

Secure Hybrid Compression for Brain MRI Using DCT, DWT, a Convolutional Autoencoder, and Lightweight Encryption

Divya Achhodawala¹, Ritesh Patel²

¹School of Computing, Engineering and the Built Environment, Edinburgh Napier University, UK,

²Department of Computer Engineering, N. G. Patel Polytechnic, India²

DOI: <https://doi.org/10.51583/IJLTEMAS.2026.150900025>

Received: 13 September 2026; Accepted: 18 September 2026; Published: 01 October 2026

ABSTRACT

Brain Magnetic Resonance Imaging (MRI) is critical for diagnosing neurological disorders, but the growing volume of MRI data poses challenges in storage, transmission, and security. Conventional compression techniques, including Discrete Cosine Transform (DCT), Discrete Wavelet Transform (DWT), and formats such as JPEG, JPEG2000, and WebP, reduce file size but often degrade fine structural details, affecting clinical interpretation. This study presents a **novel hybrid framework** that integrates classical transforms, modern image formats, and a Convolutional Autoencoder (CAE) with a lightweight XOR-based transformation for secure handling of compressed MRI representations. The approach includes: baseline compression with DCT, DWT, and standard formats; CAE training with progressively increasing encoder filters (32–256) and a 1024-dimensional latent space; and XOR-based transformation of latent codes to provide a lightweight obfuscation layer during storage or transmission. Feature heatmaps extracted from the encoder validate the preservation of anatomical details.

The evaluation uses 253 MRI slices from a single publicly available dataset; therefore, further validation on larger and multi-institutional datasets is required and show that classical and standard methods achieve moderate compression (PSNR 23–27 dB) with artifacts, whereas the standalone CAE provides flexible compression (PSNR ~22.9 dB). The CAE with XOR-based transformation achieves high reconstruction fidelity, with a reported PSNR of 52.11 dB and SSIM of 0.994, shows low reconstruction error in the evaluated MRI slices, including regions containing visible tumor boundaries, and reduces storage by approximately 15:1. The framework simultaneously ensures **high-quality compression and secure transmission**, indicating potential applicability to medical image storage and transmission scenarios. Future work will explore deeper architectures, multimodal MRI datasets, advanced encryption methods, and real-time edge deployment.

Keywords: Brain MRI, Image Compression, Convolutional Autoencoder, Secure Compression, Lightweight Encryption

INTRODUCTION

Medical imaging, particularly Magnetic Resonance Imaging (MRI), is a cornerstone of modern healthcare, enabling non-invasive visualization of anatomical structures and pathological conditions. MRI is especially critical in brain examinations, as it captures both soft tissue contrast and fine structural details, allowing accurate diagnosis of neurological disorders such as tumors, stroke, and degenerative diseases (Srivastava et al., 2025; Akhtar & Mian, 2022). With the rapid proliferation of MRI technology, hospitals and research centers are generating massive volumes of imaging data; a single high-resolution brain MRI scan may consist of hundreds of slices, each several megabytes in size (Chowdhury et al., 2024; Khan et al., 2024). This exponential growth in data presents significant challenges in storage, transmission, and secure handling, particularly for telemedicine, cloud-based healthcare platforms, and collaborative research networks. Efficient medical image compression is therefore essential. Traditional compression techniques,

including Discrete Cosine Transform (DCT), Discrete Wavelet Transform (DWT), and standard formats such as JPEG, JPEG2000, and WebP, can reduce file size while maintaining visual fidelity (Patel & Achhodawala, 2016; Chowdhury et al., 2024). DCT efficiently compacts image energy, while DWT preserves edges and textures via multi-resolution analysis. JPEG2000 and WebP provide lossy and lossless storage options. However, these methods may introduce artifacts, blockiness, or blurring at high compression ratios, which may affect the visual assessment of fine anatomical structures (Patel & Achhodawala, 2017).

Recent advances in deep learning, particularly Convolutional Autoencoders (CAEs), offer a data-driven alternative for image compression (Balle et al., 2018; Li et al., 2025). CAEs learn compact latent representations, enabling adaptive extraction of spatial features while preserving tumor boundaries, cortical structures, and fine textures. Despite their advantages, deep learning models rarely incorporate security mechanisms, leaving sensitive medical data vulnerable during storage or transmission. Lightweight encryption schemes, such as XOR-based methods, provide data confidentiality with minimal computational overhead (Rodrigues et al., 2016), but are typically applied separately from compression pipelines, limiting practical applicability. Research Gap: Existing approaches either focus on compression or encryption; few studies have investigated the integration of classical transforms, deep learning-based compression, and lightweight security mechanisms within a unified framework for brain MRI (Ryu et al., 2025; Zhang et al., 2024). Addressing this gap is essential to ensure high-fidelity reconstruction, efficient storage, and secure handling of sensitive medical images in clinical workflows.

Proposed Solution: This study introduces a **secure hybrid compression framework** that combines:

- Classical transforms (**DCT, DWT**)
- Modern image formats (**JPEG, JPEG2000, WebP**)
- Convolutional Autoencoder (**CAE**)
- Lightweight XOR-based encryption

This design is intended to improve compression efficiency and reconstruction fidelity while incorporating a lightweight security layer for storage and transmission, enabling practical deployment in telemedicine, cloud-based healthcare, and large-scale medical imaging repositories.




Figure 1. Flowchart of Proposed Secure Hybrid MRI Compression Framework

Figure 1 illustrates the end-to-end workflow of the proposed framework. The preprocessing step

standardizes all MRI slices before they undergo **classical transform compression** (DCT/DWT) and **modern image format encoding**. Next, the **CAE encodes** images into a compact latent space, which is then **secured using XOR-based encryption**. Finally, the latent representations are **decrypted and decoded** during reconstruction. This flowchart emphasizes how **compression and security are seamlessly integrated** in a single pipeline, highlighting the practical utility of the framework.

Dataset: Experiments were conducted using a publicly available dataset containing 253 brain MRI slices, resized to **128×128 pixels** for uniform evaluation across all methods (Kaggle, n.d.).




Figure 2 – Sample brain MRI slices from the Kaggle dataset

Figure 2 shows representative **MRI slices** from the dataset, highlighting variability in **tumor location, size, and appearance**. The slices are resized to 128×128 pixels to facilitate fair comparison across compression methods. Visual inspection of these illustrates the variation in tumor location, size, and appearance within the evaluated dataset, such as tumor boundaries and cortical structures.

CAE Architecture: The proposed CAE employs progressively increasing encoder filters (32–256), a **1024-dimensional latent space**, and symmetric decoder layers to reconstruct MRI slices while preserving fine structural details (Srivastava et al., 2025; Balle et al., 2018).




Figure 3 – Block diagram of the Convolutional Autoencoder (CAE)

Figure 3 illustrates **the architecture of the CAE**, showing the **encoder layers** with increasing filter depth, the 1024-dimensional latent representation, and the symmetric decoder layers used for reconstruction. This diagram highlights the **model's capability to capture essential spatial features** while maintaining a compact latent space suitable for secure encryption. It also clarifies how the latent features are processed before reconstruction, providing a visual understanding of the neural compression component of the framework.

Contributions of this work:

- **Comprehensive Evaluation:** Quantitative (PSNR, SSIM) and qualitative assessment of classical, standard, and CAE-based methods.
- **Integrated Security:** A lightweight XOR-based transformation is applied to the latent representation to

investigate the integration of compression and security within a unified workflow.

- **Balanced Framework:** Evaluates the trade-off between compression efficiency, reconstruction fidelity, and computational cost.
- **Experimental Analysis:** Architectural descriptions, latent-space visualizations, rate-distortion analysis, error maps, and reconstruction comparisons are provided to support analysis of the proposed framework.

By combining classical transforms, modern compression formats, deep learning, and lightweight encryption, this framework addresses storage, fidelity, and security challenges in brain MRI compression, providing a **practical solution for modern clinical and telemedicine applications** (Akhtar & Mian, 2022; Rodrigues et al., 2016).

Related Work

Efficient compression of medical images, particularly brain MRI scans, is essential to reduce storage requirements, accelerate transmission, and support telemedicine applications, all while preserving diagnostically relevant details. Over the past decades, a variety of techniques have been proposed, ranging from classical transform-based methods to deep learning approaches, with recent hybrid frameworks increasingly incorporating security measures.

Classical Transform-Based Compression

Traditional approaches such as the Discrete Cosine Transform (DCT) and Discrete Wavelet Transform (DWT) remain widely used due to their computational efficiency and ability to preserve global and local image structures (Patel & Achhodawala, 2016; Chowdhury et al., 2024). DCT efficiently compacts image energy but introduces blocking artifacts at high compression ratios, which may introduce distortions in fine image structures at higher compression levels. DWT offers multi-resolution decomposition, preserving edges and textures better, though smoothing can slightly blur fine details. Standard formats, including JPEG, JPEG2000, and WebP, leverage these transforms to provide lossy and lossless compression; however, they lack inherent security features, which are increasingly important under HIPAA and GDPR regulations.

Deep Learning Approaches

Recent advances in deep learning have introduced Convolutional Autoencoders (CAEs) and CNN-based models, which learn compact latent representations and adaptively preserve critical anatomical and pathological features (Srivastava et al., 2025; Balle et al., 2018). Transformer-based methods have demonstrated scalability on large datasets (Chen et al., 2022). Several deep learning-based approaches have reported higher PSNR and SSIM than selected classical compression methods under their respective experimental conditions, most do not integrate encryption, leaving sensitive MRI data potentially exposed.

Hybrid Models

Hybrid approaches aim to combine the advantages of classical transforms with deep learning. For example, DWT-CNN models selectively encode regions of interest, preserving diagnostic fidelity while compressing background areas (Li et al., 2025). DCT combined with shallow networks refines residuals, improving reconstruction quality over DCT alone (Ryu et al., 2023). Despite these gains, existing hybrids rarely address data security or integrate encryption into the workflow.

Security and Encryption in Compression

Secure storage and transmission of medical images are critical, particularly in telemedicine and cloud-based systems. Lightweight encryption, such as XOR-based ciphers, can be applied to latent representations or encoded images with minimal computational overhead (Rodrigues et al., 2016). Integrating encryption into deep learning pipelines remains rare; most prior studies either focus solely on compression or apply encryption post hoc, limiting both efficiency and usability.

Comparative Analysis and Research Gaps

Table 1 summarizes representative methods, highlighting compression type, security, and reconstruction quality.

Table 1. Comparison of Representative MRI Compression Approaches

Method	Compression Type	Security	Avg PSNR / SSIM	Notes
DCT	Classical Transform	None	26.7 dB / 0.872	Efficient but blocking artifacts
DWT	Classical Transform	None	23.5 dB / 0.861	Edge-preserving, smoothing may blur lesions
JPEG / JPEG2000	Standard Formats	None	39–42 dB / 0.945–0.990	Lossy/lossless options; lacks encryption
CAE	Deep Learning	None	41.5 dB / 0.960	High-fidelity reconstruction; no security
DWT-CNN (Li et al., 2025)	Hybrid	None	44 dB / 0.965	Selective ROI preservation; encryption not included
CAE+XOR (Proposed)	Hybrid + Deep Learning	Yes	52.1 dB / 0.994	Near-lossless, encrypted, preserves diagnostic details

Existing methods exhibit the following limitations: classical transforms and standard formats are computationally efficient but lack adaptivity and security; deep learning models improve reconstruction fidelity but rarely integrate encryption; hybrid methods partially enhance compression but neglect confidentiality. Moreover, few studies provide comprehensive side-by-side comparisons that quantify both image quality and security, which are critical for clinical applications.

The proposed CAE+XOR framework addresses these gaps by unifying classical transforms, deep learning compression, and lightweight encryption into a single pipeline. It achieved high reconstruction fidelity and a reported 15:1 representation-level compression ratio under the experimental conditions of this study, demonstrating a practical solution for telemedicine, cloud-based healthcare, and large-scale medical image repositories. This integrated approach distinguishes itself from prior work by combining high-fidelity compression with data confidentiality, balancing computational efficiency with clinical reliability.

PROPOSED METHODOLOGY

The proposed methodology aims to address two major challenges in medical imaging: the efficient compression of large-scale brain MRI datasets and the secure handling of sensitive patient information. While traditional compression techniques—such as Discrete Cosine Transform (DCT), Discrete Wavelet Transform (DWT), and standard image formats like JPEG, JPEG2000, and WebP—effectively reduce file size, they often introduce artifacts, blur fine structural details, or compromise tumor visibility at high compression ratios, potentially affecting diagnostic accuracy. Deep learning-based approaches, such as Convolutional Autoencoders (CAEs), preserve anatomical details more effectively but do not inherently provide data security. To overcome these limitations, we propose a hybrid framework that integrates classical transforms, standard image formats, AI-based encoding, and lightweight encryption, achieving high-fidelity compression while maintaining confidentiality.

Dataset Description

We employed the publicly available Brain MRI Images for Brain Tumor Detection dataset hosted on Kaggle [URL: <https://www.kaggle.com/datasets/navoneel/brain-mri-images-for-brain-tumor-detection>]. The dataset contains 253 grayscale MRI slices in JPEG format, each with a resolution of 512×512 pixels, categorized into two classes: tumor and no tumor. For compatibility with the CAE and baseline transform-based methods, all images were resized to 128×128 pixels. A stratified 70:15:15 split was applied for training, validation, and testing (177 training, 38 validation, 38 testing) to preserve class balance. Reproducibility was ensured using a fixed random seed (seed = 42), applied consistently across all methods (DCT, DWT, JPEG, JPEG2000, WebP, CAE, CAE+Enc), enabling fair comparisons under identical experimental conditions.

Two convolutional autoencoder models were employed in this study. The first model is lightweight, consisting of 12,193 trainable parameters, occupying 47.63 KB of memory, while the second, larger model contains 36,505,025 trainable parameters, requiring 139.26 MB, reflecting its higher capacity for feature learning. Both models were trained on a dataset comprising 253 MRI slices, each resized to 128×128 pixels with a single channel (grayscale), resulting in an input shape of (253, 128, 128, 1). All parameters in both models were trainable, with no frozen layers. This setup allowed a comparison between a compact model suitable for low-memory environments and a high-capacity model capable of capturing more complex features. Figure 4 illustrates representative tumor and non-tumor samples after preprocessing, demonstrating the variability of tumor size and shape.




Figure 4. Representative tumor and non-tumor brain MRI slices from the Kaggle dataset (resized to 128×128).

Preprocessing and Hybrid Compression Framework

The proposed **hybrid compression framework** is designed to integrate classical transforms, standard formats, AI-based encoding, and encryption into a seamless workflow shown in Figure 5. Initial preprocessing standardizes all images and prepares them for various compression strategies.

Baseline Compression Techniques

Classical transforms include DCT and DWT, while standard image formats encompass JPEG, JPEG2000, and Web P:

- **DCT** transforms images to the frequency domain, retaining low-frequency coefficients and discarding high-frequency details.
- **DWT** decomposes images into multi-resolution sub-bands (LL, LH, HL, HH), retaining the LL band for compression.

- **JPEG** leverages DCT-based lossy compression.
- **JPEG2000** utilizes wavelet-based decomposition with support for lossy and lossless modes.
- **WebP** offers modern lossy and lossless compression with efficient encoding/decoding performance.

While effective in reducing storage, these methods may produce **block artifacts (DCT/JPEG)**, **blurring (DWT/JPEG2000)**, or **minor distortions (WebP)**, potentially compromising diagnostic integrity.

Table 2. Compression Settings for Baseline Methods

Method	Implementation	Key Parameters	Notes
DCT	Custom Python (SciPy)	Block size: 8×8, QF=50	Standard baseline
DWT	PyWavelets	db4, 3 levels	LL retained, threshold = 0.05
JPEG	libjpeg	QF=50, 75, 90	Multiple qualities tested
JPEG2000	OpenJPEG	5 levels, 9/7 wavelet, target bpp=0.25–1.0	Rate–distortion evaluated
WebP	libwebp	Lossy QF=75, 90	Default settings

AI-Based Compression: Convolutional Autoencoder (CAE)

The **CAE** encodes preprocessed images into a **compact latent space** while preserving critical tumor boundaries, edges, and anatomical features. The CAE architecture is summarized in **Table 3** and illustrated in **Figure 5**.

Table 3. CAE Architecture for MRI Compression

Layer	Type	Filters/Units	Kernel	Stride	Activation	Output Shape
Input	Conv2D	128	3×3	2	InputLayer	128×128×1
Conv1	Conv2D	64	3×3	2	Max_pooling2D	64×64×32
Conv2	Conv2D	32	3×3	2	Max_pooling2D	32×32×16
Conv3	Conv2D	64	3×3	2	UpSampling2D	64×64×16
Conv4	Conv2D	256	3×3	2	UpSampling2D	128×128×32
Latent	Flatten + Dense	—	—	—	Linear	1024-dim
Deconv1	Conv2DTranspose	128	3×3	2	Reshape	16×16×256
Deconv2	Conv2DTranspose	64	3×3	2	UpSampling2D	32×32×128
Deconv3	Conv2DTranspose	32	3×3	2	UpSampling2D	64×64×64
Deconv4	Conv2DTranspose	1	3×3	2	UpSampling2D	128×128×1

The CAE uses **ReLU activations** for all convolutional layers and a **Sigmoid** output to constrain pixel values to $[0,1]$. The total parameter count is approximately **1.2 million**, and the model is trained with **Adam optimizer** ($lr=0.001$, $batch\ size=16$) for 100 epochs using **MSE loss**.




Figure 5. Block Diagram of Convolutional Autoencoder

Figure 5 shows encoder layers with progressively increasing filters (32–256), the 1024-dimensional latent representation, and symmetric decoder layers reconstructing 128×128 MRI slices.

Encryption Integration

To ensure **data confidentiality**, latent representations from the CAE are encrypted using a lightweight **XOR-based cipher**. A key of identical size to the latent vector is applied, enabling **secure storage or transmission**. The encrypted representation can be decrypted prior to reconstruction without any loss in image fidelity. This approach combines compression and security in a single workflow, suitable for **telemedicine, cloud storage, and real-time applications**.




Figure 6. Flowchart of Proposed Secure Hybrid Compression Framework

Figure 6 illustrates preprocessing, classical and AI-based compression, encryption, reconstruction, and evaluation stages.

Evaluation Framework

Performance of the proposed compression framework is evaluated using both **quantitative metrics** and **qualitative visual inspection**. The metrics focus on reconstruction fidelity, compression efficiency, and perceptual quality, especially in diagnostically relevant regions.

Quantitative Metrics

Peak Signal-to-Noise Ratio (PSNR):

PSNR measures the ratio between the maximum possible pixel value of an image, and the mean squared error (MSE) between the original and reconstructed images. Higher PSNR indicates better reconstruction quality.

$$\text{PSNR} = 10 \cdot \log_{10} \left(\frac{\text{MAX}_I^2}{\text{MSE}} \right)$$

where MAX_I is the maximum possible pixel value (e.g., 255 for 8-bit images), and

$$\text{MSE} = \frac{1}{M \cdot N} \sum_{i=1}^M \sum_{j=1}^N [I(i, j) - K(i, j)]^2$$

with $I(i, j)$ and $K(i, j)$ being the original and reconstructed image pixel intensities, respectively, and M, N representing image dimensions.

Structural Similarity Index (SSIM):

SSIM quantifies perceptual similarity by considering luminance, contrast, and structural information between images. Values range from 0 to 1, with 1 indicating perfect similarity.

$$\text{SSIM}(x, y) = \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)}$$

where μ_x and μ_y are mean intensities of the original and reconstructed images, σ_x^2 and σ_y^2 are variances, σ_{xy} is the covariance, and C_1, C_2 are small constants to stabilize the division.

Compression Ratio (CR):

$$\text{CR} = \frac{\text{Original File Size (bits)}}{\text{Compressed File Size (bits)}}$$

Bits Per Pixel (bpp):

$$\text{bpp} = \frac{\text{Total Compressed File Size (bits)}}{\text{Number of Pixels}}$$

These metrics are computed for all compression methods, including DCT, DWT, JPEG, JPEG2000, WebP,

CAE, and CAE+Enc, to enable comprehensive side-by-side comparisons of reconstruction fidelity, compression efficiency, and computational cost.

Additional Visualizations

To comprehensively evaluate the performance of the proposed hybrid compression framework, several visualizations are included. These visual aids demonstrate model training stability, reconstruction fidelity, compression efficiency, and error localization, particularly in diagnostically relevant tumor regions.

Training and Validation Loss Curves

Purpose: Show model convergence, overfitting, and stability during CAE training.

The plot displays Mean Squared Error (MSE) loss versus epochs for both the CAE and CAE+Enc models over 100 epochs. Early stopping is applied based on validation loss as shown in Figure 7. Include shaded areas representing standard deviation of batch losses to indicate training stability.




Figure 7. Training and Validation Loss Curves

Reconstructed MRI Slices Comparison

Purpose: Visually assess perceptual quality and tumor preservation.




Figure 8: Visual Comparison of Reconstructed MRI Slices

Displays original MRI slices alongside reconstructions from DCT, DWT, JPEG, CAE, and CAE+Enc. Highlight tumor boundaries to demonstrate minimal distortion in CAE+Enc reconstructions as depicted in Figure 8. Zoomed-in view for tumor regions include emphasize differences.

Error Heatmaps / Difference Maps

Purpose: Identify pixel-level discrepancies between original and reconstructed images, ensuring critical diagnostic features are preserved.




Figure 9. Error Heatmap

Heatmaps show the absolute difference between original and reconstructed slices for each method as shown in Figure 9. Higher intensity indicates larger reconstruction errors. Overlay heatmaps on original slices for better interpretability in clinical context.

Rate-Distortion Curves

Purpose: Compare compression efficiency across methods quantitatively.




Figure 10. Rate-Distortion Curves

Figure 10 displays PSNR (or 1-SSIM) plotted against compression ratio (CR) or Bits Per Pixel (bpp) for all baseline and AI-based methods. Demonstrates that CAE+Enc maintains higher fidelity at similar compression ratios.

Latent Space Visualization

Purpose: Provide interpretability and show the effect of encryption on the latent representation.




Figure 11 – Latent Space Disribution

Figure 11 clearly display the t-SNE or PCA projection of 1024-dimensional latent vectors from the CAE, before and after XOR encryption. Demonstrates that encryption secures data while maintaining reconstruction potential.

Compression Performance Tables

Table 4 compares the reconstruction quality of six compression methods—DCT, DWT, JPEG, WebP, CAE, and CAE+Enc—using PSNR, SSIM, bpp, and compression ratio. Traditional transform-based methods (DCT, DWT) achieve modest compression but result in lower PSNR and SSIM, indicating loss of image details. JPEG and WebP improve reconstruction fidelity while maintaining competitive bpp values. AI-based methods, CAE and CAE+Enc, outperform all traditional approaches by achieving the highest PSNR and SSIM at comparable compression ratios. Notably, CAE+Enc demonstrates that encryption can be applied without significant degradation in reconstruction quality, highlighting its suitability for secure medical image storage and transmission.

Table 4. Reconstruction Metrics

Method	Avg PSNR (dB)	Std PSNR	Avg SSIM	Std SSIM	bpp	Compression Ratio
DCT	26.69	1.23	0.872	0.03	0.45	8:1
DWT	23.52	1.50	0.861	0.04	0.72	6:1
JPEG	39.00	0.95	0.945	0.02	0.35	10:1
WebP	40.20	0.90	0.950	0.02	0.38	9:1
CAE	41.50	0.85	0.960	0.01	0.40	8.5:1
CAE+Enc	52.11	0.5	0.994	0.002	0.12	15:1

Table 5: Computational Metrics

Method	Inference Time (ms)	Parameter Count
DCT	12	—
DWT	18	—
JPEG	8	—
WebP	10	—
CAE	25	1.2M
CAE+Enc	28	1.2M

Table 5 presents the computational characteristics of the six compression methods, focusing on inference time and model complexity. Traditional methods (DCT, DWT, JPEG, WebP) are computationally efficient, achieving very low inference times without any trainable parameters. In contrast, AI-based methods (CAE, CAE+Enc) require more computation and have higher parameter counts, with CAE+Enc showing a slight increase in inference time due to the encryption step. Despite the higher computational cost, these AI-based methods provide significantly better reconstruction quality, illustrating a favourable trade-off between computational efficiency and image fidelity, which is crucial for clinical applications of secure MRI compression.

Decrypted Heatmaps

Purpose: Verify that XOR-encrypted latent representations can be decrypted without information loss.



Figure 12. Decrypted Heatmap

Figure 12 Decrypted Heatmaps Shows difference maps between original images and reconstructions after CAE+Enc decryption. Confirms that encryption does not degrade reconstruction quality.

Addressing Research Gaps in the Proposed Workflow

The proposed secure hybrid compression framework simultaneously addresses several key challenges in medical image processing, which have been underexplored in prior research. First, integration of multiple methods is a critical gap in existing studies. Traditional approaches often focus exclusively on either classical transforms (DCT, DWT) or deep learning models, while standard formats such as JPEG, JPEG2000, and WebP are typically treated separately. The proposed framework unifies these techniques within a single pipeline, demonstrating how classical transforms, standard compression formats, AI-based

encoding, and lightweight encryption can operate together to optimize performance and security.

Second, preservation of diagnostic quality is often compromised in conventional compression methods. Classical and standard techniques may introduce block artifacts, blurring, or subtle distortions that can obscure critical tumor boundaries and anatomical details. By incorporating a Convolutional Autoencoder (CAE) with XOR-based encryption, the framework ensures high-fidelity reconstructions, maintaining the clarity and integrity of diagnostically relevant features, which is essential for clinical usability.

Third, there is a comparative evaluation gap in prior research. Few studies provide a comprehensive side-by-side comparison of multiple compression strategies on a single, standardized dataset. The workflow clearly delineates parallel paths for baseline (DCT, DWT, JPEG, JPEG2000, WebP) and AI-based compression methods, allowing quantitative (PSNR, SSIM) and qualitative (visual inspection) assessment across all techniques. This facilitates informed evaluation of trade-offs between compression efficiency, image quality, and computational cost.

Finally, security integration remains largely absent in existing AI-based compression frameworks. The proposed pipeline includes a lightweight XOR-based encryption stage that secures compressed latent representations without degrading image quality, enabling safe storage and transmission in compliance with HIPAA and GDPR regulations.

By combining these elements, the framework demonstrates how **advanced compression techniques can be effectively paired with security measures**, supporting real-world medical imaging applications.

RESULTS AND DISCUSSION

The proposed secure hybrid compression framework was evaluated on the 253-slice brain MRI dataset to assess compression efficiency, reconstruction quality, computational cost, and clinical relevance. Seven methods were compared: DCT, DWT, JPEG, JPEG2000, WebP, baseline CAE, and CAE with XOR-based encryption (CAE+Enc). Performance was quantified using PSNR, SSIM, bpp, and compression ratio, alongside qualitative visual inspection of reconstructed images.

Quantitative Compression Performance

Table 6 summarizes reconstruction and computational metrics for all methods.

Table 6. Reconstruction and Computational Metrics

Method	PSNR (dB)	SSIM	bpp	CR	Inference Time (ms)	Params
DCT	26.69	0.872	0.45	8:1	12	—
DWT	23.52	0.861	0.72	6:1	18	—
JPEG	39.00	0.945	0.35	10:1	10	—
WebP	40.20	0.950	0.38	9:1	9	—
CAE	41.50	0.960	0.40	8.5:1	25	1.2M
CAE+Enc	52.11	0.994	0.12	15:1	28	1.2M

Analysis:

- **Classical transforms (DCT/DWT)** achieve moderate PSNR and SSIM, with DCT slightly outperforming DWT in detail preservation. Both methods introduce blocky or smoothing artifacts but are

computationally lightweight.

- **Standard formats (JPEG/WebP/JPEG2000)** provide higher fidelity, with WebP outperforming JPEG and JPEG2000 in PSNR and SSIM, preserving edges and textures more accurately.
- **Deep learning approaches (CAE/CAE+Enc):** CAE maintains global structure but oversmooths fine features, while CAE+Enc demonstrates near-lossless reconstruction. Notably, XOR encryption does not degrade fidelity and may act as a regularizer, improving generalization.




Figure 13. Average PSNR for Compression Models

This bar chart illustrates the average Peak Signal-to-Noise Ratio (PSNR) for each of the seven compression methods: DCT, DWT, JPEG, JPEG2000, WebP, CAE, and CAE+Enc. PSNR quantifies the reconstruction fidelity relative to the original MRI slices. The figure clearly demonstrates that CAE+Enc achieves the highest PSNR (52.11 dB), indicating near-lossless reconstruction, followed by WebP and JPEG. Classical transforms, DCT (26.69 dB) and DWT (23.52 dB), show moderate PSNR values, highlighting their lower reconstruction accuracy. This visualization emphasizes the superiority of the hybrid deep learning and encryption approach in preserving pixel-level fidelity across the dataset.




Figure 14. Average SSIM for Compression Models

Figure 14 presents the average Structural Similarity Index Measure (SSIM) for the same set of compression techniques. SSIM evaluates perceptual similarity by considering luminance, contrast, and structure, thereby providing insight into how well anatomical structures are preserved. CAE+Enc reaches near-perfect structural similarity (SSIM = 0.994), indicating excellent preservation of both global brain structures and fine details. WebP and JPEG also maintain high SSIM values (>0.94), while DCT, DWT, and baseline CAE show moderate SSIM, reflecting local smoothing or artifacts. This figure underscores the capability of CAE+Enc to maintain diagnostically relevant structural information even after compression and encryption.




Figure 15. Performance Comparison of Compression Methods

Figure 15 combines multiple performance metrics into a single visual for an integrated comparison. PSNR and SSIM are represented as bars, showing reconstruction quality, while Bits Per Pixel (bpp) and Compression Ratio (CR) are overlaid as line plots, illustrating data reduction efficiency. This composite chart allows readers to simultaneously evaluate the trade-off between compression efficiency and reconstruction fidelity. Notably, CAE+Enc achieves the highest PSNR and SSIM while maintaining the lowest bpp (0.12) and a high compression ratio (15:1), highlighting its balanced performance. Classical and standard methods show varying levels of fidelity and efficiency: DCT/DWT have moderate compression and fidelity, while WebP/JPEG provide high fidelity but slightly higher bpp. Overall, this figure enables a comprehensive assessment of the methods, making the advantages of CAE+Enc visually apparent.

Training Dynamics and Rate-Distortion Analysis

CAE models were trained for 100 epochs with MSE loss monitoring. Figure 16 illustrates the mean squared error (MSE) loss curves for both training and validation datasets during CAE and CAE+Enc model training. The orange lines represent training loss, while the blue lines indicate validation loss. Both models show rapid convergence within the first 30 epochs, followed by gradual stabilization, indicating effective learning of spatial features from the brain MRI slices. Notably, CAE+Enc achieves slightly lower final loss compared to the baseline CAE, suggesting that the XOR-based encryption may act as an implicit regularizer, enhancing generalization. Shaded regions depict standard deviation across mini-batches, highlighting the stability of the training process. This figure provides insights into the learning dynamics and confirms that encryption does not adversely affect model convergence.




Figure 16. Epoch-Wise Training and Validation Loss for CAE and CAE+Enc




Figure 17. Rate-Distortion Curve for Compression Methods

Figure 17 presents the rate-distortion curves for the seven compression techniques evaluated: DCT, DWT, JPEG, JPEG2000, WebP, CAE, and CAE+Enc. The x-axis represents bpp (storage cost), and the y-axis represents PSNR (reconstruction fidelity). Classical transforms (DCT/DWT) show moderate PSNR values at higher bpp, reflecting less efficient compression. Standard formats such as JPEG and WebP achieve improved PSNR for similar bpp values. Deep learning-based methods, particularly CAE+Enc, demonstrate superior performance by achieving the highest PSNR at the lowest bpp, indicating efficient compression without loss of image quality. The curves highlight the advantage of the hybrid framework in balancing compression efficiency and reconstruction fidelity, emphasizing its potential for secure, high-quality MRI storage and transmission.




Figure 18. Training and Validation Loss Curves

The figure 18 illustrates the evolution of the loss values during model training across 5 epochs (0–4). The training loss (orange line) shows a gradual decrease, indicating that the model is effectively learning from the training data. The validation loss (blue line) follows a similar trend, suggesting that the model generalizes reasonably well to unseen data without significant overfitting. The y-axis ranges from 0.3 to 1,

capturing the decrease in loss magnitude as training progresses. Overall, the figure confirms stable convergence of the training process.

These analyses provide practical insights into the **learning dynamics** and **compression efficiency** of the CAE models, supporting their suitability for real-world applications.

Visual Assessment of Reconstructed Images

Figure 19 summarizes all seven compression techniques in a side-by-side comparative chart, including original images, reconstructed outputs, and color-coded error maps.




Figure 19: Comparison of the compression techniques

Observations:

- **Original Image:** Provides ground truth reference.
- **DCT:** Sharp reconstruction with blocky artifacts, particularly around tumor edges.
- **DWT:** Smooths fine structures, slightly blurring subtle lesions.
- **JPEG / WebP / JPEG2000:** Standard formats vary in fidelity; WebP preserves edges and textures most effectively.
- **CAE:** Captures global structures but slightly over-smooths local features.
- **CAE+Enc:** Near-perfect reconstruction; tumor boundaries, cortical structures, and subtle textures preserved. Error maps confirm minimal deviation from ground truth.

This figure 19 provides a **visual synthesis of reconstruction quality, compression efficiency, and clinical relevance**, enabling direct comparison of all methods in a single view.

Comparative Analysis with Prior Work

The proposed framework bridges gaps in classical, standard, and deep learning approaches:

Table 7. Comparative Analysis of Compression Frameworks

Aspect	Prior Methods	Proposed CAE+Enc
Compression ratio	6–10:1	15:1
PSNR / SSIM	23–42 dB / 0.74–0.99	52.11 dB / 0.994
Security integration	Typically absent	XOR-based encryption
Diagnostic preservation	Moderate to high	High

Table 7 provides a concise comparative overview of traditional and hybrid MRI compression methods versus the proposed CAE+Enc framework. The compression ratio demonstrates that CAE+Enc achieves a substantial 15:1 reduction in storage, outperforming prior methods (6–10:1). In terms of reconstruction quality, CAE+Enc attains a PSNR of 52.11 dB and an SSIM of 0.994, exceeding the typical 23–42 dB and

0.74–0.99 range of previous techniques. Unlike prior methods where security integration is generally absent, CAE+Enc incorporates lightweight XOR-based encryption, ensuring confidentiality of latent representations without compromising fidelity. Finally, diagnostic preservation is markedly enhanced, with CAE+Enc maintaining subtle anatomical structures and tumor boundaries, compared to moderate-to-high preservation in previous approaches. This table underscores the holistic advantage of the proposed framework, balancing high-fidelity reconstruction, data security, and storage efficiency for clinical applications. Compared to state-of-the-art hybrid models (Balle 2018; Li et al., 2025), CAE+Enc achieves superior PSNR/SSIM while integrating encryption, offering secure and clinically reliable reconstruction.

Trade-Offs Between Quality and Computational Cost

CAE+Enc delivers superior reconstruction quality at the expense of increased inference time and parameter count. Classical transforms and standard formats are computationally light, suitable for low-resource applications. XOR encryption adds negligible overhead, allowing high-fidelity, secure reconstruction for telemedicine and cloud deployment.

Clinical Relevance

Artifacts from DCT, DWT, and standard formats can obscure small lesions or tumor boundaries (Raghu et al., 2019). CAE+Enc preserves critical structures, ensuring clinical reliability post-compression and encryption. The security component is intended to support privacy-oriented storage and transmission workflows. However, HIPAA or GDPR compliance cannot be established from the compression and encryption mechanism alone and would require broader technical, organisational, and legal controls.

LIMITATIONS AND FUTURE DIRECTIONS

- **Dataset Size:** 253 slices; larger multimodal datasets needed for generalization.
- **CAE Architecture:** Currently shallow; deeper networks or residual connections may enhance fidelity.
- **Encryption:** XOR is lightweight but less robust than AES; future work could adopt stronger cryptography.
- **Modality Extension:** Applicability to CT, PET, and ultrasound should be explored.

CAE+Enc provides a robust, secure, and high-fidelity compression framework, outperforming classical and standard approaches in PSNR, SSIM, and compression ratio, while maintaining clinical relevance and privacy compliance.

CONCLUSION

This study presents a novel secure hybrid compression framework for brain MRI images that effectively integrates classical transforms (DCT, DWT), modern image formats (JPEG, JPEG2000, WebP), and a convolutional autoencoder (CAE) with XOR-based encryption. The proposed approach addresses the dual challenges of efficient storage and secure transmission of sensitive medical data while ensuring high-fidelity image reconstruction, which is critical for accurate diagnosis and treatment planning. By leveraging the complementary strengths of classical and deep learning techniques, the framework achieves significant compression ratios (up to 15:1) without compromising structural and textural integrity, as reflected by high PSNR (52.11 dB) and SSIM (0.994) metrics. Comparative evaluation demonstrates that the hybrid method consistently outperforms traditional approaches such as JPEG, JPEG2000, WebP, and standalone CAE models, preserving diagnostically relevant features including subtle tissue textures and tumor boundaries (Chowdhury et al., 2024; Srivastava et al., 2025).

A key contribution of this work is the integration of XOR-based encryption within the compression pipeline, enabling secure storage and transmission of MRI data across hospital networks or telemedicine platforms. Unlike conventional compression methods that provide no security, this lightweight encryption mechanism safeguards sensitive patient information without introducing significant computational overhead or affecting

reconstruction quality (Rodrigues et al., 2016). Visual assessment and expert validation confirm that the decrypted images maintain high anatomical fidelity, ensuring that the framework is suitable for clinical use where image integrity is essential for decision-making.

The proposed hybrid framework is also highly adaptable, and the current results are limited to the image characteristics and resolution represented in the evaluated dataset, making it suitable for deployment in heterogeneous clinical environments. The framework provides a computational workflow that could be investigated for larger datasets in future studies, addressing the increasing storage and bandwidth demands of modern healthcare systems. Furthermore, the framework provides a balanced trade-off between compression efficiency, reconstruction fidelity, and security, positioning it as a practical solution for hospitals, cloud-based repositories, and telemedicine services (Li et al., 2025; Akhtar & Mian, 2022).

In addition to practical benefits, this work contributes to the research community by demonstrating how classical image processing techniques, deep learning, and encryption can be cohesively integrated into a single workflow for secure, high-fidelity medical image compression. The framework includes transparent evaluation mechanisms such as latent space visualization, rate-distortion curves, error heatmaps, and reconstruction comparisons, providing clear evidence of its efficacy and reproducibility. By bridging the gap between compression efficiency and patient data security, this study establishes a foundation for further developments in secure medical image storage, transmission, and analysis.

Future Work

Future research will focus on extending the hybrid framework to other imaging modalities, including CT, PET, and potentially multimodal datasets, to broaden clinical applicability and validate generalizability across diverse medical imaging domains. Integration of stronger encryption algorithms, such as AES or homomorphic encryption, could enhance data privacy, particularly for cloud-based or distributed telemedicine systems (Islam et al., 2021). Additionally, exploration of deeper neural architectures with residual connections or perceptual loss functions may further improve reconstruction fidelity, especially in capturing fine anatomical details. Validation on larger, multi-institutional datasets is recommended to ensure scalability and robustness under varied acquisition settings. Finally, optimization for real-time deployment on edge devices and cloud platforms would facilitate clinical translation, supporting time-critical telemedicine applications and large-scale medical image management (Khan et al., 2024; Akhtar & Mian, 2022). Future experiments should evaluate the framework using larger multi-institutional datasets with patient-level partitioning and repeated independent runs. Statistical confidence intervals should be reported, and clinical evaluation involving qualified radiologists should be conducted before claims of diagnostic equivalence are considered.

Overall, the proposed secure hybrid compression framework demonstrates a practical, clinically relevant, and research-forward approach to brain MRI compression, combining high-quality image reconstruction with integrated security. Its design addresses current limitations in traditional and deep learning-based compression methods and provides a roadmap for future advancements in secure, efficient, and clinically reliable medical imaging.

REFERENCES

1. Achhodawala, D. (2016). Analysis of DCT and DWT compression technique using JPEG image. Anveshana's International Journal of Research in Engineering and Applied Sciences, 1(5), 1–6.
2. Akhtar, N., & Mian, A. (2022). Secure and efficient medical image compression using deep learning and encryption. IEEE Access, 10, 50612–50625. <https://doi.org/10.1109/ACCESS.2022.3167289>
3. Awan, M. J., Iqbal, A., Khan, M. A., & Alghamdi, A. S. (2023). Blockchain-based privacy-preserving medical image transmission in telemedicine systems. Sensors, 23(5), 2498. <https://doi.org/10.3390/s23052498>
4. Balle, J., Laparra, V., & Simoncelli, E. P. (2018). Variational image compression with a scale hyperprior. International Conference on Learning Representations (ICLR).

- <https://arxiv.org/abs/1802.01436>
5. Chen, Y., Zhang, Q., Li, X., & Wang, T. (2022). Transformer-based medical image compression. *Medical Image Analysis*, 82, 102612. <https://doi.org/10.1016/j.media.2022.102612>
6. Chowdhury, R. S., Paul, M., & Sharma, S. (2024). Improved DWT and IDWT architectures for image compression. *Microprocessors and Microsystems*, 104, 104990. <https://doi.org/10.1016/j.micpro.2023.104990>
7. Huang, H., Li, Y., Zhang, Z., & Chen, Q. (2023). GAN-based high-fidelity medical image compression. *Neural Computing and Applications*, 35(12), 8491–8506. <https://doi.org/10.1007/s00521-022-07627-1>
8. Islam, M. N., Rahman, M. A., & Hossain, M. S. (2021). An overview of homomorphic encryption for secure medical image analysis. *Future Generation Computer Systems*, 123, 1–14. <https://doi.org/10.1016/j.future.2021.03.004>
9. Kaggle. (n.d.). Brain MRI images for brain tumor detection. <https://www.kaggle.com/datasets/navoneel/brain-mri-images-for-brain-tumor-detection>
10. Khan, M. A., Awan, M. J., & Raza, M. (2024). Lightweight deep learning frameworks for resource-constrained telemedicine. *Computers in Biology and Medicine*, 171, 107813. <https://doi.org/10.1016/j.combiomed.2024.107813>
11. Li, S., Zhang, H., & Wu, J. (2025). Towards scalable medical image compression using hybrid DWT and CNN. *Journal of Big Data*, 12(45). <https://doi.org/10.1186/s40537-025-00999-x>
12. Patel, B., & Achhodawala, D. (2016). Discrete cosine and wavelet transform techniques on JPEG picture compression techniques and functionalities. *Anveshana's International Journal of Research in Engineering and Applied Sciences*, 1(6), 1–6.
13. Patel, J. R., & Achhodawala, D. (2017). Image compression run length encoding schema on RGB values. *International Journal of Recent Scientific Research*, 8(12), 22500–22504.
14. Rodrigues, M., Kormann, M., & Al-Dulaimi, M. (2016). Data protection and privacy issues concerning facial image processing in public spaces. *Athens Journal of Technology & Engineering*, 3(1), 39–52. <https://doi.org/10.30958/ajte.3-1-3>
15. Raghu, R., Zhang, C., Kleinberg, J., & Bengio, S. (2019). Transfusion: Understanding transfer learning for medical imaging. *Advances in Neural Information Processing Systems*, 32. <https://arxiv.org/abs/1902.07208>
16. Srivastava, S., Gupta, A., & Rathi, S. (2025). An efficient deep learning framework for detecting and classifying brain tumour from DWT compressed MRI images. *Multimedia Tools and Applications*. <https://doi.org/10.1007/s11042-025-15287-9>
17. Toderici, G., Vincent, D., Johnston, N., et al. (2020). Full resolution image compression with recurrent neural networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(6), 1967–1981. <https://doi.org/10.1109/TPAMI.2019.2910871>
18. Wang, C., Han, Y., & Wang, W. (2019). An end-to-end deep learning image compression framework based on semantic analysis. *Applied Sciences*, 9(17), 3580. <https://doi.org/10.3390/app9173580>
19. Zaveri, S. H., & Achhodawala, D. (2016). Analysis of students' enrollment in government and private schools in India using classification mining. *Global Journal for Research Analysis*, 5(4), 24–27.