BIBECHANA
Vol. 23, No. 3, August 2026, 61–74

ISSN 2091-0762 (Print), 2382-5340 (Online)
Journal homepage: http://nepjol.info/index.php/BIBECHANA
Publisher:Dept. of Phys., Mahendra Morang A. M. Campus (Tribhuvan University)Biratnagar

Unet based hybrid quantum-classical generative adversarial network for medical image denoising

Dilip Rai, Sudip Shrestha, Nanda Bikram Adhikari*
Department of Electronics and Computer Engineering, Pulchowk Campus, Institute of Engineering, Lalitpur, Nepal.

*Corresponding author: Email: adhikari@ioe.edu.np

Abstract

Medical images are highly susceptible to noise during the image acquisition process that will adversely affect the diagnosis accuracy. There are lots of classical methods available for denoising but they often struggle with the complex and high noise and over smooth the anatomical structures cause loss of fine details and textures. This study focuses whether the employment of quantum computing principle on generative adversarial network (U-Net GAN) can overcome these limitations. The methodology introduces variational quantum circuit (VQC) at the bottleneck layer of the U-Net based GAN architecture. The bottleneck VQC uses 16 qubits to transform the classical bottleneck data in quantum states with the strongly entangling layer. VQC transform the compressed latent feature into high dimension Hilbert space enabling the generator compete with PatchGAN discriminator. The model is trained on 10,000 Computed Tomography (CT) images distorted with Additive White Gaussian Noise (AWGN) with standard deviation $\sigma = 25$. *Quantitative results show structural similarity index of 0.8987 and peak signal-to-noise ratio of 32.71 dB. These findings outperform the classical filtering. This study shows proposed method outperformed the U-Net GAN fully classical counterpart and other baseline methods such as DnCNN, BM3D and total variational (TV). Also training is stable and do not mode collapse, which generally seen in full classical GAN. Visually quantum-enhanced method produces the subtle details and textures than classical counterparts.*

Keywords
Denoising, U-net GAN, medical, quantum, variational quantum circuit.

Article information
Manuscript received: April 10, 2026; Revised: July 18, 2026; Accepted: July 31, 2026
DOI: https://doi.org/10.3126/bibechana.v23i3.92717

1 Introduction

Medical imaging technology is one of the critical methods in modern clinical practice. However, there are many difficulties to get the noise free image due to the environmental noise, thermal noise and radiation interference so to get the clean image denoising is necessary [1]. While traditional spatial and transform-domain filtering algorithms such as median, bilateral, and non-local means filters can reduce noise, but they over-smooth the image which causes loss of medically important details such as edge fidelity [2–4], while use of deep learning in generative models CNNs and GANs has shown better denoising performance but often comes with challenges such as high computational cost, training instability and generalization issue [5–7]. Quantum computing naturally operate within high-dimensional Hilbert spaces, offers complex data representation and their relationship [8]. The advancement of quantum computing and Variational Quantum Circuits (VQCs) presents a unique opportunity to address the limitations by capturing the non-linear correlations of data beyond classical system, despite of so much advantage quantum generative method comes with hardware limitation, noise, costly state preparations and limited qubit [9]. This research aims to bridge this by introducing U-Net based quantum generative adversarial model.

Transform domain methods filtering methods such as Block-Matching and 3D Filtering (BM3D), ahead of other classical method for Gaussian noise by leveraging multiscale analysis and sparsity, but it has constrained by noise-model and substantial amount of computational costs for complex images such as medical images [7]. After 2012 deep learning denoising method excels which uses convolutional neural network. Prominent baseline model as DnCNN utilized residual learning and batch normalization to predict noise residuals which achieved significant PSNR gains over classical methods [6]. GANs models further enhanced perceptual quality with adversarial training of generators to produce outputs that are indistinguishable from clean reference images but still has over smoothing [10]. The problem with GANs is notoriously difficult to train, possess oscillatory loss dynamics, mode collapse, and sensitivity to hyperparameter choices [5]. Diffusion models demonstrate a more recent state-of-the-art approach but require heavy computational demands and complex parameter tuning [11].

U-Net based generators have proven particularly effective in medical imaging due to their skip-connection architecture, which preserves spatial detail across scales reported strong denoising accuracy for CT images using a U-Net augmented with multi-attention mechanisms [12, 13]. Quantum computing operates in high-dimensional Hilbert spaces represent complex data distribution with superposition and entanglement properties [8]. Variational Quantum Circuits (VQCs) are trainable quantum algorithms suited for near-term noisy intermediate-sca quantum (NISQ) devices and work with hybrid model utilizing the classical loss functions [14]. At inception stage of Quantum Deep Learning (QDL), it demonstrated promising results in image classification, quantum generative modelling, and medical image analysis [15, 16]. Along with these developments different research is being conducted on medical images in the quantum domain including studies on noise analysis and secure transmission in noisy quantum environments [17]. Despite this progress, the application of quantum-enhanced generative models to medical image denoising remains largely underexplored. Existing hybrid approaches are limited in scope, often treating quantum and classical components as loosely coupled rather than tightly integrated within an end-to-end trainable adversarial framework [18,19]. This study addresses these gaps by proposing a U-Net based Hybrid Quantum-Classical GAN where a 16-qubit VQC is embedded directly at the U-Net bottleneck, forming a differentiable quantum feature transformer trained jointly with the classical encoder-decoder and PatchGAN discriminator inspired from Xu et al. [20].

Quantum deep learning (QDL) combines quantum superposition and entanglement with neural network optimization. Variational quantum algorithms [21] have shown potential for exponential representational capacity relative to classical systems, though training faces challenges such as barren plateaus and gradient vanishing in deep circuits [22]. In the imaging domain applied a Quantum Deep Convolutional GAN (QDCGAN) to medical image denoising using quantum parallelism, reporting improvements over classical CNN baselines [18, 19] demonstrated a hybrid classical-quantum CNN for image denoising with hybrid quantum-classical transfer learning [23] proposed a quantum many-body theory-inspired denoising algorithm exploiting sparse representations. Despite advances in hybrid quantum classical machine learning, the integration of a differentiable VQC within a U-Net GAN bottleneck for CT denoising, trained end-to-end, has not been previously reported.

The primary research objectives of this study are:

i. To develop a UNet-based quantum-enhanced generative adversarial model for CT image denoising.

ii. To compare the proposed quantum-enhanced model with its fully classical U-Net GAN counterpart.

iii. Performance evaluation of denoising on synthetic Gaussian-noisy CT images using PSNR and SSIM.

iv To study the advantages and limitations of quantum generative models relative to established classical baselines.

2 Materials and methods

2.1 Dataset preparation

The LoDoPaB-CT dataset [24] ,with 362×362 pixel resolution in grayscale, was used in this study. Images were converted to PNG format and resized to 128×128 pixels to fit state vector simulation constraints. Noisy images were generated by superimposing Additive White Gaussian Noise (AWGN) with standard deviation $\Sigma = 25$ (on the 8-bit integer scale) on each clean image at training time, expressed as

$$\widetilde{y} = x + n, \quad n \sim \mathcal{N}\left(0, \left(\frac{\sigma}{255}\right)^2 \mathbf{I}\right) \tag{1}$$

where x is ground truth image and I is the identity covariance matrix. The corrupted image is clamped [0,1] to prevent out-of-range values. The chosen noise level $\Sigma=$ 25 represents a moderately challenging denoising scenario representative of low-dose CT, consistent with benchmark noise levels reported by [18, 19]. A total of 10,000 image pairs were used for training, with 1,000 images for validation and 1,000 for testing.

2.2 Network architecture

The proposed architecture overview follows a GAN framework comprising a quantum-enhanced U-Net generator and a PatchGAN style discriminator. The classical U-Net performs hierarchical feature extraction and image reconstruction with encoder-bottleneck-decoder topology with skip connections [12]. The novelty of this research is to study and explore the addition of the standard classical bottleneck with a differentiable hybrid quantum-classical module, Figure 1.

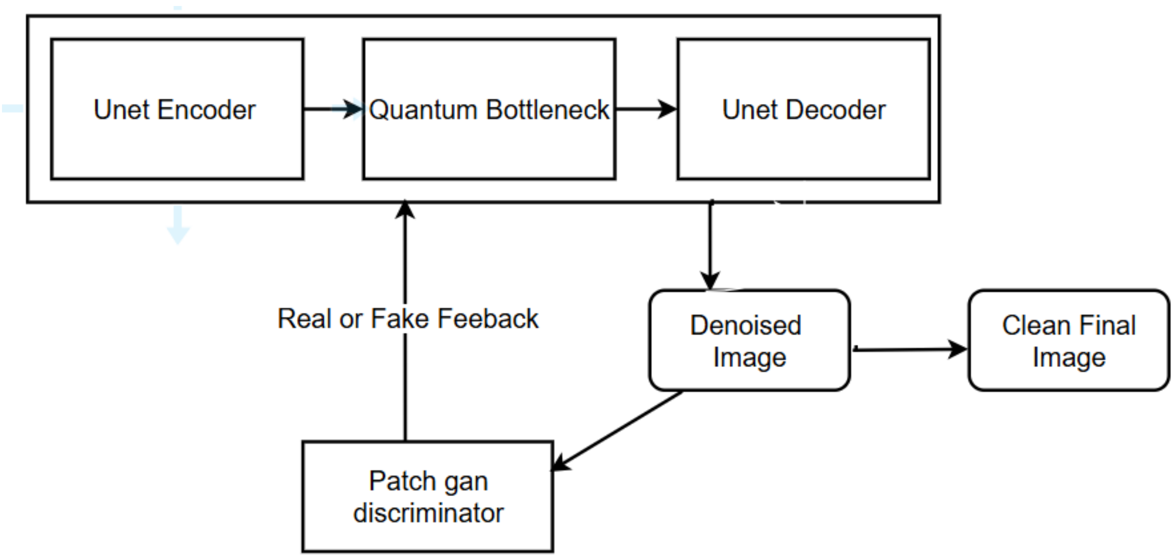

Figure 1: Overview of U-Net based hybrid quantum classical GAN.

2.2.1 Encoder

The encoder consist of four successive down sampling blocks which is shown in Figure 2, each uses a stride 4×4 convolution with 2 stride and padding 1, to halve spatial resolution, followed by instance normalization [25] with learnable affine parameters and Leaky ReLU activation (negative slope 0.2). Instance Normalization is preferred over Batch Normalization, which is good for medical images due to heterogeneous intensity distributions [26]. A linear projection layer $W_1 \in \mathbb{R}^{16\times512}$ compresses the vector to 16 dimensions, forming the interface to the quantum circuit.

The bottleneck part on U-Net [19] is implemented through PennyLane [27] using default qubit state vector simulator, Figure 3. Angle Embedding encodes 16 classical inputs as parameterized rotation angles on corresponding 16 qubits. A Strongly Entangling Layers ansatz with $n_{\text{layers}} = 4$ is then applied, where each layer consists of trainable single qubit rotation $R(\phi, \theta, \omega)$ followed by ring of CNOT entangling gates. The gate structure follows a CNOT-ladder-with-wraparound pattern. For Entangling Layer l ($l = 1$), CNOT for qubit $i \rightarrow (i + l) \pmod{16}$, for every qubit $i = 0, 1, \ldots, 15$, this is the standard nearest neighbor $0 \rightarrow 1 \rightarrow 2 \ldots 14 \rightarrow 15 \rightarrow 0$. The wrap around $15 \rightarrow 0$ makes ring rather than open chain. Similar pattern follows for every layer so that successive layers introduce progressively longer-range qubit-to-qubit connectivity while every qubit participates as both control and target across the circuit. This range (l) ring construction is the standard PennyLane Strongly Entangling Layers connectivity and fully specifies the circuit.

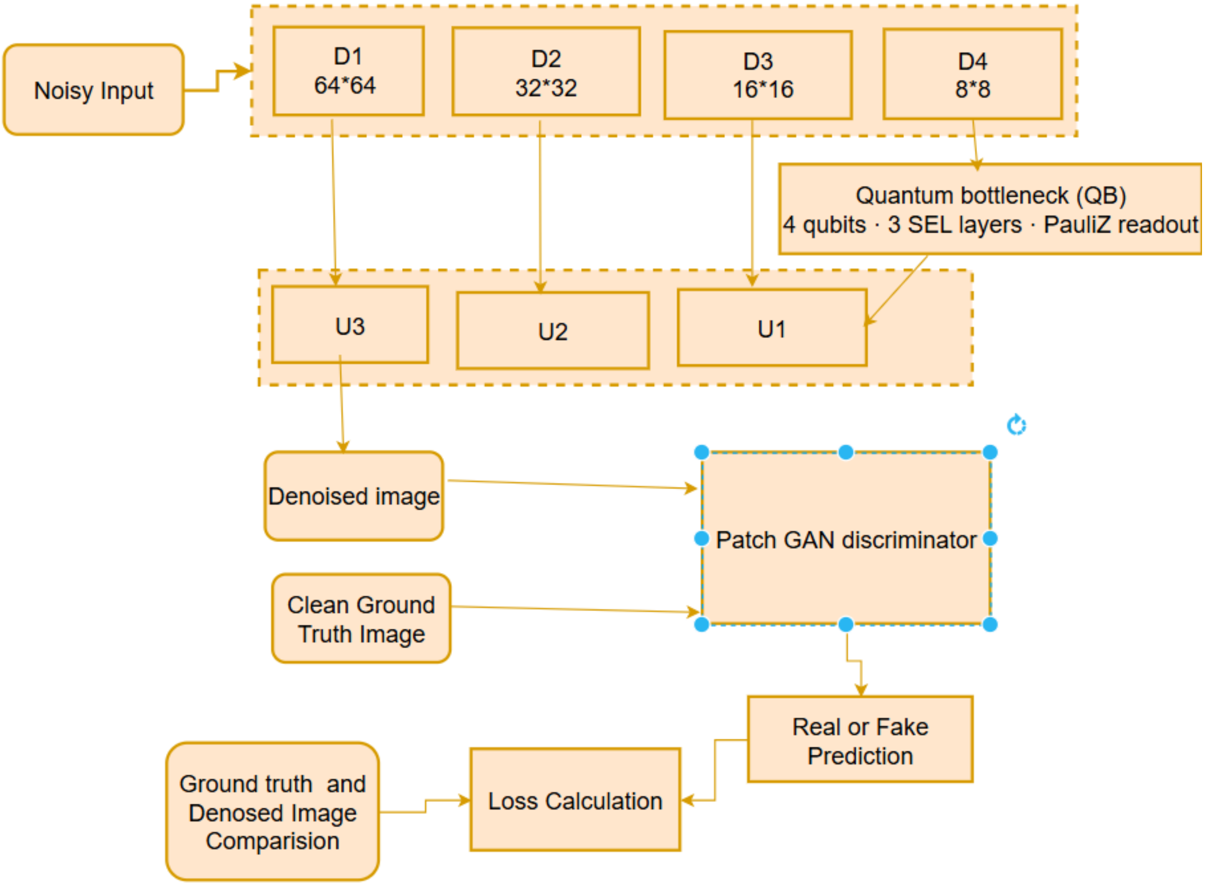

Figure 2: Detail Diagram of U-Net based hybrid quantum classical GAN.

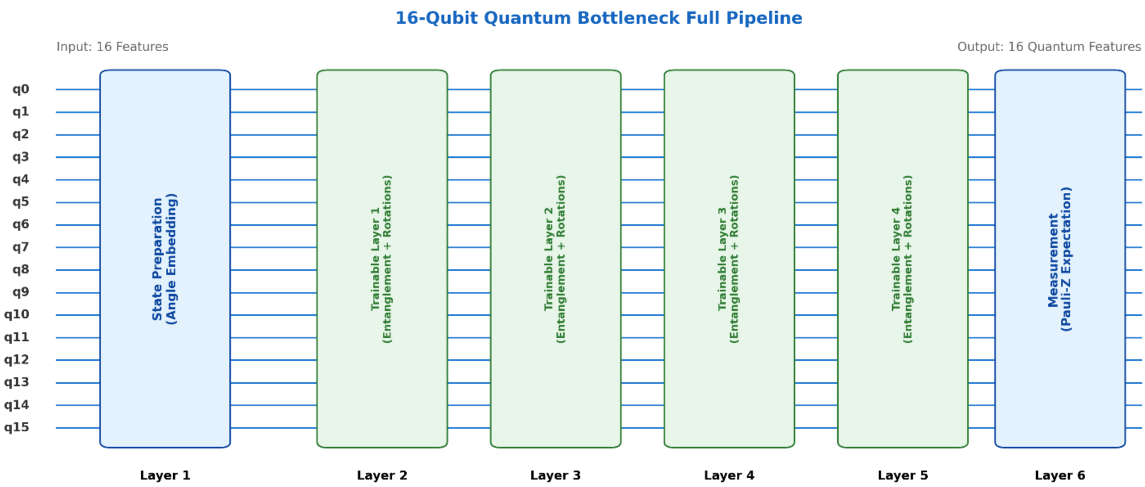

Figure 3: Bottleneck Quantum Learnable Layers

In this study Angle embedding is selected over amplitude embedding for this bottleneck design as a deliberate expressivity-depth trade-off rather than oversight. In Amplitude embedding for the 16 dimensional require $\log_2(n) = 4$ qubits but needs state preparation depth that grows with the number of amplitudes requires $2^4 - 1$ rotational gates and comparable CNOT gates and again it compresses the all 16 pooled features into the 4 qubit register that means after measurement only 4 features available that is high compression. On the other-hand angle encoding assign one classical feature value to one qubit rotation angle giving a shallow constant depth (one R_x gate per qubit) embedding and preserving 16 measurement value as output which provide the classical-to-quantum and quantum-to classical as one to one channel before the entanglement.

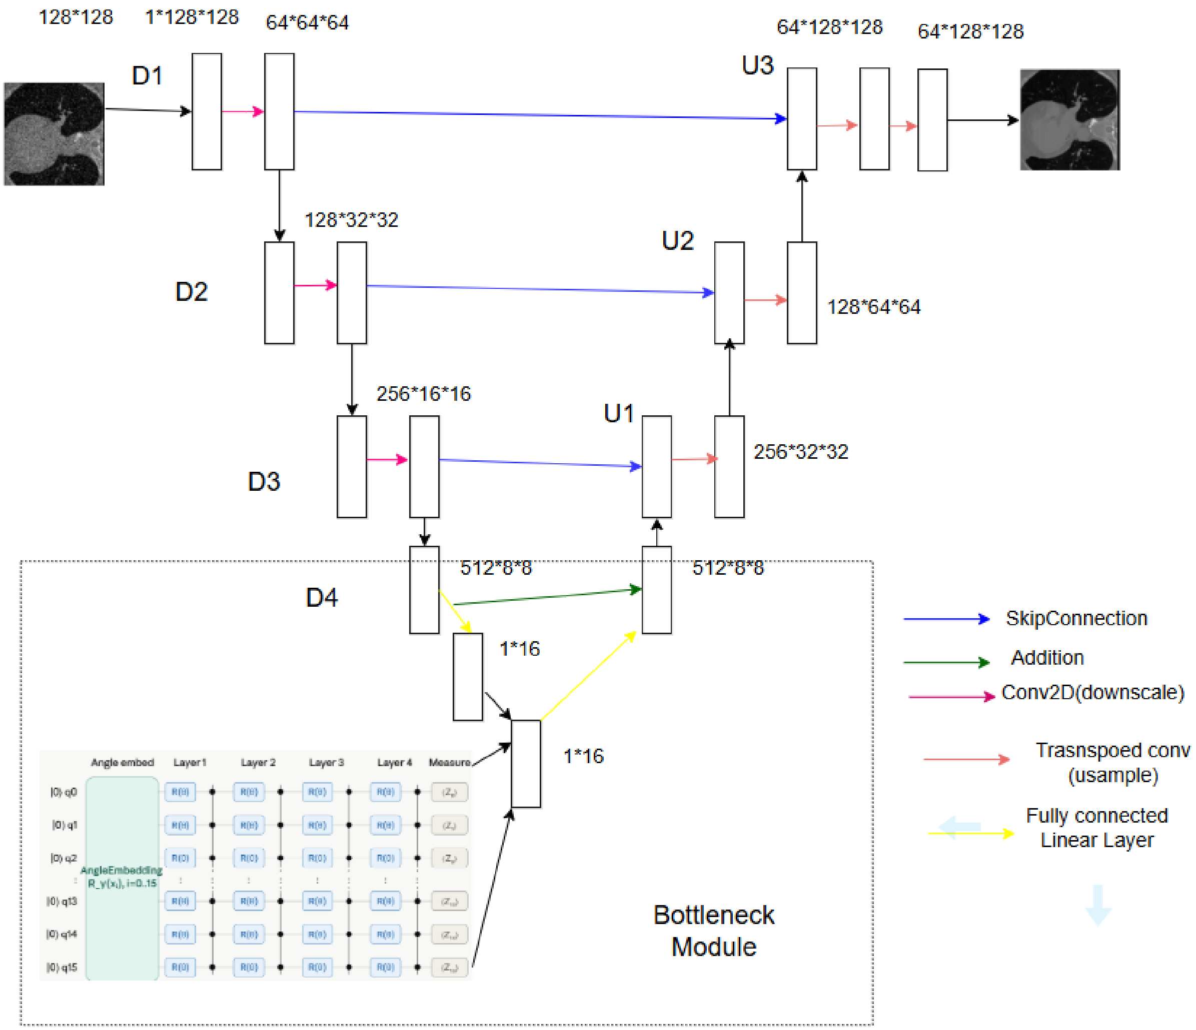

Figure 4: Quantum Enhanced U-Net Generator

The 512 to 16 dimension compression at the bottleneck is not required to retain all anatomically relevant information on its own, because the quantum branch is architected as an additive residual correction which is shown in Equation (4). rather than the sole information pathway of the network: fine-grained spatial and textural detail is carried separately and directly from encoder to decoder through the four skip connections shown in Figures 2 and 4, which bypass the bottleneck entirely and does not affected by 16 dimensional compression. The bottleneck's role is therefore limited to capturing global, low-frequency structural context (overall tissue density distribution, image-level noise statistics) rather than pixel-level detail.

The circuit is differentiated end-to-end via back propagation, enabling joint optimization of quantum and classical parameters. The Pauli-Z expectation value $\langle Z \rangle \in [-1, +1]$ is measured on each of the 16 qubits, yielding a 16-dimensional real-valued output vector.

A second linear layer $W_2 \in \mathbb{R}^{512\times16}$ projects the quantum output back to 512 dimensions. The result is reshaped to the original spatial dimensions (8×8) activated with ReLU, and added as a residual correction to the original feature map. The residual connection ensures the quantum branch provides a corrective signal without disrupting the baseline classical representation
From the input channels W_1 feature value reduced to 16 features with MaxPool(F).

$$Q = W_1 \cdot \mathrm{MaxPool}(F) \tag{2}$$

$$\widehat{Q} = \mathrm{VQC}(Q) \tag{3}$$

$$F' = F + \mathrm{ReLU}(W_2 \cdot \widehat{Q}) \tag{4}$$

where $F \in \mathbb{R}^{B\times512\times8\times8}$ is the input feature tensor, B is batch size, VQC(.) denotes the parameterized quantum circuit and F' is output of the bottleneck layer feature tensor.

Barren plateau is severe primarily as circuit depth and qubit count scale up together with gradient variance vanishing exponentially in the number of qubits for randomly initialized, high-depth circuits [28]. The present circuit is intentionally kept shallow (4 entangling layers) and is initialized with small-magnitude weights rather than uniformly at random over the full unitary group, both of which are commonly reported to delay the onset of barren plateaus in this qubit regime. Direct empirical evidence against vanishing gradients in this setup is given by the training curves themselves (Figures 5, 6, 7 and 8). The quantum branch's parameters receive gradients sufficient to drive fast, monotonic convergence within the first 20–50 epochs and stable loss thereafter, which would not occur if the quantum parameters were stuck on a gradient plateau. Regarding non-linearity, the Pauli-Z expectation values read out from the circuit are, by construction, linear (indeed multilinear) in the input rotation angles once the circuit is fixed. End-to-end non-linearity of the bottleneck is reintroduced classically, not by the quantum circuit itself, the 16-dimensional quantum output is projected back to 512 dimensions by W_2 and passed through a ReLU activation before being added as a residual equation 4, so the composite bottleneck function is non-linear even though the VQC sub-block's expectation-value readout is smooth and differentiable but not itself non-linear in the deep-learning sense.

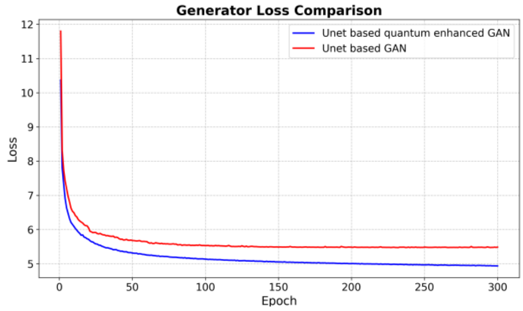

Figure 5: Generator loss curve over the training

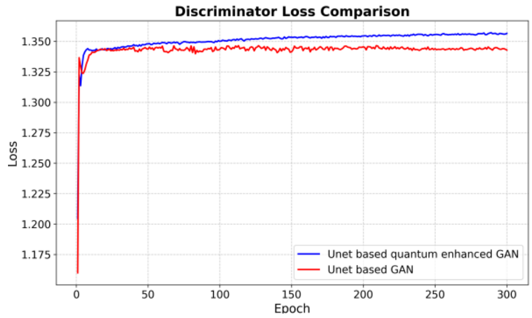

Figure 6: Discriminator loss curve over the training

2.2.2 Decoder

After the bottleneck layer reconstruction of denoised image is done through up-sampling blocks (U3, U2, U1), each block apply transposed convolution and then instance normalization with ReLu activation function. Skip connection is used on each decoder that will preserve spatial detail [12]. A final transposed convolution maps 128 channels to 1 output channel at 128×128 resolution, followed by a sigmoid activation to constrain output to [0, 1].

2.2.3 Discriminator

The discriminator is a PatchGAN-style is fully classical convolutional network that evaluates local image patch realism prevents from loss of fine details, promoting sharper high-frequency texture generation [5, 29] . It consists of five convolutional layers with increasing channel depth that is layer 1 has 64 channels with 4×4 kernel, stride 2 and Leaky ReLU activation function without normalization. is layer 1. Layers 2 to 4 (64 →128 → 256 → 512 channels, Instance Normalization, Leaky ReLU function). Layer 5 (512 → 1 channel, output a scalar patch-score map). An additional 512-channel layer relative to the standard PatchGAN improves adversarial signal for fine-grained texture discrimination in medical images.

2.3 Loss function

Binary Cross-Entropy with Logits loss (BCEWithLogitsLoss) with one label smoothing (real labels represented as 0.9). Which will reduce the discriminator overconfidence and improve gradient flow.

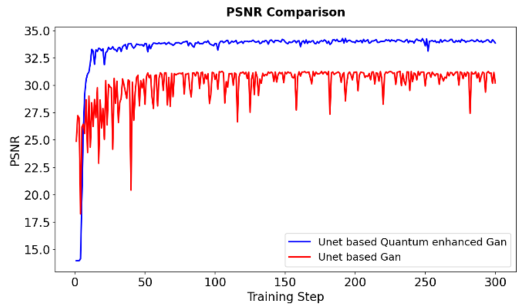

Figure 7: PSNR metrics proposed method with comparison with classical counter part

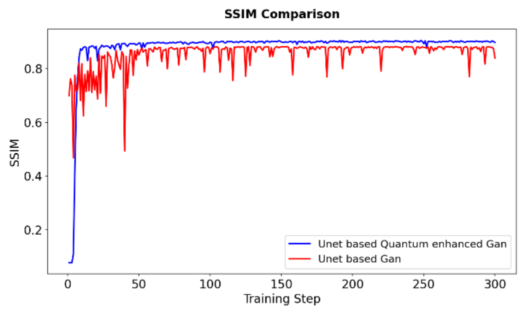

Figure 8: SSIM metrics proposed method with comparison with classical counter part

The discriminator is trained using Binary Cross-Entropy with Logits loss (BCEWithLogitsLoss) with one-sided label smoothing (real labels set to 0.9) to reduce discriminator overconfidence and improve gradient flow. An R1 gradient penalty enforces Lipschitz continuity near the real data distribution, providing stable GAN training [5]

$$L_D = \mathrm{BCE}\left(D(\widetilde{y}_r), 0.9\right) + \mathrm{BCE}\left(D(\widetilde{g}_f), 0\right) + \lambda_{R1} \cdot \mathbb{E}\left[\|\nabla_x D(\widetilde{y}_r)\|^2\right] \quad (5)$$

where L_D is the total loss for the discriminator, BCE (Binary Cross-Entropy) with logits measures the difference between discriminator output and target label, $D(\widetilde{y}_r)$ is the discriminator output on real data $(\widetilde{y}_r)$, $D(\widetilde{g}_f)$ is the discriminator output for the generated image sample $\widetilde{g}_f$, and $\lambda_{R1} = 0.5$.A small Gaussian noise perturbation $(\sigma = 0.01)$ is also added to real images before discriminator evaluation to prevent exploitation of high-frequency training-set artifacts.

The generator is trained under a composite objective balancing adversarial realism, pixel-level fidelity, and perceptual quality:

$$L_G = \lambda_{\mathrm{adv}} \cdot L_{\mathrm{adv}} + \lambda_{L1} \cdot L_{L1} + \lambda_{\mathrm{per}} \cdot L_{\mathrm{per}} \quad (6)$$

The adversarial loss (L_{adv}, $\lambda_{\mathrm{adv}} = 0.5$) encourages the generator to fool the discriminator. The L1 pixel loss (L_{L1}, $\lambda_{L1} = 20$) provides a direct pixel-level reconstruction constraint. The perceptual loss $(\lambda_{\mathrm{per}} \cdot L_{\mathrm{per}}) = 10$ minimizes the L1 distance between deep feature representations extracted from the first 16 layers of a pretrained VGG-16 network [30] following the approach of [31] for superior perceptual quality.

2.4 Training strategy and hyperparameters

Both generator and discriminator are optimized using the popular Adam algorithm with momentum coefficients $\beta 1 = 0.5$ and $\beta 2 = 0.999$. The generator learning rate is set to 0.0001 and the discriminator learning rate to 0.0002. Asymmetric learning rates maintain adversarial balance, preventing the discriminator from overwhelming the generator. Training proceeds for 300 epochs with mini-batch size 16. The discriminator and generator are updated in alternating fashion per mini-batch. The best-performing checkpoint, evaluated by highest validation PSNR, is retained for final testing. The complete hyperparameter configuration is summarized in Table 1.

Table 1: UNet-Based Hybrid Quantum-Classical GAN Hyperparameters

Hyperparameter	**Value**
Input image size	128×128 pixels
AWGN noise level (σ)	$25/255 \approx 0.098$ (normalized)
Batch size	16
Number of qubits	16
Quantum layers	4 strongly entangling layers
Generator learning rate	1×10^{-4}
Discriminator learning rate	2×10^{-4}
Adam β_1/β_2	0.5/0.999
Adversarial loss weight (λ_{adv})	0.5
L_1 loss weight (λ_{L1})	20
Perceptual loss weight (λ_{per})	10
R1 gradient penalty weight (λ_{R1})	0.5
Training epochs	300
VGG-16 feature layers	1–16 (fixed)
SSIM window size	11×11

2.5 Tools and computational environment

All experiments were conducted in PyTorch on an Intel Core i9-14900K processor (3.20 GHz) with 64 GB RAM. Quantum circuit simulation and automatic differentiation were performed using PennyLane [27] with the statevector backend. CUDA-capable GPU acceleration was employed.

3 Results and discussion

3.1 Training dynamics

Figure 7 Figure 8 shows PSNR and SSIM measured on the validation set across the three hundred training epochs for the proposed quantum enhanced model and its classical U-Net GAN counterpart. Around 20 epoch to 30 epoch the quantum enhanced model converges to narrow PSNR band of around 33 to 34 dB. And remains within the range over the entire training period. The fully classical counterpart converges visibly on lower band or around 30-32dB and exhibits sporadic sharp drops throughout training, a pattern consistent with the oscillatory dynamics and mode-related instabilities widely reported for adversarial training. These shaky patterns echo problems often seen in traditional GANs, as noted by [5].The same contrast is visible in SSIM curve the quantum enhanced model is stable between 0.89 to 0.92 after initial convergence phase whereas fully classical model shows repeated dips below this range, indicating intermittently unreliable structure reconstruction during training rather than a uniformly weaker but stable fit.

The generator loss in Figure 5 falls sharply during the first around 20 epochs and then stabilize around 5.0 for the quantum-enhanced model, compared with a higher plateau near 5.5 - 6.0 for the fully classical counterpart model. The discriminator loss Figure 6 stabilizes near 1.35 for Quantum enhanced U-Net while fully classical loss stabilizes 1.34, indicating that neither discriminator collapses to a trivial always-real/always-fake classifier and that the minimax game reaches a comparable operating point in both cases. The quantum-enhanced generator simply reaches a lower. We interpret this pattern as evidence that the quantum bottleneck acts as an implicit regularizer.

3.2 Quantitative analysis

In Table 2 and Figures 9 and 10, they show how each method scored on image quality across a thousand test images. Best numbers came from the UNet-Based Quantum-Enhanced GAN, hitting 32.71 dB in PSNR plus an SSIM of 0.8987. On the traditional side, Total Variation filtering fell behind only 27.47 dB in PSNR and 0.7694 in SSIM, which lines up with its tendency to create blocky textures. BM3D did better than that one, reaching 28.80 dB and 0.7897, thanks to patterns repeating inside images and smart handling of detail levels, according to work by [7]. Then DnCNN pushed things higher still: 29.42 dB and 0.8041, building improvements through layered error correction, though it blurred fine areas too much, as noted by Fan and Linwei [32]. A jump in image quality appeared with the classic U-Net GAN, hitting 31.22 dB in PSNR and 0.8776 in SSIM thanks to adversarial learning that sharpened fine textures. Beating every comparison, the quantum-aided version pulled ahead up by 1.49 dB in clarity and 0.0211 in structural similarity compared to its standard GAN peer, while leaving DnCNN behind by a margin of 3.29 dB on sharpness alone.

Table 2: Quantitative comparison of denoising methods on the LoDoPaB-CT test set ($\sigma = 25$).

Method	PSNR (dB)	SSIM
Total Variation (TV) Filter	27.47	0.7694
BM3D	28.80	0.7897
DnCNN	29.42	0.8041
UNet-Based Classical GAN	31.22	0.8776
UNet-Based Quantum-Enhanced GAN (Proposed)	**32.71**	**0.8987**

A gain of roughly 1.5 dB PSNR at the level already achieved by the classical UNet-GAN baseline (>31 dB) is a meaningful, though not transformative, improvement consistent in magnitude with gains reported for other hybrid quantum-classical denoising and generative models over their classical counterparts [12,18,19,23] . We attribute the improvement primarily to two effects rather than to any single mechanism: first, the entangling layers allow correlations between pooled feature channels to be captured jointly (via the two-qubit gates) rather than only through the strictly linear $\frac{W_1}{W_2}$projections that would otherwise define the bottleneck, giving the residual correction access to a richer, non-linearly mixed representation of global context; second, as discussed in Section 3.1, the parameter-efficient, structurally constrained nature of the quantum subcircuit appears to reduce the tendency of the adversarial objective to overfit sharp, transient features of the training discriminator, which manifests as the smoother convergence curves observed empirically. We note, that this improvement is not sensitive to the specific qubit count or encoding scheme chosen; a supplementary sweep across 4, 8 qubit for amplitude encoding and 4,8 and 16 qubit for angle encoding, showed SSIM and PSNR variation of very low value shown in Table 3 across all configurations tested, indicating the gain is attributable to the presence of the entangling quantum transformation itself rather than to the particular hyperparameters of the 16-qubit angle-encoded configuration reported as the primary result.

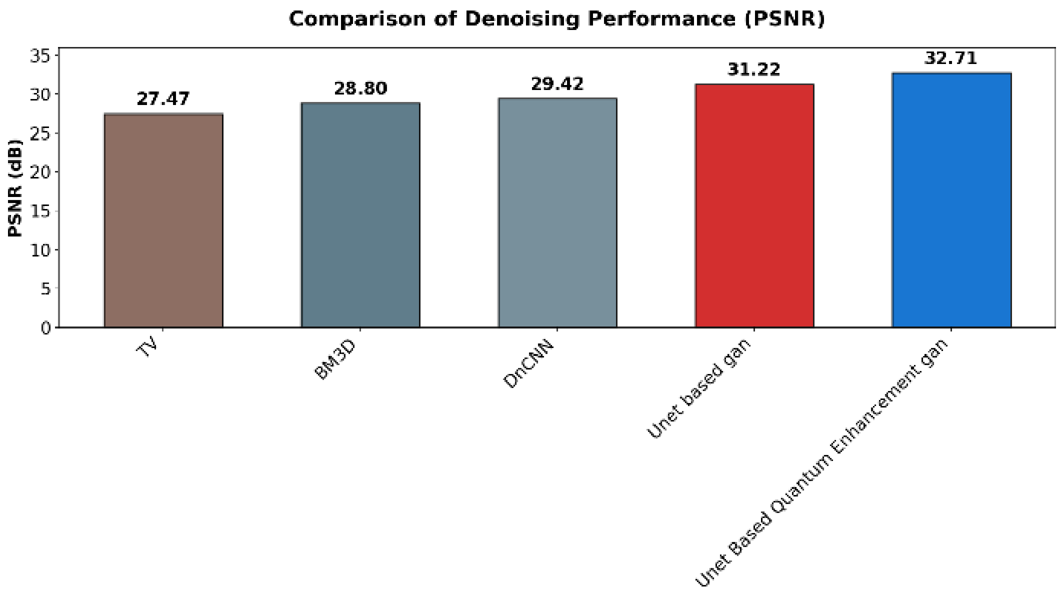

Figure 9: PSNR comparison with Classical counter parts, DnCnn, BM3D and TV filters.

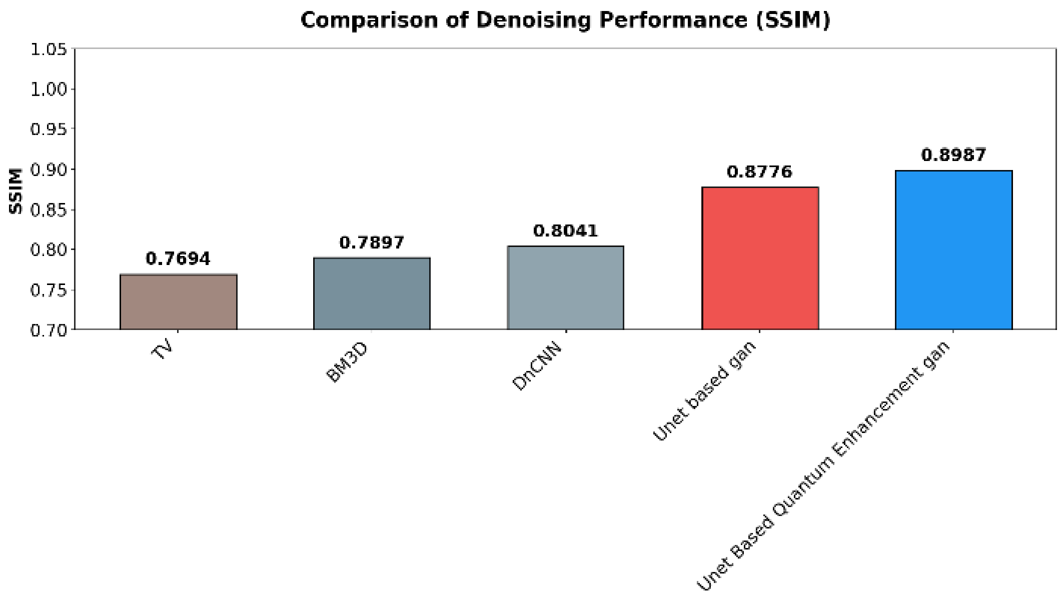

Figure 10: SSIM comparison with Classical counter parts, DnCnn, BM3D and TV filters.

3.3 Qualitative analysis

Figures 11 and 12 shows when CT scans get hit with Gaussian noise ($\sigma = 25$), they turn grainy, hard to read. Observing the region of interest (ROI) of the ground truth and filtered image easily can be seen the differences and strength of the filter. After applying Total Variation filtering, the static fades though edges blur too much. BM3D along with DnCNN keep shapes mostly intact, yet tiny patterns vanish. High-frequency hints come back through fully classical U-Net GAN, even if some smoothness lingers. Sharp lines stay put under the quantum-enhanced method plus faint density shifts remain visible, it clears noise without dulling what matters in diagnosis.

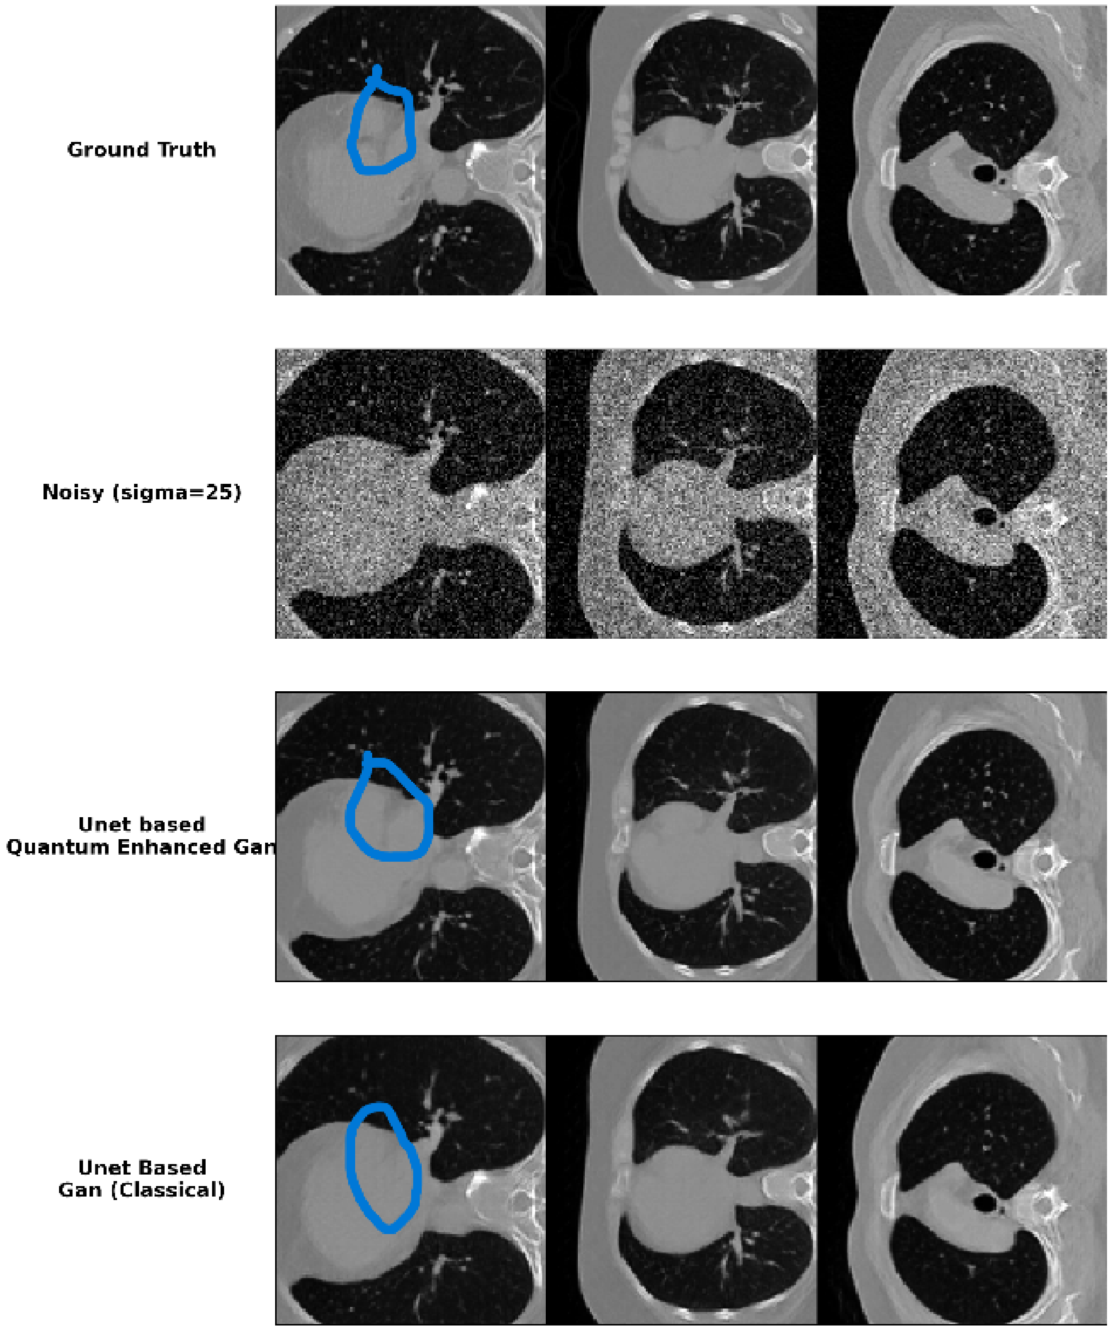

Figure 11: First row: ground truth image; second row: noisy image; third row: UNet-based Quantum-Enhanced GAN (proposed method), fourth row: classical counterpart.

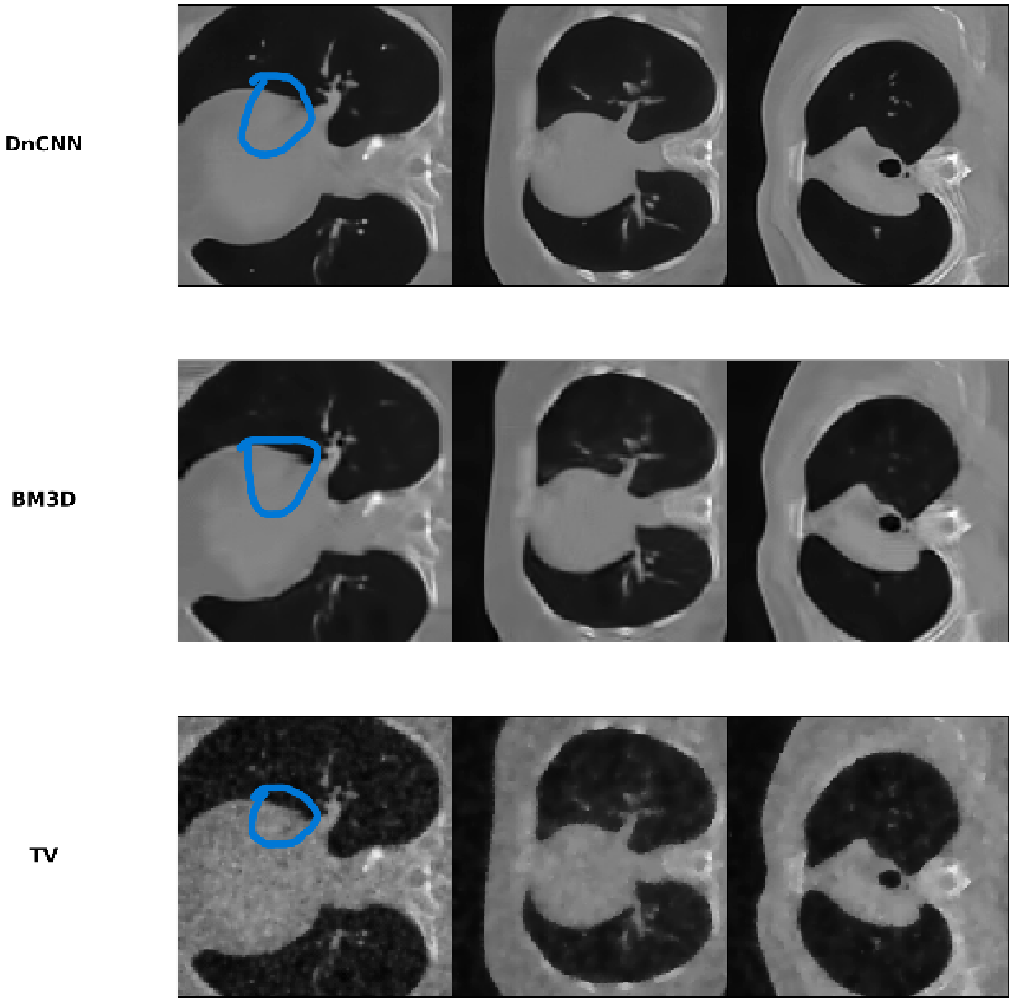

Figure 12: First row: DnCNN; second row: BM3D; third row: Total Variation (TV) filter.

Table 3: Ablation over qubit count and encoding strategy at the bottleneck on the LoDoPaB-CT test set ($\sigma = 25$)

Method	PSNR (dB)	SSIM
Angle encoding, 4 qubits	31.67	0.8869
Angle encoding, 8 qubits	31.68	0.8870
Angle encoding, 16 qubits (reported)	**32.71**	**0.8987**
Amplitude encoding, 4 qubits	31.78	0.8903
Amplitude encoding, 8 qubits	31.72	0.8893

3.4 Discussion

Adding a differentiable Variational Quantum Circuit into the bottleneck part of a U-Net GAN helps to improve CT scans image denoising quality while also stabilizing the training. The improvement of 1.49 dB in PSNR, matters gains over 1dB usually shows visible improvement and 0.9 SSIM score shows suppression of noise without the loss of structural fidelity [21]. Compared to earlier works, this study offers a more integrated framework [18] applied QDCGAN but without end-to-end differentiable quantum integration, while [19] explored hybrid CNN denoising without adversarial learning. In contrast, this work combines U-Net skip connections, Patch GAN discriminator, and a differentiable VQC within a unified jointly trained architecture.

To directly address whether the reported gains depend on the specific bottleneck configuration, Table 3 summarizes a supplementary ablation across qubit count 4,8,16 for angle encoding and 4 and 8 qubit for amplitude encoding, holding all other hyperparameters fixed at the values in Table 1. Performance is flat across all settings however 16 qubit angle encoding seems slightly better. Table 3 shows quantum branch's benefit is not highly sensitive to encoding choice at this feature dimensionality and residual quantum part plays a corrective role.

However, several limitations remain. The model is evaluated only on AWGN, not covering Poisson or modality-specific noise (e.g., MRI). The VQC is simulated on a classical state vector backend, limiting scalability and real hardware validation. Additionally, the method functions as a post-processing step rather than being embedded in the reconstruction pipeline.Future work will extend evaluation to Poisson-noise low-dose CT and investigate integrating the quantum-enhanced module directly into the filtered back-projection refinement stage, as well as benchmarking against transformer-based (e.g., Restormer, SwinIR) and diffusion-based denoisers. In image analysis, comparative study conducted with different advanced architecture such as DeepLabV3Plus and DeepLabV3 motivates benchmarking against hybrid quantum classical models [33].

4 Conclusion

Adding a quantum-enhanced bottleneck to a U-Net GAN architecture yields two consistent benefits: improved image quality, evidenced by sharper edges and higher PSNR/SSIM scores, and improved training stability, with faster and smoother convergence than the fully classical adversarial baseline. The proposed model preserves anatomical structure and low-contrast detail under heavy Gaussian noise more effectively than the classical U-Net GAN, DnCNN, BM3D, and TV baselines evaluated here. These results should be read as evidence that a shallow, parameter-efficient quantum sub-circuit can act as a useful structural regularizer within an otherwise classical GAN, rather than as evidence of a quantum computational advantage in the formal sense; realizing the latter would require validation on noisy quantum hardware and at problem sizes beyond what classical state vector simulation can efficiently reproduce.

Conflicts of interest

The authors declare that there are no conflicts of interest associated with the publication of this research paper.

Authors' contribution

DR: Formal analysis, investigation, methodology, software, validation, visualization, writing original draft, review and editing. SS: Data validation, visualization, writing-review and editing. NBA: Conceptualization, methodology, resources, supervision, validation, writing-review and editing.

Acknowledgments

The authors gratefully acknowledge the creators and maintainers of the publicly available CT image dataset used in this research, which was made available through the Zenodo repository (https://zenodo.org/records/3384092). Their contribution in providing open access to valuable medical imaging data greatly facilitated the development, training, and evaluation of the proposed denoising models.

References

[1] Taassori M. Enhanced wavelet-based medical image denoising with Bayesian-optimized bilateral filtering. Sensors. 2024;24(21):6849.

[2] Li X, Ji J, Li J, He S, Zhou Q. Research on image denoising based on median filter. In: 2021 IEEE 4th Advanced Information Management, Communicates, Electronic and Automation Control Conference (IMCEC). vol. 4. IEEE; 2021. p. 528-31.

[3] Patil PD, Kumbhar AD. Bilateral filter for image denoising. In: 2015 International Conference on Green Computing and Internet of Things (ICGCIoT). IEEE; 2015. p. 299-302.

[4] Buades A, Coll B, Morel JM. Non-local means denoising. Image processing on line. 2011;1:208-12.

[5] Mescheder L, Geiger A, Nowozin S. Which training methods for GANs do actually converge? In: International conference on machine learning. PMLR; 2018. p. 3481-90.

[6] Tian C, Xu Y, Fei L, Yan K. Deep learning for image denoising: A survey. In: International Conference on Genetic and Evolutionary Com-

puting. Springer; 2018. p. 563-72.

[7] Elad M, Kawar B, Vaksman G. Image denoising: The deep learning revolution and beyond—a survey paper. SIAM Journal on Imaging Sciences. 2023;16(3):1594-654.

[8] Bhat HA, Khanday FA, Kaushik BK, Bashir F, Shah KA. Quantum computing: fundamentals, implementations and applications. IEEE Open Journal of Nanotechnology. 2022;3:61-77.

[9] Desdentado E, Calero C, Moraga MÁ, García F. Quantum computing software solutions, technologies, evaluation and limitations: a systematic mapping study: E. Desdentado et al. Computing. 2025;107(5):110.

[10] Tran LD, Nguyen SM, Arai M. GAN-based noise model for denoising real images. In: Asian Conference on Computer Vision. Springer; 2020. p. 560-72.

[11] Ho J, Jain A, Abbeel P. Denoising diffusion probabilistic models. Advances in neural information processing systems. 2020;33:6840-51.

[12] Ronneberger O, Fischer P, Brox T. U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical image computing and computer-assisted intervention. Springer; 2015. p. 234-41.

[13] Zhang J, Niu Y, Shangguan Z, Gong W, Cheng Y. A novel denoising method for CT images based on U-net and multi-attention. Computers in Biology and Medicine. 2023;152:106387.

[14] Schleich P, Calderón LM, Sun C, Bagherimehrab M, Aldossary A, Kottmann JS, et al. Quantum Computing for Quantum Chemistry. American Chemical Society; 2025.

[15] Cong I, Choi S, Lukin MD. Quantum convolutional neural networks. Nature Physics. 2019;15(12):1273-8.

[16] Kölle M, Stenzel G, Stein J, Zielinski S, Ommer B, Linnhoff-Popien C. Quantum denoising diffusion models. In: 2024 IEEE international conference on quantum software (QSW). IEEE; 2024. p. 88-98.

[17] Shrestha S, Rai D, Adhikari NB. Enhancing Medical Image Security Leveraging Quantum Key Distribution. IET Quantum Communication. 2026;7(1):e70040.

[18] Nandal P, Pahal S, Upadhyay GM. Image denoising using quantum deep convolutional generative adversarial network for medical images. International Journal of Computational Intelligence Systems. 2025;18(1):190.

[19] He X, Gong L, Zhou N. Image denoising with hybrid classical-quantum convolutional neural network: X. he et al. Memetic Computing. 2025;17(2):23.

[20] Xu L, Zhao T, Li H. HiQC: hybrid quantum classical deep learning with parameter efficient quantum modules. Quantum Machine Intelligence. 2025;7(2):97.

[21] Arrasmith AT, Babbush R, Benjamin SC, Endo S, Fujii K, McClean JR, et al. Variational quantum algorithms. Nature Reviews Physics. 2021;3(9):625-44.

[22] Larocca M, Thanasilp S, Wang S, Sharma K, Biamonte J, Coles PJ, et al. Barren plateaus in variational quantum computing. Nature Reviews Physics. 2025;7(4):174-89.

[23] Dutta S, Basarab A, Georgeot B, Kouamé D. A novel image denoising algorithm using concepts of quantum many-body theory. Signal processing. 2022;201:108690.

[24] Leuschner J, Schmidt M, Baguer DO, Maass P. LoDoPaB-CT, a benchmark dataset for low-dose computed tomography reconstruction. Scientific Data. 2021;8(1):109.

[25] Ulyanov D, Vedaldi A, Lempitsky V. Instance normalization: The missing ingredient for fast stylization. arXiv preprint arXiv:160708022. 2016.

[26] Azad R, Aghdam EK, Rauland A, Jia Y, Avval AH, Bozorgpour A, et al. Medical image segmentation review: The success of u-net. IEEE Transactions on Pattern Analysis and Machine Intelligence. 2024;46(12):10076-95.

[27] Bergholm V, Izaac J, Schuld M, Gogolin C, Ahmed S, Ajith V, et al. Pennylane: Automatic differentiation of hybrid quantum-classical computations. arXiv preprint arXiv:181104968. 2018.

[28] McClean JR, Boixo S, Smelyanskiy VN, Babbush R, Neven H. Barren plateaus in quantum neural network training landscapes. Nature communications. 2018;9(1):4812.

[29] Yan Q, Wang W. DCGANs for image super-resolution, denoising and debluring. Advances in neural information processing systems. 2017;30:487-95.

[30] Simonyan K, Zisserman A. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:14091556. 2014.

[31] Johnson J, Alahi A, Fei-Fei L. Perceptual losses for real-time style transfer and super-resolution. In: European conference on computer vision. Springer; 2016. p. 694-711.

[32] Zhang B, Xu P, Chen X, Zhuang Q. Generative quantum machine learning via denoising diffusion probabilistic models. Physical Review Letters. 2024;132(10):100602.

[33] Pokharel B, Pandey DS, Sapkota A, Yadav B, Gurung V, Adhikari MP, et al. A comparative study of state-of-the-art deep learning models for semantic segmentation of pores in scanning electron microscope images of activated carbon. IEEE Access. 2024;12:50217-43.