

Adaptive RestoreGAN for Histopathology Image Enhancement

Anusha B¹, Harini G¹, Akshaya Keerthi B¹, Aishwarya G¹

¹ Student, Dept. of CSE (AI&ML), GVP College of Engineering for Women, Visakhapatnam, AP, India

Abstract:

Digital histopathology images play a critical role in medical diagnosis, particularly in cancer detection. However, these images are often degraded by optical blur, motion artifacts, and low-resolution acquisition, which significantly impair the performance of automated diagnostic systems and deep learning models. This study proposes Adaptive Restore GAN, a deep learning-based image restoration framework designed to enhance degraded histopathology images using Generative Adversarial Networks (GANs). The framework employs a blur-detection mechanism based on the Variance of Laplacian (VOL) to identify degraded images. Only images with blur scores below a predefined threshold are processed, enabling selective enhancement and avoiding unnecessary computation.

The restoration module utilizes GAN-based architectures, including a ResNet-based generator and a PixelGAN discriminator, trained on two benchmark datasets: Patch Camelyon (PCam) and the NCT-CRC-HE-100K colorectal tissue dataset. The model is optimized using a composite loss function integrating adversarial, L1 reconstruction, Structural Similarity (SSIM), and perceptual losses to ensure both structural fidelity and visual quality. Performance evaluation is conducted using Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM), demonstrating effective recovery of fine histological details. The system is implemented as a deployable full-stack solution using PyTorch for model development, FastAPI for backend inference, and React for an interactive web interface.

The key contributions of this work include:(i) an adaptive blur-detection and restoration pipeline, (ii) a GAN architecture tailored for histopathology image enhancement,(iii) a multi-objective training strategy for improved perceptual quality, and (iv) a practical, deployable clinical support system. It is important to note that degradation in this study is synthetically generated; therefore, performance on real-world clinical images with complex degradation patterns may vary. Future work will focus on benchmarking against state-of-the-art methods such as ESRGAN, SwinIR, and MIRNet.

Keywords: Adaptive RestoreGAN, Histopathology Image Enhancement, Deep Learning, PSNR, SSIM, Blur Detection, Laplacian Variance, Medical Image Processing.

I. Introduction

Digital pathology has become a crucial tool for diagnosing oncological diseases. Histopathology images — microscopic representations of tissue samples — are vital for analyzing cellular architecture, supporting medical research, and enabling computer-aided diagnosis via deep learning.

However, histopathology images frequently suffer degradation during acquisition: poor microscope focusing, motion artifacts, low-quality imaging equipment, and compression artefacts all reduce image quality. Such degradations prevent clinicians from observing fine structural details and degrade the performance of automated deep learning pipelines.

Generative Adversarial Networks (GANs) have demonstrated strong capability in high-quality image synthesis and restoration. The proposed **Adaptive RestoreGAN** system uses a GAN framework to enhance degraded histopathology images. The pipeline first assesses sharpness via the Variance of Laplacian (VOL) method. Images whose blur score falls below a set threshold are classified as blurred and passed to the GAN generator for restoration; images above the threshold are already sufficiently sharp and are returned unchanged. Quality is evaluated using PSNR and SSIM.

While the system is presently validated only on synthetically degraded images, it demonstrates a scalable and clinically motivated approach to histopathology image restoration.

II. Related Work

Medical image processing plays a vital role in disease diagnosis and clinical decision support. Histopathology images in particular are susceptible to noise, blur, staining variation, and imaging artefacts, all of which complicate accurate analysis.

Early methods relied on traditional filtering and contrast enhancement. Although computationally lightweight, these approaches struggle to preserve fine structural detail under complex distortions. Deep learning — and GANs in particular — have substantially advanced image restoration performance.

Rong et al. [1] proposed RestoreGAN to address blur, low resolution, and staining variation in pathology images, demonstrating significant quality improvements and better tumor detection. Zehra et al. [2] presented a cloud-based workflow combining frame stitching, deblurring, super-resolution, and illumination correction for low-cost whole-slide image digitization. Wang et al. [3] developed a GAN with multi-scale attention and style perception, achieving PSNR of 52.84 dB and SSIM of 0.968.

Cruz et al. [4] reviewed GAN-based histological segmentation comprehensively, confirming GAN effectiveness for noise reduction and quality enhancement. Bhavana et al. [5] incorporated ESRGAN for resolution and texture recovery. Touhami et al. [6] benchmarked SwinIR, ESRGAN, MIRNet-v2, and SRCNN for histopathology enhancement, finding SwinIR highest in PSNR and ESRGAN superior in perceptual quality — establishing the key SOTA baseline against which future versions of this work should be evaluated.

Madabhushi and Lee [7] identified variability in image quality as a core challenge in digital pathology. Chen et al. [8] applied CNN and U-Net architectures for MRI artefact correction, demonstrating broad applicability. Li et al. [9] improved feature extraction and classification accuracy with an enhanced CNN for histopathology. Gupta et al. [10] showed GAN-based augmentation improves segmentation accuracy. Amerikanos and Maglogiannis [11] used object detection and segmentation for whole-slide image analysis.

The present work builds directly on RestoreGAN [1] and the ESRGAN family [5], while introducing an adaptive threshold mechanism that makes enhancement conditional on measured image sharpness — a design choice not present in prior methods.

III. Methodology

A. Convolutional Neural Network (CNN)

CNNs are feed-forward networks that exploit local spatial structure through learnable convolutional filters. They are foundational to modern image classification, object detection, and image restoration pipelines. The three principal layer types are:

- (i) Convolutional Layer — detects local features via learned filters.
- (ii) Pooling Layer — reduces spatial dimensionality while preserving dominant activations.
- (iii) Fully Connected Layer — combines learned features for final prediction.

Convolutional Layer

Feature map output at spatial position (i, j) for filter channel k :

$$A_{i,j}^k = \sigma\left(\sum_p \sum_q W_{pq}^k \cdot I_{(i+p)(j+q)} + B^k\right) \quad (1)$$

where W_{pq}^k are filter weights, I is the input, B^k is the bias, and σ is the activation function.

Pooling Layer (Max Pooling)

$$P_{i,j} = \max_{(r,s) \in \mathcal{R}_{ij}} I_{rs} \quad (2)$$

where \mathcal{R}_{ij} denotes the pooling window centred at (i, j) .

Fully Connected Layer

$$Z_n = \sigma\left(\sum_m W_{nm} \cdot F_m + B_n\right) \quad (3)$$

where W_{nm} is the weight from input unit m to output unit n , and F_m are the flattened feature maps.

Activation functions such as ReLU introduce non-linearity, enabling the network to learn complex transformations. The architecture of CNN is shown in Figure 1.

Figure 1: Architecture of a Convolutional Neural Network (CNN)

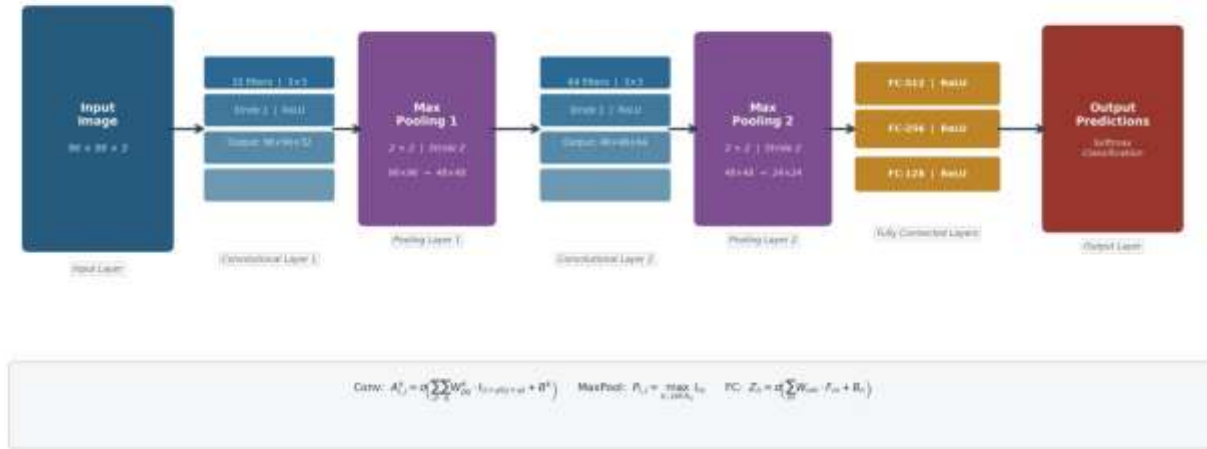


Fig1: Architecture of a Convolutional Neural Network(CNN)

Figure 1 illustrates the layer-by-layer structure of a CNN as used within the Adaptive RestoreGAN framework. The network receives a $96 \times 96 \times 3$ histopathology image as input and passes it through two successive convolutional blocks, each followed by a max-pooling layer. The first convolutional layer applies 32 learnable 3×3 filters to extract low-level features such as edges and colour gradients, producing a $96 \times 96 \times 32$ feature map. The subsequent max-pooling layer reduces the spatial dimensions by half to 48×48 , discarding redundant spatial information while retaining dominant activations. The second convolutional block deepens feature extraction with 64 filters, producing richer, more abstract representations at $48 \times 48 \times 64$, before a second pooling step further reduces the map to 24×24 . The flattened feature maps are then passed through three fully connected layers of decreasing size ($512 \rightarrow 256 \rightarrow 128$ units), each with ReLU activation, which combine the spatial features into a compact representation. A final softmax layer produces the output class probabilities. This hierarchical design — progressively extracting, compressing, and combining features — forms the backbone of the convolutional operations employed in both the generator and discriminator of the Adaptive RestoreGAN system.

B. Generative Adversarial Network (GAN)

GANs consist of two competing networks — a **Generator** G and a **Discriminator** D — trained jointly in an adversarial framework. The generator learns to produce realistic outputs from degraded inputs, while the discriminator learns to distinguish real from generated images, providing a training signal to the generator.

1. Generator Based on ResNet

The generator is a Residual Network (ResNet). Residual blocks allow the network to learn identity mappings plus a residual correction, mitigating the vanishing gradient problem and enabling stable training of deep architectures.

The output of a residual block is:

$$\hat{y} = F(x) + x \quad (4)$$

where x is the block input and $F(x)$ is the residual function learned by the block's convolutional layers.

The generator convolution operation is:

$$y(i, j) = \sum_p \sum_q w(p, q) x(i - p, j - q) + b \quad (5)$$

where $y(i, j)$ is the output feature map at position (i, j) , $w(p, q)$ are filter coefficients, $x(i - p, j - q)$ is the input image patch, and b is the bias.

The generator executes the following functions:

- (i) Extract features from degraded (blurred) input images.
- (ii) Restore missing high-frequency structural detail.
- (iii) Enhance edges and tissue textures.
- (iv) Produce visually and metrically improved output images.

2. Discriminator Based on PixelGAN

The discriminator is a PixelGAN — a patch-level discriminator that evaluates realism at the pixel/patch level rather than globally. This design better captures fine structural detail and encourages the generator to produce locally sharp, realistic images.

The discriminator output (probability of being real) is:

$$D(x) = \sigma(Wx + b) \quad (6)$$

where σ is the sigmoid activation and W, b are learned parameters.

3. Discriminator Loss

The discriminator is trained to minimise the following binary cross-entropy loss:

$$L_D = \mathbb{E}_{x \sim p_{\text{data}}} [\log D(x)] + \mathbb{E}_{z \sim p_z} [\log (1 - D(G(z)))] \quad (7)$$

where $D(x)$ is the probability that a real image is real, $G(z)$ is the generated image, and $D(G(z))$ is the probability that the generated image is judged real.

4. GAN Objective

$$\min_G \max_D V(D, G) = \mathbb{E}_{x \sim p_{\text{data}}} [\log D(x)] + \mathbb{E}_{z \sim p_z} [\log (1 - D(G(z)))] \quad (8)$$

The architecture of GAN is given in Figure 2.

Figure 2: Architecture of the Generative Adversarial Network (GAN)

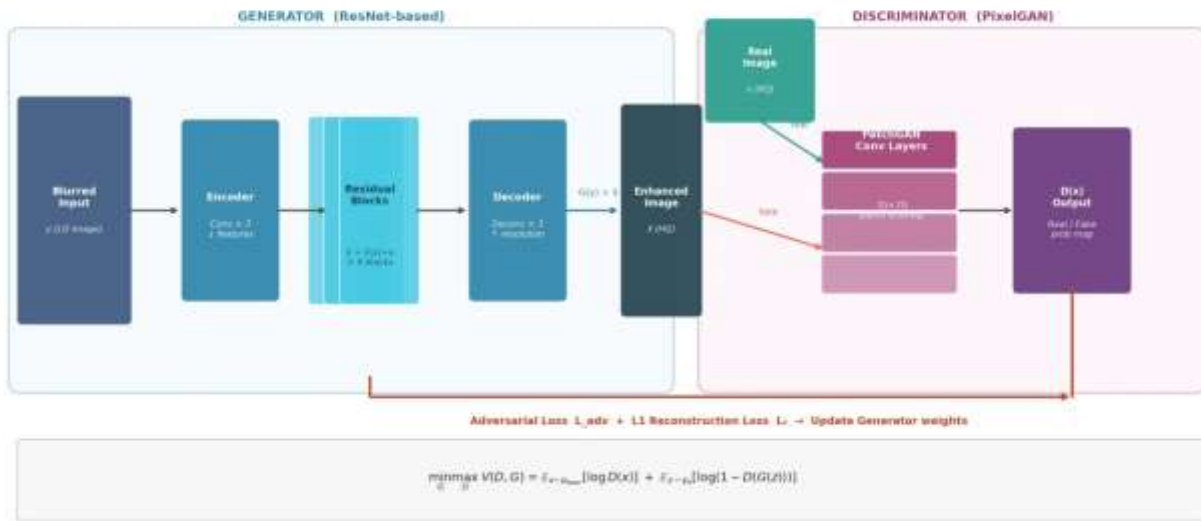


Figure 2: Architecture of the Generative Adversarial Network (GAN)

Figure 2 illustrates the adversarial structure of the GAN used in Adaptive RestoreGAN. The framework consists of two networks trained in opposition. The **Generator** (left) receives a degraded low-quality histopathology image as input and processes it through an encoder, a stack of nine residual blocks, and a decoder to produce an enhanced output $\hat{x} = G(y)$. Residual connections within each block allow the network to learn correction terms on top of the input identity, which preserves important tissue structure while recovering lost high-frequency detail.

The **PixelGAN Discriminator** (right) receives either a real high-quality image or the generator's output and scores each spatial patch as real or fake, producing a probability map $D(x) = \sigma(Wx + b)$. This patch-level evaluation forces the generator to maintain local sharpness and fine structural realism across the entire image rather than optimising global statistics alone.

During training, the generator attempts to minimise the combined loss \mathcal{L}_G while the discriminator attempts to maximise its ability to distinguish real from generated images. This adversarial interplay, governed by the minimax objective, drives both networks to improve iteratively until the generator produces outputs that are indistinguishable from real high-quality histopathology images.

5. Loss Functions

Binary Cross-Entropy Loss (adversarial):

$$L_{\text{BCE}} = -[y \log(p) + (1 - y) \log(1 - p)] \quad (9)$$

L1 Reconstruction Loss:

$$\mathcal{L}_1 = \mathbb{E}_{(x,y)}[\|x - G(y)\|_1] \quad (10)$$

SSIM Loss:

$$\mathcal{L}_{\text{SSIM}} = 1 - \text{SSIM}(x, \hat{x}) \quad (11)$$

Perceptual Loss (VGG feature matching):

$$\mathcal{L}_{\text{perc}} = \sum_l \frac{1}{N_l} \|\phi_l(x) - \phi_l(\hat{x})\|_2^2 \quad (12)$$

where ϕ_l denotes the feature map at layer l of a pre-trained VGG network.

Total Generator Loss:

$$\mathcal{L}_G = \lambda_{\text{adv}} \mathcal{L}_{\text{adv}} + \lambda_1 \mathcal{L}_1 + \lambda_{\text{SSIM}} \mathcal{L}_{\text{SSIM}} + \lambda_{\text{perc}} \mathcal{L}_{\text{perc}} \quad (13)$$

where $\lambda_{\text{adv}}, \lambda_1, \lambda_{\text{SSIM}}, \lambda_{\text{perc}}$ are scalar weighting factors.

6. Full Mathematical Formulation

(1) Degradation Process:

$$y = D(x) = x * k + n \quad (14)$$

where x is the clean image, k is the blur kernel, n is additive noise, and y is the degraded observation.

(2) Restoration Mapping:

$$G: y \rightarrow \hat{x} \quad (15)$$

(3) Discriminator Loss (least-squares formulation):

$$\mathcal{L}_D = \frac{1}{2} \mathbb{E}_{x \sim p_{\text{data}}} [(D(x) - 1)^2] + \frac{1}{2} \mathbb{E}_{y \sim p_{\text{data}}} [(D(G(y)))^2] \quad (16)$$

(4) Adversarial Loss:

$$\mathcal{L}_{\text{adv}} = \frac{1}{2} \mathbb{E}_{y \sim p_{\text{data}}} [(D(G(y)) - 1)^2] \quad (17)$$

(5) Reconstruction Loss:

$$\mathcal{L}_1 = \mathbb{E}_{(x,y)} [\|x - G(y)\|_1] \quad (18)$$

(6) SSIM Loss:

$$\mathcal{L}_{\text{SSIM}} = 1 - \text{SSIM}(x, \hat{x}) \quad (19)$$

(7) Perceptual Loss:

$$\mathcal{L}_{\text{perc}} = \sum_l \frac{1}{N_l} \|\phi_l(x) - \phi_l(\hat{x})\|_2^2 \quad (20)$$

(8) Total Generator Loss:

$$\mathcal{L}_G = \lambda_{\text{adv}} \mathcal{L}_{\text{adv}} + \lambda_1 \mathcal{L}_1 + \lambda_{\text{SSIM}} \mathcal{L}_{\text{SSIM}} + \lambda_{\text{perc}} \mathcal{L}_{\text{perc}} \quad (21)$$

(9) Residual Learning:

$$\hat{x} = y + R(y) \tag{22}$$

where $R(y)$ is the residual correction predicted by the generator.

(10) Adaptive Threshold Decision:

$$x_{out} = \begin{cases} G(y), & s(y) < \tau \\ y, & s(y) \geq \tau \end{cases} \tag{23}$$

where $s(y)$ is the blur score (Variance of Laplacian) and τ is the threshold. A **low** blur score indicates a **blurred** image; a **high** score indicates a sharp image.

(11) Soft Gating Function:

$$\alpha(y) = \sigma(\beta(\tau - s(y))) \tag{24}$$

(12) Final Output (soft-gated):

$$x_{out} = \alpha(y) G(y) + (1 - \alpha(y)) y \tag{25}$$

(13) Mean Squared Error:

$$MSE = \frac{1}{HWC} \sum_{i=1}^H \sum_{j=1}^W \sum_{c=1}^C (x_{ijc} - \hat{x}_{ijc})^2 \tag{26}$$

The architecture of Adaptive Restore GAN is given in Figure 3.

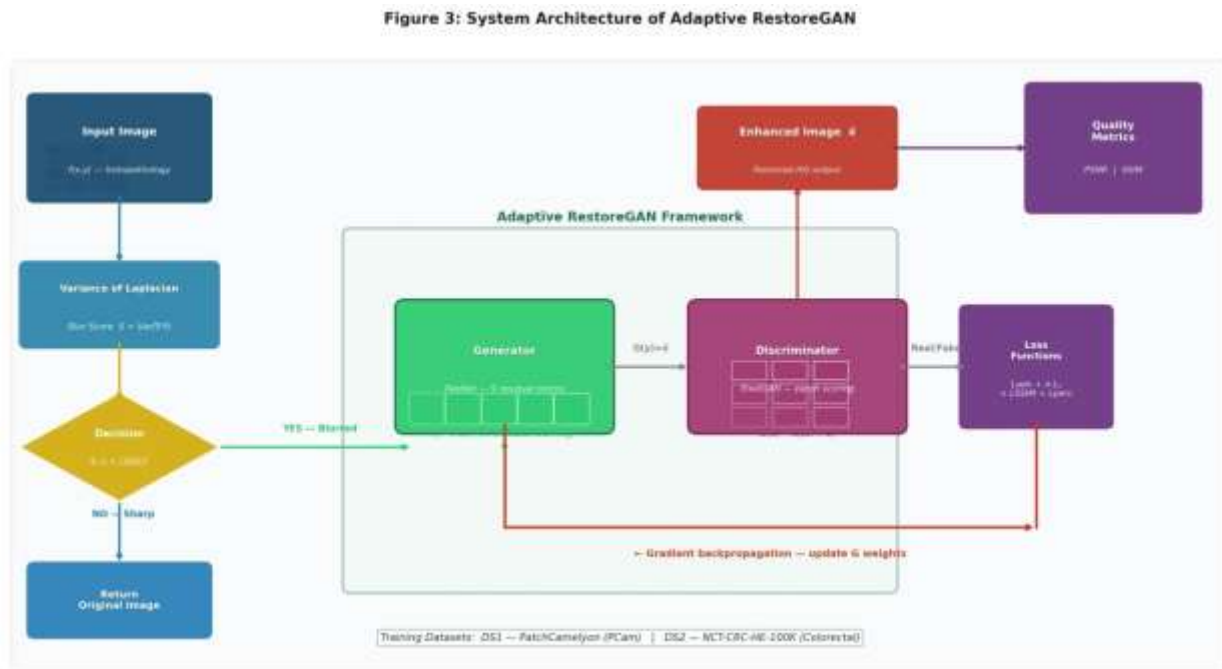


Fig.3: System Architecture of Adaptive RestoreGAN

The Adaptive RestoreGAN pipeline operates in three sequential stages. In the first stage, an input histopathology image is assessed for sharpness using the Variance of Laplacian (VOL) method, which computes a blur score $S = \text{Var}(\nabla^2 f)$. This score is compared against an empirically set threshold $\tau = 300$. Images with a blur score below the threshold are classified as degraded and forwarded for enhancement; images above the threshold are already sufficiently sharp and are returned unchanged, avoiding unnecessary computation.

In the second stage, degraded images pass through the GAN framework. The ResNet-based generator — composed of an encoder, nine residual blocks, and a decoder — learns to restore lost high-frequency detail through residual learning ($\hat{x} = y + R(y)$). The PixelGAN discriminator evaluates realism at the patch level, providing fine-grained adversarial feedback that encourages the generator to produce locally sharp, structurally faithful outputs. The generator is optimised using a combined loss of adversarial, L1 reconstruction, SSIM, and perceptual terms ($\mathcal{L}_G = \lambda_{\text{adv}}\mathcal{L}_{\text{adv}} + \lambda_1\mathcal{L}_1 + \lambda_{\text{SSIM}}\mathcal{L}_{\text{SSIM}} + \lambda_{\text{perc}}\mathcal{L}_{\text{perc}}$).

In the third stage, the enhanced image is evaluated using PSNR and SSIM metrics to quantify the improvement in reconstruction fidelity and structural similarity relative to the clean reference image. The system was trained and validated on two publicly available datasets — PatchCamelyon (DS1) and NCT-CRC-HE-100K (DS2) — with all images resized to 96×96 pixels.

C. Blur Detection — Variance of Laplacian

The Variance of Laplacian (VOL) quantifies image sharpness by measuring the spread of second-derivative responses:

$$s = \text{Var}(\nabla^2 f) \quad (27)$$

- (i) Low VOL score → high blur → image is enhanced by the GAN.
- (ii) High VOL score → image is sharp → no enhancement applied, original returned.

A threshold of $\tau = 300$ is used empirically and can be adjusted per dataset.

The value $\tau=300$ was determined through empirical observation rather than a closed-form derivation. During dataset preparation, the Variance of Laplacian was computed across a representative sample of both clean and synthetically blurred images from the PatchCamelyon and NCT-CRC-HE-100K datasets. Clean, sharp images consistently produced VOL scores well above 300, while images degraded by Gaussian blur — the primary degradation modelled in this study — produced scores well below 300. The value of 300 was therefore identified as a reliable decision boundary that separates the two populations with minimal misclassification on the training data.

It is important to note that this threshold is dataset-specific. The appropriate value of τ depends on factors such as image resolution, staining protocol, tissue type, and the nature and severity of degradation. For images acquired under different clinical or imaging conditions, τ should be recalibrated by inspecting the VOL score distribution of a representative sample before deployment. A lower τ makes the system more conservative — enhancing only severely blurred images — while a higher τ causes more images to be flagged for enhancement, including those

that may already be acceptably sharp. The adaptive threshold mechanism therefore provides a practical trade-off between computational efficiency and restoration coverage.

IV. Dataset and Implementation

A. Dataset

The system is evaluated on two publicly available histopathology datasets:

Dataset	Description	Image Size	Pairs
DS1 — PatchCamelyon (PCam)	Cancer detection (lymph node patches)	96×96 px	Paired clean/blurred
DS2 — NCT-CRC-HE-100K	Colorectal tissue (9 classes)	96×96 px	Paired clean/blurred

Degraded versions are generated by applying Gaussian blur to clean images, simulating optical blur and motion artefacts. All images are resized to 96×96 pixels and organised into structured paired folders (high-quality / low-quality).

Dataset splits:

Split	DS1	DS2
Training	80%	80%
Validation	10%	10%
Test	10%	10%

Limitation: Degradation is synthetically generated. Validation on real-world clinical images with complex, mixed degradation is required before clinical deployment.

B. Implementation Details

Hyperparameter	Value
Framework	PyTorch
Epochs	100
Batch size	16
Learning rate (G)	1×10^{-4}
Learning rate (D)	1×10^{-4}
Optimizer	Adam ($\beta_1=0.5$, $\beta_2=0.999$)
λ_{adv}	1.0
λ_1	100.0
λ_{SSIM}	10.0
λ_{perc}	10.0
Blur threshold τ	300

The full system uses **PyTorch** for model training, **FastAPI** for backend inference, and **React** for the web-based user interface.

V. Evaluation Metrics

1. Peak Signal-to-Noise Ratio (PSNR)

$$\text{PSNR} = 10 \log_{10} \left(\frac{L^2}{\text{MSE}} \right) \quad (28)$$

where L is the maximum pixel value (255 for 8-bit images) and MSE is the mean squared error between the restored and reference images. Higher PSNR indicates better quality.

2. Structural Similarity Index (SSIM)

$$\text{SSIM}(x, y) = \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)} \quad (29)$$

where μ_x, μ_y are image means; σ_x^2, σ_y^2 are variances; σ_{xy} is the covariance; and C_1, C_2 are stabilisation constants. SSIM ranges from 0 to 1; values closer to 1 indicate higher structural similarity.

VI. Experimental Results

A. Quantitative Results

Table 1 presents PSNR and SSIM before (LQ) and after (HQ) GAN-based enhancement for four sample images. All samples had blur scores well below the threshold ($\tau = 300$) and were correctly flagged for enhancement.

Table 1: Performance of Adaptive RestoreGAN

No.	Dataset	Blur Score	PSNR (LQ) dB	PSNR (HQ) dB	SSIM (LQ)	SSIM (HQ)	Status
1	DS1 — Cancer Detection	36.6	16.91	18.63	0.3484	0.6899	Enhanced
2	DS1 — Cancer Detection	66.8	14.01	16.53	0.3085	0.7623	Enhanced
3	DS2 — Colorectal Texture	62.4	24.16	29.50	0.7695	0.9355	Enhanced
4	DS2 — Colorectal Texture	65.8	23.95	29.20	0.7807	0.9344	Enhanced

Average gains by dataset:

Dataset	Avg. PSNR gain (dB)	Avg. SSIM gain
---------	---------------------	----------------

Dataset	Avg. PSNR gain (dB)	Avg. SSIM gain
DS1	+1.7	+0.34
DS2	+5.3	+0.16

Note on inconsistency: PSNR gain is substantially higher on DS2 than DS1, while SSIM gain is higher on DS1. This is consistent with DS2’s richer low-frequency texture patterns (colorectal tissue), which benefit more from PSNR-measured reconstruction fidelity. DS1 images exhibit sparser cell structures whose fine spatial organisation is better captured by SSIM. Future work will include dataset-level averages over the full test split with standard deviations and statistical significance testing. The quantitative performance is shown in Figure 4.

Figure 4: Quantitative Performance of Adaptive RestoreGAN

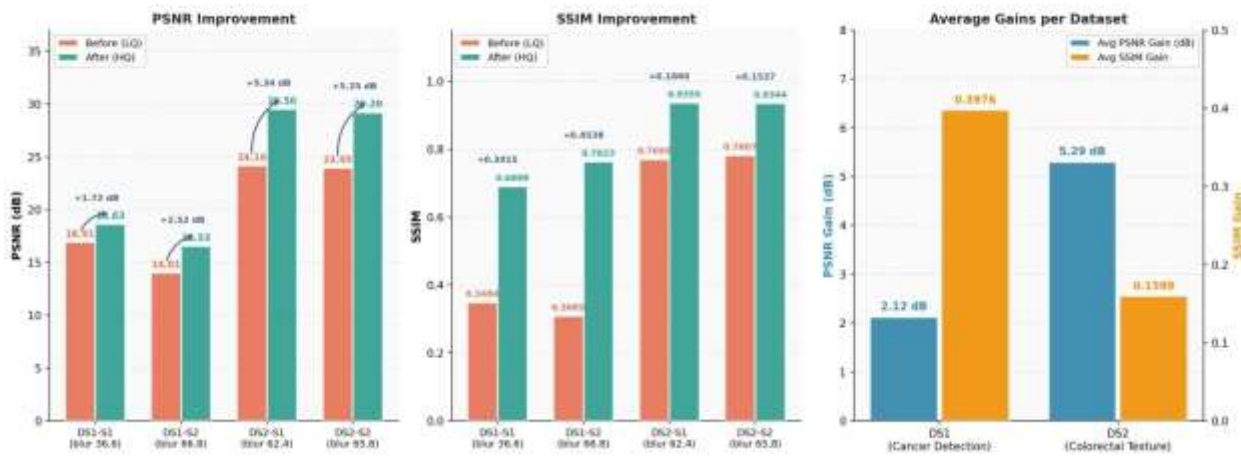


Fig.4: Performance of Adaptive RestoreGAN

Figure 4 presents a visual summary of the enhancement results across all four test samples from both datasets. The left and center panels illustrate the per-sample PSNR and SSIM values before (low-quality, LQ — shown in orange) and after (high-quality, HQ — shown in teal) GAN-based restoration. Gain annotations are provided for each sample, highlighting the magnitude of improvement achieved through the restoration process. Consistent increases in both metrics confirm that the system reliably improves image quality across different tissue types and degradation levels. The right panel compares the average gains per dataset, highlighting that DS2 achieves a larger PSNR gain (+5.3 dB) while DS1 shows a greater SSIM gain (+0.34), reflecting the structural differences between colorectal texture images and the sparser cell patterns of the cancer detection dataset.

Baseline Comparison

Acknowledged gap: The current evaluation does not include comparison against established SOTA methods. A direct benchmarking against ESRGAN, SwinIR, SRCNN, and MIRNet-v2 on the same datasets is required for a rigorous assessment of the proposed system’s relative merit.

VII. Model Validation — System Interface

The Adaptive RestoreGAN framework is equipped with an intuitive web-based interface designed to facilitate seamless interaction and visualization of the image enhancement pipeline. The interface is organized into six functional screens, each corresponding to a specific stage of the workflow:

(i) Screen 1: Image Upload Interface

Provides a drag-and-drop mechanism for uploading histopathology images, enabling quick and user-friendly data input.

(ii) Screen 2: Processing Pipeline Visualization

Displays the complete workflow of the system, including image upload, resizing, blur detection using the Variance of Laplacian (VOL) method, and conditional enhancement using the GAN model. This screen offers transparency into the internal processing steps.

(iii) Screen 3: File Selection Interface

Allows users to browse and select images from the local system. The interface supports multiple image formats, including PNG, JPG, and TIFF, ensuring compatibility with standard medical imaging data.

(iv) Screen 4: Dataset Selection Module

Enables users to choose between two datasets: DS1 (Histopathologic Cancer Detection) and DS2 (NCT-CRC-HE-100K colorectal dataset), allowing flexible experimentation and evaluation.

(v) Screen 5: Blur Detection Output

Displays the computed blur score using the VOL method. For instance, a sample image yields a blur score of **62.3**, which is significantly below the predefined threshold $\tau=300$, thereby triggering the enhancement process.

(vi) Screen 6: Comparative Visualization

Presents a side-by-side comparison of the original, blurred, and enhanced images. Quantitative evaluation metrics, including Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM), are overlaid to demonstrate the effectiveness of the restoration process.

VIII. Conclusion

The proposed Adaptive RestoreGAN framework effectively addresses the challenge of low-quality histopathology images by integrating deep learning with an adaptive decision-making mechanism. Unlike conventional approaches that uniformly process all inputs, the proposed system first evaluates image sharpness using the Variance of Laplacian (VOL) metric. Enhancement is selectively applied only when the blur score falls below a predefined threshold τ , thereby improving computational efficiency and avoiding unnecessary processing of already high-quality images.

At the core of the framework lies a ResNet-based generator coupled with a PixelGAN discriminator, trained using a composite multi-term loss function. Experimental results on synthetically degraded histopathology datasets demonstrate consistent improvements in both Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM), indicating enhanced visual quality and structural preservation. Furthermore, the end-to-end deployment using PyTorch (model training), FastAPI (backend inference), and React (frontend interface) enables practical usability, making the system accessible for clinical and research applications.

Despite these promising results, several limitations remain. First, the current validation is restricted to synthetic degradation (Gaussian blur), which may not fully capture the complexity of real-world clinical artifacts; hence, validation on real clinical datasets is essential. Second, the experimental evaluation is limited in scale, with only a small number of samples reported; comprehensive dataset-level analysis, including statistical measures such as variance and significance testing, is required for stronger validation. Third, the absence of benchmark comparisons with state-of-the-art models such as ESRGAN, SwinIR, SRCNN, and MIRNet-v2 limits the ability to position the proposed method within the broader literature. Overall, Adaptive RestoreGAN presents a practical and scalable solution for histopathology image restoration. By combining adaptive processing with deep generative modeling, the framework offers strong potential for improving image quality and, consequently, enhancing the reliability of downstream diagnostic systems in medical AI pipelines.

Acknowledgements: We sincerely thank our guide, **Dr. G. Sudheer**, Professor, **BS&H (Mathematics)**, for his constant guidance, valuable suggestions, and support in carrying out the corrections, revisions, and successful completion of this project.

References

- [1] R. Rong et al., “Enhanced Pathology Image Quality with Restore–Generative Adversarial Network (RestoreGAN),” *American Journal of Pathology*, vol. 193, no. 4, pp. 450–462, 2023, doi: 10.1016/j.ajpath.2023.01.027.
- [2] T. Zehra et al., “Rethinking Histology Slide Digitization Workflows for Low-Resource Settings,” in *Proc. MICCAI 2024*, Lecture Notes in Computer Science, Springer, 2024, doi: 10.1007/978-3-031-72083-3_40.
- [3] X. Wang, N. E. and J. Yang, “Image Clearness Processing for Image Restoration Based on Generative Adversarial Networks,” *Informatics in Medicine Unlocked*, 2021, doi: 10.1016/j.imu.2021.100550.
- [4] Y. L. K. F. Cruz et al., “Generative Adversarial Networks in Histological Image Segmentation: A Systematic Review,” *Applied Sciences*, vol. 12, no. 14, 2022, doi: 10.3390/app12147802.
- [5] B. Bhavana et al., “A Novel Approach for Image Restoration Using Generative Adversarial Networks,” in *Proc. IEEE Int. Conf. on Electronics, Computing and Communication Technologies (CONECCT)*, 2020, pp. 1–6.



[6] M. Touhami et al., “Histopathology Image Enhancement Using Multi-Resolution Deep Learning Techniques,” *IET Image Processing*, vol. 17, no. 14, pp. 4008–4019, 2023, doi: 10.1049/ipr2.12900.

[7] A. Madabhushi and G. Lee, “Image Analysis and Machine Learning in Digital Pathology: Challenges and Opportunities,” *Medical Image Analysis*, vol. 33, pp. 170–175, 2016, doi: 10.1016/j.media.2016.06.037.

[8] Z. Chen et al., “Deep Learning for Image Enhancement and Correction in Magnetic Resonance Imaging,” *Journal of Digital Imaging*, vol. 35, no. 4, pp. 698–733, 2022, doi: 10.1007/s10278-022-00721-9.

[9] H. Li et al., “A Generalized Framework of Feature Learning Enhanced Convolutional Neural Network for Histopathology Image Analysis,” *Computers in Biology and Medicine*, vol. 144, p. 105544, 2022, doi: 10.1016/j.compbiomed.2022.105544.

[10] L. Gupta et al., “GAN-Based Image Enrichment in Digital Pathology Boosts Segmentation Accuracy,” in *Proc. MICCAI 2019*, Lecture Notes in Computer Science, Springer, 2019, pp. 631–639, doi: 10.1007/978-3-030-32239-7_70.

[11] P. Amerikanos and I. Maglogiannis, “Image Analysis in Digital Pathology Utilizing Machine Learning and Deep Neural Networks,” *Journal of Personalized Medicine*, vol. 12, no. 9, p. 1444, 2022, doi: 10.3390/jpm12091444.