International Journal of Safety and Security Engineering
Vol. 16, No. 6, June, 2026, pp. 1361-1372
Journal homepage: http://iieta.org/journals/ijsse

Robust Loss-Weight Optimization for Medical Image Encryption via Bayesian UR-CycleGAN

Mays Y. Mhawi[1*], Hikmat N. Abdullah[1], Axel Sikora[2]

[1] Department of Information and Communication Engineering, College of Information Engineering, Al-Nahrain University, Baghdad 10001, Iraq
[2] Institute of Reliable Embedded Systems and Communication Electronics, University of Applied Sciences, Offenburg 77652, Germany

Corresponding Author Email: maysyousif92@gmail.com

https://doi.org/10.18280/ijsse.160617

Received: 30 April 2026
Revised: 15 June 2026
Accepted: 22 June 2026
Available online: 30 June 2026

Keywords:
CycleGAN, medical image encryption, adversarial attacks, Bayesian optimization, loss weights optimization

ABSTRACT

In modern healthcare, medical image encryption ensures privacy and secure transmission of sensitive patient data. Recently, CycleGAN-based frameworks have emerged as promising alternatives to conventional image encryption. CycleGAN is an unsupervised deep learning model for image-to-image translation. Despite these advantages, the loss weights within CycleGAN objective function significantly influence model performance, as they regulate the trade-off between image reconstruction fidelity and encryption security. In many existing studies, these weights are selected heuristically, which often leads to unstable training and suboptimal compromises between image quality and security strength. To address this limitation, this work proposes a robust loss-weight optimization framework for medical image encryption based on a Bayesian-optimized Uniformity-Regularized CycleGAN (Bayesian UR-CycleGAN). The proposed framework formulates the optimization process as a max–min problem, aiming to determine the optimal configuration of loss-weight hyperparameters that maximizes performance under worst-case conditions rather than optimizing average or best-case performance. To efficiently solve this optimization problem, Bayesian optimization (BO) is employed to explore the high-dimensional loss-weight space. The proposed framework is evaluated on medical imaging datasets using widely adopted image quality metrics, including Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM), as well as security-oriented metrics such as entropy, histogram, Chi-square test, and correlation. Experimental results demonstrate that the proposed framework consistently achieves a more robust balance between reconstruction quality and encryption security compared with baseline approaches that use fixed or manually tuned loss weights. Bayesian UR-CycleGAN is evaluated under various attack scenarios. Results show that the optimized framework is resistant to these attack models.

1. INTRODUCTION

Nowadays, the rapid development of Electronic Healthcare (e-healthcare) is both possible and popular. E-healthcare refers to an internet-based system that enables patients to access services from expert clinicians worldwide [1]. Transmission of medical images over a network is always vulnerable to attack [2]. Since medical images contain patients' confidential information, ensuring their safe storage and transmission between workstations and the Picture Archiving and Communication System (PACS), and over the internet, has become a critical issue for medical applications in real-world settings [2, 3].

Imaging Machines such as brain Magnetic Resonance Imaging (MRI), computed tomography (CT), X-rays, and dermatoscopes, etc. [4]. To secure the confidentiality, integrity, and authenticity of medical images, several encryption and decryption techniques can be used, such as the Data Encryption Standard (DES) and the Advanced Encryption Standard (AES). In these two encryption techniques, the same key is used for both encryption and decryption. DES uses a 64-bit block size, whereas AES uses 128-, 192-, and 256-bit block sizes. However, image sizes are larger than block sizes, which can lead to security issues due to slackness or disorganization [5]. Furthermore, Chaos-based image encryption techniques are widely used. Chaos-based image encryption is highly sensitive to initial conditions; a slight change in the input results in significant changes to the output. While the sensitivity to the initial condition is mainly advantageous in the encryption process, it also presents challenges in key management and error propagation [6].

Recently, deep learning techniques have been widely used to encrypt and decrypt medical images using neural networks that learn complex, convoluted patterns across multiple layers with nonlinear activation functions. Deep generative models are utilized to provide an understanding of unsupervised data.

For image encryption, deep learning is just beginning to emerge. In this field, generative adversarial network (GAN) has proven to be a powerful tool for generating images. This adaptability reassures us of their effectiveness in developing advanced encryption algorithms [7]. CycleGAN, a type of deep generative model, is designed to generate images across different modalities and to perform mappings that differ from those observed in the real training images, which is known as the image-to-image translation task [8]. It is an unsupervised model that learns to transform images from two distinct domains using unpaired training data [9].

It uses a deep convolutional neural network (CNN) to learn a mapping between two image domains, allowing it to perform tasks such as style transfer, super-resolution, domain adaptation, and image-to-image synthesis [4].

Unlike conventional techniques, the CycleGAN function does not require paired training data, enabling it to handle constrained data availability while still improving image quality and feature preservation. It has two main components: the Generator G and the discriminator D. The Generator G plays a critical role in the CycleGAN function, transforming images from the original domain to the target domain. This process is complemented by the discriminator D, which distinguishes between real images (from the target domain) and fake images (generated by the G) [4].

In practical healthcare environments, medical images are frequently transmitted via telemedicine platforms and stored in cloud-based healthcare infrastructure.

Therefore, the proposed encryption framework is designed to secure medical image transmission and storage by providing both image confidentiality and high reconstruction fidelity for diagnostic purposes.

A critical issue in CycleGAN-based medical image encryption is the need for a more systematic approach to selecting the appropriate loss weights. In CycleGAN, loss weights (λ) control the relative impact of each loss term, such as cycle-consistency, structural similarity, gradient preservation, and uniformity regularization within the overall objective function. Most existing approaches rely on fixed or experimentally chosen values of the loss weights (λ) without analyzing how these parameters relate to the dataset. This paper introduces a new architecture that combines a CycleGAN with a Bayesian optimization (BO) algorithm to address the challenge of selecting optimal loss weight combinations. By balancing exploitation and exploration within high-dimensional search spaces, the BO algorithm optimizes the loss-weight parameters of the CycleGAN objective function.

2. RELATED WORK AND CONTRIBUTIONS

The original CycleGAN framework [10] introduced loss weights (λ) to control the impact of each loss component, which are the identity mapping loss and cycle consistency loss. Despite the common use of $\lambda = 10$ for cycle consistency, there is no clear methodology for selecting these parameters, particularly in models with multiple loss terms that need balancing. The following section reviews existing works that use this configuration. Table 1 represents the summary of existing works.

Table 1. Summary of existing works

Ref.	Dataset	Proposed Loss Function	Hyperparameter Value	System Evaluation	Attack Analysis
[11]	Caltech-256	SSIM-thumbnail consistency loss	Manual tuning $\lambda_1 = 20$ $\lambda_2 = 0.9$	Key space; PSNR; SSIM	N/A
[12]	Skin cancer dataset	SSIM-cycle consistency loss	Default value $\lambda = 10$	Key space; PSNR; SSIM; histogram; correlation; entropy	Differential attacks
[13]	A color image from NASA Hubble space telescope images	BCE-reconstruction loss-feature loss	Manual tuning λ_1=2 λ_2=50	Key space; PSNR; SSIM; histogram; correlation; entropy	Differential attacks; noise attack
[4]	Skin cancer dataset	BCE	N/A	Key space; PSNR; SSIM; histogram; correlation; entropy; Chi-square	Differential attacks
[14]	Kodak dataset; Corel-1K; Lena image	Cycle consistency loss-SSIM-Block regularization loss	N/A	PSNR; SSIM; histogram; correlation; entropy	Differential attacks
[15]	Breast cancer dataset	BCE-SSIM-Cycle consistency	Default value $\lambda = 10$	PSNR; SSIM; histogram; correlation; entropy	Differential attacks
Proposed	MRI-Skin lesion	Adversarial attack-cycle consistency loss-SSIM -gradient loss-uniformity loss	Bayesian optimization algorithm	PSNR; SSIM; histogram analysis; Chi-square test; correlation; entropy; key space	Differential attacks-adversarial attacks

Recent CycleGAN-based image encryption frameworks have introduced various loss functions to improve reconstruction quality and encryption security. These approaches incorporate thumbnail consistency loss, adversarial, cycle-consistency, Structural Similarity Index (SSIM)-based reconstruction, Feature loss, and block regularization terms to preserve structural information, perceptual quality, and semantic features during the encryption and decryption processes [11-15].

In addition, several frameworks [4, 12, 15] employ multiple generators, discriminators, and decryption networks to achieve secure medical image transformation while maintaining high-quality image reconstruction.

Despite these advances, the corresponding loss-weighting parameters (λ) are generally selected empirically, directly adopted from the original CycleGAN formulation, or determined through limited experimental trials without a systematic optimization strategy [11-15]. Furthermore, several

studies either omit the values of these parameters or provide limited discussion of their influence on balancing the contributions of different loss components, making it difficult to reproduce their results and understand the relationship between loss-weight configuration, reconstruction quality, and encryption security. Consequently, determining an optimal set of loss-weight parameters remains an open research problem.

To address this limitation, this work introduces a systematic optimization framework to determine optimal loss-weight parameters. Specifically, a BO strategy is used to automatically search for the optimum combination of λ values within the CycleGAN model's loss function. This approach eliminates the need for manual tuning and provides a principled mechanism for achieving a more robust trade-off between competing objectives.

The significant contribution of this study is as follows:

(1) Introduced a BO algorithm to identify the loss weights of the cycle-consistency, SSIM, gradient-based losses, and uniformity ($\lambda_{cycle}, \lambda_{SSIM}, \lambda_{Gradient}, \lambda_{uniformity}$) in CycleGAN for medical image encryption rather than using fixed experimental values.

(2) Developed a robust loss-weight combination strategy that balances reconstruction quality and encryption security. This enables the model to trade off visual fidelity and resistance to statistical and adversarial attacks.

Compare image reconstruction quality before and after the proposed optimization algorithm, clarifying how the optimization specifically improves image retrieval performance and enhances resistance to various attack scenarios.

3. PROPOSED METHOD

The methodology followed in designing the proposed UR-CycleGAN and its architectural model, ensuring that every aspect is thoroughly considered, is presented in this section.

3.1 Composite loss function design

The CycleGAN can learn a mutual mapping between original and cipher images, enabling both encryption and decryption. The CycleGAN framework consists of two generators (G and F) and two discriminators (D_x and D_y). Generator G encrypts the image, while Generator F decrypts it. Discriminator D_y is used to discriminate between the target and the cipher images, while discriminator D_x is used to distinguish the source and the reconstructed images. The main objective of generator G is to ensure closeness between the generated ciphertext image and the target domain image as much as possible.

On the other hand, the objective of generator F is the same, but for the generated plaintext and original plaintext images. In this work, LSGAN replaces the BCE with a least-squares loss, which generally stabilizes training and produces less blurriness in generated images. The following equation shows a least-squares loss [10].

$$\mathcal{L}_{adv} = E_{x\sim p_x}\left[\left(D_y(G(X)) - 1\right)^2\right] + E_{y\sim p_y}\left[\left(D_x(F(Y)) - 1\right)^2\right] \quad (1)$$

where, X is the real image in domain X, Y is real image in domain Y, and G refers to the generator that translates image from X domain to Y domain. F denotes the generator that translates image from Y domain to X domain, D_y is the discriminator for domain Y, distinguishes real Y image from $G(X)$, and D_x is the discriminator for domain X which distinguishes real X image from $F(Y)$.

The main objective is to keep the difference between the original and the reconstructed images as small as possible to achieve the decryption process. This loss is called cycle consistency loss. Without cycle consistency, generators could map all inputs to the same output. We can calculate this loss using the following equation [10]:

$$\mathcal{L}_{cycle} = \lambda\, \mathrm{E}_{x\sim p_{data(x)}}\|F(G(X) - X)\|_1 + \lambda\, \mathrm{E}_{y\sim p_{data(y)}}\|G(F(Y) - Y)\|_1 \quad (2)$$

where, $\|.\|_1$ represents L1 paradigm, and λ is a hyperparameter used to control the impact of the cycle consistency loss.

SSIM loss plays a key role in improving the quality of the decrypted image. It serves as a metric for further assessing the perceptual similarity between the original and decrypted images. It captures structural similarity such as luminance, contrast, and texture. By using gradients influenced by local image statistics, SSIM guides the model in a crucial task: preserving spatial relationships. This encourages the model to retain the structural information between the original and reconstructed images. The SSIM loss can be expressed as [12]:

$$\mathcal{L}_{SSIM} = 1 - SSIM(X, F_{Y\to X}(G_{x\to Y}(X))) \quad (3)$$

The gradient loss function is used to ensure consistent edge sharpness. It penalizes structural inconsistency with a strong focus on edge and gradient alignment between real and reconstructed images. The Schar filter is used to compute edge gradients and to maintain edge and sharpness details, thereby enhancing the method's precision.

$$\mathcal{L}_{Gradient} = 0.7\frac{1}{N}\sum_{i,j}\left|M_{real}(i,j) - M_{gen}(i,j)\right| + 0.3\frac{1}{N}\sum_{i,j}\left|\theta_{real}(i,j) - \theta_{gen}(i,j)\right| \quad (4)$$

where, M_{real} refers to the gradient magnitude map of the real image from the Scharr filter, M_{gen} means the gradient magnitude map of the generated image from the Scharr filter, N means the total number of pixels, θ_{real} denotes the gradient orientation map of real image, and θ_{gen} denotes the gradient orientation map of generated image.

The conventional loss function of Cycle-GAN does not clearly apply a structured, uniform intensity distribution. This results in images exhibiting a Gaussian distribution of histogram values, making them susceptible to statistical attacks. To address this vulnerability, a regularization term is proposed to enhance CycleGAN performance. Specifically, it calculates the Pearson's chi-squared (χ^2) divergence between the normalized histogram of cipher images and a uniform target distribution, weighted by λ uniformity, and then incorporates this into the generator loss. This approach significantly improves the security of image encryption. The new loss function is shown below [16].

$$\mathcal{L}_{uniformity} = E_{x\sim p_x}\left[\sum_{k=1}^{K}\frac{\left(O_K(G(X)) - E_K\right)^2}{E_K}\right],\ E_K = \frac{1}{K} \quad (5)$$

where, O_K refers to the observed normalized frequency in the k-th histogram bin of the cipher image $G(x)$, E_K denotes the expected frequency in bin k under the uniform distribution.

By using Pearson's chi-squared (χ^2) divergence, the proposed loss term penalizes deviations from a uniform intensity distribution while the expected distribution is uniform across all histogram bins.

The main idea of this proposed approach is to remove intensity-based patterns from the ciphertext domain that may be exploited by malicious attacks, thereby improving encryption robustness. This regularization is applied especially to the cipher images, decoupling it from the reconstruction pathway. This approach applies adversarial loss, cycle-consistency loss, SSIM, and gradient loss to maintain semantic and structural fidelity in the reconstructed images.

The overall objective for the generator combines adversarial training with multiple regularization terms to balance image realism, reconstruction fidelity, perceptual similarity, edge consistency, and statistical uniformity, as defined in the following loss function:

$$\mathcal{L}_{cycleGAN} = \mathcal{L}_{adv} + \lambda_{cycle}\mathcal{L}_{cycle} + \lambda_{SSIM}\mathcal{L}_{SSIM} + \lambda_{Gradient}\mathcal{L}_{Gradient} + \lambda_{uniformity}\mathcal{L}_{uniformity} \quad (6)$$

where, λ_{cycle}, λ_{SSIM}, $\lambda_{Gradient}$, and $\lambda_{uniformity}$ are hyperparameters that control the influence of each corresponding loss term in the total generator objective. For two discriminators, the overall objective function is to discriminate between real and fake images. The combined discriminator loss is as follows [10, 17]:

$$\mathcal{L}_{D_x} = 0.5\begin{bmatrix} E_{x\sim P_{data(x)}}[(D_x(X)-1)^2] + \\ E_{y\sim P_{data(y)}}[(D_x(F(y))-0)^2] \end{bmatrix} \quad (7)$$

$$\mathcal{L}_{D_y} = 0.5\begin{bmatrix} \mathrm{E}_{y\sim P_{data(y)}}\left[\left(D_y(Y)-1\right)^2\right] + \\ \mathrm{E}_{x\sim P_{data(x)}}\left[\left(D_y(G(X))-0\right)^2\right] \end{bmatrix} \quad (8)$$

$$\mathcal{L}_D = \mathcal{L}_{D_x} + \mathcal{L}_{D_y} \quad (9)$$

where, $\mathcal{L}_D$ is the overall discriminator loss, $\mathcal{L}_{D_x}$ is the discriminator loss for distinguishing real image and generated image in domain *X*, and $\mathcal{L}_{D_y}$ is the discriminator loss for distinguishing real and generated images in domain *Y*.

3.2 Architecture details of CycleGAN

The CycleGAN architecture consists of two generators and two discriminators for bidirectional image translation between two domains, as shown in Figure 1. The generator network consists of an encoder, residual blocks, and a decoder. The generator network specifications are given in Table 2.

The main goal of the encoder is to encode images into a form that captures semantic content and supports firm texture mixing while remaining invertible. It consists of three down-sampling blocks, each involving a 2D convolutional layer (with 4 × 4 kernels, stride 2, and padding 1), instance normalization, and a leaky ReLU activation. The encoder structure increases feature abstraction while decreasing spatial resolution, starting from 64 to 256 filters.

Six residual blocks were composed of two convolutional layers (stride-1), ReLU activation, and instance normalization, connected via skip connections, serving as the core transformation layers. Residual connections alleviate vanishing gradients and retain high-frequency details, such as lesion boundaries in medical imagery.

The decoder is composed of three up-sampling blocks that employ transposed convolutions (4 × 4 kernels, stride 2) with ReLU activation and instance normalization, culminating in a final convolution layer with tanh activation to generate 128 × 128 × 3 output images. The objective of the decoder is to synthesize ciphertext images at the target resolution that are indistinguishable as natural to the discriminator.

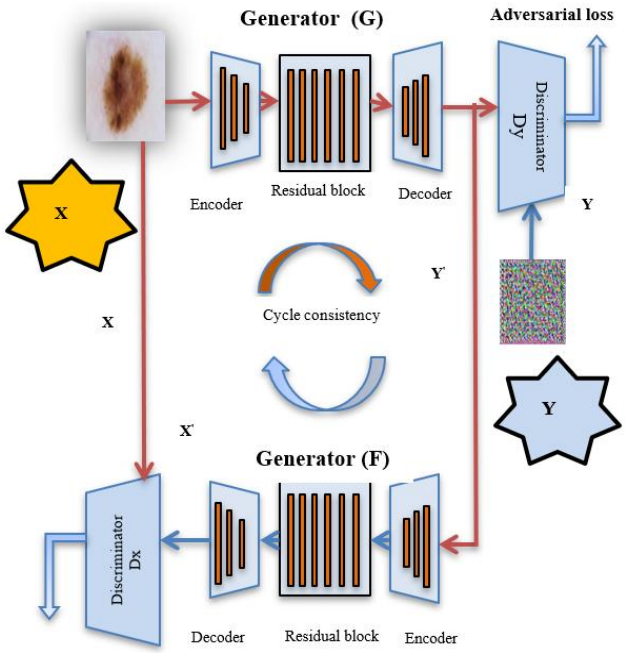

Figure 1. Framework structure of CycleGAN

Table 2. Encryption network structure

Layer Type	Filter	Kernel	Stride	Activation	Output
Encoder (down sample with three blocks)					
Conv 2D	64	4 × 4	2	LeakyReLU	64 × 64 × 64
Conv 2D	128	4 × 4	2	LeakyReLU	32 × 32 × 128
Conv 2D	256	4 × 4	2	LeakyReLU	16 × 16 × 256
Residual Blocks					
Conv2D	256	3 × 3	1	ReLU	16 × 16 × 256
Conv2D	256	3 × 3	1	ReLU	16 × 16 × 256
Decoder (up sample with three blocks)					
Transposed conv2D	128	4 × 4	2	ReLU	32 × 32 × 128
Transposed conv2D	64	4 × 4	2	ReLU	64 × 64 × 64
Transposed conv2D	3	4 × 4	2	tanh	128 × 128 × 3

Table 3. Discriminator structure

Layer	Filter	Kernel	Stride	Activation	Output
Conv 2D	64	4 × 4	2	LeakyReLU	64 × 64 × 64
Conv 2D	128	4 × 4	2	LeakyReLU	32 × 32 × 128
Conv 2D	256	4 × 4	2	LeakyReLU	16 × 16 × 256
Conv 2D	512	4 × 4	1	LeakyReLU	16 × 16 × 512
Output	1	4 × 4	1	–	16 × 16 × 1

Unlike conventional GAN discriminators, the discriminator network in this implementation, as shown in Table 3, follows a multi-scale Patch GAN architecture with four convolutional layers, where each layer, except the last one, employs 4 × 4 convolutions with stride 2, instance normalization, and leaky ReLU, progressively increasing filter depth from 64 to 512. The final layer utilized a stride-1 convolution without normalization.

A discriminator network distinguishes an input image as either a real or a fake image generated by the Generator network. An Adam optimizer with $\beta_1 = 0.5$ and $\beta_2 = 0.999$ is used to optimize both networks, a standard choice for stabilizing GAN training and reducing oscillations. A learning rate of 2×10^{-4} is chosen based on empirical findings in CycleGAN literature, balancing convergence speed and training stability.

In this work, the ISIC Challenge Dataset [18] and the Brain Tumor Detection MRI Dataset [19] have been used.

The datasets contain no personally identifiable information, and all handling complied with HIPAA, GDPR, and official usage protocols. This dataset choice instills confidence in the model's performance.

4. BAYESIAN OPTIMIZATION FRAMEWORK FOR LOSS-WEIGHTS TUNING

In the CycleGAN objective function, selecting weighted parameters λ is a challenging optimization problem. The objective function comprises several competing loss components: cycle consistency, structural similarity, gradient penalty, and uniformity regularization. Each loss contributes differently to training. Adjusting the weight of one term may affect the others. Consequently, determining a suitable combination of λ parameters requires balancing conflicting objectives, especially within a highly non-convex search space. Despite this interdependence, most studies rely on empirically selected loss weights. Researchers often use values from previous works, usually without systematic analysis for each specific application. As a result, there is no understanding of how individual λ parameters affect the CycleGAN objective as a whole or how suboptimal choices may compromise either reconstruction fidelity or security in medical image encryption. This study introduces a systematic framework for medical image encryption and decryption. The method uses a Uniformity-Regularized CycleGAN (UR-CycleGAN) with BO. Bayesian UR-CycleGAN selects the optimal loss-weight configuration for the CycleGAN objective, helping the model adapt to the target medical image dataset while balancing reconstruction quality and encryption security.

By efficiently exploring the parameter space, the BO algorithm identifies loss-weight combinations that allow the UR-CycleGAN model to adapt to the characteristics of the target skin lesion medical image dataset while achieving a balanced trade-off between image reconstruction quality and encryption security.

The BO algorithm is an iterative, model-based optimization method. In this setting, it is applied specifically to tune the loss weights in the UR-CycleGAN objective, where the problem involves optimizing a noisy, expensive, and black-box function.

BO uses its acquisition function to efficiently balance exploration and exploitation, making it well-suited for tuning the loss weights in the UR-CycleGAN objective function. The proposed approach models the objective function using a probabilistic surrogate, a statistical approximation of the true objective. It iteratively selects new loss-weight configurations that maximize an acquisition function, a rule that balances exploration and exploitation, thereby efficiently exploring the high-dimensional search space with several evaluations.

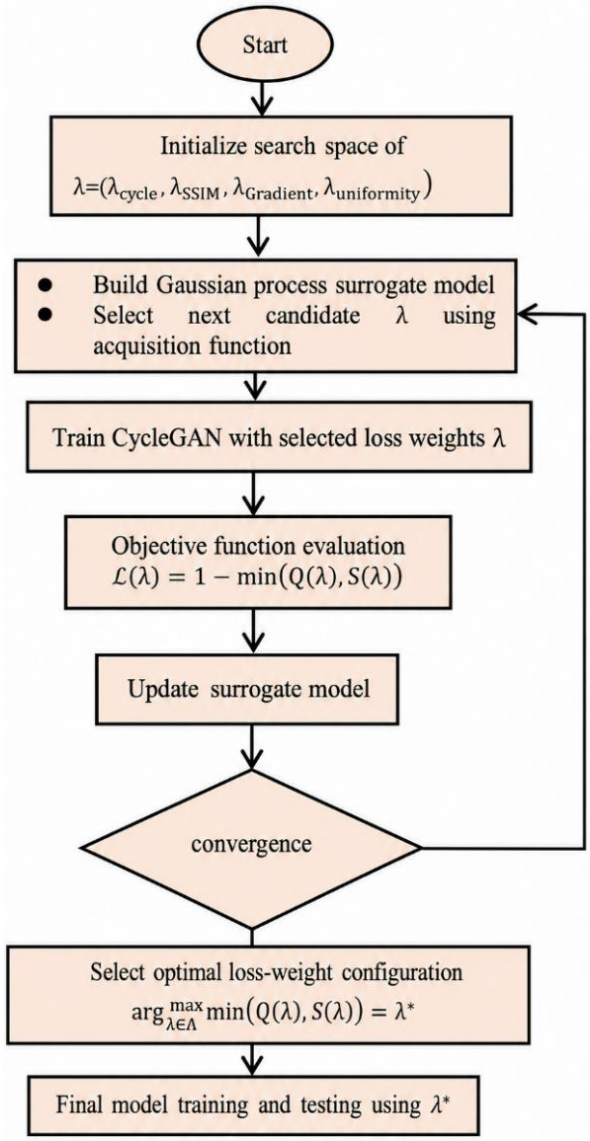

Figure 2. The overall optimization workflow

As shown in Figure 2, the proposed optimization strategy adopts a max–min objective, aiming to optimize for the best possible worst-case outcome. This approach aims to achieve a balanced trade-off between reconstruction quality and encryption security by maximizing the minimum of the two criteria. The optimal set of loss-weight parameters is therefore obtained via a max–min optimization formulation, as expressed in the equation below:

$$\arg\max_{\lambda\subset\Lambda} min\ \big(Q(\lambda), S(\lambda)\big) = \lambda^* \tag{10}$$

where, λ denotes the set of loss-weight parameters, $\lambda = (\lambda_{cycle}, \lambda_{SSIM}, \lambda_{Gradient}, \lambda_{uniformity})$ represents the set of loss weights of cycle consistency, structural similarity index measure, gradient loss, and uniformity regularization, respectively. Q(λ)denotes the reconstruction quality score corresponding to the loss weight parameters, as shown below.

$$Q(\lambda) = \frac{1}{N}\sum_{i=1}^{N}\big(0.7\ SSIM(x,\hat{x}) + 0.3\ PSNR(x,\hat{x})\big) \tag{11}$$

where,

N is the number of medical images.

i = 1, 2, 3, ……. N denotes the index of input images.

x_i refers to the *i-th* original image.

$\hat{x}_i$ denotes the *i-th* reconstructed image using UR-CycleGAN.

$SSIM(x_i, \hat{x}_i)$ is the Structural Similarity Index between the original image x_i and its reconstructed image $\hat{x}_i$.

PSNR$(x_i, \hat{x}_i)$ is the Peak Signal-to-Noise Ratio (PSNR) between the original image x_i and its reconstructed image $\hat{x}_i$.

S(λ) denotes the security score corresponding to the loss weight parameters as shown below:

$$S(\lambda) = \frac{1}{N}\sum_{i=1}^{N}[0.20(1-|Corr_i|) + 0.20\ Ent_i + 0.20(1-NCorr_i) + 0.20\ NPCR_i + 0.20\ UACI_i] \tag{12}$$

where,

$Corr_i$ is the correlation coefficient between the original image and its encrypted image for the *i*-th image,

Ent_i is the entropy of the encrypted image for the *i*-th image.

$NCorr_i$ is the average neighbor-pixel correlation within the encrypted image for the *i*-th image.

$NPCR_i$ is the Number of Pixels Change Rate computed for the i-th image, measuring sensitivity to plaintext changes.

$UACI_i$ is the Unified Average Changing Intensity for the *i*-th image, measuring the average intensity difference caused by plaintext perturbations.

Given the non-convex nature of CycleGAN training and the absence of a closed-form solution for this objective, BO is employed to approximate the optimal solution. To facilitate optimization, the objective is reformulated as a loss function expressed in terms of the model's trainable parameters and measurable error components.

To simplify optimization, the objective can be rewritten as a loss function:

$$\mathcal{L}(\lambda) = 1 - \min\big(Q(\lambda), S(\lambda)\big) \tag{13}$$

The optimization process addresses imbalanced solutions by penalizing parameter configurations that prioritize high-quality performance at the expense of security, and conversely. The specific hyperparameters and their respective ranges used to configure the BO process are presented in Table 4. This configuration is well-suited to the proposed framework.

Table 4. Configuration of Bayesian optimization (BO) algorithm for loss weight

Term	Description
Optimization method	BO
Surrogate model	Gaussian process GP
Acquisition function	Expected improvement plus
Search variable	$\lambda = (\lambda_{cycle}, \lambda_{SSIM}, \lambda_{Gradient}, \lambda_{uniformity})$
Search space range	$\lambda_{cycle}\in[1,120]$, $\lambda_{SSIM}\in[3,20]$, $\lambda_{Gradient}\in[0.3,8]$, $\lambda_{uniformity}\in[1,60]$
Maximum evaluation	60
Final output	Optimal loss weight configuration λ^*

5. EXPERIMENTAL RESULTS

This section shows the empirical evaluation of the proposed Bayesian UR-CycleGAN framework for medical image encryption. The effectiveness of the optimized loss-weight algorithm is assessed in terms of both reconstruction quality and encryption security. Additionally, comparative results against fixed and manually tuned loss weights are used to demonstrate the advantages of the proposed optimization approach. On the other hand, robustness is evaluated under different attack scenarios.

5.1 Analysis of optimal loss weights

This section reports the experimental results for the proposed BO of UR-CycleGAN loss weights. The main goal is to achieve a balanced trade-off between reconstruction quality and encryption security. The results highlight the selected optimal loss-weight configuration and describe how these values directly influence the system's performance metrics for both reconstruction quality and encryption security. Figure 3 illustrates the result of the optimization objective across evaluations. Each blue point represents the objective value obtained for a specific loss-weight configuration. The BO algorithm used 60 evaluation points. The red marker denotes that the optimal solution occurs at evaluation 57, achieving a minimum objective value of 0.22329.

The optimization criterion explicitly targets configurations that minimize the discrepancy between reconstruction quality and encryption security. This design ensures that neither objective is improved at the expense of the other, which is the core goal of the proposed optimization framework. The optimal loss weights are reported in Table 5.

Figure 4 illustrates the convergence of the BO algorithm, showing a rapid decrease in the objective function value from initial evaluations, followed by stabilization in later evaluations. The blue curve represents the minimum observed objective value obtained from the actual evaluations, while the green curve denotes the minimum objective value estimated by the Gaussian Process (GP) surrogate model. Moreover, the mean gap between observed and estimated objective values across 60 evaluations demonstrates the GP model's reliability. Therefore, additional optimization budget may increase computational cost with only marginal performance improvement.

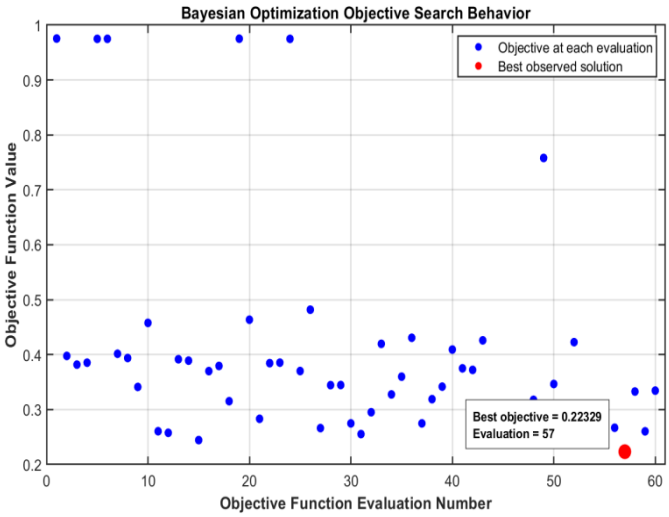

Figure 3. Objective function values obtained during the Bayesian optimization (BO) over 60 evaluations

Table 5. Optimum values of λ parameters using Bayesian Optimization (BO)

Parameters	Manually Setting λ Hyperparameters Values	λ Hyperparameters Values Using BO
λ_{cycle}	10	24.499
λ_{SSIM}	1	6.1743
$\lambda_{Gradient}$	1	0.73384
$\lambda_{uniformity}$	1	1.0004

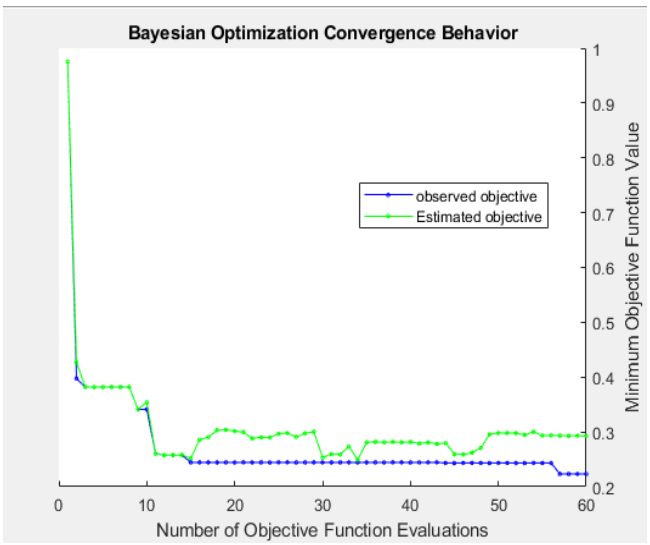

Figure 4. The convergence behavior of the Bayesian optimization (BO) over 60 evaluations

From the experimental results, we conclude that the cycle consistency loss plays an important role in stabilizing the bidirectional mapping and significantly influences image reconstruction quality. Still, it does not constrain the structural details, texture, or perceptual fidelity of the reconstructed image. So this is the main role of the SSIM loss λ_{SSIM}. In practice, SSIM plays a very significant role in reconstruction quality, while the cycle loss contributes only when SSIM is strong enough to provide a meaningful structural prior. Consequently, the impact of λ_{cycle} becomes visible only when λ_{SSIM} is adequately appropriate; otherwise, its contribution remains peripheral. The gradient loss weight $\lambda_{Gradient}$ has only a very minor influence on both reconstruction quality and security.

This is because the model relies on the cycle-consistency loss and the structural SSIM loss, while the gradient term contributes only a small local smoothing effect. A small uniformity loss weight ($\lambda_{uniformity} \approx 1$) demonstrates the best balance between security and reconstruction quality. At this low value, the uniformity term gently improves histogram uniformity without disrupting the natural randomness produced by the adversarial generator, allowing NPCR and entropy to remain high.

As shown in Table 6, BO outperformed Random Search under the same optimization budget of 60 evaluations and within the same search space range. The lower objective value achieved by BO confirms its superior ability to identify the optimal loss-weight hyperparameter configuration.

Table 6. Comparison of random search and Bayesian optimization (BO)

Method	Number of Evaluation	Minimum Objective Value
Random search	60	0.4007
BO	60	0.22329

5.2 Experimental evaluation of the optimized CycleGAN framework

This section evaluates the UR-CycleGAN framework with and without BO.

5.2.1 Quality of recovered image

PSNR serves as a metric to evaluate the quality of an image relative to the original. A higher PSNR value indicates superior image quality. Conversely, a lower PSNR value highlights the difference between the encrypted and original images. PSNR is calculated using the formula provided in Eq. (14) [20].

$$PSNR = 10\log_{10}(2^n - 1)^2 / MSE \tag{14}$$

where, n is the number of bits per pixel. MSE is the mean squared error. The value of $PSNR \in [0, \infty]$ [20].

The SSIM is a technique used to assess the similarity between two images. SSIM incorporates a description of structural information and perceived image quality, including luminance, contrast, and structural comparison. SSIM is calculated using the formula in Eq. (15) [4].

$$SSIM = [l(x,y)]^{\alpha}[c(x,y)]^{\beta}[s(x,y)]^{\gamma} \tag{15}$$

where, α, β, and γ are the positive constants that must be greater than zero. $l(x,y)$ represents the brightness evaluation, $c(x, y)$ indicates the contrast comparison, and $s(x, y)$ denotes the structure comparison [4].

Table 7 quantitatively compares the heuristic and optimized hyperparameters for the loss weight. Optimization increases average PSNR from 28.65 dB to 34 dB, substantially enhancing reconstruction quality. The SSIM metric increases from 0.85 to 0.90, indicating enhanced preservation of structural and perceptual features. The average PSNR and SSIM values are computed over the entire test set of medical images, consisting of 200 images. For each test image, PSNR

and SSIM are computed between the original and reconstructed images.

Such quantitative gains are particularly crucial in medical imaging, where anatomical accuracy and visual consistency are imperative. Importantly, this enhancement in quality is achieved without compromising the encryption performance, as the statistical properties of the encrypted images, such as entropy, remain stable, as well as a uniform histogram distribution. This confirms that the proposed optimization approach effectively improves reconstruction quality while preserving the system's security characteristics. Figure 5 shows examples of the original image, the reconstructed image without optimization, and the reconstructed image obtained with optimum hyperparameter values.

Table 7. Average Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM) values of reconstructed image

Avg. PSNR	Avg. SSIM
Without Optimization	
28.65 dB	0.85
Using Bayesian Optimization (BO) Algorithm	
34 dB	0.90

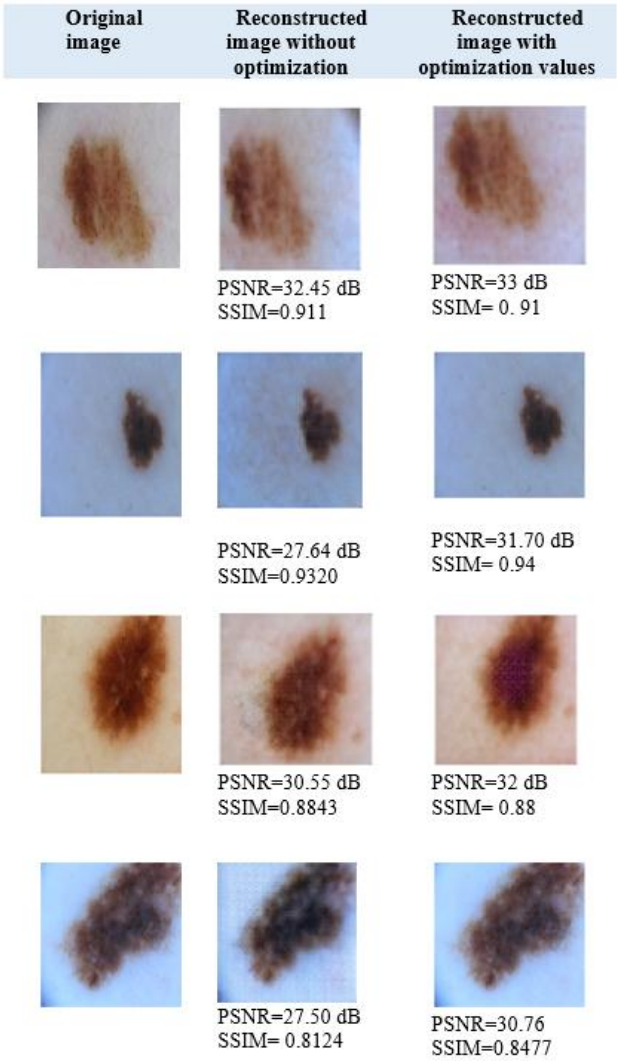

Figure 5. Assessment of the reconstructed image with/without optimization using Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM)

5.2.2 Information entropy

It is a metric of image randomness used to verify the security of the proposed UR-CycleGAN. It demonstrates the distribution values of the encrypted image. It is 8 for an 8-bit image when all values are equally probable. Information entropy is close to 8 for highly random image information. The entropy of the proposed method is calculated using Eq. (16).

$$H(x) = -\sum_{i=1}^{n} p(x_i) \log_2 p(x_i) \tag{16}$$

where, $p(x_i)$ is the pixel probability in the encrypted image [4]. As shown in Table 8, the entropy values of the encrypted image remain relatively stable before and after optimization, indicating that its statistical randomness is preserved while the reconstruction quality is significantly improved.

Table 8. Comparison of average entropy values with/without Bayesian Optimization (BO) and existing image encryption method

Method	Average Entropy
Without BO algorithm	7.836
With BO algorithm	7.74
Ref. [4]	7.21
Ref. [12]	7.36
Ref. [15]	7.38

The optimization process significantly improved reconstruction quality by adjusting the associated loss weights, including the cycle-consistency and SSIM terms. In contrast, the optimum value of the uniformity loss weight remained very close to its manually selected value. Therefore, only marginal changes in entropy were observed.

5.2.3 Correlation coefficient analysis

The correlation, a key factor in the effectiveness of encryption, measures adjacent pixel values in the image. In a plaintext image, adjacent pixels are firmly correlated. This correlation is crucial, as attackers can exploit it to deduce the next gray-scale pixel value. The correlation of the cipher image demonstrates the encryption model's strength in disrupting correlations of the original image. A good encryption scheme results in an encrypted image in which the correlation between adjacent pixels is reduced, thereby providing robustness against statistical attacks. As shown in Table 9, the results indicate that the correlation between adjacent pixels in the cipher images is significantly reduced in all directions (horizontal, vertical, and diagonal) following optimization. This pronounced reduction suggests that the optimized model produces more decorrelated ciphertext, thereby enhancing resistance to statistical and correlation-based attacks. Figure 6 shows the correlation in detail.

Table 9. Correlation values of cipher images with/without Bayesian optimization (BO) and existing image encryption method

Method	Horizontal	Vertical	Diagonal
Without BO algorithm	0.2188	0.3819	0.1016
With BO algorithm	0.0765	0.1273	0.0745
Ref. [4]	0.5920	0.4555	0.4326
Ref. [12]	0.5203	0.4147	0.182

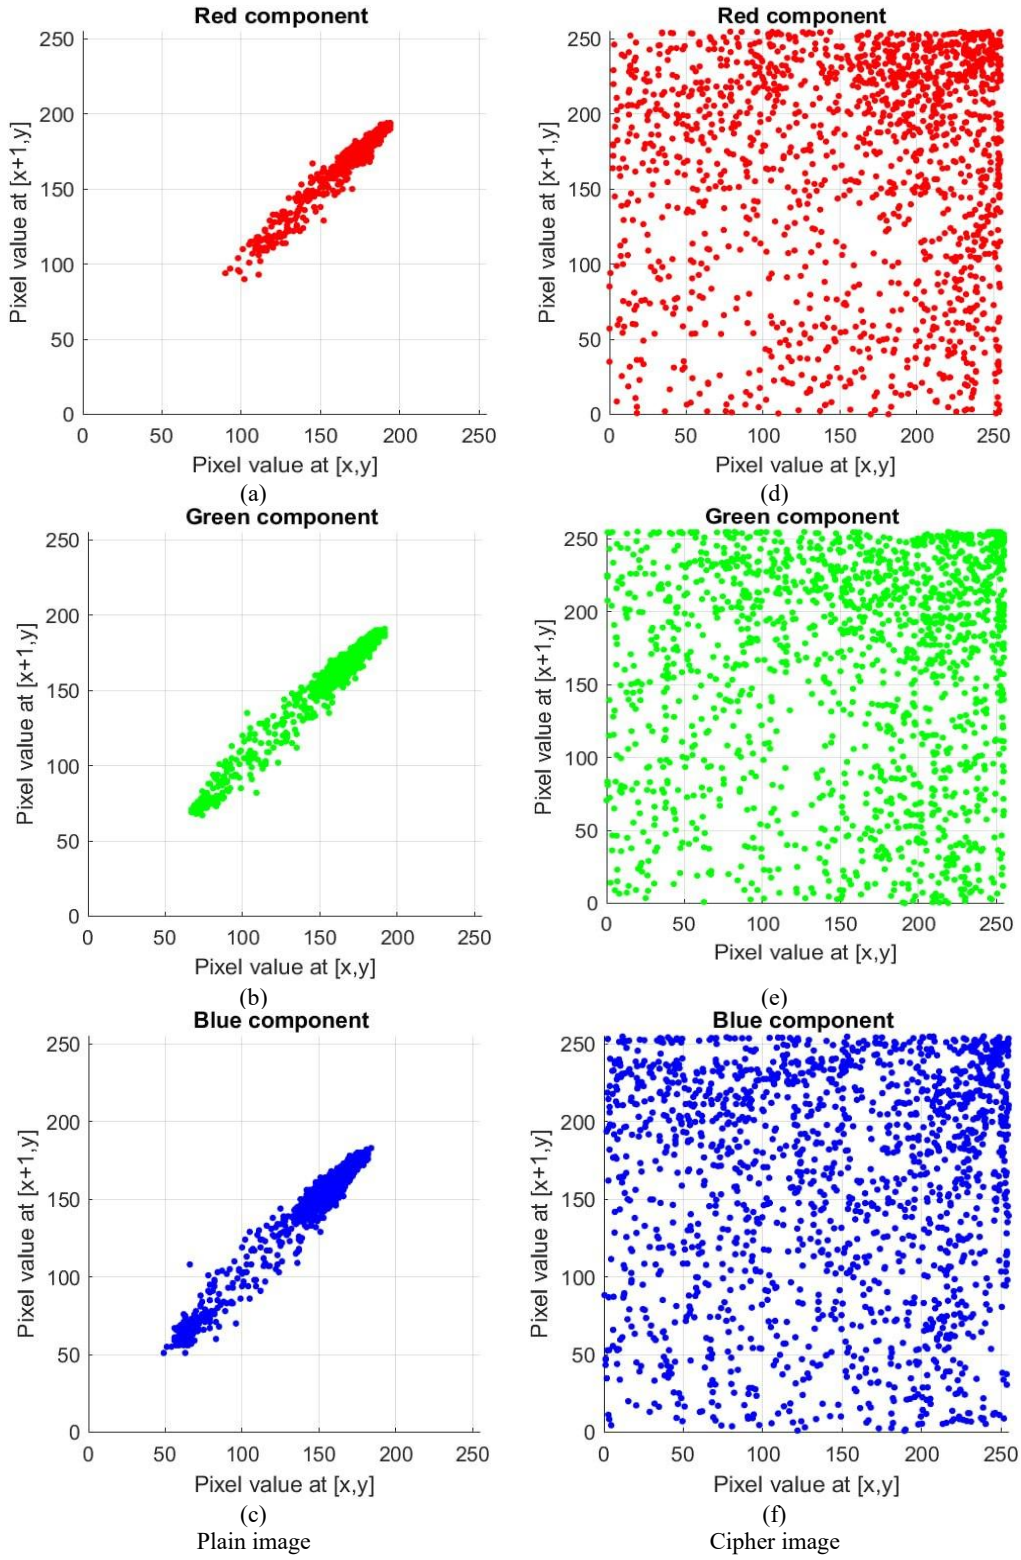

Figure 6. Pixel correlations (horizontal, vertical and diagonal) for original image (a-c) and analogous cipher images (d-f)

5.2.4 Histogram analysis of an encrypted image

In prior work, the histograms of cipher images generated by CycleGAN without uniformity regularization appear to follow a Gaussian distribution, reflecting statistical regularities. As shown in Figure 7, these smooth, bell-shaped histograms reflect structured intensity patterns that may maintain information from the original images. The integration of the proposed uniform histogram loss term into the CycleGAN objective function is a significant aspect. Compared with the uniform histograms of the encrypted outputs from the proposed UR-CycleGAN shown in Figure 8, the influence of the uniformity-regularized loss on eliminating such statistical traces becomes obvious. This emphasizes that the encryption process disrupts pixel value distributions, contributing to enhanced protection against histogram-based analysis.

Table 10 demonstrates the impact of the proposed uniformity regularization loss term on the histogram of cipher images. After incorporating the proposed uniformity loss, the average chi-square value decreases substantially to approximately 599. This reduction confirms that the proposed loss term effectively reduces statistical information that statistical attacks could exploit.

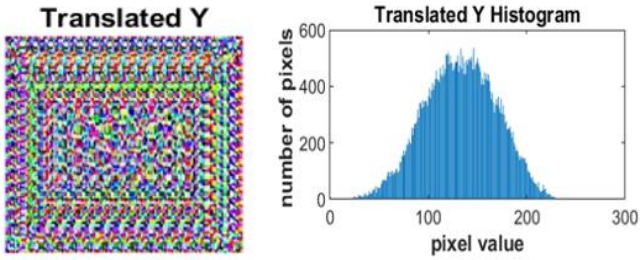

Figure 7. Histogram of encrypted images without uniformity regularization loss function

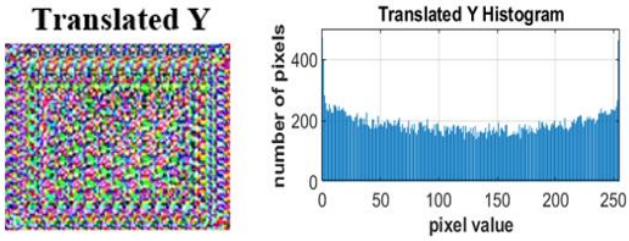

Figure 8. Histogram of encrypted images with uniformity-regularized loss

Table 10. Average Chi-square values of encrypted image

Configuration	Average Chi-square
Without uniformity loss	25070
With uniformity loss	599

5.2.5 Security analysis under differential attacks

Chosen-plaintext attacks are used to test the encryption algorithm's sensitivity to small changes in the plaintext. The attacker tries to make a small change to the plaintext image, then encrypts the original and modified images with the same key, to find the relationship between the plaintext image and the encrypted image of the modified image. If a small change in the plaintext image can cause a huge change in the ciphertext image, this type of attack usually fails to take effect. This means the encryption algorithm is resistant to a chosen-plaintext attack. NPCR is used to measure system resistance against differential attacks. The ideal value of NPCR is 100%. The definition of NPCR can be represented mathematically [21]:

$$NPCR = \frac{\sum_{i=0}^{W}\sum_{j=0}^{H} D(i,j)}{W \times H} \tag{17}$$

where, W denotes the image width, and H represents the image height.

$$D(i,j) = \begin{cases} 1, D_1(i,j) \neq D_2(i,j) \\ 0, D_1(i,j) = D_2(i,j) \end{cases} \tag{18}$$

To evaluate the encryption model's sensitivity, the original medical image and its modified version (with only one pixel changed) are encrypted using the same keys. NPCR is used to compare two encrypted images, and the experimental results are 97.84%, demonstrating a substantial enhancement in diffusion strength compared with NPCR values before using the optimization strategy, which are 92.163%.

This improvement confirms that the optimized model is more sensitive to small changes in plaintext, leading to widespread variations in the ciphertext. Consequently, the optimized system exhibits stronger resistance to chosen-plaintext attacks, as it better satisfies the avalanche property required for secure encryption.

A chosen-ciphertext attack (CCA) is a cryptanalysis model in which the attacker can obtain information by decrypting chosen ciphertexts. Using this information, the attacker can attempt to recover the secret decryption key.

In CCA, the attacker can access the decryption model. In this experiment, the ciphertext image is the input to the decryption model. NPCR is used to measure the difference between two decrypted images. When the input ciphertext image is modified by a single pixel, the NPCR between the two decrypted images is 92.30%. The NPCR value decreases slightly from 94.83% before optimization to 92.30% after optimization. This slight reduction may be attributed to the lower uniformity-loss weight, which decreases excessive statistical randomness in the encrypted domain.

By reducing the influence of the uniformity constraint, the optimized model becomes less sensitive to ciphertext perturbations during decryption, resulting in more controlled error propagation. This behavior suggests improved reconstruction stability under ciphertext perturbation.

In a ciphertext-only attack, the adversary has access only to a string of ciphertexts (encrypted medical images) without any knowledge of the corresponding plaintexts. In this proposed model, the key space, which can be expressed as $(2^{32})^{11,313,160}$, is challenging for attackers to break down the proposed algorithm.

5.2.6 Robustness against adversarial attacks

The proposed encryption framework is evaluated using Projected Gradient Descent (PGD), a strong gradient-based adversarial attack. This attack perturbs a clean example X for T steps with a smaller step size. After each perturbation, PGD projects the adversarial example back onto X's ε-ball if it exceeds it, ensuring it remains within a set distance of the original input [22].

$$x^i = x^{i-1} + \alpha \, sign\left(\nabla_x l\left(h(x^{i-1})\right), y\right) \tag{19}$$

where, α is the step size, and x^i is the adversarial example at the i-th step ($x^0 = x$). The step size is usually set to $\varepsilon/T \leq \alpha < \varepsilon$ for overall T steps of perturbations. $\nabla_x l$ denotes the gradient of the loss function with respect to the input image x. $h(x)$ denotes the output of the model given input image x. y is the ground-truth class label.

We evaluate the stability of the encryption/decryption process under adversarial conditions and analyze the impact of loss-weight optimization on system resilience.

Table 11. Comparison of the baseline and optimized configurations

Evaluation Results Without Optimization	Evaluation Results Using Proposed Bayesian Optimization (BO)
PSNR drop = 1.64	PSNR drop = 1.3056
SSIM drop = 0.094	SSIM drop = 0.028
MSE increase = 0.00082	MSE increase = 0.0006
ASR by SSIM = 22.5%	ASR by SSIM = 2.5%
ASR by PSNR = 0	ASR by PSNR = 0

Note: Structural Similarity Index = SSIM; Peak Signal-to-Noise Ratio = PSNR; Attack Success Rate = ASR

Table 11 shows that BO increases system robustness to

PGD attacks. It reduces the proportion of images that fail reconstruction under adversarial perturbations. The observed attack success rate indicates that adversarial degradation primarily causes structural distortion rather than pixel-level errors. This means that perceptual integrity is more affected than raw intensity values. After optimization, the number of images with structural reconstruction failure drops significantly. This confirms that BO enhances resistance to adversarial perturbations while preserving encryption robustness and reconstruction quality.

6. EVALUATION ON BRAIN MAGNETIC RESONANCE IMAGING DATASET

To further evaluate the applicability of the proposed framework beyond the ISIC dataset, a brain MRI dataset was employed. A total of 500 MRI images were used for training, and 80 for testing. To ensure a comprehensive assessment, PSNR and SSIM metrics are used to evaluate the reconstructed image, whereas entropy is used to assess the encrypted image, as illustrated in Table 12.

As shown in Figure 9, the proposed framework maintains reconstruction fidelity and encryption security on the MRI brain image.

Table 12. Evaluation results on brain Magnetic Resonance Imaging (MRI) dataset

Dataset	Peak Signal-to-Noise Ratio (PSNR) (dB)	Structural Similarity Index (SSIM)	Entropy
MRI brain dataset	27.5	0.86	7.85

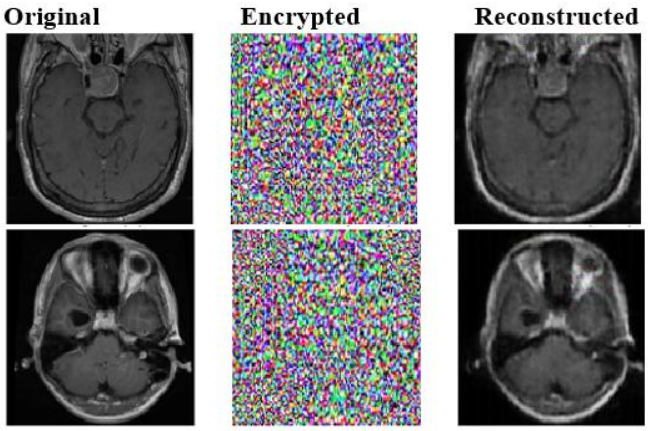

Figure 9. Original, encrypted, and reconstructed brain Magnetic Resonance Imaging (MRI) images

7. STABILITY ANALYSIS AND COMPUTATIONAL COMPLEXITY ANALYSIS

To assess the stability of the proposed method under the stochastic nature of CycleGAN training, the model was trained in two independent runs within identical experimental settings.

The mean and standard deviation values of PSNR, SSIM, and entropy are summarized in Table 13. The results from the two independent runs suggest a limited performance variability, providing preliminary evidence of the robustness of the proposed framework.

All experiments were conducted in MATLAB using an 11th Gen Intel Core i7-1165G7 CPU, 16 GB RAM, and an NVIDIA GeForce MX450 GPU (2 GB). The CycleGAN training took approximately 58 minutes over 700 epochs, while BO took approximately 65 hours and was performed only once during the offline hyperparameter search stage.

Table 13. Stability analysis

Metrics	Mean $\pm$ Std
Peak Signal-to-Noise Ratio (PSNR)	30.98 ± 1.77
Structural Similarity Index (SSIM)	0.8934 ± 0.0242
Entropy	7.8876 ± 0.0204

8. CONCLUSION

This work proposes a loss-weight optimization framework for medical image encryption based on a Bayesian-optimized Uniformity-Regularized CycleGAN (Bayesian UR-CycleGAN), where CycleGAN is a type of GAN designed for unpaired image-to-image translation, and BO is a probabilistic model-based method for optimizing complex functions. This optimization process is formulated as a Max–Min Optimization, an approach that seeks the optimal loss-weight hyperparameters for the CycleGAN objective function by maximizing the minimum performance across multiple competing criteria, rather than maximizing only the best-case performance. This optimization strategy systematically prioritizes the minimum component of the objective function, promoting balanced optimization performance. The proposed framework is evaluated on medical imaging datasets using widely adopted image quality metrics, including PSNR and SSIM.

The experimental findings demonstrate that the optimized hyperparameter configuration for loss weights significantly improves reconstruction quality. The PSNR increased to approximately 34 dB, while the SSIM reached around 0.9. These improvements highlight the influence of the Cycle-consistency and SSIM loss weights on the reconstruction fidelity. Using BO, the cycle-consistency and the SSIM loss weight were adjusted to 24.499 and 6.1743, respectively. In comparison, when the cycle consistency and the SSIM loss were manually set to 10 and 1, respectively, the average PSNR and SSIM were 28.65 dB and 0.85, respectively.

Additionally, entropy and correlation metrics were used to assess the security characteristics of the optimized framework. The proposed optimization framework maintains an acceptable level of security. The framework exhibited resilience under the evaluated attack scenarios. Overall, the experimental results indicate that the proposed Bayesian UR-CycleGAN framework provides a balanced trade-off between reconstruction quality and encryption security under the considered dataset and experimental conditions. These experimental findings motivate further evaluation of more diverse datasets and suggest their potential applicability to medical image encryption in a broader experimental setting.

The proposed framework applies to the secure transmission and storage of medical images in telemedicine and cloud-based healthcare environments. It supports both patient privacy and secure diagnostic usability of recovered images.

On the other hand, the main limitation of this work is that it is evaluated on a limited medical image dataset. Furthermore, this work focused primarily on image-level security analysis

and did not investigate threats such as model extraction, replay attacks, key leakage, or data tampering.

Future work will consider these threat models and evaluate the framework on additional medical imaging modalities to further assess its robustness and generalization capabilities.

REFERENCES

[1] Chai, X., Zhang, J., Gan, Z., Zhang, Y. (2019). Medical image encryption algorithm based on Latin square and memristive chaotic system. Multimedia Tools and Applications, 78: 35419-35453. https://doi.org/10.1007/s11042-019-08168-x

[2] Prabhu, P., Manjunath, K.N. (2019). Secured Image Transmission in Medical Imaging Applications—A survey. Computer Aided Intervention and Diagnostics in Clinical and Medical Images, 125-133. https://doi.org/10.1007/978-3-030-04061-1_12

[3] Belazi, A., Talha, M., Kharbech, S., Xiang, W. (2019). Novel medical image encryption scheme based on chaos and DNA encoding. IEEE Access, 7: 36667-36688. https://doi.org/10.1109/ACCESS.2019.2906292

[4] Inam, S., Kanwal, S., Anwar, A., Mirza, N., Alfraihi, H. (2024). Security of end-to-end medical images encryption system using trained deep learning encryption and decryption network. Egyptian Informatics Journal, 25(1): 15-26. https://doi.org/10.1016/j.eij.2024.100541

[5] Chang, K., Chen, Y., Hsieh, C., Huang, C., Chang, C. (2009). Embedded a low area 32-bit AES for image encryption / decryption application. In 2009 IEEE International Symposium on Circuits and Systems, Taipei, China, pp. 1922-1925. https://doi.org/10.1109/ISCAS.2009.5118160

[6] Preishuber, M., Hütter, T., Katzenbeisser, S., Uhl, A. (2018). Depreciating motivation and empirical security analysis of chaos-based image and video encryption. IEEE Transactions on Information Forensics and Security, 13: 2137-2150. https://doi.org/10.1109/TIFS.2018.2812080

[7] Bhat, R., Nanjundegowda, R. (2024). Comparative analysis of CryptoGAN: Evaluating quality metrics and security in GAN-based image encryption. Journal of Robotics and Control (JRC), 5(5): 1557-1569. https://doi.org/10.18196/jrc.v5i5.23096

[8] Saad, M.M., O'Reilly, R., Rehmani, M.H. (2024). A survey on training challenges in generative adversarial networks for biomedical image analysis. Artificial Intelligence Review, 57(2): 62. https://doi.org/10.1007/s10462-023-10624-y

[9] Tong, Z. (2024). Exploring the impact of hyperparameters on the generation quality of CycleGAN. Transactions on Computer Science and Intelligent Systems Research, 5: 265-271. https://doi.org/10.62051/01m93a63

[10] Zhu, J.Y., Park, T., Isola, P., Efros, A.A. (2017). Unpaired image-to-image translation using cycle-consistent adversarial networks. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), Venice, Italy, pp. 2242-2251. https://doi.org/10.1109/ICCV.2017.244

[11] Chai, X., Wang, Y., Chen, X., Gan, Z., Zhang, Y. (2022). TPE-GAN: Thumbnail preserving encryption based on GAN with key. IEEE Signal Processing Letters, 29: 972-976. https://ieeexplore.ieee.org/abstract/document/9745296.

[12] Panwar, K., Singh, A., Kukreja, S., Singh, K.K., Shakhovska, N., Boichuk, A. (2023). Encipher GAN: An end-to-end color image encryption system using a deep generative model. Systems, 11(8): 373. https://doi.org/10.3390/systems11010036

[13] Liu, Z., Xue, R. (2024). Visual image encryption based on compressed sensing and Cycle-GAN. The Visual Computer, 40: 5857-5870. https://doi.org/10.1007/s00371-023-03140-1

[14] Nan, Y., Wo, Y. (2025). RIE-GAN: A retrievable image encryption method based on GAN. Digital Signal Processing, 163: 104384. https://doi.org/10.1016/j.dsp.2025.105202

[15] Inam, S., Kanwal, S., Hajjej, F., Alluhaidan, A.S. (2026). Secure breast cancer imaging: A novel advanced generative model encryption approach. Engineering Applications of Artificial Intelligence, 163: 112747. https://doi.org/10.1016/j.engappai.2025.112747

[16] Banu, S.A., Amirtharajan, R. (2020). A robust medical image encryption in dual domain: Chaos DNA-IWT combined approach. Medical & Biological Engineering & Computing, 58(7): 1445-1458. https://doi.org/10.1007/s11517-020-02178-w

[17] Mao, X., Li, Q., Xie, H., Lau, R.Y., Wang, Z., Smolley, S.P. (2016). Least squares generative adversarial networks. In 2017 IEEE International Conference on Computer Vision (ICCV), Venice, Italy, pp. 2813-2821. https://doi.org/10.1109/ICCV.2017.304

[18] Gutman, D., Codella, N.C.F., Celebi, M.E., Helba, B., Marchetti, M., Mishra, N., Halpern, A. (2016). Skin lesion analysis toward melanoma detection: A challenge at the International Symposium on Biomedical Imaging (ISBI) 2016, hosted by the International Skin Imaging Collaboration (ISIC). arXiv preprint arXiv: 1605.01397. https://arxiv.org/abs/1605.01397.

[19] Kaggle. Brain Tumor Detection MRI Dataset. https://www.kaggle.com/abhranta/brain-tumor-detection-mri.

[20] Kaur, M., Kumar, V. (2020). A comprehensive review on image encryption techniques. Archives of Computational Methods in Engineering, 27: 15-43. https://doi.org/10.1007/s11831-018-9298-8

[21] Ding, Y., Wu, G., Chen, D., Zhang, N., Gong, L., Cao, M., Qin, Z. (2021). DeepEDN: A deep-learning-based image encryption and decryption network for internet of medical things. IEEE Internet of Things Journal, 8(13): 10562-10575. https://doi.org/10.1109/JIOT.2020.3012452

[22] Mohandas, S., Manwani, N., Dhulipudi, D. (2022). Momentum iterative gradient sign method outperforms PGD attacks. In Proceedings of the 14th International Conference on Agents and Artificial Intelligence (ICAART), 3: 913-916. https://doi.org/10.5220/0010938400003116