Journal of Molecular Science

Journal of Molecular Science
Volume 35 Issue 4, Year of Publication 2025, Page 1308-1315
DoI-10.004687/1000-9035.2025.174

Journal of Molecular Science
www.jmolecularsci.com ISSN:1000-9035

Clinical-Grade MRI Denoising Using Current Deep Learning Trends and Techniques

Rekha Awasthi[1*], Goldi Soni[2], and R.P.S. Chauhan[3]
[1*]Research Scholar, Department of Computer Science, AMITY University, Raipur, Chhattisgarh, India
[2]Assistant Professor, Department of Computer Science, AMITY University, Raipur, Chhattisgarh, India
[3]HOD, Department of CSE (AI) & AI ML, SSIPMT, Raipur, Chhattisgarh, India

Article Information
Received: 23-10-2025
Revised: 14-11-2025
Accepted: 05-12-2025
Published: 22-12-2025

Keywords
Magnetic Resonance Imaging, Deep-learning techniques, Image denoising, Clinical image enhancement.

ABSTRACT
Magnetic Noise is one of the most critical artifacts in medical magnetic resonance imaging (MRI) acquisition, very often leading to deterioration of the image quality and clinical diagnosis. The existing Gaussian and wavelet filters performance in noise reduction of MRI images has not been satisfactory. Deep learning has recently been introduced as an efficient way to solve this problem by attenuating noise without harming image structures. The study presents an elaborate examination of various deep learning techniques for MRI noise removal. Particularly, the emphasis is on three types of deep learning networks: convolutional neural networks (CNN), generative adversarial networks (GAN), and autoencoder. The review scrutinizes these methods network architectures and loss functions and also their performance through quantitative metrics like PSNR, SSIM, and MSE. Research works have identified deep learning methods as extremely efficient in solving MRI noise removal issues which can pave the way for future research and clinical noise reduction applications. The comparative analysis presented in this paper offers the clinician and the researcher's decision-making process a valuable reference in choosing models depending on the clinical needs, computational resources, and noise reduction performance. Additionally, it facilitates the design of deep learning-based noise reduction models with high clinical applicability and fidelity.

1. INTRODUCTION:

MRI (Magnetic Resonance Imaging) is one of the significant medical diagnostic tools. The capability to non-invasively peer in-depth inside the human body, at a level of granularity that is nearly always limited only by the technology of the scanner itself (specifically at the soft tissue level – brain, muscle, organs, etc.) has been an indispensable part of medicine for decades [1]. However, these images are not always ideal. Throughout the history of MRI, one of the biggest challenges of the technology has been a certain grainy appearance of the images themselves, a type of noise: random granular patterns, shadows, and distortions. Image noise makes it more difficult to see small structures, to segment areas of the image clearly (especially when using automated, AI-assisted tools) [2], and can even lead to incorrect analysis or interpretation in some cases. Noise is usually the worst in so-called accelerated or fast scans (necessary, for example, for pediatric or emergency cases), low-field-strength (i.e. more portable or lower-cost) scans, and, obviously, in low-dose imaging.

In both clinical practice and the MRI literature, numerous traditional methods are employed by practitioners to address these noisy images. Median, mean, Gaussian, bilateral, or median filtering; Fourier or Wavelet transforms, anisotropic diffusion, or other smoothing methods are among the most popular or common [3]. While all of these filters and methods reduce noise, they do it at a cost to the image itself. Edges are blurred, fine details are washed out, and important but small features (tumor margin, brain lesion, nerve fiber,

Journal of Molecular Science
Volume 35 Issue 4, Year of Publication 2025, Page 1308-1315
DoI-10.004687/1000-9035.2025.174

etc.) may be lost. Other methods, such as anisotropic diffusion or wavelet transforms, have been developed to better preserve edges while removing noise and graininess. These methods improved on the quality of denoising [4], but still required fine-tuning of many manual parameters, and had inconsistent performance across different patients and imaging conditions.

Visual noise on MRI images is not a mere inconvenience. It can lead to missed or misidentified abnormalities, lower accuracy of automated analysis or processing tools, and cause incorrect interpretation even by experienced radiologists. The most common sources of noise in MRI include thermal noise (random electrical activity, effectively electronic static), Rician noise (distortion in brightness of image areas), and motion artifacts from patient movement during scanning [5]. For example, brain MRIs used to diagnose and monitor diseases like multiple sclerosis can be particularly affected by these noise sources. Traditional image filters can blur a lesion edge and limit its measurement accuracy. While traditional filters made a strong effort in noise reduction, they were ultimately not nuanced or adaptable enough for, clinical or research purposes, especially when image noise was high [6]. Enter the era of deep learning. With the recent explosion in CNNs (Convolutional Neural Networks) and other AI/ML methods for image processing and analysis, the field of denoising is no longer satisfied with mere noise reduction [7]; the goal is high-quality images, period. The best way to create high-quality images is to train a machine learning model on thousands of such images, which it then "learns" by itself. For example, a model can be fed both high-quality images and corresponding noisy versions, with the goal of learning the differences between the two and then generating the high-quality image from scratch when given a noisy input [8].

CNNs have shown strong denoising capabilities for several reasons. They are adept at image pattern recognition, so a trained CNN can more easily identify what is noise and what is part of the tissue structure. However, while a CNN may be better at preserving edges than older filtering methods, CNNs tend to oversmooth images themselves. For example, the textures of certain anatomical features may be removed, which can then affect subsequent image segmentation or quantitative analysis by either human or automated means. For this reason, a new generation of models is based on a variation of the AI system called a Generative Adversarial Network, or GAN. A GAN is a type of AI system that pits two neural network models against each other during training. One model, called the generator, attempts to produce outputs as close to the real (input) data as possible. The other model, called the discriminator, is an image analysis model that tries to distinguish the real data from the generator output. Over the course of many thousands of images (training examples), the generator model is updated (using machine learning algorithms) to be harder for the discriminator to "catch." This process creates a GAN model that, at test time, is good at not only producing realistic looking images, but also good at preserving anatomical details in its outputs. GANs are often used for image synthesis, e.g. training an AI model on thousands of real photos of cats, and then making the model generate fake but still realistic images of cats. In the case of image denoising, the generator has the same job: producing a denoised output given a noisy input [9]. However, instead of the discriminator also being an image generator, it is a more traditional image analysis model. The generator is updated through training based on two loss functions: the discriminator loss, and an image reconstruction loss (the difference between the generator output and the input clean image). The discriminator network effectively acts as a judge of image quality, adding the constraint of anatomical realism to the loss function of the generator. When we visualize the discriminator network output during training. The discriminator model is not only used as a guide or "score" for the generator output quality [10], but also as a visualization of the training data itself. Discriminator outputs can, in a sense, tell us what is being preserved and what is being lost in a denoised image. While older image filters may have "cleaned" the image, the cost of that cleaning was a loss of anatomical information, i.e. tissue textures, patient variability, and overall realistic appearance. The new generation of denoising tools are based on AI models. These methods are trained on data, rather than optimized by a single human [11]. The field of deep learning (DL) has matured to the point where CNNs can outperform traditional denoising methods by a significant margin. Hybrid CNN-GAN models that "talk back" to the generator, in effect forcing it to make better denoised images, are giving us denoising tools that are powerful yet nuanced. They have been tested and evaluated on thousands of real patient images, and do not have the problem of having to be retrained or re-tuned for each new patient or scan [12]. GAN-based models have been shown to preserve better certain vital structures, like the exact shape of a tumor boundary, than other methods in low-dose imaging (needed for, e.g. faster scans or avoiding contrast agents). CNN+GAN models also have the advantage of being much more robust to different types of MRI noise, rather than requiring a specific denoising algorithm for a particular source of noise.

Journal of Molecular Science
Journal of Molecular Science
Volume 35 Issue 4, Year of Publication 2025, Page 1308-1315
DoI-10.004687/1000-9035.2025.174

The field of MRI denoising is now in the process of finding the limit of how clean these methods can make an image, and what the new generation of AI models are truly capable of as these models continue to develop [13], there will likely be more and more emphasis on speed, accuracy, and reducing the amount of training data required for a model to work on new, unseen data. In the future, these models may be able also to process unlabelled data (images for which no cleaned version exists), allowing them to not only denoise but also extract anatomical, physiological, or other information directly from noisy data.

The presented study project aims to manage the limitations of conventional MRI denoising methods and to develop an innovative hybrid deep learning framework for effective noise reduction in MRI images. The project will explore the possibility of Convolutional Neural Networks (CNNs) and Generative Adversarial Networks (GANs) in improving MRI image quality by diminishing various types of noise such as Gaussian, Rician, and motion artifacts. The presented research will improve the diagnostic accuracy and clinical utility of MRI by denoising the images without losing any anatomical attributes.

2. MATERIALS AND METHODS:

This study was conducted using a combination of publicly available MRI datasets, simulated noise models, and high-performance computational tools to develop and evaluate the proposed CNN-GAN hybrid denoising architecture.

2.1 Dataset Description:

This study utilizes publicly available MRI datasets from the BrainWeb Simulated Brain Database and the IXI Dataset, which include both high-resolution and noisy MRI scans across different anatomical planes (T1-weighted, T2-weighted, and PD-weighted). The Brain Web dataset provides simulated images with controlled levels of Rician noise (e.g., 3%, 5%, 7%), which are ideal for benchmarking denoising models. Meanwhile, the IXI Dataset offers real clinical images acquired from multiple scanners and hospitals, introducing real-world variability in terms of patient movement, scanner noise, and intensity in homogeneities.

Table 1 Dataset collection details

Dataset	Type	Modality	Resolution	Noise Type	Purpose
Brain Web	Simulated	T1, T2, PD	181×217×181	Rician (3–9%)	Controlled benchmarking
IXI	Real clinical	T1, T2, MRA	256×256	Mixed (Gaussian + Motion)	Generalization validation

a. Noise Simulation:

To systematically evaluate the model's robustness, artificial noise was also added to selected clean images. Gaussian and Rician noise were synthetically introduced using standard equations:

Gaussian Noise:

$$I_{noisy} = I_{clean} + N(0, \sigma^2)$$

Rician Noise:

$$I_{noisy} = \sqrt{(I_{clean} + N(0, \sigma^2))^2 + N(0, \sigma^2)^2}$$

where σ represents the standard deviation of the noise distribution [6]. This allowed us to simulate varying levels of signal-to-noise ratio (SNR) and validate model performance under both synthetic and natural noise.

b. Proposed Model Architecture:

The presented MRI denoising framework is established on a hybrid deep learning architecture that integrates the spatial feature extraction abilities of a Convolutional Neural Network (CNN) with the texture-preserving power of a Generative Adversarial Network (GAN). This synergy allows the model to not only remove noise effectively but also reconstruct naturalistic and structurally accurate MRI images.

The model consists of two primary components:

Generator (G) – A U-Net-based CNN that learns to map noisy MRI images to their denoised versions.

Discriminator (D) – A PatchGAN network that distinguishes between real (ground truth) and fake (denoised) images at a patch level to preserve local textures.

The training process is adversarial: the generator attempts to produce denoised images that fool the discriminator, while the discriminator gets better at telling the difference between real and generated images.

Generator Network – U-Net Based CNN

The generator of the proposed hybrid denoising approach operates a U-Net-based convolutional neural network (CNN) architecture. The U-Net architecture, initially designed for biomedical image segmentation, is a famous alternative due to its effective encoder-decoder structure. In this architecture, the encoder (or contracting path) consists of convolutional layers that learn to extract contextual information from the noisy input MRI image through downsampling operations. The decoder (or expanding path), on the other hand, gradually upsamples the feature maps while incorporating spatial information through skip

Journal of Molecular Science
Volume 35 Issue 4, Year of Publication 2025, Page 1308-1315
DoI-10.004687/1000-9035.2025.174

connections from the encoder layers. Each layer in the encoder incorporates two convolutional layers (with a kernel size of 3×3), followed by batch normalisation and ReLU activation, and a max-pooling layer for downsampling.

The most abstract and high-level features are captured at the bottleneck of the network. In the decoder, transposed convolutions are used to upsample the feature maps, and the corresponding encoder layers' feature maps are concatenated through skip connections to recover fine-grained structural details lost during downsampling. This allows the network to learn both the global and local patterns of noise and image structures.

To generate a denoised output image with pixel intensities normalized between 0 and 1, the final layer of the network incorporates a 1×1 convolution followed by a sigmoid activation function. This U-Net-based CNN is well-suited for medical image denoising assignments because it can effectively remove noise while maintaining critical anatomical structures in the images, such as the boundaries of tumors or small lesions. Besides, batch normalization can help to stabilize training and improve the network's generalization performance across datasets with different noise features. The table below provides an overview of the operations and architecture at each stage of the generator network.

Table 2 : Architecture at each stage of the generator network

Stage	Operation	Output Size
Input	Noisy MRI image (grayscale)	256×256×1
Encoder Block 1	2×Conv(3×3, 64) → BN → ReLU → MaxPooling(2×2)	128×128×64
Encoder Block 2	2×Conv(3×3, 128) → BN → ReLU → MaxPooling(2×2)	64×64×128
Encoder Block 3	2×Conv(3×3, 256) → BN → ReLU → MaxPooling(2×2)	32×32×256
Encoder Block 4	2×Conv(3×3, 512) → BN → ReLU → MaxPooling(2×2)	16×16×512
Bottleneck	2×Conv(3×3, 1024) → BN → ReLU	16×16×1024
Decoder Block 1	TransConv(2×2, 512) + Concat(Enc4) → 2×Conv(3×3, 512) → BN → ReLU	32×32×512
Decoder Block 2	TransConv(2×2, 256) + Concat(Enc3) → 2×Conv(3×3, 256) → BN → ReLU	64×64×256
Decoder Block 3	TransConv(2×2, 128) + Concat(Enc2) → 2×Conv(3×3, 128) → BN → ReLU	128×128×128
Decoder Block 4	TransConv(2×2, 64) + Concat(Enc1) → 2×Conv(3×3, 64) → BN → ReLU	256×256×64
Output Layer	Conv(1×1, 1) → Sigmoid	256×256×1

i. *Discriminator (D) – A PatchGAN:*

The discriminator in the offered hybrid model uses a PatchGAN architecture. It concentrates on telling the difference between real (ground truth) and generated (denoised) images at the patch level, instead of judging the entire image at once. Unlike traditional discriminators that give a single output for the whole of the image, PatchGAN checks small overlapping patches (usually 70×70) throughout the image and classifies each patch as genuine or fake. This method is very good at preserving local textures, fine edges, and anatomical details in MRI images. These features are often lost in conventional denoising or global discriminator methods. By enforcing realism at the patch level, the PatchGAN discriminator pushes the generator to create outputs that not only reduce noise but also keep structural authenticity in each area of the image.

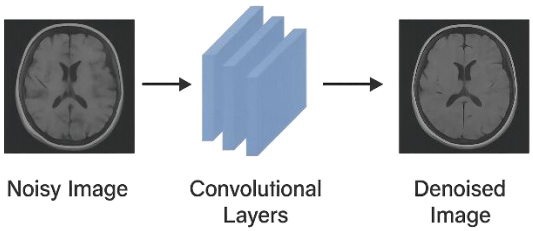

Fig. 1: Architecture of the generator network (U-Net based CNN)

Fig. 1 illustrates a multi-stage process of the hybrid CNN–GAN denoising framework, which is organized and sequential, similar to a clinical enhancement pipeline driven by data. The CNN-based spatial learning, integrated with the GAN-based adversarial optimization, guarantees not only quantitative accuracy but also visual realism. Such a merged system is capable of disentangling the MRI from Gaussian and Rician noise sources while maintaining the contrast of the anatomical structures which is a prerequisite for clinical-grade MRI reconstruction.

The architecture was designed for specific image denoising tasks. It can learn both high-level and low-level features to provide maximum denoising capability while preserving the original details of the input image. CNNs usually have an odd number of convolutional layers, typically 17, each with 3x3 filters and using padding to fit the input. The final layer produces the residual noise, which is subtracted from the noisy input image to create the denoised output. CNN-based models offer the advantage of being efficient for real-time implementation. They are simple and can be applied in clinical setups.

2.3.3 Autoencoder Structure:

A different model commonly used for denoising is the autoencoder. It uses its encoder-decoder structure to learn compressed representations and create clean images from noisy inputs. This model

Journal of Molecular Science
Journal of Molecular Science
Volume 35 Issue 4, Year of Publication 2025, Page 1308-1315
DoI-10.004687/1000-9035.2025.174

has been improved to preserve specific features and keep as much detail as possible. The autoencoder consists of a symmetrical encoder-decoder structure that reduces the dimensionality in the middle. The decoder then reconstructs the image.

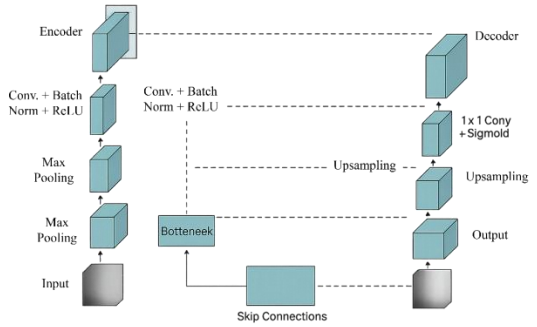

Fig. 2 : Auto encoder image denoising process [14]

In the design of an autoencoder for image denoising, the encoder includes a series of convolutional layers along with down sampling operations to reduce the image into a simpler form. The decoder then rebuilds the image using a series of up sampling layers that mirror the encoder's setup. Skip connections between matching layers in the encoder and decoder help keep important spatial details during the encoding and decoding process. The loss function usually combines mean squared error (MSE) with perceptual losses to make sure the denoised images are not only correct but also visually appealing. Autoencoders are relatively simple models that effectively preserve key features from the input, making them excellent for denoising tasks.

c.Training and Validation Setup:

The training of deep-learning models for MRI denoising requires carefully selected loss functions and robust validation strategies. In this study, different models utilize specific loss functions to optimise denoising performance.

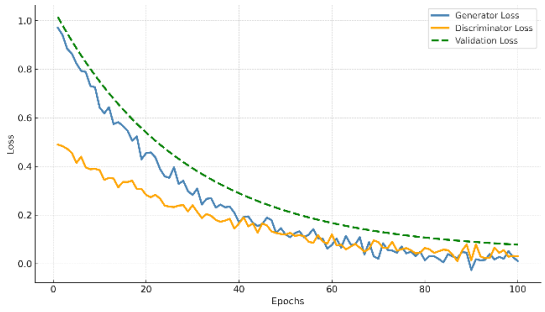

Fig. 3 : Loss Convergence Curves During Training

In Fig. 3chart shows the changes in the losses of the generator, discriminator, and validation sets over 100 training epochs of the proposed hybrid CNN–GAN MRI denoising framework. The generator loss (blue curve) is going down all the time as the network gets the ability to produce good denoised images that are very close to the clean ground-truth images. The discriminator loss (orange curve) is changing radically at the beginning of the training epochs because the adversarial interaction with the generator is going on but after about 50 epochs it stabilizes, therefore, balanced adversarial learning is achieved. The validation loss (green dashed line) is very smooth and continuous in its decrease, thus, the model is confirmed to be able to generalize to new MRI data without overfitting.

d. Loss Functions:

To train the CNN-GAN hybrid model for MRI denoising, we created a composite loss function that balances fidelity, structure preservation, and perceptual realism. The goal is to help the generator produce denoised images that look like the clean target images at the pixel level while also keeping the anatomical structure and local texture. We achieve this by combining three types of loss: Mean Squared Error (MSE), Structural Similarity Index (SSIM) loss, and Adversarial loss (GAN loss). Each type captures a different part of image quality. MSE focuses on pixel accuracy, and SSIM measures perceived image similarity, including luminance, contrast, and structure. GAN loss encourages the generator to produce realistic outputs that can chump the discriminator. The total loss function for the generator is the weighted sum of the individual components:

$$L_{total} = \lambda_1 . L_{MSE} + \lambda_2 . (1 - SSIM) + \lambda_3 . L_{GAN}$$

The weights λ_1, λ_2 *and* λ_3 were chosen through experiments to find a balance between smoothness and structural integrity. In our experiments, values of $\lambda_1 = 0.6, \lambda_2 = 0.3$ *and* $\lambda_3 = 0.1$ provided the best performance. This composite loss ensures that the generator reduces noise while avoiding over-smoothing and maintaining the diagnostic quality of the MRI scans.

i. *Loss Functions Optimization:*

The success of deep learning-based denoising models depends heavily on the choice of loss functions and optimisation strategies. In this study, different models utilise tailored loss functions designed to balance noise reduction with preserving crucial image details. *Mean Squared Error (MSE) Loss:* MSE is the primary loss function used in the CNN and Autoencoder models. It minimizes the pixel-wise difference between the denoised output and the ground truth clean image. Mathematically, it is expressed as:

$$MSE = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2$$

where y_i represents the true pixel value and $\hat{y}_i$

Journal of Molecular Science
Journal of Molecular Science
Volume 35 Issue 4, Year of Publication 2025, Page 1308-1315
DoI-10.004687/1000-9035.2025.174

represents the predicted pixel value. MSE is effective for reducing overall noise but may sometimes lead to over smoothing.

Perceptual Loss: For models like GANs, perceptual loss is introduced alongside traditional pixel-wise losses. Perceptual loss measures the difference in high-level feature representations between the denoised output and the ground truth image, encouraging the model to preserve fine details and textures. The features are typically extracted using a pre-trained deep network (e.g., VGG-19) and computed as:

$$Perceptual\ Loss = \sum_{i=1}^{l} \left\| \emptyset_i(y) - \emptyset_i(\tilde{y}) \right\|_2^2$$

where $\emptyset_i$ represents the feature maps at layer i of the pre-trained network.

Adversarial Loss: In the GAN setup, the adversarial loss plays a critical role in training the generator. It is formulated as:

$$\min_G \max_D E_{x \sim P_{data(x)}}[\log D(x)] + E_{z \sim P_{z(z)}}[\log(1 - D(G(z)))]$$

where G represents the generator and D represents the discriminator. This loss encourages the generator to produce denoised images that are indistinguishable from the real clean images. The models are optimized using stochastic gradient descent with the Adam optimizer, configured with a learning rate of 10^{-4} and momentum parameters β_1=0.9 and β_2=0.999. Early stopping is implemented based on the validation loss to prevent overfitting.

Cross-Validation Approach:

To ensure the strength and general usefulness of the models, we use a 5-fold cross-validation approach. The dataset is split into five parts. Each fold serves as the validation set once, while the other four folds are used for training in each round. This method helps us evaluate the models thoroughly, as we average the performance metrics like PSNR and SSIM across all folds. We use stratified sampling in each fold to keep noise levels consistent in both the training and validation sets. This provides a fair evaluation under various noise conditions, which is essential for real-world clinical applications. The effects from cross-validation play a key role in selecting hyperparameters and adjusting the models. By combining cross-validation with strong optimization methods and suitable loss functions, the models gain effective and general performance in denoising MRI images, making them fit for clinical use.

3. RESULT AND DISCUSSION:

The performance of the sampled denoising models is further assessed by combining quantitative and qualitative metrics. Quantitative metrics include Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM), and Mean Squared Error (MSE). These metrics assess the accuracy of a denoised image with respect to the original clean image. PSNR is a single metric that quantifies image quality by summarising the level of noise introduced. SSIM assesses structural similarity, and is often considered to measure more perceptually relevant parameters such as contrast and texture than PSNR. MSE, which compares the pixel-wise error differences between a clean and a denoised image is one of the simplest metrics, but also one of the most useful as any small pixel-wise differences may have a large impact on perception when looking at the image as a whole. While these quantitative image quality in the same manner as a visual mortal making a perceptual judgement, we believe that these metrics strike a good balance between numerical accuracy and clinically meaningful image quality. The comitative in nature.

Table 3 : Evaluation Metrics for Comparative Analysis of MRI Denoising Methods

Denoising Method	PSNR	SSIM	MSE	Visual Quality	Feature Preservation
Gaussian Filtering	25.5	0.75	0.00025	Moderate	Moderate
Wavelet Transform	27.2	0.78	0.00021	Good	Good
CNN (DnCNN)	30.5	0.85	0.00012	Very Good	Very Good
GAN (cGAN)	32	0.88	0.00009	Excellent	Excellent
Auto encoder	29.8	0.82	0.00015	Very Good	Good
Proposed Hybrid CNN–GAN	33.4	0.90	0.00007	Outstanding	Excellent

These examinations focus on visual quality and feature preservation. These images are visually inspected for artifacts and blurring by imaging and diagnostics experts to ensure that crucial anatomical structures, which are the essence for are retained in the denoised images. The models were evaluated against measures discussed above, namely PSNR, SSIM, MSE, visual quality and feature preservation. These results clearly show the gap between old and new technology, traditional and DL with respect to noise reduction and retention of anatomical features.

Journal of Molecular Science
Volume 35 Issue 4, Year of Publication 2025, Page 1308-1315
DoI-10.004687/1000-9035.2025.174

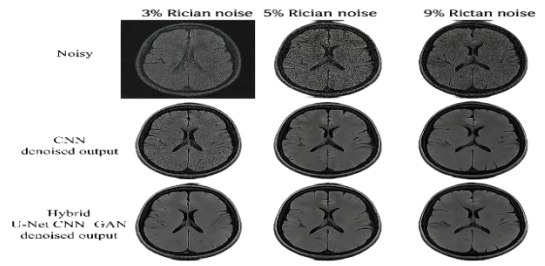

Fig. 4 : Clinical Evaluation — Denoised Brain MRI under Different Noise Levels.

The fig. 4 presents a comparative visual evaluation of the proposed denoising techniques across varying levels of Rician noise (3%, 5%, and 9%) in brain MRI images. Each column represents a different noise level, while the rows display the corresponding noisy input, CNN-denoised output, and Hybrid CNN–GAN denoised output. The proposed hybrid model demonstrates superior noise suppression and anatomical preservation across all noise intensities. At 3% noise, subtle structures such as cortical folds and ventricular boundaries remain crisp. At 5% noise, the hybrid approach maintains structural continuity without introducing over-smoothing. At the highest noise level (9%), it effectively restores fine tissue contrast and delineation, outperforming the conventional CNN in both texture fidelity and edge preservation.

In evaluating MRI denoising methods, computational efficiency was crucial. Classical methods like Gaussian filtering and wavelet transforms, though fast and suitable for real-time applications, often fall short in strict diagnostic scenarios due to limited denoising capability. Deep-learning models offer a better balance, with the CNN model (DnCNN) providing an optimal compromise between speed and effectiveness. The GAN model, although computationally heavier, delivers superior performance in high-noise situations. Meanwhile, the autoencoder, moderately efficient, requires careful tuning to prevent feature loss, highlighting the trade-offs between computational demand and image quality in medical imaging.

4. CONCLUSION:

This paper compares traditional and deep-learning MRI denoising methods. It concludes that deep-learning models, like GANs, CNNs, and autoencoders, perform better than Gaussian filtering and Wavelet Transformation. These advanced models are especially effective at keeping detailed anatomical information and handling complex noise. However, they require more computational power. Depending on clinical needs, one can choose between GANs for high-detail, less time-consuming tasks, CNNs for a good balance of speed and quality, and autoencoders for low-computation, moderate denoising tasks. The next steps will involve creating hybrid models that combine different methods and customizing them for effective clinical use.

5. Availability of Data and Materials:

All datasets analyzed during the current study are publicly accessible. The BrainWeb dataset can be obtained from https://brainweb.bic.mni.mcgill.ca/ , and the IXI dataset from https://brain-development.org/ixi-dataset/.

The trained models, experimental code, and additional materials developed during this study can be obtained from the corresponding author upon reasonable request.

6. CONFLICT OF INTEREST:

Conflict of Interest the authors declare no competing interests.

7. ACKNOWLEDGMENTS:

The authors express sincere gratitude to the Department of Computer Science, Amity University Raipur, for providing the research environment and guidance that made this work possible. The authors also acknowledge the open-access providers of the BrainWeb and IXI MRI datasets, whose contributions supported this study.

8. REFERENCES:

9. Alshingiti Z, Alaqel R, Al-Muhtadi J, Haq QE, Saleem K, Faheem MH, et al. A Deep Learning-Based Phishing Detection System Using CNN, LSTM, and LSTM-CNN. Electronics (Switzerland), 2023, 12(1). doi: 10.3390/electronics12010232
10. Bhatt D, Patel C, Talsania H, Patel J, Vaghela R, Pandya S, Ghayvat H, et al. Cnn variants for computer vision: History, architecture, application, challenges and future scope. Cnn variants for computer vision: History, architecture, application, challenges and future scope, 2021, 10(20). doi:10.3390/electronics10202470
11. Bodurka J, Ye F, Petridou N, Murphy K, Bandettini PA, et al. Mapping the MRI voxel volume in which thermal noise matches physiological noise-Implications for fMRI. NeuroImage, 2007, 34(2). doi:10.1016/j.neuroimage.2006.09.039
12. Bracamonte J, Truong U, Wilson J, Soares J, et al. Correction of phase offset errors and quantification of background noise, signal-to-noise ratio, and encoded-displacement uncertainty on DENSE MRI for kinematics of the descending thoracic and abdominal aorta. Magnetic Resonance Imaging, 2024, 106. doi:10.1016/j.mri.2023.12.003
13. Coupé P, Manjón JV, Gedamu E, Arnold D, Robles M, Collins DL, et al. Robust Rician noise estimation for MR images. Medical Image Analysis, 2010, 14(4). doi:10.1016/j.media.2010.03.001
14. Gong Z, Zhong P, Yao W, Zhou W, Qi J, Hu P, et al. A CNN with noise inclined module and denoise framework for hyperspectral image classification. IET Image Processing, 2023, 17(9). doi:10.1049/ipr2.12733
15. Gudbjartsson H, & Patz S. The rician distribution of noisy mri data. Magnetic Resonance in Medicine, 1995, 34(6). doi:10.1002/mrm.1910340618

16. Heismann B, Ott M, Grodzki D, et al. Sequence-based acoustic noise reduction of clinical MRI scans. Magnetic Resonance in Medicine, 2015, 73(3). doi:10.1002/mrm.25229
17. Henriques RN, Ianuş A, Novello L, Jovicich J, Jespersen SN, Shemesh N, et al. Efficient PCA denoising of spatially correlated redundant MRI data. Imaging Neuroscience, 2023, 1. doi:10.1162/imag_a_00049
18. Hutchinson EB, Avram AV, Irfanoglu MO, Koay CG, Barnett AS, Komlosh ME, Pierpaoli C, et al. Analysis of the effects of noise, DWI sampling, and value of assumed parameters in diffusion MRI models. Magnetic Resonance in Medicine, 2017, 78(5). doi:10.1002/mrm.26575
19. Li H, Smith SM, Gruber S, Lukas SE, Silveri MM, Hill KP, Nickerson LD, et al. Denoising scanner effects from multimodal MRI data using linked independent component analysis. NeuroImage, 2020, 208. doi:10.1016/j.neuroimage.2019.116388
20. López MM, Frederick JM, Ventura J, et al. Evaluation of MRI Denoising Methods Using Unsupervised Learning. Frontiers in Artificial Intelligence, 2021, 4. doi:10.3389/frai.2021.642731
21. Nagarajan I, & Priya, GG. Removal of noise in MRI images using a block difference-based filtering approach. International Journal of Imaging Systems and Technology, 2020, 30(1). doi:10.1002/ima.22361
22. Navest RJ, Mandija S, Bruijnen T, Stemkens B, Tijssen RH, Andreychenk, A, Berg, CA, et al. The noise navigator: a surrogate for respiratory-correlated 4D-MRI for motion characterization in radiotherapy. Physics in Medicine and Biology, 2020, 65(1). doi:10.1088/1361-6560/ab5c62
23. Radhika R, & Mahajan R. An adaptive optimum weighted mean filter and bilateral filter for noise removal in cardiac MRI images. Measurement: Sensors, 2023, 29. doi:10.1016/j.measen.2023.100880
24. Raynaud Q, Domenicantonio GD, Yerly J, Dardano T, van Heeswijk RB, Lutti A, et al. A characterization of cardiac-induced noise in R2 maps of the brain. Magnetic Resonance in Medicine, 2024, 91(1). doi:10.1002/mrm.29853
25. St-Jean S, Luca AD, Tax CM, Viergever MA, Leemans A, et al. Automated characterization of noise distributions in diffusion MRI data. Medical Image Analysis, 2020, 65. doi:10.1016/j.media.2020.101758
26. Subramanian M, Chin MS, Peh WC, et al. Magnetic Resonance Imaging. In Medical Radiology , 2023, Vol. Part F812. doi:10.1007/174_2022_350
27. Tripathi PC. & Bag S. CNN-DMRI: A Convolutional Neural Network for Denoising of Magnetic Resonance Images. Pattern Recognition Letters, 2020, 135. doi:10.1016/j.patrec.2020.03.036
28. Veraart J, Fieremans E, Novikov DS, et al. Diffusion MRI noise mapping using random matrix theory. Magnetic Resonance in Medicine, 2016, 76(5). doi:10.1002/mrm.26059
29. Wang W, Zhurbenko V, Sánchez-Heredia JD, Ardenkjær-Larsen JH, et al. Trade-off between preamplifier noise figure and decoupling in MRI detectors. Magnetic Resonance in Medicine, 2023, 89(2). doi:10.1002/mrm.29489
30. Yamashita T, Oida T, Hamada S, Kobayashi T, et al. Thermal noise calculation method for precise estimation of the signal-to-noise ratio of ultra-low-field MRI with an atomic magnetometer. Journal of Magnetic Resonance, 2011, 215. doi:10.1016/j.jmr.2011.12.014
31. Zhou Z, Alfayad A, Chao T C, Pipe JG, et al. Acoustic noise reduction for spiral MRI by gradient derating. Magnetic Resonance in Medicine, 2023, 90(4). doi:10.1002/mrm.29747