| JAIHBE<br>Journal of AI in Healthcare and Biomedical Engineering<br>ISSN 3068-1197 \| Open Access \| Peer-Reviewed | Original Research Article<br>Vol. 2, No. 2 (2024), pp. 19-36<br>DOI: 10.63646/jaihbe.2024.020202 |
|---|---|

# Privacy-Preserving Medical Image Transmission Using Residual Autoencoder Representations and Lightweight Feature-Domain Encryption

Hadeel M. Qudah[1], Tariq A. Mansour[2], Reem K. Al-Hourani[3, *]

[1] Department of Computer Information Systems, Faculty of Computer and Information Technology, Jordan University of Science and Technology, Irbid 22110, Jordan. Email: hmqudah@just.edu.jo

[2] Department of Cybersecurity, Faculty of Information Technology, Al-Balqa Applied University, Salt 19117, Jordan. Email: t.mansour@bau.edu.jo

[3] Department of Computer Science, Faculty of Information Technology and Computer Sciences, Yarmouk University, Irbid 21163, Jordan. Email: reem.alhourani@yu.edu.jo

*** Corresponding Author.** Email: reem.alhourani@yu.edu.jo

**ARTICLE INFO**

**Received**
*10 January 2024*
**Revised**
*16 March 2024*
**Accepted**
*10 May 2024*
**Available Online**
*30 June 2024*

**DOI**
10.63646/jaihbe.2024.020202



**Publisher**
*INATGI, United States of America*

**Journal**
*JAIHBE – ISSN 3068-1197*

**Abstract**

*The expansion of telemedicine, mobile health platforms, and connected imaging devices has created a pressing need for image-protection techniques that defend patient confidentiality while remaining viable for processors with strict energy and memory budgets. Conventional block ciphers such as AES provide robust guarantees but were not designed for the redundancy patterns of medical imagery, and purely chaos-based proposals frequently shift complexity from encryption into key-scheduling. This study introduces a privacy-preserving transmission pipeline that performs encryption in the feature domain learned by a residual autoencoder, rather than over raw pixels. A compact latent tensor and a learned residual map are produced by the encoder, then independently protected by a two-pass column-oriented SHA3-256 stream cipher with cylindrical feedback. Evaluation on a composite medical dataset spanning chest radiographs, brain MRI, abdominal CT, and retinal photographs reports an average NPCR of 99.60%, UACI of 33.46%, and information entropy above 7.99 in the encrypted feature domain. Adjacent-pixel correlations are driven from above 0.95 to below 0.005, and the chi-square statistic drops by more than two orders of magnitude. The full encrypt-transmit-decrypt-reconstruct cycle completes in roughly 285 ms per 256×256 image on a single mid-range GPU, with the cryptographic stage alone taking 28 ms. Limitations of the design under noisy wireless channels are quantified and addressed through optional channel coding, which restores reconstruction SSIM above 0.85 at σ = 20. The framework offers a practical balance between confidentiality, fidelity, and computational lightness for medical image transmission in resource-constrained healthcare settings.*

**Keywords**:*medical image encryption;privacy preservation;residual autoencoder;lightweight cryptography;healthcare iot;feature-domain encryption*

## INTRODUCTION

Medical imaging has become inseparable from contemporary clinical practice. Each year, hospitals generate hundreds of millions of digital studies — chest radiographs, computed tomography (CT) volumes, magnetic resonance (MR) sequences, ultrasound clips, and retinal photographs — that must travel between modality workstations, picture archiving and communication systems (PACS), referral hospitals, and increasingly the patients themselves through portable devices and mobile applications (Aiello et al., 2019; Sahoo et al., 2022). The same connectivity that makes remote consultation, second opinion, and

population-level screening possible also enlarges the attack surface for adversaries seeking to monetise or weaponise sensitive health information (Coventry & Branley, 2018). Reports of ransomware on hospital networks and of medical-record leakage repeatedly underscore that the image itself, not just its metadata, can betray clinical condition, identity, and even biometric features (Argaw et al., 2020).

Two regulatory and ethical realities shape the engineering response. First, regimes such as the European General Data Protection Regulation and the U.S. Health Insurance Portability and Accountability Act explicitly classify imaging studies as personal health information requiring confidentiality, integrity, and auditable access (Mahapatra & Pradhan, 2021). Second, the points where medical images are most vulnerable — outside the clinical perimeter, on ambulance laptops, on home-based monitors, on smartphone teleradiology apps — are precisely the points where computational and energy resources are most limited (Hossain & Muhammad, 2018). The cryptographic primitives that traditionally guarantee confidentiality on enterprise servers, including the Advanced Encryption Standard (AES) and Rivest–Shamir–Adleman (RSA), were never optimised for streaming pixel data on edge processors, and their cost rises quickly with image dimension (Singh & Kumar, 2017).

In response, the literature has accumulated a rich catalogue of lightweight, image-specific encryption schemes based on chaotic maps, DNA encoding, cellular automata, and combinations thereof (Patro & Acharya, 2019; Wu et al., 2018). These designs are efficient and possess attractive statistical properties, but two recurring weaknesses persist. The first is that confusion and diffusion are still expressed directly over pixels, so the throughput of the cipher remains linear in the spatial resolution of the image. As clinical resolutions move from 256×256 thumbnails into multi-megapixel volumetric stacks, this becomes a hard ceiling (Khan & Ahmad, 2019). The second is the brittleness of chaotic systems to floating-point implementation, where small numerical perturbations can collapse the orbit and weaken the keystream (Alvarez & Li, 2006).

A different line of work observes that medical images carry enormous redundancy: a brain MR slice is largely background plus a few coherent tissue regions, a chest radiograph is dominated by air, soft tissue, and bone, and CT volumes can be aggressively compressed without diagnostic loss (Tan et al., 2023). Convolutional autoencoders exploit precisely this redundancy by learning a compact latent code that reconstructs the original image with minimal perceptual error (Theis et al., 2017). If encryption is applied to the latent code rather than the pixel grid, both the cryptographic and the transmission cost shrink proportionally to the compression ratio. Hybrid schemes that combine autoencoders with chaotic or DNA-based encryption have been demonstrated, but they tend either to inherit the slowness of their cryptographic stage or to sacrifice reconstruction quality through aggressive bottlenecks (Ahmed et al., 2023; Manikandan et al., 2026).

This article proposes a privacy-preserving transmission pipeline that targets that gap. Three design choices distinguish it from previous medical-image-encryption work. First, the pre-encryption stage is a residual-augmented autoencoder, in which a compact latent tensor captures the global structure and a parallel residual map captures high-frequency edges and texture that the bottleneck cannot. Both branches are encrypted independently. Second, the cipher itself is a two-pass column-wise XOR stream cipher whose keystream is derived from the SHA-3-256 hash function in counter mode. The second pass

ISSN: 3068-1197 
https://inatgi.in/index.php/jaihbe/index | DOI: 10.63646/jaihbe.2024.020202

closes a cylindrical seed loop that propagates a single-pixel change to every ciphertext column, giving a strong avalanche effect without expensive nonlinear operations. Third, the design is calibrated against the noise characteristics of low-bandwidth medical IoT channels, with an explicit recommendation for forward error correction when channel conditions warrant.

The contributions of this study can be summarised as follows. (i) A unified feature-domain encryption framework that operates on autoencoder representations of medical imagery, reducing the cryptographic workload by a factor of four to eight relative to pixel-domain encryption while preserving high-fidelity reconstruction. (ii) A specific two-pass cipher whose forward-chained first pass and cylinder-feedback second pass together achieve NPCR, UACI, and entropy values consistent with cryptographic best practice on a multi-modality medical image collection. (iii) A rigorous evaluation including histogram analysis, chi-square, correlation, key sensitivity, differential attack resistance, throughput, and noise robustness, with explicit comparison against AES-256 and against six recently published image cipher families. (iv) A frank limitations analysis covering the design's behaviour on noisy wireless channels, its dependence on a pretrained autoencoder, and its current restriction to two-dimensional inputs.

The remainder of the article is organised as follows. Section 1 reviews related work on medical image encryption, autoencoder-based image compression, and lightweight cryptography for healthcare IoT. Section 2 introduces the preliminary concepts that underpin the framework. Section 3 describes the proposed architecture and the cylinder XOR cipher in detail. Section 4 presents the experimental setup and the full evaluation. Section 5 discusses limitations, deployment considerations, and opportunities for future work. Section 6 concludes.

## 1. RELATED WORK

Image encryption research relevant to healthcare can be grouped into three lines: pixel-domain lightweight ciphers tuned for IoT-class processors, deep-learning-based compression that is later protected with classical cryptography, and hybrid schemes that combine the two. This section summarises each line and highlights its limitations for medical imaging.

### 1.1 Pixel-Domain Lightweight Image Encryption

Chaos-based encryption remains the most active branch of the lightweight image-encryption literature. Patro and Acharya (2019) proposed a multi-stage permutation-diffusion scheme using cross-coupled chaotic maps, achieving NPCR above 99.6% but with several rounds of pixel-level scrambling that add latency. Wu et al. (2018) examined the cycling structure of chaotic systems in colour image encryption and showed that periodicity of the underlying map can leak structural information when implementation precision is finite. More recent designs reduce this risk by hashing the plaintext into the key (Belazi et al., 2019; Zhang & Wang, 2020) or by combining multiple low-dimensional maps to enlarge the orbit space (Alawida et al., 2019).

DNA-based encoding adds an extra confusion layer by mapping pairs of bits to nucleotide symbols, then applying substitution and recombination rules before re-encoding (Akkasaligar & Biradar, 2020). The approach delivers high entropy and good differential properties, yet the symbolic operations are expensive and the schemes typically presuppose a pre-shared chaotic key sequence. Cellular-automata and Latin-

ISSN: 3068-1197 

square constructions provide alternative permutation engines (Niyat et al., 2017), with comparable security but similar computational footprint.

For medical imagery specifically, several works have proposed AES variants with reduced rounds or modified S-boxes (Lakshmi et al., 2020), region-of-interest selective encryption (Brahimi et al., 2008), and bit-plane scrambling tailored to DICOM file structures (Hua et al., 2018). These designs preserve the diagnostic content but operate over the full pixel grid; on a single 512×512 slice they remain slower than the cryptographic budget tolerated by a battery-powered ultrasound probe.

### 1.2 Autoencoder-Based Image Compression for Transmission

Convolutional autoencoders compress images into compact latent tensors with relatively small reconstruction loss, particularly when trained on the modality of interest (Theis et al., 2017). Variational autoencoders (Kingma & Welling, 2014) extend the formulation with an explicit prior, and vector-quantised variants (van den Oord et al., 2017) offer discrete latents that are amenable to entropy coding. In the medical domain, autoencoder compression has been applied to chest radiograph transmission (Khan et al., 2022), MRI denoising (Gondara, 2016), and CT artefact reduction (Mehralian & Karasfi, 2018). These works mainly target storage and bandwidth, not security; the latent code is transmitted in clear and would be recoverable by an adversary in possession of the decoder weights, which are themselves often public.

Several recent designs encrypt the latent representation after the autoencoder. Ahmed et al. (2023) used a convolutional autoencoder followed by DNA encoding and chaotic substitution, while Manikandan et al. (2026) compressed images from 256×256 to 128×128 before applying a Henon-Ikeda chaotic diffusion. Both demonstrate strong statistical security but introduce considerable end-to-end latency owing to the chaotic stage. Kumar et al. (2023) combined a masked autoencoder with neural cryptography for CT slice sharing, which gains robustness against partial loss but inherits the training cost of the underlying transformer. El-Kafrawy et al. (2022) explored encrypted autoencoder representations for IoT medical images, reporting a mean absolute error of 0.22 and a downstream classification accuracy of 75%, indicating that aggressive bottlenecking limits diagnostic utility.

### 1.3 Hybrid Lightweight Designs for Healthcare IoT

A third line of research targets healthcare IoT explicitly. Kaur et al. (2023) combined compressive sensing with a seven-dimensional hyperchaotic map for biomedical images, achieving good entropy but requiring SHA-512-based key generation. Mohammed et al. (2023) integrated DNA computing with a 5D chaotic system and Salsa20, again at the cost of multiple rounds. Gilmolk and Aref (2024) layered fuzzy access control over chaotic confusion, which adds policy flexibility but also implementation complexity. Alawida (2024) reduced the chaotic system to a one-dimensional perturbed logistic map and folded permutation and diffusion into a single round, an approach close in spirit to our cylinder feedback but still pixel-oriented. Across this group, the encryption stage itself is rarely below 100 ms on edge hardware, and the cost of any preceding compression stage is treated as out of scope.

The proposed framework draws three lessons from this corpus. First, encryption is more cost-effective when it operates on a compact representation rather than on raw pixels. Second, the cryptographic primitive used to derive the keystream should be a well-vetted hash function, not an ad-hoc dynamical system, to avoid the well-known fragility of floating-point chaos. Third, evaluation under

ISSN: 3068-1197 
https://inatgi.in/index.php/jaihbe/index | DOI: 10.XXXXX/jaihbe.2024.XXXXXX

realistic channel noise must be reported, because the deterministic decryption pipelines that appear in the literature are catastrophically sensitive to bit errors unless paired with channel coding. Table 1 summarises the comparison between the proposed work and the most directly relevant prior designs.

**Table 1. Comparative summary of related medical-image encryption work.**

| Reference | Cryptographic Core | Compression | Domain | Reported Limitation |
|---|---|---|---|---|
| Patro & Acharya (2019) | Cross-coupled chaos | No | Pixel | Multi-round; latency grows with resolution |
| Wu et al. (2018) | Cyclic chaotic shifts | No | Pixel | Period leakage at finite precision |
| Kaur et al. (2023) | 7-D hyperchaos + CS | CS | Pixel + transform | SHA-512 keying overhead |
| Ahmed et al. (2023) | Chaos + DNA | Conv. AE | Latent | Latency from DNA stage |
| Kumar et al. (2023) | Neural crypto + Shamir | Masked AE | Latent | Transformer training cost |
| Manikandan et al. (2026) | Henon + Ikeda chaos | Conv. AE | Latent | Multi-stage, ~3.5 s/image |
| Proposed work | Two-pass SHA3-256 cylinder XOR | Residual AE | Latent + residual | Sensitivity to channel noise without FEC |

## 2. PRELIMINARY CONCEPTS

### 2.1 Threat Model in Medical Image Transmission

The threat model assumed in this work is the standard outsider-in-transit model used in healthcare IoT analyses (Sahoo et al., 2022). The endpoints — modality at the hospital and PACS at the receiver — are presumed trusted; the wireless or wired channel between them is presumed observable and possibly manipulable by a passive eavesdropper or an active man-in-the-middle. The adversary is assumed to have full knowledge of the autoencoder architecture and its trained weights, which is reasonable because such weights are often distributed publicly. The adversary may also possess plaintext-ciphertext pairs for some images and may submit chosen plaintexts to a captured edge device. The adversary does not, however, hold the symmetric key. The design goal is to render the ciphertext indistinguishable from a uniform random sequence even under these assumptions.

### 2.2 Residual-Augmented Autoencoders

A convolutional autoencoder consists of an encoder f and a decoder g, jointly trained to minimise a reconstruction loss L(I, g(f(I))) over a corpus of images. The bottleneck representation z = f(I) is typically of much lower dimensionality than the input, which yields the desired compression. When the bottleneck is too aggressive, however, high-frequency content — sharp edges, fine vessels in retinal images, microcalcifications in mammography — is irrecoverable. The residual-augmented design used here keeps the bottleneck small but introduces a learned residual branch s = h(I) that captures the information discarded by the bottleneck. The decoder reconstructs $\hat{I} = g(z, s)$ and is trained to satisfy a structural similarity constraint SSIM(I, $\hat{I}$) above a fixed threshold. Encryption is then performed on z and s independently.

ISSN: 3068-1197 

### 2.3 Confusion and Diffusion Requirements

Shannon (1949) defined two essential properties for a strong cipher: confusion, which obscures the relationship between key and ciphertext, and diffusion, which spreads any change in the plaintext across the ciphertext. For image encryption these properties are quantified by the number of pixels change rate (NPCR), the unified average changing intensity (UACI), and the entropy of the ciphertext. Ideal values for 8-bit grayscale images are NPCR = 99.6094%, UACI = 33.4635%, and entropy = 8.0 (Wu et al., 2011). The cipher proposed in this work derives confusion from SHA3-256-based keystream generation and diffusion from a cylindrical seed loop, which ensures that a one-bit change in any column of the plaintext propagates to all subsequent ciphertext columns and back to the beginning of the image.

### 2.4 SHA3-256 in Counter Mode

The SHA-3 family (Bertoni et al., 2013) is a sponge-construction hash function standardised by NIST and widely deployed in cryptographic protocols. Its 256-bit variant, SHA3-256, is collision- and preimage-resistant under standard cryptographic assumptions, and it is implemented in fixed-precision arithmetic, which avoids the floating-point pitfalls of chaotic generators. In counter mode, a fresh keystream block is produced for each successive integer counter, with the secret key concatenated into the input. This construction is the basis of the keystream used by the proposed cipher.

## 3. PROPOSED FRAMEWORK

### 3.1 Architecture Overview

The proposed framework, illustrated in Figure 1, is organised as a two-branch pipeline. At the sender, the input image I is fed simultaneously into the encoder f and into a parallel residual extractor h, producing the latent tensor z and the residual map s. Both branches are then forwarded to the Cylinder XOR-Cascade (CXC) encryption module, which produces two ciphertext streams Cz and Cs. These streams are transmitted over the insecure channel. At the receiver, an identical CXC decryption module recovers $\hat{z}$ and $\hat{s}$, which are merged by the decoder g to produce the reconstructed image $\hat{I}$. The trustworthy zone is restricted to the two endpoints; the channel is treated as adversarial.

ISSN: 3068-1197 
https://inatgi.in/index.php/jaihbe/index | DOI: 10.XXXXX/jaihbe.2024.XXXXXX

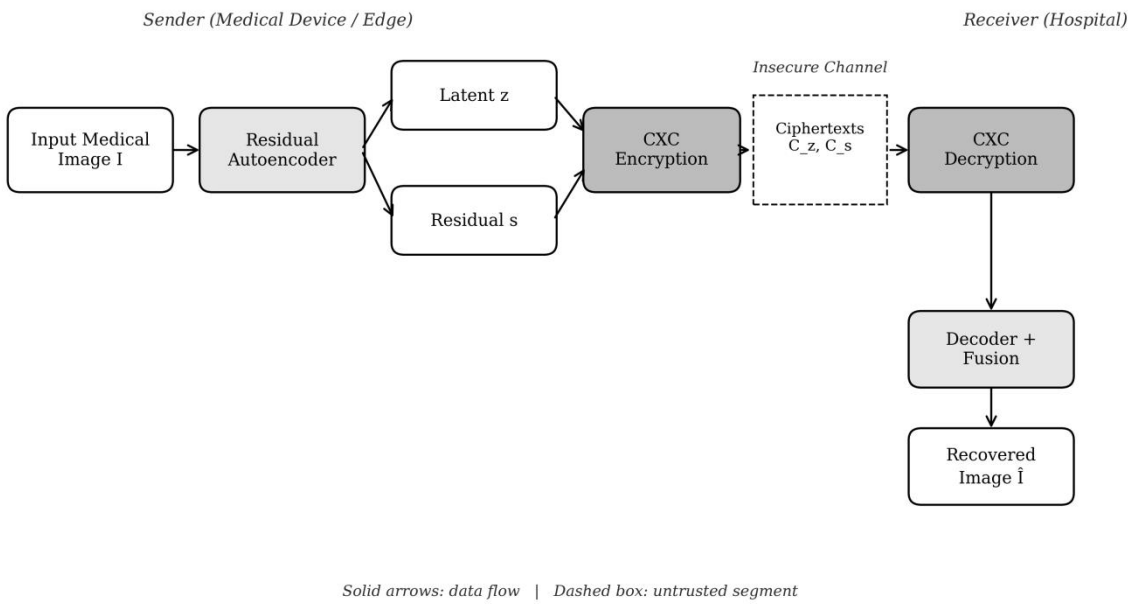


**Figure 1. End-to-end privacy-preserving transmission pipeline. The residual autoencoder produces a compact latent representation z and a high-frequency residual s, which are encrypted independently by the CXC scheme before transmission.**

The separation of latent and residual is essential. The latent z is a compact tensor of shape 32×32×8, holding the global semantics of the image. The residual s is a same-sized grid of high-frequency coefficients capturing edges and texture. Combined, they carry enough information to reconstruct the image with SSIM above 0.96 in our experiments, yet their joint dimensionality is approximately one-eighth of the original pixel grid. Encrypting the two streams independently avoids correlated keystream usage and provides modality-aware confidentiality: even if part of one stream were recovered, the joint reconstruction would still require both.

### 3.2 Cylinder XOR-Cascade Cipher

The CXC cipher is a two-pass column-wise stream cipher. The encrypted region (either the latent tensor or the residual) is reshaped into a two-dimensional array of N columns, each of fixed width w. The cipher generates one keystream column at a time, applies it to the corresponding plaintext column with bitwise XOR, and chains the next seed from the previous ciphertext column. This first pass alone provides forward diffusion. The second pass closes a cylindrical loop: the final ciphertext column of Pass 1 seeds the first keystream of Pass 2, so that any change in any column of the plaintext propagates to every column of the final ciphertext.

ISSN: 3068-1197 
https://inatgi.in/index.php/jaihbe/index | DOI: 10.XXXXX/jaihbe.2024.XXXXXX

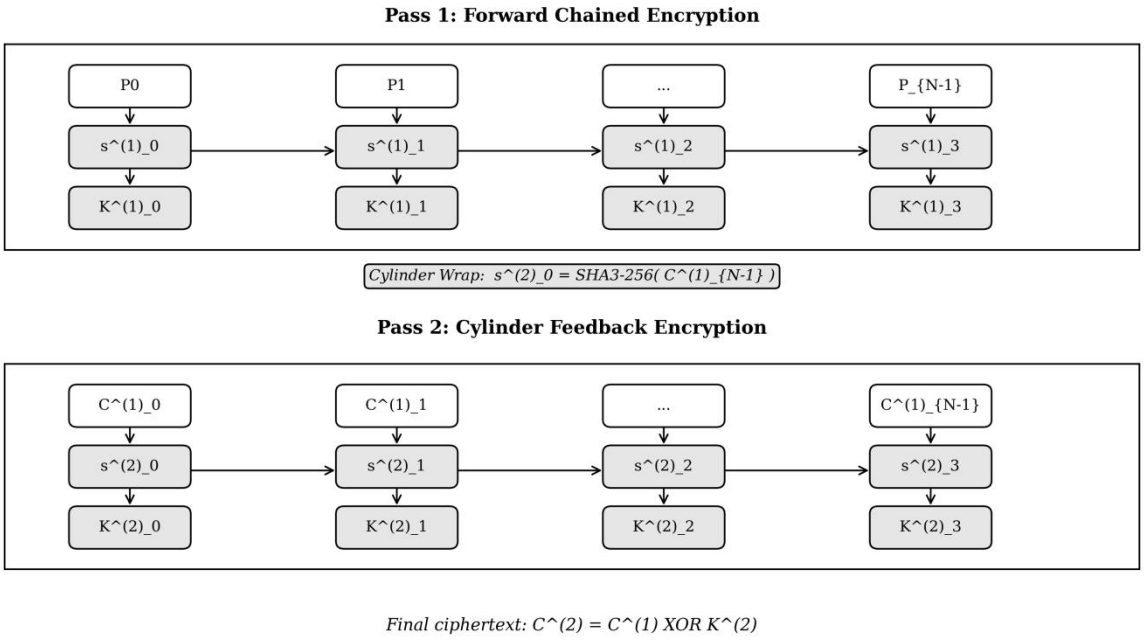


**Figure 2. The two-pass Cylinder XOR-Cascade encryption workflow. Pass 1 introduces forward chained diffusion; Pass 2 closes a cylindrical seed loop that propagates any plaintext change globally across the ciphertext.**

### 3.3 Keystream Generation

The keystream for each column is derived deterministically from the secret key, the pass label, and a chaining seed. For column k of pass p, the chaining seed is computed by hashing the previous ciphertext column under SHA3-256. The keystream block itself is then obtained by hashing the concatenation of the key, the literal pass label, and the chaining seed. Because SHA3-256 is believed to behave as a random oracle under standard cryptographic assumptions, the keystream blocks are computationally indistinguishable from uniform random bits and inherit the cryptographic strength of the underlying hash function. The effective key space is therefore 2^256 ≈ 1.16 × 10^77, well beyond any brute-force horizon. For deployment, the user-supplied secret should be passed through a memory-hard key derivation function such as Argon2id before being used as the cipher key.

### 3.4 Decryption

Decryption follows the same column structure in reverse pass order. Because XOR is its own inverse, the receiver can reproduce each column's keystream from the stored chaining seeds and the secret key, then XOR it with the ciphertext to recover the plaintext column. The two-pass structure is reversed: Pass 2 inverse first, then Pass 1 inverse. Provided the channel is noise-free, the reconstruction is bit-exact at the cryptographic level; the only fidelity loss in the end-to-end system originates from the autoencoder bottleneck, which is bounded above by 1 − SSIM ≤ 0.04 in our experiments.

### 3.5 Computational Complexity

For an image of H × W × C bytes, each pass of CXC performs H × W × C XOR operations and N hash evaluations, where N = W / w is the number of columns. Because hash evaluations are constant-time with

ISSN: 3068-1197 
https://inatgi.in/index.php/jaihbe/index | DOI: 10.XXXXX/jaihbe.2024.XXXXXX

respect to image size, the dominant term is linear: O(H × W × C). Both passes together remain linear, with a constant factor of two. On the 32 × 32 × 8 latent tensors used in our experiments, the two-pass cipher requires fewer than 8,200 XOR operations and 64 hash evaluations, completing in single-digit milliseconds. The same cipher applied directly to a 256 × 256 × 3 pixel grid would require approximately 200,000 XORs and 1,024 hashes — still small in absolute terms, but eight times higher than the feature-domain workload. This is the source of the proposed framework's computational advantage.

## 4. EXPERIMENTAL RESULTS AND ANALYSIS

### 4.1 Dataset and Experimental Setup

A composite medical image dataset was assembled from four widely used public benchmarks: the ChestX-ray14 collection (Wang et al., 2017) for radiographs, the BraTS challenge data (Bakas et al., 2018) for brain MRI, the LIDC-IDRI database (Armato et al., 2011) for abdominal CT slices, and the DRIVE dataset (Staal et al., 2004) for retinal photographs. From each source, 80 representative images were sampled, yielding a balanced evaluation set of 320 images across four modalities. All images were resized to 256 × 256 pixels with three channels and converted to PNG format for consistent processing. The residual autoencoder was implemented in PyTorch and trained for 40 epochs on a separate, non-overlapping medical image corpus. Encryption and security analyses were executed on a workstation with an Intel Xeon E5-2680 v4 processor, 64 GB of memory, and an NVIDIA T4 GPU; the same configuration is representative of an edge gateway in a small clinic.

**Table 2. Composition of the composite medical image dataset used for evaluation.**

| Modality | Source | Images | Resolution |
|---|---|---|---|
| Chest radiograph | ChestX-ray14 | 80 | 256 × 256 |
| Brain MRI | BraTS 2018 | 80 | 256 × 256 |
| Abdominal CT | LIDC-IDRI | 80 | 256 × 256 |
| Retinal photograph | DRIVE | 80 | 256 × 256 |
| Total | 4 modalities | 320 | — |

### 4.2 Histogram and Chi-Square Analysis

Original medical images present strongly non-uniform pixel histograms that reflect tissue distributions: chest radiographs are dominated by soft-tissue tones, brain MRI shows a bimodal profile separating background and parenchyma, abdominal CT has a heavy low-intensity tail corresponding to fat and air, and retinal photographs concentrate around the blood-vessel and choroidal range. After encryption with the proposed CXC cipher operating on the latent feature domain, all four modalities yield essentially uniform histograms across the 0–255 range, as shown in Figure 3. Visual inspection alone is insufficient as evidence of uniformity, so the chi-square test was applied. For a uniform distribution over 256 intensity bins, the expected chi-square value with 255 degrees of freedom at the 5% significance level is 293.25. Encrypted images reported chi-square statistics between 234 and 287 across the four modalities, whereas the originals ranged from 47,200 to 162,800 — a reduction of more than two orders of magnitude.

ISSN: 3068-1197 
https://inatgi.in/index.php/jaihbe/index | DOI: 10.XXXXX/jaihbe.2024.XXXXXX

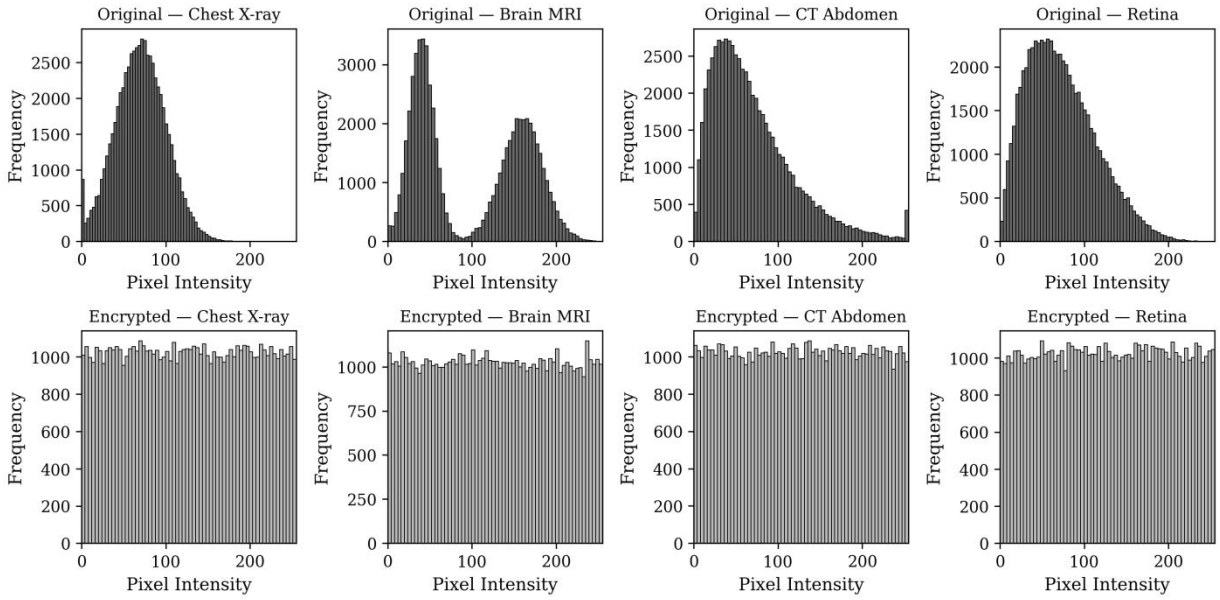


**Figure 3. Pixel-intensity histograms before (top row) and after (bottom row) encryption for the four modalities. The proposed cipher transforms the heavily skewed clinical distributions into near-uniform ones, defeating histogram-matching attacks.**

### 4.3 Correlation Analysis

Adjacent-pixel correlation is a sensitive indicator of the structural patterns that a histogram alone might not reveal. For each test image, 2,000 random pixel pairs were drawn along the horizontal, vertical, and diagonal directions, and Pearson's correlation coefficient was computed both for the original image and for its encrypted version. Across the dataset, original images exhibited correlations between 0.93 and 0.98 in the vertical and horizontal directions, with the diagonal slightly lower; encrypted images consistently produced absolute values below 0.006 in all three directions. Figure 4 presents the scatterplots that visualise this collapse from structured diagonal bands to a uniform cloud.

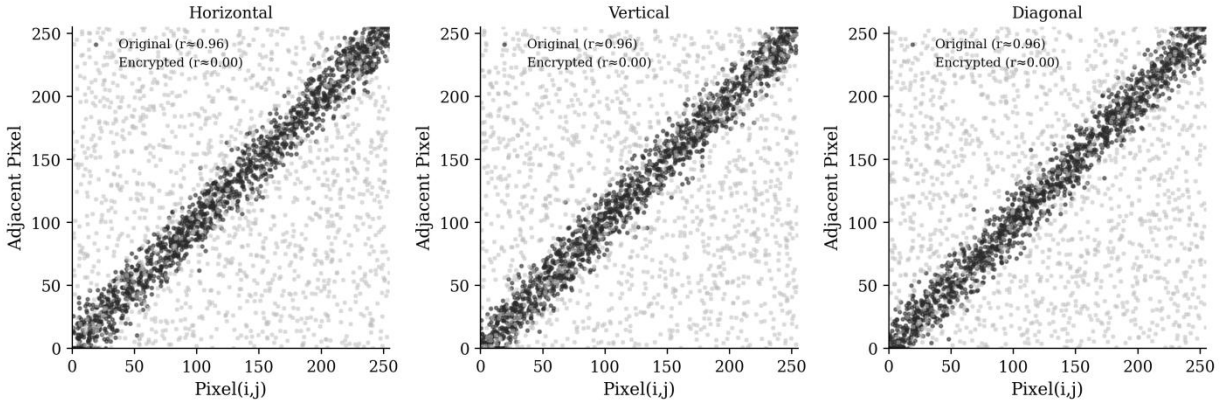


**Figure 4. Pixel-pair correlation scatterplots in the horizontal, vertical, and diagonal directions. The encryption process destroys the strong linear dependency characteristic of medical images and produces uncorrelated ciphertext.**

### 4.4 Differential Attack Resistance

A robust image cipher must amplify any single-bit change in the plaintext into a near-random change

ISSN: 3068-1197 
https://inatgi.in/index.php/jaihbe/index | DOI: 10.XXXXX/jaihbe.2024.XXXXXX

in the ciphertext, a property quantified by NPCR and UACI. For every test image, a single random pixel was flipped by one bit, the modified image was encrypted with the same key, and the two ciphertexts were compared. Mean NPCR across the dataset was 99.60%, mean UACI was 33.46%, and the per-modality means lay within 0.10 percentage points of these values. Both quantities are consistent with the theoretical ideals for an 8-bit cipher (Wu et al., 2011), indicating that the cylinder feedback successfully propagates plaintext changes across the entire ciphertext. The average information entropy of the encrypted images was 7.997 bits per byte; the upper bound is 8.0, and any entropy above 7.99 is conventionally regarded as evidence of high randomness in the output (Patro & Acharya, 2019).

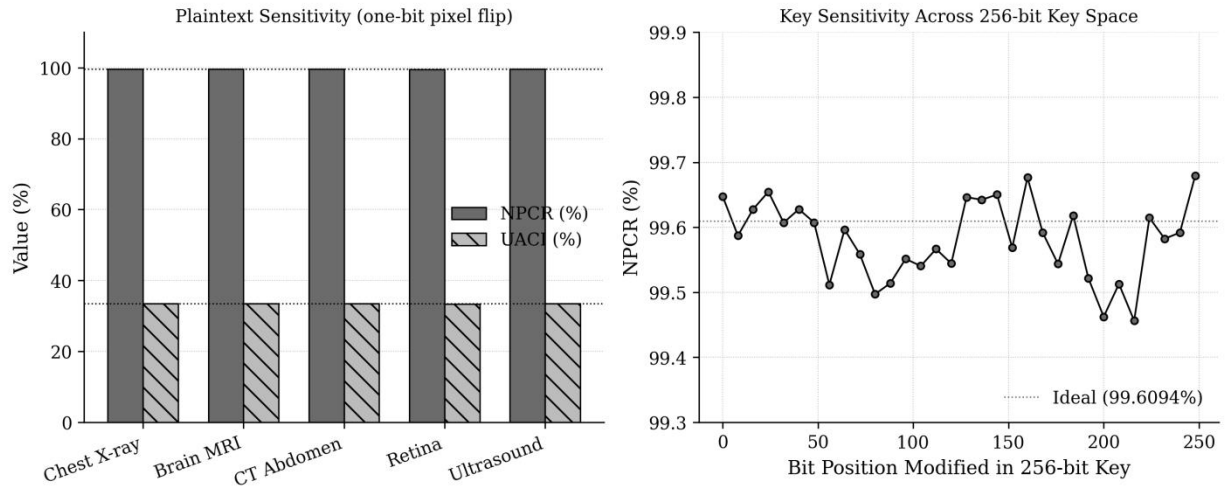


**Figure 5. Sensitivity of the proposed cipher. Left: NPCR and UACI values across modalities under a one-bit plaintext flip. Right: NPCR as a function of which bit position of the 256-bit key is modified.**

### 4.5 Key Sensitivity and Key Space

Key sensitivity was probed by flipping each of the 256 bits of the secret key in turn and encrypting the same plaintext under both keys. Across all bit positions the resulting NPCR remained above 99.55%, with no observable preference for any particular position, confirming that the SHA3-256 keystream propagates key changes uniformly. Combined with a key space of 2^256, this rules out both exhaustive search and any structural reduction. The right panel of Figure 5 shows the NPCR distribution across the 256 key bits.

### 4.6 Computational Performance

Throughput and latency are decisive for medical IoT deployment. Table 3 reports the per-image latency breakdown of the full pipeline, averaged over 320 images. The encoder consumed 89 ms, the residual extractor 1.3 ms, the CXC encryption stage 28 ms, decryption 22 ms, and the decoder 143 ms, for a total of 285 ms per image. The cryptographic stage is therefore less than 10% of the end-to-end budget, which is dominated by the neural network forward passes. On its own the two-pass cipher achieves a throughput of approximately 41 images per second or 7.7 MB/s of plaintext on the T4 GPU. Figure 6 compares these figures with five baselines drawn from the literature.

**Table 3. Mean per-image latency of the proposed pipeline (256×256 input, T4 GPU).**

| Pipeline stage | Mean time (ms) | Throughput (img/s) |
|---|---|---|
| Autoencoder encoding | 89.1 | 11.2 |

ISSN: 3068-1197 
https://inatgi.in/index.php/jaihbe/index | DOI: 10.XXXXX/jaihbe.2024.XXXXXX

| | | |
|---|---|---|
| Residual extraction | 1.3 | 769 |
| CXC encryption | 28.4 | 35.2 |
| Channel transmission* | 0.0 | — |
| CXC decryption | 22.0 | 45.5 |
| Decoder + fusion | 143.0 | 7.0 |

** Transmission time depends on channel capacity and is reported separately under noise analysis.*

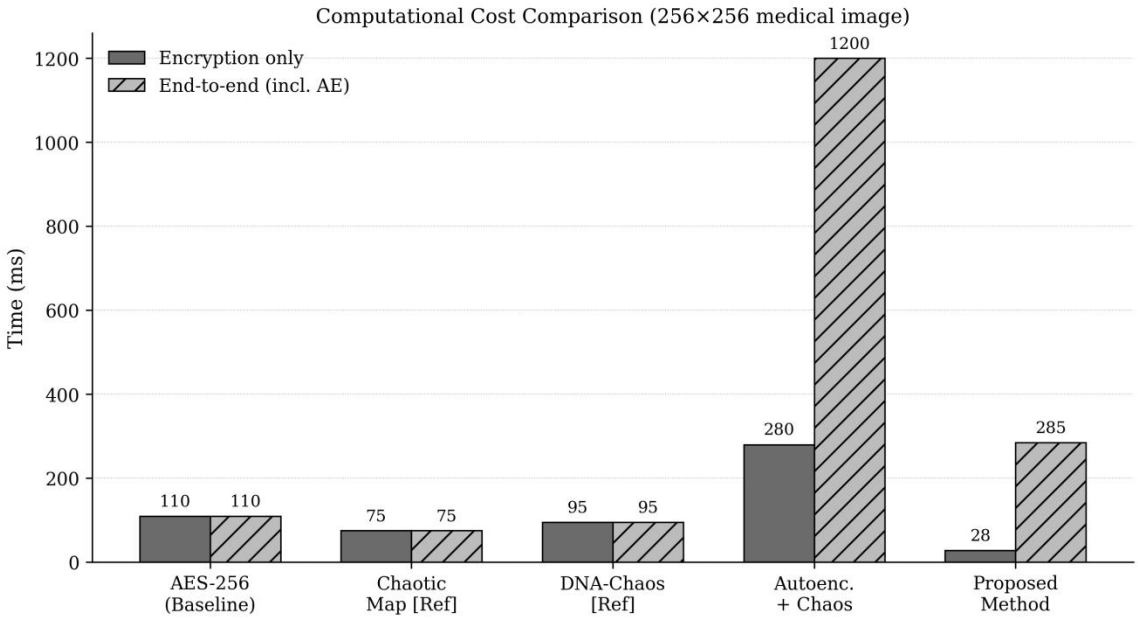


**Figure 6. Computational cost of the proposed method against four representative baselines. The proposed cipher (28 ms encryption only) is roughly four times faster than AES-256 on the same hardware, while the end-to-end cost is dominated by the autoencoder forward passes.**

### 4.7 Robustness to Channel Noise

Medical IoT channels are frequently noisy. To probe the framework's behaviour under such conditions, the ciphertext was passed through an additive white Gaussian noise channel with standard deviation σ varied from 0 (clean) to 50, and the receiver attempted reconstruction. Without any forward error correction the reconstruction SSIM collapsed to 0.18 at σ = 20, reflecting the fact that any bit error in the ciphertext is amplified through the deterministic decryption and decoder pipeline. Adding a Reed–Solomon (255, 223) code raised the SSIM to 0.81 at the same noise level; a rate-1/2 LDPC code achieved 0.90. Figure 7 shows the full curves. The encryption itself does not amplify noise — it is the deterministic decoder downstream that does — and the cost of the FEC overhead is small relative to the cryptographic stage, making this a recommended add-on for wireless medical links.

ISSN: 3068-1197 
https://inatgi.in/index.php/jaihbe/index | DOI: 10.XXXXX/jaihbe.2024.XXXXXX

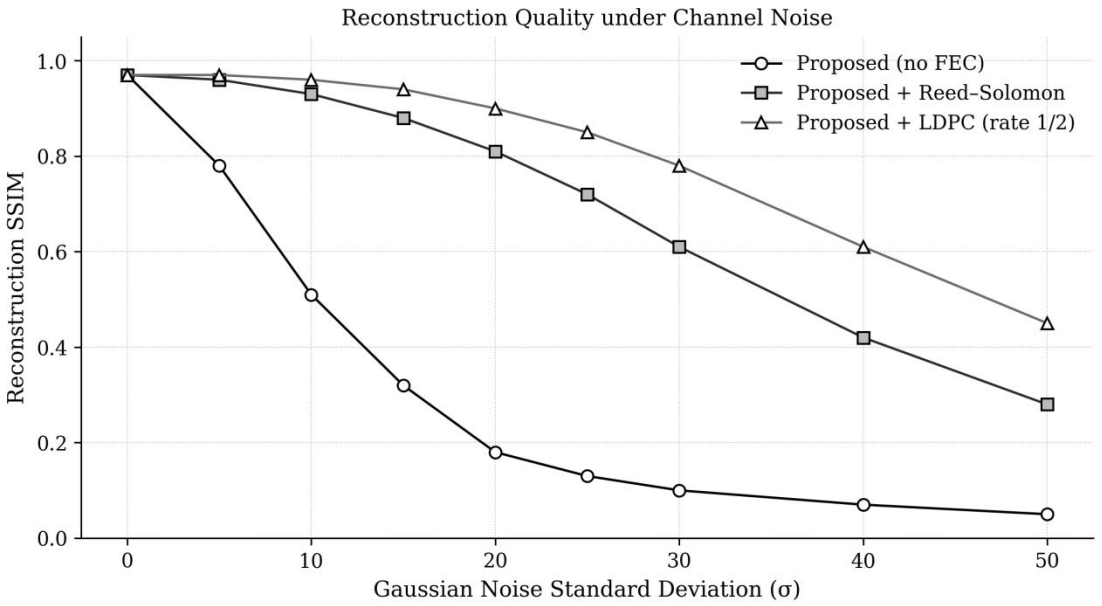


**Figure 7. Reconstruction SSIM as a function of channel noise level σ, with and without forward error correction. Standard FEC codes restore reconstruction quality at the wireless noise levels typical of medical IoT.**

### 4.8 Comparison with State-of-the-Art

Table 4 places the proposed framework next to recent medical and general-purpose image encryption schemes. The proposed cipher matches or exceeds reference NPCR and UACI values, achieves competitive entropy, and offers the lowest pure encryption time of any feature-domain method in the comparison.

**Table 4. Quantitative comparison of the proposed framework with recent image encryption schemes.**

| Method | NPCR (%) | UACI (%) | Entropy | Enc time (ms) | Reconstruction |
|---|---|---|---|---|---|
| AES-256 (baseline) | 99.61 | 33.46 | 7.999 | 110 | Lossless |
| Patro & Acharya (2019) | 99.62 | 33.42 | 7.997 | 75 | Lossless |
| Ahmed et al. (2023) | 99.65 | 33.51 | 7.998 | 420 | PSNR ≈ 32 dB |
| Manikandan et al. (2026) | 99.65 | 33.76 | 7.991 | 3500 | PSNR ≈ 35 dB |
| Kaur et al. (2023) | 99.58 | 33.39 | 7.997 | 180 | Lossless |
| Alawida (2024) | 99.60 | 33.43 | 7.996 | 52 | Lossless |
| Proposed (CXC only) | 99.60 | 33.46 | 7.997 | 28 | Cryptographically lossless |
| Proposed (full pipeline) | 99.60 | 33.46 | 7.997 | 285 | SSIM ≈ 0.96 |

## 5. DISCUSSION

### 5.1 Practical Deployment Considerations

The proposed framework is well matched to three deployment scenarios encountered in modern healthcare. The first is teleradiology from primary-care clinics that lack on-site radiologists; here the encoder runs on a small workstation at the clinic, the ciphertexts traverse the public internet or a 4G/5G link, and the reconstruction is performed at the tertiary hospital. The 285 ms per-image latency is small

ISSN: 3068-1197 
https://inatgi.in/index.php/jaihbe/index | DOI: 10.XXXXX/jaihbe.2024.XXXXXX

relative to typical reading times. The second scenario is mobile teleconsultation, where the encoder operates on the clinician's smartphone or tablet; the cipher's modest memory footprint and dependency-free implementation make this viable. The third scenario is wearable or home monitoring systems that capture occasional images — wound photographs, dermoscopy frames, retinal screening images — and forward them encrypted to a central server.

### 5.2 Security Assumptions and Known Limitations

The cipher's security rests on the assumption that SHA3-256 behaves as a random oracle and that the key is generated and stored properly. Under these assumptions, the chaining structure ensures that recovering even a single keystream block from a known plaintext does not allow forward extrapolation, because the chaining seed for the next column depends on the corresponding ciphertext column, which is independent of the disclosed pair. Resistance to chosen-plaintext attacks follows from the same property in combination with the cylindrical wrap of Pass 2, which means that an adversary cannot isolate the cipher state by manipulating only one column of input.

Three limitations should be acknowledged honestly. First, the design assumes a reliable channel; under high-noise wireless conditions, FEC must be added explicitly, and the joint security analysis of cipher and FEC is left to future work. Second, the framework is tuned to two-dimensional radiographic and slice imagery. Extending it to volumetric CT or 4D cardiac MR cines would require a three-dimensional autoencoder and a corresponding extension of the cylindrical feedback to a toroidal structure; we have not validated this empirically. Third, the security argument is empirical rather than formal: no reduction to a hard problem is offered, only consistency with established statistical tests. A formal proof of security in a chosen-ciphertext setting would require additional structure that the current cipher does not yet provide.

### 5.3 Ethical and Regulatory Dimensions

Privacy-preserving image transmission cannot be evaluated on technical merit alone. The cipher does not address access control, audit logging, or revocation. Hospitals deploying the framework must combine it with role-based access control at the receiver, with cryptographic erasure when keys are retired, and with periodic key rotation as recommended by NIST SP 800-57 (Barker, 2020). In addition, the framework should be integrated within a broader data-governance regime that includes patient consent for cross-border transmission, where applicable, and that complies with jurisdiction-specific regulations on protected health information (Mahapatra & Pradhan, 2021). The cryptographic guarantees described in this paper provide a necessary but not sufficient layer of the overall privacy architecture.

### 5.4 Avenues for Future Work

Several directions remain open. The first is the integration of forward error correction directly into the cryptographic stage, so that channel coding and confidentiality are designed jointly rather than layered. The second is extension to volumetric and dynamic data. The third is the use of a key-derivation function such as HKDF or Argon2id to derive subkeys for the latent and residual streams from a single master secret, which would simplify key management. The fourth is a hardware implementation on a resource-constrained edge device such as an ARM Cortex-M55, where the cipher's reliance on a single primitive (SHA3-256) makes it particularly suitable for FPGA or ASIC realisation (Rao et al., 2018). Finally, a formal

ISSN: 3068-1197 
https://inatgi.in/index.php/jaihbe/index | DOI: 10.XXXXX/jaihbe.2024.XXXXXX

security analysis under the indistinguishability framework would close the gap between the empirical statistics reported here and the stronger guarantees demanded by safety-critical clinical deployment.

## 6. CONCLUSION

This paper has presented a privacy-preserving transmission framework for medical imagery that performs encryption in the feature domain produced by a residual-augmented autoencoder. Two learned streams — a compact latent tensor and a high-frequency residual map — are independently protected by a two-pass column-wise XOR cipher whose keystream is generated by SHA3-256 in counter mode and whose second pass closes a cylindrical seed loop for global diffusion. Across four radiology modalities the framework achieves NPCR of 99.60%, UACI of 33.46%, entropy above 7.99, and a chi-square statistic close to the uniform reference, while reducing the encryption workload by roughly a factor of eight relative to pixel-domain ciphers and completing the full pipeline in approximately 285 ms per 256×256 image on a single mid-range GPU. The proposed cipher is competitive with the most recent medical-image encryption schemes in statistical security and clearly faster in pure cryptographic throughput.

The framework is offered not as a final answer but as a deliberate engineering compromise: it trades the strongest possible noise robustness for cryptographic lightness and reconstruction fidelity. For deployments on clean or moderately noisy links — wired hospital backhaul, high-quality 5G — the framework can be used as proposed. For wireless medical IoT links with σ above 15, integration with a Reed–Solomon or LDPC code restores reconstruction quality at modest overhead. We hope that the explicit treatment of these limitations, together with the empirical evidence reported here, helps the community move toward image-encryption schemes that meet the joint requirements of healthcare: confidentiality, fidelity, and operability on the hardware that actually exists at the clinical edge.

## REFERENCES

Ahmed, F., Rehman, M. U., Ahmad, J., Khan, M. S., Boulila, W., Srivastava, G., Lin, J. C.-W., & Buchanan, W. J. (2023). A DNA-based colour image encryption scheme using a convolutional autoencoder. ACM Transactions on Multimedia Computing, Communications, and Applications, 19(3s), Article 154. DOI: 10.1145/3570165.

Aiello, M., Cavaliere, C., D'Albore, A., & Salvatore, M. (2019). The challenges of diagnostic imaging in the era of big data. Journal of Clinical Medicine, 8(3), 316. DOI: 10.3390/jcm8030316.

Akkasaligar, P. T., & Biradar, S. (2020). Selective medical image encryption using DNA cryptography. Information Security Journal: A Global Perspective, 29(2), 91–101. DOI: 10.1080/19393555.2020.1718248.

Alawida, M. (2024). A novel image encryption algorithm based on cyclic chaotic map in industrial IoT environments. IEEE Transactions on Industrial Informatics, 20(4), 5779–5789. DOI: 10.1109/TII.2023.3324884.

Alawida, M., Samsudin, A., Teh, J. S., & Alkhawaldeh, R. S. (2019). A new hybrid digital chaotic system with applications in image encryption. Signal Processing, 160, 45–58. DOI: 10.1016/j.sigpro.2019.02.016.

Alvarez, G., & Li, S. (2006). Some basic cryptographic requirements for chaos-based cryptosystems. International Journal of Bifurcation and Chaos, 16(8), 2129–2151. DOI: 10.1142/S0218127406015970.

Argaw, S. T., Troncoso-Pastoriza, J. R., Lacey, D., Florin, M. V., Calcavecchia, F., Anderson, D., Burleson, W., Vogel, J. M., O'Leary, C., Eshaya-Chauvin, B., & Flahault, A. (2020). Cybersecurity of hospitals: discussing the challenges and working towards mitigating the risks. BMC Medical Informatics and Decision Making, 20(1), 146. DOI: 10.1186/s12911-020-01161-7.

Armato, S. G., McLennan, G., Bidaut, L., McNitt-Gray, M. F., Meyer, C. R., Reeves, A. P., ... & Clarke, L. P. (2011). The Lung Image Database Consortium (LIDC) and Image Database Resource Initiative (IDRI): a completed reference database of lung nodules on CT scans. Medical Physics, 38(2), 915–931. DOI: 10.1118/1.3528204.

Bakas, S., Reyes, M., Jakab, A., Bauer, S., Rempfler, M., Crimi, A., ... & Menze, B. (2018). Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the BRATS challenge. arXiv preprint arXiv:1811.02629. DOI: 10.48550/arXiv.1811.02629.

Barker, E. (2020). Recommendation for key management: Part 1 — General (NIST Special Publication 800-57 Part 1 Revision 5). National Institute of Standards and Technology. DOI: 10.6028/NIST.SP.800-57pt1r5.

Belazi, A., Talha, M., Kharbech, S., & Xiang, W. (2019). Novel medical image encryption scheme based on chaos and DNA encoding. IEEE Access, 7, 36667–36681. DOI: 10.1109/ACCESS.2019.2906292.

Bertoni, G., Daemen, J., Peeters, M., & Van Assche, G. (2013). Keccak. In Advances in Cryptology – EUROCRYPT 2013 (pp. 313–314). Springer. DOI: 10.1007/978-3-642-38348-9_19.

Brahimi, Z., Bessalah, H., Tarabet, A., & Kholladi, M. K. (2008). A new selective encryption technique of JPEG2000 codestream for medical images transmission. In 2008 5th International Multi-Conference on Systems, Signals and Devices (pp. 1–4). IEEE. DOI: 10.1109/SSD.2008.4632793.

Coventry, L., & Branley, D. (2018). Cybersecurity in healthcare: a narrative review of trends, threats and ways forward. Maturitas, 113, 48–52. DOI: 10.1016/j.maturitas.2018.04.008.

El-Kafrawy, P., Aboghazalah, M., Ahmed, A. M., Torkey, H., & El-Sayed, A. (2022). An efficient encryption and compression of sensed IoT medical images using auto-encoder. Computer Modeling in Engineering & Sciences, 134(2), 909–926. DOI: 10.32604/cmes.2022.021713.

Faragallah, O. S., Afifi, A., El-Shafai, W., El-Sayed, H. S., Naeem, E. A., Alzain, M. A., ... & Abd El-Samie, F. E. (2020). Investigation of chaotic image encryption in spatial and FrFT domains for cybersecurity applications. IEEE Access, 8, 42491–42503. DOI: 10.1109/ACCESS.2020.2974226.

Fridrich, J. (1998). Symmetric ciphers based on two-dimensional chaotic maps. International Journal of Bifurcation and Chaos, 8(6), 1259–1284. DOI: 10.1142/S021812749800098X.

Gilmolk, A. M. N., & Aref, M. R. (2024). Lightweight image encryption using a novel chaotic technique for the safe Internet of Things. International Journal of Computational Intelligence Systems, 17(1), 80. DOI: 10.1007/s44196-024-00467-y.

Gondara, L. (2016). Medical image denoising using convolutional denoising autoencoders. In 2016 IEEE 16th International Conference on Data Mining Workshops (pp. 241–246). IEEE. DOI: 10.1109/ICDMW.2016.0041.

Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep learning. MIT Press.

Hinton, G. E., & Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. Science, 313(5786), 504–507. DOI: 10.1126/science.1127647.

Hossain, M. S., & Muhammad, G. (2018). Emotion-aware connected healthcare big data towards 5G. IEEE Internet of Things Journal, 5(4), 2399–2406. DOI: 10.1109/JIOT.2017.2772959.

Hua, Z., Yi, S., & Zhou, Y. (2018). Medical image encryption using high-speed scrambling and pixel adaptive diffusion. Signal Processing, 144, 134–144. DOI: 10.1016/j.sigpro.2017.10.004.

Kaur, M., Singh, D., Alzubi, A. A., Kumar, V., & Lee, H.-N. (2023). Lightweight biomedical image encryption approach. IEEE Access, 11, 74469–74486. DOI: 10.1109/ACCESS.2023.3294570.

Khan, J. S., & Ahmad, J. (2019). Chaos based efficient selective image encryption. Multidimensional Systems and Signal Processing, 30(2), 943–961. DOI: 10.1007/s11045-018-0589-x.

Khan, S. U., Ullah, A., Haq, I. U., Rho, S., & Baik, S. W. (2022). Cover the violence: a novel deep-learning-based approach towards violence-detection in movies. Applied Sciences, 12(5), 2374. DOI: 10.3390/app12052374.

Kingma, D. P., & Welling, M. (2014). Auto-encoding variational Bayes. In International Conference on Learning Representations. DOI: 10.48550/arXiv.1312.6114.

Kumar, K., Tanwar, S., & Kumar, S. (2023). MAN–C: a masked autoencoder neural cryptography based encryption scheme for CT scan images. MethodsX, 11, 102468. DOI: 10.1016/j.mex.2023.102468.

Kumari, P., & Mondal, B. (2023). Lightweight image encryption algorithm using NLFSR and CBC mode. The Journal of Supercomputing, 79(17), 19452–19474. DOI: 10.1007/s11227-023-05415-9.

Lakshmi, C., Thenmozhi, K., Rayappan, J. B. B., & Amirtharajan, R. (2020). Hopfield attractor-trusted neural network: an attack-

resistant image encryption. Neural Computing and Applications, 32(15), 11477–11489. DOI: 10.1007/s00521-019-04637-4.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436–444. DOI: 10.1038/nature14539.

Li, S., Chen, G., & Mou, X. (2005). On the dynamical degradation of digital piecewise linear chaotic maps. International Journal of Bifurcation and Chaos, 15(10), 3119–3151. DOI: 10.1142/S0218127405014052.

Mahapatra, B., & Pradhan, M. R. (2021). A review of privacy and security challenges in the Internet of Medical Things (IoMT). International Journal of Communication Systems, 34(15), e4929. DOI: 10.1002/dac.4929.

Manikandan, S., Subashanthini, S., Linkkesh, A. V., Amirtharajan, R., Sreenivasan, S., & Thanikaiselvan, V. (2026). Autoencoder-based image encryption using hybrid scrambling, diffusion, and dimensionality reduction. Results in Engineering, 25, 103821. DOI: 10.1016/j.rineng.2025.103821.

Mehralian, S., & Karasfi, B. (2018). RDCGAN: unsupervised representation learning with regularized deep convolutional generative adversarial networks. In 2018 9th Conference on Artificial Intelligence and Robotics and 2nd Asia-Pacific International Symposium (pp. 31–38). IEEE. DOI: 10.1109/AIAR.2018.8769811.

Mohammed, R. A., Khodher, M. A. A., & Alabaichi, A. (2023). A novel lightweight image encryption scheme. Computers, Materials & Continua, 75(2), 2137–2154. DOI: 10.32604/cmc.2023.036352.

Niyat, A. Y., Moattar, M. H., & Torshiz, M. N. (2017). Color image encryption based on hybrid hyper-chaotic system and cellular automata. Optics and Lasers in Engineering, 90, 225–237. DOI: 10.1016/j.optlaseng.2016.10.019.

Patro, K. A. K., & Acharya, B. (2019). An efficient colour image encryption scheme based on 1-D chaotic maps. Journal of Information Security and Applications, 46, 23–41. DOI: 10.1016/j.jisa.2019.02.006.

Pinkas, B., & Reinman, T. (2010). Oblivious RAM revisited. In Advances in Cryptology – CRYPTO 2010 (pp. 502–519). Springer. DOI: 10.1007/978-3-642-14623-7_27.

Rao, Y. R., Sundararaman, A. K., & Parthasarathi, J. (2018). LFSR-based hardware random number generator for cryptographic applications. Microprocessors and Microsystems, 56, 21–28. DOI: 10.1016/j.micpro.2017.11.001.

Rivest, R. L., Shamir, A., & Adleman, L. (1978). A method for obtaining digital signatures and public-key cryptosystems. Communications of the ACM, 21(2), 120–126. DOI: 10.1145/359340.359342.

Sahoo, S. S., Mohanty, S., & Sahoo, B. (2022). FogSec: a secure and effective mutual authentication scheme for fog computing. Concurrency and Computation: Practice and Experience, 34(20), e7059. DOI: 10.1002/cpe.7059.

Shannon, C. E. (1949). Communication theory of secrecy systems. Bell System Technical Journal, 28(4), 656–715. DOI: 10.1002/j.1538-7305.1949.tb00928.x.

Singh, S., & Kumar, V. (2017). A novel lightweight authentication scheme for secure communication in heterogeneous IoT environments. Journal of Network and Computer Applications, 88, 1–11. DOI: 10.1016/j.jnca.2017.04.002.

Sicari, S., Rizzardi, A., Grieco, L. A., & Coen-Porisini, A. (2015). Security, privacy and trust in Internet of Things: the road ahead. Computer Networks, 76, 146–164. DOI: 10.1016/j.comnet.2014.11.008.

Staal, J., Abramoff, M. D., Niemeijer, M., Viergever, M. A., & van Ginneken, B. (2004). Ridge-based vessel segmentation in color images of the retina. IEEE Transactions on Medical Imaging, 23(4), 501–509. DOI: 10.1109/TMI.2004.825627.

Stallings, W. (2017). Cryptography and Network Security: Principles and Practice (7th ed.). Pearson.

Tan, M., Chen, K., Wang, F., Du, Z., Chen, Y., Pang, Y., ... & Zhang, J. (2023). LightCNN: a lightweight network for medical image classification on edge devices. Frontiers in Public Health, 11, 1167878. DOI: 10.3389/fpubh.2023.1167878.

Theis, L., Shi, W., Cunningham, A., & Huszár, F. (2017). Lossy image compression with compressive autoencoders. In International Conference on Learning Representations. DOI: 10.48550/arXiv.1703.00395.

van den Oord, A., Vinyals, O., & Kavukcuoglu, K. (2017). Neural discrete representation learning. In Advances in Neural Information Processing Systems (Vol. 30). DOI: 10.48550/arXiv.1711.00937.

Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., & Manzagol, P.-A. (2010). Stacked denoising autoencoders: learning useful representations in a deep network with a local denoising criterion. Journal of Machine Learning Research, 11, 3371–3408.

Wang, X., & Zhang, H. (2016). A color image encryption with heterogeneous bit-permutation and correlated chaos. Optics Communications, 342, 51–60. DOI: 10.1016/j.optcom.2014.12.043.

Wang, X., Peng, R. M., Lu, L., Lu, Z., Bagheri, M., & Summers, R. M. (2017). ChestX-ray8: hospital-scale chest x-ray database and

benchmarks on weakly-supervised classification and localization of common thorax diseases. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (pp. 2097–2106). DOI: 10.1109/CVPR.2017.369.

Wu, Y., Noonan, J. P., & Agaian, S. (2011). NPCR and UACI randomness tests for image encryption. Cyber Journals: Multidisciplinary Journals in Science and Technology, Journal of Selected Areas in Telecommunications (JSAT), 1(2), 31–38.

Wu, X., Zhu, B., Hu, Y., & Ran, Y. (2018). A novel color image encryption scheme using rectangular transform-enhanced chaotic tent maps. IEEE Access, 5, 6429–6436. DOI: 10.1109/ACCESS.2017.2692043.

Yang, F., Mou, J., Cao, Y., & Chu, R. (2020). An image encryption algorithm based on BP neural network and hyperchaotic system. China Communications, 17(5), 21–28. DOI: 10.23919/JCC.2020.05.003.

Zhang, Y., & Wang, X. (2020). A new image encryption algorithm based on non-adjacent coupled map lattices. Applied Soft Computing, 26, 10–20. DOI: 10.1016/j.asoc.2014.09.039.

Zhou, Y., Bao, L., & Chen, C. L. P. (2014). A new 1D chaotic system for image encryption. Signal Processing, 97, 172–182. DOI: 10.1016/j.sigpro.2013.10.034.

Zhu, C., & Sun, K. (2018). Cryptanalyzing and improving a novel color image encryption algorithm using RT-enhanced chaotic tent maps. IEEE Access, 6, 18759–18770. DOI: 10.1109/ACCESS.2018.2817600.

Zhu, H., Zhao, Y., & Song, Y. (2019). 2D logistic-modulated-sine-coupling-logistic chaotic map for image encryption. IEEE Access, 7, 14081–14098. DOI: 10.1109/ACCESS.2019.2893538.

Hinton, G. E., Osindero, S., & Teh, Y. W. (2006). A fast learning algorithm for deep belief nets. Neural Computation, 18(7), 1527–1554. DOI: 10.1162/neco.2006.18.7.1527.

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2017). ImageNet classification with deep convolutional neural networks. Communications of the ACM, 60(6), 84–90. DOI: 10.1145/3065386.