iJRASET
International Journal For Research in
Applied Science and Engineering Technology

INTERNATIONAL JOURNAL
FOR RESEARCH
IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** II **Month of publication:** February 2026

DOI: https://doi.org/10.22214/ijraset.2026.77691

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue II Feb 2026- Available at www.ijraset.com

Chest X-Ray Image Denoising Using Edge-Enhanced U-Net and Generative Adversarial Networks

Danish Wilson[1], Alsa Elizabeth Sali[2], Anna S P[3], Athulya M[3], Aneesa H A[5]
Department of Computer Science and Engineering, St. Thomas Institute for Science and Technology, Trivandrum, India

Abstract: Chest X-ray image quality is crucial for accurate medical diagnosis, yet acquired images are often degraded by noise, which can obscure important anatomical details. This work presents a deep learning–based denoising framework for chest X-ray images using a customized U-Net-like architecture combined with adversarial learning. Normal chest X-ray images from the Kaggle Chest X-ray Pneumonia dataset are preprocessed, resized, and synthetically corrupted with Gaussian noise to generate paired noisy and clean training samples. The proposed generator network, termed SharpXRLikeUNet, integrates Laplacian filtering to enhance edge preservation during denoising. Initially, the model is trained using a composite loss function consisting of L1 loss, structural similarity index measure (SSIM) loss, and perceptual loss derived from a pretrained VGG16 network. To further improve visual realism, a generative adversarial network (GAN) framework is employed by introducing a patch-based discriminator, and an adversarial loss component is incorporated into the generator objective. Experimental results demonstrate effective noise reduction while preserving fine structural details in chest X-ray images. The trained models are saved, and the generator is exported to ONNX format to support efficient deployment and cross-platform inference.

Keywords: Chest X-ray denoising, Deep learning, U-Net architecture, Generative Adversarial Network (GAN), Medical image enhancement, Gaussian noise, Perceptual loss, Structural Similarity Index (SSIM), Patch discriminator, ONNX deployment.

I. INTRODUCTION

Chest X-ray imaging is one of the most widely used diagnostic tools in clinical practice due to its low cost, rapid acquisition, and effectiveness in identifying thoracic diseases such as pneumonia, tuberculosis, and lung abnormalities. However, the diagnostic reliability of chest X-ray images is often affected by various sources of noise introduced during image acquisition, transmission, or low-dose imaging protocols. Such noise can obscure fine anatomical structures, reduce image contrast, and adversely impact both manual interpretation by radiologists and automated computer-aided diagnosis (CAD) systems.

Traditional image denoising techniques, including spatial filtering, frequency-domain methods, and model-based approaches, have been extensively studied to mitigate noise in medical images. While these methods can suppress noise to a certain extent, they often suffer from over-smoothing, leading to loss of critical edge information and structural details that are essential for accurate diagnosis. In recent years, deep learning– based methods have demonstrated superior performance in image restoration tasks by learning complex, data-driven mappings between degraded and clean images.

Convolutional neural networks (CNNs), particularly encoder–decoder architectures such as U-Net, have become popular for medical image enhancement due to their ability to capture multi-scale contextual information while preserving spatial resolution through skip connections. Nevertheless, standard U-Net architectures optimized solely with pixel-wise loss functions may produce overly smooth outputs and fail to recover perceptually realistic textures. To address these limitations, perceptual loss functions and adversarial learning frameworks have been incorporated into image restoration pipelines to improve visual fidelity and structural realism.

In this work, we propose a deep learning–based chest X-ray denoising framework that combines a customized U-Net-like generator architecture with generative adversarial learning. The proposed generator, termed SharpXRLikeUNet, integrates Laplacian filtering to explicitly emphasize edge and high-frequency components during reconstruction. The model is first trained using a composite loss function comprising L1 loss, structural similarity index measure (SSIM) loss, and perceptual loss derived from a pretrained VGG16 network, enabling both pixel-level accuracy and perceptual consistency. Subsequently, a patch-based discriminator is introduced to form a generative adversarial network (GAN), encouraging the generator to produce denoised images that are indistinguishable from real clean X-ray images.

The effectiveness of the proposed approach is evaluated through qualitative visual analysis by comparing noisy input images with their denoised counterparts. The trained models are saved for reuse, and the generator is exported to the ONNX format to facilitate deployment across different platforms. The proposed framework demonstrates the potential of combining structural priors, perceptual learning, and adversarial training for enhancing the quality of chest X-ray images..

II. LITERATURE SURVEY

Medical image denoising has been extensively studied due to its importance in improving image quality and diagnostic accuracy. Early denoising approaches relied on traditional spatial-domain filters such as mean, median, Gaussian, and Wiener filters [1]. Although these techniques are computationally efficient, they tend to blur important anatomical structures and edges, making them unsuitable for medical imaging applications where fine details are critical. To address the shortcomings of basic filtering methods, transform-domain techniques such as wavelet-based denoising were introduced [2]. These methods perform multi-resolution decomposition of images and selectively suppress noise components. Non-local means (NLM) filtering further improved denoising performance by exploiting self-similarity within images [3]. Despite better edge preservation, such methods are highly sensitive to parameter selection and often struggle with varying noise distributions in real-world medical images. Sparse representation and dictionary learning techniques marked the transition toward data-driven denoising approaches [4]. By learning image priors from training data, these methods achieved improved adaptability compared to handcrafted filters. However, their reliance on iterative optimization and handcrafted feature extraction limited scalability and real-time applicability. With the rise of deep learning, convolutional neural networks (CNNs) have become the dominant approach for image denoising. Early CNN-based models demonstrated superior performance over classical methods by learning end-to-end mappings between noisy and clean images [5]. Encoder– decoder architectures, particularly U-Net, gained widespread adoption in medical image processing tasks due to their ability to capture multi-scale contextual information while preserving spatial resolution through skip connections [6]. Several studies have reported the effectiveness of U-Net-based architectures for medical image restoration tasks, including X-ray and CT image denoising [7]. However, CNN-based models trained solely with pixel-wise loss functions such as mean squared error (MSE) or L1 loss often produce overly smooth outputs that lack perceptual realism [8]. To mitigate this issue, perceptual loss functions based on pretrained networks such as VGG were introduced to enforce feature-level similarity between restored and ground truth images [9]. Structural similarity index measure (SSIM)– based losses have also been employed to better preserve luminance, contrast, and structural information in denoised images [10]. Generative Adversarial Networks (GANs) further advanced image denoising by incorporating adversarial learning, where a discriminator network encourages the generator to produce visually realistic outputs [11]. GAN-based denoising methods have shown improved texture recovery and sharper edges compared to non-adversarial CNN models [12]. Patch-based discriminators have been particularly effective in enforcing local realism and reducing blurring artifacts [13]. However, GAN training introduces stability challenges and requires careful balancing between reconstruction and adversarial losses. Although GAN-based approaches have been explored for general image denoising, their application to chest X-ray denoising remains relatively limited. Furthermore, existing works often lack explicit edge-enhancement mechanisms and rarely consider deployment-oriented aspects such as model portability. The proposed work addresses these gaps by integrating Laplacian-based edge enhancement within a U-Net-like generator architecture and combining perceptual and adversarial learning, while also enabling deployment through ONNX model export.

III. METHODOLOGY

The proposed chest X-ray denoising framework follows a multi-stage methodology that combines data preprocessing, deep convolutional learning, and adversarial training. The overall workflow consists of dataset preparation, generator network design and training, and adversarial refinement using a generative adversarial network (GAN). The complete methodology is described in the following subsections.

A. Dataset Preparation and Preprocessing

The dataset used in this work is derived from the Kaggle Chest X-ray Pneumonia dataset. Only normal chest X-ray images are selected to avoid pathological bias during denoising and to ensure consistent anatomical structures. All images are resized to a fixed resolution and normalized to standardize pixel intensity values, thereby facilitating stable training. To simulate real-world imaging noise, Gaussian noise is artificially added to the clean images. This process generates paired noisy–clean image samples, where the noisy images serve as inputs to the model and the corresponding clean images act as ground truth references. The dataset is randomly split into training and validation subsets to evaluate the generalization performance of the proposed approach.

Fig 1. An image in dataset

B. Noise Modeling Using Gaussian Distribution

Gaussian noise is chosen as it closely represents electronic noise commonly observed in medical imaging systems. Noise is introduced by sampling from a zero-mean Gaussian distribution with controlled variance. By varying the noise intensity, the model is trained to handle different levels of degradation, improving robustness and adaptability to real-world scenarios.

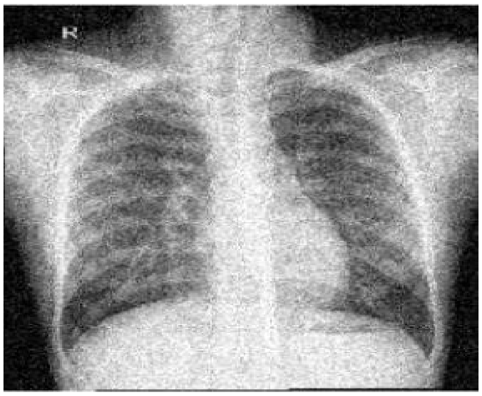

Fig 2. Xray after adding gaussian noise.

C. SharpXRLikeUNet Generator Architecture

The denoising model is based on a customized U-Net-like encoder–decoder architecture, referred to as SharpXRLikeUNet. The encoder path progressively extracts hierarchical features using convolutional layers and downsampling operations, while the decoder path reconstructs the denoised image using upsampling layers. Skip connections are employed between corresponding encoder and decoder stages to preserve spatial details and prevent information loss. To enhance structural sharpness, a Laplacian filtering mechanism is integrated into the architecture to emphasize high-frequency and edge-related features. This explicit edge-aware enhancement allows the network to retain fine anatomical boundaries that are typically lost during denoising.

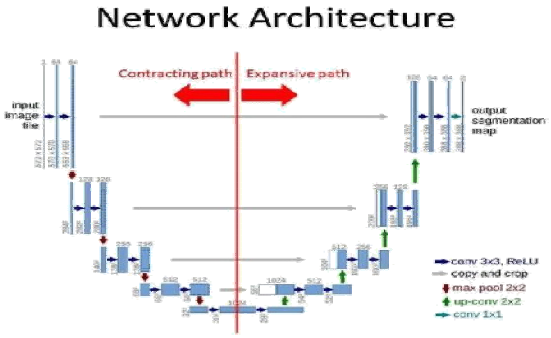

Fig 3. SharpXR like U-net architecture

D. Composite Loss

Function for Generator Training During the initial training phase, the generator is optimized using a composite loss function designed to balance reconstruction accuracy and perceptual quality. The total loss consists of three components: L1 loss, Structural Similarity Index Measure (SSIM) loss, and perceptual loss computed using feature maps extracted from a pretrained VGG16 network.

The L1 loss ensures pixel-level fidelity between denoised and clean images, SSIM loss preserves structural consistency, and perceptual loss enforces high-level feature similarity. The weighted combination of these losses enables the generator to produce visually sharp and structurally accurate outputs.

E. Adversarial Learning Using Patch Discriminator

To further improve the realism of the denoised images, a Generative Adversarial Network (GAN) framework is employed. A patch-based discriminator is introduced to distinguish between real clean X-ray images and generated denoised images at the local patch level. This approach encourages the generator to focus on fine-grained texture details rather than global appearance alone.

During adversarial training, the generator and discriminator are trained alternately. The generator loss is augmented with an adversarial loss term, while the discriminator is optimized using a binary classification objective. This adversarial learning strategy significantly enhances edge sharpness and reduces over-smoothing artifacts.

F. Model Training Strategy and Optimization

The training process is conducted in two stages. In the first stage, the generator is trained independently using the composite reconstruction loss to ensure stable convergence. In the second stage, adversarial training is introduced by jointly training the generator and discriminator. Optimization is performed using gradient-based methods with carefully selected learning rates to maintain training stability and prevent mode collapse.

G. Model Evaluation and Visualization

The performance of the proposed denoising framework is evaluated qualitatively through visual inspection. Noisy input images are compared with the corresponding denoised outputs to assess noise suppression and structural preservation. Visual results demonstrate effective removal of Gaussian noise while maintaining critical anatomical details in chest X-ray images.

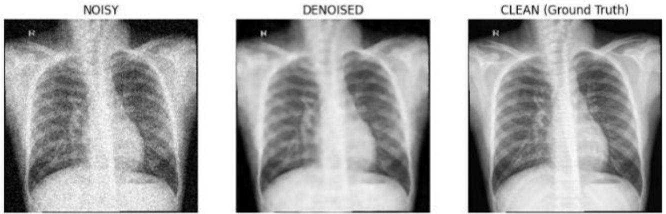

Fig 4. Performance of the model

H. Model Saving and ONNX Export

After training, the generator and discriminator models are saved for future inference and reuse. To facilitate deployment and cross-platform compatibility, the trained generator is exported to the Open Neural Network Exchange (ONNX) format. This enables efficient inference across different deep learning frameworks and hardware environments, supporting real-world medical imaging applications.

I. React-Based Frontend for Demonstration

To demonstrate the practical usability of the proposed denoising model, a web-based frontend is developed using React. The frontend provides an interactive interface that allows users to upload noisy chest X-ray images and visualize the corresponding denoised outputs generated by the deployed model.

The React application communicates with the backend inference pipeline through API calls, enabling real-time image processing and result display. This frontend serves as a proof-of-concept demonstration, showcasing how the trained denoising model can be integrated into a user-friendly clinical or research-oriented application.

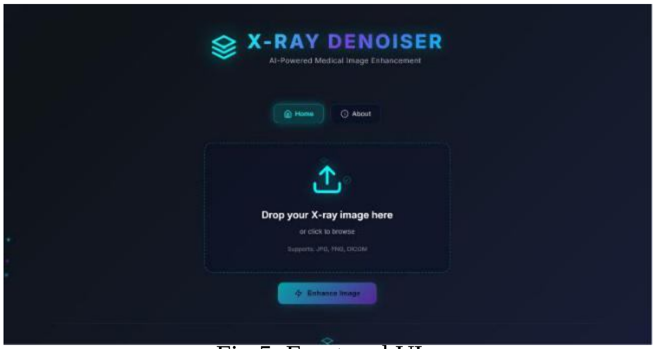

Fig 5. Front end UI

J. End-to-End System Integration

The complete system integrates the trained deep learning model, ONNX-based inference pipeline, and React frontend into a unified workflow. This end-to-end design highlights the feasibility of deploying advanced deep learning–based medical image denoising systems in real-world environments, bridging the gap between research and practical application.

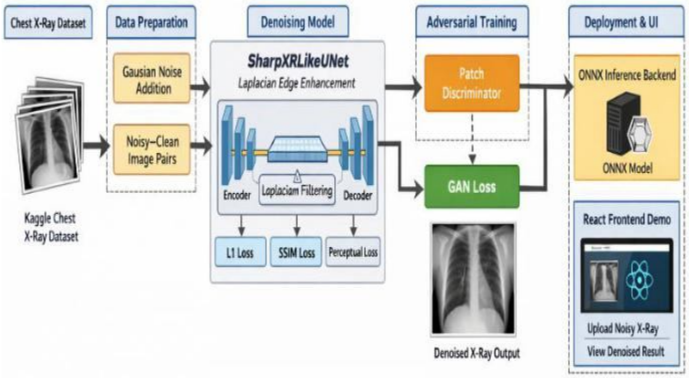

Fig 6. Overall System architecture

IV. EXPERIMENTAL SETUP AND EXPECTED RESULT

A. Experimental Setup

The experiments are conducted using the Kaggle Chest X-ray Pneumonia dataset, from which only normal chest X-ray images are selected. All images are resized to a fixed resolution and normalized before processing. Gaussian noise with predefined variance levels is added to generate noisy–clean image pairs for supervised learning.

The denoising generator, SharpXRLikeUNet, is implemented using a deep learning framework and trained in two stages. In the first stage, the generator is trained independently using a composite reconstruction loss consisting of L1 loss, SSIM loss, and perceptual loss derived from a pretrained VGG16 network. This stage focuses on stabilizing the reconstruction capability of the model and preserving anatomical structures.

In the second stage, adversarial training is introduced by incorporating a patch-based discriminator to form a generative adversarial network (GAN). The discriminator is trained to distinguish between real clean X-ray images and generated denoised images, while the generator is optimized using a weighted combination of reconstruction loss and adversarial loss. Training is performed using gradient-based optimization techniques with carefully selected learning rates to ensure stable convergence.

After training, the generator and discriminator models are saved for inference. The trained generator is exported to the Open Neural Network Exchange (ONNX) format to support efficient deployment. For demonstration purposes, a React-based frontend is developed, which communicates with the backend inference pipeline through APIs, allowing users to upload noisy chest X-ray images and visualize denoised outputs.

B. Expected Results and Performance Analysis

The proposed denoising framework is expected to effectively suppress Gaussian noise while preserving fine anatomical structures in chest X-ray images. Compared to the noisy inputs, the denoised outputs are anticipated to exhibit improved visual clarity, enhanced edge sharpness, and better contrast, particularly around lung boundaries and rib structures.

The integration of Laplacian edge enhancement within the generator architecture is expected to reduce over-smoothing and maintain high-frequency details that are commonly lost in conventional CNN-based denoising models. The use of SSIM and perceptual loss is anticipated to improve structural consistency and perceptual quality, while adversarial training with a patch discriminator is expected to produce more realistic and sharper outputs.

Qualitative visual comparisons between noisy input images and denoised results are expected to demonstrate noticeable improvements in image quality. The React-based frontend further validates the practical applicability of the proposed approach by enabling real-time visualization of denoised X-ray images, highlighting the feasibility of deploying the model in real-world clinical or research environments.

V. RESULTS AND DISCUSSION

The proposed chest X-ray denoising framework was evaluated through a two-stage training process consisting of initial reconstruction-based training followed by adversarial refinement. The performance of the model is analyzed based on training behavior, loss convergence, and qualitative visual improvements.

During the initial training phase, the SharpXRLikeUNet generator was optimized using a composite loss function comprising L1 loss, SSIM loss, and perceptual loss. After 10 epochs, the model achieved a final training loss of 0.0247 and a validation loss of 0.0221. The close alignment between training and validation losses indicates stable convergence and good generalization, with no significant signs of overfitting. These results demonstrate the effectiveness of the combined reconstruction losses in learning accurate pixel-level mappings while preserving structural similarity and perceptual features in chest X-ray images.

Following this phase, adversarial training was introduced by incorporating a patch-based discriminator. After 10 epochs of GAN training, the generator achieved a final loss of 0.1302, while the discriminator loss converged to 0.1880. The balanced behavior of generator and discriminator losses suggests stable adversarial learning, where the generator progressively improves its ability to produce realistic denoised images while the discriminator maintains sufficient discriminative power. This interplay indicates that the adversarial component effectively complements the reconstruction losses without causing training instability.

Qualitative visual inspection reveals that the non-GAN model successfully suppresses Gaussian noise but exhibits mild smoothing in certain regions. In contrast, the GAN-enhanced model produces denoised outputs with sharper edges and improved preservation of fine anatomical structures, particularly around lung boundaries and rib contours. The integration of Laplacian edge enhancement further contributes to maintaining high-frequency details, reducing the loss of structural information commonly observed in conventional denoising networks.

Although quantitative metrics such as PSNR and SSIM were not explicitly computed, the observed reduction in validation loss and the improved visual clarity of denoised outputs indicate a positive impact on both reconstruction accuracy and perceptual quality. The results demonstrate that combining edge-aware architectural design with perceptual and adversarial learning leads to more visually realistic and diagnostically relevant chest X-ray denoising.

Overall, the experimental results confirm that the proposed SharpXRLikeUNet-based GAN framework effectively balances noise suppression and detail preservation, making it a promising approach for medical image denoising applications.

VI. CONCLUSION AND FUTURE WORK

This work presented a deep learning–based framework for denoising chest X-ray images using a customized U-Net-like architecture combined with adversarial learning. The proposed SharpXRLikeUNet integrates Laplacian-based edge enhancement to preserve fine anatomical details while effectively suppressing noise. The model was trained in two stages, initially using a composite reconstruction loss consisting of L1, SSIM, and perceptual loss, followed by adversarial refinement using a patch-based discriminator.

Experimental results demonstrate stable training behavior and effective convergence, with low training and validation losses achieved during the reconstruction phase and balanced generator–discriminator losses during adversarial training. Qualitative evaluation shows that the GAN-enhanced model produces denoised outputs with improved sharpness and structural fidelity compared to non-adversarial reconstruction alone. These findings indicate that the proposed approach successfully balances noise reduction and detail preservation, making it suitable for medical image enhancement applications.

A. Future Work

Several directions can be explored to further enhance the proposed framework. First, quantitative evaluation using standard image quality metrics such as PSNR and SSIM can be incorporated to provide a more comprehensive performance assessment. Second, the robustness of the model can be improved by training with multiple noise levels and different noise types, enabling better generalization to real-world imaging conditions.

Future work may also explore alternative or more advanced generator architectures, attention mechanisms, or multi-scale feature fusion strategies to further enhance denoising performance. Additionally, experimenting with different GAN variants and discriminator designs could improve training stability and visual realism. Finally, deploying the model in a real-world clinical environment and integrating it with hospital imaging workflows would be an important step toward practical adoption.

REFERENCES

[1] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, “Image denoising by sparse 3-D transform-domain collaborative filtering,” IEEE Transactions on Image Processing, vol. 16, no. 8, pp. 2080–2095, Aug. 2007.

[2] A. Buades, B. Coll, and J.-M. Morel, “A non-local algorithm for image denoising,” in Proc. IEEE Computer Society Conf. Computer Vision and Pattern Recognition (CVPR), 2005, pp. 60–65.

[3] O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional networks for biomedical image segmentation,” in Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI), 2015, pp. 234–241.

[4] C. Ledig et al., “Photo-realistic single image super-resolution using a generative adversarial network,” in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2017, pp. 4681–4690.

[5] J. Johnson, A. Alahi, and L. Fei-Fei, “Perceptual losses for real-time style transfer and super-resolution,” in Proc. European Conf. Computer Vision (ECCV), 2016, pp. 694–711.

[6] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, “Image quality assessment: From error visibility to structural similarity,” IEEE Transactions on Image Processing, vol. 13, no. 4, pp. 600–612, Apr. 2004.

[7] I. Goodfellow et al., “Generative adversarial nets,” in Proc. Advances in Neural Information Processing Systems (NeurIPS), 2014, pp. 2672–2680.

[8] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, “Image-to-image translation with conditional adversarial networks,” in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2017, pp. 1125–1134.

[9] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770–778.

[10] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, “Densely connected convolutional networks,” in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2017, pp. 4700–4708.

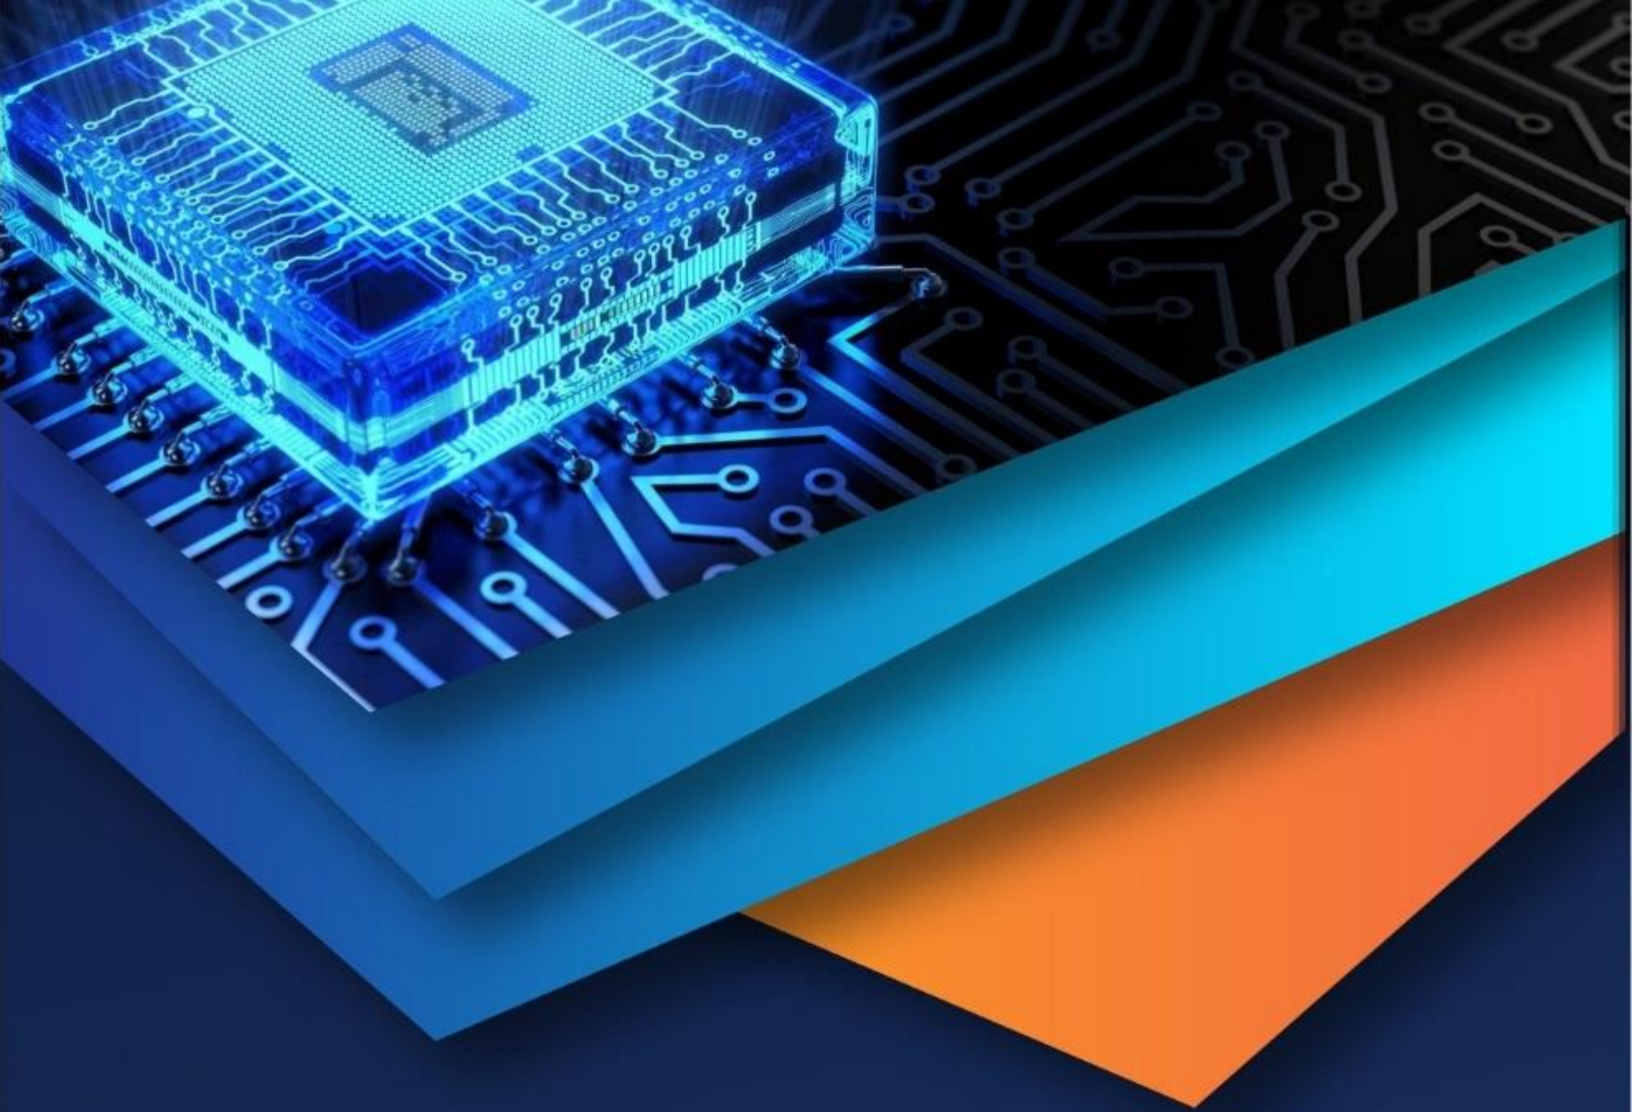

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY