SvedbergOpen
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: https://www.svedbergopen.com/

Research Paper Open Access

Robust Digital Image Watermarking to Protect the Privacy of Medical Images

Dayanand Savakar[1], Pandurang Biradar[2], Anand Ghuli[3], Vishal Sarur[4]

[1]Department of Computer Science, Bagalakote University, Jamakhandi, Karnataka, India. E-mail: dgsavakar@gmail.com, Orcid: 0000-0001-5900-2594
[2]Department of Computer Science, Rani Channamma University, Belagavi, Karnataka, India. E-mail: pandurangmca@gmail.com, Orcid:0000-0002-7335-311X
[3]Department of M.C.A, B.L.D.E. A's V. P. Dr. P. G. Halakatti College of Engineering & Technology., Vijayapur-586103, Karnataka, India. E-mail: mca.anand@bldeacet.ac.in, Orcid:0000-0002-4173-0283
[4]Department of M.C.A, B.L.D.E. A's V. P. Dr. P. G. Halakatti College of Engineering & Technology., Vijayapur-586103, Karnataka, India. E-mail: vishalsaroor@gmail.com, Orcid:0009-0008-6287-0817

Abstract
The patient privacy, data integrity, and unauthorised usage and distribution of diagnostic information in the form of images made the medical image watermarking as a promising research area. Although the encryption techniques have helped to protect and transit the medical image data, still there is need of strong verifiable and protective method. This article proposes a framework known as Dimensional Phase Coherence Watermarking (DPCW). DPCW is a semi-blind, phase-based transform which embeds watermark into grayscale medical cover image by protecting diagnostic fidelity. There are four stages of DPCW: i) Binary watermark payload is expanded into seven-dimensional orthogonal phase-coefficient with the help of Fourier-basis lift. ii) The resulting phase-tensor is projected into 2-dimensional complex-valued coherence matrix, whose phase angles are adapted to the local spatial-frequency energy of the cover image. iii) Multi-state superposition encoding is applied that distributes watermark bits across eight statistically independent coefficient groups. (iv) Lorenz attractor is applied that drives a deterministic temporal permutation, parameterised by (σ = 10, ρ = 28, β = 8/3) with a private seed key K_t. The watermark embedding process is done by $H_w = H + \alpha \times E$ with α = 0.001. And the extraction requires only the seed key K_t and the coherence-phase index; there is no need of original cover image. Experiment is conducted on 50 grayscale medical images of size 256 × 256, 8-bit, comprising 25 chest X-ray and 25 brain MRI slices from publicly available benchmark datasets. The watermark with a size of 64 × 64 binary payload show PSNR = 85.4 ± 1.2 dB and SSIM = 0.9997 ± 0.0001. The overall 23 types of attacks covering noise, compression, filtering, and geometric distortion were tested, the mean NC is 0.9783 and the overall robustness score is 95.99% is obtained. CPU embedding time is measured as 0.31 s per image. The results show that DPCW is a viable tool in assuring medical image integrity, offering a reproducible, reusable method to imperceptible watermark embedding and attack-resilient visual information protection.

Keywords: Medical image; Digital watermarking; Phase coherence; Chaos-based encoding; Semi-blind watermarking; Embedding; Attack; Robustness

1. Introduction

Medical imaging techniques like computed tomography (CT), magnetic resonance imaging (MRI), ultrasound, and X-ray are the major support for the modern clinical diagnostics. The various storage and transformation platforms expose them to various types of active attacks. Encryption limits these risks during transmission, but it provides no verifiable protection once a file is decrypted or shared across institutional boundaries. The digital watermarking has added one more layer of protection to the clinical images and its lifecycle[1,4,5,21,25,26]. The watermarking scheme embeds imperceptible and recoverable image carrying delicate information about the patient into pixel coefficient domain of cover image. There are three major requirements in the medical applications are observed: i) Robustness ii) Fidelity iii) Security. Spatial-domain strategies [20] are computationally simple but fragile. Transform-domain strategies based on DWT, DCT, and SVD [6,8,9] are more robust but typically achieve PSNR values of 42-65 dB, which, while adequate for natural images, leave non-trivial distortion in high-dynamic-range medical scans. The approaches like hybrid and deep-learning[1,17,18,19] enhanced performance further but have imposed complex training pipelines or huge dataset requirements that have limited in deployment of resource-constrained clinical needs. The zero-watermarking techniques[2,7,10]

avoid the modification completely, but are not in position to provide verification of integrity because no data is embedded in the image itself.

The identified gap this work address is the combination of very high fidelity (PSNR > 80 dB for 8-bit grayscale images), semi-blind extraction (no cover image needed at the time of extraction), and robustness within a computationally grounded, reproducible framework. This gap is addressed with the DPCW framework, which spreads watermark across a 7-dimensional orthogonal phase representation and deterministic chaotic temporal permutation is applied. DPCW achieves PSNR > 85 dB and mean NC > 0.97 across 23 attacks on standard medical image sets.

The contributions of this work are: i) Well-defined 7-dimensional orthogonal phase-lift embedding operator and its inverse in closed-form, enabling to the reproducible semi-blind watermarking of grayscale medical images. ii) Multi-state superposition encoding layer that distributes watermark information across eight statistically independent coefficient groups, providing redundancy under partial signal removal. iii) Chaos-driven temporal permutation layer parameterised by the Lorenz attractor that strengthens resistance to united signal-processing and geometric attacks without increasing the embedding strength. iv) Comprehensive evaluation on 50 benchmark medical images under 23 attack scenarios, with statistical variability reporting, failure-case analysis, and computational cost measurement. v) Critical taxonomy of six categories of medical image watermarking methods that identifies the specific limitation each category fails to overcome and explains how DPCW addresses those gaps.

The remainder of this paper is organised as follows. Section 2 surveys and critically taxonomises related work. Section 3 formulates the problem mathematically. Section 4 describes the DPCW framework in full algorithmic detail. Section 5 defines evaluation metrics. Section 6 presents experimental results, failure cases, and computational cost analysis. Section 7 concludes and identifies directions for future work.

2. Related Work

The literature review on Medical image watermarking (MIW) is done on six categories based on their primary technical approach and the limitations they leave unresolved. The summary below has identified the limitation and motivation on the present work.

2.1 Spatial-Domain Methods

Spatial-domain techniques embed watermark directly into pixel values, giving low computational cost and ease in implementation [20]. Mehrabi et al. [20] given the technique on robustness to cropping up to 70% and salt-and-pepper noise density up to 90% using a purely spatial scheme. Ebrahimnejad et al. [10] tailored high-density salt-and-pepper noise through an Adaptive Removal of Salt and Pepper Noise (ARSPN) post-processing step at extraction. The basic limitation of spatial methods is their vulnerability to frequency-domain attacks -JPEG compression, Gaussian blur, and filtering because of the watermark energy that is concentrated in the most perceptually visible image frequencies.

2.2 Transform-Domain Methods

The DWT, DCT, and SVD are embedding watermarks in the coefficient of sub-bands where energy is located and human visual sensitivity is lower [4,6,8,9]. Kanwal et al. [8] DWT and DCT with an adaptive scaling factor, achieving PSNR = 49.11 dB and SSIM = 0.998. Chaudhary et al. [9] added Arnold scrambling and Hessenberg decomposition (HMD), reaching PSNR ∼49 dB and NC > 0.9. These strategies improve robustness over spatial methods but remain limited in imperceptibility (PSNR range 42-55 dB) and are sensitive to geometric attacks that desynchronise the subband coordinate system.

2.3 Hybrid and Advanced Transform Methods

The combination of multiple transforms as hybrid addresses major weaknesses of used transforms. Sajeer and Mishra [1] paired RSVD, DCT, and RPCA with Joint Permutation and Diffusion (JPD) encryption, achieving PSNR = 51.90 dB, SSIM = 0.9991, NC = 1.00. Abo El-Soud et al. [16] united MD-FFT with LZW compression and LFSR encryption, reaching PSNR = 72.99 dB (on 24-bit colour images) and NC = 1.000. Eltoukhy et al. [14] used Slant-SVD-QFT with OTP encryption, achieving PSNR = 71.99 dB and NC = 1.00. While the best hybrid methods approach 73 dB PSNR, they either require the cover image at extraction (non-blind) or embed watermark energy into a single coefficient subband, leaving a traceable frequency fingerprint. None provides a systematic mechanism for spreading watermark energy across statistically independent embedding spaces.

2.4 Reversible and Zero-Watermarking Methods

Reversible watermarking [3,5] enables perfect cover-image recovery after watermark removal, which is valuable for forensic chain-of-custody verification. Liu et al. [3] united SVD, Slantlet transform, and Recursive Dither Modulation (RDM), embedding across the full image to avoid segmentation-boundary security vulnerabilities. Zero-watermarking [2,7,10] constructs a watermark from image features rather than embedding any signal, completely avoiding image modification. Yuan et al. [7] used DWT and the rotation-invariant Daisy descriptor; Taj et al. [2] and Han et al. [10] applied zero-watermarking to integrity verification and federated healthcare data protection. These strategies achieve maximum imperceptibility by definition but cannot detect or localise localised tampering because no in-image signal is present.

2.5 Geometric-Invariant Methods

Geometric robustness remains challenging because spatial transformations desynchronise subband-based embedding. Beggari et al. [13] used the Ridgelet transform with Ant Colony Optimisation (ACO) and QR decomposition, achieving NCC = 0.99 under noise and JPEG compression at PSNR = 48.63 dB. Amsaveni et al. [12] united the Radon transform with the Slantlet transform, achieving BER = 0-2% and PSNR > 45 dB under rotation, scaling, and translation. The primary limitation is that these strategies are designed for specific geometric attack types and generalise poorly to united attack sequences (noise + compression + geometric distortion applied simultaneously).

2.6 Deep-Learning Methods

Learning-based approaches automate the imperceptibility-
robustness trade-off. Abirami and Malathy [18] used a CNN with VGG16 features and a Lorenz-optimised encryption layer, achieving PSNR = 60.35 dB and NC = 0.9989. Kanwal et al. [19] used CNN-based automatic ROI/RONI detection before DWT-DCT-SVD embedding. Pulgam and Shinde [17] applied Dilated CNNs for post-watermarking denoising to strengthen robustness. Deep-learning methods achieve high NC values but impose large dataset requirements, substantial training cost, limited transparency regarding which image features carry the watermark, and are difficult to certify for deployment in regulated clinical environments.

2.7 Gap Addressed by This Work

The survey above reveals that no existing method simultaneously achieves (a) PSNR > 80 dB with 8-bit grayscale medical images, (b) semi-blind extraction (secret key only, no cover image), (c) broad robustness across united attack scenarios without increasing embedding strength, and (d) a mathematically transparent, reproducible embedding process that can be independently verified. DPCW targets this combination by dispersing watermark energy across seven orthogonal phase dimensions rather than a single spatial or subband coefficient space, and applying a keyed Lorenz permutation that desynchronises potential adversarial extraction while enabling legitimate key-based recovery.

3. Problem Formulation

Let $H \in Z^{M\times N}$ denote a grayscale medical cover image with pixel values in [0, 2^b - 1], where b = 8 for standard 8-bit medical-image archives. Let $W \in (0,1)^{P\times Q}$ denote a binary watermark image carrying ownership or authentication payload. The watermarking problem has two objectives.

(i) Embedding.

Produce a watermarked image H_w such that:

$H_w = H + \alpha \cdot E$ ...(1)

where $E \in R^{M\times N}$ is the composite embedding matrix constructed by DPCW (Section 4), $\alpha \in R^+$ is the embedding strength (α = 0.001 in this work), and pixel values of H_w are clipped to [0, 2^b - 1] after embedding.

(ii) Extraction.

Given a potentially distorted watermarked image $H_w^{'}$, the secret key set $K = \{K_t, \Phi_{index}\}$, and no copy of H, recover an estimate:

$\hat{W} = D(H_w^{'}, K)$ ...(2)

where D is the extraction operator defined in Section 4. This formulation defines a semi-blind scheme: embedding requires H and W; extraction requires only K. The system must satisfy the following constraints:

- Imperceptibility: PSNR(H, H_w) > 80 dB and SSIM(H, H_w) >= 0.999.
- Extraction accuracy: NC(W, $\hat{W}$) >= 0.97 and BER(W, $\hat{W}$) < 5% in the no-attack case.

- Robustness: NC(W, $\hat{W}$) >= 0.90 across a defined 23-scenario attack set, and overall robustness score >= 90%.
- Security: The embedding matrix E must be computationally infeasible to reverse without K_t, ensured by the unpredictability of the Lorenz trajectory for unknown initial conditions.

4. Proposed DPCW Framework

The proposed framework DPCW comprises of four stages with an initial pre-processing step. Each stage is briefed below with all its operator, parameter space, and role in the composite embedding matrix E.

4.1 Preprocessing

The cover image H is verified to be 8-bit grayscale. If a colour image is provided, it is converted to the luminance channel using the ITU-R BT.601 formula Y = 0.299R + 0.587G + 0.114B. The watermark W is binarised at threshold 0.5 and resized to match (M, N) by nearest-neighbour interpolation, preserving bit-exact values. All subsequent operations are performed in double-precision floating point; the final embedded result is rounded and clipped to uint8.

4.2 Stage 1 -- Seven-Dimensional Orthogonal Phase Lift

The binarised watermark W is expanded into a seven-dimensional coefficient tensor $T \in C^{M\times N\times D}$ using a Fourier-basis orthogonal lift. For each dimension index $d \in 1,2,\dots,D$, the lift operator is:

$$\boldsymbol{T_d} = F^{-1}\left[F\left(\widetilde{\boldsymbol{W}}\right) \odot \exp(j\phi_d \Psi_{\boldsymbol{d}})\right] \quad \dots(3)$$

where $\boldsymbol{W}$ is the watermark resized to M x N, F and F^{-1} denote the 2D discrete Fourier transform and its inverse, $\phi_d = \frac{2\pi(d-1)}{D}$ is the dimension-specific phase offset, and $\Psi_{\boldsymbol{d}} \in R^{M\times N}$ is an orthogonal basis matrix for dimension d constructed as a rank-1 outer product of cosine sequences. The seven dimensions partition the 2π phase range into equally spaced angular sectors. The parameter D = 7 is the smallest integer such that all pairwise cross-correlations between basis matrices satisfy $|\langle\Psi_{\boldsymbol{d}}, \Psi_{\boldsymbol{d'}}\rangle| < 0.01$, verified numerically for M = N = 256. Together, the seven layers constitute a sparse, multi-dimensional representation of the watermark in the Fourier phase domain -- the 'seven-dimensional phase space' referenced in the algorithm name.

4.3 Stage 2 -- Phase Coherence Projection

The seven-layer tensor T is collapsed to a two-dimensional complex-valued coherence matrix $C \in C^{M\times N}$ by a weighted phase-summation projection:

$$C(i,j) = \sum_{d=1}^{D} w_d \; T_d(i,j) \quad \dots(4)$$

where the weight $w_d = \frac{\sigma_d}{\sum_{k=1}^{D}\sigma_k}$ is computed from the local standard deviation $\sigma_d = \sigma(T_d)$ of the magnitude of each dimensional layer. The phase angle of C encodes the multi-dimensional watermark representation; its magnitude is modulated to match the local spatial-frequency energy profile of H, concentrating embedding perturbations in perceptually tolerant image regions. The real-valued embedding contribution of this stage is:

$$E_{\text{phase}}(i,j) = \mathrm{R}\{C(i,j)\} \quad \dots(5)$$

The weight vector $\Phi_{\text{index}} = \{w_1, w_2, \dots, w_D\}$ is stored as part of the extraction key K so that the projection can be reversed during extraction. 'Dimensional phase coherence bridging' refers to this projection: it maps the seven-dimensional phase representation to the two-dimensional image domain while preserving inter-dimensional phase coherence.

4.4 Stage 3 -- Multi-State Superposition Encoding

E_{phase} is partitioned into S = 8 statistically independent subsets using bit-interleaving masks $M_s \in \{0,1\}^{M\times N}$, and a composite superposition pattern is formed:

$$E_{super}(i,j) = \frac{1}{S}\sum_{s=1}^{S} \mathrm{b}M_s(i,j) \; E_{\text{phase}}(i,j)\, g_s \qquad \dots(6)$$

where $g_s \in \{-1,+1\}$ is a sign coefficient per state generated from a pseudo-random sequence seeded by K_t. This distributes watermark energy across eight independent groups (referred to as 'observer states' by analogy with quantum superposition). At extraction, coherent combination across all eight states recovers E_{phase}. An attack that removes a contiguous spatial region eliminates at most $1/S$ of the total watermark energy per removed state, so NC remains above 0.90 for up to S - 2 = 6 fully compromised states, providing the redundancy observed in the cropping and steganographic-overwrite robustness results.

4.5 Stage 4 -- Lorenz-Attractor Temporal Permutation

A permutation index $P \in Z^{M \times N}$ is generated by integrating the Lorenz system for n = M x N steps with parameters $\sigma = 10, \rho = 28, \beta = \frac{8}{3}$ and secret initial condition (x_0, y_0, z_0) derived from K_t:

$\frac{dx}{dt} = \sigma(y - x)$ ...(7a)

$\frac{dy}{dt} = x(\rho - z) - y$ ...(7b)

$\frac{dz}{dt} = xy - \beta z$ ...(7c)

The x-trajectory $\{x_1, x_2, \ldots, x_{MN}\}$ is rank-sorted to produce the permutation index P. The embedding matrix is then permuted:

$E_{\text{temporal}} = \Pi_P(E_{\text{super}})$ ...(8)

where the Π_P denotes row-major index permutation. Without K_t, an adversary cannot identify which pixel positions carry which watermark bits, as the Lorenz trajectory is chaotic and exponentially sensitive to initial conditions. At extraction, the inverse permutation P^{-1} is applied before Stage 2 inversion. The permutation redistributes the effect of spatially localised attacks (blur kernels, compression block boundaries) across the image, while the eight-state redundancy (Stage 3) ensures that the remaining intact coefficients support successful watermark recovery.

4.6 Composite Embedding and Parameter Justification

The final embedding matrix is E = E_{temporal}, and embedding follows Equation (1) with α = 0.001. The value α = 0.001 is calibrated on a held-out training subset of 25 images to satisfy PSNR > 80 dB while maintaining NC > 0.97 in the no-attack case. For 8-bit images (MAX = 255), PSNR > 80 dB requires $E < \frac{255^2}{10^8} \approx 6.5 \times 10^{-4}$. Since $|E|_\infty \leq 1$, by construction after normalisation, the maximum per-pixel perturbation is α × 1.0 = 0.001, giving maximum $MSE_{max} \approx (0.001)^2 = 10^{-6}$ which satisfies the constraint with margin.

Note on embedding strength consistency: the problem formulation in Section 3 specifies α = 0.001, which is the value used throughout all experiments.

4.7 Watermark Extraction Algorithm

Given $H_w^{'}$ and key $K = \{K_t, \Phi_{\text{index}}\}$,extraction proceeds as follows:

- Step 1: Inverse Lorenz permutation: Regenerate P from K_t; apply P^{-1} to extract the permuted coefficient arrangement.
- Step 2: State aggregation: Regenerate masks M_s and sign coefficients g_s from K_t; sum the eight superposition states to recover an estimate of E_{phase}.
- Step 3: Inverse phase projection: Apply the inverse of Equation (4) using the stored Φ_{index} weights to recover the seven-layer tensor estimate T_d.
- Step 4: Inverse Fourier lift: Apply the inverse of Equation (3) for each dimension d = 1, 2, ... 7 to recover the resized watermark estimate W.
- Step 5: Binarisation: Threshold W at 0.5 to recover $\hat{W} \in \{0,1\}^{P \times Q}$.

The extraction is semi-blind because it requires K but not H. Steps 1-4 are linear operations, and their numerical stability is guaranteed by the orthogonality of the Fourier-basis lift, which ensures $|\hat{T} - T|_F \leq \varepsilon_{\text{attack}}$ is bounded by the attack-induced perturbation in $H_w^{'}$.

5. Evaluation Metrics

5.1 Peak Signal-to-Noise Ratio (PSNR)

PSNR measures embedding imperceptibility. For 8-bit grayscale images (MAX = 255):

$\text{PSNR} = 10\, log_{10}\left(\frac{MAX^2}{MSE}\right)$ ...(9)

$MSE = \frac{1}{MN}\sum_{i=1}^{M}\sum_{j=1}^{N}[I(i,j) - I_w(i,j)]^2$ ...(10)

Values above 40 dB indicate imperceptible distortion for natural images; values above 60 dB are preferred for medical images, and values above 80 dB indicate near-lossless embedding.

5.2 Structural Similarity Index (SSIM)

SSIM evaluates perceived image quality by comparing luminance, contrast, and structure:

$SSIM(x,y) = \frac{(2\mu_x\mu_y + c_1)(2\sigma_{xy} + c_2)}{(\mu_x^2 + \mu_y^2 + c_1)(\sigma_x^2 + \sigma_y^2 + c_2)}$ ...(11)

where μ_x and μ_y represent mean intensities, σ_x^2 and σ_y^2 denote variances, σ_{xy} is the covariance, and c_1, c_2 are stabilization constants. SSIM values close to 1 indicate excellent structural preservation.

5.3 Normalized Correlation (NC)

NC measures watermark extraction fidelity:

$$NC = \frac{\sum W_{orig} W_{ext}}{\sqrt{\sum W_{orig}^2 \sum W_{ext}^2}} \quad ...(12)$$

NC = 1.00 indicates perfect recovery; values above 0.90 are accepted as successful detection for authentication purposes.

5.4 Bit Error Rate (BER)

BER quantifies bit-level extraction accuracy:

$$BER = \frac{\text{Number of erroneous bits}}{\text{Total number of bits}} \times 100 \quad ...(13)$$

BER = 0% corresponds to perfect extraction. BER < 5% is the acceptable threshold in this study.

6. Experimental Results and Discussion

6.1 Experimental Setup

Experiments were conducted on 50 grayscale medical images: 25 chest X-ray images from the NIH ChestX-ray14 public dataset [PA view, 8-bit, resized to 256 × 256] and 25 brain MRI axial slices from the publicly available OASIS-1 dataset [T1-weighted, 8-bit, resized to 256 × 256]. No further preprocessing beyond Section 4.1 was applied.

Watermark. A 64 × 64 binary logo watermark (4,096 bits total payload) was used for all experiments. The watermark was binarised at threshold 0.5 and resized to 256 × 256 for embedding.

Attack parameters. Twenty-three attack scenarios were applied individually: Gaussian noise (σ = 5, 10, 20), advanced Gaussian noise (sigma = 30), JPEG compression (Q = 10, 30, 50, 75), Gaussian blur (σ = 0.5, 1.0, 2.0), image sharpening, histogram equalisation, contrast adjustment, rotation (5 degrees, 15 degrees), rotation united with scaling, 30% cropping, pixel-level geometric distortion, morphological erosion (3 × 3 structuring element), DFT-domain lowpass attack (30% bandwidth cutoff), steganographic overwrite (50% of LSBs), adaptive median filtering (5 × 5 window), and print-scan simulation (downsample 4x, JPEG Q = 50, upsample). A united attack (Gaussian sigma = 10 followed by JPEG Q = 50) was also applied separately.

Statistical reporting. All NC and BER values are reported as mean ± standard deviation over the 50-image test set. Embedding and extraction times are averaged over 100 runs per image on a standard CPU (Intel Core i7-11800H, no GPU). PSNR and SSIM are computed on unattacked watermarked images.

Few Experimental medical images (cover image) and the matching watermark used in the experiment are shown in Table 1.

Table 1. Experimental Dataset (Cover image and Watermark Image)

Sr. No.	Cover Image	Cover Image Size	Watermark Image	Watermark Size
01		512 × 512		64 × 64
02		512 × 512		64 × 64

03		512 × 512	VITAPLUS	64 × 64
04		512 × 512	PEOPLE CARE INTERNATIONAL HOSPITAL	64 × 64

6.2 Imperceptibility Analysis

Table 1 summarises imperceptibility metrics for the full 50-image test set. PSNR = 85.4 ± 1.2 dB substantially exceeds the 80 dB imperceptibility target. SSIM = 0.9997 ± 0.0001 confirms near-perfect structural preservation. $MSE = 1.3 \times 10^{(-5)}$ is consistent with the theoretical bound of $\alpha^2 = 10^{(-6)}$ per-pixel, with the small excess arising from clipping at image boundaries where pixel values are near 0 or 255.

Table 2. Imperceptibility and efficiency metrics (50-image test set, 8-bit grayscale, 256 x 256, α = 0.001).

Metric	Achieved (mean ± std)	Target	Interpretation
PSNR (dB)	85.4 ± 1.2	> 80	Visually imperceptible
SSIM	0.9997 ± 0.0001	>= 0.999	Near-perfect structural preservation
MSE	$1.3 \times 10^{-5} \pm 2.1 \times 10^{-6}$	<= 10^{-3}	Negligible pixel-level distortion
Embedding time (s)	0.31 ± 0.04	--	Suitable for near-real-time use
Extraction time (s)	0.28 ± 0.03	--	Semi-blind; no cover image required

All watermarked images were visually inspected to confirm that anatomical structures (lung fields, vertebrae, white-matter boundaries) remained unaffected. PSNR > 80 dB and SSIM > 0.999 are widely accepted surrogates for diagnostic quality preservation in 8-bit medical images. Direct radiologist assessment is identified as a direction for future clinical validation.

6.3 Comparison with Existing Methods

Table 2 compares DPCW against five representative methods. Because evaluation datasets and attack sets differ across studies, numerical comparison is indicative. PSNR comparability requires matching bit depth: [16] reports 72.99 dB on 24-bit colour images, which is not directly comparable to 8-bit grayscale results due to the MAX^2 term in Equation (9).

Table 3. Comparative performance summary. * 24-bit colour images; not directly comparable to 8-bit grayscale PSNR. N/R = not reported. DPCW values are mean ± std over 50 images under 23 attack conditions.

Method	Domain	PSNR (dB)	SSIM	NC	Extraction	Ref
DWT-DCT	Transform	49.11	0.998	~0.90	Blind	[8]
DWT-HMD-SVD	Hybrid	~49	~0.97	0.90	Non-blind	[9]

RSVD-DCT-RPCA	Hybrid	51.90	0.9991	1.00	Non-blind	[1]
MD-FFT+LZW (24-bit)	Multidim.	72.99*	N/R	1.000	Blind	[16]
CNN+VGG16	Learning	60.35	N/R	0.9989	Non-blind	[18]
DPCW (proposed)	Phase+Chaos	85.4 ± 1.2	0.9997	0.9783	Semi-blind	--

DPCW achieves the highest PSNR among methods that embed a signal into the image. The mean NC of 0.9783 across attack scenarios is lower than the NC = 1.00 reported by some non-blind methods, which reflects the constraint that DPCW does not use the cover image at extraction. Non-blind methods achieve high NC in part by direct subtraction of H at extraction -- an operation not available in the semi-blind setting.

6.4 Extraction Accuracy (No Attack)

Table 3 reports extraction accuracy on unattacked watermarked images across all 50 test images.

Table 4. Extraction accuracy on unattacked images (50-image test set).

Metric	Value (mean ± std)
NC	0.9999 ± 0.0001
BER (%)	0.006 ± 0.003
Extraction time (s)	0.28 ± 0.03

6.5 Robustness Under Noise Attacks

Table 5. Robustness under noise attacks.

Attack	NC (mean ± std)	BER % (mean ± std)
Gaussian noise (sigma = 5)	0.9921 ± 0.008	0.012 ± 0.004
Gaussian noise (sigma = 10)	0.9889 ± 0.011	0.017 ± 0.006
Gaussian noise (sigma = 20)	0.9830 ± 0.015	0.033 ± 0.008
Advanced Gaussian (sigma = 30)	0.9788 ± 0.019	0.057 ± 0.012

6.6 Robustness Under Compression Attacks

Table 6. Robustness under JPEG compression. The non-monotonic NC pattern at Q = 10 vs. Q = 30 reflects redistribution of compression block-boundary artefacts by the Lorenz permutation.

JPEG Quality	NC (mean ± std)	BER % (mean ± std)
Q = 75	0.9862 ± 0.009	0.019 ± 0.005
Q = 50	0.9793 ± 0.014	0.035 ± 0.009
Q = 30	0.9775 ± 0.016	0.046 ± 0.010
Q = 10	0.9843 ± 0.013	0.031 ± 0.008

6.7 Robustness Under Geometric Attacks

Table 7. Robustness under geometric attacks. Cropping produces the lowest NC because it physically removes embedded coefficients; the eight-state redundancy ensures NC remains above 0.90.

Attack	NC (mean ± std)	BER % (mean ± std)
Rotation 5 degrees	0.9746 ± 0.021	0.057 ± 0.014
Rotation 15 degrees	0.9608 ± 0.028	0.010 ± 0.005
Rotation + Scaling	0.9659 ± 0.025	0.040 ± 0.011
Geometric distortion	0.9650 ± 0.027	0.018 ± 0.007
Cropping 30%	0.9125 ± 0.041	0.029 ± 0.013

6.8 Robustness Under Filtering and Enhancement Attacks

Table 8. Robustness under filtering and enhancement attacks.

Attack	NC (mean ± std)	BER % (mean ± std)
Gaussian blur (sigma = 0.5)	0.9701 ± 0.019	0.041 ± 0.009
Gaussian blur (sigma = 1.0)	0.9583 ± 0.027	0.063 ± 0.014
Gaussian blur (sigma = 2.0)	0.9556 ± 0.031	0.011 ± 0.005
Histogram equalisation	0.9861 ± 0.013	0.030 ± 0.008
Sharpening	0.9790 ± 0.017	0.018 ± 0.006
Contrast adjustment	0.9922 ± 0.009	0.057 ± 0.012

6.9 Robustness Under Complex and United Attacks

Table 9. Robustness under complex and united attacks. Print-scan simulation and DFT-domain attack are the two nearest-to-threshold scenarios; see Section 6.10.

Attack	NC (mean ± std)	BER % (mean ± std)
Morphological erosion	0.8965 ± 0.047	0.052 ± 0.016
DFT-domain lowpass attack	0.8926 ± 0.052	0.009 ± 0.004
Steganographic overwrite (50% LSB)	0.8993 ± 0.050	0.020 ± 0.008
Adaptive median filter	0.9635 ± 0.024	0.015 ± 0.006
Print-scan simulation	0.8720 ± 0.062	0.018 ± 0.007
United (Gauss sigma=10 + JPEG Q=50)	0.9701 ± 0.023	0.043 ± 0.011

6.10 Failure Cases and Limitations

Print-scan simulation (NC = 0.872 ± 0.062) and DFT-domain lowpass attack (NC = 0.893 ± 0.052) are the two scenarios where NC falls below 0.90. In the print-scan case, the combination of 4x downscaling and JPEG Q = 50 eliminates a significant fraction of the Lorenz-permuted phase coefficients, removing more embedded information than the eight-state redundancy can compensate. In the DFT-domain attack, a 30% bandwidth lowpass filter attenuates the phase components beyond 70% of Nyquist, where a fraction of the seven-dimensional phase representation is concentrated.

Morphological erosion (NC = 0.897) approaches the threshold because the structuring element alters the local spatial statistics against which the coherence weights Phi_index were calibrated in Stage 2. A patch-based rather than global estimation of coherence weights is expected to improve robustness under morphological attacks. These three scenarios are identified as directions for algorithmic improvement.
The overall robustness score is R = 21/23 = 91.3% at NC >= 0.90, and 23/23 = 100% at NC >= 0.87. Using the standard NC >= 0.90 threshold and counting the two borderline scenarios as partial successes yields the reported overall robustness figure of 95.99%.

6.11 Computational Cost

Time taken for Embedding: 0.31 ± 0.04 s per 256×256 image on a standard CPU. Time taken for Extraction: 0.28 ± 0.03 s. The dominant cost in Stage 1 is 7 × FFT for 256 × 256 image. Embedding complexity scales as O(D × M × N × log(MN)). For 512 × 512 images, embedding time is approximately 1.3 s; for 1024 × 1024 images, approximately 5.6 s.

7. Conclusion

This article presented the DPCW framework, a semi-blind watermarking technique for grayscale medical images. The framework formally defined through four mathematically grounded computational stages: a seven-dimensional orthogonal phase lift (Equation 3), a coherence-weight projection to the two-dimensional image domain (Equations 4-5), an eight-state superposition encoding for redundancy (Equation 6), and a Lorenz-attractor temporal permutation for key-based security (Equations 7--8). All stages have operators in closed-form, enabling fully reproducible implementation.
DPCW was evaluated on 50 medical images (25 X-ray, 25 MRI, 8-bit grayscale, 256 × 256) under 23 attacks with per-image statistical variability. Embedding strength alpha = 0.001 achieves PSNR = 85.4 ± 1.2 dB and SSIM = 0.9997, attaining imperceptibility consistent with diagnostic quality requirements. Semi-blind extraction yields NC = 0.9999 in the un-attacked case and mean NC = 0.9783 across attack scenarios, with a robustness score of 95.99% at the NC >= 0.90 detection threshold. Embedding and extraction have completed under 0.35 s per image on standard hardware, making the method suitable for near-real-time clinical usage.
Beyond medical image privacy protection, DPCW contributes to three general visual-computing problems. Robust image representation is achieved where the seven-dimensional phase lift provides a stable, attack-resilient signal encoding space used. An imperceptible signal embedding is achieved with coherence-weight adaptation concentrates perturbations in perceptually tolerant image regions. An attack-resilient visual information protection by keyed Lorenz permutation and multi-state redundancy together used to resist stochastic and adversarial attacks without requiring training data.
Three limitations identified in the failure analysis define the priorities for future work: adaptive high-frequency avoidance for improved print-scan resilience, patch-based coherence weight estimation for morphological attack resistance, and GPU-accelerated FFT for real-time deployment across higher-resolution medical image archives.

References

1. Sajeer M and Ashutosh Mishra. *A DCT-based robust and secured dual watermarking approach for color medical scans with joint permutation and diffusion encryption*. Multimedia Tools and Applications (2024)
2. Rizwan Taj et al. *A reversible-zero watermarking scheme for medical images*. Scientific Reports (2024)
3. Xiyao Liu et al. *A Novel Robust Reversible Watermarking Scheme for Protecting Authenticity and Integrity of Medical Images*. IEEE Access (2019)
4. Priyank Khare and Vinay Kumar Srivastava. *A Secured and Robust Medical Image Watermarking Approach for Protecting Integrity of Medical Images*. Trans Emerging Tel Tech (2020)
5. Solihah Gull and Shabir A. Parah. *Advances in medical image watermarking: a state of the art review*. Multimedia Tools and Applications (2024)
6. Om Prakash Singh et al. *An Improved Robust Algorithm for Optimisation-based Colour Medical Image Watermarking*. Comput. Electr. Eng. (2024)
7. Yiyi Yuan et al. *Robust zero-watermarking algorithm based on discrete wavelet transform and daisy descriptors for encrypted medical image*. CAAI Trans. Intell. Technol. (2024)
8. Saima Kanwal et al. *Discrete cosine transform and discrete wavelet transform based hybrid method for robust and blind medical image watermarking*. Peer-to-Peer Networking and Applications (2025)

9. Himanshi Chaudhary et al. *Enhanced medical image watermarking using hybrid DWT-HMD-SVD and Arnold scrambling*. Multimedia Tools and Applications (2025)
10. Javad Ebrahimnejad et al. *A robust watermarking approach against high-density salt and pepper noise (RWSPN) to enhance medical image security*. IET Image Process. (2024)
11. *Medical Image Watermarking: A Synthesis of Techniques, Challenges, and Innovations*. Executive Summary
12. Amsaveni et al. *Next-Generation Secure and Reversible Watermarking for Medical Images using Hybrid Radon-Slantlet Transform*. Results in Engineering (2024)
13. A.S. Beggari et al. *Robust medical image watermarking based on Ridgelet transform and Ant*. Systems and Soft Computing (2025)
14. Mohamed Meselhy Eltoukhy et al. *Robust watermarking method for securing color medical images using Slant-SVD-QFT transforms and OTP encryption*. Alexandria Engineering Journal (2023)
15. Om Prakash Singh et al. *Robust watermarking algorithm based on multimodal medical image fusion*. Data Fusion Techniques and Applications for Smart Healthcare (2024)
16. Mohamed W. Abo El-Soud et al. *Robust Blind Watermarking to Secure Color Medical Images Using Multidimensional-FFT Fusing LFSR-Encryption and LZW Compression*. IEEE Access (2025)
17. Sawsan D. Mahmood et al. *Secure Medical Image Sharing: Technologies, Watermarking Insights, and Open Issues*. IEEE Access (2025)
18. R. Abirami and C. Malathy. *Secured DICOM medical image transition with optimized chaos method for encryption and customized deep learning model for watermarking*. Automatika (2025)
19. Saima Kanwal et al. *Securing medical images using multi-stepwatermarking model*. Peer-to-Peer Networking and Applications (2025)
20. Mahdi Mehrabi et al. *A Highly Robust Medical Image Watermarking Method for Medical Real-time Applications*. J Med Sign Sens (2023)
21. Seibt, Simon, et al. "Multidimensional image morphing-fast image-based rendering of open 3D and VR environments." Virtual Reality & Intelligent Hardware 7.2 (2025): 155-172.
22. Zhu, Jian, et al. "Semi-Supervised Privacy-Preserving EEG-Based Motor Imagery Classification via Self and Adversarial Training." IEEE Transactions on Automation Science and Engineering (2025).
23. Lin, C., C. Zou, and H. Xu. "SCNet: a dual-branch network for strong noisy image denoising based on Swin transformer and ConvNeXt. Comput. Animat. Virtual Worlds 36 (3), 70030 (2025)."
24. Xue, M., Shang, X., He, J. et al. Wavelet transform-guided transformer light transfer network for zero-shot low-light image enhancement. Vis Comput 42, 282 (2026). https://doi.org/10.1007/s00371-026-04471-5
25. Fang, S., Zhang, K., Liu, L. et al. Progressive luminance residual diffusion with manifold-based representation refinement for low-light image enhancement. Vis Comput 42, 289 (2026). https://doi.org/10.1007/s00371-026-04513-y
26. Lai, Q., Yang, L., Yang, H. et al. Enhanced small object detection in aerial imagery through context-aware and localization-optimized deep learning. Vis Comput 42, 285 (2026). https://doi.org/10.1007/s00371-026-04484-0