Pneumonia Detection from Chest X-Ray Images Using Deep Learning with Smart Validation

Kuushal Kamboji
Department of Computer Science and Engineering
JAIN (Deemed-to-be University)
Ramanagara, India.

Meda Rakshith Kumar
Department of Computer Science and Engineering
JAIN (Deemed-to-be University)
Ramanagara, India

M Sai Hashwanth Reddy
Department of Computer Science and Engineering
JAIN (Deemed-to-be University)
Ramanagara, India

P Sai Venkat
Department of Computer Science and Engineering
JAIN (Deemed-to-be University)
Ramanagara, India

Ajay Kumar Singh
Department of Computer Science and Engineering
JAIN (Deemed-to-be University)
Ramanagara, India

Abstract— **Pneumonia is among the most significant causes of death in the world and early diagnosis is highly valued for effective treatment. This paper presents a deep learning-based system for pneumonia detection using chest X-ray images with Convolutional Neural Networks (CNN), ResNet, and MobileNet models. The primary objective is to correctly classify chest X-ray images as Pneumonia or Normal. A smart validation system is incorporated to eliminate irrelevant images (e.g., non-medical X-rays) and enhance system reliability. MobileNet is applied to identify non-X-ray images and validate chest X-rays. The multi-step methodology demonstrates the capability to detect pneumonia while filtering irrelevant images that could cause false classifications. The models were trained and verified on a large sample of chest X-ray images with high classification accuracy. The proposed system shows strong potential to assist healthcare providers in rapidly and accurately identifying pneumonia while processing only valid medical images.**

Keywords: Pneumonia diagnosis, Chest X-ray, Deep learning, CNN, ResNet, MobileNet, Image classification, Smart validation, Irrelevant image detection, Medical imaging.

I. INTRODUCTION

Pneumonia is a severe respiratory disease that causes significant morbidity and mortality worldwide. The World Health Organization (WHO) estimates that pneumonia among children under the age of 5 is among the leading causes of death in children, accounting for approximately 15 percent of child deaths, making it a significant global health issue. Early and accurate diagnosis is critical to successful treatment and early intervention, particularly in low-resource facilities where availability of health workers and diagnostic equipment may be limited. Clinical examination and laboratory tests, the most prevalent pneumonia diagnostic methods, are not only time-consuming and invasive, but also subject to human error. Thus, the need for automated systems capable of diagnosing pneumonia faster and more accurately is increasingly important.

Chest X-ray imaging has been regarded as the gold standard in pneumonia diagnosis due to its non-invasive nature and its ability to image the inner structures of the lungs. However, manual X-ray reading by radiologists is subjective, time-consuming, and cannot always be applied in settings with limited medical expertise. The increasing availability of large medical image datasets and advances in deep learning technologies have enabled the development of automated image analysis systems that can effectively detect pneumonia in chest X-rays with a high degree of accuracy.

Convolutional neural networks (CNNs) and deep learning have proven extremely effective in medical imaging tasks, including pneumonia detection. CNNs are highly suitable for image classification as they can learn hierarchical features and patterns directly from image data. Recent advances such as ResNet and MobileNet architectures have further increased the performance of CNN-based models through deeper network architectures and efficient computation, enabling their application in analyzing large-scale medical images.

Despite their effectiveness, deep learning models for pneumonia detection face several challenges. One significant issue is the presence of irrelevant images, such as non-medical images or poor-quality X-rays, which can adversely affect model performance through misclassifications and reduced reliability. To address this, the proposed research introduces a multi-step pipeline that incorporates a smart validation system to filter out irrelevant images. Specifically, MobileNet — a lightweight and efficient deep learning model — is used to distinguish between valid chest X-ray images and non-relevant images, ensuring that only medical images are passed to the pneumonia detection models.

The primary goal of this study is to develop a viable pneumonia detection system using CNN, ResNet, and MobileNet. The system is designed to classify chest X-ray images as either Pneumonia or Normal, while also handling the problem of irrelevant image inputs through an intelligent validation step. The models are trained and tested on a large chest X-ray dataset to achieve high and consistent classification accuracy. [1]

II. RELATED WORK

Recent advances in deep learning have transformed pneumonia detection from chest X-ray (CXR) images, making it more precise and efficient. Several studies have deployed Convolutional Neural Networks (CNNs) to automate the detection process. For instance, CNN models such as VGG16 and ResNet have been used for pneumonia classification, demonstrating high accuracy in differentiating between pneumonia and normal CXRs. These models, trained on large datasets, are capable of learning hierarchical image features. ResNet-50 and DenseNet models have become especially popular for processing complex image characteristics while avoiding overfitting. [2]

More recent research in pneumonia detection has embraced transfer learning, wherein models such as VGG16, ResNet, and InceptionNet — pre-trained on general image datasets — are fine-tuned for medical image classification. This technique has been effective in addressing the problem of limited medical training data by leveraging knowledge from large general datasets. Furthermore, ensemble methods combining multiple CNNs, including GoogLeNet, ResNet-18, and DenseNet-121, have been investigated to improve classification rates and model stability. [3]

MobileNet has gained recognition in recent years as a lightweight CNN architecture that offers faster inference time and lower computational cost compared to traditional CNN models, making it suitable for real-time applications. Several studies have applied MobileNet to predict pneumonia from CXR images with competitive accuracy and efficiency. Additionally, attention-based deep learning models and Vision Transformers (ViTs) have been proposed to enhance pneumonia detection by focusing on the most significant regions of X-ray images, thereby improving prediction precision. Hybrid combinations of CNNs and ViTs have shown strong performance on complex image classification tasks.

To optimize pneumonia detection, various researchers have proposed validation steps to eliminate irrelevant or poor-quality images, ensuring that only valid medical images are processed by deep learning models. Such multi-stage approaches, in which models like MobileNet are employed to identify and discard irrelevant images, have successfully reduced false classifications caused by non-medical inputs. [4]

III. DATASET

1) Description

The dataset utilized in this study was obtained from Kaggle and specifically targets pneumonia identification in chest X-ray images. It comprises two main categories — PNEUMONIA and NORMAL — based on the presence or absence of pneumonia in the X-ray images. The dataset contains 5,216 training images, 16 validation images, and 624 test images, providing sufficient data for training and evaluating the model.

The training set is the largest subset, consisting of 4,273 images labeled as PNEUMONIA and 1,343 labeled as NORMAL. This distribution reflects a significant class imbalance, with the PNEUMONIA class having considerably more samples than the NORMAL class. Class imbalance is a common issue in medical image classification and may affect model generalization, particularly the ability to correctly classify both classes equally well.

The data is divided into three subsets: training, validation, and testing. The majority of the images reside in the training set, with 5,216 images used to train the deep learning models. The validation set is small (16 images) but serves an important role in assessing model performance during training and preventing overfitting.

IV. ARCHITECTURE DETAILS

The system architecture is designed to process pneumonia detection based on chest X-ray images through several key modules. The process begins when a user uploads a chest X-ray image to the system. The uploaded image is then verified by the MobileNet model, which determines whether the image is relevant for pneumonia identification. If the image is found to be relevant, it is forwarded to the image preprocessing stage, where it is prepared for analysis. Data augmentation follows preprocessing and is applied to improve model performance and generalization, as shown in Figure 1.

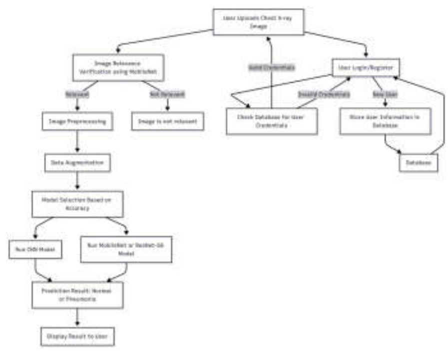

Figure 1: Architecture Diagram

For a relevant image, the system selects from two pre-trained models: the CNN model or the ResNet-50 model. The appropriate model is chosen based on prior performance and its suitability for the specific characteristics of the chest X-ray image. The selected model is then applied to determine whether the image shows evidence of pneumonia or is normal. The final output is the prediction result — either Pneumonia or Normal — which is returned to the user. [5]

In the case where the uploaded image is deemed irrelevant (i.e., it is not a chest X-ray), the system discards it without further processing. This intelligent validation step ensures that the detection models process only legitimate medical images, thereby improving overall system accuracy and preventing unnecessary computations. [6]

V. PROPOSED METHODOLOGY

1) Overview:

The system is designed to support pneumonia detection using deep learning models applied to chest X-ray images. When a user inputs a chest X-ray image, its relevance is confirmed by a MobileNet model. If the image is relevant, it is preprocessed and augmented to prepare it for classification. The system then selects a suitable model — either CNN or ResNet-50 — based on accuracy to detect pneumonia. The model processes the image and produces a prediction indicating whether pneumonia is present or the image is normal, which is then delivered to the user. If the image is irrelevant (e.g., a non-medical image), it is skipped with no further processing, resulting in an efficient and accurate system.

The system also includes a user login and registration module alongside the image analysis process. Users can log in with existing credentials or create a new account, with information stored in a database. This feature enables individual access to the system and protects user data. Overall, the system combines image relevance validation, deep learning-based pneumonia detection, and user management to deliver a secure, efficient, and reliable automated pneumonia detection service. [7]

2) MobileNet:

MobileNet is a lightweight, high-performing neural network model designed for edge computing and mobile devices. It is particularly suitable when computational resources are limited, making it an ideal choice for the image relevance checking component of the pneumonia detection system. The primary contribution of MobileNet in this system is filtering out unwanted images before further processing. It is trained to distinguish between medical chest X-ray images and non-medical or irrelevant images. MobileNet achieves this with fewer parameters and computations compared to classical CNN architectures, resulting in reduced inference time without compromising result quality.

Once trained, the model classifies new images as relevant (i.e., medical chest X-rays) or irrelevant (e.g., non-medical images). This pre-processing step ensures that only valid X-ray images are introduced to the pneumonia detection models, minimizing unnecessary computation and improving overall system efficiency and reliability. Images classified as irrelevant are discarded, while relevant images proceed to preprocessing, data augmentation, and model selection. [8], [9]

3) ResNet-50:

ResNet-50 is a deep convolutional neural network based on the Residual Network (ResNet) architecture. The key innovation of ResNet is the use of residual blocks with skip connections, which enable the model to preserve meaningful features across many layers without encountering the vanishing gradient problem. By learning residual mappings rather than direct mappings, the network can train more effectively. ResNet-50, with its 50 layers, has been widely applied in image classification tasks and is well-suited for analyzing medical images to identify patterns indicative of pneumonia.

The system uses a ResNet-50 model pre-trained on the large ImageNet dataset, which is then fine-tuned on chest X-ray data for pneumonia classification. Fine-tuning involves adjusting the final layers of the pre-trained model to adapt it to the specific task of medical image classification. [10], [11]

4) CNN Model:

The CNN model forms the foundational deep learning architecture employed to classify chest X-ray images in the pneumonia detection system. CNNs are widely used in image classification due to their ability to automatically learn spatial hierarchies of features from image data without requiring manual feature extraction. The CNN in this system is designed to detect the presence of pneumonia-specific features in chest X-rays or classify them as normal. It is trained on a large number of chest X-ray images categorized as either PNEUMONIA or NORMAL, using multiple convolutional layers, activation functions (ReLU), pooling layers, and fully connected layers. [12], [13]

The convolutional layers automatically learn patterns in the images such as edges, textures, and more complex characteristics such as pneumonia-related anomalies in the lung area. The learned features are passed to pooling layers that downsample the spatial dimensions to reduce computational load, and then to fully connected layers that generate the final classification output. Dropout layers may also be applied during training to prevent overfitting by randomly deactivating a percentage of neurons at each iteration. [14], [15]

VI. RESULTS AND DISCUSSION

MobileNet:

The MobileNet model achieves an accuracy of 94% on the pneumonia detection task, as shown in its confusion matrix. In the matrix, actual labels are shown on the vertical axis and predicted labels on the horizontal axis. The upper-left cell shows 8 True Negatives (TN), indicating images correctly classified as Normal. The lower-right cell shows 7 True Positives (TP), indicating pneumonia images correctly identified. There is 1 False Negative (FN), where the model incorrectly classified a pneumonia image as Normal. These minor errors contribute a slight reduction in accuracy, but the high TP and TN values demonstrate strong overall model performance, as shown in Figure 2.

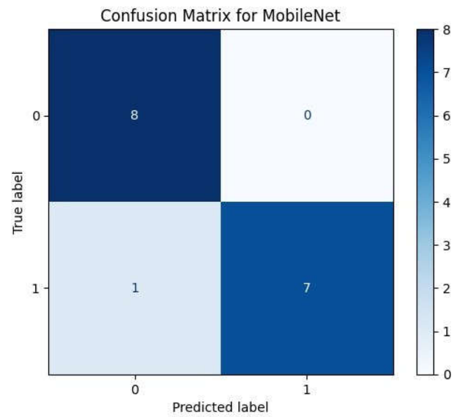

Figure 2: Confusion Matrix for MobileNet

ResNet-50:

The ResNet-50 model also achieves an accuracy of 94%, matching the performance of MobileNet. In its confusion matrix, the upper-left cell (TN = 8) shows the number of chest X-ray images correctly classified as Normal. The lower-right cell (TP = 7) represents pneumonia images correctly identified. The model has two misclassifications: one False Positive (FP), where a normal image was classified as pneumonia, and one False Negative (FN), where a pneumonia image was classified as normal. Despite these errors, the model demonstrates high overall accuracy, as shown in Figure 3.

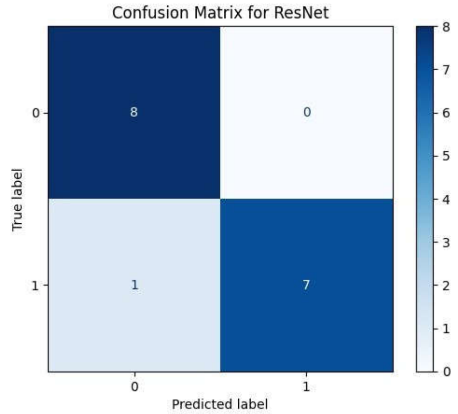

Figure 3: Confusion Matrix for ResNet-50

CNN Model:

The CNN model with downsampling achieves a substantially lower accuracy of 56%, indicating that this model struggles to correctly classify chest X-ray images for pneumonia detection. In its confusion matrix, the TN (upper-left cell) = 8, showing that the model correctly classified 8 normal images as Normal. The TP (lower-right cell) = 1, meaning only 1 pneumonia image was correctly identified. The False Negative (FN) count is 7, indicating that the model failed to detect the majority of pneumonia cases. There are no False Positives (FP), meaning the model never incorrectly labeled a normal image as pneumonia. This high FN rate indicates poor sensitivity and has critical implications in medical diagnosis, where missed detections can delay treatment, as shown in Figure 4.

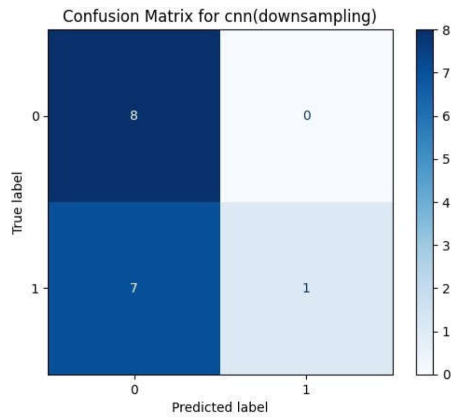

Figure 4: Confusion Matrix for CNN Model

Model Comparison Table:

Model	Accuracy	Precision	Recall	F1-Score
ResNet	0.94	1.00	0.88	0.93
MobileNet	0.94	1.00	0.88	0.93
CNN	0.56	0.88	0.14	0.24

Model Comparison:

The bar chart comparison in Figure 5 illustrates the performance differences among the three deep learning models used for pneumonia detection: ResNet, MobileNet, and CNN. Both ResNet and MobileNet achieve an accuracy of 94%, demonstrating their effectiveness in distinguishing between pneumonia and normal chest X-ray images. ResNet benefits from its deep architecture and residual connections, which enable learning of complex features without performance degradation. MobileNet employs depthwise separable convolutions that offer an optimal trade-off between accuracy and computational cost.

In contrast, the CNN model performs significantly worse with an accuracy of only 56%. This sharp drop indicates that the standard CNN with downsampling is unable to generate reliable differentiating features for pneumonia detection, as evidenced by its high false negative count in the confusion matrix. This limitation is particularly critical in medical diagnosis, where failure to detect pneumonia can result in delayed treatment.

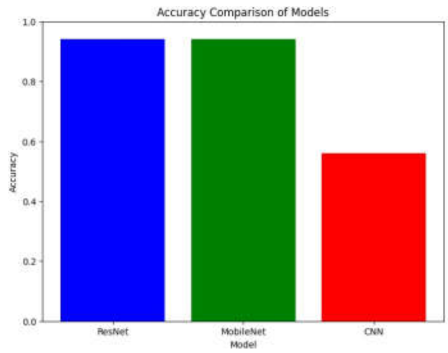

Figure 5: Accuracy Comparison for All Models

VII. CONCLUSION

This paper has explored the application of deep learning models for pneumonia detection in chest X-ray images. Leveraging sophisticated convolutional neural networks, we have demonstrated how automated systems can assist medical professionals in diagnosing pneumonia with greater precision and efficiency. The models evaluated in this study — ResNet, MobileNet, and CNN — were assessed on their ability to classify chest X-ray images as either PNEUMONIA or NORMAL. The comparative results reveal significant performance differences, confirming that model architecture choice has a strong impact on classification effectiveness.

Both ResNet and MobileNet achieved an impressive accuracy of 94%, demonstrating the value of selecting appropriate deep learning architectures for chest X-ray analysis. ResNet's deep architecture and residual connections are highly effective for extracting complex image features, making it well-suited for the intricate problem of pneumonia detection. MobileNet, on the other hand, is a less resource-intensive and lightweight model that achieves the same level of accuracy while consuming fewer computational resources. This makes MobileNet particularly appropriate for real-time systems where speed and resource usage are of primary concern. Both models demonstrated strong generalization by classifying pneumonia and normal cases with minimal misclassifications, and their high recall rates confirm their usefulness in clinical settings.

VIII. FUTURE SCOPE

The future of this research on pneumonia detection using deep learning models is large and promising, with several directions for enhancement in model capability, functionality, and practical application. One primary area requiring improvement is the class imbalance problem. Although data augmentation strategies and class weighting were considered, further research could explore more advanced techniques such as synthetic data generation using Generative Adversarial Networks (GANs) or more sophisticated sampling methods (e.g., SMOTE). These approaches would help equalize the dataset, reduce model bias, and improve the detection of pneumonia cases that resemble normal images, particularly when the class ratio is highly skewed.

Another promising direction is the investigation of advanced deep learning architectures. While ResNet and MobileNet have produced satisfactory results, state-of-the-art architectures such as Vision Transformers (ViTs), DenseNets, and hybrid CNN-ViT models could offer even greater classification accuracy. These architectures are particularly well-suited for complex medical image analysis where fine image details are crucial for accurate diagnosis. Ensemble methods combining predictions from multiple models could further improve performance by leveraging the strengths of individual architectures to reduce misclassification.

Additionally, future research could extend the dataset to cover medical conditions beyond pneumonia, such as tuberculosis or lung cancer, enabling the construction of multi-disease detection systems. The integration of multimodal data — combining clinical information, patient history, and lab results with X-ray images — could create a more comprehensive diagnostic framework. Another important direction is the deployment of these models in real-time, resource-constrained environments. Optimized models like MobileNet, designed for edge computing and mobile devices, can provide fast and reliable diagnoses without requiring large computational infrastructure, potentially expanding diagnostic access to underserved healthcare settings.

IX. REFERENCES

[1] Q. An, W. Chen, and W. Shao, “A Deep Convolutional Neural Network for Pneumonia Detection in X-ray Images with Attention Ensemble,” Diagnostics (Basel), vol. 14, no. 4, p. 390, Feb. 2024, doi: 10.3390/diagnostics14040390.

[2] N. Sebastian and B. Ankayarkanni, “Enhanced ResNet-50 with Multi-Feature Fusion for Robust Detection of Pneumonia in Chest X-Ray Images,” Diagnostics (Basel), vol. 15, Aug. 2025, doi: 10.3390/diagnostics15162041.

[3] Md. R. Hasan, S. M. Azmat Ullah, and S. Md. Rabiul Islam, “Recent advancement of deep learning techniques for pneumonia prediction from chest X-ray image,” Medical Reports, vol. 7, p. 100106, Oct. 2024, doi: 10.1016/j.hmedic.2024.100106.

[4] F. Alshanketi, A. Alharbi, M. Kuruvilla, and V. Mahzoon, “Pneumonia Detection from Chest X-Ray Images Using Deep Learning and Transfer Learning for Imbalanced Datasets,” Journal of Imaging Informatics in Medicine, vol. 38, no. 4, pp. 2021–2040, Nov. 2024, doi: 10.1007/s10278-024-01334-0.

[5] R. Siddiqi and S. Javaid, “Deep Learning for Pneumonia Detection in Chest X-ray Images: A Comprehensive Survey,” Journal of Imaging, vol. 10, no. 8, p. 176, Jul. 2024, doi: 10.3390/jimaging10080176.

[6] C.-T. Yen and C.-Y. Tsao, “Lightweight convolutional neural network for chest X-ray images classification,” Scientific Reports, vol. 14, p. 29759, Nov. 2024, doi: 10.1038/s41598-024-80826-z.

[7] S. Siddiqi and S. Javaid, “Interpretable Deep Learning for Pneumonia Detection Using Chest X-Ray Images,” Information, vol. 16, no. 1, p. 53, Jan. 2025, doi: 10.3390/info16010053.

[8] J. Lee and H. Kim, “Deep Transfer Learning Using Real-World Image Features for Medical Image Classification, with a Case Study on Pneumonia X-ray Images,” Bioengineering, vol. 11, no. 4, p. 406, Apr. 2024, doi: 10.3390/bioengineering11040406.

[9] N. Izdihar, S. B. Rahayu, and K. Venkatesan, “Comparison Analysis of CXR Images in Detecting Pneumonia Using VGG16 and ResNet50 Convolution Neural Network Model,” JOIV: International Journal on Informatics Visualization, vol. 8, no. 1, pp. 326–332, Mar. 2024, doi: 10.62527/joiv.8.1.2258.

[10] M. Abutaleb, H. Al-Aqrabi, and R. Hill, “Hybrid Deep Learning and Machine Learning Framework for Automated Pneumonia Detection in Chest X-ray Images,” MethodsX, Nov. 2025, doi: 10.1016/j.mex.2025.103205. [In Press / Accepted Manuscript]

[11] M. Ben Nasr, R. Bchir, and M. Tlili, “PELM: A Deep Learning Model for Early Detection of Pneumonia in Chest Radiography,” Applied Sciences, vol. 15, no. 12, p. 6487, Jun. 2025, doi: 10.3390/app15126487.

[12] A. Kaya et al., “Trustworthy Pneumonia Detection in Chest X-ray Imaging through Attention-Guided Deep Learning,” Scientific Reports, vol. 15, Nov. 2025, doi: 10.1038/s41598-025-23664-x.

[13] R. Sriramkumar, “Enhancing Diagnostic Accuracy: Leveraging Deep Transfer Learning for Disease Detection in Chest X-Ray Images,” International Journal of Intelligent Systems and Applications in Engineering, vol. 12, no. 22s, pp. 1582–1591, 2024.

[14] J. Rabbah, M. Ridouani, and L. Hassouni, “Improving Pneumonia Diagnosis with High-Accuracy CNN-Based Chest X-Ray Image Classification and Integrated Gradient,” Biomedical Signal Processing and Control, vol. 101, p. 107239, 2025, doi: 10.1016/j.bspc.2025.107239.

[15] S. Priyatharsini, S. Sivaneasan, P. R. Kshirsagar, and P. Chakrabarti, “Revolutionizing Lung Disease Diagnosis: A Unique Hybrid Deep Learning Framework for Explainable Chest X-Ray Analysis,” SSRN, preprint, 2025, doi: 10.2139/ssrn.5110794.