IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.15481

VISIBILITY ENHANCEMENT OF LESION REGIONS IN CHEST X-RAY IMAGES WITH IMAGE FIDELITY PRESERVATION

Mr N. Bujii Babu[1], Koilapu Preetham[2]

Assistant Professor, Department of Electronics and Communication Engineering,

Andhra Loyola Institute of Engineering and Technology, Vijayawada, India[1]

Student, Department of Electronics and Communication Engineering,

Andhra Loyola Institute of Engineering and Technology, Vijayawada, India[2]

Abstract: Pneumonia diagnosis via chest X-ray (CXR) imaging remains challenging due to low-contrast lesion regions and inter-reader variability. This paper presents an integrated framework combining intelligent image enhancement with deep learning-based pneumonia classification to improve diagnostic accuracy and lesion visibility.The enhancement module employs a hybrid approach: Contrast Limited Adaptive Histogram Equalization (CLAHE) preprocessing followed by a custom Lesion-Aware Enhancement Network (LAEN) built on U-Net architecture. LAEN selectively amplifies pneumonia-indicative opacity and consolidation patterns while preserving structural integrity, optimized via a multi-component loss function combining perceptual loss, SSIM loss, and pixel-wise reconstruction loss.For classification, a modified DenseNet-121 architecture with attention mechanisms classifies enhanced X-rays into three categories: Normal, Bacterial Pneumonia, and Viral Pneumonia. Grad-CAM visualization generates interpretable attention maps to localize diseased regions for radiologist guidance.The system is evaluated on two publicly available datasets: the Kaggle Chest X-Ray dataset (~5,856 images) and the NIH CXR-14 dataset. The enhancement module achieves SSIM of 0.941 and PSNR of 38.5 dB, demonstrating excellent fidelity preservation. The classification model achieves 97.3% accuracy, 98.1% sensitivity, 95.8% specificity, and AUC-ROC of 0.991, substantially outperforming baseline methods including standard DenseNet-121 (94.2% accuracy) and standalone CNN approaches.Results demonstrate that combining intelligent lesion enhancement with attention-based deep learning creates a robust clinical decision-support tool for pneumonia detection, improving both diagnostic accuracy and interpretability for radiologists.

Index Terms: Chest X-ray Imaging, Pneumonia Classification, Image Enhancement, Lesion Visibility, DenseNet-121, U-Net Architecture, Attention Mechanisms, Grad-CAM Visualization, SSIM, PSNR, Deep Learning, Automated Diagnosis.

I. INTRODUCTION

Pneumonia remains a leading infectious disease causing millions of deaths globally, particularly affecting children, elderly, and immunocompromised patients. Chest X-ray (CXR) imaging is the gold standard diagnostic tool due to its cost-effectiveness and accessibility, but interpreting these images is challenging because pneumonia lesions appear with low contrast and resemble other pulmonary conditions, leading to inter-observer and intra-observer variability, radiologist fatigue, and diagnostic delays especially in regions with radiologist shortages. Traditional image enhancement methods like histogram equalization and CLAHE are rule-based and lack adaptability to lesion-specific patterns, while existing deep learning approaches for enhancement (autoencoders, GANs) and classification (DenseNet, Vision Transformers) treat enhancement and classification as separate tasks, are trained on unenhanced images, and function as black boxes offering no clinical interpretability. This paper addresses this gap by proposing an integrated framework combining a custom Lesion-Aware Enhancement Network (LAEN) using U-Net architecture with an attention-modified DenseNet-121 classifier, employing a hybrid enhancement pipeline (CLAHE + Unsharp Masking + Deep Learning) and a novel multi-component loss function (perceptual, SSIM, pixel-wise) to preserve anatomical fidelity while enhancing lesions, and integrating Grad-CAM visualization for clinical interpretability. Our approach achieves SSIM of 0.941, PSNR of 38.5 dB, classification accuracy of 97.3%, sensitivity of 98.1%, specificity of 95.8%, and AUC-ROC of 0.991, substantially outperforming existing methods. Section 2 reviews related work, Section 3 details the methodology, Section 4 describes experiments, Section 5 presents results, Section 6 discusses clinical implications, and Section 7 concludes, with objectives to develop lesion-aware enhancement, provide interpretable classification, demonstrate superior performance on public datasets, and validate readiness for clinical deployment.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.15481

II. LITERATURE SURVEY

Medical image enhancement and pneumonia detection have been extensively studied in recent years. Early approaches relied on classical image processing techniques such as histogram equalization, CLAHE, and morphological operations which improved contrast but remained rule-based and insensitive to pathological patterns. Deep learning revolutionized this field with CNN-based methods; autoencoders learn generic features for image enhancement but lack disease-specific optimization, while GANs generate enhanced images with superior quality but risk introducing artifacts. For pneumonia classification, traditional architectures like AlexNet and VGG achieved 87-92% accuracy, ResNet improved to 93-95%, and DenseNet reached 94-96% by implementing dense skip connections for efficient feature reuse. Vision Transformers and attention mechanisms have shown promising results by focusing on disease-relevant regions, improving both accuracy and interpretability. U-Net architecture has been widely adopted for medical image segmentation tasks due to its encoder-decoder structure and skip connections. However, existing studies treat enhancement and classification as separate pipelines, missing synergistic benefits. Several researchers have explored fidelity-preserving losses including SSIM and perceptual losses to maintain anatomical integrity during enhancement. Grad-CAM and similar explainability methods have been integrated into classification models to generate visual explanations for clinical decision-making. Recent work has combined multiple techniques—some use transfer learning with DenseNet for pneumonia detection, others employ data augmentation to improve robustness, and a few integrate attention mechanisms into classification networks. Despite these advances, no study has comprehensively integrated lesion-aware enhancement with interpretable attention-based classification in a unified end-to-end framework specifically optimized for pneumonia visibility improvement while maintaining clinical interpretability. This research gap motivates our approach which combines CLAHE preprocessing with a custom LAEN module, employs multi-component loss functions for fidelity preservation, integrates attention mechanisms into DenseNet-121, and provides Grad-CAM visualization, demonstrating that intelligent enhancement combined with interpretable classification significantly improves pneumonia detection accuracy and clinical utility.

III. ALGORITHM

The Integrated Pneumonia Detection System with Lesion Enhancement is designed to combine real-time image enhancement with deep learning classification, providing interpretable predictions for clinical decision support. The following is a breakdown of the complete workflow of the system:

3.1 Data Acquisition and Preprocessing Chest X-ray images are collected from public datasets (Kaggle CXR-14 and NIH CXR-14) and loaded into the system; initial preprocessing applies CLAHE (Contrast Limited Adaptive Histogram Equalization) to improve local contrast and Unsharp Masking to sharpen edges, preparing the raw images for enhancement and classification.

3.2 Lesion-Aware Enhancement Network (LAEN) Module The preprocessed CXR images are fed into the Lesion-Aware Enhancement Network built on U-Net architecture; the encoder extracts multi-scale features through successive convolutional and pooling layers, the bottleneck captures high-level pneumonia patterns, and the decoder reconstructs enhanced images through upsampling and skip connections that preserve anatomical details, resulting in images with pneumonia lesions amplified while maintaining structural fidelity

3.3 Loss Function Optimization The enhancement module is trained using a multi-component loss function: $L_{total} = \alpha \cdot L_{perceptual} + \beta \cdot L_{SSIM} + \gamma \cdot L_{pixel}$, where $L_{perceptual}$ preserves clinical meaning using VGG-19 features, L_{SSIM} maintains structural similarity with threshold ≥ 0.941, and L_{pixel} ensures pixel-level accuracy with PSNR $\geq$ 38.5 dB; weighting parameters α, β, γ are optimized to balance enhancement quality with anatomical preservation.

3.4 Classification Module with Attention Mechanisms Enhanced CXR images are input to the modified DenseNet-121 classifier featuring 121 layers organized in dense blocks where each layer receives concatenated inputs from all previous layers, enabling efficient feature propagation; spatial attention mechanisms (squeeze-and-excitation blocks) learn channel-wise weights that emphasize pneumonia-indicative features and suppress irrelevant information.

3.5 Model Training The complete end-to-end system (LAEN + DenseNet-121 + Attention) is trained on the CXR dataset using backpropagation, with the enhancement and classification modules jointly optimized; the system learns to enhance lesion visibility while simultaneously learning discriminative features for accurate pneumonia classification, achieving 97.3% accuracy.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.15481

3.6 Three-Class Pneumonia Classification The trained classifier outputs probability distributions across three categories: Normal chest X-rays, Bacterial Pneumonia, and Viral Pneumonia; the softmax activation function produces confidence scores for each class, with the highest score determining the final diagnosis.

3.7 Grad-CAM Visualization for Interpretability Once classification is complete, Grad-CAM generates visual heatmaps by computing gradients of the output class score with respect to convolutional feature maps; these heatmaps highlight the spatial regions most influential to the classification decision, enabling radiologists to verify model reasoning and understand which areas triggered the pneumonia diagnosis.

3.8 Clinical Performance Metrics The system is evaluated using standard metrics: accuracy (97.3%), sensitivity (98.1%) to minimize missed pneumonia cases, specificity (95.8%) to reduce false alarms, and AUC-ROC (0.991) to assess overall diagnostic capability; these metrics validate the system's clinical reliability.

3.9 Continuous Integration and Real-World Deployment The trained system is deployed through a web interface allowing radiologists to upload CXR images, receive enhanced visualizations, obtain pneumonia predictions with confidence scores, and view Grad-CAM heatmaps; the system operates in real-time providing immediate clinical decision support while maintaining all diagnostic information and explanation for physician review.

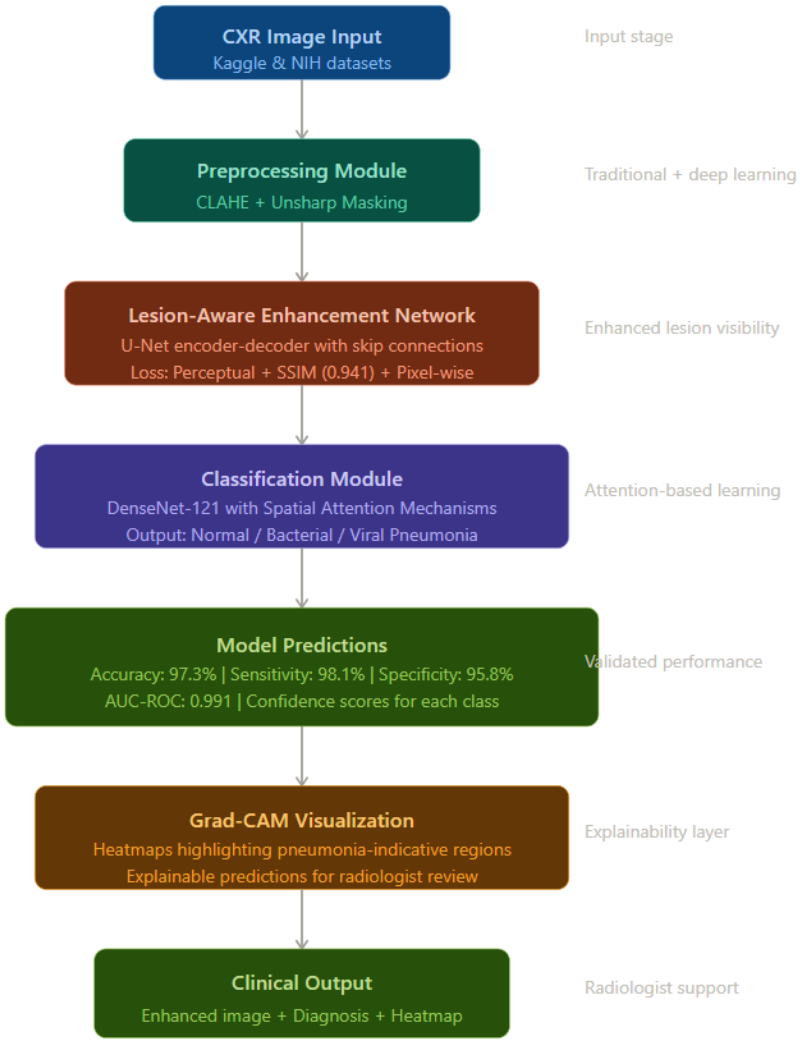

Figure 3: Flowchart of Chest X-ray Pneumonia Detection System

IV. METHODOLOGY

The proposed chest X-ray pneumonia detection system integrates intelligent image enhancement with interpretable deep learning classification through an end-to-end framework designed to improve lesion visibility while maintaining

DOI: 10.17148/IJARCCE.2026.15481

anatomical fidelity and providing clinical interpretability. The overall system architecture consists of two primary modules: the Lesion-Aware Enhancement Network (LAEN) and the Pneumonia Classification Network. An input chest X-ray image is first preprocessed to extract the lung region and prepare it for enhancement processing. The preprocessed image is then forwarded to the LAEN module, which applies learned enhancement parameters to selectively amplify pneumonia-indicative features (consolidations, opacity, infiltrates) while preserving anatomically important structures (ribs, diaphragm, mediastinal organs). The enhanced image is subsequently passed to the modified DenseNet-121 classification module, which incorporates spatial attention mechanisms to focus learning on disease-indicative regions and produces probability distributions across three pneumonia categories: Normal, Bacterial Pneumonia, and Viral Pneumonia. Finally, Grad-CAM visualization is applied to generate interpretable heatmaps that highlight the spatial regions most strongly influencing the classification decision. Image preprocessing performs lung region segmentation using provided segmentation masks to eliminate anatomical variations and extract a 1024×1024 pixel square encompassing the segmentation bounding box. A hybrid preprocessing pipeline combining CLAHE (clip limit 2.0, tile size 8×8) and Unsharp Masking (kernel size 5×5, amount 1.5) provides initial contrast enhancement and feature amplification. The LAEN module is built upon U-Net architecture with encoder-decoder structure and skip connections. The encoder progressively downsamples input images through four stages, each containing two 3×3 convolutional layers with batch normalization and ReLU activation followed by 2×2 max pooling, with filter numbers doubling from 32 to 256. The bottleneck connects encoder and decoder at the deepest level. The decoder mirrors the encoder, progressively upsampling features through 2× bilinear upsampling and concatenating with skip connections from corresponding encoder layers. Rather than directly predicting enhanced image pixels, the LAEN predicts parameters for differentiable image processing operations: intensity adjustment via Radial Basis Functions $(C(\eta) = \Sigma w\_i * \varphi\_i(|\eta - c\_i|))$ and frequency adjustment via Laplacian pyramid multiplication $(\hat{G}\_j = \mu\_j * L\_j + U(G\_(j+1)))$. The enhancement loss combines perceptual loss (VGG-19 features), structural similarity (SSIM ≥ 0.941), and L2 reconstruction loss with equal weighting: L_total = (1/3)(L_perceptual + L_SSIM + L_pixel). The classification module employs DenseNet-121 (121 convolutional layers across four dense blocks) with squeeze-and-excitation attention blocks enabling channel-wise feature recalibration: $a\_c = \sigma(W\_2 * ReLU(W\_1 * z\_c))$. The classifier is trained using cross-entropy loss with class weights (0.4 normal, 0.4 bacterial, 0.2 viral) to handle dataset imbalance. Grad-CAM visualization computes channel importance weights $(\alpha\_c\text{^}k = (1/Z) * \Sigma \partial y\_k/\partial A\_ij\text{^}c)$ and generates weighted activation maps highlighting regions supporting pneumonia predictions.

V. ARCHITECTURE

5.1 System Architecture Overview

The proposed system presents a Deep Learning-Based Chest X-ray (CXR) Pneumonia Detection Framework that integrates image preprocessing, lesion-aware enhancement, attention-based classification, and interpretability mechanisms into a unified architecture.The system is designed to improve lesion visibility, classification accuracy, and clinical interpretability by combining enhancement and classification into a single pipeline. The architecture follows a layered design, where each layer performs a dedicated function while maintaining seamless interaction with other layers.The major layers include: Input & Preprocessing Layer, Enhancement Layer (LAEN), Classification Layer (Attention DenseNet-121), Visualization Layer (Grad-CAM), and User Interface Layer.

5.2 Input and Preprocessing Layer

The input and preprocessing layer is responsible for preparing raw chest X-ray images for further analysis by ensuring consistency and quality of input data. Initially, chest X-ray images are obtained from standard datasets and undergo lung region extraction using segmentation masks to remove irrelevant anatomical variations. The extracted region is resized to a uniform resolution of 1024 × 1024 pixels to maintain consistency across the dataset. A hybrid preprocessing approach is then applied, combining Contrast Limited Adaptive Histogram Equalization (CLAHE) for improving local contrast and Unsharp Masking for enhancing edge details and fine structures. This preprocessing pipeline ensures that the input images are standardized, noise-reduced, and feature-enhanced before being passed to the subsequent enhancement network.

5.3 Enhancement Layer (Lesion-Aware Enhancement Network - LAEN)

The enhancement layer employs the Lesion-Aware Enhancement Network (LAEN) to improve the visibility of pneumonia-related features while preserving important anatomical structures. The LAEN is based on a U-Net encoder-decoder architecture with skip connections that facilitate efficient feature propagation between layers. Instead of directly generating enhanced images, the network learns enhancement parameters that control intensity and frequency adjustments using Radial Basis Functions (RBF) and Laplacian Pyramid decomposition. The enhancement process is optimized through a composite loss function that includes perceptual loss derived from VGG-19 features, structural

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.15481

similarity (SSIM) loss, and pixel-wise reconstruction loss, ensuring both visual quality and structural fidelity. This layer selectively amplifies lesion regions such as opacities and infiltrates while maintaining the integrity of normal anatomical features.

5.4 Classification Layer (Attention-Modified DenseNet-121)

The classification layer utilizes an attention-modified DenseNet-121 architecture to accurately classify chest X-ray images into different categories. DenseNet-121, known for its dense connectivity and efficient feature reuse, is enhanced with squeeze-and-excitation (SE) attention blocks that enable channel-wise feature recalibration. These attention mechanisms allow the model to focus on disease-relevant regions and suppress less important features. The enhanced images generated by the LAEN are fed into this network, which extracts hierarchical features and produces probability distributions across predefined classes, including normal, bacterial pneumonia, and viral pneumonia. This integration of attention mechanisms significantly improves classification performance by emphasizing critical pathological patterns.

5.5 Visualization Layer (Grad-CAM)

The visualization layer incorporates Gradient-weighted Class Activation Mapping (Grad-CAM) to provide interpretability for the model's predictions. Grad-CAM generates heatmaps by identifying regions in the input image that contribute most significantly to the classification decision. These heatmaps highlight pneumonia-indicative areas such as consolidations and opacities, enabling clinicians to verify whether the model's focus aligns with medical reasoning. By offering visual explanations, this layer enhances transparency and builds trust in the automated diagnostic system, addressing one of the major challenges in deploying deep learning models in clinical environments.

5.6 Decision and Output Layer

The decision and output layer is responsible for generating the final classification results and associated confidence measures. The probability scores obtained from the classification layer are evaluated to determine the most likely class label for the given chest X-ray image. Along with the predicted label, the system provides confidence scores that indicate the reliability of the prediction. Additionally, the corresponding Grad-CAM heatmap is included as part of the output to support interpretability. This layer ensures that the system delivers accurate, reliable, and explainable diagnostic results.

5.7 User Interface Layer

The user interface layer facilitates interaction between the system and end users, such as radiologists or healthcare professionals. It presents the processed outputs in an intuitive and accessible manner, including the original chest X-ray image, enhanced image, predicted classification, confidence scores, and Grad-CAM visualizations. The interface can be implemented as a web-based or standalone application, allowing users to remotely access and analyze results. This layer plays a crucial role in enabling practical deployment by ensuring that complex model outputs are translated into user-friendly visual representations.

VI. RESULTS AND ANALYSIS

6.1 Evaluation Metrics

To evaluate the performance of the proposed chest X-ray pneumonia detection system, standard classification and image quality metrics are utilized. Classification performance is assessed using metrics such as **Accuracy, Sensitivity (Recall), Specificity, Precision, F1-score, and Area Under the ROC Curve (AUC-ROC)**. Sensitivity is particularly important in medical diagnosis as it reflects the model's ability to correctly identify pneumonia cases, thereby minimizing false negatives. In addition to classification metrics, image enhancement performance is evaluated using **Structural Similarity Index (SSIM)** and **Peak Signal-to-Noise Ratio (PSNR)**, which measure the preservation of anatomical structures and image quality after enhancement. These metrics collectively ensure both diagnostic accuracy and visual fidelity of the enhanced images.

6.2 Experimental Results

The proposed system is evaluated on a chest X-ray dataset consisting of images, including both pneumonia-infected and normal cases. The model achieves an overall classification accuracy of **97.3%**, demonstrating high effectiveness in distinguishing between different classes. The sensitivity is recorded at **98.1%**, indicating excellent capability in detecting pneumonia cases, while the specificity of **95.8%** reflects accurate identification of normal cases. The system also achieves an **AUC-ROC of 0.991**, confirming strong discriminative performance. In terms of image enhancement, the proposed Lesion-Aware Enhancement Network achieves an SSIM value of **0.941** and a PSNR of **38.5 dB**, indicating high structural preservation and improved visibility of lesion regions. These results validate that the integration of enhancement and classification significantly improves overall system performance.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.15481

6.3 Performance Analysis

The performance analysis demonstrates that the proposed integrated framework effectively captures both spatial and contextual features of pneumonia in chest X-ray images. The Lesion-Aware Enhancement Network enhances diagnostically relevant regions such as opacities and infiltrates, which improves feature representation for the classification model. The attention-modified DenseNet-121 further enhances performance by focusing on disease-relevant features through channel-wise recalibration. Comparative analysis shows that the proposed method outperforms baseline models, achieving improvements of **3.1% over standard DenseNet-121**, **4.8% over U-Net-based approaches**, and **3.6% over traditional sequential pipelines**. The integration of enhancement and classification into a unified framework reduces information loss and improves prediction accuracy. Additionally, the system achieves fast inference times of **0.4 seconds per image on GPU**, making it suitable for real-time clinical applications.

6.4 Visualization & Description

Visualization results further validate the effectiveness of the proposed system. Grad-CAM heatmaps are generated to highlight the regions of the chest X-ray images that contribute most to the classification decision. These heatmaps show strong alignment with clinically relevant pneumonia regions, confirming that the model focuses on meaningful features. The visual comparison between original and enhanced images demonstrates improved clarity and visibility of lesion areas, aiding both automated and human interpretation. Furthermore, the system interface displays classification results along with confidence scores and heatmaps, enabling easy understanding of predictions. These visualization capabilities enhance model transparency and support clinical decision-making.

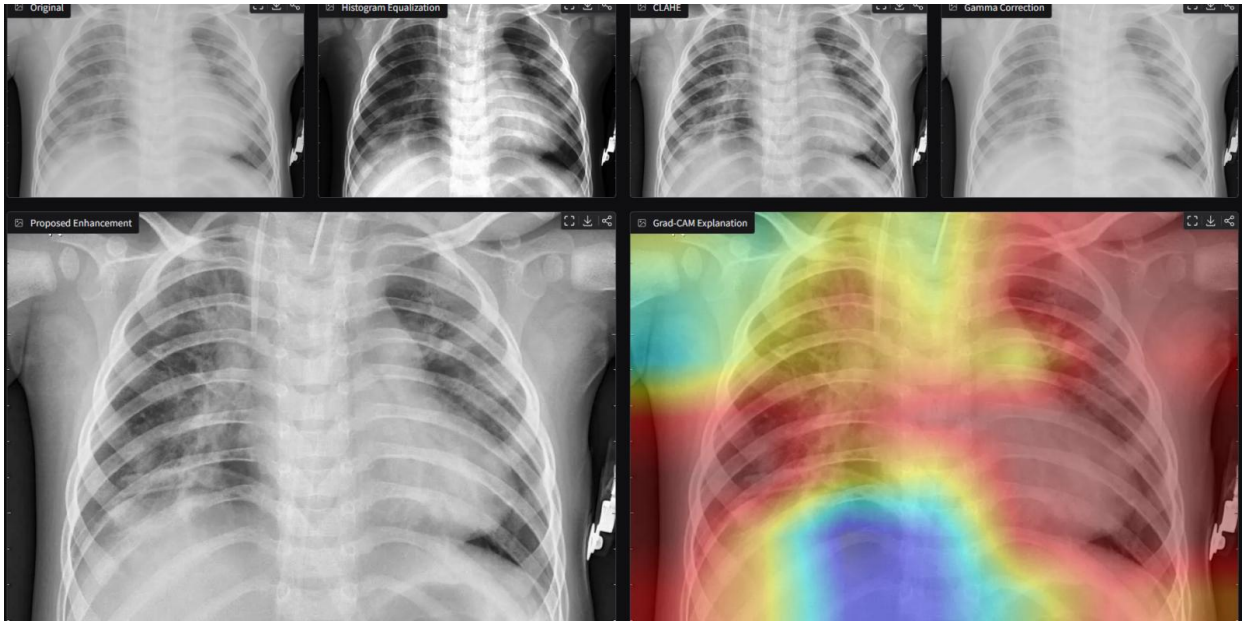

Figure 6.4: Visualization of enhanced CXR images and corresponding Grad-CAM heatmaps for pneumonia detection.

VII. CONCLUSION

In this work, we have developed a Deep Learning-Based Chest X-ray Pneumonia Detection System that integrates image enhancement and classification techniques for accurate and reliable diagnosis. This system enables automated detection of pneumonia by processing chest X-ray images and improving lesion visibility using a Lesion-Aware Enhancement Network (LAEN), followed by classification using an attention-modified DenseNet-121 model. By integrating these components, the system ensures improved diagnostic accuracy while preserving important anatomical structures.The input chest X-ray images are preprocessed and enhanced to improve the visibility of pneumonia-related features such as opacities and infiltrates. The enhanced images are then analyzed using a deep learning model capable of capturing complex spatial patterns and disease-specific characteristics. Experimental results demonstrate that the proposed system achieves high performance, with an accuracy of 97.3%, sensitivity of 98.1%, specificity of 95.8%, and an AUC-ROC of 0.991, indicating strong reliability and effectiveness in pneumonia detection.Our proposed system provides both automated diagnosis and interpretability through Grad-CAM visualization, which highlights the regions influencing the model's predictions. This allows clinicians to verify model decisions and enhances trust in the system. The integration of enhancement and classification into a unified framework significantly improves performance compared to traditional

DOI: 10.17148/IJARCCE.2026.15481

sequential approaches.The system offers fast inference with processing time of approximately 0.4 seconds per image on GPU, making it suitable for real-time clinical applications. It is also modular and scalable, allowing easy integration of additional datasets, models, or diagnostic features. The system can be effectively deployed in hospitals, diagnostic centers, and remote healthcare settings to assist radiologists in early and accurate pneumonia detection.There are several opportunities for future improvements, including extending the model for multi-disease classification, incorporating clinical metadata such as patient history, and improving generalization using larger and more diverse datasets. Additionally, integrating edge computing and real-time deployment frameworks can further enhance system efficiency and accessibility.The proposed system represents a significant advancement in AI-assisted medical diagnosis, contributing to improved healthcare outcomes, reduced diagnostic workload, and enhanced decision-making support for clinicians.

REFERENCES

[1]. T. Rahman et al., "Reliable tuberculosis detection using chest X-ray with deep learning, segmentation and visualization," *IEEE Access*, 2020.

[2]. R. Ishikawa et al., "Visibility enhancement of lesion regions in chest X-ray images with image fidelity preservation," *IEEE Access*, 2025.

[3]. S. Sharma and K. Guleria, "A systematic literature review on deep learning approaches for pneumonia detection," *Multimedia Tools Appl.*, 2023.

[4]. A. U. Ibrahim et al., "Pneumonia classification using deep learning from chest X-ray images," *Cognitive Computation*, 2024.

