Libyan Open University Journal of Applied Sciences (LOUJAS)
Volume 2 - Issue 2 - 2026 - Pages 87-96

A Comparative Analysis of Enhanced Residual U-Net and DnCNN Architectures for Blind Medical Image Denoising

Farag J. Zbeda[1], Abdurahman E. Saleh[2] , Salem A. Salem[3]
[1]Computer Engineering Department, Azzaytuna University, Tarhuna, Libya.
[2]Computer Science Department, Higher Institute Awlad Ali, Tarhuna, Libya.
[3]General Science Department, Higher Institute Awlad Ali, Tarhuna, Libya.
***Email (for reference researcher):** fzb_79@azu.edu.ly

تحليل مقارن لأداء معماريتي U-Net المتبقية المحسّنة و DnCNN في إزالة الضوضاء العمياء من الصور الطبية

فرج جمعة زبيدة[1*]، عبدالرحمن اشتيوي صالح[2]، سالم علي سالم[3]
[1] قسم هندسة الحاسب الألي، كلية الهندسة، جامعة الزيتونة، ترهونة، ليبيا.
[2] قسم علوم الحاسب، المعهد العالي، أولاد علي، ترهونة، ليبيا.
[3] القسم العام، المعهد العالي، أولاد علي، ترهونة، ليبيا.

Received: 09-05-2026; Accepted: 10-07--2026; Published: 15-07-2026

Abstract

High resolution CT imaging plays a sigificant role in today's diagnostic practices; yet, the images may be degraded by quantum noise and acquisition artifacts, which hide important anatomical features. Although there have been many studies on denoising using deep learning methods including DnCNN and U-Net models showing outstanding performance, the traditional approaches lack the ability to deal with different noise conditions and preserve fine structures in blind denoising. This paper compares two modified versions of deep learning networks: ER-U-Net and ER-DnCNN.

For the purposes of architectural fairness, both models employ the same optimization techniques, namely optimized local residual blocks, dynamic stochastic noise augmentation pipeline with σ noise values between 0.05 and 0.50, and a combined loss function based on the use of Mean Squared Error (MSE) (L2 loss) and Mean Absolute Error (MAE) (L1 loss) to enhance numerical and structural preservation of images. The experimental analysis of the two models conducted on a publicly available brain CT dataset has shown that both of them demonstrate strong capabilities of blind denoising at all noise levels. Specifically, the ER-U-Net outperforms in PSNR and SSIM scores, especially in conditions of severe noise, due to the presence of multi-scale feature extraction and skip connection techniques in its architecture. On the other hand, the ER-DnCNN proves to be a good alternative of high quality with simple architecture for medical imaging.

Keywords: Blind Image Denoising, U-Net, DnCNN, Internal Residual Learning, Hybrid Loss Function.

الملخص

التصوير المقطعي المحوسب (Computed Tomography) عالي الجودة يعد جزءاً أساسياً في التشخيص السريري الحالي، ولكن وجود الضوضاء الكمية (Quantum Noise) والتشوهات الناجمة من عملية الاستقبال قد يؤدي إلى طمس بعض التفاصيل التشريحية الدقيقة ذات الصلة بعملية التشخيص. وعلى الرغم من قدرة النماذج القائمة على التعلم العميق لإزالة الضوضاء، وخصوصاً معماريتا DnCNN وU-Net، في استعادة جودة الصورة، فإن استخدامات النماذج التقليدية عادةً ما تواجه تحديات في التعامل مع مستويات متفاوتة من الضوضاء، كما أنها قد لا تحمي البنيات التشريحية الدقيقة أثناء إزالة الضوضاء العمياء (Blind Denoising).

هذه الدراسة تقترح مقارنة بين اثنين من النماذج المستخدمة في التعلم العميق هما (ER-U-Net) و (ER-DnCNN). ولضمان مقارنة معمارية عادلة بين النموذجين، تم تزويدهما بعدد من استراتيجيات التحسين، مثل الكتل المحلية المحسَّنة (Optimized Local Residual Blocks)، والمعالجة الديناميكية لإزالة الضوضاء العشوائية، حيث تتراوح مستويات الضوضاء بين σ = 0.05 و σ = 0.50، بالإضافة إلى استخدام دالة خسارة هجينة والتي تدمج دالة خسارة متوسط مربع الخطأ (MSE) مع دالة خسارة متوسط الخطأ المطلق (MAE).

وقد أظهرت نتائج التقييم باستخدام قاعدة بيانات عامة لصور الأشعة المقطعية للدماغ، أن كل من المعماريتين حققتا أداءً قويًا في إزالة الضوضاء العمياء لجميع مستويات الضوضاء المستخدمة في هذه البحث. كما أن معمارية ER-U-Net حققت قيمًا أعلى لكلٍّ من نسبة الإشارة إلى الضوضاء القصوى (PSNR) ومؤشر التشابه البنيوي(SSIM) ، ولا سيما في ظروف الضوضاء الشديدة، وذلك بفضل قدرتها على استخراج السمات متعددة المقاييس وآلية الوصلات التخطية (Skip Connections). في المقابل، قدّمت قدمت ER-DnCNNجودة استعادة تنافسية مع الحفاظ على بنية معمارية أبسط، مما يجعلها خيارًا مناسبًا للتطبيقات الطبية ذات الموارد الحاسوبية المحدودة أو التي تتطلب معالجة آنية(Real-Time) .

الكلمات المفتاحية: إزالة الضوضاء العمياء للصور، شبكة U-Net، شبكة DnCNN، التعلم المتبقي الداخلي، دالة الخسارة الهجينة.

Introduction

Computed Tomography (CT) remains an essential diagnostic modality in modern medicine, providing cross-sectional anatomical detail crucial for oncology, cardiology, and emergency trauma assessment [1]. However, it is necessary to preserve the best image fidelity for a wide range of scanning protocols to obtain accurate diagnostic results [2]. The main problem in high resolution CT imaging is the intrinsic degradation of image quality due to photon starvation, electronic noise and systemic artefacts. This appears as highly structured, non-uniform noise and streak artefacts that obscure subtle tissue lesions, confound automated segmentation algorithms and reduce diagnostic confidence [3]. Therefore, the development of effective post-processing image denoising algorithms is essential to alleviate the discrepancy between practical hardware acquisition limitations and high-quality diagnostic reconstruction

In the past, the main methods used in medical image restoration include classical analytical methods such as Non-Local Means (NLM), Block-Matching and 3D Filtering (BM3D), and Total Variation (TV) regularization [4]. Though these techniques are mathematically based, they are computationally expensive and are often subject to over-smoothing artifacts that eliminate important high-frequency clinical features such as fine bone trabeculae and soft-tissue boundaries. In addition, classical algorithms are highly sensitive to manual tuning of the hyperparameters and lack the ability to generalize across different scanner geometries and varying noise levels [5].

In recent years, Deep Learning (DL) based on Convolutional Neural Networks (CNNs) has revolutionized medical image denoising by shifting the paradigm toward data-driven feature extraction [6]. Among deep learning approaches for medical image denoising, two representative architectural philosophies have emerged:

1. The Plain Residual Paradigm (DnCNN): Exemplified by the Denoising Convolutional Neural Network, this approach preserves the spatial resolution of the image across all hidden layers, relying on deep residual learning to separate latent clean anatomical signals from background noise [7].

1.The Encoder-Decoder Paradigm (U-Net): Characterized by a symmetric contracting and expanding path interconnected by skip connections, this architecture progressively downsamples feature maps to capture multi-scale global context before upsampling to reconstruct the image [8].

Despite their individual successes, standard implementations of these networks possess critical limitations when applied directly to clinical settings. Traditional training paradigms are typically "non-blind," optimized for a single, fixed noise level (σ), which causes models to fail when encountering unpredictable noise profiles in real-world clinical operations [9]. Additionally, relying solely on standard Mean Squared Error (L2 loss) often yields structurally blurry outputs due to its statistical focus on pixel averages, neglecting structural similarity [10].

To address these limitations, this paper introduces a unified evaluation framework for comparing two enhanced deep learning architectures that incorporate identical optimization strategies while preserving their distinctive structural characteristics. We apply identical, state-of-the-art enhancements to both a U-Net and a DnCNN architecture. Specifically, we integrate internal residual blocks to stabilize deep gradient propagation, implement a dynamic stochastic noise augmentation pipeline during training to enforce "blind" multi-noise robust inference, and employ an edge-aware hybrid (L1 + L2) loss function.

The main contributions of this study are summarized as follows:

1. A fair comparative framework is established for evaluating two representative deep learning paradigms, namely the Enhanced Residual U-Net (ER-U-Net) and the Enhanced Residual DnCNN (ER-DnCNN), under identical training conditions and optimization settings.
2. A unified enhancement strategy is applied to both architectures through the integration of optimized local residual blocks, a dynamic stochastic noise augmentation pipeline, and a hybrid loss function that combines MSE and MAE, thereby ensuring that performance differences primarily reflect architectural characteristics rather than training configurations.
3. A comprehensive blind denoising evaluation is performed over multiple Gaussian noise levels using publicly available brain CT images, enabling an objective assessment of restoration performance under varying degradation conditions.

2. Methodology

2.1 Problem Formulation and Dynamic Noise Augmentation

Let $X \in \mathbb{R}^{H*W*C}$ represent a clean, latent medical image sampled from the target domain, where H, W, and C denote the height, width, and channel dimensions, respectively. In a clinical environment, the observed degraded image y is corrupted by an additive noise artifact [7]. To enforce a blind denoising paradigm capable of generalizing across fluctuating scanner parameters without prior noise-level knowledge, we implement a dynamic stochastic noise simulation pipeline during the operational training phase.
For every training sample iteration, a unique noise standard deviation σ is drawn uniformly from a continuous distribution bounded by clinical extremes:

$$\sigma \sim \mathrm{U}\,(\,\sigma_{min}\,,\sigma_{max}\,) \qquad 1$$

Where σ_{min}= 0.05 and σ_{max} = 0.5.
Sampling the noise level independently for every training iteration exposes the network to continuously varying degradation conditions, thereby improving its ability to generalize to previously unseen noise levels during blind denoising.
The additive white Gaussian noise (AWGN) matrix n is stochastically generated such that [11]:

$$n \sim N(\,0, \sigma^2 . I) \qquad 2$$

Where I represents the identity tensor.
The simulated observation tensor y passed to the networks is formally expressed via a clipping function Ω to preserve the localized physical bounds of the digitized clean image [9]:

$$y = \,\Omega\,(\,x + n\,) \qquad 3$$

2.2 Global Residual Learning Strategy

Instead of directly estimating the clean image, both architectures are trained to predict the underlying noise residual. Residual learning considerably simplifies the optimization problem because the statistical complexity of the noise component is substantially lower than that of the anatomical structures contained in medical images. Consequently, the networks converge faster while preserving diagnostically relevant image details [12].

Let $F(\mathrm{y}\,;\,\theta\,)$represent the forward mapping function of the deep network parameterized by weights and biases θ. The network is specifically constrained to isolate and predict the latent noise field $\hat{n}$:

$$\hat{n} \approx F(\mathrm{y}\,;\,\theta\,) \qquad 4$$

Consequently, the final clean restored image estimation $\hat{x}$is structurally recovered during the inference phase by subtracting the isolated residual noise from the corrupted input [7]:

$$\hat{x} = y - F(\mathrm{y}\,;\,\theta\,) \qquad 5$$

This global residual framing minimizes optimization bottlenecks, allowing gradients to propagate more effectively since the background noise patterns typically present lower structural complexity than the underlying human anatomy [11].

2.3 Hybrid Loss Function

Standard image restoration methods mostly use the Mean Squared Error (MSE). The problem with MSE optimization is that it makes the overall Peak Signal-to-Noise Ratio (PSNR) better, It does not really punish structural dislocations enough. This leads to a loss of high-frequency textures and the output looks blurry[13]. To fix this issue we suggest a Hybrid Loss function (Lhybrid) that keeps the anatomy and edges intact while balancing structural alignment and high-frequency boundary preservation.

The part of the Hybrid Loss function that calculates the Mean Squared Error, for each pixel is defined as:

$$L_{MSE} = \frac{1}{N}\sum_{i=1}^{N} \|(y_i - F(\mathrm{y}\,;\,\theta\,) - x_i)\|_2^2 \qquad 6$$

Where N is the batch size, yi is the input noisy image, xi is the ground truth clean image, and F(y ; θ) is the noise residual map output by the network.

While the MSE (L_2) loss function reduces the pixel-wise reconstruction error efficiently and results in increased PSNR values, it is more sensitive to outliers than the L_1 loss and generates overly smooth outputs. On the other hand, the MAE (L_1) loss is robust to outliers and maintains sharp anatomical edges and fine details. This hybrid loss function enhances both numerical and visual aspects of image reconstruction [13].

To prevent over-smoothing at sharp tissue interfaces (such as bone-to-soft-tissue transitions), we integrate a Mean Absolute Error (MAE) (L_1 loss) regularizer, which provides a constant gradient for small errors and preserves sharp edges:

$$L_{MAE} = \frac{1}{N}\sum_{i=1}^{N} \|(y_i - F(\mathrm{y}\,;\,\theta\,) - x_i)\|_1 \qquad 7$$

The total structural objective function minimized during the network optimization phase is formulated as a convex combination:

$$L_{hybrid} = \alpha \,.\, L_{MSE} + (1 - \alpha).\, L_{MAE} \qquad 8$$

where $\alpha \in [0,1]$ controls the relative contribution of the MSE and MAE terms. A larger value of α emphasizes numerical reconstruction accuracy, whereas a smaller value increases the influence of edge preservation [13].

In this work α =0.8 was selected empirically to provide an effective balance between quantitative accuracy and structural fidelity.

2.4 Network Architectures

To analyze the structural trade-off between spatial-resolution preservation and deep multi-scale context aggregation, we evaluate two structurally distinct deep topologies under the exact same optimization constraints. Both networks are constructed utilizing an elementary Local Residual Block (LRB) to facilitate smooth, deep gradient flow [12].

The operational mapping of a Local Residual Block receiving an input tensor z_l is defined as:

$$z_{l+1} = z_l + W_s\left(C_2\left(BN\left(ReLU\left(C_1\left(BN(ReLU(z_l))\right)\right)\right)\right)\right) \quad 9$$

where C_1 and C_2 denote 3x3 2D convolutional layers, BN represents Batch Normalization, ReLU is the Rectified Linear Unit activation function, and W_s represents a 1x1 linear projection shortcut used to match channel dimensions when necessary [10].

The overall architectures of the proposed ER-DnCNN and ER-U-Net are illustrated in Figures 1 and 2, respectively. Both models incorporate identical optimization strategies and local residual blocks, thereby enabling a controlled comparison in which the primary difference lies in the architectural design rather than the optimization procedure.

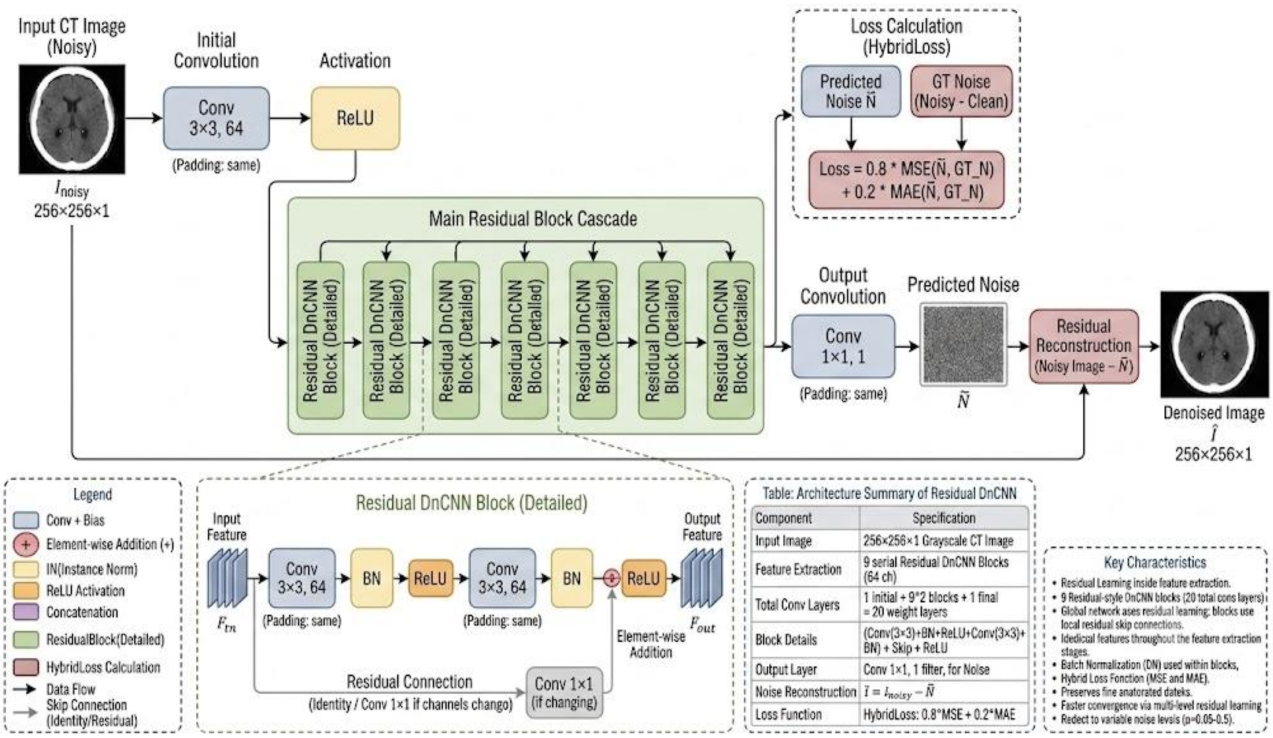

Figure 1: Architecture of ER-DnCNN

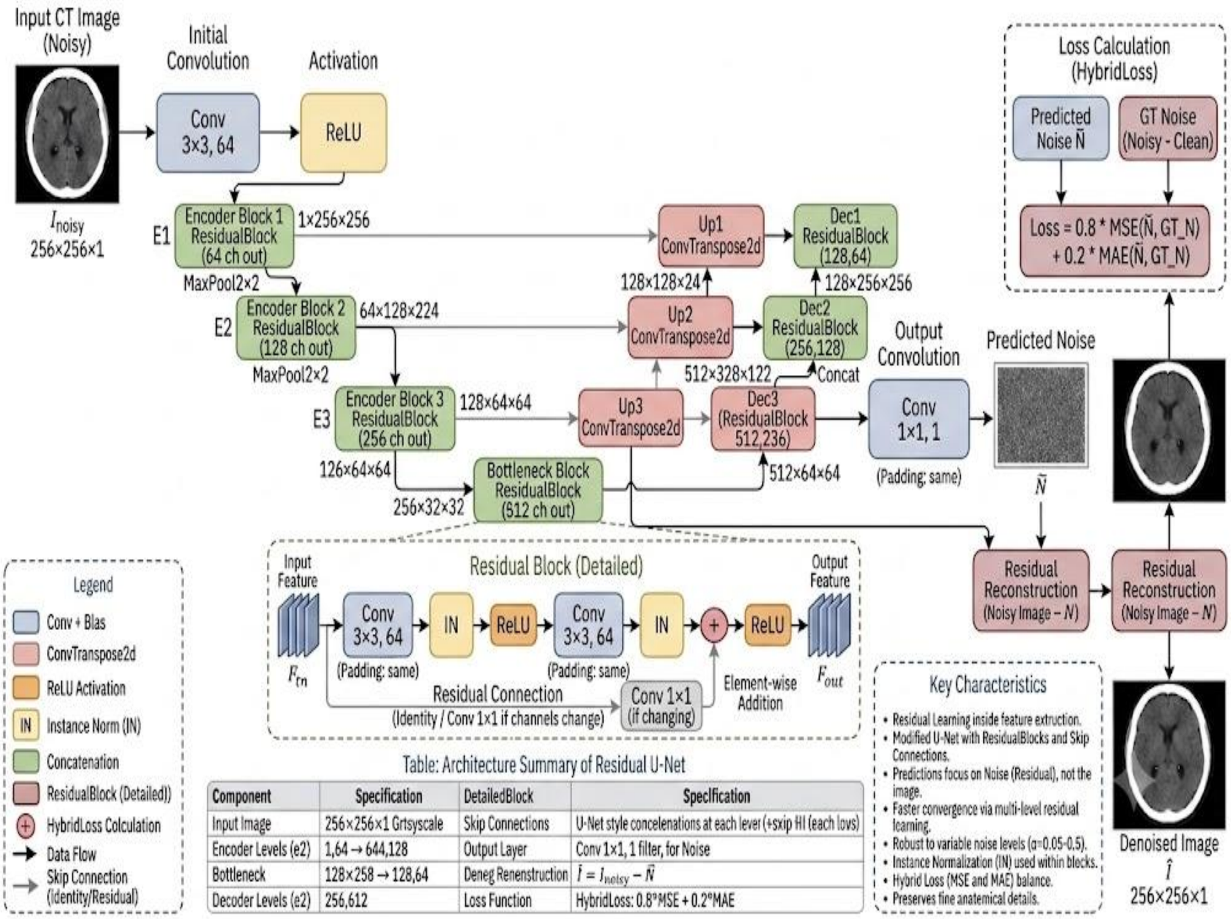

Figure 2: Architecture of ER-U-Net

3. Simulation Setup

3.1 Dataset Characteristics and Preprocessing

The experimental evaluation was conducted using a publicly available brain CT image dataset obtained from the Kaggle platform [14]. The dataset consists of 4,427 individual CT image slices without explicit patient identifiers. Therefore, the dataset partitioning was performed at the image level. Prior to training, all images were resized to 256x256 pixels and normalized to the intensity range [0,1] to improve numerical stability and accelerate network convergence.

To simulate realistic blind denoising conditions, additive white Gaussian noise (AWGN) was synthetically injected into each clean image during data preparation. The noise standard deviation was randomly selected from the $\sigma \in [0.05, 0.10, 0.20, 0.30, 0.40, 0.50]$, enabling the proposed models to learn a generalized mapping across multiple degradation levels rather than overfitting to a single predefined noise intensity. This dynamic noise generation strategy improves the robustness and generalization capability of the trained networks when encountering previously unseen noise conditions.

The dataset was partitioned into 80% training, 10% validation, and 10% testing. The validation subset was used to monitor the optimization process and prevent overfitting, whereas the independent test subset was reserved exclusively for the final quantitative evaluation to ensure an unbiased performance assessment.

3.2 Network Optimization and Hyperparameters

To ensure a fair comparison between the proposed architectures, identical optimization settings were employed during training. All hyperparameters were kept constant across both networks so that any observed performance differences could be attributed to architectural characteristics rather than differences in the optimization procedure. The complete training configuration is summarized in table 1.

Hyperparameter	Operational Value
Input image size	256 x 256
Number of layers	20
Initial learning rate	1 x 10-3
learning rate scheduler	ReduceLROnPlateau
LR reduction factor	0.5
Scheduler patience	5
Batch size	8
Epochs	100
Loss function	Hybrid (MSE+MAE)
Optimizer	Adam
GPU	NVIDIA Tesla P100
Framework	TensorFlow/Keras

Table 1: Hyper-parameter Optimization Configuration

During training, the network parameters were optimized using the Adam optimizer owing to its adaptive learning capability and fast convergence characteristics. The learning rate was automatically adjusted using the ReduceLROnPlateau scheduler, which reduced the learning rate by a factor of 0.5 whenever the validation loss failed to improve for five consecutive epochs. This adaptive strategy stabilized the optimization process, accelerated convergence, and reduced the risk of oscillatory training behavior near local minima.

3.3 Evaluation Metrics

The denoising performance of the proposed architectures was quantitatively evaluated using the Peak Signal-to-Noise Ratio (PSNR) and the Structural Similarity Index Measure (SSIM), which are widely adopted for assessing image restoration quality.

$$\mathrm{PSNR} = 10\log_{10}\left(\frac{\mathrm{MAXi}^2}{\mathrm{MSE}}\right) \qquad 10$$

Where MAXi is the maximum possible pixel value of the image. Higher PSNR values indicate lower reconstruction error and improved numerical fidelity.

$$\mathrm{SSIM}(x, y) = \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2+\mu_y^2 + C_1)(\sigma_x^2+\sigma_y^2 + C_2)} \qquad 11$$

Where μ_x and μ_y are the local mean intensities, σ_x^2 and σ_y^2 are the local variances, σ_{xy} is the covariance between images x and y, and c_1 and c_2 are stability constants. Unlike PSNR, SSIM evaluates luminance, contrast, and structural similarity between images, making it particularly suitable for assessing perceptual image quality and preservation of anatomical structures in medical imaging applications [15].
To ensure a fair comparison, both ER-DnCNN and ER-U-Net were trained and evaluated using the identical dataset, optimization strategy, hyperparameter configuration, and evaluation protocol**.**

4. Simulation Results and Discussion

All experiments were performed using the TensorFlow/Keras deep learning framework on an NVIDIA Tesla P100 GPU to efficiently train and evaluate the proposed models. The restoration performance of the Enhanced Residual DnCNN and Enhanced Residual U-Net is systematically measured using two standard mathematical figures of merit: Peak Signal-to-Noise Ratio (PSNR) and the Structural Similarity Index Measure (SSIM). Testing is evaluated systematically across six distinct additive white Gaussian noise regimes $\sigma \in [0.05, 0.1, 0.2, 0.3, 0.4, 0.5]$ The quantitative performance comparison between the Enhanced Residual DnCNN and the Enhanced Residual U-Net is summarized in Table 2. Both architectures demonstrate robust blind denoising capability over the complete range of evaluated Gaussian noise levels (σ= 0.05 - 0.5), consistently producing high PSNR and SSIM values despite progressively increasing image degradation.

Libyan Open University Journal of Applied Sciences (LOUJAS)
Volume 2 - Issue 2 - 2026 - Pages 87-96

Method	Evaluation Metrics	Noise Level					
		0.05	0.1	0.2	0.3	0.4	0.5
ER-U-Net	PSNR	36.503	33.482	30.662	28.924	27.618	26.543
	SSIM	0.9637	0.9352	0.9043	0.8849	0.8698	0.8564
ER- DnCNN	PSNR	34.843	33.24	30.353	28.512	27.120	25.889
	SSIM	0.9591	0.9320	0.8991	0.8768	0.8586	0.8389

Table 2: Quantitative Performance Comparison

With the smallest noise (σ = 0.05), the ER-U-Net provides a PSNR value of 36.503 dB and an SSIM value of 0.9637, which is better than the PSNR of 34.843 dB and the SSIM of 0.9591 of the ER-DnCNN. Therefore, the encoder-decoder-based architecture can be considered to preserve fine anatomical structures in images in which there is no much corruption to remove.

While increasing the corruption level, both the tested architectures show the desired decrease in their PSNR and SSIM metrics due to growing corruption of images which becomes harder to restore. Still, in spite of the gradual increase in noise, the ER-U-Net keeps better quantitative results. For instance, for the noise level equal to 0.3, the ER-U-Net demonstrates 28.924 dB and 0.8849 SSIM, and the ER-DnCNN shows 28.512 dB and 0.8768 SSIM.

The gap in performance becomes apparent with increased severity in the degradation process. At the highest level of tested noise intensity (σ= 0.5), ER-U-Net has a PSNR value of 26.543 dB and SSIM of 0.8564 compared to ER-DnCNN's PSNR of 25.889 dB and SSIM of 0.8389. This indicates that the encoder-decoder structure has the advantage of having multiscale feature extraction and skip connections, enabling it to be capable of detecting both large context and small anatomical structures at once.

On the other hand, ER-DnCNN exhibits consistent and competitive performance levels across all noise intensities regardless of its relatively simple structure.

By preserving spatial resolution throughout the network and learning the residual noise directly, the model effectively restores image quality.

Overall, the simulation results, as shown in table 2 and figures 3, 4 and 5, indicate that ER-U-Net provides the highest restoration accuracy under all tested noise conditions. The consistent superiority of ER-U-Net in both PSNR and SSIM demonstrates the effectiveness of multi-scale contextual feature learning for blind medical image denoising, while the strong performance of ER-DnCNN confirms that residual learning remains a highly effective strategy for practical clinical image restoration.

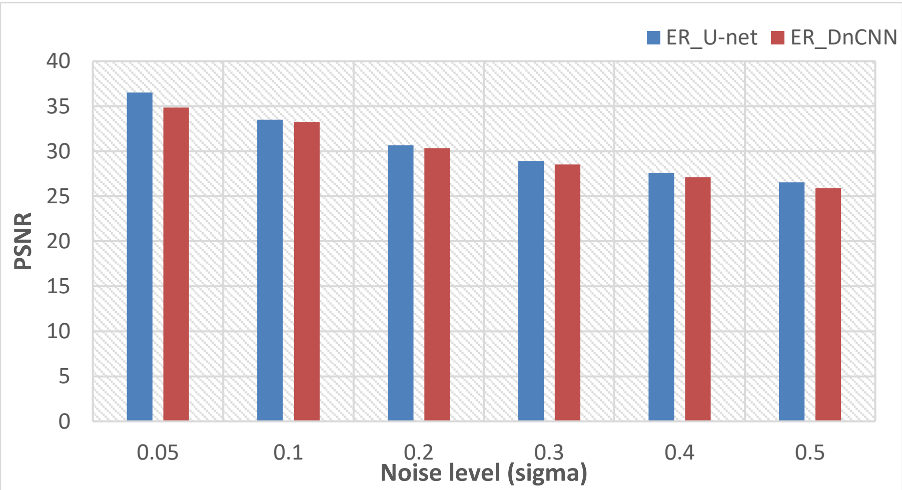

Figure 3. Comparison of PSNR values obtained by ER-DnCNN and ER-U-Net under different Gaussian noise levels.

Libyan Open University Journal of Applied Sciences (LOUJAS)
Volume 2 - Issue 2 - 2026 - Pages 87-96

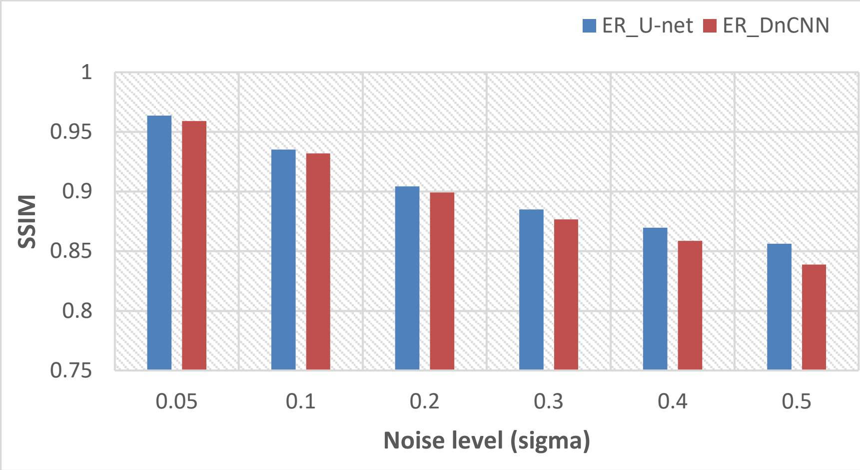

Figure 4. Comparison of SSIM values obtained by ER-DnCNN and ER-U-Net under different Gaussian noise levels.

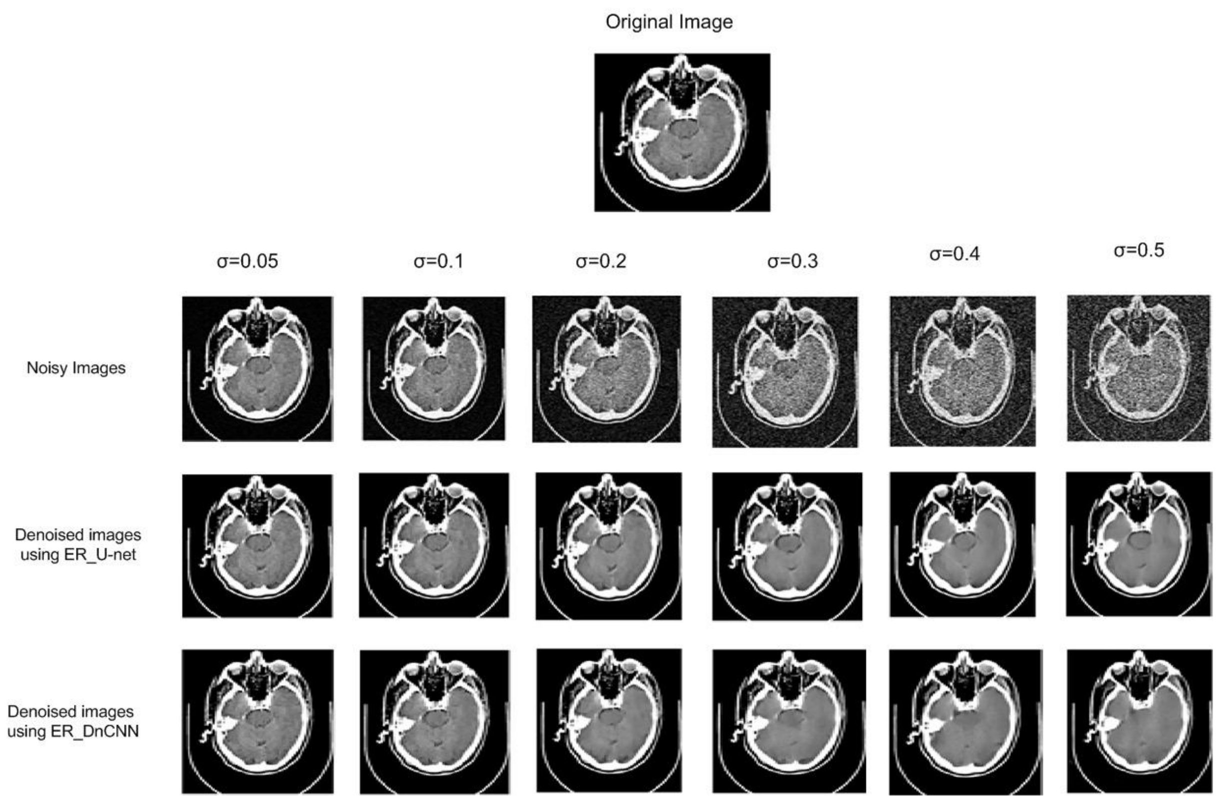

Figure 5: Sample of Visual Results

5. Conclusion

This study provided a rigorous comparative evaluation between two prominent deep learning paradigms-the plain residual spatial-preservation architecture (DnCNN) and the multi-scale encoder-decoder structure (U-Net)-both incorporating identical enhancement strategies for blind medical image denoising. The integration of local residual blocks, a stochastic blind training pipeline, and an edge-aware hybrid (L_1+L_2) loss function enabled a fair comparison of the two architectural paradigms under identical optimization conditions.

From the simulation results, it can be concluded that the ER-U-Net network is superior in terms of higher PSNR and SSIM values compared to all others at different levels of noise added. This clearly depicts its effectiveness in preserving anatomical structures in both mild and extreme conditions of noise presence. On the other hand, ER-DnCNN shows high restoration ability by means of a simpler architecture, which makes it applicable in clinical settings. Future research can explore this comparison using 3D CT volumes and MRIs.

References

1. Kalender, W. A. (2015). Computed tomography: Principles, design, artifacts, and recent advances (2nd ed.). SPIE Press.
2. McCollough, C. H., Bruesewitz, M. R., & Kofler, J. M., Jr. (2006). CT dose reduction and dose management tools: Overview of available options. Radiographics, 26(2), 503–512.
3. Ye, J. C., Han, Y., & Cha, E. (2018). Deep learning in medical imaging: General overview. Annual Review of Biomedical Engineering, 20, 221–248.
4. Buades, A., Coll, B., & Morel, J.-M. (2005). A non-local algorithm for image denoising. In Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (Vol. 2, pp. 60–65). .
5. Singh, S. K., & Gupta, S. (2021). A survey on image denoising techniques, challenges and future directions. Signal Processing, 184, Article 108051.
6. Litjens, G., Kooi, T., Bejnordi, B. E., Setio, A. A. A., Ciompi, F., Ghafoorian, M., Van der Laak, J. A. W. M., Van Ginneken, B., & Sánchez, C. I. (2017). A survey on deep learning in medical image analysis. Medical Image Analysis, 42, 60–88.
7. Zhang, K., Zuo, W., Chen, Y., Meng, D., & Zhang, L. (2017). Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising. IEEE Transactions on Image Processing, 26(7), 3142–3155.
8. Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. In N. Navab, J. Hornegger, W. M. Wells, & A. F. Frangi (Eds.), Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015 (Vol. 9351, pp. 234–241). Springer.
9. Chen, H., Zhang, Y., Kalra, M. K., Lin, F., Chen, Y., Liao, P., Zhou, J., & Wang, G. (2017). Low-dose CT with a residual encoder–decoder convolutional neural network. IEEE Transactions on Medical Imaging, 36(12), 2524–2535.
10. Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. In F. Bach & D. Blei (Eds.), Proceedings of the 32nd International Conference on Machine Learning (Vol. 37, pp. 448–456). PMLR.
11. Zhang, K., Zuo, W., & Zhang, L. (2018). FFDNet: Toward a fast and flexible solution for CNN-based image denoising. IEEE Transactions on Image Processing, 27(9), 4608–4622.
12. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 770–778.
13. Zhao, H., Gallo, O., Frosio, I., & Kautz, J. (2017). Loss functions for image restoration with neural networks. IEEE Transactions on Computational Imaging, 3(1), 47–57.
14. Koç, U., Akçapınar Sezer, E., Özkaya, Y. A., et al. (2022). Artificial intelligence in healthcare competition (TEKNOFEST-2021): Stroke dataset. *Eurasian Journal of Medicine*, *54*(3), 248–258.
15. Huynh-Thu, Q., & Ghanbari, M. (2008). Scope of validity of PSNR in image/video quality assessment. Electronics Letters, 44(13), 800–801.