Received: 29 November 2025 | Revised: 8 April 2026 | Accepted: 21 May 2026 | Published online: 3 June 2026

RESEARCH ARTICLE

Artificial Intelligence and Applications
2026, Vol. 00(00) 1–13
DOI: 10.47852/bonviewAIA62028511

BON VIEW PUBLISHING

A Novel PneumoniaNet Framework Integrating Explainable AI for Pediatric Pneumonia Detection from Chest X-rays

Shahriar Siddique Arjon[1,*], Tamanna Yasmin[1], Ankur Kumar Mondol[1], Nakib Aman[1] and Shabbir Mahmood[1]

[1] *Department of Computer Science and Engineering, Pabna University of Science and Technology, Bangladesh*

Abstract: Pneumonia remains one of the leading causes of death among children worldwide, and therefore, it is necessary to use reliable and efficient tools to detect the disease early. In this study, a novel hybrid residual-dense deep learning architecture called PneumoniaNet is presented that detects pediatric pneumonia from chest X-ray images with high accuracy and interpretability. The proposed model is designed on the basis of the MobileNetV2 framework and utilizes fine-tuned deep layers with frozen shallow layers. This design enables learning high-level radiological features while preserving low-level visual information. The architecture has residual-dense connectivity to improve the propagation of gradients, feature reuse, and model generalization. The model was evaluated with a widely used benchmark pediatric chest X-ray dataset, and its results were comparatively analyzed with the performance of a number of state-of-the-art CNNs, such as DenseNet121, InceptionV3, ResNet50, EfficientNetB0, and AlexNet. PneumoniaNet yielded the best results with an accuracy of 93.11%, precision of 92.84%, recall of 93.09%, specificity of 92.43%, and F1-score of 94.29%. In addition, Gradient-weighted Class Activation Mapping visualizations showed focus on clinically relevant regions, improving interpretability and transparency. An ablation study of four PneumoniaNet variants, trained in the same manner, suggested that residual-dense connections along with selective layer freezing provide the most generalizable and effective feature representations. This work is unique because it combines residual-dense connectivity with selective layer freezing in a lightweight MobileNetV2 architecture, enabling effective feature reuse, facilitating gradient flow, and providing interpretable predictions—all within a computationally efficient framework.

Keywords: pneumonia detection, pediatric chest X-ray, deep learning, explainable AI

1. Introduction

Pneumonia is one of the major causes of death worldwide, and it affects infants, children, and immunocompromised individuals in a disproportionate manner. It involves the inflammation of pulmonary alveoli that can be filled with fluid or pus, causing such symptoms as cough, fever, chest pain, and shortness of breath. It may escalate to severe respiratory failure, sepsis, or death within a short time unless diagnosed and treated in time [1]. Global health reports have given pneumonia as the leading infectious cause of death in children below the age of five years, and this has necessitated a greater focus on developing more effective ways of early diagnosis and treatment. Early and precise diagnosis is essential, as it provides a chance to initiate treatment on a timely basis and avoid complications, and to a considerable degree, this reduces morbidity and mortality [2].

The conventional diagnostic test of pneumonia is the chest X-ray (CXR). It is commonly applied because it is affordable, available, and effective in the diagnosis. Although it is widely used, the interpretation of CXRs by clinicians and radiologists remains challenging and subjective. The fact that radiographic images are complex and variable complicates accurate diagnosis because they overlap, such as the ribs, blood vessels, and bronchi, so much that they may obscure or resemble features of pneumonia [3]. Moreover, the existence of early-stage or abnormal cases, changes in patient positioning, and discrepancies in image quality are also factors leading to diagnostic uncertainty. Clinical assessment outcomes are often characterized by a lack of reliability and inter-rater variability, which is a characteristic of human judgment subject to individual interpretation due to experience, as well as fatigue of a radiologist. These constraints demonstrate the necessity of sophisticated automated methods of diagnosing to boost interpretability, reduce diagnostic uncertainty, and enhance the precision and reproducibility of pneumonia detection on the basis of CXR findings [4]. This means that automated diagnostic technologies that can promptly, accurately, and consistently detect pneumonia are becoming more important to aid clinical decisions and enhance patient outcomes [5].

Deep learning and artificial intelligence (AI) have emerged as technologies that have revolutionized the analysis of medical images by providing strong solutions to the automated diagnosis of the disease. Convolutional neural networks (CNNs), as one of the core elements of this change, have proven to be particularly effective at tasks like pneumonia detection on CXRs, hierarchic feature representations learned directly on pixel data [6].

***Corresponding author:** Shahriar Siddique Arjon, Department of Computer Science and Engineering, Pabna University of Science and Technology, Bangladesh. Email: arjon.200141@s.pust.ac.bd

The benefits of AI-driven systems are that they can handle such complex analyses and make a rapid assessment of a wide variety of imaging features that some might be missed in the manual evaluation process, making them useful computer-aided diagnostic (CAD) tools that improve the quality and consistency of diagnosis [7]. The technologies can be used in different fields of medical imaging, such as skin cancer classification, diabetic retinopathy detection, and early lung cancer identification. They can be applied beyond pneumonia, which highlights the scope and ever-growing presence of AI in contemporary healthcare analytics [8].

The fact that the underlying deep learning models are not transparent and hence cannot be easily understood in their decision-making process, however, remains an obstacle to adopting and accepting them in clinical practice. Explainable AI (XAI) methods help with this problem by trying to explain the black box and provide transparent and interpretable explanations for model decisions. To take an example, the existing XAI methods can be incorporated into the framework to produce visual heatmaps that reveal certain areas of a CXR that have the most significant effect on the model prediction [9]. This will enhance transparency and interpretability by demonstrating that the model focuses on clinically relevant areas, including those with pulmonary opacities. Moreover, XAI methods can identify the most crucial pixels to diagnose the patient by manipulating single pixels in an image systematically and measuring the impact of the alterations on the predictions of the model [10]. XAI presents the decision made by the AI in a form of visualization that reflects the expertise of radiologists and makes these models more transparent instead of opaque and a trustworthy diagnostic assistant that can be safely introduced into clinical practice.

Although significant progress has been made, most existing studies focus on either improving classification accuracy or enhancing interpretability, but not both simultaneously. The combined optimization of residual-dense feature learning and selective layer freezing within a single lightweight and explainable system for analyzing pediatric CXRs—particularly building upon prior research with this dataset—remains largely unexplored [11]. This gap motivates the present work. This paper proposes PneumoniaNet, a hybrid residual-dense convolutional network with selective layer freezing and integrated XAI, designed to detect pediatric pneumonia with outstanding interpretability and accuracy. The framework addresses the limitations of existing methods by jointly optimizing feature reuse, gradient propagation, and model transparency, enabling the model to capture subtle radiographic variations and maintain diagnostic consistency. The main contributions of this work are as follows:

a. A novel PneumoniaNet architecture integrating residual-dense connectivity with a selectively frozen MobileNetV2 backbone for enhanced feature reuse and gradient propagation.
b. An explainable AI pipeline using Gradient-weighted Class Activation Mapping (Grad-CAM) to provide clinically relevant visual interpretations.
c. A comprehensive ablation analysis (A1–A4) isolating each architectural component.
d. A reproducible pipeline with extensive augmentation and experimental settings.

The rest of this paper is organized in the following manner: Section 2 will contain a review of the available literature on the topic of deep learning-based pediatric pneumonia detection. Section 3 outlines the research methodology, including datasets, preprocessing steps, and CNN architectures employed. Section 4 reports experimental results, performance evaluation, and Grad-CAM-based interpretability analyses. Section 5 summarizes the study and gives its limitations and future work directions.

2. Literature Review

In recent years, profound progress in detecting pneumonia in CXRs with the help of deep learning has been achieved. Research has focused on enhancing both diagnostic accuracy and clinical interpretability with explainable, hybrid, and convolutional AI models.

Sharma and Guleria have proposed a method of transfer learning to automatically identify pneumonia based on the CXRs [12]. They employed VGG-16 for feature extraction along with a classifier. The model was trained and evaluated on two Kaggle datasets with a total of 12,000 images and reached accuracies of 92.15% and 95.4%, outperforming conventional algorithms such as Support Vector Machine (SVM) and K-Nearest Neighbors (KNN). However, the approach is more dependent on the quality of the dataset than on more recent architectures, and it might use more diversity and augmentation. This paper shows that integrating deep transfer learning with small neural classifiers may prove a successful approach for pneumonia detection.

Siddiqi and Javaid reviewed 140 research studies published between 2020 and 2023 to provide an overview of deep learning approaches for detecting pneumonia from CXR images [13]. They considered publicly available datasets such as the COVID-19 Radiography Database and the NIH CXR dataset and analyzed approaches like CNNs, Vision Transformers (ViTs), and hybrid or ensemble models. Segmentation and data augmentation were used in most studies to improve feature learning, and the performance was measured by such measures as accuracy, F1-score, and Area Under the Curve (AUC). ViTs have higher representational strength than CNN-based techniques, but they still face challenges with dataset bias and interpretability. Overall, the survey offers a comparative evaluation of the strengths and limitations of different model families.

Wu et al. made a move to eliminate the necessity of manual adjustment of the anchors, as is typical of the detectors of the traditional type, in their attempt to introduce an anchor-free deep learning architecture to the task of locating pneumonia in the CXRs [14]. It is based on their model, trained on the RSNA data with more than 26,000 annotated images, which consists of a ResNet-50 backbone and a feature pyramid network with a dual-branch detection head optimized with the focus loss. It had a mean average precision of 51.5%, which is better than numerous anchor-based baseline models. This approach improves robustness and simplifies design when compared to anchor-based approaches, but performance is still dependent on the choice of feature layer.

Bakir et al. combined a bespoke artificial neural network (ANN) with pretrained convolutional models to develop a hybrid deep learning framework for the classification of pneumonia in CXRs [15]. They utilized ResNet, Inception, and MobileNet as feature extractors on a publicly available Kaggle dataset. Following image scaling and normalization, an ANN classifier was used. When it came to binary classification accuracy, the MobileNet-ANN combo was the best (95.67%). By integrating robust feature extraction with adaptive decision bounds, their hybrid method outperforms standalone CNNs. However, distinguishing between bacterial and viral pneumonia remained difficult.

Yi et al. have concentrated on the end-to-end learning that does not depend on transfer learning by training a 52-layer deep CNN to identify the presence of pneumonia using CXRs [16]. To solve the problem of class imbalance, they processed a Kaggle dataset with data augmentation, contrast-limited adaptive

histogram equalization, and MinMax normalization. The model had an accuracy of 96.09% and specificity of 98.61%, which is higher than other pretrained architectures and features extraction and classification in a single step. This model captures more task-specific features than transfer learning approaches, but it requires more computational resources.

Kaushik et al. proposed the Beyond Radiology approach, which consists of combining multimodal data and advanced machine learning methods to enhance the detection of pneumonia [17]. Their research aimed at decreasing diagnostic subjectivity and reliance on radiological interpretation through a rigorous dataset that comprised CXRs, electronic health records, and laboratory outcomes. In order to make the model more robust and generalized, they used multimodal fusion, data augmentation, and standardization methods. Their model achieved 90.36% accuracy. By adding additional clinical data, this method improves generalization in comparison with unimodal imaging models, but it also enhances system complexity.

Prakash et al. created the Artificial Intelligence-based Learning for Pneumonia Detection (AILPD) model to automatize the process of detecting pneumonia through the use of CXRs [18]. They integrate gradient boosting and CNN-based designs in their system to enhance the accuracy of the classification and the diagnostic delay. AI learning was superior to traditional CNN models, as AILPD produced 97.86% accuracy with 500 epochs of training. Although this hybrid strategy performs better than many traditional CNN models, it has higher computational complexity.

Colin integrated four explainability techniques—Layer-wise Relevance Propagation (LRP), adversarial training, Class Activation Maps (CAMs), and spatial attention mechanisms—with ResNet50 to develop an interpretable deep learning framework for pneumonia detection from CXR images [19]. Without compromising performance, LRP achieved the best results on a dataset of 5863 images, reaching 91% accuracy, 90% sensitivity, 92% specificity, and an interpretability score Magnetic Resonance Spectroscopy (MRS) of 0.85. In contrast, SAM produced the lowest interpretability (MRS 0.12), while adversarial training and CAMs showed performance trade-offs, with specificity decreasing to 79% and 83%, respectively. Overall, LRP proved to be the most suitable for clinical applications despite its higher computational cost.

Ruchika et al. proposed a CNN that was a VGG-19-based one, which was used to diagnose pneumonia based on CXR images [20]. To address the limitation in the quantity of data accessible, their approaches involved data augmentation by using transfer learning on a Kaggle dataset. The first VGG-19 frozen blocks were trained, and the model was further developed with deeper layers being tuned with batch normalization and dropout regularization. Their approach demonstrated a significant difference in performance as they had an accuracy of 95% using augmented data against 78% using no augmentation. Further validation in multi-source clinical data is yet to be done in order to use it in a broader clinical setting.

Ren et al. developed an interpretable framework, MulNet, that combines a DenseNet-121 neural network model of CXR processing by a Bayesian network model using seven medical report clinical symptoms, with an AUC of 0.87 [21]. Application of a medically informed Bayesian model coupled with conditional probability tables that describe the combined effect of both symptoms and imaging results on diagnosis is one of the key contributions of the model because of its inherent explainability. Nevertheless, the extraction of symptoms in terms of textual processing is rule based and possibly lacks the richness of clinical language, which might influence its robustness, and the picture model was not contrasted with more advanced structures.

Sheu et al. created an explainable AI system called XAI-ICP to classify pneumonia infections, which is based on the DenseNet-121 model, pretrained on the NIH CXR dataset, and then fine-tuned on hospital and VinDr data [22]. In the case of transfer learning, the system attained 93.29% accuracy. The use of SHAP with Grad-CAM to provide both feature-level and region-level interpretability is another valuable contribution of the study, as well as a human-in-the-loop method of medical labeling and model optimization. Nevertheless, the study is limited by the fact that it has focused on only three radiological characteristics, namely, infiltrate, cardiomegaly, and effusion, which may have ignored other clinical indicators that can be important in pneumonia.

Papadimitriou et al. proposed a new deep CNN for pneumonia detection that scored a high accuracy of 96.6%, sensitivity of 98.1%, and specificity of 92.4% [23]. The systematic examination of different architectural structures and batch sizes in an attempt to establish the best model that balances the various architectural structures is one of the strengths of the study. Despite strong performance, the lack of explainability limits its applicability in clinical decision-making.

Ali et al. designed a new deep learning architecture to detect pneumonia using a dataset of child CXRs and demonstrated that the EfficientNetV2L model performed better than others with 94.02% accuracy [24]. A major contribution of the study is the comparison of the study with five popular models with k-fold cross-validation and data augmentation. It lacks interpretability and real-time validation, which are essential for deployment in resource-constrained conditions, even though performance is competitive, similar to other advanced models.

Overall, current studies either focus on improving interpretability or enhancing classification accuracy, with limited efforts addressing both simultaneously. Moreover, there is still much to explore regarding the use of lightweight architectures combined with explainable AI and residual-dense feature learning, particularly in pediatric CXR analysis.

3. Research Methodology

This study employs a robust and interpretable deep learning method to automatically detect pediatric pneumonia from CXR images. To enhance feature reuse, gradient flow, and model generalization for accurate pneumonia detection, it introduces a novel architecture, PneumoniaNet, built upon the MobileNetV2 backbone, integrating fine-tuned deep layers with frozen shallow layers. The framework also incorporates six baseline CNN architectures, including MobileNetV2, AlexNet, DenseNet121, ResNet50, InceptionV3, and EfficientNetB0. The pediatric CXR dataset is used for training and validation of all models, with comprehensive preprocessing including balanced class stratification, resizing to 224 × 224 pixels, and image normalization. Furthermore, advanced data augmentation techniques—such as translation, rotation, zoom, shear, brightness adjustment, and flipping—were applied to ensure robustness and reduce overfitting. To ensure statistical reliability, the models were optimized using the Adam optimizer with categorical cross-entropy loss. Quantitative performance was evaluated using accuracy, precision, recall, F1-score, specificity, and AUC metrics. To address the gap between interpretability and clinical transparency, Grad-CAM visualizations were used to provide explanations, highlighting the key image regions that influence the model's predictions. Furthermore, an ablation study was conducted to systematically evaluate the contribution of each architectural component of PneumoniaNet under identical training conditions. The proposed workflow diagram is illustrated in Figure 1.

Figure 1
Proposed workflow for pediatric pneumonia detection using PneumoniaNet

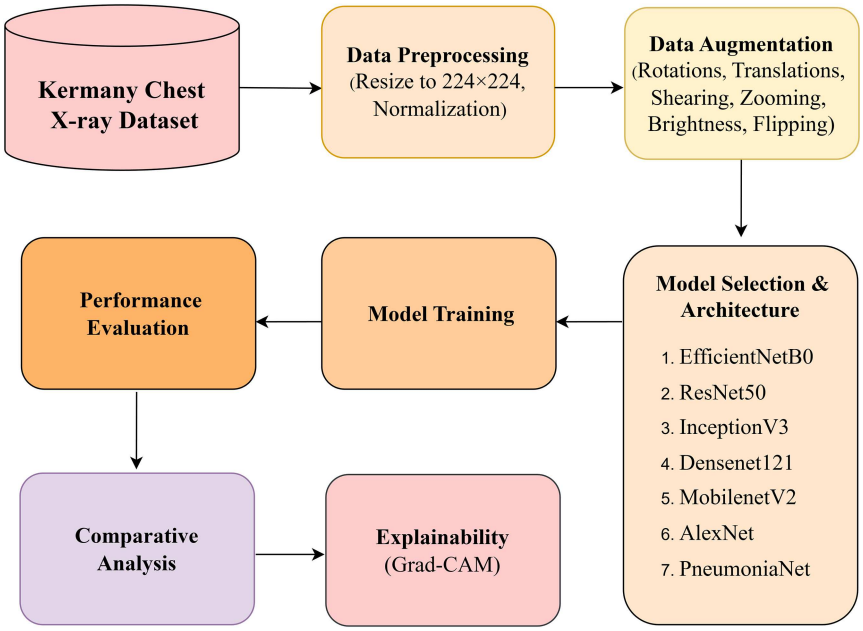

3.1. Dataset description

The dataset used in this study is publicly accessible on Kaggle and was initially presented by Kermany et al. [25]. It contains 5856 CXR images of children aged between 1 and 5 years in China, and they were formatted retrospectively at Guangzhou Women and Children Medical Center. The images are classified into two, that is, pneumonia and normal. The dataset has gained popularity as a standard on which deep learning and machine learning models are designed and trained to detect pneumonia. The imbalance in clinical classes is accurately reflected by the significant disparity in the pneumonia cases to the normal cases. The training set consists of 3883 pneumonia and 1349 normal images, which gives a rich source of labeled data to develop the models. The test set, used to evaluate model performance and generalization, contains 390 pneumonia and 234 normal images. The dataset distribution is summarized in Table 1, and representative images are presented in Figure 2.

Table 1
Class distribution of the Kermany chest X-ray dataset (training/test splits)

Class	Training image	Testing image	Total images
Normal	1349	234	1583
Pneumonia	3883	390	4273
Total images	5233	624	5856

3.2. Data preprocessing

The implementation of a structured pipeline during data preprocessing ensured consistency, reliability, and robustness across the dataset. CXR images were all resized to 224 × 224 pixel size to match the input size of the deep learning models. Pixel values were normalized to the [0, 1] range by dividing by 255, which stabilizes training and accelerates convergence. The preprocessing was applied to both the training and test sets to ensure a consistent input pipeline.

Figure 2
Illustration of normal and pneumonia images from the Kermany dataset

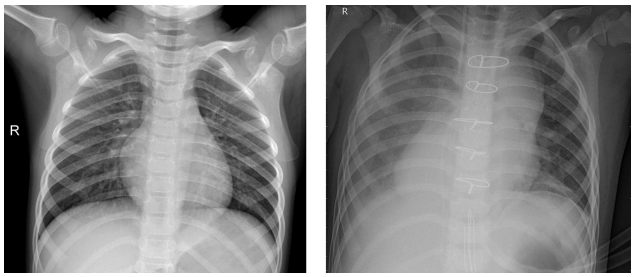

To reduce overfitting and improve model generalization, an extensive data augmentation strategy was employed on the training set. This involved a comprehensive set of real-time geometric and photometric transformations, including random rotations (up to 30 degrees), horizontal and vertical shifts (± 20% of the image width/height), shearing, and zooming. Furthermore, photometric variations were introduced via horizontal flipping, adjustments in brightness (range [0.8, 1.2]), and channel shifting. This robust augmentation protocol simulates a wide range of clinical imaging conditions, such as variations in patient positioning and X-ray machine settings, thereby forcing the model to learn invariant, pathognomonic features of pneumonia rather than memorizing spurious correlations in the training data. Figure 3 shows representative examples of the augmented images.

Table 2 summarizes the model hyperparameters and outlines the experimental setup used in this study.

3.3. CNN architectures

3.3.1. Benchmark CNN models

This study examines six widely used CNN architectures—DenseNet121, ResNet50, InceptionV3, EfficientNetB0, Mobile

Figure 3
Examples of augmented chest X-ray images used during training

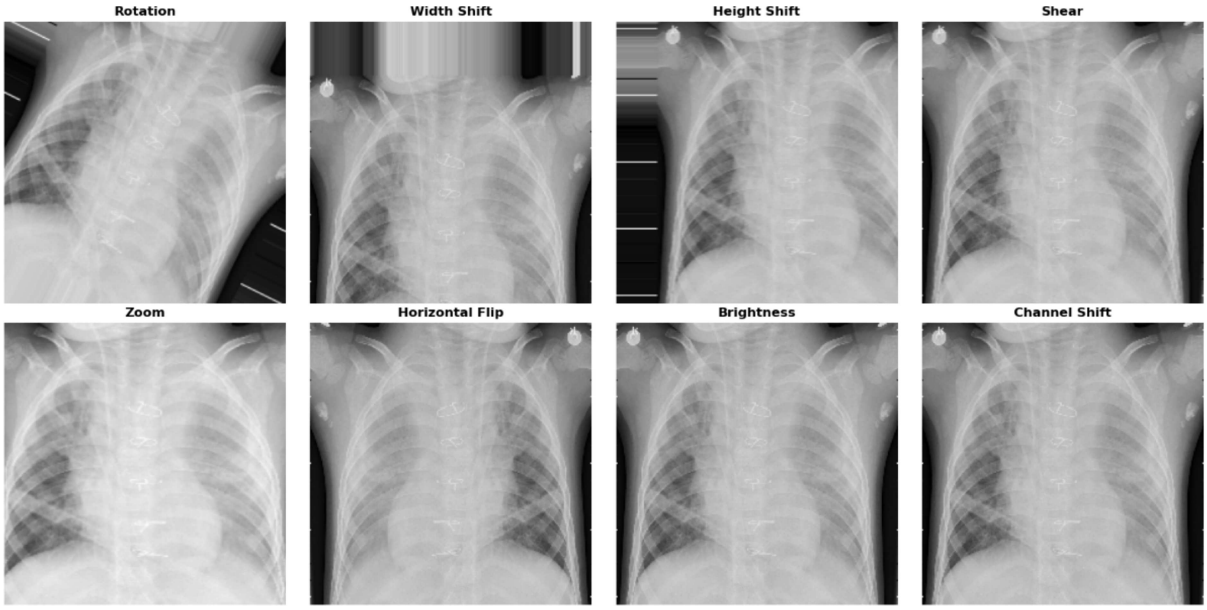

Table 2
Experimental setup and hyperparameters used for all models

Parameter/setup	Value/description
Dataset	Chest X-Ray Image Dataset [25]
Input image size	224 × 224 pixels
Preprocessing	Pixel normalization (0–1), resizing
Data augmentation	Rotation (±30°), W/H shift (±20%), shearing, zooming, horizontal flip, brightness adjustment ([0.8, 1.2])
Models evaluated	DenseNet121, ResNet50, InceptionV3, EfficientNetB0, MobileNetV2, AlexNet, proposed PneumoniaNet
Batch size	32
Number of epochs	30
Optimizer	Adam
Learning rate	1e-4
Loss function	Categorical cross-entropy
Hardware	NVIDIA Tesla V100 (16GB VRAM)
Software framework	TensorFlow 2.10/Keras 2.10

NetV2, and AlexNet—for classifying pneumonia from CXR images. Transfer learning was applied to ResNet50 by freezing the ImageNet-pretrained backbone, selectively fine-tuning the top residual blocks for XAI interpretability, and adding a task-specific classification head. While ResNet50's residual skip connections support deeper network training, DenseNet121's dense connectivity promotes feature reuse and mitigates vanishing gradient issues. InceptionV3 employs factorized convolutions for computational efficiency and parallel convolutional paths with varying kernel sizes to capture multi-scale features. Similarly, EfficientNetB0 was trained using transfer learning with the ImageNet-pretrained backbone frozen and the top blocks partially unfrozen for fine-tuning and XAI analysis, followed by a task-specific classification head. To achieve high accuracy with fewer parameters, EfficientNetB0 uses compound scaling to balance network width, depth, and resolution. For lightweight and efficient inference, MobileNetV2 employs inverted residual blocks with linear bottlenecks and depthwise separable convolutions. AlexNet, one of the earliest CNN architectures, consists of five convolutional layers followed by three fully connected layers. The additional classification head includes global average pooling, batch normalization, dropout, one or two dense layers, and a softmax layer for binary classification. This consistent approach enables a fair assessment of each architecture's ability to extract pathological features for accurate pneumonia detection.

3.3.2. Proposed PneumoniaNet

The proposed PneumoniaNet, illustrated in Figure 4, applies residual-dense connectivity immediately after the MobileNetV2 backbone. A three-layer convolutional module (Conv2D 512, 256, and 128) processes the 7 × 7 × 1280 feature map from MobileNetV2. Each layer receives the concatenated outputs of all

Figure 4
Layer architecture of the proposed PneumoniaNet model

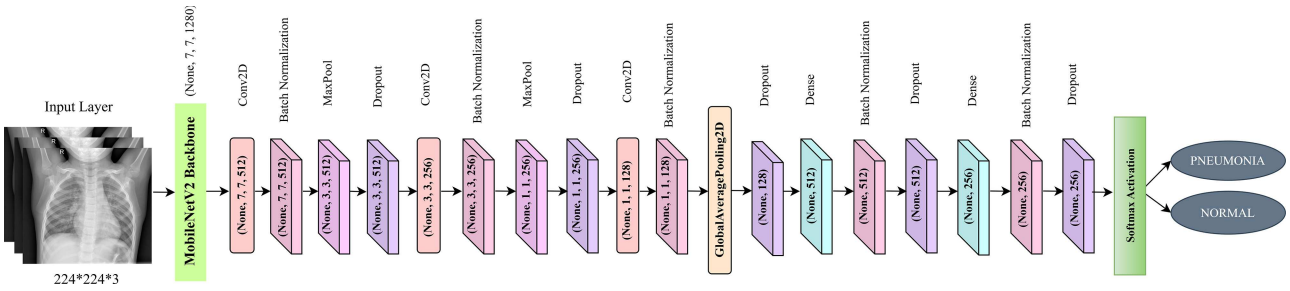

previous layers, forming a dense block that boosts feature reuse and strengthens local pattern extraction. To preserve semantic information and ensure smooth gradient flow, a residual shortcut adds the original MobileNetV2 output to the block's final output, bypassing the dense structure. Before global average pooling and the subsequent fully connected layers for classification, this residual-dense block enriches high-level feature representations.

3.4. Equations

The proposed PneumoniaNet architecture combines convolutional operations, dense feature connections, and residual learning to improve feature reuse and ensure smoother gradient flow throughout the network. The main mathematical formulations used in the model are outlined below.

Convolution Operation:

$$Z(i,j,k) = \Sigma_m \Sigma_n \Sigma_c X(i+m,\ j+n,\ c) \cdot W(m,n,c,k) + b_k \quad (1)$$

ReLU Activation:

$$A(i,\ j,\ k) = \max(0,\ Z(i,\ j,\ k)) \quad (2)$$

Dense Feature Aggregation:

$$F_l = H_l([F_0,\ \ F_1,\ F_2,\ \ldots,\ \ F_{l-1}]) \quad (3)$$

Residual-Dense Fusion:

$$F_{out} = F_{dense} + F_{MobileNetV2} \quad (4)$$

Selective Feature Learning:

$$F_{selected} = \alpha \cdot F_{low} + \beta \cdot F_{high} \quad (5)$$

4. Results and Performance Analysis

This section evaluates the proposed PneumoniaNet model, focusing on its efficiency, discriminative capability, and predictive performance in the automated detection of pediatric pneumonia.

4.1. Performance analysis

The proposed PneumoniaNet showed high performance relative to six trained CNN frameworks—DenseNet121, MobileNetV2, ResNet50, InceptionV3, EfficientNetB0, and AlexNet—to detect pneumonia in children by using CXR images. The overall results were best in PneumoniaNet, with an accuracy of 93.11% and an AUC of 0.9505. Comparatively, DenseNet121 had an accuracy of 89.10% and an AUC of 0.9165. ResNet50 achieved an accuracy of 87.18% and an AUC of 0.9018, while EfficientNetB0 had a slightly lower accuracy of 84.62% and an AUC of 0.8801. MobileNetV2 had an accuracy of 81.25% and an AUC of 0.8977. InceptionV3 had 79.33% accuracy and an AUC of 0.9199, while AlexNet had much lower values, 65.38% and 0.7732, respectively. In addition to its performance, PneumoniaNet was also found to have good computational efficiency with 570.56 MB of memory, 9.8 million trainable parameters, 1157.78 s of training time, and 1177.82 s of overall run time. DenseNet121 was able to achieve competitive accuracy but required much more resources, with 12.1 million parameters and 2257.51 s of training time. Lightweight architectures like AlexNet were also able to finish training faster but with much less accuracy.

These results demonstrate that PneumoniaNet achieves a strong balance of high accuracy and computational efficiency, making it practical for real-world deployment. Its superior AUC indicates robust feature representation and clear class

Table 3
Comparison of CNN architectures for pneumonia detection

Model	Accuracy (%)	AUC	Training time (s)	Execution time	Peak memory(MB)	Parameter(M)
AlexNet	65.38	0.7732	1017.61	1033.94	513.60	4,019,202
InceptionV3	79.33	0.9199	1719.84	1746.73	633.98	31,640,098
MobileNetV2	81.25	0.8977	1065.40	1085.80	548.90	8,556,354
EfficientNetB0	84.62	0.8801	2268.39	2291.39	629.35	10,347,941
ResNet50	87.18	0.9018	2705.06	2725.78	644.40	33,425,026
DenseNet121	89.10	0.9165	2257.51	2289.40	520.81	12,156,226
Proposed PneumoniaNet	93.11	0.9505	1157.78	1177.82	570.56	9,836,226

separation. Compared with baseline CNNs, the improvements suggest that the residual-dense block and selective fine-tuning effectively enhance pneumonia-specific feature extraction without imposing excessive computational cost. Table 3 presents the comparative performance and computational efficiency of CNN architectures for pneumonia detection.

Figures 5 and 6 depict the accuracy and loss of each model during both training and validation across 30 epochs. These figures provide a clear comparison of how each model's performance improves over the course of training.

4.1.1. Ablation study

The PneumoniaNet was tested with four different variants under the same training conditions to determine the contribution of each major architectural component. The unmodified backbone of MobileNetV2 (A1) had an accuracy of 81.25%, an AUC of 0.8977, and an F1-score of 86.83%, indicating that it alone could not effectively extract pneumonia-specific features. Adding the residual-dense block (A2) improved feature extraction compared to the baseline (A1), which demonstrates the significance of dense connectivity in the architecture. The performance increased to 86.38% accuracy, 0.915 AUC, and 88.15% F1, demonstrating the improvement achieved by incorporating the residual-dense block. Performance further rose to 89.10% accuracy, 0.928 AUC, and 90.60% F1 when shallow layers were frozen and deeper layers fine-tuned (A3), with deeper fine-tuning enhancing discriminative capability but introducing moderate overfitting. The overall performance of the complete PneumoniaNet model (A4) has the highest accuracy of 93.11%, AUC of 0.9505, and F1-score of 94.29%, indicating that residual-dense fusion combined with layer freezing yields the most generalizable feature representations. This shows that the residual-dense fusion and layer freezing together yield the best and most generalizable feature representations. The findings are a good justification of the architectural design decisions that have been taken in PneumoniaNet in the detection of pneumonia in children. Table 4 shows the ablation study of PneumoniaNet, summarizing the performance metrics of four architectural variants.

4.2. Receiver Operating Characteristic (ROC) analysis

The ROC was used to test the discriminative capability of the proposed PneumoniaNet model against a number of

Figure 5
Training and validation accuracy curves over 30 epochs for all models

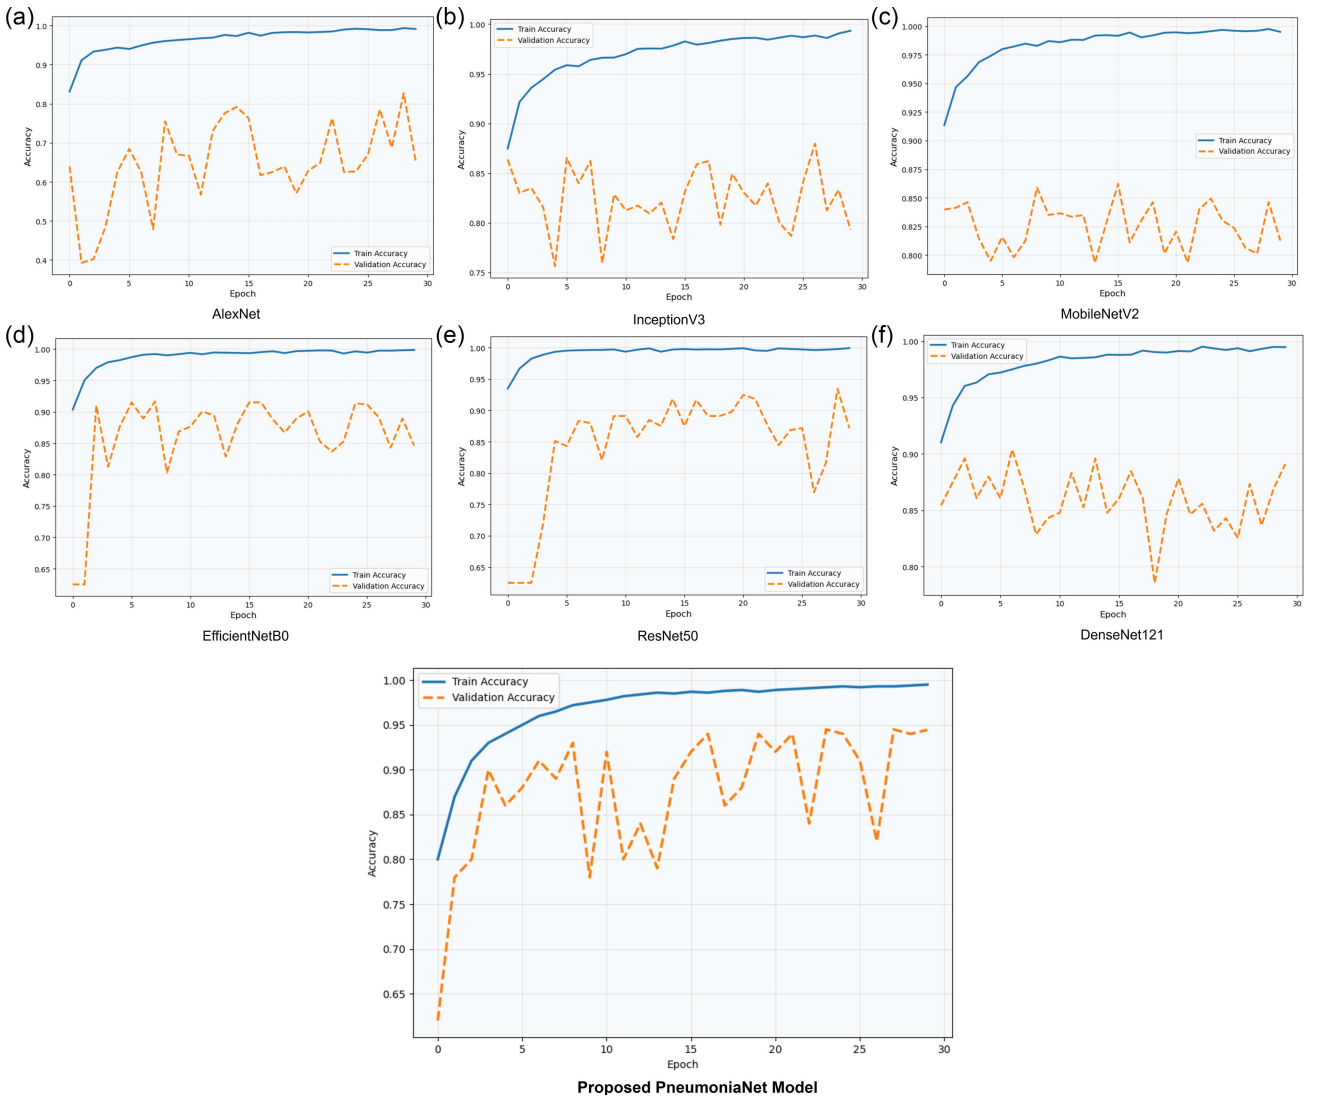

Figure 6
Training and validation loss curves over 30 epochs for all models

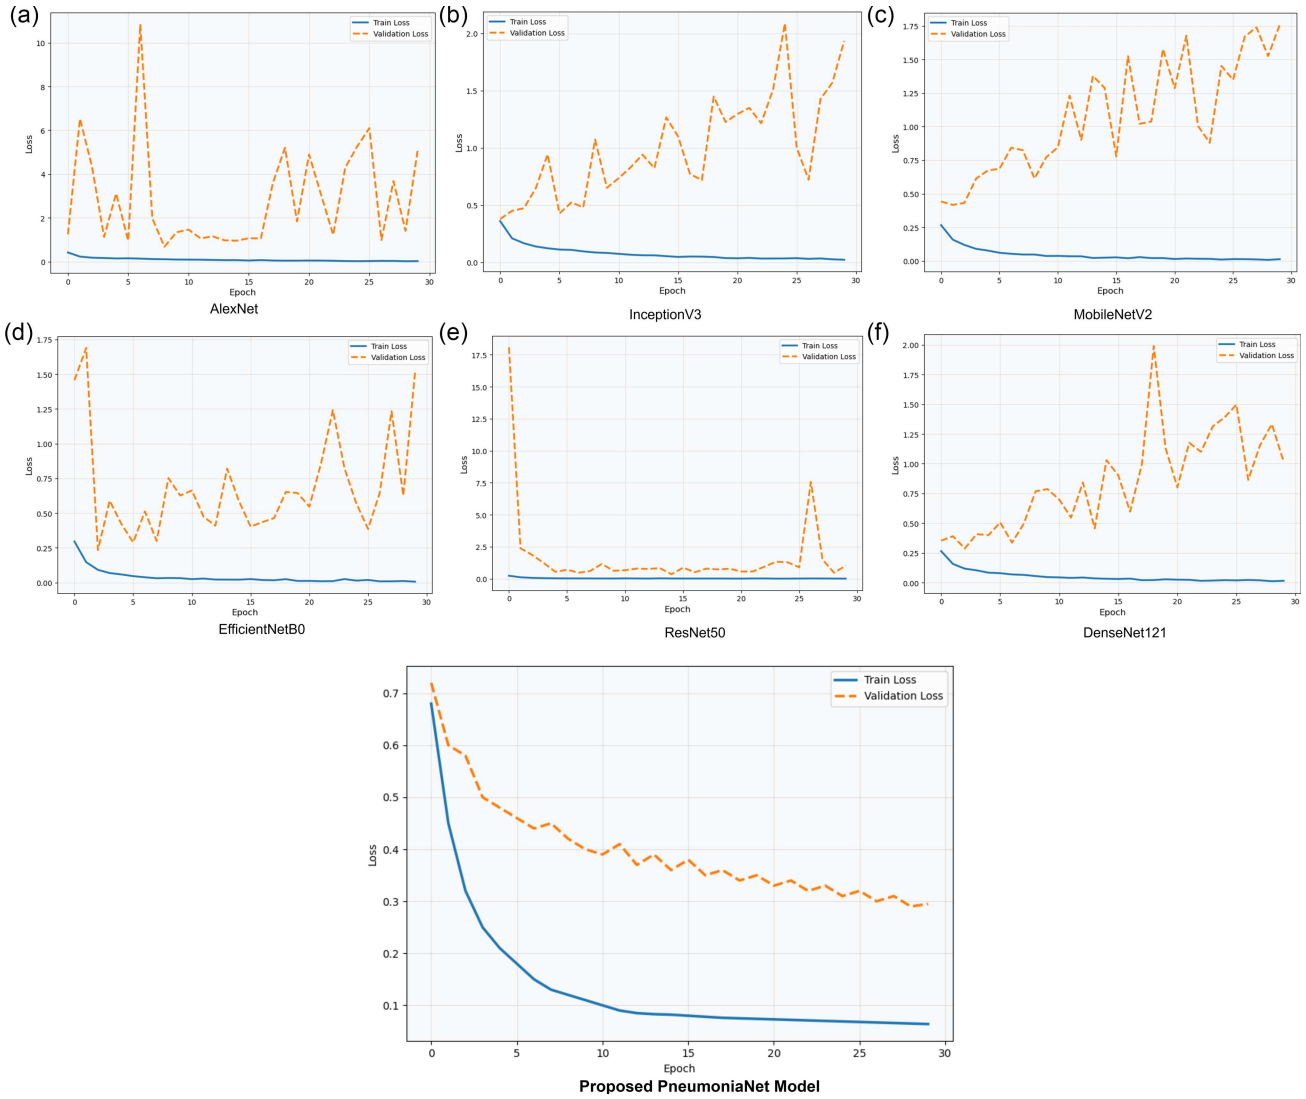

Table 4
Ablation study results for four variants of PneumoniaNet

Variant	Description	Accuracy (%)	F1-score (%)	AUC
A1	Baseline MobileNetV2 (no freezing, no dense block)	81.25	86.83	0.8977
A2	Residual-dense block (MobileNetV2 + classifier)	86.38	88.15	0.9150
A3	Frozen shallow layers (fine-tune all MobileNetV2 layers)	89.10	90.60	0.928
A4	Proposed PneumoniaNet (frozen shallow layers + residual-dense block)	93.11	94.29	0.9505

state-of-the-art CNN architectures, as illustrated in Figure 7. Each of the models showed different levels of class separability between normal and pneumonia CXR images in the curves of the ROC. PneumoniaNet had the best AUC value of 0.9505, showing the best feature representation and high diagnostic reliability. Comparatively, DenseNet121 and InceptionV3 achieved scores of 0.9165 and 0.9199 on the AUC, respectively, which is good but slightly worse at the discriminative level. ResNet50 achieved an AUC of 0.9018, while MobileNetV2 attained 0.8977. EfficientNetB0 had a slightly lower AUC of 0.8801, and AlexNet had a much lower AUC of 0.7732. The greater accuracy of pneumonia and normal cases with no overlap of PneumoniaNet is supported by the greater AUC that supports the strength, consistency, and clinical relevance of the tool in automated detection of pneumonia.

4.3. Confusion matrix analysis

The confusion matrix analysis, as shown in Figure 8, gives an overall assessment of how each model can differentiate between

Figure 7
ROC curves and AUC values for all evaluated models

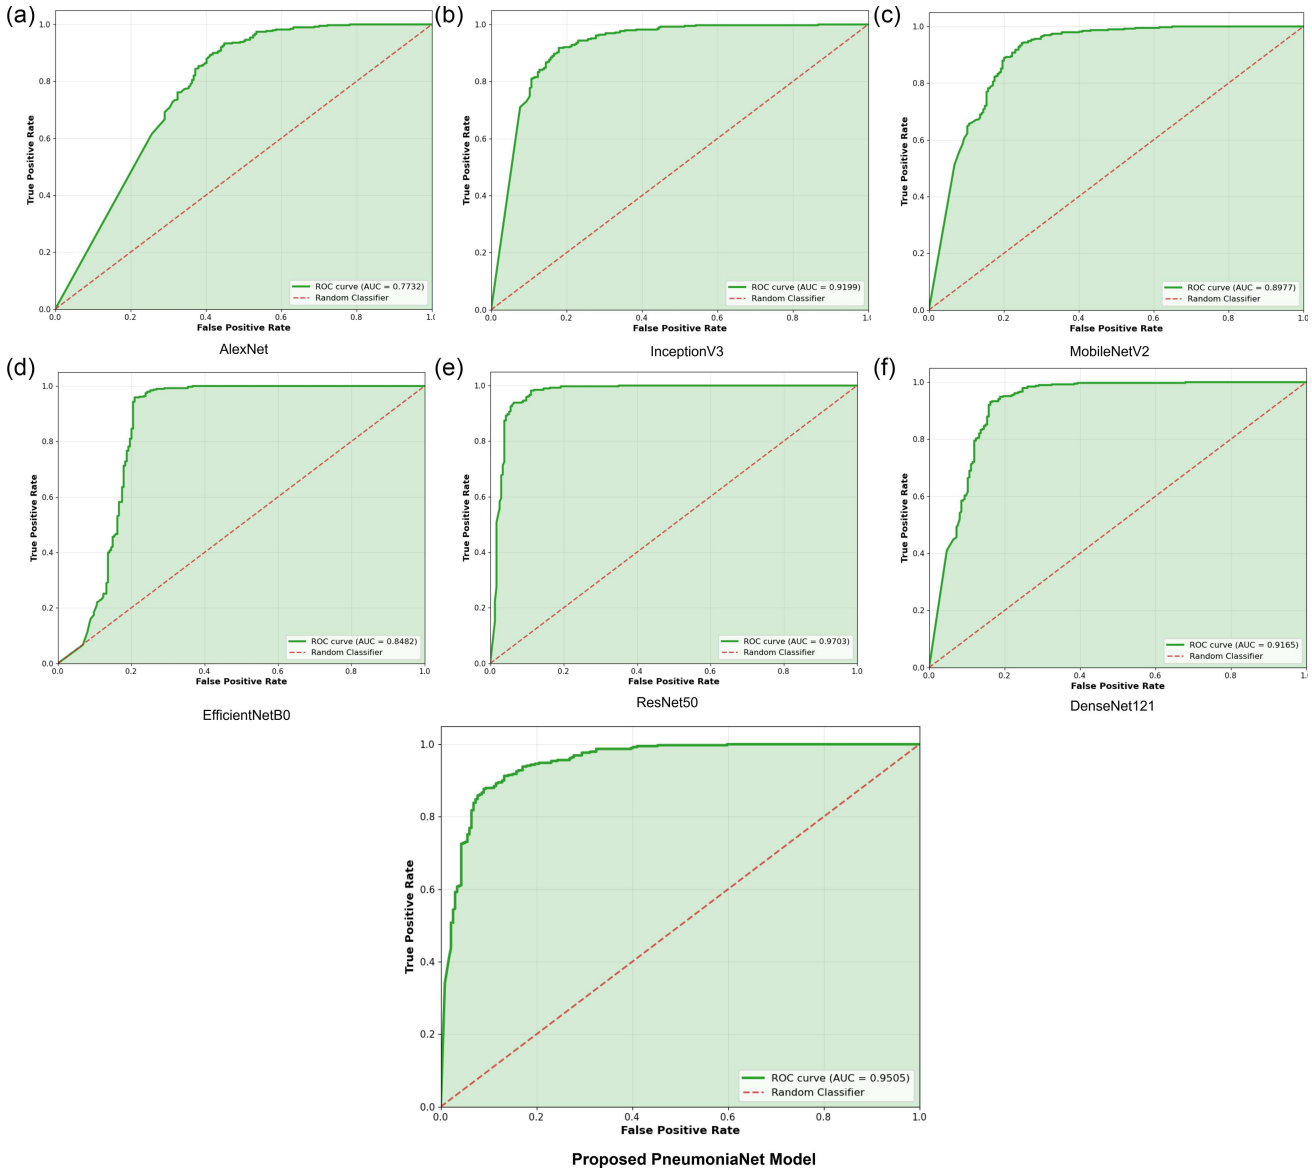

the two classes of CXR images. DenseNet121 performed well on generalization, where 173 normal and 383 pneumonia cases were correctly classified with the least amount of misclassification. Conversely, EfficientNetB0 correctly classified 138 pneumonia and 390 normal samples, with most misclassifications occurring as pneumonia samples being identified as normal, while ResNet50 correctly classified 155 pneumonia and 389 normal samples, with a small number of misidentified cases, mainly pneumonia samples as normal. Correspondingly, AlexNet also exhibited the same imbalance, with the model generating significantly more false pneumonia cases. Although InceptionV3 had a higher sensitivity and stability between categories, MobileNetV2 had a moderate balance as it correctly identified both categories with lesser false identifications. The most consistent and discriminative results were provided by the proposed PneumoniaNet that was able to classify 225 normal and 356 pneumonia samples with minimum errors. PneumoniaNet had outperformed all its baseline CNN models in balanced class recognition, reliable diagnostic accuracy, and precise feature representation.

4.4. Classification report analysis

The proposed PneumoniaNet exhibited a strong and balanced ability to differentiate between pneumonia and normal CXR images of pediatric patients, outperforming all comparative models. It scored highest in all the evaluation metrics with an accuracy of 93.11%, precision of 97.54%, recall of 91.26%, specificity of 96.16%, and F1-score of 94.29%, which indicates its effectiveness and reliability in pneumonia detection. DenseNet121 fared better with the highest recall of 98.21% though with low specificity of 73.93%. ResNet50 and EfficientNetB0 moderate results with accuracies of 87.18% and 84.62%, respectively, while AlexNet showed comparatively lower performance with an accuracy of 65.38%. MobileNetV2 and InceptionV3 showed moderate

Figure 8
Confusion matrices of all evaluated models

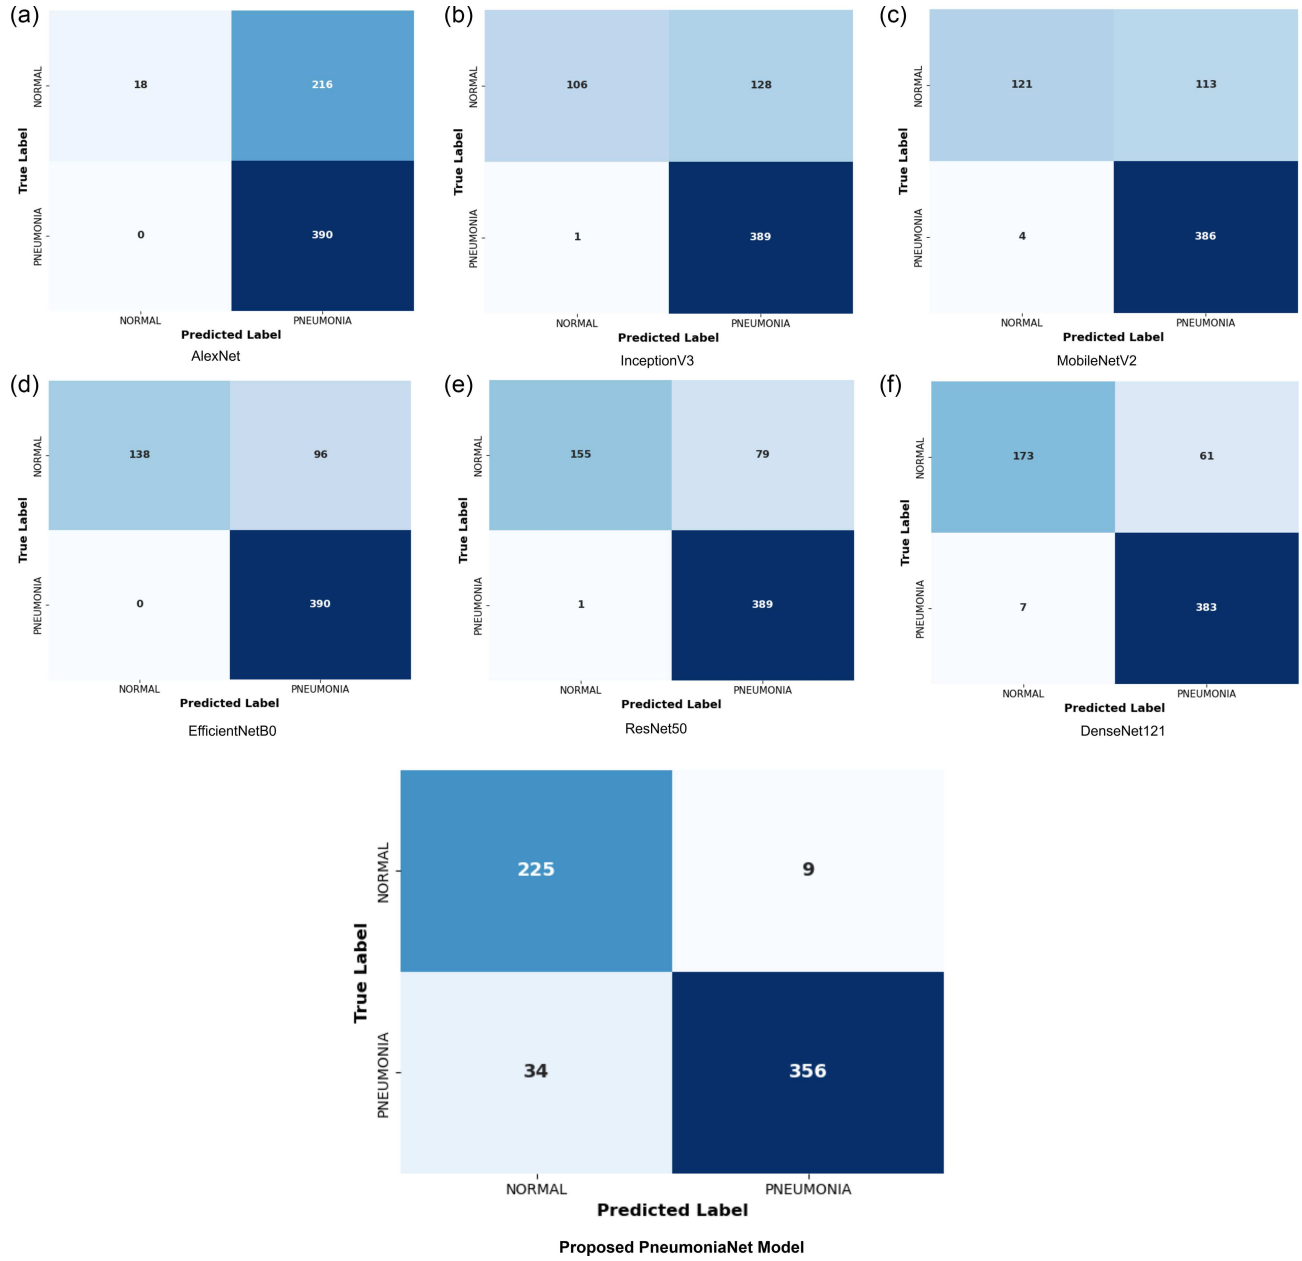

and more balanced results of the evaluation metrics. On balance, these findings indicate that PneumoniaNet can be used to automatically diagnose pediatric pneumonia because it is more accurate than baseline architectures and offers consistent and reliable differentiation. The classification overview of all models is shown in Table 5.

4.5. Comparison with existing works

The PneumoniaNet architecture was tested on the same pediatric CXR dataset and compared to various state-of-the-art models, as summarized in Table 6. CNN-XGBoost [26], Lightweight CNN [27], ResNet50-v2 [28], and Simple CNN [29] achieved accuracies of 87.00%, 89.89%, 90.87%, and 92.00%, respectively, while ResNet50 [30] reached 93.06%. In comparison, PneumoniaNet achieved the highest accuracy of 93.11%, demonstrating its strong ability to extract discriminative features for pneumonia detection. This outcome can be explained by the fact that its optimized deep architecture and explainable AI techniques have been implemented, which further improve the accuracy of diagnostics and explainability. Overall, the findings indicate that PneumoniaNet is an effective and reliable automated pediatric pneumonia diagnostic system.

4.6. Model interpretability with Grad-CAM visualizations

Grad-CAM was utilized to visualize the regions of CXRs to which the model was most sensitive, enhancing interpretability and clinical confidence in PneumoniaNet. Figure 9 presents Grad-CAM heatmaps overlaid on images for both correctly and incorrectly classified cases. Regions showing the highest activations are highlighted in red and yellow, pointing out the features that made the most contributions to decisions made by the model.

Table 5
Classification metrics (accuracy, precision, recall, specificity, F1) for all models

Model	Accuracy (%)	Precision (%)	Recall (%)	Specificity (%)	F1-score (%)
AlexNet	65.38	64.35	100.00	7.69	78.30
InceptionV3	79.33	75.22	99.74	45.30	85.67
MobileNetV2	81.25	77.35	98.97	51.71	86.83
EfficientNetB0	84.62	80.25	100.00	58.97	89.03
ResNet50	87.18	83.12	99.74	66.24	90.73
DenseNet121	89.10	86.26	98.21	73.93	91.81
Proposed PneumoniaNet	93.11	97.54	91.26	96.16	94.29

Table 6
Performance comparison with prior studies using the same dataset

Model	Accuracy (%)
CNN-XGBoost [26]	87.00
Lightweight CNN [27]	89.89
ResNet50-v2 [28]	90.87
Simple CNN [29]	92.00
ResNet50 [30]	93.06
Proposed PneumoniaNet model	93.11

Figure 9
Grad-CAM heatmaps for correctly and incorrectly classified cases by PneumoniaNet

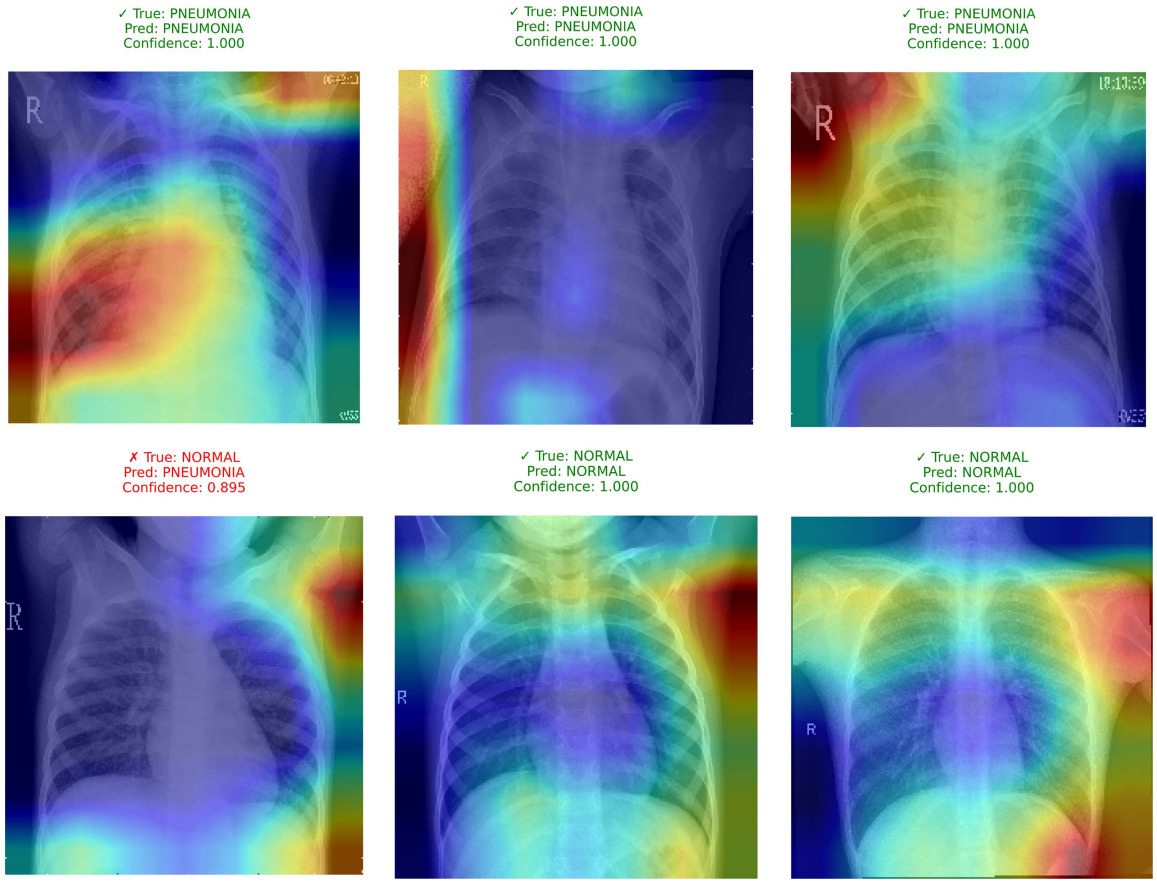

In true positive and true negative prediction, these erected areas are close to areas of the lung which are usually studied by radiologists, indicating that the model is capable of truly localizing normal pulmonary structures as well as the opacities that are caused by pneumonia. Misclassified samples, on the other hand, would exhibit more diffuse or off-target activations, meaning there is some ambiguity in the anatomy or imaging artifact that makes proper classification challenging. These visualizations confirm that PneumoniaNet focuses on clinically relevant radiographic patterns rather than irrelevant background features.

5. Conclusion

This study describes PneumoniaNet, a refined deep learning architecture based on MobileNetV2, which is trained to automatically detect pediatric pneumonia through CXRs. The model was trained and evaluated with the publicly available Kermany CXR dataset, and by combining frozen shallow and fine-tuned deeper layers, it extracts more disease-specific image features while sacrificing some general image features. Comparisons of the performance with a number of CNN baselines such as DenseNet121, InceptionV3, EfficientNetB0, ResNet50, and AlexNet showed that PneumoniaNet performed better than all of them in fundamental classification metrics, such as accuracy, precision, recall, specificity, and the F1-score. Moreover, Grad-CAM-based visualizations ensured that the network focused on clinically relevant areas rather than irrelevant background. This would offer reliable and easy-to-interpret diagnostic assistance to medical practitioners.

This study is associated with various limitations despite its promising results. First, the model was trained and evaluated using a single publicly available dataset, which might not be enough to reflect the diversity of real-world clinical practice. The differences in the quality of radiographs, demographics of the patients, the prevalence of the disease, and imaging equipment in healthcare institutions may affect the model's generalizability. Second, the existing framework treatment is limited to binary classification (normal vs. pneumonia) and does not deal with multi-class distinction between other respiratory conditions, potentially limiting its application in more complex cases. Lastly, the dataset is not large in comparison to large-scale clinical repositories, which may limit the model's ability to detect subtle pathological differences.

Future studies could explore several approaches to enhance the clinical usability and robustness of PneumoniaNet. Increasing the data of multi-institutional and real-world clinical images would enhance the generalizability of the model. Transformer-based hybrid structures and superior attention models would further improve the feature representation and localization accuracy. The use of multimodal data, including patient history, laboratory findings, and clinical records, may assist in a more complex decision-making system. Further research should be conducted to test the model on external clinical data to prove its strength further. Future research will also employ cross-validation to strengthen experimental results and validate the model on additional external datasets. Finally, the PneumoniaNet application in the clinical context, where continuous learning frameworks and prospective validation will be implemented, will be necessary in achieving correct, interpretable, and reliable AI-assisted pneumonia detection.

Ethical Statement

The authors declare that this study did not require formal ethical approval because it was conducted using a publicly available and anonymized pediatric CXR dataset. According to institutional and national research guidelines in Bangladesh, studies based solely on secondary analysis of publicly accessible data without direct human or animal involvement do not require Institutional Review Board approval. The dataset used in this research was originally collected and ethically approved by the respective medical institution and is publicly available on the Kaggle platform.

Conflicts of Interest

The authors declare that they have no conflicts of interest to this work.

Data Availability Statement

The data that support the findings of this study are openly available in Kermany's CXR Images Datasets [Kaggle] at https://www.kaggle.com/datasets/riyadhhalmosawi/kermanys-cxr-images-datasets.

Author Contribution Statement

Shahriar Siddique Arjon: Conceptualization, Methodology, Software, Investigation, Data curation, Writing – original draft, Writing – review & editing. **Tamanna Yasmin:** Software, Investigation, Resources, Writing – original draft, Writing – review & editing, Visualization. **Ankur Kumar Mondol:** Software, Formal analysis. **Nakib Aman:** Supervision, Project administration. **Shabbir Mahmood:** Validation.

References

[1] Kundur, N. C., Anil, B. C., Dhulavvagol, P. M., Ganiger, R., & Ramadoss, B. (2023). Pneumonia detection in chest x-rays using transfer learning and TPUs. *Engineering, Technology & Applied Science Research*, *13*(5), 11878–11883. https://doi.org/10.48084/etasr.6335

[2] Bal, U., Bal, A., Moral, Ö. T., Düzgün, F., & Gürbüz, N. (2024). A deep learning feature extraction-based hybrid approach for detecting pediatric pneumonia in chest X-ray images. *Physical and Engineering Sciences in Medicine*, *47*(1), 109–117. https://doi.org/10.1007/s13246-023-01347-z

[3] Singh, S., Kumar, M., Kumar, A., Verma, B. K., & Shitharth, S. (2023). Pneumonia detection with QCSA network on chest X-ray. *Scientific Reports*, *13*(1), 9025. https://doi.org/10.1038/s41598-023-35922-x

[4] Singh, S., Rawat, S., Gupta, M., Tripathi, B., Alanzi, F., Majumdar, A., . . . , & Thinnukool, O. (2023). Deep attention network for pneumonia detection using chest X-ray images. *Computers, Materials, & Continua*, *74*(1), 1673–1691. http://dx.doi.org/10.32604/cmc.2023.032364

[5] Ibrahim, A. U., Ozsoz, M., Serte, S., Al-Turjman, F., & Yakoi, P. S. (2024). Pneumonia classification using deep learning from chest X-ray images during COVID-19. *Cognitive Computation*, *16*(4), 1589–1601. https://doi.org/10.1007/s12559-020-09787-5

[6] Bhatt, H., & Shah, M. (2023). A convolutional neural network ensemble model for pneumonia detection using chest x-ray images. *Healthcare Analytics*, *3*, 100176. https://doi.org/10.1016/j.health.2023.100176

[7] Maniruzzaman, M. (2024). Pneumonia prediction using deep learning in chest X-ray Images. *International Journal of Science and Research Archive*, *12*(1), 767–773. https://doi.org/10.30574/ijsra.2024.12.1.0880

[8] Zhong, Y., Liu, Y., Gao, E., Wei, C., Wang, Z., & Yan, C. (2024). Deep learning solutions for pneumonia detection: Performance comparison of custom and transfer learning models. In *International Conference on Automation and Intelligent Technology*, *13401*, 95–100. https://doi.org/10.1117/12.3053114

[9] Antunes, C., Rodrigues, J. M., & Cunha, A. (2025). PneumoNet: Artificial intelligence assistance for pneumonia detection on X-rays. *Applied Sciences*, *15*(13), 7605. https://doi.org/10.3390/app15137605

[10] Novokshanov, A. (2024). Using explainable artificial intelligence to locate pneumonia. *Research Archive of Rising Scholars*, 1–14. https://doi.org/10.58445/rars.1724

[11] Yasmin, T., Arjon, S. S., Mondol, A. K., Aman, N., & Mahmood, S. (2025). Pneumonia detection from chest x-ray images based on deep learning. In *International Conference on Data Science, AI and Applications*, 252–266. https://doi.org/10.1007/978-3-032-11352-8_18

[12] Sharma, S., & Guleria, K. (2023). A deep learning based model for the detection of pneumonia from chest x-ray images using VGG-16 and neural networks. *Procedia Computer Science*, *218*, 357–366. https://doi.org/10.1016/j.procs.2023.01.018

[13] Siddiqi, R., & Javaid, S. (2024). Deep learning for pneumonia detection in chest x-ray images: A comprehensive survey. *Journal of Imaging*, *10*(8), 176. https://doi.org/10.3390/jimaging10080176

[14] Wu, L., Zhang, J., Wang, Y., Ding, R., Cao, Y., Liu, G., . . . , & Guan, F. (2024). Pneumonia detection based on RSNA dataset and anchor-free deep learning detector. *Scientific Reports*, *14*(1), 1929. https://doi.org/10.1038/s41598-024-52156-7

[15] Bakır, H., Oktay, S., & Tabaru, E. (2023). Detection of pneumonia from x-ray images using deep learning techniques. *Journal of Scientific Reports-A*, *052*, 419–440. https://doi.org/10.59313/jsr-a.1219363

[16] Yi, R., Tang, L., Tian, Y., Liu, J., & Wu, Z. (2023). Identification and classification of pneumonia disease using a deep learning-based intelligent computational framework. *Neural Computing and Applications*, *35*(20), 14473–14486. https://doi.org/10.1007/s00521-021-06102-7

[17] Kaushik, P., Arora, M., Sharma, Y., Poonia, M., Kumawat, P., & Charak, R. S. (2024). PneumoAI: Redefining accuracy in pneumonia detection using advanced machine learning. In *2024 IEEE International Conference on Interdisciplinary Approaches in Technology and Management for Social Innovation*, *2*, 1–6. https://doi.org/10.1109/IATMSI60426.2024.10503052

[18] Prakash, V. S., Aiyyasamy, V., Shobana, D., Jayadurga, R., Kandaswamy, V. A., & Murugan, P. S. B. (2023). Unleashing the potential of artificial intelligence and deep learning in pneumonia detection systems. In *2023 9th International Conference on Smart Structures and Systems*, 1–6. https://doi.org/10.1109/ICSSS58085.2023.10407052

[19] Colin, J., & Surantha, N. (2025). Interpretable deep learning for pneumonia detection using chest x-ray images. *Information*, *16*(1), 53. https://doi.org/10.3390/info16010053

[20] Das, R., Nayak, D. S. K., Rout, C. P., Jena, L., & Swarnkar, T. (2024). Deep learning techniques for identification of pneumonia: A CNN approach. In *2024 International Conference on Advancements in Smart, Secure and Intelligent Computing*, 1–5. https://doi.org/10.1109/ASSIC60049.2024.10507933

[21] Ren, H., Wong, A. B., Lian, W., Cheng, W., Zhang, Y., He, J., . . . , & Zhang, H. (2021). Interpretable pneumonia detection by combining deep learning and explainable models with multisource data. *IEEE Access*, *9*, 95872–95883. https://doi.org/10.1109/ACCESS.2021.3090215

[22] Sheu, R. K., Pardeshi, M. S., Pai, K. C., Chen, L. C., Wu, C. L., & Chen, W. C. (2023). Interpretable classification of pneumonia infection using eXplainable AI (XAI-ICP). *IEEE Access*, *11*, 28896–28919. https://doi.org/10.1109/ACCESS.2023.3255403

[23] Papadimitriou, O., Kanavos, A., & Maragoudakis, M. (2023). Automated pneumonia detection from chest x-ray images using deep convolutional neural networks. In *2023 14th International Conference on Information, Intelligence, Systems & Applications (IISA)*, 1–4. https://doi.org/10.1109/IISA59645.2023.10345859

[24] Ali, M., Shahroz, M., Akram, U., Mushtaq, M. F., Altamiranda, S. C., Obregon, S. A., . . . , & Ashraf, I. (2024). Pneumonia detection using chest radiographs with novel efficientnetv2l model. *IEEE Access*, *12*, 34691–34707. https://doi.org/10.1109/ACCESS.2024.3372588

[25] Kermany, D. S., Goldbaum, M., Cai, W., Valentim, C. C., Liang, H., Baxter, S. L., . . . , & Zhang, K. (2018). Identifying medical diagnoses and treatable diseases by image-based deep learning. *Cell*, *172*(5), 1122–1131. https://doi.org/10.1016/j.cell.2018.02.010

[26] Hariri, M., & Avşar, E. (2023). COVID-19 and pneumonia diagnosis from chest X-ray images using convolutional neural networks. *Network Modeling Analysis in Health Informatics and Bioinformatics*, *12*(1), 17. https://doi.org/10.1007/s13721-023-00413-6

[27] Iparraguirre-Villanueva, O., Robles-Espiritu, W., Suxe-Ramírez, M., & Flores-Castañeda, R. O. (2025). Use of CNN with transfer learning to improve the accuracy of pneumonia diagnosis. *International Journal of Engineering Trends and Technology*, *73*(6), 382–395. https://doi.org/10.14445/22315381/IJETT-V73I6P131

[28] Manickam, A., Jiang, J., Zhou, Y., Sagar, A., Soundrapandiyan, R., & Samuel, R. D. J. (2021). Automated pneumonia detection on chest X-ray images: A deep learning approach with different optimizers and transfer learning architectures. *Measurement*, *184*, 109953. https://doi.org/10.1016/j.measurement.2021.109953

[29] Nageye, A. Y., Jimale, A. D., Abdullahi, M. O., Ahmed, Y. A., & Addow, M. A. (2025). Enhancing deep learning for pneumonia detection: Developing web based solution for Dr. Sumait Hospital in Mogadishu Somalia. *Discover Applied Sciences*, *7*(4), 309. https://doi.org/10.1007/s42452-025-06735-6

[30] Hedhoud, Y., Mekhaznia, T., & Amroune, M. (2023). An improvement of the CNN-XGboost model for pneumonia disease classification. *Polish Journal of Radiology*, *88*, 483–493. https://doi.org/10.5114/pjr.2023.132533

How to Cite: Arjon, S. S., Yasmin, T., Mondol, A. K., Aman, N., & Mahmood, S. (2026). A Novel PneumoniaNet Framework Integrating Explainable AI for Pediatric Pneumonia Detection from Chest X-rays. *Artificial Intelligence and Applications*. https://doi.org/10.47852/bonviewAIA62028511