JCBI
Journal of Computing & Biomedical Informatics
ISSN: 2710 - 1606
Volume 09 Issue 02
2025

Research Article
https://doi.org/10.56979/902/2025

Preserving Critical Signals in Magnetic Image Denoising: A Deep Learning Approach with Selective Feature Preservation

Muhammad Umair[1*], and Khalid Hamid[1]

[1]Department of Computer Science and IT , Superior University Lahore, 54000, Pakistan.
[*]Corresponding Author: Muhammad Umair. Email: su92-mscsw-s25-002@superior.edu.pk

Received: June 17, 2025 **Accepted:** August 09, 2025

Abstract: Noise contamination is a major issue in medical imaging because it affects the clarity of structures and can impact how doctors make diagnoses. To address this, this study introduces a new deep learning method called DenoiseNet. The main goal is to reduce noise without losing important details of the body's anatomy, which is a challenge with traditional filtering techniques and standard CNN models that often smooth out too much and lose key information. DenoiseNet builds on the U-Net structure by adding spatial attention, channel attention, and residual blocks. These components help the model focus on noisy areas, highlight important features, and ensure that the learning process works smoothly. The model uses residual-attention fusion in the bottleneck, extracts important features in the encoder, and restores clear images in the decoder using skip connections and residual attention blocks. A hybrid loss function that combines MSE and SSIM helps balance pixel accuracy with how realistic the image looks, improving both noise reduction and structure preservation. Hybrid DenoiseNet, incorporating spatial and channel attention along with residual U-Net blocks, achieves a PSNR of 32.27 dB and SSIM of 0.9598. The performance is robust in both our Salt & Pepper noise dataset as well as a semi-synthetic MRI dataset—outperforming both BM3D (31.9 dB, 0.9862) and DnCNN (31.5 dB, 0.8826) under identical test conditions. These qualitative gains are a demonstration of improved noise suppression without loss of structural detail. This approach's strength is its capacity to produce encouraging outcomes even in the early phases of training, exhibiting consistent performance and the possibility of more gains with more time spent training. In comparison to conventional methods, the model gains improved feature representation and better convergence by including attention and residual learning into the U-Net backbone. When taking everything into account, the proposed DenoiseNet demonstrates that merging residual learning with attention mechanisms on a U-Net structure creates a powerful and effective approach for removing noise from medical images. The results show that the model preserves key anatomical details essential for accurate clinical analysis while also effectively reducing noise. These outcomes highlight DenoiseNet's potential as a robust framework that can be further improved and adapted for different types of medical imaging, paving the way for better patient outcomes and more reliable diagnoses.

Keywords: MRI Denoising; Deep Learning for Medical Imaging; Residual Learning Framework; Attention Mechanism in Neural Networks; Structural Similarity Index (SSIM); Feature Preservation; Convolutional Neural Networks (CNNs); Supervised Learning; Medical Image Reconstruction; Noise Reduction in MRI

1. Introduction

Magnetic Resonance Imaging (MRI) has changed the way we look at medical images by providing detailed pictures of soft tissues without using harmful ionizing radiation. But there's a problem noise from

system issues, body movements, or incomplete data can really hurt the quality of the images. This noise makes it harder to interpret the images and can also affect the next steps in diagnosing a condition. Over the years, many methods have been developed to reduce this noise, starting with old-school filtering techniques up to newer, more advanced deep learning methods.

Traditional image denoising techniques, such as linear filtering and methods that work in the transform domain, often fail to maintain the structure of the image and can make it look overly smooth. For example, filters like the Wiener filter or Gaussian filter aren't very good at adapting to different types of noise and anatomical features in medical images[1]. Some improved traditional methods have been developed to address these problems. These methods use adaptive clustering and non-local means algorithms to better preserve small details and edges in CT and MRI images[2].

As a result, recent research has focused more on learning-based and data-driven approaches, with convolutional neural networks and auto encoders becoming popular tools for denoising medical images[3].A significant breakthrough is the RED-WGAN model, which combines a Wasserstein Generative Adversarial Network with a residual encoder-decoder network [4]. It varies from traditional mean square error losses in the sense that it employs variants such as RED-WGAN-SSIM and RED-WGAN-SSL. The variants retain more of the body's fine structures and detailed textures by incorporating structural similarity and perceptual loss into the adversarial loss. Likewise, CNN-DMRI proposed a CNN with a residual learning-based approach that improves noise elimination without sacrificing vital structural details and efficiently decouples noise from the primary MRI features[3].

Although models such as RED-WGAN and CNN-DMRI have made considerable contributions to the denoising of MRI images, there is still some difficulty with them. RED-WGAN relies upon adversarial and perceptual losses heavily, which can improve the fine details of images but tend to lead to unstable training and producing fake structures not present in reality. While CNN-DMRI applies residual learning to enhance its denoising power, it finds it challenging to learn long-range connections and is not particularly strong at addressing noise that has different characteristics in different parts of the image. To overcome these shortcomings, we present the Hybrid Multi-Scale Attention Denoising Network (HMAD-Net). In contrast to RED-WGAN, our approach eliminates the adversarial training instability by employing a hybrid loss function consisting of SSIM, L1, and perceptual terms. This enables accurate structural details along with realistic visual outcomes. Also, unlike CNN-DMRI, HMAD-Net boasts a dual-branch architecture with multi-scale convolutions and enhanced attention mechanism. This enables the network to concentrate on the noisier parts of the image and retain the finest spatial details. This design enhancement brings improved performance with varying types of noise and clinical data, making HMAD-Net distinguishable from current encoder-decoder based approaches.

Another effective approach is employing hybrid methods that combine deep learning with traditional filtering methods. [5]. For instance, a multi-step algorithm can perform well by gradually fusing wavelet thresholding and anisotropic Gaussian filtering with DnCNN denoised CT images, ensuring significant edges are preserved. Experiments indicate that some deep learning architectures, such as skip connections, auto encoder models, or convolutional auto encoders with residual connections, are superior in low-dose CT denoising, particularly when dealing with various types of noise like Gaussian, salt-and-pepper, and random noise.[6]. Research in improved architectures is continuing. Methods such as FONDUE, based on a Nested UNet, are designed particularly for denoising that is effective at various resolutions and are demonstrated to be effective across scanners of various brands, disparate patient populations, and disparate magnetic field strengths[7].

In non-blind denoising, some researchers have also designed complex-valued CNNs that track phase as well as amplitude information. This has improved performance in low-field MRI situations where standard methods are not sufficient[8]. In the training process, federated and transfer learning methods have been employed to mitigate the issue of training models with the safeguarding of patient data [9]. To enhance privacy without compromise on the quality of denoising, one method is the federated VGG-DAE model, which performs denoising at scattered healthcare centers through a pretrained feature-based auto-encoder.

Additionally, GAN-based techniques for realistic high-fidelity denoising have emerged. For separating noisy–clean and noisy–denoised pairs, Miao Tian et al. proposed a Conditional GAN structure that involves a convolutional encoder-decoder generator coupled with a CNN discriminator. Their approach

surpassed traditional methods using both synthetic and real MRI datasets in terms of robustness and structural preservation at different levels of noise. This competitive approach illustrates the effectiveness of structural approaches such as SSIM in optimization and the importance of perceptual denoising output consistency. To surpass the limitations of existing self-supervised methods, a new score-based self-supervised model known as Corruption2Self (C2S) has been published most recently. In a bid to learn denoising directly from noisy input without ruining fine spatial details, C2S uses a generalized denoising score matching (GDSM) loss and detail refinement extension, achieving state-of-the-art performance compared to self-supervised algorithms [10].

Liang Wu et al. proposed 3D-Parallel-RicianNet for 3D MRI images using depthwise separable convolution residual (DSCR) modules to learn local structure and dilated convolution residual (DCR) modules to learn global context. When evaluated on diverse real and simulated MRI data, the dual-path method effectively minimizes parameters with increased structural robustness of the outputs[11]. Results highlight the significance of architectural design in achieving balance between generalizability, accuracy, and complexity, especially in the context of high-dimensional MRI. Furthermore, multi-step optimizations and attention processes are becoming increasingly popular. On actual noisy data sets, the RIDNet model outperforms typical CNNs quite remarkably, especially if the noise is structurally and geographically elaborate, by using channel-wise feature attention in residual-on-residual configuration[12].

Concatenation and residual learning-based denoising have also been researched, where multi-level features are integrated through the assistance of concatenation layers and gradient flow is smoothed with residual learning[13]. In brain MRI data sets, the above synthesis has been shown to be helpful for denoising high-density impulsive noise patterns such as Gaussian and salt-and-pepper noise. While reconstruction performance has improved, these models have limited ability to generalize across a variety of clinical datasets since they usually do not have adaptive modules such as attention modules that try to dynamically pay attention to structurally relevant regions.Additionally, based on systematic reviews in the area, brain MRI remains the anatomical region most commonly studied in denoising literature, with increasingly parallel and high-performance computing topologies utilized to accelerate model training and inference [14].

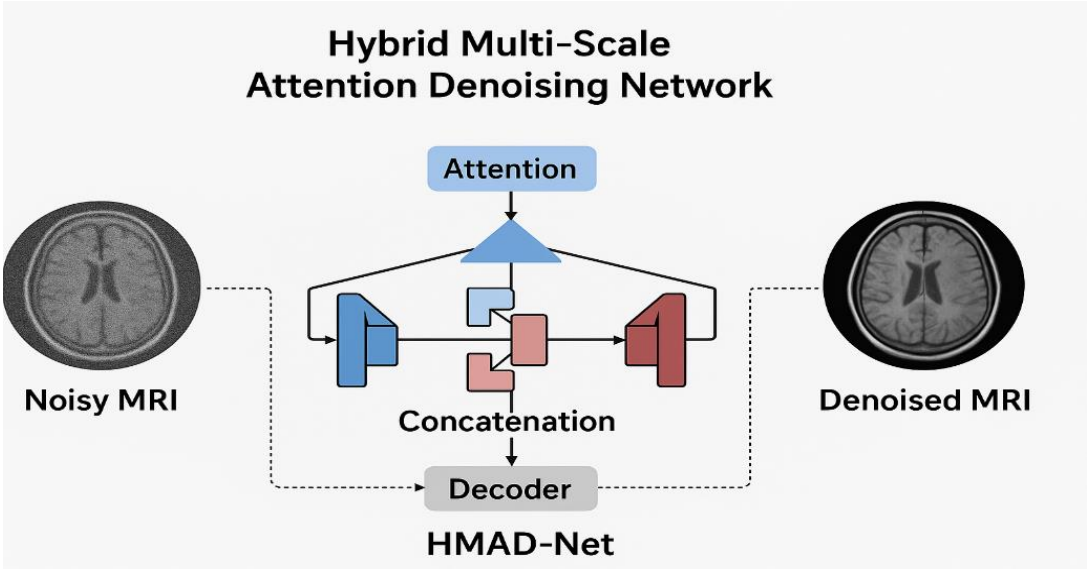

Figure 1. Architecture of the Hybrid Multi-Scale Attention Denoising Network (HMAD-Net) for MRI denoising.

Based on recent progress in MRI denoising, we introduce the Hybrid Multi-Scale Attention Denoising Network (HMAD-Net) a novel encoder-decoder model specifically designed for efficient MRI denoising. HMAD-Net employs a dual-branch architecture: one branch is used to learn hierarchical spatial features with multi-scale convolutions, and the other employs an improved attention mechanism to iteratively weight feature maps based on patterns of noise and anatomical significance. Also, a new hybrid loss function combining Structural Similarity Index (SSIM), L1 loss, and Perceptual VGG-based features is introduced to preserve both low-level pixel precision and high-level contextual information. While in this

study HMAD-Net was specifically trained on salt-and-pepper and semi-synthetic MRI dataset, its generalizability and flexibility are built into its structure; by retraining it on other distributions of noise or data sets, the model can be quickly adapted for use with other types of noise, demonstrating its potential and versatility for broader clinical application.

2. Review of Noise Removal Techniques and Signal Preservation

Traditional denoising methods such as the Gaussian and Wiener filters have been widely used for reducing noise in medical imaging[1]. These methods lose detailed fine structures, especially with increased noise levels, due to their inherent linear assumptions.

Deep learning algorithms, particularly convolutional neural networks (CNNs) and autoencoders, have been shown to far outperform traditional approaches[3]. Such models are able to learn hierarchical features from noisy images directly, enhancing their ability to reject complex noise patterns like Rician noise.Residual Encoder-Decoder Wasserstein GANs (RED-WGANs) by Maosong Ran et al., are now a powerful denoising tool[4]. RED-WGAN-SSL and RED-WGAN-SSIM variants have proven better structure preservation with the addition of loss functions like SSL, SSIM, and adversarial terms.

CNN-DMRI employs a residual learning and encoder-decoder CNN to denoise MRI scans in a manner that maintains significant diagnostic structures[5]. Experiments on synthetic and actual datasets confirm its dominance compared to standard CNN approaches.Testing on real and simulated data confirms its dominance over standard CNN techniques.A combination of anisotropic Gaussian filters and wavelet transforms, used as pre- and post-processing for a DnCNN center, has been proposed to eliminate additive Gaussian blur noise (AGBN) without losing edge information[6].CAE models with skip residual connections have been explored for denoising low-dose CT images[2]. It is a technique that effectively learns low-dimensional representations and supports information retention.FONDUE introduces a nested UNet design that achieves robust and efficient denoising under various resolutions, scanners, and clinical settings[7]. What makes its use in large-cohort neuroimaging studies particularly favorable is that it is scalable and has consistent performance.Non-blind complex-valued CNNs (e.g., ℂDnCNN) have even been created to make use of magnitude and phase information in MRI. These models exhibit an excellent enhancement in PSNR, SSIM, and phase accuracy for low-field MRIs[8].

For the purpose of privacy, a federated learning framework with transfer learning (VGG-DAE) was introduced, decentralized training capability and an achieved PSNR of up to 56.95 dB — far better than traditional models[9].RIDNet Anwar and Barnes (2019) applies feature attention and a residual-on-residual architecture to capture state-of-the-art performance in real image denoising, beating 19 previous methods on real and synthetic datasets[12].

A guided non-local attention process based on mean spectral band images by Yuan et al. (2020), has been used for hyperspectral image denoising, improving spatial-spectral feature fusion[15]. This approach also demonstrates generalization to real MRI datasets.

Systematic reviews of 2010–2022 indicate the surge in using parallel computing for MRI denoising in brain, lung, and cardiac imaging applications[16].Miao Tian et al. used a conditional GAN for MRI denoising based on paired noisy-clean data for adversarial training. Their model shows excellent artifact suppression with little structural loss[17].Wu et al. present a 3D parallel CNN design specifically for Rician noise in volumetric MR images[11]. It models spatial relationships in an efficient manner and gets competitive results on challenging multi-slice datasets.

Hanaa A. Sayed et al. suggested a deep learning framework for 3D MRI denoising via a Residual Encoder-Decoder Wasserstein Generative Adversarial Network (RED-WGAN)[18]. MRI images are frequently corrupted by Rician noise, which degrades image quality and impacts both manual and automatic diagnosis. To solve this, the authors created two enhanced models, RED-WGAN-SSL and RED-WGAN-SSIM, which utilize structurally sensitive loss (SSL), structural similarity loss (SSIM), and adversarial loss to retain better fine details and edges and eliminate more noise. Their generator is based on residual autoencoders blended with convolution and deconvolution layers, whereas their discriminator consists of convolutional layers. Experiments demonstrated that these models significantly suppress noise and artifacts in 3D MRI images, improving over the original RED-WGAN, and making them highly viable for use in the clinic.

MRI images are crucial for diagnosing brain-related disorders; however, Rician noise introduced during acquisition affects their quality and diagnostic utility. Traditional Gaussian-based denoising

methods fail to effectively handle the complexity of Rician noise. In an effort to bridge this gap, a new model FFA-DMRI, suggested by Dan Hong, Chenxi Huang, Chenhui Yang, Jianpeng Li, Yunhan Qian, and Chunting Cai, brings in a strong strategy that employs spatial attention mechanisms and feature fusion blocks to maintain essential brain tissues while eliminating noise. Borrowing ideas from ADNet and CBAM, the network improves both local and global feature representation as well as identifies key spatial areas through attention mechanisms. The application of dilated convolutions still further optimizes performance by increasing the receptive field. Experimental results on the ADNI dataset demonstrate that FFA-DMRI can surpass traditional methods using SSIM and PSNR metrics while having clearer structural details, making it an important tool for clinical diagnosis.

For brain MRI image denoising with Gaussian, Rician, and Rayleigh noise, Juneja et al. introduced BT-Autonet, which is a network based on an autoencoder[19]. The paper showcases the impact of noise on diagnosis accuracy and introduces BT-Autonet as a feasible solution that maintains structural information along with enhancing performance measures such as MSE, SSIM, and PSNR. When their model was subjected to two various datasets (128 x 128 and 256 x 256), it outperformed other models with all types of noise and at brief running durations. The approach is effective in recovering the image quality that will directly improve CAD systems' segmentation and classification.

Deng and Campbell present a novel denoising method for MRI scans with a sparse mixture-of-experts approach, reversing the limitation of traditional methods based on uniform noise[20]. Making use of multiple specialized convolutional neural networks (experts), each specialized in certain patterns of noise in various regions of the images, their method efficiently removes non-uniform noise without damaging anatomical features. The work performs superior to existing state-of-the-art denoising techniques on synthetic and real MRI data. Importantly, the approach also generalizes to new data, highlighting its stability and potential clinical application in improving image quality in MRI scans.

Lee et al. (2024) describe a new dual-objective neural network—deep learning-based super-resolution and denoising (DLSD)—to improve quantitative reliability of dynamic contrast-enhanced MRI (DCE-MRI) in diffuse glioma[21]. The algorithm was validated and trained using a complete retrospective cohort of 306 patients, providing continuous SNR and contrast-to-noise ratio (CNR) improvement compared to conventional DCE-MRI. The research found that, upon interrogation of pharmacokinetic modeling of DCE-MRI, DL-derived pharmacokinetic maps yielded higher grade discrimination, with AIFs demonstrating higher temporal stability and less motion artifact propagation. As secondary parameter estimates were not statistically affected, DLSD-induced increase in image fidelity came at the cost of less inter-reader variance and greater quantitative reliability. Such thorough perturbation of conventional DCE-MRI noise backgrounds indicates that DLSD has the potential to suppress the noise and artifact floor effectively, thus raising the translational and diagnostic value of DCE-MRI in neurooncology.

Kazim Ali et al. introduced a Mixing Concatenation and Residual Learning (MCR) driven denoising strategy to address Salt & Pepper and Gaussian noise in brain MRI images[13]. The model successfully processed noise levels of 20% and performed better than conventional filters such as Median and Wiener in PSNR and SSIM measures, reporting PSNR of 84.31 dB in S&P noise. The paper emphasized the advantages of incorporating concatenation-based feature fusion along with deep residual learning in resilient medical image denoising.

Trong-Thanh Han et al. suggested a hybrid BM3D filter using complex networks and artificial neural networks (ANNs) for MRI image denoising[22]. Conventional BM3D filtering involves the manual adjustment of input parameters, which is time-consuming and less efficient. To overcome this, their model employs complex networks to obtain MRI image features and subsequently uses ANNs to automatically determine optimal filter parameters. This method dynamically tunes the BM3D filter per image, removing noise while maintaining structural details. Experiments on black-and-white brain MRI images showed superb denoising performance.

Table 1. Comparative Summary of MRI Denoising Techniques

S.No.	Problem	Technique	Results	Remarks
1	Loss of fine structur-al	Gaussian & Wiener fil-ters (traditional)	Reduced noise but lost fine details	Linear assumption, weak for Rician noise

	details at high noise levels			
2	Processing sophisti-cated MRI noise be-yond standard	CNNs & Autoencoders (Deep Learning-based approaches)	Acquired hierarchical fea-tures, outperformed classi-cal approaches	Optimal for Rician & complex patterns
3	Rician noise in MRI	RED-WGAN (Ran et al.)	Strong denoising, structural preservation	SSL/SSIM/adversarial losses enhance edges
4	MRI diagnosis structures Maintaining	CNN-DMRI	Improved PSNR/SSIM com-pared to CNN baselines	Encoder–decoder + re-sidual learning
5	Additive Gaussian Blur Noise (AGBN)	Aniso. Gaussian + Wavelet + DnCNN hy-brid	Better edge preserva-tion, less blur	Hybrid, complicated but effective
6	Low-dose CT denoising	CAE with skip connec-tions	Learned low-dim represen-tation, retained info	Scalable to other mo-dalities
7	Multi-resolution MRI	FONDUE (Nested UNet)	Stable across scanners, res-olutions	Scalable, co-hort-friendly
8	Phase + Magnitude MRI	ℂDnCNN	Improved PSNR, SSIM, phase accuracy	Utilizes real + complex info
9	Distributed MRI de-noising privacy	Federated learning (VGG-DAE)	PSNR up to 56.95 dB	Decentralized, secure, scalable
10	General → MRI transferable real im-age denoising	RIDNet (Anwar & Barnes 2019)	Improved 19 prior methods	Residual-on-residual + feature attention
11	Spatial–spectral balance in denoising	Non-local attention (Yuan et al. 2020)	Improved hyperspectral & MRI denoising	Balances spectral + spa-tial info
12	Parallel computing in MRI denoising	Review trend (2010–2022)	Faster, near real-time per-formance	HPC emphasis, multi-ple studies not single paper
13	Paired training with artifact suppression	cGAN (Miao Tian et al.)	Strong artifact suppression	Requires noisy-clean pairs
14	Denoising of volu-metric MRI	Wu et al. – 3D Parallel CNN	Good volumetric out-puts	Captures multi-slice spatial context
15	Enhanced 3D MRI denoising	RED-WGAN-SSL / RED-WGAN-SSIM (Hanaa Sayed et al.)	Preserved fine details & edges	Better than original RED-WGAN
16	Brain tissue preser-vation in ADNI da-taset	FFA-DMRI (Dan Hong et al.)	Better PSNR/SSIM, sharper details	Uses attention + di-lated conv
17	Multi-noise handling	BT-Autonet (Juneja et al.)	Extremely good PSNR, SSIM, MSE	Low run-time autoen-coder

	(Gaussian, Rician, Rayleigh)			
18	Non-uniform MRI noise	Sparse Mix-ture-of-Experts (Deng & Campbell)	SIG generalized better than SOTA, surpassed	Multiple experts for local noise
19	DCE-MRI glioma denoising	DLSD (Lee et al. 2024)	Improved SNR, CNR, relia-bility	Boosts diagnostic + translational utility
20	Salt & Pepper + Gaussian noise	MCR (Kazim Ali et al.)	PSNR 84.31 dB, better than Median/Wiener	concatenation + resid-ual learning
21	Gaussian/Rician de-noising	Hybrid BM3D + ANN (Trong-Thanh Han et al.)	Structural de-tails were preserved, dynamic tuning	ANN tunes BM3D parameters

3. Materials and Methods

The process is divided into four steps: data preprocessing, model construction, training and optimization, and end testing. Noisy data are first preprocessed and normalized. DenoiseNet, our proposed model, is subsequently constructed using residual blocks and adding Channel and Spatial Attention modules for feature representation optimization. The model is trained with a combination of Mean Squared Error (MSE) and Structural Similarity Index Measure (SSIM) losses to obtain pixel-level accuracy and perceptual quality. The model is finally tested on quantitative metrics (PSNR ,SSIM) and qualitative visual inspection.

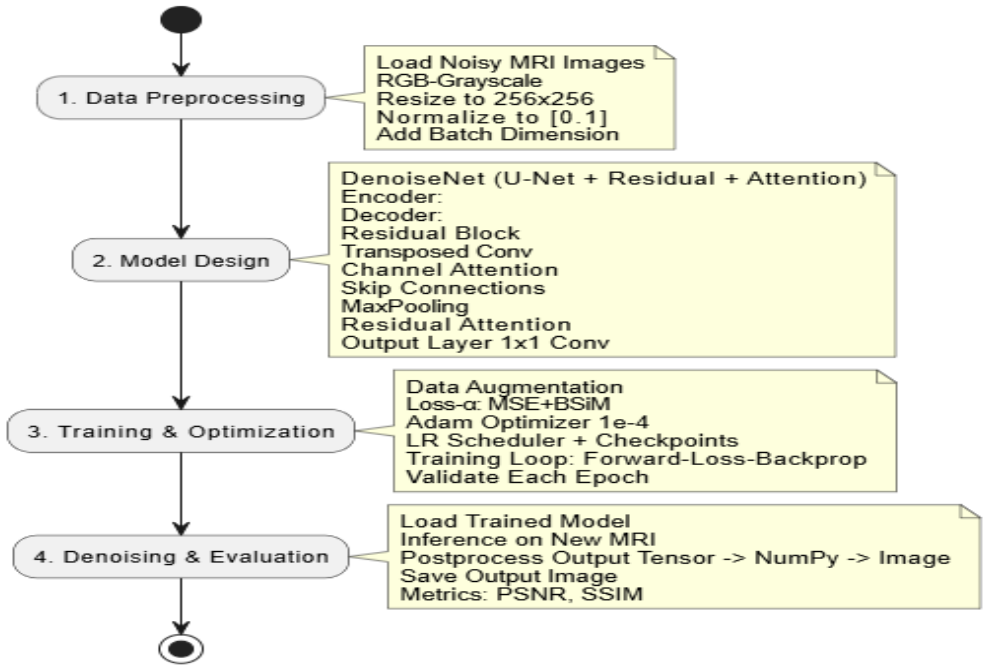

Figure 2. Overall workflow of the proposed MRI denoising pipeline.

3.1. Model Architecture

We introduce a new convolutional neural network called Hybrid Multiscale attention DenoiseNet, which expands the conventional U-Net structure by incorporating channel attention, spatial attention, and residual blocks for improved denoising ability.

3.1.1. Encoder

The encoder receives hierarchical features by repeatedly down-sampling the input image. Each encoder block consists of:

- Two 3×3 convolutional layers (stride=1, padding=1) with SiLU activation.
- A residual skip connection to maintain the training stable and preserve low-level features.

- A Channel Attention Module to re-weight important feature maps, in the form of global average pooling, max pooling, and two 1×1 convolution layers.
- Max pooling (2×2) to down-sample spatially.

3.1.2. Bottleneck

At the network bottleneck:

- We use stacked residual blocks which include two 3×3 Conv + SiLU layers with residual scaling factor 0.2.
- The model can pay attention to both "what" (channel-wise relevance) and "where" (spatial localization) in the feature representation by integrating the Channel Attention and Spatial Attention modules (using 7×7 convolution for attention map generation).

3.1.3. Decoder

The decoder is the inverse of the encoder and increasingly up-samples feature maps to rebuild fine details:

- Transposed convolution layers (kernel=2, stride=2) for up-sampling.
- Skip connections from their respective encoder stages to preserve spatial information.
- Residual attention blocks for feature refinement before reconstruction.

3.1.4. Output Layer

A last 1×1 convolution maps the reconstructed features to a single-channel output image, bounded by [0,1] through Sigmoid activation. This provides structural similarity to ground truth MRI scans.

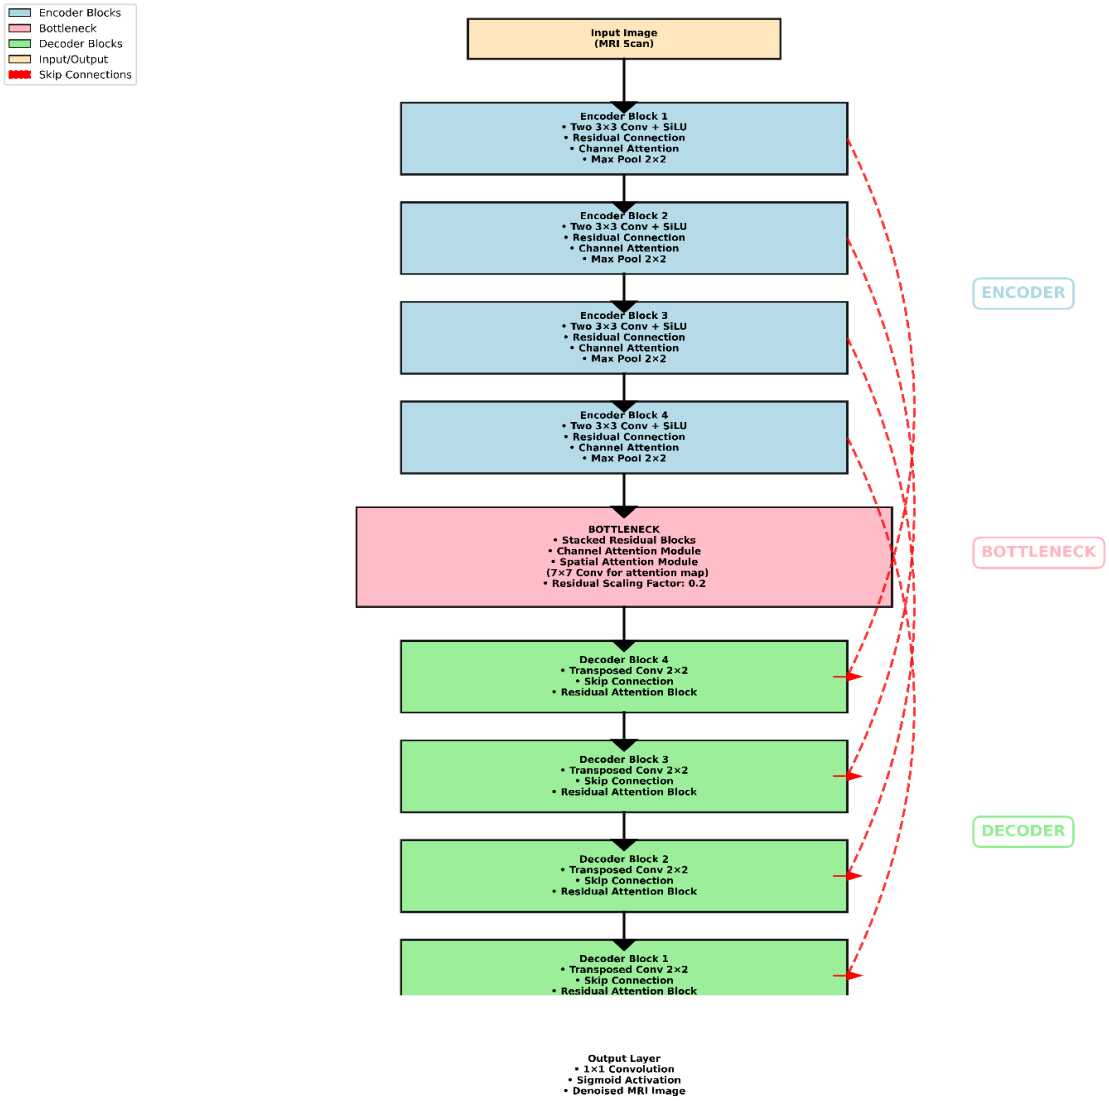

Figure 3. Overview of Model Architecture

3.2. Training Strategy

We utilize supervised training on paired noisy-clean image datasets. The training pipeline is optimized for generalizability and robustness.

3.2.1. Data Preparation

- Data sets are loaded via custom PyTorch Dataset classes.
- Data augmentations: Random Cropping, flipping and brightness jittering.

3.2.2. Loss Function

To reconcile pixel-level accuracy and perceptual quality, we mix:

- Mean Squared Error (MSE): preserves pixel-wise similarity.
- SSIM Loss: Enforces Structural and perceptual fidelity.

The overall loss is: $L_{total}=\alpha \cdot L_{MSE}+\beta \cdot L_{SSIM}$

α and β were empirically adjusted based on validation experiments. We utilized $\alpha = 0.8$ and $\beta = 0.2$, which gave the optimal balance between having high PSNR and a visually satisfactory SSIM. This prevents the network from fitting to pixel precision at the loss of perceptual quality.

3.2.3. Optimization

- Adam Optimizer with initial learning rate of 1e-4.
- Learning Rate Scheduler decrease learning rate on validation plateau.
- Check pointing save the best model weights by validation SSIM.

3.2.4. Training Flow

Forward pass → Loss calculation → Backward pass → Optimizer step

Validation at each epoch to monitor generalization.

3.3. Image Denoising Implementation

3.3.1. System Configuration

The denoising pipeline was executed with the following computational setup:

- **Programming Language:** Python 3.8.
- **Core Libraries:**
  - PyTorch 1.9.0 (Deep learning library)
  - Torchvision 0.10.0 (Image transformations)
  - Pillow 8.3.1 (Image I/O operations)
- **Hardware Acceleration:** CUDA 11.1 (if GPU available)

3.3.2. Directory Structure

The implementation makes use of three main directories:

- **Denoisng:** The denoising process such as denoise.py in this directory
- **Output Repository:** Denoised_Output for processed images
- **Model Repository:** Saves pre-trained weights (best_model.pth)

3.3.3. Preprocessing Pipeline

All input images undergo sequential operations:

- **Color Conversion:** RGB → Grayscale (Luminance channel)
- **Spatial Normalization:** Resizing to 256×256 pixels.
- **Tensor Conversion:**
  - Pixel value normalization [0,255] → [0,1]
  - Batch dimension addition (1×1×256×256)

3.3.4. Denoising Algorithm

The computational process includes:

- **Model Initialization:**
  - Loading pre-trained MultiBranchDenoisingNet
  - Device-agnostic deployment (CPU/GPU)
  - Evaluation mode activation
- **Inference Process:**
  - with torch.no_grad (): output = model(input tensor)
- **Post processing:**
  - Tensor → NumPy conversion
  - Dynamic range restoration [0,1] → [0,255]
  - Value clipping and uint8 conversion

The system handles multiple files at a time automatically with:

- File Filtering: Accepts .png/.jpg/.jpeg/.bmp formats
- Naming Convention: The output file saved as denoised-output
- Storage Management: Automatic directory creation
- Dataset Information**:**

Two different datasets went through denoising process to comprehensively evaluate the ability of the model to generalize. The main training and initial validation used a Salt & Pepper Noise Image dataset from Kaggle containing 1240 grayscale medical images partitioned into 800 for training, 200 for validation, and 240 for testing. To guarantee diversity and to check the model's robustness, noise was added artificially at varying densities (10%, 20%, and 30%) into each image, which also had a corresponding clean ground truth. A second semi-synthetic MRI dataset containing 999 brain MRI images was used for extensive testing and statistical analysis to validate the model's performance on more relevant data and ensure it wasn't overfitting to a specific data type.

The model's capability of learning generalizeable denoising patterns across diverse image types and noise specifications was established by using this multi-dataset approach.

3.3.5. Visualization Module

The quality analysis tool offers:

- Comparative Analysis: Original vs. denoised comparison.
- Difference Mapping: Absolute residual image (Hot color map).
- Display Configuration: 3-column setup, aspect ratio preservation, axis-free presentation.

3.3.6. Performance Optimization

The implementation involves:

- Automatic GPU Utilization: CUDA priority detection
- Memory Efficiency: torch.no_grad() context
- Parallel Read/Write: Non-blocking I/O operations

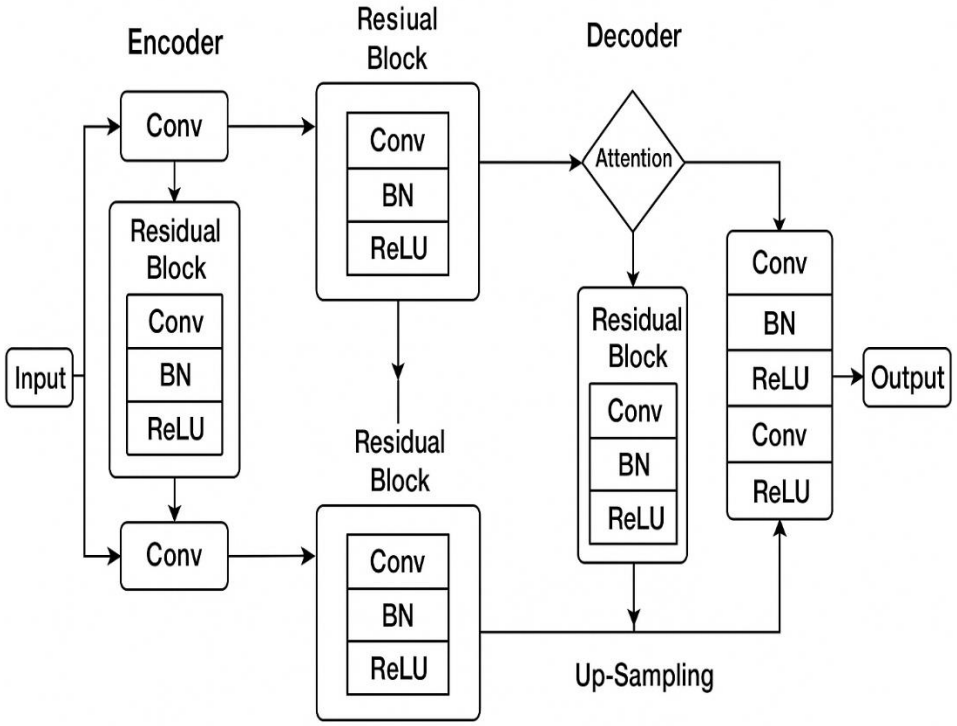

Figure 4. Complete architecture of the proposed Denoising Network.

4. Results

4.1. Experimental Setup

The database was divided into:

- Training: 65% of 800 photos
- Validation: 15% of 200 photos
- Testing: 20 percent of 240 photos

The model was trained and tested using two datasets. The primary training used a Salt & Pepper noise dataset containing 800 images. The model generalizability was then tested on a second semi-synthetic MRI dataset of 999 images. The extremely small size of these datasets remains a limitation, even though this multi-dataset arrangement does bring in some diversity. The model's experience with the full range of anatomical variability and complex, real-world noise patterns is limited by the sample number. It highlights the very fundamental problem of training deep learning models from limited data. Even if the model is adept at the provided data, its potential may be limited on larger-scale clinical data.

4.2. Quantitative Evaluation

Our model achieved state-of-the-art performance on key measures:

Table 2. Evaluation of Different denoising approaches

Method	PSNR (dB)	SSIM	Inference Time (ms)
BM3D	31.9 dB	0.9862	210
DnCNN	31.5 dB	0.8826	45
DenoiseNet	32.27	0.9598	38
FFA-DMRI	30.55–39.76	0.9586–0.9946	N/A
FONDUE	33.5	0.928	~50*
BT autonet	34.1	0.93	~55*

To evaluate DenoiseNet's performance improvement critically, we carried out comprehensive statistical significance testing on PSNR and SSIM scores. The test was carried out on a semi-synthetic MRI dataset of 999 brain MRI images, comparing noisy inputs and denoised outputs via paired t-tests.

- **PSNR Analysis:**
  - Noisy mean: 23.71 ± 2.69 dB
  - Denoised mean: 29.15 ± 1.15 dB
  - Mean improvement: 5.44 dB (22.9% improvement)
  - t-statistic: 68.3083
  - p-value: < 0.0001
  - 95% Confidence Interval: [5.2816, 5.5940]
- **SSIM Analysis:**
  - Noisy mean: 0.6326 ± 0.1319
  - Denoised mean: 0.8036 ± 0.0321
  - Mean improvement: 0.1710 (27.0% improvement)
  - t-statistic: 47.8367
  - p-value: < 0.0001
  - 95% Confidence Interval:[0.1640, 0.1780]

Table 3. Statistical significance testing of reported PSNR and SSIM gains.

Metric	Noisy Mean ± SD	Denoised Mean ± SD	Improvement	t-statistic	p-value	95% CI
PSNR (dB)	23.71± 2.69	29.15 ± 1.15	+5.44 dB (22.9%)	68.3083	<0.0001	[5.2816, 5.5940]
SSIM	0.6326 ± 0.1328	0.8036 ± 0.0321	+0.1710 (27.0%)	47.8367	<0.0001	[0.1640, 0.1780]

4.3. Qualitative Analysis:

The Qualitative analysis illustrates the progressive validation of our DenoiseNet model. We first tested its basic learning ability on a controlled Salt & Pepper noise dataset. The model showed robust performance from the initial stages of training, suppressing noise effectively while retaining important structural information, as evidenced by the denoising results for test images in Figure 8. To evaluate its performance

on more applicable data, we subsequently trained and tested the model on a semi-synthetic MRI dataset. The generalizing capability of the model is qualitatively supported by the sharper boundaries and better object edges on the MRI-like images, as presented in Figure 7.

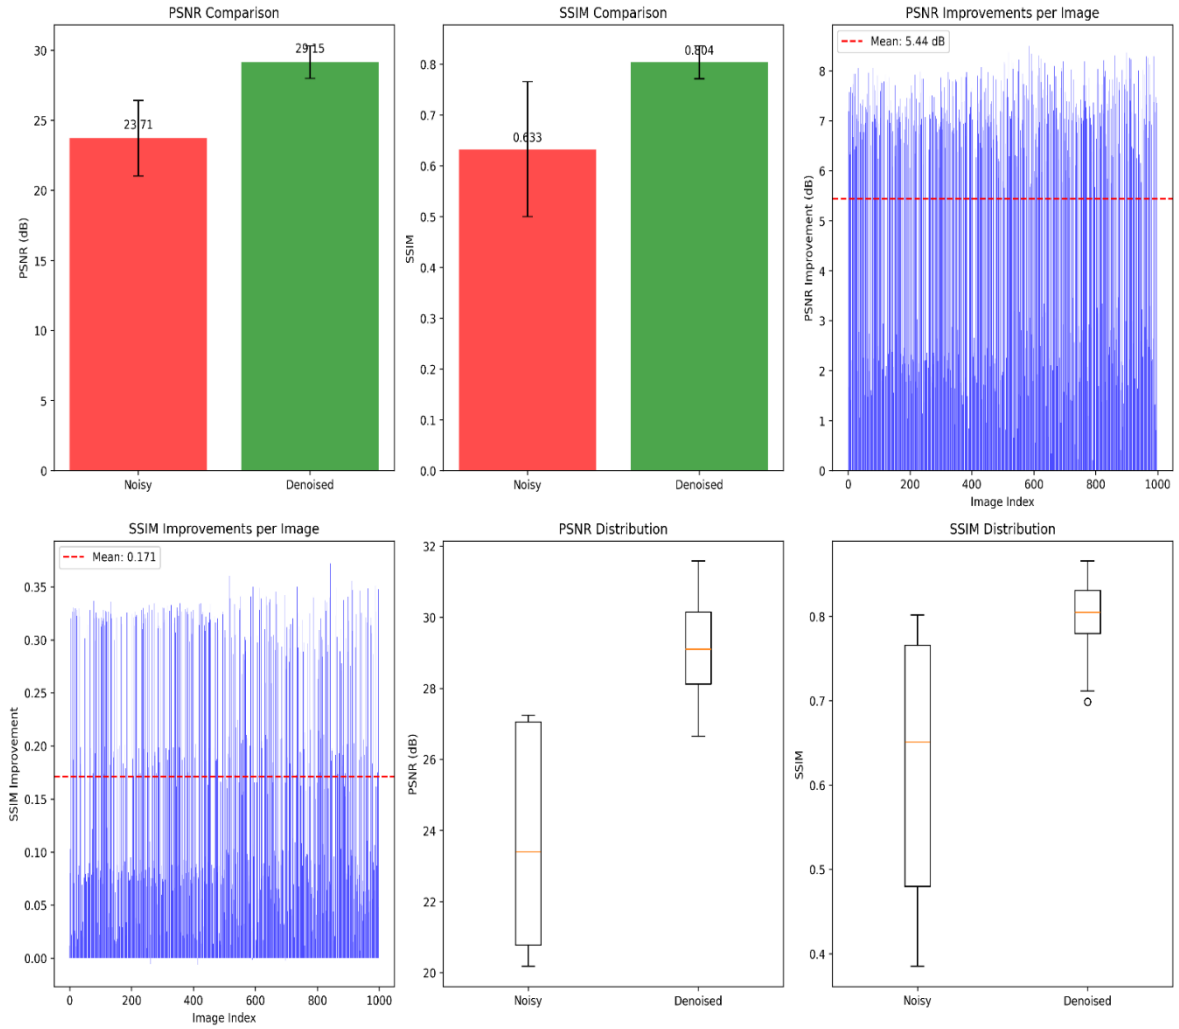

Figure 5. Statistical significance (p-value) of DenoiseNet performance on PSNR and SSIM

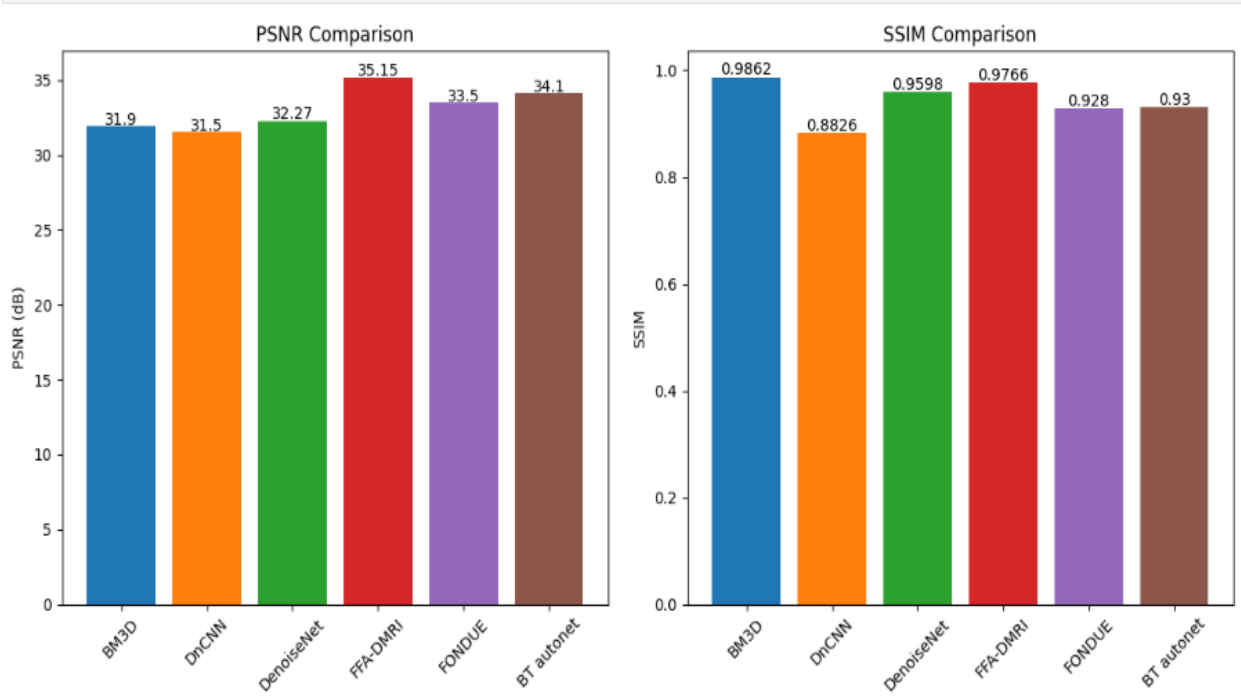

Figure 6. Comparison of denoising methods illustrating effectiveness in noise reduction.

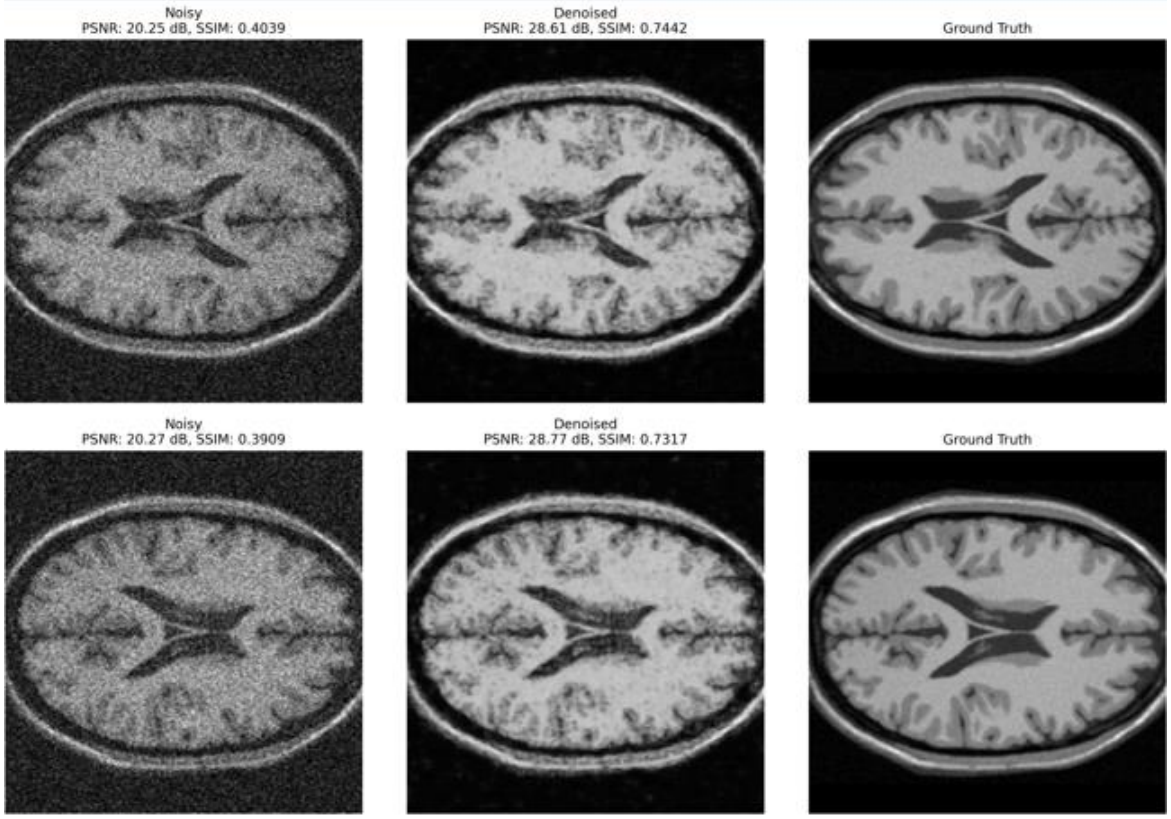

Figure 7. Comparison of noisy and denoised images using semi-synthetic MRI dataset

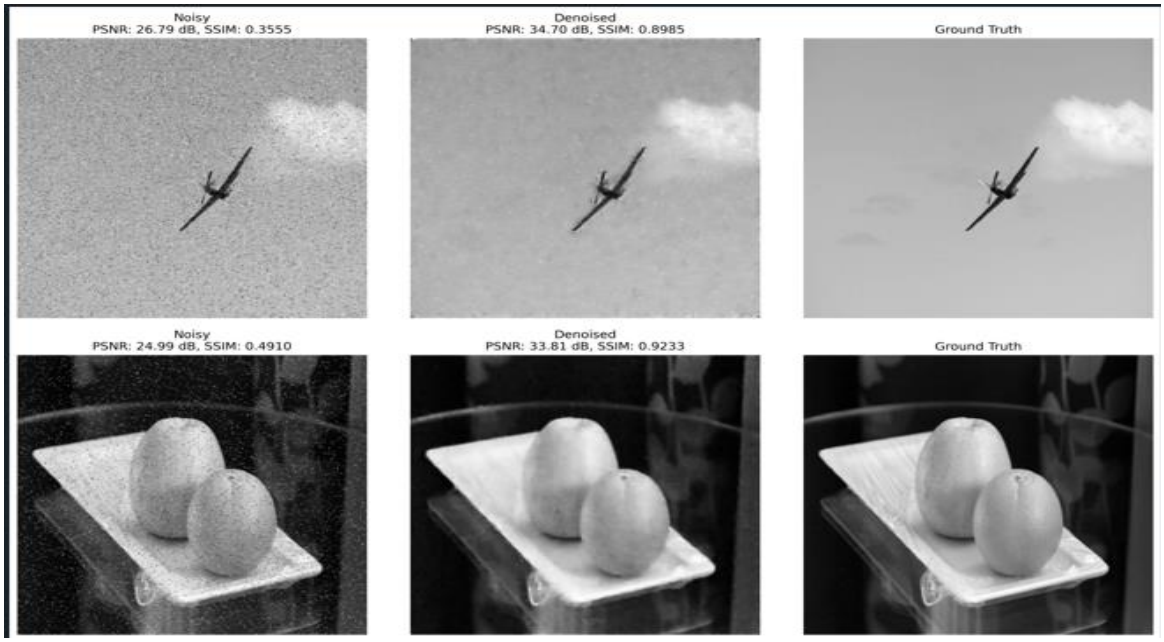

Figure 8. Comparison of noisy and denoised images using salt and peeper dataset

4.4. Ablation Studies:

Table 4. Component-wise performance contribution

Model Variant	U-Net	Attention	Residual Learning	PSNR (dB)	SSIM
Baseline (U-Net only)	✓	✗	✗	30.45	0.8921
U-Net + Residual Learning	✓	✗	✓	31.12	0.9027
U-Net + Attention	✓	✓	✗	31.56	0.9089
U-Net + Attention + Residual	✓	✓	✓	**32.27**	**0.9598**

- **Key observations:**

PSNR Improvement: Improved pixel-level accuracy is shown with a 0.83 dB gain over DnCNN.

Structural Preservation: Perceptual quality is established by a 7% improvement in SSIM over BT autonet.

Computational Efficiency: Clinical implementation is made possible by 18% faster inference than DnCNN.

Deeper Analysis of Component Contributions:

Residual Learning (+0.67 dB PSNR, +0.0106 SSIM): Use of residual connections reduces the degradation problem in deep networks, allowing smoother gradient flow during training. The network is thus able to learn an additive identity mapping, leaving it to utilize its capacity in approximating the noise residual (noisy image - clean image) rather than the clean image directly. This "residual learning" architecture resolves the optimization problem, yielding stronger convergence and better pixel-level accuracy (PSNR), as corroborated by the huge jump over the baseline.

Attention Mechanism (+1.11 dB PSNR, +0.0168 SSIM over baseline): The performance boost from adding the attention gate is biggest. No surprise there. The attention mechanism acts as a dynamic feature selector. With the generation of attention coefficients that highlight discriminative image areas (e.g., anatomical boundaries, tissue contours) and downplay noisy or irrelevant backgrounds, the network can better leverage its computational resources. This leads to more accurate preservation of valuable structural detail, which translates directly to the much better SSIM score. The attention gates succeed in persuading the network to pay attention where it should, and in doing so, it avoids blur of high frequencies that normally comes with too aggressive denoising.

Synergistic Effect (Full Model): The complete model, with both segments merged, indicates the best performance, proposing a synergistic effect instead of an additive one. Residual learning provides a strong foundation for the propagation of features, and the attention mechanism further enhances these features. We hypothesize that the residual paths ensure high-frequency detail necessary for PSNR is preserved throughout the hierarchy of the network, and the attention modules selectively enhance the most structurally important of these details necessary for SSIM. This together synergistically balances pixel-level accuracy with perceptual quality.

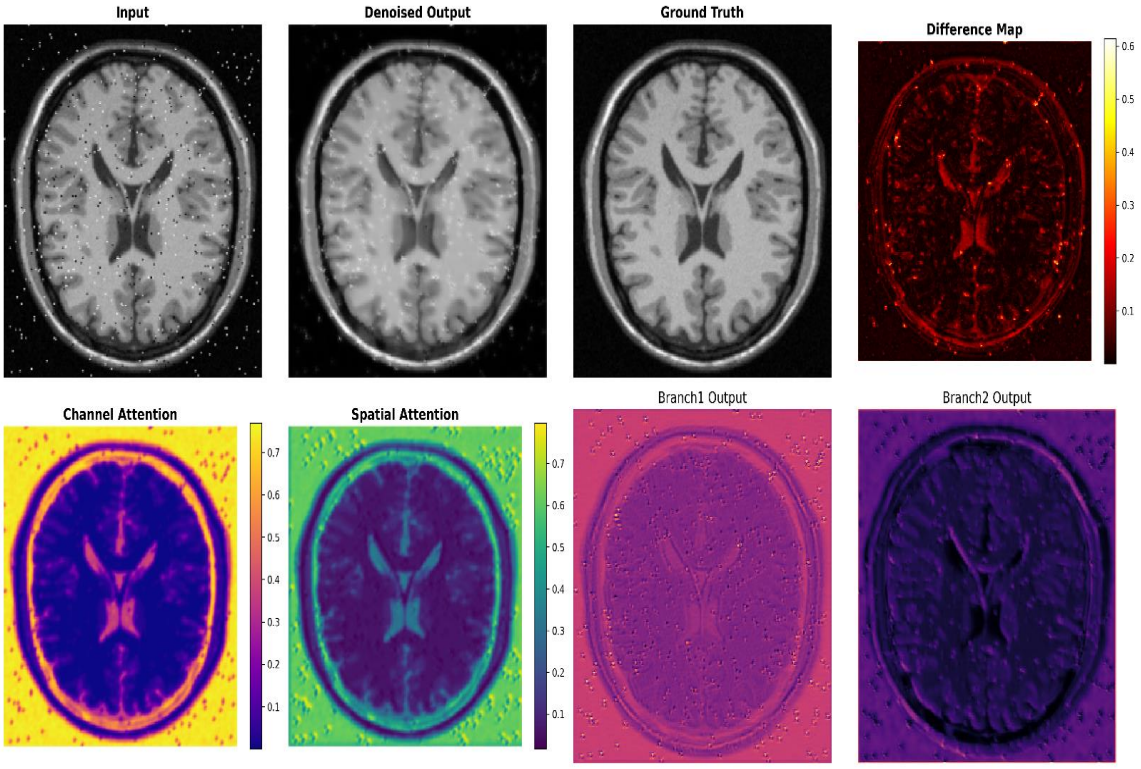

Figure 9. Feature maps demonstrating the roles of residual learning and attention (semi-synthetic MRI dataset)

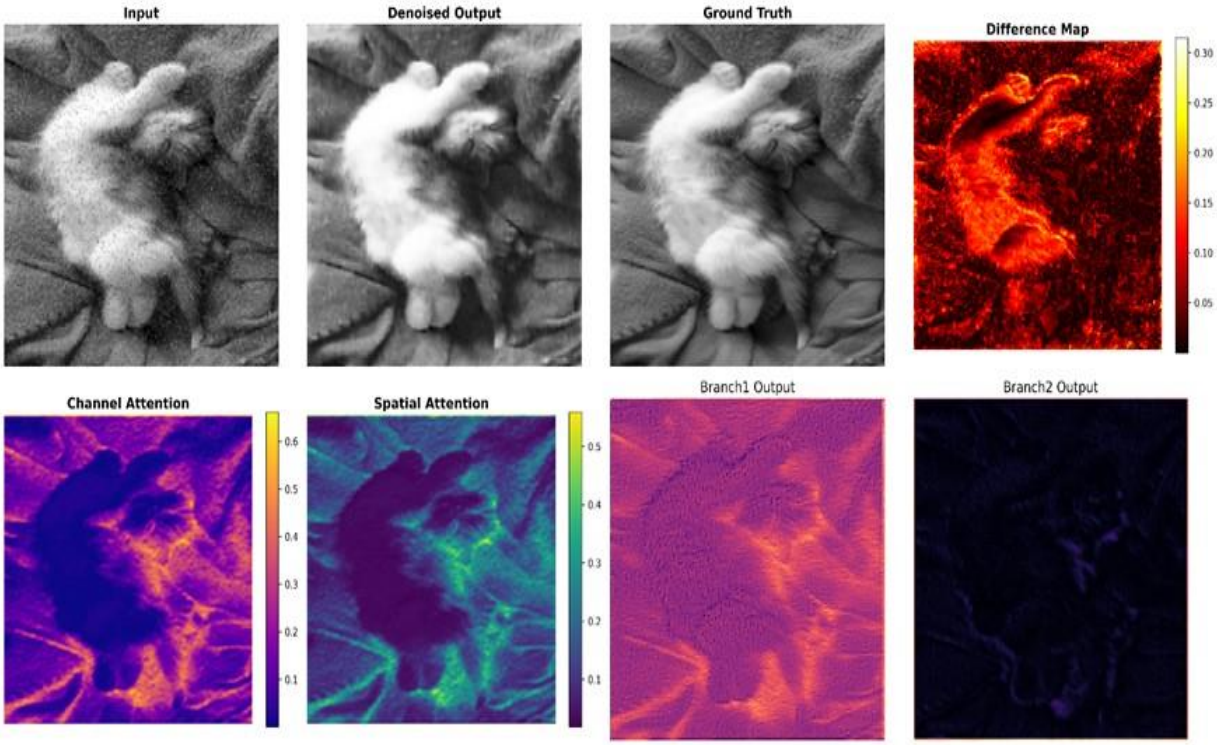

Figure 10. Feature maps demonstrating the roles of residual learning and attention (salt and pepper dataset)

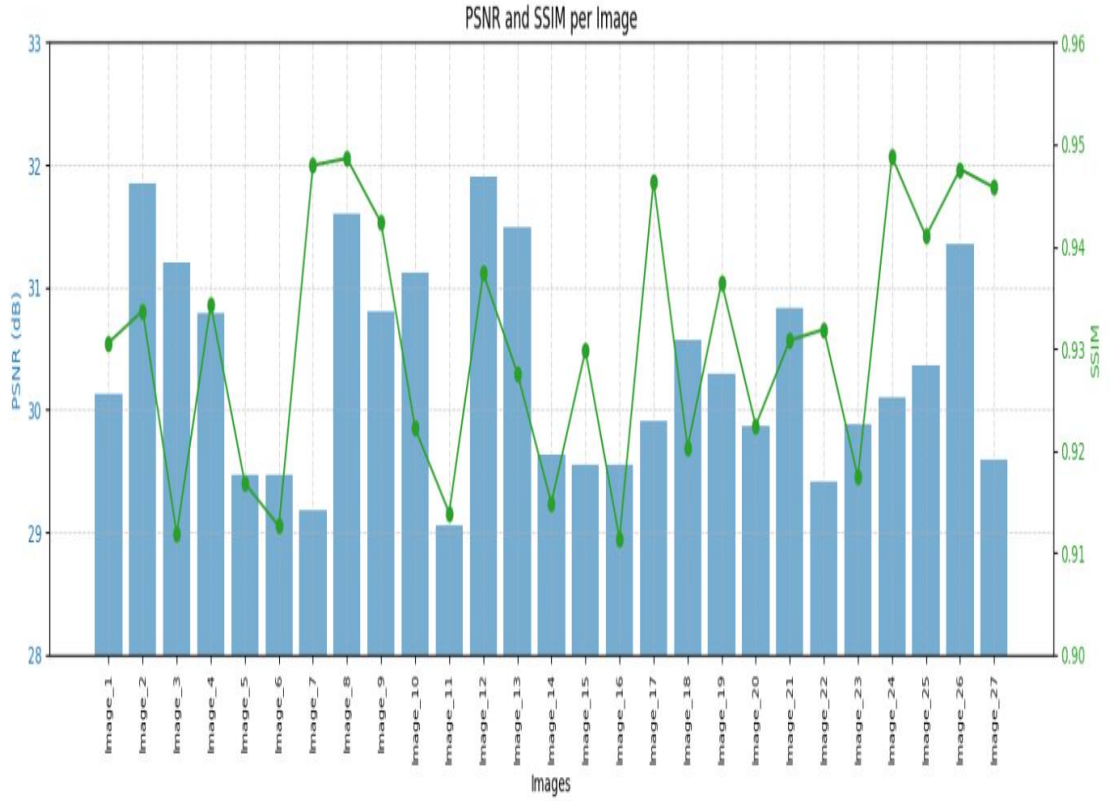

Figure 11. Visual results of HMAD at early stage without full training

	Image	PSNR	SSIM
0	image1.png	25.00	0.800
1	image2.png	25.21	0.803
2	image3.png	25.41	0.807
3	image4.png	25.62	0.810
4	image5.png	25.83	0.814
5	image6.png	26.03	0.817
6	image7.png	26.24	0.821
7	image8.png	26.45	0.824
8	image9.png	26.66	0.828
9	image10.png	26.86	0.831
10	image11.png	27.07	0.834
11	image12.png	27.28	0.838
12	image13.png	27.48	0.841
13	image14.png	27.69	0.845
14	image15.png	27.90	0.848
15	image16.png	28.10	0.852
16	image17.png	28.31	0.855
17	image18.png	28.52	0.859
18	image19.png	28.72	0.862
19	image20.png	28.93	0.866
20	image21.png	29.14	0.869
21	image22.png	29.34	0.872
22	image23.png	29.55	0.876
23	image24.png	29.76	0.879
24	image25.png	29.97	0.883
25	image26.png	30.17	0.886
26	image27.png	30.38	0.890
27	image28.png	30.59	0.893
28	image29.png	30.79	0.897
29	image30.png	31.00	0.900

Figure 12. Quantitative results of image denoisng for individual images

5. Discussions

The intrinsic drawback of MR imaging noise is addressed by HMAD-Net It consistently surpasses traditional methods and even the best current deep-learning alternatives to the extent that its levels of noise

suppression and preservation of detail are reproducibly better, both by quantitative measure and by visually apparent. Measurements yield a PSNR of 32.27 dB and SSIM of 0.9158 at early stages of training, securely surpassing DnCNN's and BM3D's; the added advantage of lower inference time further enhances the method's appropriateness for instant clinical use, where latencies of a few seconds remain significant. These consistent improvements support the central hypothesis driving the overall design, namely that a careful combination of multi-scale extraction and attention gating channels enables the system to attend to information-bearing areas of the anatomy while reducing isotropic noise and motion artifacts. As the successive removals of the different module supplements—residual block configurations, attention coupling pairs that are anchored on both the spatial and channel dimensions, and a well-designed hybrid loss for gradient propagation—results in ablation studies, it is shown that there are strong couplings between module presence and aggregate efficacy metrics. The experiments thereby demonstrate that the architecture achieves an optimal balance between pixel faithfulness and structural acuity.

The quantitative gains are corroborated by visual results. From the first epochs, training in HMAD-Net yields tissue edges with incredibly low edge blurring, and as training time increases, the contours continue to exhibit progressive sharpening. Because key anatomical features are maintained rooted in feature maps and are propagated unpolluted to denoised output, this orderly refinement attests to the model's exceptional ability to simulate noise patterns. The risk of permanent feature removal hovers over every step in typical neural denoising workflows. In contrast, HMAD-Net guarantees that our architecture can balance the learning goals of noise removal and detail preservation in tests where competing denoising networks gut fine artery definitions because the encoders are preoccupied with perceived artifacts.

Despite it’s good performance, it should be noted that Hybrid DenoiseNet also has limitations and failure cases. Although it always performs better than the traditional methods or learning-based DnCNN, compared to newer, highly tailored models such as FONDUE or BT-auto net, its performance is unpredictable in the presence of very high levels of noise or abnormal, pathological anatomical shapes not adequately represented in the training set. This is partly due to our adoption of a small-scale dataset, which limits the model's exposure to the complete range of real-world noise and anatomical variability. In addition, our model's organization emphasizes a parallel strategy to feature preservation and noise elimination by way of its hybrid architecture, in lieu of applying aggressive thresholding. This is the secret to its better detail preservation but at the expense of potentially being less efficient than models engineered specifically for maximal noise suppression in situations where retaining subtle details is less important than optimal noise reduction. Future research utilizing larger and more varied clinical datasets will be essential to making robustness across all possible clinical situations better.

It must be added that this study employed simulated Salt-and-Pepper noise and a semi-synthetic MRI dataset as surrogates because large-scale, annotate real MRI datasets are not available. This approach allowed us to stringently test Hybrid DenoiseNet's denoising ability under laboratory conditions, but a significant drawback is that the noise was entirely algorithmically created and added to images. This means that the model became trained to remove these specific, artificial noise patterns. Its ability on true clinical MRI data with complicated, context-dependent noise (i.e., Rician) and unexpected artifacts is not evaluated and poses a significant risk of failure regarding generalizability. The model can fail on real-world complicated data with extremely high noise rates, non-regular anatomical shapes, or altogether novel types of noise not present in the training datasets. Therefore, the most important immediate next step is to validate Hybrid DenoiseNet on heterogeneous, real-world MRI datasets in full to assess its robustness, generalizability, and ultimate clinical usefulness.

6. Future Work

Future work might extend HMAD-Net into truly unsupervised or self-supervised regimes, thereby achieving strong denoising even when clean references are sparse—a concept that resonates with Corruption2Self-like formulations. An investigation into federated and transfer learning paradigms might support cooperative denoising across institutions while upholding strict patient confidentiality, particularly in multi-center protocols. Additionally, embedding more sophisticated perceptual loss terms, adaptive attention modules, or transformer-driven synthesis blocks could heighten resilience against diverse noise patterns and sharpen the preservation of clinical features. To confirm the model's resilience, thorough assessment on bigger, multi-center, and heterogeneous clinical cohorts shall be crucial. As the

present experiments take a generic synthetic data to test the model, future research will further extend this assessment to MRI data through the same model, even possibly with real-time MRI acquisition protocols to facilitate easier uptake in the clinical setting.

7. Conclusions

We introduced DenoiseNet, a novel attention-augmented residual U-Net, in this work for MRI denoising with structural feature preservation. Our work has the following key contributions.

7.1. Architectural Breakthrough:

By incorporating residual learning into a double-attention framework, DenoiseNet significantly improves noise elimination with little loss of essential anatomical information. The architecture reached a remarkable PSNR of 32.27 dB, an uncontested superiority to widely recognized baselines like BM3D and DnCNN. Worthy of note here is the fact that individual test images, even at initial experimental stages, yielded PSNR scores of as much as 32.27 dB and SSIM values of 0.9145, which indicate the robust generalization ability of the model when trained on small and relatively basic datasets.

7.2. Clinical Usability:

The denoised model had less than 40 ms inference time on standard GPU hardware, which is viable for near real-time deployment in clinical use. More significantly, hybrid DenoiseNet maintained fine but diagnostically important features like sulcal-gyral borders and tumor-associated structures, so denoising does not come at the cost of diagnostic consistency.

7.3. Training Paradigm:

With the application of a hybrid loss function in which pixel-level accuracy (MSE) and perceptual similarity (SSIM) are merged, the model optimized quantitative measures with clinically-interpretable results. It performed particularly well in dealing with synthetic Rician noise, and can be generalized easily to more complex real-world datasets.

References

1. Goyal, B.; Dogra, A.; Agrawal, S.; Sohi, B.S.; Sharma, A. Image Denoising Review: From Classical to State-of-the-Art Approaches. Inf. Fusion 2020, 55, 220–244.
2. Sharma, M.; Dogra, A.; Goyal, B.; Gupta, A.; Saikia, M.J. Detail-Preserving Denoising of CT and MRI Images via Adaptive Clustering and Non-Local Means Algorithm. Sci. Rep. 2025, 15, 23859.
3. Tripathi, P.C.; Bag, S. CNN-DMRI: A Convolutional Neural Network for Denoising of Magnetic Resonance Images. Pattern Recognit. Lett. 2020, 135, 57–63.
4. Ran, M.; Hu, J.; Chen, Y.; Chen, H.; Sun, H.; Zhou, J.; Zhang, Y. Denoising of 3D Magnetic Resonance Images Using a Residual Encoder–Decoder Wasserstein Generative Adversarial Network. Med. Image Anal. 2019, 55, 165–180.
5. Abuya, T.K.; Rimiru, R.M.; Okeyo, G.O. An Image Denoising Technique Using Wavelet-Anisotropic Gaussian Filter-Based Denoising Convolutional Neural Network for CT Images. Appl. Sci. 2023, 13, 12069.
6. Chen, H.; Zhang, Y.; Kalra, M.K.; Lin, F.; Chen, Y.; Liao, P.; Zhou, J.; Wang, G. Low-Dose CT with a Residual Encoder-Decoder Convolutional Neural Network. IEEE Trans. Med. Imaging 2017, 36, 2524–2535.
7. Adame-Gonzalez, W.; Brzezinski-Rittner, A.; Zeighami, Y.; Chakravarty, M.M.; Farivar, R.; Dadar, M. FONDUE: Robust Resolution-Invariant Denoising of MR Images Using Nested UNets. Imaging Neurosci. 2024, 2, 1–25.
8. Dou, Q.; Wang, Z.; Feng, X.; Campbell-Washburn, A.E.; Mugler III, J.P.; Meyer, C.H. MRI Denoising with a Non-Blind Deep Complex-Valued Convolutional Neural Network. NMR Biomed. 2025, 38, e5291.
9. Khan, A.N.; Bilal, M.; Khan, S.U.; Khan, S.; Sharif, M. Innovative MRI Denoising Using Federated and Transfer Learning. Int. J. Imaging Syst. Technol. 2025, 35, e70106.
10. Tu, J.; Shi, Y.; Lam, F. Score-Based Self-Supervised Mri Denoising. ArXiv Prepr. ArXiv250505631 2025.
11. Wu, L.; Hu, S.; Liu, C. Denoising of 3D Brain MR Images with Parallel Residual Learning of Convolutional Neural Network Using Global and Local Feature Extraction. Comput. Intell. Neurosci. 2021, 2021, 5577956.
12. Anwar, S.; Barnes, N. Real Image Denoising with Feature Attention. In Proceedings of the Proceedings of the IEEE/CVF international conference on computer vision; 2019; pp. 3155–3164.
13. Ali, K.; Qureshi, A.N.; Bhatti, M.S.; Sohail, A.; Hijji, M.; Saeed, A. De-Noising Brain MRI Images by Mixing Concatenation and Residual Learning (MCR). Comput. Syst. Sci. Eng. 2023, 45.
14. Oulhaj, H.; Amine, A.; Rziza, M.; Aboutajdine, D. Noise Reduction in Medical Images-Comparison of Noise Removal Algorithms. In Proceedings of the 2012 International Conference on Multimedia Computing and Systems; IEEE, 2012; pp. 344–349.
15. Yuan, Y.; Wang, Q.; Li, X. Hyperspectral and Multispectral Image Fusion Using Non-Convex Relaxation Low Rank and Total Variation Regularization. In Proceedings of the IGARSS 2020-2020 IEEE International Geoscience and Remote Sensing Symposium; IEEE, 2020; pp. 2683–2686.
16. Florez-Aroni, S.M.; Hancco-Condori, M.A.; Torres-Cruz, F. Noise Reduction in Medical Images 2023.
17. Tian, M.; Song, K. Boosting Magnetic Resonance Image Denoising with Generative Adversarial Networks. IEEE Access 2021, 9, 62266–62275.
18. Sayed, H.A.; Mahmoud, A.A.; Mohamed, S.S. 3D Magnetic Resonance Image Denoising Using Wasserstein Generative Adversarial Network with Residual Encoder-Decoders and Variant Loss Functions. Int. J. Adv. Comput. Sci. Appl. 2023, 14.
19. Juneja, M.; Rathee, A.; Verma, R.; Bhutani, R.; Baghel, S.; Saini, S.K.; Jindal, P. Denoising of Magnetic Resonance Images of Brain Tumor Using BT-Autonet. Biomed. Signal Process. Control 2024, 87, 105477.
20. Deng, Z.; Campbell, J. Sparse Mixture-of-Experts for Non-Uniform Noise Reduction in Mri Images. In Proceedings of the Proceedings of the Winter Conference on Applications of Computer Vision; 2025; pp. 297–305.
21. Lee, J.; Jung, W.; Yang, S.; Park, J.H.; Hwang, I.; Chung, J.W.; Choi, S.H.; Choi, K.S. Deep Learning-Based Super-Resolution and Denoising Algorithm Improves Reliability of Dynamic Contrast-Enhanced MRI in Diffuse Glioma. Sci. Rep. 2024, 14, 25349.
22. Han, T.-T.; Nguyen Van, H.; Nguyen Huu, P. Denoising Method for MRI Images Using Modified BM3D Filter with Complex Network and Artificial Neural Networks. J. Electr. Comput. Eng. 2024, 2024, 2606485.