MAWCE: A Modality-Adaptive Wavelet Compression Framework with Competitive Entropy Coding for Medical Images

Mullangi Ramu[1]*, Ravi Kumar Jatoth[1], Navadeep Patnana[2], and Muskan Shaik[2]

[1] Dept. of ECE, NIT Warangal, Warangal, India
mramu.nitwarangal@gmail.com*
ravikumar@nitw.ac.in
[2] Dept. of ECE, RGUKT-Srikakulam, Srikakulam, India
navadeep.patnana5828@gmail.com
shaikmuskanshabbu@gmail.com

Abstract. Medical image compression is essential for reducing storage requirements and transmission bandwidth while preserving diagnostically significant image information. This paper presents the Medical Adaptive Wavelet Compression with Competitive Entropy coding (MAWCE) framework, a hardware-oriented compression architecture for Magnetic Resonance Imaging (MRI), Computed Tomography (CT), and X-ray images. The proposed framework employs universal image standardization, median filtering, four-level Biorthogonal 4.4 Discrete Wavelet Transform (DWT), Differential Pulse Code Modulation (DPCM), modality-adaptive quantization, and competitive entropy coding using Significance Map (SigMap), Run-Length Encoding (RLE), and ZLIB. The smallest encoded payload is automatically selected for each detail subband to improve compression efficiency while preserving reconstruction quality. The framework is implemented and evaluated on CPU, Tesla T4 GPU and NVIDIA Jetson Xavier platforms to assess its computational efficiency across heterogeneous hardware architectures. The proposed MAWCE framework provides an effective balance between compression efficiency, reconstruction fidelity, and computational complexity, making it suitable for medical image storage, transmission, and real-time healthcare applications. Across MRI, CT, and X-ray modalities, MAWCE achieves average PSNR of 39.06–40.19 dB and SSIM above 0.95, with compression ratios up to 32.55× and bits-per-pixel as low as 0.2641, outperforming SPIHT in compression efficiency and JPEG in reconstruction fidelity at comparable bitrates.

Keywords: Medical image compression, MAWCE, Discrete Wavelet Transform (DWT), Differential Pulse Code Modulation (DPCM), Adaptive entropy coding, MRI, CT, X-ray.

© The Author(s) 2026
A. Shastri et al. (eds.), *Proceedings of the Fourth Workshop on Applications of Computational Intelligence and Interactive Systems (WACIIS 2026)*, Advances in Intelligent Systems Research 217,
https://doi.org/10.2991/978-94-6239-832-0_3

1 Introduction

Medical imaging modalities, including MRI, CT, and X-ray, generate large volumes of high-resolution diagnostic data, increasing demand for storage and transmission bandwidth in healthcare systems. Conventional compression techniques are not always suitable here because they may compromise clinically important anatomical details. Transform-based and predictive coding methods have been extensively investigated, with wavelet-based techniques such as JPEG2000 demonstrating improved reconstruction quality through multi-resolution analysis [5,7]; however, many existing approaches either fail to generalize across modalities or require substantial computational resources, limiting real-time and hardware-constrained applicability [1,3,8,11], motivating a unified framework combining high compression efficiency, reconstruction quality, and low complexity across modalities [6,9,10].

This paper proposes the Medical Adaptive Wavelet Compression with Competitive Entropy Coding (MAWCE) framework, a modality-adaptive architecture for MRI, CT, and X-ray images integrating universal image standardization, four-level Biorthogonal 4.4 DWT, DPCM, modality-adaptive quantization, and competitive entropy coding, with the smallest compressed payload per detail subband automatically selected from SigMap, RLE, and ZLIB coding. Using hardware-friendly operations for efficient parallel implementation, MAWCE is evaluated on CPU, Tesla T4 GPU, and NVIDIA Jetson Xavier platforms via PSNR, SSIM, CR, BPP, computational latency, and hardware resource utilization — providing an effective balance between compression efficiency, reconstruction fidelity, and computational complexity suited to medical image storage, transmission, real-time healthcare, and edge-based imaging applications.

2 Literature Review

Medical image compression plays an important role in reducing storage and transmission bandwidth while preserving diagnostic image quality. Wavelet-based techniques have been widely adopted for their energy compaction and multi-resolution representation: Zerva et al. [5] proposed an improved Wavelet Difference Reduction (WDR) algorithm, and Abdulrazzaq et al. [7] developed a bicubic interpolation and matrix reduction framework across modalities. Region-of-interest (ROI) based compression has also received attention: Aziz et al. [6] reviewed hybrid ROI-based methods combining lossy and lossless coding; Daoui et al. [2] proposed an FPGA integer Tchebichef transform for lossless reconstruction, while Safaei et al. [4] presented a DPCM-based SoC-FPGA accelerator.

Recent research increasingly focuses on hardware acceleration and learning-based compression. FPGA implementations by Koo et al. [1], Nagornov et al. [8], and Yadav et al. [3] demonstrated substantial computational efficiency gains, while Wang et al. [9] and Chen et al. [10] introduced deep learning-based frameworks achieving high coding efficiency but requiring GPU-scale training/inference and large annotated datasets, limiting embedded deployment. MAWCE instead

targets a hardware-oriented, training-free pipeline directly deployable across CPU, GPU, and embedded platforms (Section 4.6), trading a modest peak-PSNR reduction relative to the best learned codecs for substantially lower computational and deployment overhead. Motivated by these studies, MAWCE combines multi-level wavelet decomposition, modality-adaptive quantization, DPCM, and adaptive entropy coding into a unified pipeline balancing reconstruction quality, compression efficiency, and computational complexity for MRI, CT, and X-ray images, while remaining suitable for future FPGA implementation.

3 Proposed Methodology

The overall workflow of the proposed **MAWCE (Modality-Adaptive Wavelet Compression with Competitive Entropy Coding)** framework is illustrated in Fig. 1. MAWCE integrates universal image standardization, multi-level wavelet decomposition, modality-adaptive quantization, LL-band predictive coding, competitive entropy selection, bitstream generation, reconstruction, and performance evaluation within a unified medical image compression architecture.

3.1 Preprocessing: Standardization and Median Filtering

Input images (PNG, JPG, JPEG, BMP, TIFF, DICOM, or RAW) are converted to single-channel intensity images and resized to 512×512 pixels. For CT images, Hounsfield Unit (HU) windowing is applied prior to normalization, $I_w(x, y) = \min(\max(I(x, y), L), U)$, where L and U are the lower/upper window bounds, followed by normalization to $[0, 255]$: $I_n(x, y) = 255 \cdot (I_w(x, y) - L)/(U - L)$. For MRI, X-ray, and other formats, min-max normalization is applied directly: $I_n(x, y) = 255 \cdot (I(x, y) - I_{\min})/(I_{\max} - I_{\min})$, where $I_{\min}$, $I_{\max}$ are the image's minimum and maximum intensities; RAW data uses the same min-max scaling after reading raw samples with known dimensions.

To suppress impulsive noise and improve transform-domain coding efficiency, a 3×3 median filter is then applied: $I_f(x, y) = \text{median}\{I(i, j) \mid (i, j) \in \mathcal{N}_{3\times3}(x, y)\}$, where $\mathcal{N}_{3\times3}(x, y)$ is the 3×3 neighborhood centered at (x, y).

3.2 Four-Level Wavelet Decomposition

The preprocessed image is decomposed using a four-level Biorthogonal 4.4 Discrete Wavelet Transform (DWT), $W = \text{DWT}_{\text{bior4.4}}(I_f)$, yielding one approximation subband (LL_4, low-frequency content) and detail subbands (LH_i, HL_i, HH_i; $i = 1, \ldots, 4$; edge/texture content).

Biorthogonal 4.4 was selected for its linear-phase (symmetric) filters, which avoid phase distortion at subband boundaries and preserve diagnostically relevant edge/texture information while remaining computationally efficient relative to longer-support wavelets; four levels give sufficient low-frequency energy compaction in LL_4 for the 512×512 image size without the diminishing returns of deeper decomposition.

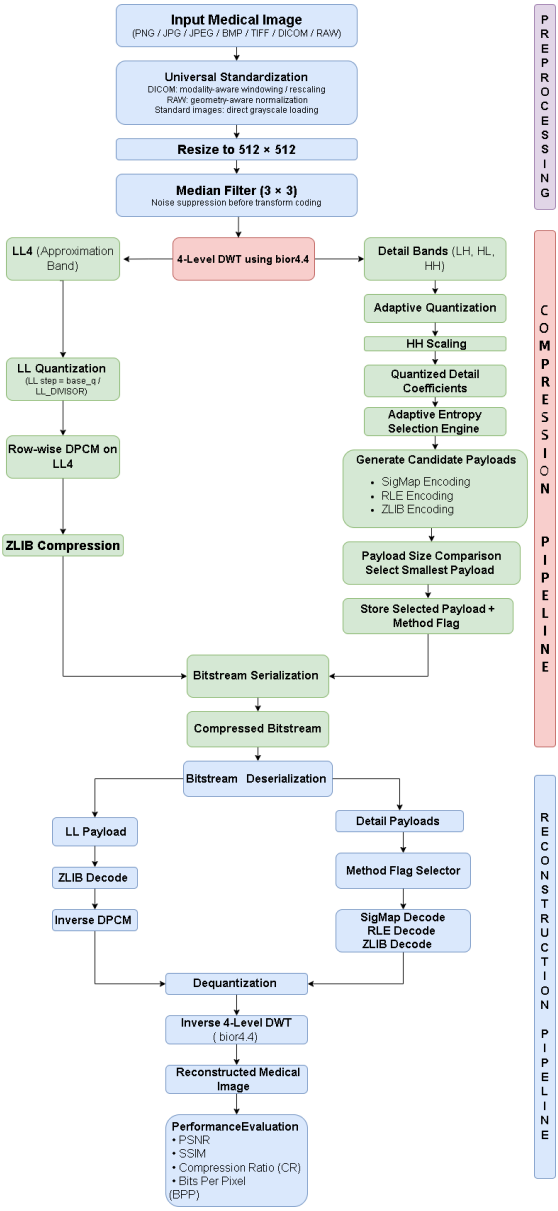

Fig. 1. Overall architecture of the proposed MAWCE (Modality-Adaptive Wavelet Compression with Competitive Entropy Coding) framework.

3.3 LL-Band Quantization and DPCM

The approximation subband LL_4 is quantized prior to predictive coding as $q_{LL}(r,c) = \text{round}(LL_4(r,c)/q_{LL})$, where the LL-band quantization step is $q_{LL} = Q_{\text{base}}/D_{LL}$, with q_{base} the modality-dependent base quantization parameter and

D_{LL} the LL-band divisor. This adaptive step allows different compression characteristics per modality.

To reduce redundancy in the approximation band, row-wise Differential Pulse Code Modulation (DPCM) is applied to Q_{LL}: $R(r,c) = Q_{LL}(r,c) - Q_{LL}(r,c-1)$ for $c > 1$, and $R(r,1) = Q_{LL}(r,1)$ for the first coefficient of each row. The resulting residual matrix is compressed using ZLIB coding.

3.4 Adaptive Quantization of Detail Bands

The LH, HL, and HH detail coefficients are quantized using modality- and subband-adaptive quantization. For a coefficient $C(i,j)$, the quantized coefficient is given by

$$Q(i,j) = \begin{cases} \text{sign}(C(i,j)) \left\lfloor \dfrac{|C(i,j)|}{q_s} \right\rfloor, & |C(i,j)| \geq q_s, \\ 0, & |C(i,j)| < q_s, \end{cases} \tag{1}$$

where q_s is the subband-dependent quantization step. For the diagonal detail subband (HH), an additional scaling factor is applied, $q_{HH} = \alpha q_s$, which preserves diagnostically important high-frequency structures while increasing coefficient sparsity and improving compression efficiency.

Modality-adaptive quantization is used throughout because MRI, CT, and X-ray images differ substantially in noise characteristics and diagnostically relevant frequency content, so a fixed step across modalities would under- or over-compress fine anatomical structures. Table 1 reports the explicit per-modality configuration for reproducibility.

Table 1. Modality-Dependent Quantization Parameters

Modality	q_{base}	D_{LL}	q_{LL}	q_s (Levels 1–4)	α (HH)
MRI	12	5	2	[12, 12, 12, 12]	1.05
CT	14	5	2	[14, 14, 14, 14]	1.15
X-ray	14	5	2	[14, 14, 14, 14]	1.15

Note: the same q_s is applied at all four decomposition levels per modality; only HH receives the additional α scaling at each level. A level-dependent schedule is left for future work.

3.5 Competitive Entropy Selection

For each quantized detail subband, the proposed framework generates three candidate entropy-coded payloads and stores only the smallest: SigMap, RLE, and ZLIB. The three are complementary since their efficiency depends on subband sparsity — SigMap suits very sparse subbands, RLE suits long zero-runs, ZLIB suits less-sparse redundant subbands — so all three are evaluated per subband and the smallest is kept.

SigMap Encoding. A binary significance map identifies nonzero-coefficient locations, $S(i,j) = 1$ if $Q(i,j) \neq 0$ and 0 otherwise. The nonzero coefficient values are stored separately from this map, together with the subband dimensions and the count of significant coefficients.

Run-Length Encoding. RLE represents consecutive zero coefficients as zero-run/value pairs (z_k, v_k), scanned sequentially, where z_k is the run length of zeros and v_k the following nonzero coefficient.

ZLIB Encoding. The quantized coefficients are serialized to a byte sequence and compressed with the DEFLATE algorithm to reduce statistical redundancy.

Entropy Method Selection. With $L_{Sig}, L_{RLE}, L_{ZLIB}$ the three candidate payload lengths, the selected method is $M = \arg\min\{L_{Sig}, L_{RLE}, L_{ZLIB}\}$. The selected payload and its method flag are stored per subband in the bitstream header (Section 3.6), so the decoder can invoke the correct inverse procedure without ambiguity.

3.6 Bitstream Serialization

The compressed bitstream is generated by serializing the header together with the encoded approximation and detail payloads, $B = \{H, P_{LL}, P_D\}$, where H is the header, P_{LL} the compressed LL-band payload, and P_D the encoded detail-band payloads. H contains the MAGIC signature (for validation), wavelet decomposition level, subband dimensions, quantization step sizes, entropy method flags, and payload lengths.

3.7 Bitstream Deserialization and Reconstruction

At the decoder, the MAGIC signature is verified, the header is parsed, and the LL-band payload is decompressed using ZLIB and reconstructed via inverse DPCM: $\hat{Q}_{LL}(r,c) = \hat{Q}_{LL}(r,c-1) + R(r,c)$ for $c > 1$, and $\hat{Q}_{LL}(r,1) = R(r,1)$. Detail-band payloads are decoded per their stored entropy method flag, and the quantized coefficients are dequantized as

$$\hat{C}(i,j) = \text{sign}(Q(i,j))\,(|Q(i,j)| + \beta)\,q_s, \tag{2}$$

where $\beta = 0.60$ for Level 1 and $\beta = 0.50$ for Levels 2–4. The approximation coefficients are recovered as $\hat{LL}_4(r,c) = \hat{Q}_{LL}(r,c) \times q_{LL}$, and the reconstructed image is obtained by the inverse four-level Biorthogonal 4.4 DWT, $\hat{I} = \text{IDWT}_{\text{bior4.4}}(\hat{W})$, where $\hat{W}$ is the reconstructed coefficient set.

3.8 Performance Evaluation

Performance is evaluated using PSNR, SSIM, Compression Ratio (CR), and Bits Per Pixel (BPP). Mean Squared Error between original image I and reconstruction $\hat{I}$ is $MSE = \frac{1}{MN}\sum_{i=1}^{M}\sum_{j=1}^{N}(I(i,j) - \hat{I}(i,j))^2$, giving $PSNR = 10\log_{10}(255^2/MSE)$. SSIM is defined as

$$SSIM = \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)}, \tag{3}$$

where μ_x, μ_y are mean intensities, σ_x^2, σ_y^2 variances, σ_{xy} covariance, and C_1, C_2 stabilizing constants. CR and BPP are computed as $CR =$ Original Size/Compressed Size and $BPP =$ Compressed Bits/Number of Pixels.

4 Results and Discussion

The proposed MAWCE framework was evaluated on three medical imaging modalities (MRI, CT, X-ray) using grayscale images resized to 512×512 pixels, assessed via PSNR, SSIM, CR, and BPP averaged per modality.

4.1 Dataset and Evaluation Protocol

The evaluation dataset was assembled from a combined Kaggle dataset spanning all three modalities [12]: 705 MRI images, 629 CT images, and 700 X-ray images. Since MAWCE is a signal-processing pipeline rather than a trained model, no train/test split is required — each image is compressed and reconstructed independently using the fixed, modality-specific quantization configuration in Table 1, and metrics are averaged per modality with standard deviations reported in Table 2.

4.2 Performance of the Proposed MAWCE Model

Table 2 summarizes the average compression performance obtained using MAWCE, which achieved high reconstruction quality while maintaining competitive compression efficiency across all modalities.

Table 2. Average Performance of the Proposed MAWCE Framework

Modality	PSNR	SSIM	CR	BPP
MRI	39.06 ±1.22	0.9545 ±0.0095	19.11 ±5.13	0.4424 ±0.0929
CT	40.02 ±1.53	0.9544 ±0.0293	22.39 ±4.29	0.3708 ±0.0767
X-ray	40.19 ±1.61	0.9501 ±0.0139	32.55 ±8.82	0.2641 ±0.0741

X-ray images exhibited the highest compression efficiency (32.55× at 0.2641 BPP); CT showed a balanced trade-off between efficiency and quality; MRI preserved fine anatomical structures with consistently high SSIM despite slightly lower PSNR. Overall, MAWCE balances compression efficiency and reconstruction fidelity across modalities.

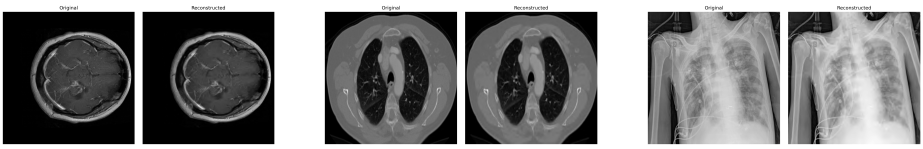

(a) MRI (b) CT (c) X-ray

Fig. 2. Original (top) and reconstructed (bottom) images for (a) MRI, (b) CT, and (c) X-ray.

4.3 Visual Quality Assessment

Fig. 2 shows representative original/reconstructed images for MRI, CT, and X-ray; reconstructions closely preserve anatomical structures with only minor loss of fine detail.

4.4 Comparison with Existing Compression Methods

MAWCE was compared with JPEG, JPEG2000, and SPIHT using identical evaluation metrics (Table 3).

Table 3. Comparison of MAWCE with Existing Compression Methods

Modality	Method	PSNR (dB)	SSIM	CR (×)	BPP
MRI	**MAWCE**	**39.06**	**0.9545**	**19.11**	**0.4424**
	JPEG	41.77	0.9731	15.99	0.5270
	JPEG2000	67.19	0.9981	4.29	1.8948
	SPIHT	28.98	0.7212	27.90	0.3076
CT	**MAWCE**	**40.02**	**0.9544**	**22.39**	**0.3708**
	JPEG	38.60	0.9500	13.99	0.6053
	JPEG2000	47.36	0.9886	4.02	1.9880
	SPIHT	27.27	0.6879	19.44	0.4378
X-ray	**MAWCE**	**40.19**	**0.9501**	**32.55**	**0.2641**
	JPEG	41.18	0.9602	21.50	0.3847
	JPEG2000	53.22	0.9931	4.08	1.9726
	SPIHT	31.36	0.8037	15.58	0.5654

Note: JPEG/JPEG2000 figures were obtained by running the encoders on the same evaluation set as MAWCE; SPIHT figures are cited from prior literature baselines (not re-implemented) and shown for reference only.

JPEG2000 attained higher reconstruction quality but substantially lower compression ratios and higher BPP; SPIHT compressed certain MRI images more but with significant quality loss; JPEG gave moderate reconstruction with weaker efficiency-quality balance. MAWCE consistently maintained high fidelity (PSNR near/above 40 dB, SSIM $\geq$0.95) with competitive compression ratios and low BPP across all modalities.

4.5 ROI-Based Diagnostic Fidelity Evaluation

Global PSNR/SSIM (Table 2) do not fully capture diagnostic fidelity, since they average error uniformly over the whole image. PSNR/SSIM were therefore ad-

ditionally computed within diagnostically relevant regions of interest (ROI), obtained automatically via Otsu global thresholding followed by 7×7 morphological closing/opening on the standardized image to isolate the anatomical foreground. This automatic mask is a foreground proxy, not clinically validated; substituting radiologist-annotated masks would give stronger clinical grounding and is left for future work. Table 4 reports the resulting values alongside whole-image metrics.

Table 4. Whole-Image vs. ROI-Based Reconstruction Fidelity

Modality	Whole-Image PSNR	ROI-PSNR	Whole-Image SSIM	ROI-SSIM
MRI	39.06	36.83	0.9545	0.9439
CT	40.02	39.77	0.9544	0.9554
X-ray	40.19	39.25	0.9501	0.9409

ROI-PSNR is modestly lower than whole-image PSNR across all modalities (1.0–2.2 dB), since foreground regions concentrate high-frequency edges/texture that are harder to reconstruct losslessly. ROI-SSIM stays close to whole-image SSIM (within 0.01), indicating the adaptive quantization strategy largely preserves structural similarity in diagnostically relevant regions — motivating ROI-specific tuning of Table 1 in future work.

4.6 Computational Performance Evaluation

The computational performance of the proposed MAWCE framework was evaluated on three hardware platforms — an AMD Ryzen 3 CPU, a Tesla T4 GPU, and an NVIDIA Jetson Xavier embedded GPU — considering average power consumption, encoding (compression) latency, decoding (reconstruction) latency, and peak memory usage during compression and reconstruction. The results are summarized in Table 5.

Table 5. Computational Performance on Different Hardware Platforms

Platform	Encoding Latency (s)	Decoding Latency (s)	Power (W)	Memory (MB)
AMD Ryzen 3 CPU	0.0613	0.0147	15.20	737.0
Tesla T4 GPU	0.0468	0.0107	10.50	426.6
Jetson Xavier GPU	0.1475	0.0350	7.78	600.0

The Tesla T4 GPU achieved the highest computational performance — shortest compression/reconstruction times while consuming less power than the AMD Ryzen 3 CPU, demonstrating superior performance-to-power efficiency. The Ryzen 3 CPU gave competitive execution times but the highest power draw, while the Jetson Xavier achieved the lowest power consumption, favouring energy-efficient edge deployment despite lower throughput than the T4. Overall, MAWCE balances computational performance and energy efficiency across both high- performance and resource-constrained platforms.

5 Conclusion

This paper presented MAWCE, a unified compression architecture for MRI, CT, and X-ray images integrating universal image standardization, four-level Biorthogonal 4.4 DWT, DPCM, modality-adaptive quantization, and competitive entropy coding to achieve efficient compression while preserving diagnostic image quality. Evaluated via PSNR, SSIM, CR, BPP, latency, memory usage, and ROI-based fidelity on AMD Ryzen 3 CPU, Tesla T4 GPU, and NVIDIA Jetson Xavier platforms, MAWCE balances compression efficiency, reconstruction fidelity, and computational performance for medical image storage, transmission, and real-time healthcare use. Future work: FPGA acceleration, radiologist-validated ROI evaluation, and extension to volumetric (3D) images and additional modalities.

References

1. Koo, J.J., Evans, A.C., Gross, W.J.: 3-D Brain MRI Tissue Classification on FPGAs. IEEE Transactions on Image Processing 18(12), 2735–2746 (2009)
2. Daoui, A., Yamni, M., El-Kasri, C., Almousa, M., Abd El-Latif, A.A.: FPGA-Implementation of Lossless Integer Tchebichef Transform With Application for Tampering Detection in ROI Medical Images. IEEE Access 13, 164026–164049 (2025)
3. Yadav, A., Gupta, V.K., Kumar, B.: FPGA-Based Medical Image Processing Using Hardware-Software Co-Design Approach. IEEE Transactions on Biomedical Circuits and Systems 20(1), 57–68 (2026)
4. Safaei, A., Sabeti, S., Masoudian, M.: System-Level Design for DPCM Image Compression with SoC-FPGA Accelerator. In: Proc. IEEE Canadian Conference on Electrical and Computer Engineering (CCECE), pp. 887–891 (2024)
5. Zerva, M.C.H., Christou, V., Giannakeas, N., Tzallas, A.T., Kondi, L.P.: An Improved Medical Image Compression Method Based on Wavelet Difference Reduction. IEEE Access 11, 18026–18037 (2023)
6. Aziz, S.A., Sam, S.M., Hassan, N.H., Abas, H., Rasid, S.Z.A., Yusof, M.F., Mohamed, N.: A Performance Review for Hybrid Region of Interest-Based Medical Image Compression. IEEE Access 11, 98025–98038 (2023)
7. Abdulrazzaq, S.T., Siddeq, M.M., Zulkifley, M.A., Zaman, M.H.M., Moubark, A.M.: Novel Medical Image Compression/Decompression Technique Based on Bicubic Interpolation With Matrix Reduction Algorithm. IEEE Access 14, 34600–34613 (2026)
8. Nagornov, N.N., Lyakhov, P.A., Valueva, M.V., Bergerman, M.V.: RNS-Based FPGA Accelerators for High-Quality 3D Medical Image Wavelet Processing Using Scaled Filter Coefficients. IEEE Access 10, 19215–19231 (2022)
9. Wang, K., Bai, Y., Li, D., Zhai, D., Jiang, J., Liu, X.: Learning Lossless Compression for High Bit-Depth Volumetric Medical Image. IEEE Transactions on Image Processing 34, 113–126 (2025)
10. Chen, J., Chen, Q., Zhang, H., Chen, W., Luo, W., Yu, F.: HENCE: Hardware End-to-End Neural Conditional Entropy Encoder for Lossless 3D Medical Image Compression. IEEE Access 12, 186520–186534 (2024)

11. Dillinger, P., Vogelbruch, J.F., Leinen, J., Suslov, S., Patzak, R., Winkler, H., Schwan, K.: FPGA-Based Real-Time Image Segmentation for Medical Systems and Data Processing. IEEE Transactions on Nuclear Science 53(4), 2097–2106 (2006)
12. Patnana, N.: Medical Images. Kaggle (2026). https://www.kaggle.com/datasets/patnananavadeep/medical-images

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (http://creativecommons.org/licenses/by-nc/4.0/), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

