



Healthcare Diagnosis System Using Deep Learning-Based Image Analysis for Automated Pneumonia Detection from Chest X-Ray Images

Mayank Kumar¹, Ansh², Malika³, Tanisha⁴

Department of CSE (AI&ML)^{1,2,3,4}

Panipat Institute of Engineering and Technology, Samalkha, Panipat, India^{1,2,3,4}

Abstract. *Pneumonia still constitutes one of the major causes of mortality across the globe, specifically targeting children, elderly people, and immunocompromised patients. Diagnosis and timely intervention can help reduce disease burden and patient morbidity. Traditional methods of diagnosing pneumonia using chest X-ray images depend largely on the involvement of radiologists, making the entire process slow and prone to observer variability, especially in limited health care facilities. This paper introduces an automated pneumonia diagnosis system based on deep learning using ResNet34 architecture. The model uses image pre-processing techniques, transfer learning, and data augmentation to boost feature extraction capabilities while reducing overfitting. The trained system classifies chest X-ray images into pneumonia and non-pneumonia classes using a user-friendly web application built on Streamlit platform. The designed system provides a good computer-aided diagnosis system which may help the medical professionals to conduct initial screening for pneumonia, reduce diagnostic efforts, and increase the accessibility of health care facilities in rural areas.*

Keywords:- Pneumonia Detection, Deep Learning, Convolutional Neural Network, ResNet34, Transfer Learning, Chest X-ray, Medical Image Analysis, Artificial Intelligence, Computer-Aided Diagnosis.

Introduction

Pneumonia is an acute disease affecting the respiratory tract with the lungs being the main organs affected, making it one of the greatest health threats in the world today. In global health statistics, pneumonia causes many deaths and admissions annually, especially in children less than five years old, senior citizens, and individuals with impaired immunity systems. It is necessary to diagnose pneumonia early to reduce any chances of deaths. But accurate diagnosis requires experience from radiologists, who interpret chest x-ray images to detect abnormal changes in the lungs.

While chest X-ray is one of the commonly used diagnostic modalities for the detection of pneumonia infection, there exist certain difficulties when conducting a manual interpretation of the chest radiograph. There may be variations in accuracy between radiologists due to disparities in expertise, tiredness, work load and subjectivity. Moreover, lack of skilled radiologists in rural healthcare facilities and developing countries hampers the diagnostic process leading to delays in treatment. Such issues underscore the importance of developing intelligent computer aided diagnostic systems that can deliver efficient



pneumonia detection [1]. Recent developments in the field of artificial intelligence especially deep learning have greatly contributed to medical image processing. The convolutional neural networks have shown an outstanding performance in automatic extraction of hierarchical features of images without needing any feature engineering manually. Unlike other machine learning approaches, the convolutional neural networks learn high-level representations through their interaction with the input images. In many existing CNN architectures, the ResNet34 model has been found to be efficient due to its residual learning algorithm that solves the problem of vanishing gradient for deep neural network models. ResNet34 model allows deep feature learning and at the same time provides stable training; therefore, it performs better in the classification of images in medical datasets [2]. Transfer learning can make the model even more effective since it uses information learned in large image datasets, and thus the training period becomes shorter.

This study suggests the development of an automated pneumonia detection system based on the ResNet34 architecture that performs the classification of chest X-ray images to normal and pneumonia images. This system involves the implementation of some steps such as image preprocessing, transfer learning, and data augmentation to make the model more robust to changes in input images and overfitting. Moreover, a web application based on Streamlit will be created to provide real-time predictions [3].

The primary contributions of this research are summarized as follows:

1. Development of an automated pneumonia detection framework using the ResNet34 deep convolutional neural network.
2. Integration of transfer learning and data augmentation to improve classification accuracy and model generalization.
3. Deployment of the trained model as a lightweight Streamlit web application for real-time clinical assistance.
4. Comprehensive evaluation using standard performance metrics including accuracy, precision, recall, F1-score, confusion matrix, and loss curves.
5. Provision of a scalable and cost-effective computer-aided diagnostic system suitable for healthcare facilities with limited radiological expertise.

The proposed system aims to support healthcare professionals by providing fast, accurate, and reliable pneumonia screening while reducing diagnostic workload and improving healthcare accessibility in underserved regions.

Related Work

Artificial intelligence has significantly improved medical image analysis, particularly for pneumonia detection using chest X-ray images. Earlier computer-aided diagnosis systems relied on conventional image processing techniques involving preprocessing, handcrafted feature extraction, and machine learning classifiers such as Support Vector Machine (SVM), K-Nearest Neighbor (KNN), and Naïve Bayes. Although these methods performed reasonably well on small datasets, they required manual feature engineering and showed limited robustness to variations in image quality [4]. Recent advances in deep learning, especially Convolutional Neural Networks



(CNNs), have overcome these limitations by automatically learning hierarchical image features. Popular architectures including AlexNet, VGG, ResNet, DenseNet, and EfficientNet have demonstrated excellent performance in pneumonia classification. Among these, ResNet has gained significant attention because its residual learning mechanism enables deeper networks to be trained efficiently while avoiding the vanishing gradient problem. Transfer learning has become a widely adopted technique in medical imaging due to the limited availability of annotated chest X-ray datasets [5]. Pre-trained CNN models fine-tuned on medical images improve classification accuracy while reducing training time and minimizing overfitting. In addition, data augmentation techniques such as image rotation, flipping, cropping, and brightness adjustment further enhance model generalization. Several researchers have also explored attention mechanisms and ensemble learning to improve diagnostic accuracy. While these approaches achieve promising results, they often increase computational complexity and are difficult to deploy in real-time clinical environments. Lightweight architectures such as MobileNet and EfficientNet-Lite have therefore been proposed for deployment on portable and low-resource devices [6]. Despite these advancements, several challenges remain. Most existing studies focus primarily on improving classification accuracy using benchmark datasets, with limited emphasis on real-world deployment, scalability, user-friendly interfaces, and comprehensive evaluation [7]. Furthermore, many models require high computational resources, limiting their applicability in resource-constrained healthcare settings.

Table 1: Summarizes representative studies on pneumonia detection using deep learning techniques.

Title	Author	Technologies Used	Findings	Limitations
“Review on Pneumonia Image Detection: Machine Learning Approach”	Kareem et al. (2022)	CNN, KNN, ANN, ResNet, CheXNet	Proposed Approach: Surveyed multiple ML/DL techniques for Pneumonia detection using chest X-rays. Highlighted CNN and ResNet as highly effective for automated diagnosis.	Lacks implementation of a single optimized model; mainly Theoretical comparison of techniques.
“Pneumonia Detection Using Deep Learning: A CNN-Based Approach”	Lakshmi et al. (2025)	CNN, Image Preprocessing, Deep Learning	Proposed Approach: Developed CNN based model for automatic pneumonia detection from chest X-rays, improving speed and reducing manual diagnosis errors.	Limited dataset and lacks generalization across diverse real world hospital data.
“Pneumonia Detection using CNN”	Pramoth et al. (2022)	CNN, Image Classification	Proposed Approach: Used CNN architecture to classify pneumonia cases from X-ray images with good accuracy.	Limited evaluation metrics and lacks comparison with advanced architectures like ResNet.



“Pneumonia Detection using Transfer Learning in CNN”	Rahman et al. (2020)	Transfer Learning, AlexNet, ResNet18, DenseNet201	Proposed Approach: Used pretrained CNN models with transfer learning achieving up to 98% accuracy in pneumonia classification.	Requires large Computational resources. performance depends on dataset quality.
--	----------------------	---	--	---

A. Research Gap

Although considerable progress has been achieved in deep learning-based pneumonia detection, several research gaps continue to hinder practical clinical adoption.

1. Most existing studies focus exclusively on maximizing classification accuracy while neglecting real-time deployment.
2. Many proposed models require substantial computational resources, making deployment in rural healthcare environments difficult.
3. Existing systems frequently rely on carefully curated benchmark datasets and therefore demonstrate reduced generalization when applied to real-world hospital data.
4. Few studies provide complete end-to-end diagnostic systems that combine preprocessing, deep learning inference, and an interactive web-based interface for clinical use.
5. Several published works evaluate performance using only accuracy, without comprehensive analysis using precision, recall, F1-score, ROC curves, and confusion matrices.
6. Interpretability and usability remain important challenges, limiting clinician trust and widespread adoption.

B. Motivation of the Proposed Work

To address these limitations, this research develops an end-to-end automated pneumonia diagnosis framework based on the ResNet34 architecture [8]. The proposed system integrates image preprocessing, transfer learning, data augmentation, and a lightweight Streamlit-based deployment platform to provide accurate and real-time pneumonia detection. In addition to improving diagnostic performance, the proposed framework emphasizes usability, scalability, and accessibility for healthcare professionals working in both urban hospitals and resource-constrained medical facilities [9].

Proposed Methodology

Here proposed methodology consists of five major stages: data collection, preprocessing, model development, evaluation, and deployment. Chest X-ray images belonging to the Normal and Pneumonia classes are collected from a publicly available dataset and divided into training, validation, and testing sets. During preprocessing, all images are resized, normalized, and augmented using techniques such as rotation, flipping, and zooming to improve model generalization and reduce overfitting [10].

For classification, the ResNet34 convolutional neural network is employed through transfer learning. Its residual connections enable effective feature extraction while mitigating the vanishing gradient problem.



The final fully connected layer is modified to perform binary classification between normal and pneumonia cases.

The model is trained using the Cross-Entropy Loss function and optimized with the Adam optimizer over multiple epochs. Performance is continuously monitored using training and validation accuracy and loss. The trained model is evaluated on unseen test data using accuracy, precision, recall, F1-score, and confusion matrix to assess its classification performance and generalization capability [11].

Finally, the trained model is deployed as a lightweight diagnostic application using Streamlit, allowing users to upload chest X-ray images and obtain real-time predictions [12]. The implementation utilizes PyTorch for deep learning, NumPy for numerical computation, Matplotlib for visualization, and Google Colab for GPU-accelerated model training.




Figure 1: Methodology Flowchart.



Some dataset samples are shown below in Figure 2:

NORMAL



PNEUMONIA




Figure 2: Dataset Sample images.

Results and Discussion

The proposed ResNet34-based transfer learning model was evaluated using a multi-class chest X-ray dataset comprising COVID-19, Lung Opacity, Viral Pneumonia, and Normal images [13]. Prior to training, the correctness of the data loading and preprocessing pipeline was verified through batch visualization using Matplotlib, ensuring that image transformations, grayscale conversion, normalization, and augmentation were applied consistently across the dataset.

A. Model Training Performance

The model was trained using the FastAI framework with the Adam optimizer and Cross-Entropy Loss function. An optimal learning rate was determined using the learning rate finder (`lr_find()`), which identified the learning rate corresponding to the steepest decline in training loss [14]. Training was conducted for 14 epochs, balancing convergence with the risk of overfitting on the relatively small medical imaging dataset.



The training process demonstrated stable convergence throughout the learning phase. As shown in Fig. 6.4, the training loss decreased significantly from 0.61 to 0.14, indicating effective optimization of the network parameters [15]. Although the validation loss fluctuated during the initial epochs, it stabilized below 0.20 after the fourth epoch, suggesting good generalization without noticeable overfitting. Simultaneously, the validation accuracy improved steadily, reaching 84.2% at the final epoch.

B. Performance Evaluation

The trained model was evaluated using standard classification metrics, including accuracy, precision, recall, F1-score, and confusion matrix. The learning curves presented in Fig. 6.6 demonstrate a consistent reduction in training and validation loss, accompanied by a corresponding increase in validation accuracy. The convergence of both curves indicates that the proposed model effectively learned discriminative features while maintaining stable performance on unseen data [16].

The confusion matrix further illustrates the classification capability of the proposed framework. The model achieved class-wise prediction accuracies of 85% for COVID-19, 84% for Viral Pneumonia, 81% for Lung Opacity, and 83% for Normal chest X-ray images. Notably, the Normal class exhibited a very low misclassification rate (less than 1%), demonstrating the model's ability to distinguish healthy lungs from pathological conditions with high reliability.

C. Discussion

The experimental results confirm the effectiveness of transfer learning using the ResNet34 architecture for automated chest X-ray disease classification. The incorporation of image preprocessing, grayscale conversion, normalization, and data augmentation improved feature extraction and enhanced the model's robustness against variations in image quality. Furthermore, adapting the pretrained ResNet34 architecture to accept grayscale images enabled efficient knowledge transfer while reducing training time.

The stable convergence of the training process and the high classification accuracy indicate that the proposed framework successfully learned meaningful radiographic features for distinguishing among the four disease categories. The low false-positive rate for the Normal class is particularly significant in clinical practice, as it minimizes the likelihood of incorrectly classifying infected patients as healthy.

Overall, the obtained results demonstrate that the proposed deep learning framework is capable of providing accurate and reliable multi-class chest X-ray classification. Combined with its lightweight Streamlit-based deployment, the proposed system offers a practical computer-aided diagnostic solution that can assist healthcare professionals in rapid and real-time disease screening.

Conclusion

The proposed deep learning framework for automated disease diagnosis using chest X-ray images has been successfully developed and presented in this paper. The paper demonstrates the complete development of an artificial intelligence-based diagnostic system, encompassing data collection, preprocessing, normalization, and data augmentation to enhance dataset quality and improve model generalization across multiple disease categories. Transfer learning was employed using the pre-trained ResNet34 architecture implemented in Python, enabling efficient feature extraction and reducing training time while improving classification



performance. The proposed model was trained to classify chest X-ray images into four categories: COVID-19, Lung Opacity, Viral Pneumonia, and Normal. The model was evaluated using standard performance metrics, including accuracy, precision, recall, and F1-score, demonstrating its effectiveness in automated disease classification. Following training and validation, the model was serialized and deployed through a Streamlit-based web application, enabling users to upload chest X-ray images and receive diagnostic predictions along with confidence scores in real time.

Furthermore, the development workflow, including source code management and version control, was systematically maintained using Git and GitHub, ensuring reproducibility and collaborative development. The experimental results demonstrate that deep learning-based approaches, particularly transfer learning with ResNet34, provide an efficient and reliable solution for automated chest X-ray analysis. The proposed framework has the potential to support healthcare professionals by facilitating rapid, accurate, and accessible disease diagnosis, especially in resource-constrained clinical environments.

References

- [1] J. Deng et al., "ImageNet: A large-scale hierarchical image database," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2009.
- [2] Fast.ai Documentation, "Fast.ai Deep Learning Library." Available: <https://docs.fast.ai/>
- [3] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016.
- [4] A. Paszke et al., "PyTorch: An Imperative Style, High-Performance Deep Learning Library," in Advances in Neural Information Processing Systems (NeurIPS), 2019.
- [5] P. Rajpurkar et al., "CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning," arXiv preprint arXiv:1611.04224, 2017.
- [6] Kareem et al., "Review on Pneumonia Image Detection: A Machine Learning Approach," International Journal / Conference, 2022.
- [7] Lakshmi et al., "Pneumonia Detection Using Deep Learning: A CNN-Based Approach," International Journal / Conference, 2025.
- [8] Rajasenbagm et al., "Detection of Pneumonia Infection using Deep CNN," International Journal / Conference, 2021.
- [9] Aljawarneh et al., "Enhanced CNN Model for Pneumonia Detection," International Journal / Conference, 2023.
- [10] Akbar et al., "Comparative Study of CNN Architectures for Pneumonia Detection," International Journal / Conference, 2024.
- [11] Mabrouk et al., "Ensemble Deep Learning for Pneumonia Detection," International Journal / Conference, 2023.
- [12] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," in International Conference on Learning Representations (ICLR), 2015.
- [13] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely Connected Convolutional Networks," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017.



-
- [14] M. Tan and Q. Le, “EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks,” in International Conference on Machine Learning (ICML), 2019.
- [15] C. Szegedy et al., “Going Deeper with Convolutions,” in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015.
- [16] O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional Networks for Biomedical Image Segmentation,” in International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI), 2015.