

Optimizing Pneumonia Identification in Chest X-Rays Using Deep Learning Pre-Trained Architecture for Image Reconstruction in Medical Imaging

Rahul P. Mahajan

Research and Development Department
Healthcare and Medical Device Development Industry
College of Engineering, Pune, India
rpm.mahajan@gmail.com
0009-0003-3038-4391

Abstract: *The rapid accumulation of fluid in the lungs is the hallmark of the fatal illness known as pneumonia. Therefore, it is crucial to get a diagnosis and medication as soon as possible in order to stop the condition from getting worse. In order to diagnose pneumonia, chest X-rays (CXR) are typically used. This study assesses the efficacy of ResNet50 and other pre-trained DL models in classifying chest X-ray images as evidence of pneumonia. The proposed ResNet50 model achieves an accuracy of 93.06%, precision of 88.97%, recall of 96.78%, and an F1-score of 92.71%, surpassing MobileNet, EfficientNetB0, and Xception across all performance metrics. The model has been further tested and shown to be reliable and effective in differentiating between normal and pneumonia patients utilizing ROC curve analysis, accuracy-loss trends, and a confusion matrix. The study highlights the superiority of ResNet50 in automated pneumonia detection, offering a promising tool for early diagnosis and clinical decision support. The results highlight how deep learning-based methods have the ability to improve radiological evaluations, which in turn can decrease diagnostic mistakes and increase patient outcomes. This research contributes to developing AI-driven medical imaging solutions, facilitating more accurate and scalable pneumonia detection in real-world healthcare settings.*

Keywords: Healthcare, Pneumonia detection, medical imaging, Image Reconstruction, Chest X-ray, Deep Learning, Chest X-ray images dataset.

I. INTRODUCTION

Healthcare plays a vital role in improving human well-being by diagnosing, preventing, and treating diseases. Advancements in medical technology have significantly enhanced disease detection and management, leading to better patient outcomes [1][2]. However, many diseases, particularly respiratory infections, continue to pose serious health risks worldwide [3]. Among them, pneumonia continues to rank high in terms of mortality and disability, particularly among the elderly, the young, and those with underlying medical issues [4]. Symptoms of pneumonia include a high temperature, coughing up mucus, chest discomfort, and trouble breathing; the infection is severe and attacks the alveoli in the lungs [5]. If not detected and treated early, pneumonia can lead to severe complications and even death. According to global health reports, pneumonia accounts for millions of hospitalizations and thousands of deaths annually, highlighting the urgent need for early and accurate detection methods [6].

When diagnosing pneumonia, chest X-rays (CXRs) are preferred over other imaging modalities, including CT scans and MRI, because of their reduced radiation exposure, affordability, and ease of accessibility [7][8]. Despite their advantages, interpreting CXRs requires expert radiologists, and misdiagnosis can occur due to image quality issues,



variations in disease presentation, and human error [9]. To address these challenges, advancements in AI and DL have paved the way for automated pneumonia detection with high accuracy and efficiency.

ML and DL have revolutionized medical imaging by enabling automated feature extraction and disease classification [10]. Additionally, image reconstruction techniques are crucial in improving the quality of medical images by reducing noise, enhancing contrast, and mitigating artifacts [11]. Diagnostic accuracy and clinical decision-making are both enhanced by high-quality image reconstruction, which allows for clearer visualization of diseased characteristics. The proposed approach leverages advanced DL models to refine CXR images, ensuring improved visualization and more accurate pneumonia detection, ultimately contributing to better patient outcomes in healthcare.

A. Motivation and Contribution of the Study

The motivation behind this study stems from Pneumonia is a significant global health concern, necessitating accurate and early detection through chest X-ray images. Traditional diagnostic methods rely on radiologists, leading to variability and delays. Deep learning, particularly pre-trained models for image reconstruction, offers a promising solution by enhancing feature extraction and improving classification accuracy. This study leverages pre-trained deep networks to optimize pneumonia detection, reducing the need for large annotated datasets while improving robustness against noise and image artifacts. Here is a rundown of what this study has accomplished:

- Collect Chest X-ray images dataset for pneumonia detection.
- Integrates essential preprocessing techniques such as data cleaning, CLAHE and grayscale conversion Data augmentation, one hot encoding and improved feature extraction across the dataset
- Its pre-trained DL models for pneumonia detection in chest X-rays, enhancing feature extraction and classification accuracy for improved medical diagnosis
- Use ResNet50 and other pre-trained models to identify pneumonia and evaluate their performance against other models.
- Analysis based on accuracy, precision, recall, and F1-score for better classification insights.
- Scalable pre-trained models are evaluated for effective pneumonia identification in chest X-rays.

B. Novelty of paper

This study presents a novel approach to pneumonia detection by leveraging the ResNet50 deep learning model with an optimized preprocessing pipeline, including histogram equalization, grayscale conversion, normalization, and data augmentation, to enhance image quality and improve classification performance. Unlike existing studies that rely on conventional models with limited generalization, this work demonstrates superior accuracy, ensuring minimal false negatives—a critical aspect in medical diagnostics. Furthermore, a comprehensive comparative analysis against MobileNet, EfficientNetB0, and Xception establishes the robustness of ResNet50. The study also explores ROC curve evaluation and confusion matrix analysis, offering deeper insights into model reliability for clinical applications.

C. Structure of the paper

The sections of the paper are as follows: Section II gives the necessary context for understanding pneumonia detection, and Section III lays out the research strategy that will be used to conduct this study. Section IV details the findings from the experiment as well as an evaluation of the model's efficiency. Here, it provides the conclusion and future research for Section V.

II. LITERATURE REVIEW

The study examines the body of research on the use of chest X-ray imaging for pneumonia detection in this area. Most of the reviewed works focused on classification techniques. Some reviews are:

Usman et al. (2025) created a CNN model and pretrained it with ResNet-50. The original dataset from the RSNA pneumonia detection challenge was utilized to train a ML model. The dataset contained 26,684 chest array pictures from distinct individuals, with 56% being male and 44% female. They then compared the models' ability to detect



pneumonia, looked at how augmentation and dataset size affected their performance, and concluded that the models were generally good. The CNN model trained from scratch, without any augmentation, had the best overall accuracy, which was 0.79 [12].

Roy et al. (2024) used DenseNet-121 as the basis for the normal and pneumonia binary classes and ResNet50 as the basis for the bacterial, viral, and normal multi-class classification procedures. Highlighted the precise spatial areas of relevant channels and removed unnecessary ones from the backbone's feature extraction using a channel-specific spatial attention technique called FCSSAM. Using binary and multi-class classification settings on a chestX-ray dataset that is accessible to the public. Their suggested approach yields 79.79% accuracy rates for the binary [13].

Jain and Kumar (2024) Pneumonia is a serious infection of the respiratory system; it is also a significant health risk, more so in vulnerable populations. According to this, performance metrics indicate performance in robust capabilities of classification, with overall accuracy at 93%. CNN can be a helpful tool within a clinical set-up for timely diagnosis of pneumonia, hence doing timely medical interventions for better patient outcomes [14].

Shrimali (2024) investigated an application of DL to enhance chest X-ray-based pneumonia diagnosis, offering a potentially transformative approach. Five pre-trained CNNs – MobileNetV2, VGG16, DenseNet121, EfficientNetB0, and ResNet50 – were evaluated for their effectiveness. VGG16 had the highest accuracy, reaching 90.87%, whereas DenseNet121 achieved a well-balanced F1 score of 74.52% [15].

Wang et al. (2023) chest X-ray picture identification using a DL style. Using a classic convolutional network architecture, it retrieves features that are rich in semantic information. The investigation called for a library of 6,189 children's X-ray films, which included 3,319 typical instances and 2,870 cases of pneumonia. After selecting 20% as the test data set and comparing 11 typical models using 4 assessment measures, the accuracy percentage achieved 89.1% [16]

Ramli et al. (2023) seek to create a model for automated pneumonia identification using X-ray images so that patients may get the help they need faster, with less strain on healthcare providers. Noise reduction in X-ray images is achieved using Discrete Wavelet Transform in conjunction with a bilateral filter. Classification is performed using CNN, with testing using 5,840 X-ray images of both healthy and sick chests with pneumonia. Classifying patients as either pneumonia-positive or pneumonia-negative, the trained CNN model demonstrates remarkable accuracy, recall, and an F1-score of 96.0% [17].

Hasan, Ullah and Hasan (2023) devised an automated system for pneumonia detection from CXR images. To address the data scarcity issue, they employed deep transfer learning techniques on VGG16 and ResNet50 architectures. Their proposed approach achieved impressive accuracy rates of 89.23% and 88.80% for VGG16 and ResNet50, respectively. This performance surpasses that of existing methods, underscoring the effectiveness of their approach [18].

Table I presents a comparison of background studies according to their results, limitations, and potential future research

Table 1: Comparative Analysis for Deep Learning based Pneumonia Detection

Author	Source	Methodology	Findings	limitation	Future work
Usman et al. (2025)	RSNA pneumonia detection challenge dataset	CNN from scratch & ResNet-50, impact of augmentation	CNN (no augmentation) achieved 79% accuracy	Limited dataset size impact analysis	Further exploration of augmentation and dataset size effects
Roy et al. (2024)	Public chest X-ray dataset	DenseNet-121 & ResNet-50 with Fuzzy Channel Selective Spatial Attention Module	79.79% accuracy for binary classification	Lacks detailed comparison with other attention-based models	Optimization of FCSSAM for better feature extraction
Jain & Kumar (2024)	Dataset with 624 images	CNN model trained for pneumonia detection	93% accuracy	Small dataset size	Testing on a larger dataset for improved



					generalization
Shrimali (2024)	Chest X-ray images	Multiple pre-trained CNNs	VGG16 achieved 90.87% accuracy, DenseNet121 had an F1-score of 74.52%	No external validation dataset	Further, implement different datasets and ensemble models.
Wang et al. (2023)	Children's X-ray dataset (6,189 images)	CNN with adjacency matrix and graph inference for global modeling	89.1% accuracy across 11 models	Requires more complex graph-based reasoning	Integration of more advanced graph-based techniques
Ramli et al. (2023)	5,840 X-ray images	CNN with Discrete Wavelet Transform bilateral filtering for noise reduction	96% accuracy with excellent recall and F1-score	Limited generalization beyond dataset	Testing with different preprocessing techniques
Hasan, Ullah & Hasan (2023)	Transfer learning on VGG16 & ResNet50.	Deep transfer learning for data scarcity issue	89.23% (VGG16) and 88.80% (ResNet50) accuracy	Performance not significantly different from baseline methods	Exploring other transfer learning architectures

III. METHODOLOGY

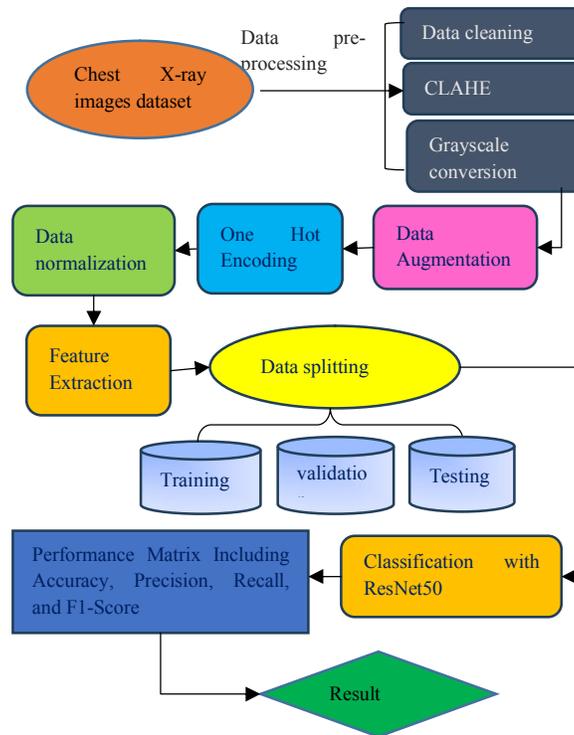


Figure 1: Flowchart of Chest X-Ray Images for Pneumonia Detection

This study aims to evaluate pre-trained models for pneumonia detection. The following steps of research design are shown in Figure 1 flowchart. The methodology starts with the collected chest X-ray images dataset. Data preparation to



improve image quality is the first step in the suggested pipeline for pneumonia identification in chest X-rays. This includes data cleaning to remove irrelevant images, histogram equalization for contrast enhancement, and grayscale conversion for intensity adjustments. The next step in improving the model's performance and preventing overfitting is to apply data augmentation methods, one-hot encoding, feature extraction, and normalization. The dataset is then split into training (80%), validation (10%), and testing (10%) sets to ensure a balanced evaluation. Deep learning models like ResNet50 are employed for categorization. It is possible to measure the model's efficacy using a performance matrix that includes the F1-score, recall, accuracy, and precision.

The following is a flow chart showing the main procedures for detecting pneumonia:

A. Data Collection

In this research work used chest X-ray image dataset form the Kaggle. The information on the 5,232 chest X-ray pictures with RGB channels that are part of this collection. Among them, there are 3883 photos depicting pneumonia and 1346 images bearing normal dimensions of 227×227 . It used chest X-rays taken from children in Guangzhou between the ages of one and five from retrospective cohorts. The following histogram plots of data are shown in Figure 2.

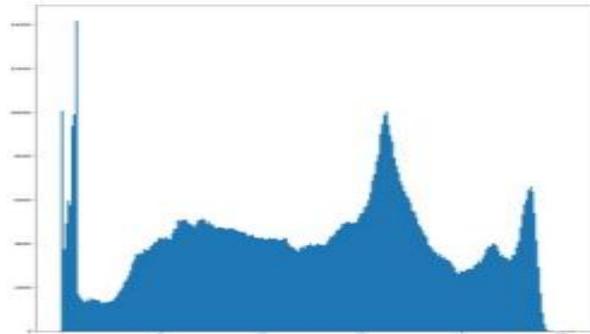


Figure 2: Histogram Plot

Figure 2 displays the Histogram plot which presents distributed histograms and density plots across 30 numerical features indicating various distribution patterns of a grayscale medical image like chest X-rays. Pixel frequencies are displayed on a y-axis, while intensity values from 0 (black) to 255 (white) are shown on the x-axis. The histogram shows multiple peaks, with significant intensity concentrations in both the darker and brighter regions, indicating a broad contrast range.

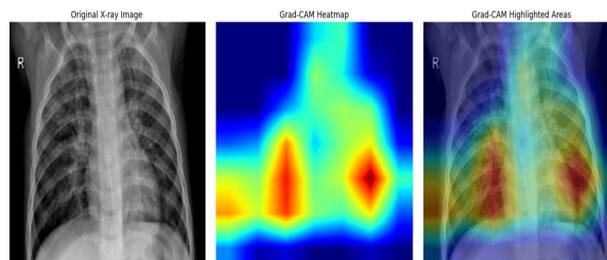


Figure 3: Heatmap visualization of chest X-ray images for pneumonia detection

Figure 3 displays the heatmap representation of chest X-ray pictures used to identify pneumonia. Grad-CAM helps models make better decisions by graphically representing the parts of a chest X-ray that matter the most. The consolidated lung areas and other pneumonia-related regions are constantly visible on heatmaps, and Grad-CAM uses these heatmaps to reveal which lung regions affect the model's pneumonia prediction. This increases the model's credibility in medical situations. Raw predictive performance measurements are slightly lower, but model predictions are more transparent.



B. Data Preprocessing

Data preprocessing for pneumonia detection in chest X-rays includes data cleaning to remove irrelevant images, histogram equalization for contrast enhancement, and grayscale conversion to adjust contrast inversion. Data augmentation (rotation, flipping) prevents overfitting, while one-hot encoding and normalization prepare the data for training. Here are the stages involved in pre-processing:

- **Data cleaning:** Data cleaning for chest X-rays involves identifying and removing irrelevant or poor-quality images from a dataset, typically including steps like filtering out images with incorrect patient information and detecting and removing images that aren't chest X-rays.
- **Contrast Limited Adaptive Histogram Equalization (CLAHE):** CLAHE is applied to enhance a contrast of the grayscale image. Especially in low-contrast areas, CLAHE successfully raises the image's feature visibility. The initial range of the pixel's intensity is increased from 0 to 255 using this procedure. The result is a somewhat greater contrast ratio and a broader intensity range in the improved image.
- **Grayscale conversion:** Grayscale conversion in chest X-rays, also known as grayscale inversion (GSI), is a post-processing technique that changes the image's contrast inversion, inverse digested imaging, or "bones black" imaging

C. Data Augmentation

Data augmentation has been widely utilized; it boosts picture count by performing a sequence of adjustments while keeping class labels. A dataset regulator and image diversity booster, data augmentation is used to train pneumonia class pictures. In addition, to establish parity in the datasets, data augmentation was employed to increase the number of normal or non-pneumonia photos in the first dataset. The number of training pictures for pneumonia classes versus non-pneumonia. There is no augmentation done on the dataset's testing and validation photos. Images of data augmentation are shown in Figure 4.

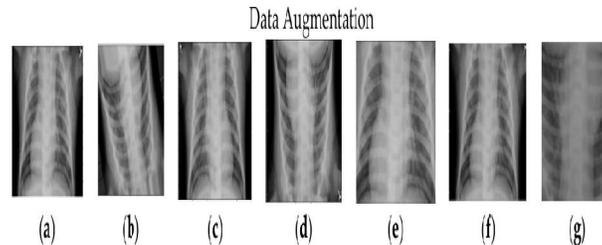


Figure 4: Data Augmentation Images

Figure 4 displays the results of using data augmentation techniques on chest X-ray images to identify pneumonia. Among the techniques are the following: scaling the image, random brightness modification, horizontal flipping, vertical rotation, zooming, and random rotation. These transformations enhance model generalization by introducing variability in the training data, decreasing overfitting, and enhancing robustness in pneumonia classification.

D. One Hot Encoding

An often-used method for transforming binary labels into categories is one-hot encoding. This technique compares each level of the categorical component to a predetermined reference level by treating each category as an independent binary variable. The procedure converts a single variable with n observations and d potential values into two binary variables with n observations each. For pneumonia detection, one-hot encoding converts categorical labels (e.g., "normal," "pneumonia") into binary vectors where only one element is set to 1, allowing the model to classify chest X-ray images based on these labels effectively.

E. Data Normalization

As a preprocessing step, normalization ensures that an input data is scaled to a standard range, such as [0,1] or [-1,1], in order to improve stability and convergence of a model. In medical imaging, it enhances contrast, reduces intensity



variations, and ensures consistent feature representation for deep learning models. Normalize pixel values (e.g., scale to [0,1] or [-1,1]) for stable training. Equation 1 represents the scaling.

$$X_{nor} = \frac{x - x_{min}}{x_{max} - x_{min}} \quad (1)$$

The original feature is represented by X , the minimum value by x_{min} and the highest value by x_{max} .

F. Feature Extraction

The feature extraction capability of these models simultaneously operates as an image classifier. Input features get captured through initial layers of DL models that extract edges and contours and such features are essential for medical picture analyses. Its purpose is to achieve comprehensive feature extraction. In addition, multi-head attention methods were used to the output of each feature extraction network in order to enhance the representation of critical locations.

G. Data Splitting

Divides the dataset into training (e.g., 80%), validation (10%), and test (10%) sets to ensure effective learning and unbiased evaluation. This helps prevent overfitting and ensures the model generalizes well for pneumonia detection in chest X-rays

H. Classification with ResNet50 Model

There are other varieties of ResNet, but the most significant are ResNet-18 and ResNet-50, which are the most sophisticated and widely used for image classification [19]. The ReLU activation method, 48 coiled layers, 1 MaxPool, and an intermediate pool layer make up the ResNet-50 model. ResNet is known for its ability to prevent overfitting, which happens when a front-fed neural network has a hidden layer with too many parameters and too many specific neurones to train with certain data[20][21]. ResNet also succeeded in resolving the vanishing gradient issue, which occurs when the derivative value of the initial layers is iteratively multiplied, leading to a fall in the derivative value and, thus, a decline in the network's accuracy for deep learning. The following Equations of ResNet is (2, 3).

$$y_1 = h(x_1) + f(x_1, w_1) \quad (2)$$

$$x_{l+1} = f(y_1) \quad (3)$$

where x_i is the input feature w_i , is a collection of biases and weights, and f represents the function that remains after two 3×3 convolutional layers have been stacked. The operation that follows element-wise addition is denoted as f , and it is associated with ReLU. In order to compute the features according to Equation (4), the function h is defined as an identity mapping.

$$x_L = x_l + \sum_{i=l}^{L-1} F(x_i, w_i) \quad (4)$$

A sequence of matrix-vector products makes up the feature x_L where L is the total of the outputs of all previous residual functions. The loss equation is derived from the chain rule of backpropagation and works as Equation (5).

$$\frac{\partial E}{\partial x_l} = \frac{\partial E}{\partial x_L} \frac{\partial x_L}{\partial x_l} = \frac{\partial E}{\partial x_L} \left(1 + \frac{\partial}{\partial x_l} \sum_{i=l}^{L-1} F(x_i, w_i)\right) \quad (5)$$

The gradient $\frac{\partial E}{\partial x_l}$ can be broken down into two additive terms, as shown in Equation 5. One term, $\frac{\partial E}{\partial x_L}$ spreads information directly without considering any weight layers. The other term, $\frac{\partial E}{\partial x_L} \left(\frac{\partial}{\partial x_l} \sum_{i=l}^{L-1} F(x_i, w_i)\right)$ propagates through the weight layers. The information is immediately transmitted back to any shallower because of the additive component of $\frac{\partial E}{\partial x_l}$. It is highly improbable that the gradient $\frac{\partial E}{\partial x_l}$ will be cancelled out for a mini batch, according to Equation 5. This is because, in general, the term $\frac{\partial}{\partial x_l} \sum_{i=l}^{L-1} F(x_i, w_i)$ cannot equal -1 for every sample in a mini-batch.

I. Performance Metrics

An evaluation of the model's performance was carried out using a confusion matrix, a table that describes the accuracy with which a classification model handled a dataset when the true values were revealed [22]. By comparing the model's predicted values with the actual target values, the matrix gives information on the classifications' accuracy. The



primary measures that are gathered are F1-score, recall, accuracy, and precision. The confusion matrix made use of the following terms:

- **True Positive (TP):** A model properly forecast that the observation belongs to the class of pneumonia.
- **True Negative (TN):** A model accurately forecast that the observation is in the normal class, meaning that there is no pneumonia.
- **False Positive (FP):** The model mistakenly forecasts it as pneumonia, even though the data is really in the usual class.
- **False Negative (FN):** The observation actually belongs to the pneumonia class, but the model thinks it's normal.

Accuracy: The accuracy of the model is evaluated by determining the proportion of occurrences for which the predictions were accurate relative to the total number of instances. It is expressed Equation (6):

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (6)$$

Precision: It assesses how well the model can distinguish real positive cases from all other positive cases. In medical applications, ensuring pneumonia instances is very important. It is represented as Equation (7):

$$Precision = \frac{TP}{TP + FP} \quad (7)$$

Recall: Sensitivity is another name for this metric, which calculates the percentage of positive observations anticipated to occur out of all the observations in the actual class. Equation (8) shows the formula for calculating the recall:

$$Recall = \frac{TP}{TP + FN} \quad (8)$$

F1-score: The mean of the two metrics, recall and precision scaled together. As a result, this score considers both the good and negative outcomes. The following is a definition of the F1 score Equation (9):

$$F1 = \frac{2 * (precision * recall)}{precision + recall} \quad (9)$$

ROC: A graphical representation of a classification model's performance across different choice criteria is the ROC curve. The TPR and the FPR are plotted against one other to show how sensitivity and specificity could be compromised.

IV. RESULT & DISCUSSION

Results for the proposed ResNet50 model are detailed below with respect to recall, accuracy, precision, and f1-score.

The implementation work conducted on Python programming language and leveraging an NVIDIA RTX 4090 GPU and AMD EPYC 7453 CPU ensured efficient deep learning computations. Table II displays the results of the evaluation of the suggested model, ResNet50, using the dataset consisting of chest X-ray pictures.

ResNet50 model Performance on Chest X-Ray Images Dataset for Pneumonia Detection

Evaluation Matrix	Residual Network-50 (ResNet50)
Accuracy	93.06
precision	88.97
Recall	96.78
F1-score	92.71



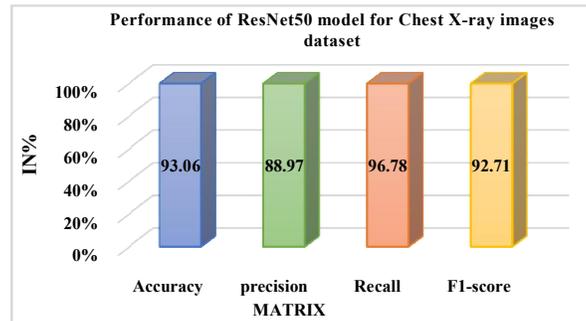


Figure 5: Bar Graph for ResNet50 Model Performance

The ResNet50 model's performance is illustrated by the bar graph in Figure 5 and Table II. The ResNet50 exhibits outstanding performance metrics, achieving an accuracy of 93.06%, a precision of 88.97%, and a recall of 96.78%. The F1-score of 92.71% shows that the model has a good balance between recall and precision, indicating that it can find instances properly with few false positives and negatives. These results highlight the robustness and reliability of the ResNet50 models.

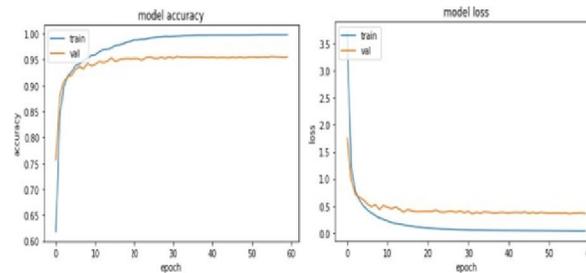


Figure 6: Accuracy and Loss Graphs for ResNet50 Model

Figure 6 Accuracy and Loss Trends of the ResNet50 Model Over 60 Epochs. The accuracy plot shows training accuracy (blue) increasing rapidly and converging near 1.0, while validation accuracy (orange) stabilizes around 0.95, indicating slight generalization. The loss plot illustrates training loss (blue) decreasing as learning progresses, whereas validation loss (orange) initially follows a similar trend but may plateau or rise, suggesting potential overfitting.

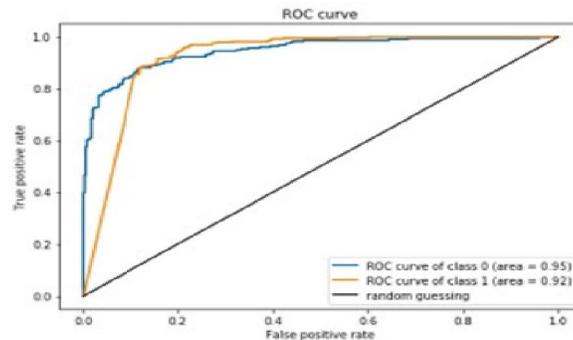


Figure 7: Roc Curve of ResNet50 Model

Figure 7 shows the ROC curve, which compares the TPR with the FPR to indicate how well the classification model performed. The model achieves high AUC values of 0.95 for class 0 and 0.92 for class 1, indicating strong discrimination between classes. The ROC curves significantly outperform the random guessing baseline (AUC = 0.5), demonstrating the model's effective classification capability.



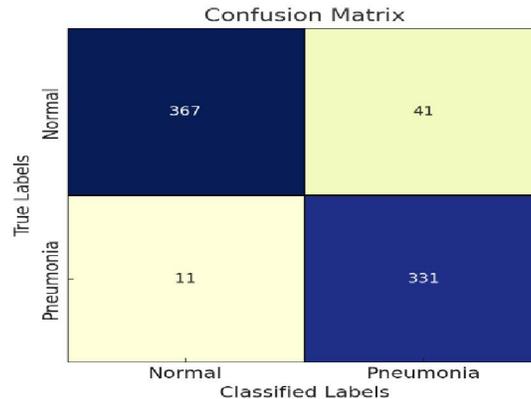


Figure 8: Confusion Matrix of the ResNet50 Model

Figure 8 shows the confusion matrix evaluating a classification model for distinguishing between normal and pneumonia cases. The model correctly classified 367 normal cases and 331 pneumonia cases while misclassifying 41 normal cases as pneumonia and 11 pneumonia cases as normal. The high number of correctly classified instances suggests strong model performance with relatively low misclassification rates. This shows that the model does a good job of distinguishing between typical instances and those with pneumonia.

A. Comparative Analysis

This section provides a comparison between the proposed ResNet50 against existing models including Mobile Net [23], EfficientNetB0 [24], and Xception [25] across the evaluation matrix, shown in Table III.

Comparison between ResNet50 and Existing model Performance for Chest X-ray images dataset

Matrix	Mobile Net[23]	EfficientNetB0[24]	Xception [25]	ResNet50
Accuracy	89.22	73.71	82	93.06
precision	94.71	74.35	86	88.97
Recall	89.54	88.46	65	96.78
F1-score	92.06	80.79	74	92.71

The performance comparison between the proposed ResNet50, Mobile Net, EfficientNetB0 and Xception models is shown in Table III. In this comparison, ResNet50 outperforms in all key metrics. With an accuracy of 93.06%, ResNet50 surpasses Mobile Net (89.22%), EfficientNetB0 (73.71) and Xception (82%). It also leads in precision (88.97%), recall (96.78%) and F1-score (92.71%), indicating that it not only classifies more accurately but also identifies and correctly classifies (Normal and Pneumonia) more effectively, In contrast, Mobile Net, EfficientNetB0 and Xception show lower performance across this measure, suggesting that the ResNet50 model offer superior generalization and classification power for this task.

The proposed ResNet50-based pneumonia detection model offers several advantages, including high accuracy (93.06%), superior recall (96.78%), and strong generalization capability compared to MobileNet, EfficientNetB0, and Xception. Its deep hierarchical structure enhances feature extraction, improving classification performance. Additionally, data augmentation and preprocessing techniques mitigate overfitting, ensuring robust predictions. The model's high recall reduces FN, making it suitable for critical clinical applications. However, potential challenges such as dataset bias and computational complexity must be addressed. Overall, the study demonstrates the feasibility of DL for pneumonia detection, highlighting its potential to enhance radiological assessments and improve patient outcomes.

Conclusion & Future Work

Pneumonia is a serious respiratory illness that can infect either lung and cause inflammation of the air sacs, which can then become filled with pus or fluid. Chest pain, fever, chills, and difficulty breathing are typical symptoms. Anyone can get it; however, those with compromised immune systems, young children, and the elderly are at increased risk.



The microorganisms that cause the illness might be fungus, viruses, or bacteria. The most common way to diagnose pneumonia is by employing a chest X-ray, which can detect lung abnormalities. This study demonstrates that the ResNet50 model can effectively identify instances of pneumonia using chest X-ray pictures. The model outperformed MobileNet, EfficientNetB0, and Xception, with results of 93.06% accuracy, 88.97% precision, 96.78% re-call, and an F1score of 92.71%. The findings provide credence to the model's prospective clinical uses by demonstrating its resilience in differentiating between normal and pneumonia instances. There are certain restrictions, though, such as the possibility of overfitting because of the size of the dataset, the restricted capacity to be applied to a variety of populations, and the requirement for explainability when making decisions on clinical adoption. Additionally, class imbalance may impact model performance, necessitating further enhancements. Exploring ensemble learning approaches to improve prediction accuracy, increasing the dataset's interpretability using explainable AI techniques, and incorporating multi-source and real-world clinical data will be the primary goals of future study.

REFERENCES

- [1] S. S. S. Neeli, "The Convergence of AI and Database Administration in Revolutionizing Healthcare," *ESP Int. J. Adv. Comput. Technol.*, vol. 2, no. 4, pp. 150–153, 2024, doi: 10.56472/25838628/IJACT-V2I4P119.
- [2] Suhag Pandya, "Integrating Smart IoT and AI-Enhanced Systems for Predictive Diagnostics Disease in Healthcare," *Int. J. Sci. Res. Comput. Sci. Eng. Inf. Technol.*, vol. 10, no. 6, pp. 2093–2105, Dec. 2023, doi: 10.32628/CSEIT2410612406.
- [3] M. H. A. S. Ashish Shiwlani, Sooraj Kumar, Samesh Kumar, Syed Umer Hasan, "Transforming Healthcare Economics: Machine Learning Impact on Cost Effectiveness and Value-Based Care," *Pakistan J. Life Soc. Sci.*, 2024.
- [4] V. Kolluri, "AI for Personalized Medicine: Analyzing How AI Contributes to Tailoring Medical Treatment to the Individual Characteristics of Each Patient," *IJRAR - Int. J. Res. Anal. Rev. (IJRAR)*, E-ISSN 2349-5138, 2023.
- [5] D. J. Alapat, M. V. Menon, and S. Ashok, "A Review on Detection of Pneumonia in Chest X-ray Images Using Neural Networks," *J. Biomed. Phys. Eng.*, vol. 12, no. 6, pp. 551–558, Dec. 2022, doi: 10.31661/jbpe.v0i0.2202-1461.
- [6] R. Kundu, R. Das, Z. W. Geem, G. T. Han, and R. Sarkar, "Pneumonia detection in chest X-ray images using an ensemble of deep learning models," *PLoS One*, 2021, doi: 10.1371/journal.pone.0256630.
- [7] D. Li, "Attention-enhanced architecture for improved pneumonia detection in chest X-ray images," *BMC Med. Imaging*, 2024, doi: 10.1186/s12880-023-01177-1.
- [8] A. J. Rahul Dattangire, Divya Biradar, "AI-Enhanced U-Net for Accurate Low-Grade Glioma Segmentation in Brain MRI: Transforming Healthcare Imaging," *2024 Third Int. Conf. Electr. Electron. Inf. Commun. Technol.*, vol. 1–6.
- [9] R. Siddiqi and S. Javaid, "Deep Learning for Pneumonia Detection in Chest X-ray Images: A Comprehensive Survey," *J. Imaging*, vol. 10, no. 8, 2024, doi: 10.3390/jimaging10080176.
- [10] S. Pandya, "Predictive Modeling for Cancer Detection Based on Machine Learning Algorithms and AI in the Healthcare Sector," *TIJER – Int. Res. J.*, vol. 11, no. 12, 2024.
- [11] M. Yaqub *et al.*, "Deep Learning-Based Image Reconstruction for Different Medical Imaging Modalities," *Comput. Math. Methods Med.*, vol. 2022, p. 8750648, 2022, doi: 10.1155/2022/8750648.
- [12] C. Usman, S. U. Rehman, A. Ali, A. M. Khan, and B. Ahmad, "Pneumonia Disease Detection Using Chest X-Rays and Machine Learning," *Algorithms*, vol. 18, no. 2, 2025, doi: 10.3390/a18020082.
- [13] A. Roy, A. Bhattacharjee, D. Oliva, O. Ramos-Soto, F. J. Alvarez-Padilla, and R. Sarkar, "FA-Net: A Fuzzy Attention-aided Deep Neural Network for Pneumonia Detection in Chest X-Rays," in *2024 IEEE 37th International Symposium on Computer-Based Medical Systems (CBMS)*, 2024, pp. 338–343. doi: 10.1109/CBMS61543.2024.00063.
- [14] E. Jain and R. Kumar, "Deep Learning in Medical Imaging: CNN Performance in Pneumonia Detection from X-rays," in *2024 2nd International Conference on Self Sustainable Artificial Intelligence Systems (ICSSAS)*, 2024, pp. 182–186. doi: 10.1109/ICSSAS64001.2024.10760397.
- [15] S. Shrimali, "Optimizing Chest X-ray Analysis for Pneumonia Detection Through Comparative Evaluation of Transfer Learning CNN Architectures," in *2024 IEEE 9th International Conference for Convergence in Technology (I2CT)*, Apr. 2024, pp. 1–5. doi: 10.1109/I2CT61223.2024.10543957.



- [16] C. Wang, C. Xu, Y. Zhang, and P. Lu, "Diagnosis of Chest Pneumonia with X-ray Images Based on Graph Reasoning," *Diagnostics*, vol. 13, no. 12, 2023, doi: 10.3390/diagnostics13122125.
- [17] A. A. B. Ramli, Z. Zulkifli, S. Ahmad, and N. Ghazali, "Automatic Pneumonia Detection Through Chest X-Ray Image-Based," in *2023 4th International Conference on Artificial Intelligence and Data Sciences: Discovering Technological Advancement in Artificial Intelligence and Data Science, AiDAS 2023 - Proceedings*, 2023. doi: 10.1109/AiDAS60501.2023.10284669.
- [18] M. R. Hasan, S. M. A. Ullah, and M. Hasan, "Deep Learning in Radiology: A Transfer-Learning Based Approach for the Identification and Classification of COVID-19 and Pneumonia in Chest X-ray Images," in *2023 Fourth International Conference on Smart Technologies in Computing, Electrical and Electronics (ICSTCEE)*, Dec. 2023, pp. 1–6. doi: 10.1109/ICSTCEE60504.2023.10585226.
- [19] S. Nokhwal, P. Chilakalapudi, P. Donekal, S. Nokhwal, S. Pahune, and A. Chaudhary, "Accelerating Neural Network Training: A Brief Review," *ACM Int. Conf. Proceeding Ser.*, pp. 31–35, 2024, doi: 10.1145/3665065.3665071.
- [20] A. Çınar, M. Yıldırım, and Y. Eroğlu, "Classification of pneumonia cell images using improved ResNet50 model," *Trait. du Signal*, 2021, doi: 10.18280/TS.380117.
- [21] J. Q. Gandhi Krishna, "Implementation Problems Facing Network Function Virtualization and Solutions," *IARIA*, pp. 70–76, 2018.
- [22] S. Singh, M. Kumar, A. Kumar, B. K. Verma, K. Abhishek, and S. Selvarajan, "Efficient pneumonia detection using Vision Transformers on chest X-rays," *Sci. Rep.*, vol. 14, no. 1, p. 2487, 2024, doi: 10.1038/s41598-024-52703-2.
- [23] M. S. Al Reshan *et al.*, "Detection of Pneumonia from Chest X-ray Images Utilizing MobileNet Model," *Healthc.*, 2023, doi: 10.3390/healthcare11111561.
- [24] Q. An, W. Chen, and W. Shao, "A Deep Convolutional Neural Network for Pneumonia Detection in X-ray Images with Attention Ensemble," *Diagnostics*, 2024, doi: 10.3390/diagnostics14040390.
- [25] E. Ayan and H. M. Ünver, "Diagnosis of pneumonia from chest X-ray images using deep learning," in *2019 Scientific Meeting on Electrical-Electronics and Biomedical Engineering and Computer Science, EBBT 2019*, 2019. doi: 10.1109/EBBT.2019.8741582.

