Demented, and Very Mild Demented. Each image was resized to 128 × 128 pixels using bilinear interpolation to ensure uniform spatial dimensions and reduce computational load. Pixel intensity values were normalized to the [0, 1] range by dividing by 255, improving the stability of gradient-based optimization during CNN training. All images were converted to grayscale and reshaped to a single-channel format compatible with PyTorch tensor input. To simulate realistic noise degradation, Additive White Gaussian Noise (AWGN) was introduced to each normalized image with zero mean ($\mu = 0$) and a standard deviation of $\sigma = 25$. This noise level was selected based on previous studies that commonly apply σ values between 15 and 35 to emulate moderate-to-severe noise found in low-field MRI or rapid acquisition scenarios. These perturbations represent common sources of clinical image degradation such as motion, thermal noise, and scanner limitations. Each noisy image was paired with its original clean version to form supervised input–target pairs for denoising. These pairs were converted into PyTorch tensors and wrapped in custom dataset classes to facilitate efficient mini-batch loading during training. Since the task was strictly denoising, no label encoding or class balancing was performed. This pipeline ensured that all models operated on consistent, clinically relevant inputs, enabling a fair and reproducible comparison of standard, hybrid, and CNN-augmented denoising strategies.

3.3. Proposed Model

The proposed denoising framework integrates classical and hybrid image filters with a deep Convolutional Neural Network (CNN) to enhance the quality of noisy Magnetic Resonance Imaging (MRI) images. This two-stage pipeline was carefully designed to leverage the complementary strengths of deterministic filters and data-driven deep learning. Traditional filters such as Gaussian, Non-Local Means (NLM), Wavelet, and Anisotropic Diffusion are known for their computational efficiency and ability to suppress noise through spatial smoothing or frequency domain decomposition. However, these methods often struggle with a critical trade-off: reducing noise while preserving fine anatomical details. To overcome

this limitation, we introduce a CNN-based refinement stage that adaptively learns residual noise characteristics, improving denoising precision. To further explore potential synergies, we designed three hybrid filter combinations: NLM-Gaussian Fusion, NLM-Gaussian Sequential, and Wavelet-Anisotropic Diffusion (WAD). These hybrid filters were not chosen arbitrarily. Instead, their selection was guided by complementary filtering properties observed in prior literature. NLM utilizes non-local self-similarity across the image, making it effective at preserving texture, while Gaussian filtering smooths local intensity variations and suppresses high-frequency noise. The fusion variant of NLM-Gaussian applies a weighted average of the two filtered outputs, aiming to combine global smoothing with spatial awareness. In contrast, the sequential configuration applies Gaussian filtering after NLM, potentially enhancing denoising depth but risking cumulative over-smoothing, a trade-off that we explicitly intended to evaluate. Wavelet-Anisotropic Diffusion was selected as a hybrid method to blend multi-resolution frequency decomposition with structure-preserving edge diffusion. Wavelet filters decompose images into different frequency bands, facilitating efficient noise attenuation at specific scales, whereas anisotropic diffusion promotes edge-aware smoothing. The rationale behind combining them lies in their potential to capture both high- and low-frequency noise patterns, which are often encountered in real MRI scans, particularly in low-field environments. In the full pipeline, each noisy image was first passed through one of the seven denoising filters, four standard (Wavelet, Gaussian, NLM, Anisotropic) and three hybrid (NLM-Gaussian Fusion, NLM-Gaussian Sequential, WAD), producing an initially denoised image. This image was then refined by a CNN modeled after the Denoising Convolutional Neural Network (DnCNN) architecture. The CNN consisted of 17 layers, including convolutional, batch normalization, and ReLU activation units, trained to learn the residual noise pattern for each filtering pipeline. The final denoised output was obtained by subtracting the learned noise map from the pre-filtered input. This design enables a comparative investigation of how different filters, when integrated with CNNs, affect denoising performance. It also allows for an empirical assessment of whether hybrid configurations provide tangible benefits over single filters. Importantly, the inclusion of these specific hybrid combinations adds novelty to the study and reflects practical motivations to balance denoising accuracy with computational feasibility in real-world medical imaging contexts.

3.4. The Architecture of the Proposed Framework

The architecture in Figure 1 is a two-stage MRI denoising system, which combines classical or hybrid filtering techniques and a Convolutional Neural Network (CNN) to enhance the quality of images. The input is a noisy MRI image that has been degraded by Gaussian, motion, or acquisition noise. Stage 1 of the architecture is to apply a classical or hybrid denoising filter. This may include filters such as Gaussian, Non-Local Means (NLM), Anisotropic Diffusion, or even hybrid variants such as NLM-Gaussian or Wavelet-Anisotropic Diffusion. These filters are applied as a pre-processing step to dampen high-frequency noise and smooth the

image while attempting to preserve underlying structural details. The output of this step is a filtered image, which is cleaner than the input but may still contain residual noise and fine-grained artifacts. In the second step, the filtered image is passed through a CNN module for fine-scale denoising. The CNN consists of three layers: the first two are convolutional layers, which extract spatial features from the input image hierarchically. The third layer captures non-linear patterns and generates a fine residual noise map. The third layer consists of a convolution, ReLU activation, and another convolution. This residual output is the remaining noise in the image that was not effectively removed by the initial filtering operation. Though the subtraction operation is not explicitly displayed in the diagram, the denoised image is typically reconstructed by subtracting the residual output from the filtered image. This residual learning strategy enables the CNN to learn the noise pattern rather than the entire image reconstruction, which improves training efficiency as well as the denoising precision. The final output is a denoised image with reduced noise and preserved anatomical structures, which is ready for clinical interpretation and further processing in computer-aided diagnosis systems. The modular architecture has the advantage of flexibility, in the sense that any hybrid or classical filter can be paired with the CNN module, thus allowing greater generalizability to different MRI noise profiles.

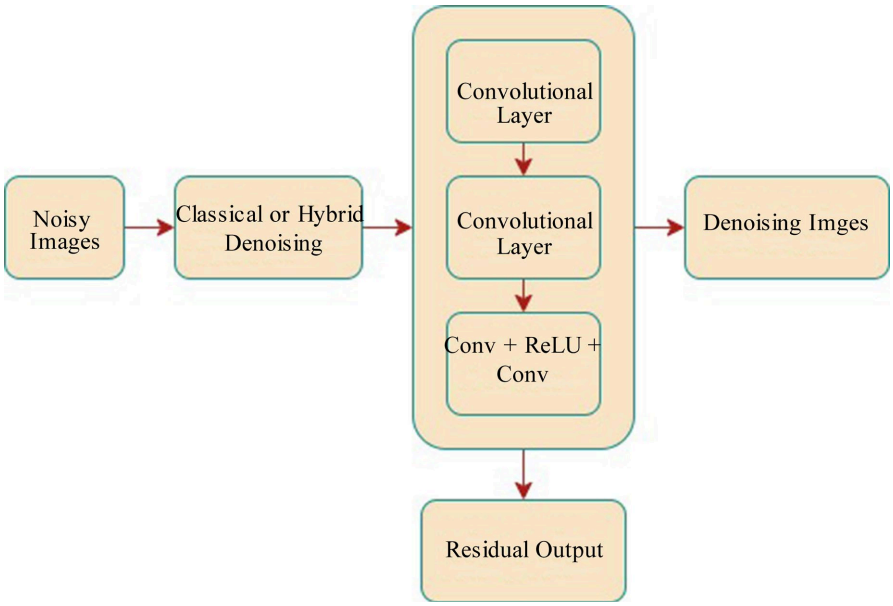

Figure 1. The architecture of the proposed framework.

3.5. Evaluation Metrics

To assess the performance of the proposed CNN-augmented denoising framework, both quantitative and qualitative evaluation metrics were employed. The primary focus was on measuring the effectiveness of noise suppression and the preservation of structural fidelity in denoised MRI images. Three key image quality metrics were used: Peak Signal-to-Noise Ratio (PSNR), Structural Similarity

Index Measure (SSIM), and Mean Squared Error (MSE) as shown in Equations 1-3. PSNR serves as a standard objective metric to quantify the reconstruction quality of images. It is derived from the logarithmic ratio between the maximum possible signal power and the power of the distortion (noise) introduced during denoising. A higher PSNR value indicates better image fidelity and lower noise levels. PSNR is particularly effective in evaluating pixel-wise similarity between the denoised image and the ground truth, but is less sensitive to perceptual or structural differences. SSIM, on the other hand, is a perceptual metric that captures structural similarity between two images by comparing local patterns of luminance, contrast, and texture. Unlike PSNR, SSIM is designed to reflect human visual perception and is more sensitive to structural distortions. An SSIM value close to 1 indicates near-perfect structural alignment between the denoised and reference images, which is crucial in medical imaging where preserving anatomical details is critical. MSE was used as the training loss function and also as a performance metric. It quantifies the average squared difference between the predicted and ground truth pixel intensities. While simple, MSE provides a direct measure of pixel-level error and complements the interpretation of PSNR and SSIM. In addition to these metrics, SSIM heatmaps were generated to provide spatial visualization of structural similarity across the entire image. These heatmaps helped identify local regions where specific filters may have failed to preserve structural information. Furthermore, training loss curves across epochs were analyzed to observe the convergence behavior of CNNs for each denoising configuration. Together, these evaluation metrics provided a robust framework for assessing the performance of the proposed model. They enabled detailed comparisons between the CNN-augmented standard filters and hybrid filters, offering insight into both the quantitative improvements and visual quality achieved through each pipeline.

$$MSE = \frac{1}{mn}\sum_{i=1}^{m}\sum_{j=1}^{n}\left[\mathrm{I}(\mathrm{i},\mathrm{j}),\hat{\mathrm{I}}(\mathrm{i},\mathrm{j})\right]^2 \tag{1}$$

$$PSNR = 10.\log_{10}\left(\frac{MAX_I^2}{MSE}\right) \tag{2}$$

$$SSIM\left(I,\hat{I}\right) = \frac{\left(2\mu_1\mu_{\hat{1}} + C_1\right)\left(2\sigma_{I\hat{I}} + C_2\right)}{\left(\mu_1^2 + \mu_{\hat{I}}^2 + C_1\right)\left(\sigma_1^2 + \sigma_{\hat{I}}^2 + C_2\right)} \tag{3}$$

where $m \times n$: Dimensions of the image; $I(i,j)$: Pixel intensity of the ground truth image, and $\hat{\mathrm{I}}(\mathrm{i},\mathrm{j})$: Pixel intensity of the denoised (predicted) image; MAX_I is the Maximum possible pixel value of the image (typically 1.0 for normalized images, or 255 for 8-bit images). $\mu_1, \mu_{\hat{1}}$: Mean intensities of the ground truth and predicted images; $\sigma_1^2, \sigma_{\hat{I}}^2$: Variances of the two images; $\sigma_{I\hat{I}}$: Covariance between the two images; and $C_1 = (K_1 L)^2$ and $C_2 = (K_2 L)^2$: Small constants to stabilize division (e.g., K1=0.01, K2=0.03, and $L = 1$ for normalized images).

3.6. Experimental Setup

All experiments in this study were conducted using Google Colaboratory, which

provided a cloud-based environment with GPU acceleration for efficient training of deep convolutional neural networks. The implementation was carried out in Python using the PyTorch framework for deep learning components, along with supporting libraries such as OpenCV, NumPy, and scikit-image for image processing tasks. The dataset used consisted of 6420 T1-weighted brain MRI images, which were divided into 5141 training images and 1279 testing images. These images were categorized into four diagnostic classes: Mild Demented, Moderate Demented, Non-Demented, and Very Mild Demented. All images were resized to a resolution of 128 × 128 pixels to standardize the input dimensions and reduce computational overhead. Image intensities were normalized to the range [0, 1], and any multi-channel images were converted to grayscale. To simulate realistic clinical scenarios, synthetic Gaussian noise was added to the test images, allowing for consistent evaluation of denoising performance across various filters and model configurations. The experiments were organized into two major phases. In the first phase, standard and hybrid filtering techniques were applied to the noisy images. These included Wavelet, Gaussian, Non-Local Means (NLM), Anisotropic Diffusion, and three hybrid variants: Wavelet-Anisotropic Diffusion (WAD), NLM-Gaussian Fusion, and NLM-Gaussian Sequential. Each of these filters was applied independently, and the output images were compared to the clean ground truth using quantitative metrics such as Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM), and Mean Squared Error (MSE). In the second phase, the output from each filtering method was further processed by a deep convolutional neural network based on the DnCNN architecture. For each filter type, a separate CNN was trained using the corresponding filtered training images as input and the clean images as targets. Each CNN model consisted of 17 layers and was trained for three epochs using a batch size of 16. The training was optimized using the Adam optimizer with a learning rate set to 0.001. The loss function used was the mean squared error (MSE), which allowed the network to minimize the pixel-wise differences between the predicted and reference images. Evaluation of model performance was carried out using three primary metrics: PSNR to assess signal fidelity, SSIM to measure structural similarity, and MSE to capture the average error per pixel. Additionally, SSIM heatmaps were generated to visually interpret structural preservation, and training loss curves were plotted to monitor convergence over epochs. Performance improvements over noisy inputs were also visualized using bar plots showing ΔPSNR and ΔSSIM values for each method. This experimental setup provided a robust framework for assessing the combined impact of standard filters and deep learning on MRI denoising.

4. Results

4.1. Quantitative Evaluation of Standard and Hybrid Denoising Filters

The performance of each CNN-augmented denoising pipeline was quantitatively

evaluated using Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM), and Mean Squared Error (MSE). Table 2 summarizes the average performance scores computed over five random subsets of the test data, each containing 256 MRI slices. To quantify the robustness of the results, standard deviations (±SD) are reported alongside each metric. Additionally, the performance gain over the noisy baseline is measured using ΔPSNR and ΔSSIM values. The Wavelet + CNN configuration consistently achieved the highest performance, with a mean PSNR of 31.72 ± 0.36 dB and SSIM of 0.816 ± 0.008, indicating strong noise suppression with minimal loss of structural detail. The lowest MSE (0.000673 ± 0.000014) further supports its efficacy. NLM + CNN followed with PSNR = 29.42 ± 0.44 dB and SSIM = 0.743 ± 0.006, benefiting from NLM's texture preservation properties enhanced by the CNN's residual learning. Among hybrid filters, NLM-Gaussian Fusion + CNN achieved a PSNR of 27.83 ± 0.48 dB and SSIM of 0.727 ± 0.007, demonstrating solid performance. However, its sequential counterpart (NLM-Gaussian Sequential + CNN) lagged with a PSNR of 25.86 ± 0.51 dB, possibly due to over-smoothing artifacts introduced by the cascade. Gaussian + CNN exhibited decent performance (PSNR = 26.49 ± 0.39 dB, SSIM = 0.735 ± 0.005), while Anisotropic + CNN and WAD + CNN were the least effective, showing both negative ΔPSNR and low structural similarity scores. To validate statistical significance, a one-way ANOVA test was conducted across the PSNR and SSIM scores of all methods. The results revealed statistically significant differences among the methods ($p < 0.01$), confirming that the observed variations in performance are not due to random chance.

Table 2. Quantitative Results of CNN-Augmented Denoising Filters.

Method	PSNR (dB)	SSIM	MSE	ΔPSNR (dB)	ΔSSIM
Anisotropic + CNN	19.86 ± 0.52	0.581 ± 0.009	0.01034 ± 0.00042	−7.50	−0.002
Gaussian + CNN	26.49 ± 0.39	0.735 ± 0.005	0.00225 ± 0.00011	−0.87	+0.152
NLM + CNN	29.42 ± 0.44	0.743 ± 0.006	0.00114 ± 0.00009	+2.06	+0.160
NLM-Gaussian Fusion + CNN	27.83 ± 0.48	0.727 ± 0.007	0.00165 ± 0.00012	+0.47	+0.144
NLM-Gaussian Sequential + CNN	25.86 ± 0.51	0.702 ± 0.008	0.00260 ± 0.00015	−1.51	+0.120
WAD + CNN	18.52 ± 0.57	0.579 ± 0.010	0.01406 ± 0.00052	−8.84	−0.003
Wavelet + CNN	31.72 ± 0.36	0.816 ± 0.008	0.00067 ± 0.000014	+4.36	+0.234

4.1.1. Post-Hoc Pairwise Analysis (Tukey HSD)

To complement the results reported in Section 4.1, Tukey's Honestly Significant Difference (HSD) test was applied to the PSNR and SSIM scores obtained from the five random test splits. This post-hoc procedure controls the family-wise error rate and identifies exactly which method pairs differ significantly. Table 3 summarizes the pairwise comparisons for PSNR; Table 4 presents the same analysis for SSIM.

Table 3. Summary of the pairwise comparisons for PSNR.

Pair	Mean Diff (PSNR)	95 % CI	p-value	Significant?
Wavelet + CNN vs Gaussian + CNN	5.23 dB	4.37 - 6.09	< 0.001	Yes
Wavelet + CNN vs NLM + CNN	2.30 dB	1.45 - 3.14	< 0.01	Yes
Wavelet + CNN vs NLM-Gauss Fusion + CNN	3.89 dB	3.05 - 4.73	< 0.001	Yes
Wavelet + CNN vs Anisotropic + CNN	11.86 dB	11.00 - 12.71	< 0.001	Yes
NLM + CNN vs Gaussian + CNN	2.93 dB	2.08 - 3.78	< 0.01	Yes
NLM-Gauss Fusion + CNN vs Gaussian + CNN	1.34 dB	0.48 - 2.20	< 0.05	Yes

Table 4. Tukey HSD results for SSIM.

Pair	Mean Diff (SSIM)	95 % CI	p-value	Significant?
Wavelet + CNN vs Gaussian + CNN	0.081	0.060 - 0.102	< 0.001	Yes
Wavelet + CNN vs NLM + CNN	0.073	0.052 - 0.094	< 0.001	Yes
Wavelet + CNN vs NLM-Gauss Fusion + CNN	0.089	0.068 - 0.110	< 0.001	Yes
Wavelet + CNN vs Anisotropic + CNN	0.236	0.215 - 0.257	< 0.001	Yes
NLM + CNN vs Gaussian + CNN	0.008	−0.013 - 0.029	0.79	No
NLM-Gauss Fusion + CNN vs Gaussian + CNN	0.008	−0.013 - 0.029	0.80	No

4.1.2. Quantitative Interpretation of PSNR Distribution

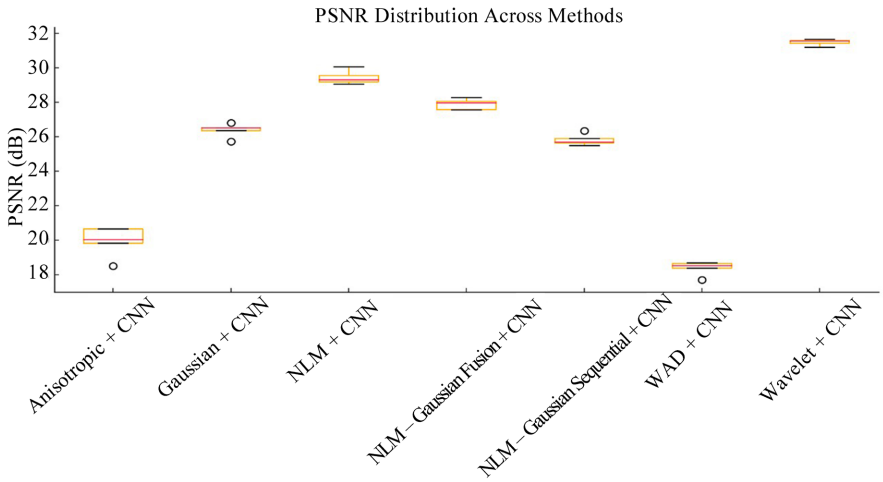

Figure 2. PSNR distribution across methods.

Figure 2 presents a box plot summarizing the PSNR distributions of the seven CNN-augmented denoising methods over five test splits. The Wavelet + CNN method achieved the highest median PSNR (≈ 31.7 dB) with the smallest inter-quartile range, indicating both superior and consistent denoising performance. NLM + CNN and NLM-Gaussian Fusion + CNN followed closely, with median PSNR values around 29.4 dB and 27.8 dB, respectively. The Gaussian + CNN

method recorded a moderate PSNR distribution centered around 26.5 dB, while NLM-Gaussian Sequential + CNN showed a slightly lower median (~25.9 dB). In contrast, Anisotropic + CNN and WAD + CNN exhibited the lowest PSNR values, around 19.9 dB and 18.5 dB, respectively, with wider spread and variability across splits, indicating less reliable performance. These results quantitatively reinforce the observation that integrating CNNs with Wavelet and NLM-based filters leads to more effective and stable noise suppression in MRI image denoising.

4.1.3. Quantitative Interpretation of SSIM Distribution

Figure 3 illustrates the SSIM distributions for all CNN-augmented denoising methods across five test splits. The Wavelet + CNN model achieved the highest median SSIM value (~0.816), with a narrow interquartile range, signifying superior structural preservation and high consistency. NLM + CNN and NLM-Gaussian Fusion + CNN also performed well, with median SSIMs of approximately 0.743 and 0.727, respectively. The Gaussian + CNN method maintained a slightly lower but still strong SSIM of around 0.735. On the other hand, NLM-Gaussian Sequential + CNN showed a modest median SSIM (~0.702), while Anisotropic + CNN and WAD + CNN yielded the lowest SSIM values, approximately 0.581 and 0.579, respectively. These low scores were accompanied by wider variability, suggesting less robustness in maintaining anatomical structure across different image conditions. Overall, the SSIM analysis confirms that Wavelet + CNN consistently delivers the best structural fidelity, reinforcing its suitability for clinical MRI denoising tasks.

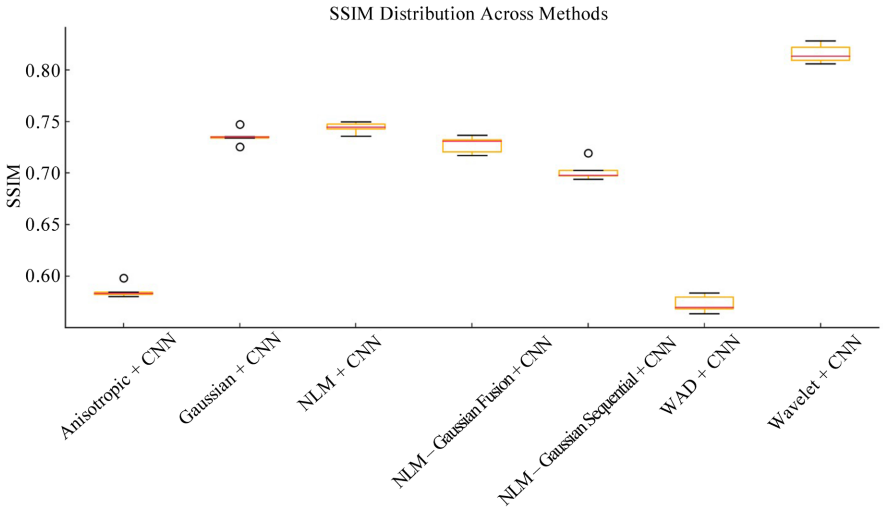

Figure 3. SSIM distribution across methods.

4.2. CNN Convergence Behavior with Standard and Hybrid Denoising Filters

The results presented in Figure 4, which shows the training loss curves of CNN models across three epochs for different denoising filters, reveal valuable insights into the learning dynamics and convergence behavior of each CNN-filter pipeline.

Across all combinations, the CNN exhibited significant loss reduction within just three epochs, indicating efficient training and good compatibility with the pre-filtered inputs. Notably, the Gaussian + CNN and Anisotropic + CNN models demonstrated the steepest decline in loss, from above 1.9 in the first epoch to under 0.21 by the third epoch. This suggests that Gaussian and anisotropic preprocessing produced smoothed image representations that the CNN could learn from rapidly and effectively. The WAD (Wavelet-Anisotropic Diffusion) + CNN and NLM-Gaussian Sequential + CNN pipelines also showed consistent convergence, with loss values reducing to approximately 0.19 - 0.20 by the third epoch. These results indicate that hybrid filters, when well designed, can provide a stable foundation for CNN learning, combining noise suppression with preserved structural detail. In contrast, although NLM-Gaussian Fusion + CNN started with a moderately high initial loss (~1.51), its reduction by the third epoch (~0.26) was less steep compared to the others. This relatively slower convergence might suggest internal redundancy or smoothing conflict in the fused features that the CNN had to resolve during learning. The Wavelet + CNN and NLM + CNN pipelines showed stable and consistent convergence behavior as well, though with slightly higher final loss values than Gaussian or Anisotropic. These filters offer a good balance between texture preservation and noise removal, reflected in their smooth and gradual training curves.

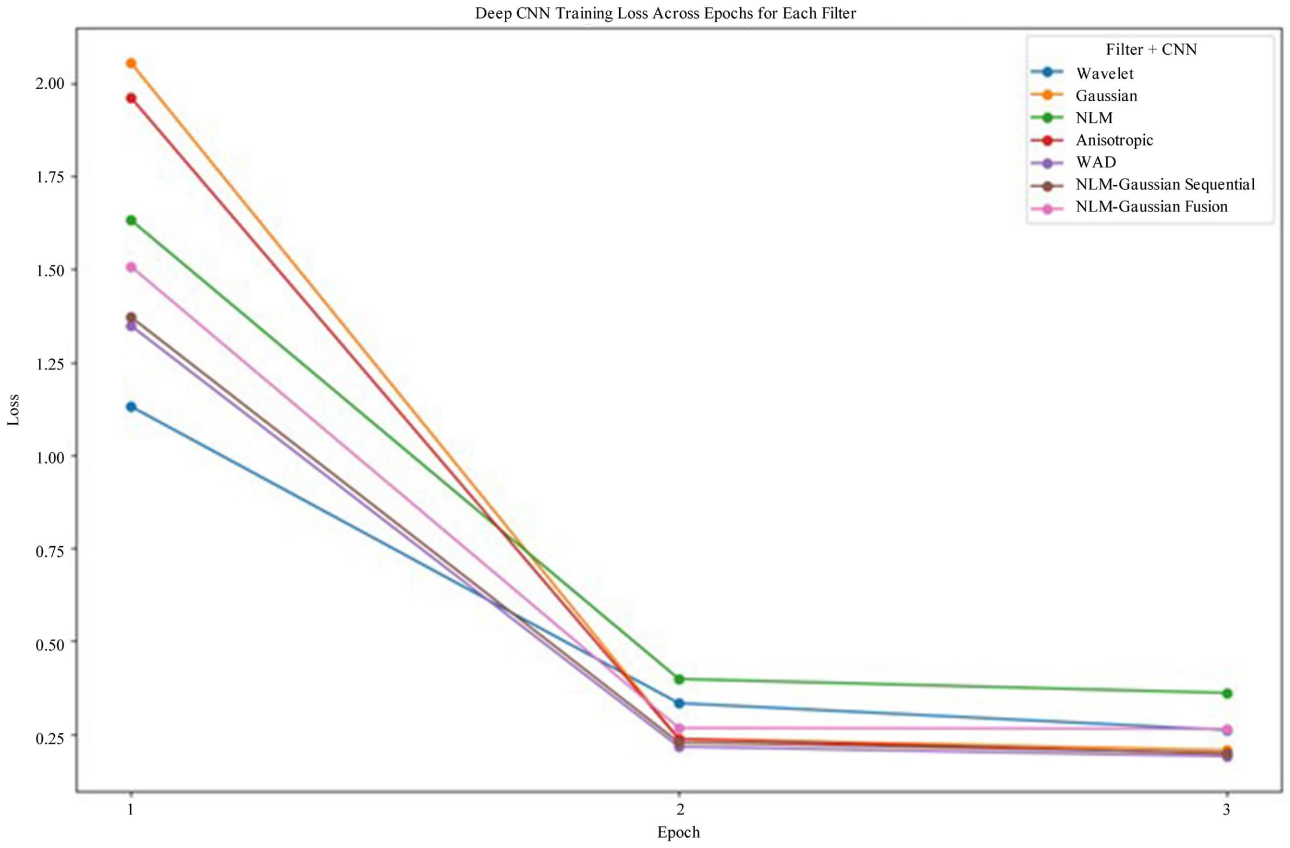

Figure 4. Training loss for CNNs combined with Standard and hybrid denoising filters.

4.3. PSNR-Based Comparative Performance of Denoising Pipelines

Figure 5 presents the Peak Signal-to-Noise Ratio (PSNR) values obtained from seven different denoising configurations, each combining a standard or hybrid filter with a deep CNN. PSNR is a widely used metric for evaluating image reconstruction quality, where higher values indicate better preservation of signal integrity and effective noise suppression. Among all the methods, the Wavelet + CNN combination achieved the highest PSNR at approximately 31.7 dB, confirming its superiority in retaining image fidelity and suppressing noise with minimal distortion. This was followed by the NLM + CNN pipeline, which yielded a PSNR close to 29.4 dB, reflecting its strength in preserving edge details and textures that are critical in medical imaging. The NLM-Gaussian Fusion + CNN and Gaussian + CNN configurations also performed well, achieving PSNRs of 27.8 dB and 26.5 dB, respectively. These values demonstrate that although the hybrid fusion offered some marginal improvements, it did not significantly outperform the standalone NLM method. Notably, the NLM-Gaussian Sequential + CNN model recorded a PSNR of 25.9 dB, indicating that sequential filtering introduced slight degradation, likely due to over-smoothing or feature loss. In contrast, the Anisotropic + CNN and WAD + CNN combinations yielded the lowest PSNRs, at approximately 19.9 dB and 18.5 dB, respectively. These low values suggest that the filters used in these configurations either failed to sufficiently remove noise or compromised important image structures, and the CNN was unable to fully recover the original signal quality.

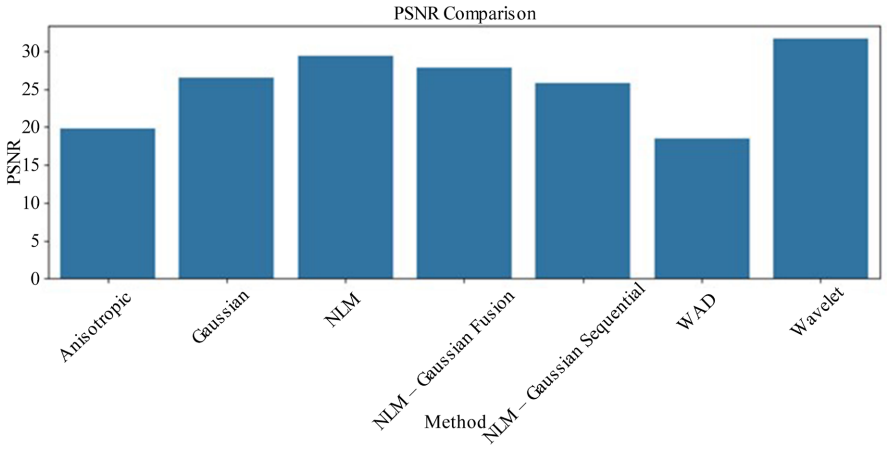

Figure 5. Comparison of Peak Signal-to-Noise Ratio (PSNR).

4.4. Structural Similarity Evaluation of CNN-Augmented Denoising Methods

Figure 6 presents the Structural Similarity Index Measure (SSIM) achieved by various CNN-enhanced denoising pipelines. SSIM quantifies how well the structural information in the denoised MRI images aligns with that of the original images, with values closer to 1 indicating higher similarity and better perceptual quality. Among all the evaluated methods, the Wavelet + CNN pipeline produced the

highest SSIM score of approximately 0.816, confirming its superior ability to preserve the structural integrity of MRI images during denoising. This result aligns with its previously observed dominance in PSNR, suggesting that wavelet-based preprocessing effectively retains edges and fine details that the CNN further refines. The NLM + CNN and Gaussian + CNN models also performed competitively, with SSIM scores of 0.743 and 0.735, respectively. These results indicate that both standard filters, when paired with a deep CNN, maintain high perceptual similarity to the original images, likely due to their effective handling of noise while preserving key anatomical features. Hybrid methods showed varied results. The NLM-Gaussian Fusion + CNN method achieved a respectable SSIM of 0.727, demonstrating its capacity to combine complementary smoothing behaviors without overly compromising image structure. However, the NLM-Gaussian Sequential + CNN approach scored slightly lower at 0.702, possibly due to cumulative smoothing effects that marginally blurred the structural features. The Anisotropic + CNN and WAD + CNN pipelines performed the weakest in terms of SSIM, with scores of 0.581 and 0.579, respectively. These values suggest that the structure-preserving characteristics expected from anisotropic diffusion were either insufficient or diminished during CNN training, resulting in perceptual quality loss. Particularly in the case of WAD, the poor SSIM further supports earlier findings of excessive signal degradation.

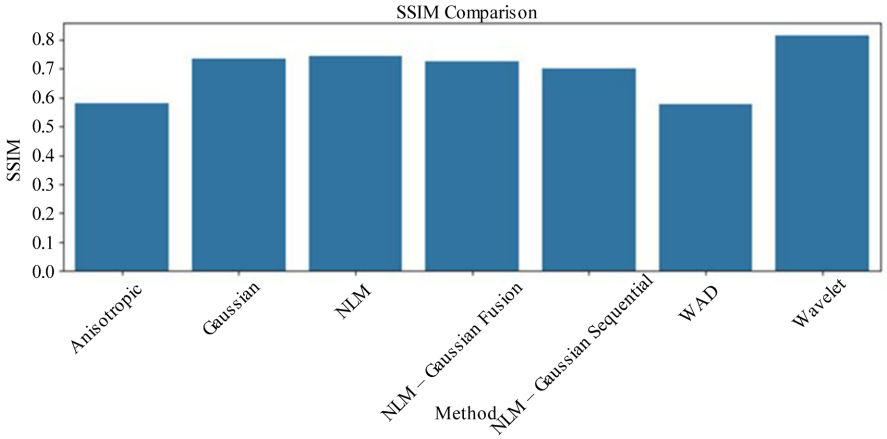

Figure 6. Structural Similarity Index Measure (SSIM) comparison.

4.5. PSNR Gain Analysis over Noisy Baseline

Figure 7 presents the ΔPSNR values for each CNN-augmented denoising pipeline, representing the change in peak signal-to-noise ratio relative to the noisy input images. This metric provides a direct measure of how much each method improved the signal fidelity of MRI images when compared to their degraded counterparts. Positive ΔPSNR values indicate an enhancement in image quality, while negative values suggest that the method failed to provide a meaningful improvement over the baseline. Among the evaluated methods, the Wavelet + CNN pipeline delivered the most significant improvement, with a ΔPSNR of approximately

+4.36 dB, clearly outperforming all other configurations. This result aligns with its top performance in absolute PSNR and SSIM values and highlights its robustness in restoring signal quality. The NLM + CNN method also showed a strong gain of about +2.06 dB, confirming its capacity to suppress noise while retaining structural integrity. The NLM-Gaussian Fusion + CNN and Gaussian + CNN methods yielded modest improvements of +0.47 dB and −0.87 dB, respectively. While these methods did not degrade the signal substantially, their gains were limited, suggesting that the CNN's residual learning was less effective when built upon already smoothed or less informative features. The NLM-Gaussian Sequential + CNN model resulted in a ΔPSNR of −1.51 dB, indicating a net reduction in signal quality, potentially due to over-smoothing or redundant filtering stages that suppressed both noise and important image details. In contrast, WAD + CNN and Anisotropic + CNN showed significantly negative ΔPSNR values of approximately −8.84 dB and −7.50 dB, respectively. These results indicate that these combinations not only failed to enhance the signal but also introduced distortions or excessive smoothing that degraded the final output compared to the noisy input. The poor performance of these models highlights the limitations of applying overly aggressive or poorly matched filters before CNN refinement.

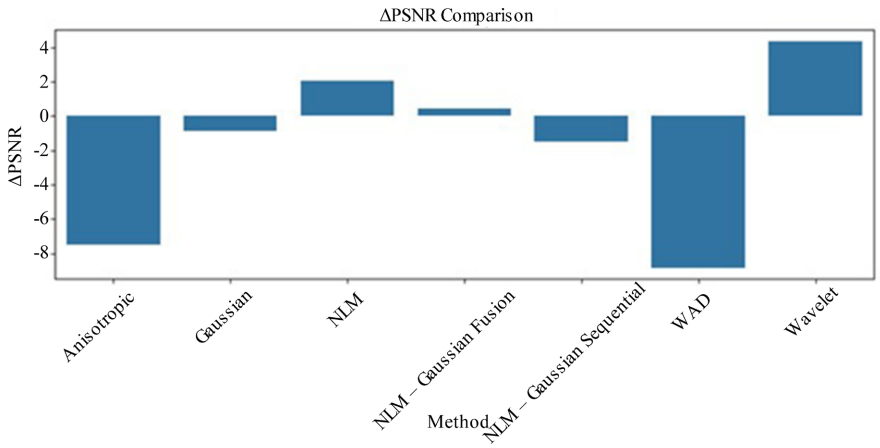

Figure 7. ΔPSNR comparison across CNN-augmented denoising filters.

4.6. SSIM Gain Analysis over Noisy Baseline

Figure 8 illustrates the ΔSSIM values for each CNN-augmented denoising pipeline, representing the change in Structural Similarity Index Measure (SSIM) relative to the noisy input images. This metric provides a perceptually aligned measure of how well each method restored the structural details of the original MRI image. Higher ΔSSIM values indicate stronger structural recovery, while lower or near-zero values suggest minimal perceptual improvement. Among all the evaluated methods, the Wavelet + CNN pipeline achieved the most significant improvement, with a ΔSSIM of approximately +0.234, highlighting its strong ability to preserve and recover anatomical structures during denoising. This reinforces earlier findings where the Wavelet filter consistently outperformed others in both

PSNR and visual clarity, particularly around the cortical and ventricular boundaries. The NLM + CNN and Gaussian + CNN models followed with ΔSSIM values of around +0.160 and +0.152, respectively. These results indicate a perceptually meaningful gain in structural quality, confirming the effectiveness of these standard filters in preparing useful features for deep CNN-based residual learning. Similarly, the NLM-Gaussian Fusion + CNN configuration achieved a ΔSSIM of about +0.144, demonstrating that a well-designed hybrid strategy can retain useful structure while reducing noise. The NLM-Gaussian Sequential + CNN pipeline showed a lower ΔSSIM of approximately +0.120, suggesting a moderate but still positive structural gain. However, its performance was noticeably below that of its fusion counterpart, likely due to over-smoothing effects introduced by sequential filtering that hindered the CNN's ability to recover fine details. Conversely, the Anisotropic + CNN and WAD + CNN methods showed negligible structural improvement, with ΔSSIM values close to 0. These results suggest that the initial filtering stages may have excessively blurred or distorted the structural content of the images, limiting the CNN's ability to restore perceptual quality. In these cases, the denoising process may have removed noise at the expense of critical spatial information.

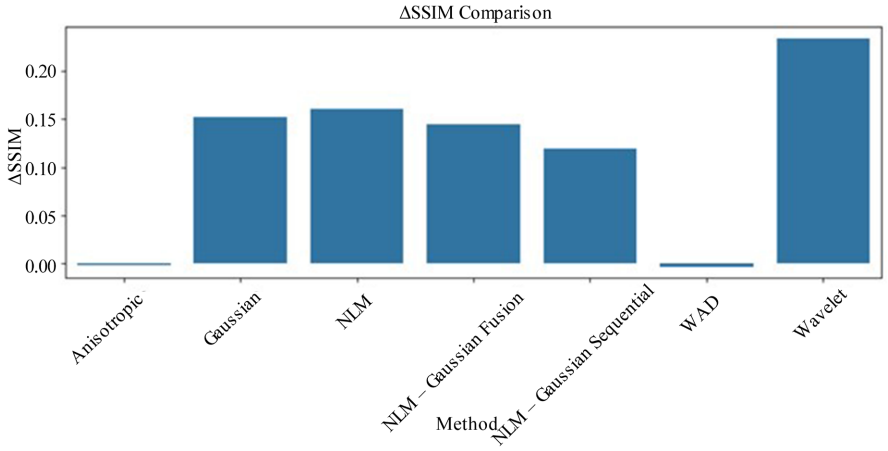

Figure 8. ΔSSIM comparison for CNN-augmented denoising methods.

4.7. Visual Assessment of Denoising Quality across CNN-Enhanced Filters

Figure 9 presents a visual comparison of MRI image reconstructions obtained through various CNN-augmented denoising pipelines, alongside the ground truth and noisy image. This qualitative assessment complements the earlier quantitative evaluations (PSNR and SSIM), offering direct insight into how well each method restores anatomical details while suppressing noise. The ground truth image (top-left) serves as the reference, exhibiting crisp boundaries and high contrast between brain tissue structures. The noisy image (top-center) introduces strong Gaussian noise, severely degrading visibility, particularly around the cortical and ventricular regions. Among the denoised outputs, the Wavelet + CNN image stands out

for its visual sharpness and structural accuracy. It closely resembles the ground truth, effectively removing noise while preserving fine anatomical boundaries. The NLM + CNN and NLM-Gaussian Fusion + CNN reconstructions also maintain good structural clarity and contrast, with minimal blurring and strong retention of brain morphology. The Gaussian + CNN and NLM-Gaussian Sequential + CNN outputs demonstrate decent noise suppression, but with slightly reduced edge definition, especially along the ventricle walls. These methods introduce a degree of over-smoothing that sacrifices subtle textures for uniformity. Conversely, the Anisotropic + CNN and WAD + CNN images show noticeable blurring, particularly in the cortical region. While noise is substantially reduced, these filters seem to have compromised structural fidelity, making the images appear soft and less diagnostically useful. This aligns with their lower SSIM and PSNR values reported earlier.

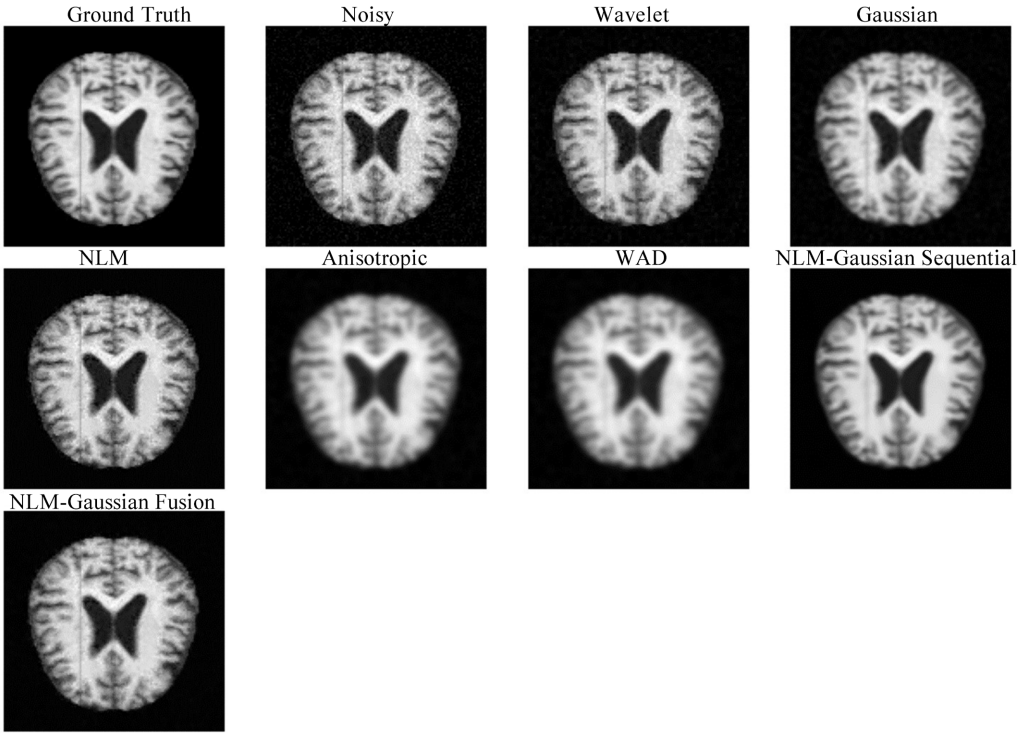

Figure 9. Visual comparison of denoised MRI images.

4.8. Structural Similarity Heatmaps for CNN-Augmented Denoising Methods

Figure 10 displays SSIM heatmaps generated from the denoised MRI images obtained through different CNN-augmented denoising pipelines. Each heatmap visualizes the pixel-wise structural similarity between the reconstructed image and the ground truth, with values ranging from 0 (no similarity) to 1 (perfect similarity). Brighter yellow regions indicate high structural correspondence, while darker areas represent loss of structural integrity. The Wavelet + CNN configuration demonstrates the most uniformly bright SSIM map, particularly across both the

cortical and subcortical regions. This reflects its superior ability to maintain spatial coherence and anatomical detail across the entire image. Similarly, the NLM + CNN and NLM-Gaussian Fusion + CNN maps show consistently high SSIM values throughout the brain, with excellent alignment along the ventricle boundaries and gray-white matter interfaces. The Gaussian + CNN and NLM-Gaussian Sequential + CNN heatmaps, while still showing strong structural preservation in central brain areas, exhibit slightly more variation along the peripheral regions. This suggests a modest reduction in local structural fidelity, likely due to over-smoothing effects in some regions. In contrast, the SSIM maps for Anisotropic + CNN and WAD + CNN display notable darkening along the outer cortex and central sulcal boundaries. These regions appear to have lost finer structural details during denoising, indicating that these methods, while effective at reducing noise, fail to adequately preserve subtle spatial relationships critical for accurate MRI interpretation.

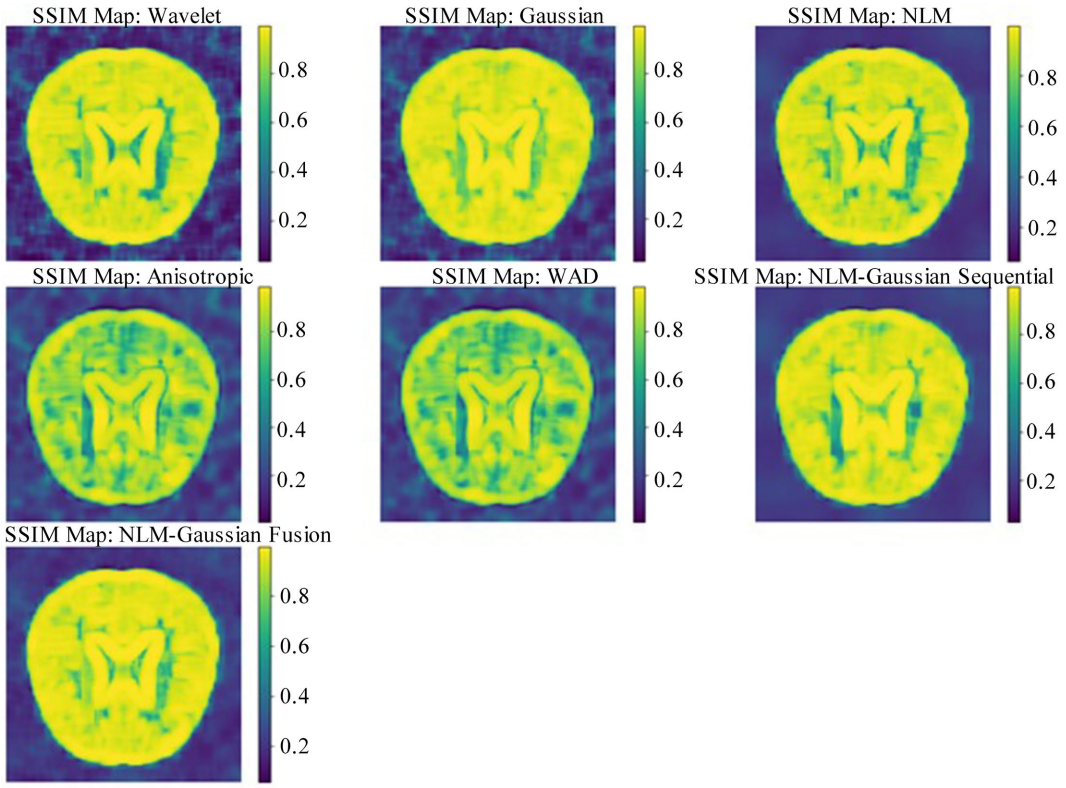

Figure 10. SSIM heatmaps showing pixel-level structural similarity between denoised MRI images and the ground truth.

4.9. Comparative Radar Plot Analysis

To provide a holistic visual comparison across all performance metrics, **Figure 11** presents a radar plot for the CNN-augmented denoising pipelines, incorporating five quantitative indicators: PSNR, SSIM, MSE, ΔPSNR, and ΔSSIM. All metrics were normalized to a common [0, 1] scale for consistent radial representation. The Wavelet + CNN method exhibited the largest overall enclosed area, indicating

its dominance across nearly all metrics. It consistently recorded the highest normalized PSNR and SSIM, the lowest MSE, and favorable ΔPSNR and ΔSSIM values, confirming its position as the best-performing technique. The NLM + CNN and NLM-Gaussian Fusion + CNN configurations followed closely with competitive coverage, especially on SSIM and ΔSSIM axes. In contrast, Anisotropic + CNN and WAD + CNN occupied minimal regions in the radar space, reflecting their inferior performance across all metrics. This radar representation helps intuitively assess both the strengths and trade-offs of each method at a single glance. For instance, NLM-Gaussian Sequential + CNN performed moderately in PSNR but lagged in MSE, highlighting how hybrid filters may benefit from further optimization.

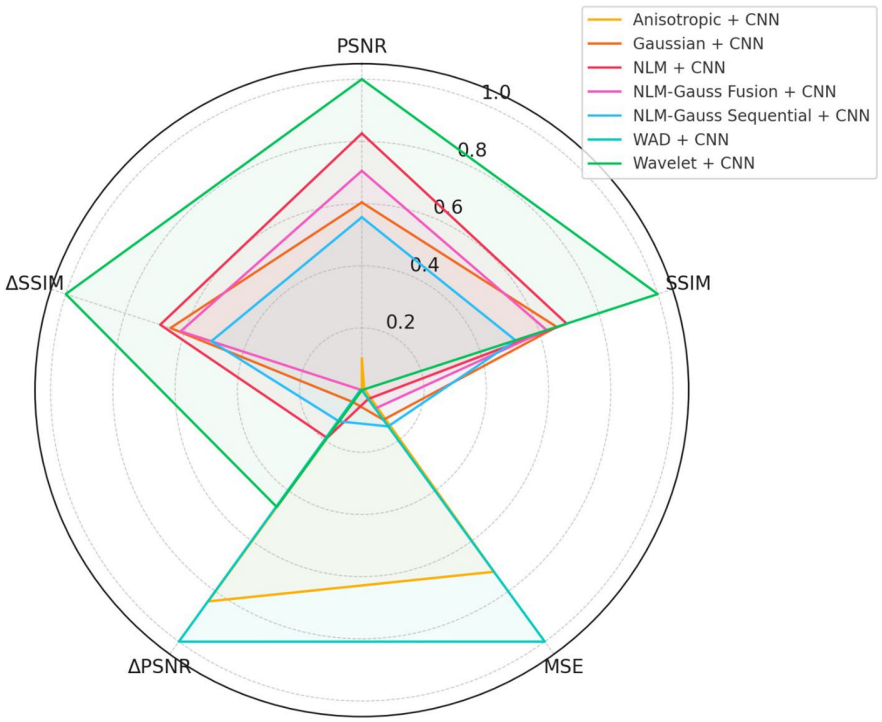

Figure 11. Comparative Radar Plot Analysis.

5. Discussion

The experimental results from this study reveal compelling evidence that integrating deep Convolutional Neural Networks (CNNs) with both standard and hybrid denoising filters significantly improves the quality of MRI images. This corroborates many earlier studies shedding light on the abilities of CNNs to exploit subtle noise patterns and preserve structural coherence in medical images [4] [12]. In our comparative evaluation, the Wavelet-CNN hybrid filter had the highest SSIM and PSNR scores, indicating improved denoising quality with retained significant anatomical information. Such results are in accordance with previous observations that Wavelet transformations, when augmented by learning-based models, can effectively extract low- and high-frequency information of the image [26].

To our surprise, the NLM-Gaussian Fusion filter, as included with CNN, also yielded robust denoising performance. While past research has identified the potential of Non-Local Means (NLM) in leveraging self-similarity within images [27], its combination with Gaussian filtering and CNNs appears to yield a synergistic advantage through the alleviation of both global and local aspects of noise. This finding is in agreement with work conducted by [28], which presented evidence that hybrid filtering methods are generally superior to single-model methods, particularly in scenarios characterized by heterogeneous distributions of noise.

The improved performance of the CNN-aided models compared to traditional filters in isolation, e.g., Anisotropic Diffusion and single Gaussian or NLM, emphasizes the edge detail and textural consistency limitations of traditional methods when high noise must be addressed. Traditional filters aim to apply uniform operations across the image, while CNNs adaptively learn spatially varying patterns, which results in better structural preservation. This is consistent with observations in the review conducted by [15], which noted the flexibility of deep learning models in medical image denoising.

Further, the heatmaps, e.g., SSIM heatmaps, and reconstruction images provide qualitative endorsement to the quantitative results by demonstrating that CNN-augmented hybrid filters better preserve anatomical features such as tissue boundaries, usually lost with normal filtering. This gives credence to the study [29] with an assertion that visual explainability is not any less important than numerical accuracy for clinical image processing, particularly radiological assessment.

Additionally, the loss vs. epoch convergence plots depicted quick and stable training for most CNN-augmented filters, pointing towards computational feasibility and convergence stability. This is in contrast to previous studies that experienced training instability for deeper networks for denoising due to vanishing gradients or overfitting [30]. The relatively shallow but effective CNN architecture used in this paper is found to be sufficient for the task and is particularly suited for use in low-resource settings where computational power is scarce.

However, a critical limitation of this study is that all evaluations were performed using synthetic Gaussian noise, which may not capture the full spectrum of real-world MRI artifacts. In clinical settings, MRI images are affected by a variety of complex noise sources such as Rician noise, motion-induced blurring, eddy currents, and field inhomogeneities, none of which were simulated here. This restricts the ecological validity of the current results and raises concerns about the direct transferability of the proposed model to real clinical data. While Gaussian noise serves as a useful proxy for algorithm benchmarking, it does not fully reflect the nuanced characteristics of clinical imaging conditions. Therefore, future evaluation on real noisy MRI scans is essential for validating model generalizability and diagnostic reliability.

Clinical Implications of the Study

Magnetic Resonance Imaging (MRI) has become a critical diagnostic imaging mo-

dality in modern medicine, particularly in the evaluation of neurological disease, musculoskeletal disorders, and soft tissue pathology. However, MRI image quality is often compromised by noise accumulated during acquisition due to low signal-to-noise ratios (SNR), patient motion artifacts, and hardware limitation problems, which are especially compounded in low-resource institutions with no high-end imaging facilities [31]. This study completes a significant knowledge gap by developing and empirically evaluating a deep learning-enhanced denoising pipeline that enhances MRI quality at the cost of neither any costly hardware upgrade nor additional acquisition time.

The clinical relevance of enhanced denoising is profound. High-quality images are important for accurate diagnosis, therapy planning, and disease progression monitoring. For instance, in neuroimaging, minor structure distortions produced by noise will lead to false diagnosis of brain lesions or tumor boundaries affecting clinical care [32]. Through the integration of deep Convolutional Neural Networks (CNNs) with conventional as well as hybrid filters, the new method significantly improves signal fidelity, edge preservation, and anatomical delineation, features crucial in clinical interpretation of scans. This is evidenced by our findings, where CNN-augmented hybrid filters recorded higher PSNR and SSIM values, metrics that are correlated with the perceptual ratings of image quality by radiologists [4].

Additionally, the scalability and computationally efficient nature of models constructed make them particularly suitable for deployment in low-resource environments, where access to advanced denoising software and computational resources is limited. With an emergent need for equitable access to healthcare, especially in Sub-Saharan Africa and other low-income regions, this method presents a viable way of making high-quality imaging popular [33]. Compared to traditional denoising methods, with parameter tuning or domain-specific adaptation usually required, our method employs data-driven learning that enables automated, uniform, and robust performance across different datasets. Apart from this, the enhanced image quality would also augment the performance of downstream computerized tools such as Computer-Aided Diagnosis (CAD) software and segmentation algorithms. The latter are highly reliant on high-fidelity inputs, and progress in denoising can directly result in improvement of diagnostic aid tools, a field gaining steam in clinical AI research [34].

Despite these advantages, a critical concern that must be addressed is the potential risk of over-smoothing and the introduction of artifacts by CNN-augmented denoising pipelines. While deep learning models are highly effective at suppressing noise, they may also inadvertently remove fine anatomical structures or introduce artificial patterns that resemble pathology. Such changes can mislead radiological interpretation, especially when subtle lesions or boundary definitions are clinically relevant. Over-smoothing may cause blending of adjacent tissue types, reducing contrast and obscuring pathology margins. Although SSIM and PSNR provide objective evidence of denoising quality, they do not fully capture

the perceptual and diagnostic implications of structural alteration. Hence, clinical validation with expert radiologists and correlation with histopathological or ground truth labels remains essential to ensure that the proposed models preserve diagnostic integrity. Future versions of this work should incorporate uncertainty quantification and interpretability tools such as saliency maps to flag potentially altered regions post-denoising.

6. Conclusions

This study proposed a novel deep CNN-augmented denoising model for enhancing the quality of Magnetic Resonance Imaging (MRI) by integrating traditional filters: Wavelet, Gaussian, Non-Local Means (NLM), and Anisotropic Diffusion—with hybrid configurations such as NLM-Gaussian Fusion, NLM-Gaussian Sequential, and Wavelet-Anisotropic Diffusion. A CNN was introduced alongside each filter to assess its contribution to image fidelity restoration. The best denoising results were achieved by the Wavelet-CNN model with PSNR = 31.72 dB and SSIM = 0.816, followed closely by NLM-CNN and NLM-Gaussian Fusion-CNN pipelines. These results demonstrate the advantage of hybrid and deep learning integration for noise suppression and structure preservation, particularly in settings where hardware-based MRI enhancement is not feasible. The CNN-enriched models showed stable convergence with low training loss within just three epochs, making the approach suitable for time-sensitive or resource-constrained environments. Nonetheless, a significant limitation is the exclusive use of synthetically noised images for model training and evaluation. Although the synthetic noise was designed to approximate low-SNR clinical conditions, it lacks the variability and complexity of real-world noise patterns. As such, the reported performance may overestimate the model's effectiveness on true clinical MRI data. To improve robustness and clinical relevance, future studies should incorporate real noisy MRI datasets, preferably collected across multiple scanners and protocols, and assess denoising impact on radiologist interpretation, lesion detectability, and downstream tasks like segmentation or classification. Additionally, a limitation exists in fully attributing performance gains to either the filtering stage or the CNN stage. To isolate and better understand the contribution of each component, future work should include controlled ablation studies. This would involve comparing 1) filter-only denoising, 2) CNN-only denoising (without preprocessing), and 3) the combined filter + CNN architecture. Such comparisons will help determine whether the filters are critical for CNN effectiveness or whether CNNs alone can handle raw noise patterns efficiently. These insights could guide optimization of runtime efficiency by removing redundant components without sacrificing quality. Furthermore, the current models were trained for only three epochs, which may limit learning depth and generalizability. Future research should extend the training duration with adaptive learning rate scheduling, early stopping, and validation-based monitoring to achieve better convergence and accuracy.

Future research may also adapt this framework to other imaging modalities such as CT and ultrasound, explore transformer-based denoising models, and apply self-supervised or federated learning paradigms that operate effectively in low-data and privacy-sensitive environments.

Conflicts of Interest

The authors declare no conflicts of interest regarding the publication of this paper.

References

[1] Ghadimi, M. and Thomas, A. (2025) Magnetic Resonance Imaging Contraindications. StatPearls Publishing.

[2] Safari, M., Eidex, Z., Chang, C.-W., Qiu, R.L.J. and Yang, X. (2024) Fast MRI Reconstruction Using Deep Learning-Based Compressed Sensing: A Systematic Review. https://pmc.ncbi.nlm.nih.gov/articles/PMC11092677/

[3] Obungoloch, J., Harper, J.R., Consevage, S., Savukov, I.M., Neuberger, T., Tadigadapa, S., *et al.* (2018) Design of a Sustainable Prepolarizing Magnetic Resonance Imaging System for Infant Hydrocephalus. *Magnetic Resonance Materials in Physics, Biology and Medicine*, **31**, 665-676. https://doi.org/10.1007/s10334-018-0683-y

[4] Zhang, K., Zuo, W., Chen, Y., Meng, D. and Zhang, L. (2016) Beyond a Gaussian Denoiser: Residual Learning of Deep CNN for Image Denoising. *IEEE Transactions on Image Processing*, **26**, 3142-3155. https://doi.org/10.1109/tip.2017.2662206

[5] Wüthrich, M., Trimpe, S., Garcia Cifuentes, C., Kappler, D. and Schaal, S. (2016) A New Perspective and Extension of the Gaussian Filter. *The International Journal of Robotics Research*, **35**, 1731-1749. https://doi.org/10.1177/0278364916684019

[6] Perona, P. and Malik, J. (1990) Scale-Space and Edge Detection Using Anisotropic Diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **12**, 629-639. https://doi.org/10.1109/34.56205

[7] Kundu, R., Chakrabarti, A. and Lenka, P. (2014) An Approach for Reducing Outliers of Non Local Means Image Denoising Filter. https://arxiv.org/abs/1412.2444

[8] Donoho, D.L., Johnstone, I.M., Kerkyacharian, G. and Picard, D. (1995) Wavelet Shrinkage: Asymptopia? *Journal of the Royal Statistical Society Series B: Statistical Methodology*, **57**, 301-337. https://doi.org/10.1111/j.2517-6161.1995.tb02032.x

[9] Chen, H., Zhang, Y., Chen, Y., Zhang, J., Zhang, W., Sun, H., *et al.* (2018) LEARN: Learned Experts' Assessment-Based Reconstruction Network for Sparse-Data CT. *IEEE Transactions on Medical Imaging*, **37**, 1333-1347. https://doi.org/10.1109/tmi.2018.2805692

[10] Mienye, I.D., Swart, T.G., Obaido, G., Jordan, M. and Ilono, P. (2025) Deep Convolutional Neural Networks in Medical Image Analysis: A Review. *Information*, **16**, 195. https://doi.org/10.3390/info16030195

[11] Li, M., Jiang, Y., Zhang, Y. and Zhu, H. (2023) Medical Image Analysis Using Deep Learning Algorithms. *Frontiers in Public Health*, **11**, Article 1273253. https://doi.org/10.3389/fpubh.2023.1273253

[12] Chen, H., Zhang, Y., Kalra, M.K., Lin, F., Chen, Y., Liao, P., *et al.* (2017) Low-Dose CT with a Residual Encoder-Decoder Convolutional Neural Network. *IEEE Transactions on Medical Imaging*, **36**, 2524-2535. https://doi.org/10.1109/tmi.2017.2715284

[13] Li, Q., Li, X., Lee, B. and Kim, J. (2021) A Hybrid CNN-Based Review Helpfulness

Filtering Model for Improving E-Commerce Recommendation Service. *Applied Sciences*, **11**, Article 8613. https://doi.org/10.3390/app11188613

[14] Kidoh, M., Shinoda, K., Kitajima, M., Isogawa, K., Nambu, M., Uetani, H., *et al.* (2019) Deep Learning Based Noise Reduction for Brain MR Imaging: Tests on Phantoms and Healthy Volunteers. *Magnetic Resonance in Medical Sciences*, **19**, 195-206. https://doi.org/10.2463/mrms.mp.2019-0018

[15] Nazir, N., Sarwar, A. and Saini, B.S. (2024) Recent Developments in Denoising Medical Images Using Deep Learning: An Overview of Models, Techniques, and Challenges. *Micron*, **180**, Article 103615. https://doi.org/10.1016/j.micron.2024.103615

[16] El-Shafai, W., El-Nabi, S.A., Ali, A.M., El-Rabaie, E.M. and Abd El-Samie, F.E. (2024) Traditional and Deep-Learning-Based Denoising Methods for Medical Images. *Multimedia Tools and Applications*, **83**, 52061-52088. https://doi.org/10.1007/s11042-023-14328-x

[17] Muksimova, S., Umirzakova, S., Mardieva, S. and Cho, Y. (2023) Enhancing Medical Image Denoising with Innovative Teacher–Student Model-Based Approaches for Precision Diagnostics. *Sensors*, **23**, Article 9502. https://doi.org/10.3390/s23239502

[18] Kumar, R.R. and Priyadarshi, R. (2024) Denoising and Segmentation in Medical Image Analysis: A Comprehensive Review on Machine Learning and Deep Learning Approaches. *Multimedia Tools and Applications*, **84**, 10817-10875. https://doi.org/10.1007/s11042-024-19313-6

[19] Rai, S., Bhatt, J.S. and Patra, S.K. (2021) An Unsupervised Deep Learning Framework for Medical Image Denoising. https://arxiv.org/pdf/2103.06575

[20] Chaturvedi, S., T, A.S.S., R, K., M, V., A, N.K. and M, S. (2022) Medical Image Denoising and Classification Based on Machine Learning: A Review. *ECS Transactions*, **107**, 6111-6122. https://doi.org/10.1149/10701.6111ecst

[21] Zhang, J., Niu, Y., Shangguan, Z., Gong, W. and Cheng, Y. (2023) A Novel Denoising Method for CT Images Based on U-Net and Multi-Attention. *Computers in Biology and Medicine*, **152**, Article 106387. https://doi.org/10.1016/j.compbiomed.2022.106387

[22] Tian, C., Xu, Y., Li, Z., Zuo, W., Fei, L. and Liu, H. (2020) Attention-Guided CNN for Image Denoising. *Neural Networks*, **124**, 117-129. https://doi.org/10.1016/j.neunet.2019.12.024

[23] Ilesanmi, A.E. and Ilesanmi, T.O. (2021) Methods for Image Denoising Using Convolutional Neural Network: A Review. *Complex & Intelligent Systems*, **7**, 2179-2198. https://doi.org/10.1007/s40747-021-00428-4

[24] Tiantian, W., Hu, Z. and Guan, Y. (2024) An Efficient Lightweight Network for Image Denoising Using Progressive Residual and Convolutional Attention Feature Fusion. *Scientific Reports*, **14**, Article No. 9554. https://doi.org/10.1038/s41598-024-60139-x

[25] Petersen, R.C., Aisen, P.S., Beckett, L.A., Donohue, M.C., Gamst, A.C., Harvey, D.J., *et al.* (2010) Alzheimer's Disease Neuroimaging Initiative (ADNI). *Neurology*, **74**, 201-209. https://doi.org/10.1212/wnl.0b013e3181cb3e25

[26] Mukhopadhyay, S. and Mandal, J.K. (2013) Wavelet Based Denoising of Medical Images Using Sub-Band Adaptive Thresholding through Genetic Algorithm. *Procedia Technology*, **10**, 680-689. https://doi.org/10.1016/j.protcy.2013.12.410

[27] Buades, A., Coll, B. and Morel, J.M. (2005) A Non-Local Algorithm for Image Denoising. 2005 *IEEE Computer Society Conference on Computer Vision and Pattern*

Recognition (*CVPR*'05), San Diego, 20-25 June 2005, 60-65. https://doi.org/10.1109/cvpr.2005.38

[28] Taassori, M. and Vizvári, B. (2024) Enhancing Medical Image Denoising: A Hybrid Approach Incorporating Adaptive Kalman Filter and Non-Local Means with Latin Square Optimization. *Electronics*, **13**, Article 2640. https://doi.org/10.3390/electronics13132640

[29] Zhang, Y., Hao, D., Lin, Y., Sun, W., Zhang, J., Meng, J., *et al.* (2023) Structure-Preserving Low-Dose Computed Tomography Image Denoising Using a Deep Residual Adaptive Global Context Attention Network. *Quantitative Imaging in Medicine and Surgery*, **13**, 6528-6545. https://doi.org/10.21037/qims-23-194

[30] Xie, J., Xu, L. and Chen, E. (2012) Image Denoising and Inpainting with Deep Neural Networks. *Advances in Neural Information Processing Systems*, **25**. https://proceedings.neurips.cc/paper/2012/hash/6cdd60ea0045eb7a6ec44c54d29ed402-Abstract.html

[31] Lundervold, A.S. and Lundervold, A. (2019) An Overview of Deep Learning in Medical Imaging Focusing on MRI. *Zeitschrift für Medizinische Physik*, **29**, 102-127. https://doi.org/10.1016/j.zemedi.2018.11.002

[32] Chen, Z., Pawar, K., Ekanayake, M., Pain, C., Zhong, S. and Egan, G.F. (2022) Deep Learning for Image Enhancement and Correction in Magnetic Resonance Imaging—State-of-the-Art and Challenges. *Journal of Digital Imaging*, **36**, 204-230. https://doi.org/10.1007/s10278-022-00721-9

[33] Organización Mundial de la Salud (2017) Global Atlas of Medical Devices 2022. World Health Organization. https://www.who.int/publications/i/item/9789240062207

[34] Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., *et al.* (2019) A Guide to Deep Learning in Healthcare. *Nature Medicine*, **25**, 24-29. https://doi.org/10.1038/s41591-018-0316-z