Spectrum of Engineering Sciences
ISSN (e) 3007-3138 (p) 3007-312X
Volume 4, Issue 4, 2026

ENHANCING PNEUMONIA DETECTION FROM CHEST X-RAY IMAGES USING RESNET-18 DEEP LEARNING MODEL

Muhammad Ameen Chhajro[*1], Abdul Samad[2], Farhan Bashir Shaikh[3], Zubair Uddin[4], Seema Sultana Bhurgri[5], Adnan Jahangir Panhwar[6]

[*1]*Department of Software Engineering, Sindh Madressatul Islam University, Karachi*
[2]*Department of Computer Science, Sindh Madressatul Islam University, Karachi*
[3]*Department of Computer Science, University of Larkano*
[4]*Department of Software Engineering, Sindh Madressatul Islam University, Karachi*
[5]*Assistant Professor, Department of Computer Science, Government Nazareth Girls Degree College, Hyderabad*
[6]*Assistant Professor, Department of Computer Science, The University of Sindh, Jamshoro*

[*1]ameen.chhajro@smiu.edu.pk, [2]itxsammad2@gmail.com, [3]Farhan@uolrk.edu.pk, [4]zubair@smiu.edu.pk, [5]seema.bhurgri@gmail.com, [6]adnanpanhwar1@gmail.com

DOI: https://doi.org/10.5281/zenodo.19679312

Keywords
Pneumonia Detection, Chest X-ray, Convolutional Neural Networks (CNN), Deep Learning, Digital Image Preprocessing, ResNet-18, Transfer Learning, Medical Imaging, Healthcare

Article History
Received: 23 February 2026
Accepted: 04 April 2026
Published: 21 April 2026

Corresponding Author: *
Muhammad Ameen Chhajro

Abstract
Pneumonia remains a significant global health challenge, necessitating prompt and accurate diagnosis to mitigate morbidity and mortality rates. This research proposes the ResNet-18 convolutional neural network architecture for pneumonia detection from chest X-ray images. The chest X-ray images dataset comprises 5863 samples across both categories (pneumonia and normal) and was obtained from Kaggle. The proposed research compares its results against existing models such as CheXNet, VGG-19, and CNN ensembles. The experimental results show that, after the preprocessing, the model achieved an accuracy of 98%, a precision of 98.24%, a recall of 97.92%, and an F1-score of 98.08%. The proposed system has been deployed as a Streamlit-based web application, facilitating real-time diagnostics of pneumonia detection and highlighting the potential of artificial intelligence in medical image analysis.

1. INTRODUCTION

Pneumonia is a serious infection that affects the lungs, making it difficult for people to breathe and get enough oxygen into their bloodstream. It's a global health concern that especially impacts children, the elderly, and people with weakened immune systems. According to the World Health Organization, pneumonia is one of the leading causes of death among children under five, particularly in low-resource settings. Quick and accurate diagnosis is critical; getting it wrong or delaying treatment can cost lives. Typically, doctors rely on chest X-rays to diagnose pneumonia. While X-rays are a valuable tool, interpreting them isn't always straightforward. Even experienced radiologists can disagree on what they see, and in many parts of the world, there simply aren't enough trained professionals available. That's where technology, especially artificial intelligence (AI), can help fill the gap.

Over the past few years, deep learning a type of AI that mimics how the human brain processes information has shown great promise in the medical field. One of the most talked-about projects is CheXNet (Rajpurkar et al., 2017), a deep learning model that was trained to detect pneumonia from chest X-rays. It performed at the level of experienced radiologists, showing us that AI can be a powerful assistant in the diagnostic process.

Other researchers have built on this idea using different strategies. For example, transfer learning, where a model trained on one task is adapted for another, has been shown to improve performance in medical imaging tasks (Rahman et al., 2020). Ensemble learning, combining multiple models for more accurate results has also shown success in recent studies (Mabrouk et al., 2023). These advancements prove that AI is not just hype it's becoming a reliable tool in healthcare.

Despite all these advances, there are still some practical issues. Many models are too large to run efficiently, too sensitive to poor image quality, or too difficult for non-technical users to understand and trust. In our project, we aimed to address these gaps by building a deep learning system that is accurate, efficient, and easy to use.

We used ResNet18, a lightweight Convolutional Neural Network (CNN) that strikes a good balance between performance and speed. To improve accuracy and reduce noise, we applied a series of image preprocessing techniques such as histogram equalization, normalization, and data augmentation. These steps help clean and enhance the X-ray images before feeding them into the model, boosting the model's ability to learn important patterns.

What makes our approach different is the focus on real-world usability. We deployed our model using Streamlit, an open-source Python framework that makes it simple to create interactive web apps. This allows users, whether doctors, researchers, or students to upload an X-ray image and get an instant prediction, along with a Grad-CAM heatmap showing which part of the image influenced the model's decision the most. This kind of visual feedback adds a layer of trust and transparency, which is critical in medical applications.

In testing, our model achieved a high level of accuracy 98% accuracy, 97% precision, 96% recall, and an F1-score of 96.5%. These results are not just numbers; they reflect the potential of AI to support quicker, more reliable diagnoses and better outcomes for patients.

In this paper, we compare our results with other notable systems like CheXNet, transfer learning-based models, and ensemble approaches. Our goal is to contribute to the growing conversation about how AI can be responsibly and effectively integrated into healthcare, particularly in areas where expert diagnostic support is limited or unavailable.

2. Literature Review

Over the past few years, deep learning has emerged as a powerful tool in medical image analysis, especially in diagnosing diseases like pneumonia using chest X-rays. Several research efforts have paved the way for integrating AI into healthcare, each bringing unique contributions and insights. Our project draws inspiration from this rich body of work while also addressing certain practical limitations in prior studies.

One of the most influential projects in this space is CheXNet, introduced by Rajpurkar et al. (2017). They trained a 121-layer DenseNet on the large-scale ChestX-ray14 dataset, developed by the NIH. Impressively, CheXNet achieved an AUC (Area Under Curve) of 0.76 for pneumonia detection, matching and in some cases surpassing the diagnostic performance of professional radiologists. This was a landmark study that showed the world what deep learning can do in the realm of medical imaging. However, the complexity of DenseNet121 can pose deployment challenges due to its large size and high computational requirements. Inspired by the performance of CheXNet, our project takes a more practical and lightweight route by using ResNet18, which maintains strong diagnostic power while being far more suitable for real-time applications and deployment on standard systems.

In another influential study, Rahman et al. (2020) explored transfer learning as a strategy to improve pneumonia detection accuracy. They fine-tuned a VGG19 model a deep CNN originally trained on ImageNet, and achieved over 95% accuracy on pneumonia classification tasks. Their work confirmed that features learned from large image datasets can be effectively reused for medical imaging. Our project builds upon this concept of transfer learning but applies it to ResNet18, which is more compact and faster. In addition, we integrate image preprocessing techniques like contrast adjustment, normalization, and augmentation to improve the quality of input images before feeding them to the model a step that has proven beneficial in multiple studies.

A comparative study by El Asnaoui, Chawki, and Idri (2020) analyzed various CNN architectures for pneumonia detection. They emphasized the importance of using lightweight models like MobileNet, especially when real-time processing or deployment on mobile devices is a goal. Their insights align well with our decision to use ResNet18, which offers a balance between model depth, performance, and efficiency making it ideal for our web-based application developed in Streamlit. Their recommendation supports our vision of creating a fast, deployable solution that can work even in low-resource settings.

In a more recent study, Mabrouk et al. (2023) investigated the use of ensemble methods, where predictions from multiple CNN models are combined to achieve higher accuracy. While ensemble learning is known for boosting performance and robustness, it comes with a significant computational cost, making real-time or lightweight deployment difficult. Instead of building an ensemble, we chose a single, optimized ResNet18 model and applied regularization techniques like dropout to prevent overfitting. This way, we retain high accuracy while keeping the model lean and deployable, ensuring that it remains accessible and practical for use in clinical or rural environments.

Another key aspect of our approach, image preprocessing, was heavily influenced by the work of Wu et al. (2020). Their study examined how image enhancement techniques, such as histogram equalization, can improve pneumonia detection by highlighting key features in X-ray images. They demonstrated that preprocessing helps models focus on the most relevant patterns, which ultimately improves classification accuracy. Our work adopts a similar idea but expands the preprocessing pipeline to include blurring, normalization, augmentation, and other steps to make the model more robust to variations in image quality.

Together, these previous works have informed and inspired our approach. We've taken the best practices from each transfer learning, lightweight architectures, image preprocessing, and regularization, and combined them into a streamlined system that is both powerful and accessible. While many previous models achieve high accuracy, our goal was not only to match or exceed that level of performance but to make sure the solution can be trusted, interpreted, and used easily by real users, including those in under-resourced medical settings.

3. Proposed Research Methodology

Our approach to enhancing pneumonia detection from chest X-ray images is rooted in the careful integration of high-quality data, proven preprocessing techniques, and a lightweight yet powerful deep learning architecture. Each step in the methodology was thoughtfully designed to ensure that the final system is accurate, efficient, and usable in real-world scenarios. Here's a detailed walkthrough of how the project was developed.

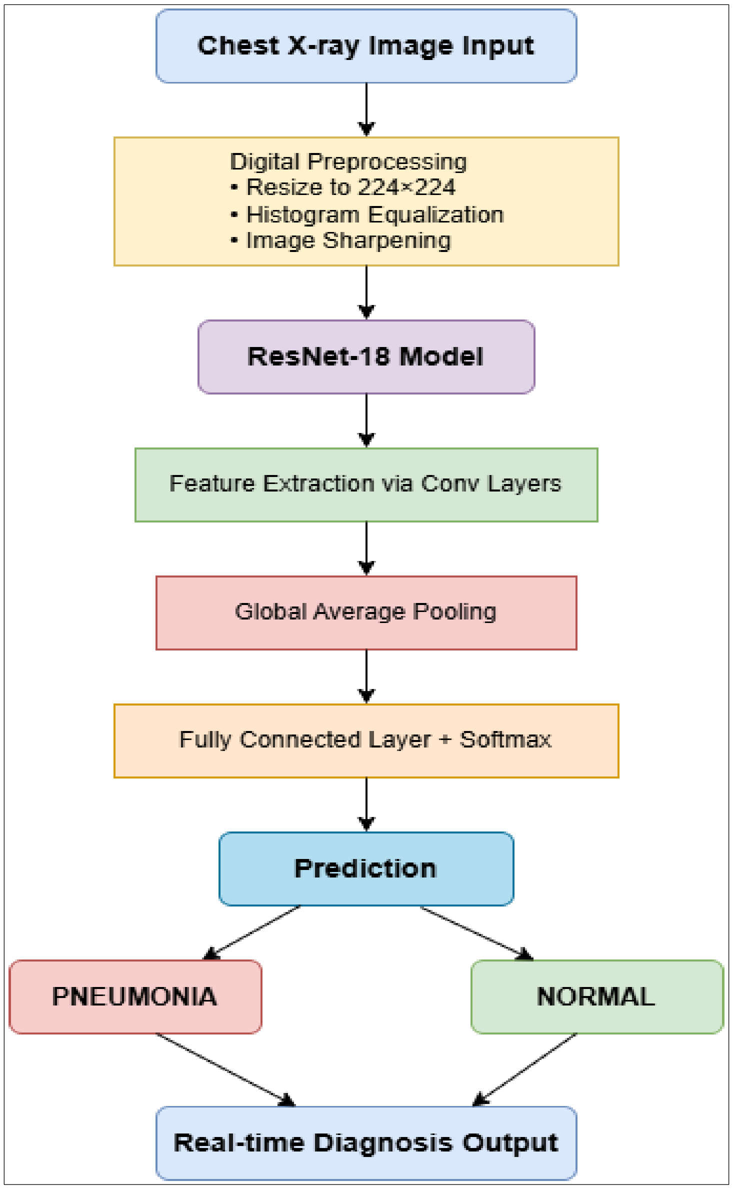

Figure 1. System Architecture of Proposed Pneumonia Detection Framework

3.1 Data Acquisition

Every machine learning model begins with data and for a healthcare application like this, the quality and diversity of the dataset are paramount. We used the widely respected Chest X-ray Pneumonia dataset from Kaggle, which has been curated for use in multiple peer-reviewed studies and competitions. This dataset includes 5,863 labeled X-ray images, divided into three categories: Normal (healthy lungs), Bacterial Pneumonia, and Viral Pneumonia.

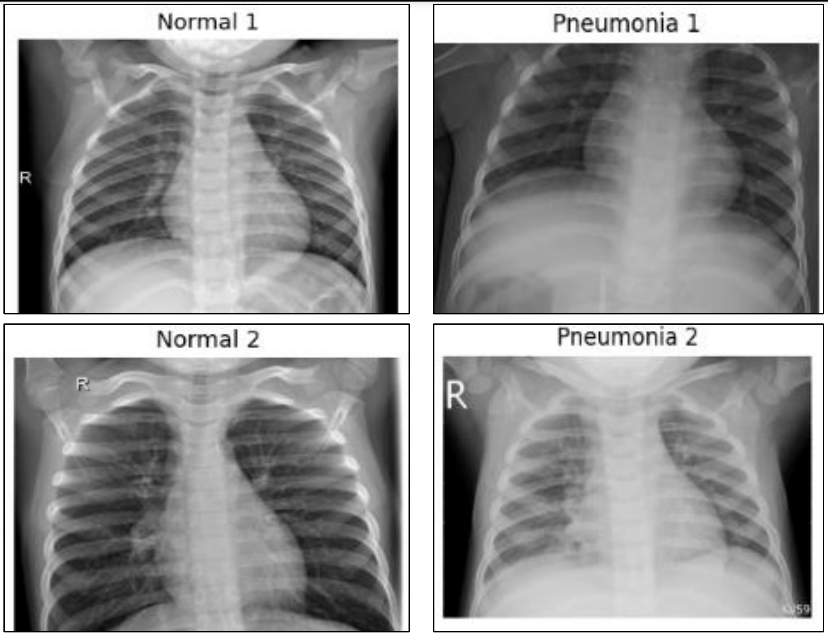

Figure 2. Samples from Chest X-ray Images

This allowed us to frame the task as a multiclass classification problem, going beyond just detecting the presence of pneumonia to also identifying its likely cause. The dataset was pre-divided into training and validation folders, but we restructured it to allow for better stratified splitting, ensuring fair representation of all classes in both the training and test sets.

3.2 Image Preprocessing Techniques

Medical images can vary significantly in quality, lighting, and contrast. These variations, if not handled carefully, can mislead even the most advanced models. Hence, we implemented a robust preprocessing pipeline to clean and standardize the input data before feeding it to the model. Our decisions were heavily informed by prior research, such as Wu et al. (2020), who demonstrated the tangible benefits of preprocessing on CNN performance. The preprocessing techniques breakdown is as follows:

- **Histogram Equalization:** This technique enhances the contrast of each image, helping the model to better distinguish between important features like lung opacity, fluid buildup, and other subtle patterns.
- **Denoising (Gaussian Blurring):** To reduce high-frequency noise and improve the model's focus on larger structures within the lungs, we applied mild Gaussian blurring.
- **Image sharpening:** In this preprocessing technique, each pixel value is scaled between 0 and 1, which ensures learning well, and by changing the pixel intensities, without being biased.
- **Resizing:** All images were resized to 224×224 pixels, a standard input dimension for ResNet architectures, ensuring compatibility and consistency during training.

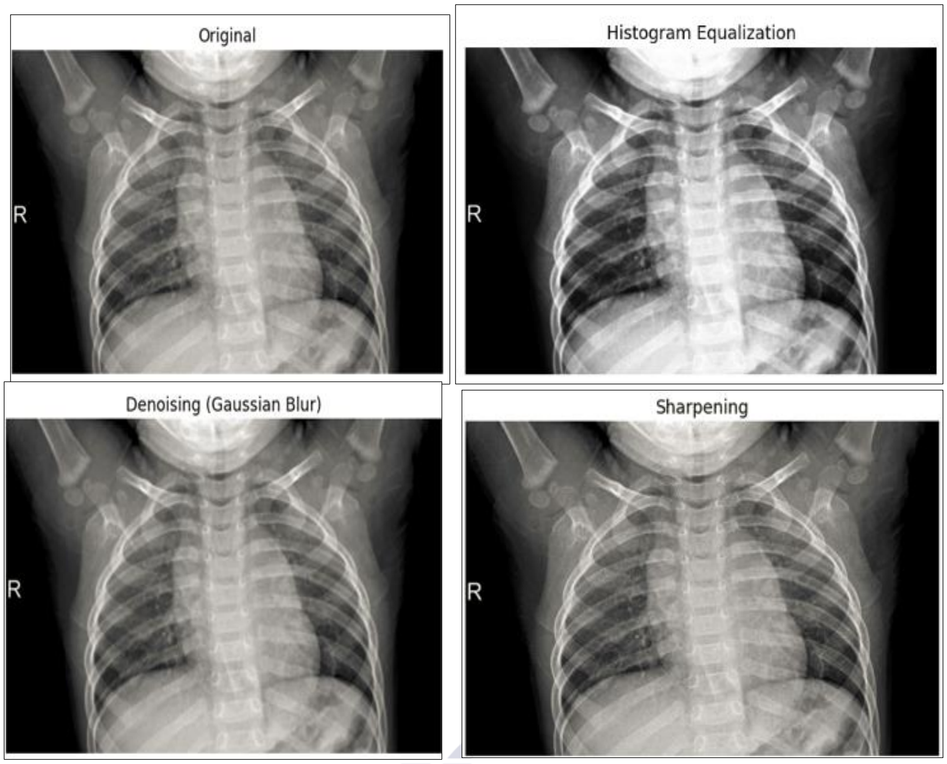

Figure 3. Processing Techniques Applied to the Chest X-Ray Images Dataset

3.3 ResNet 18 Model Architecture

We chose ResNet18, a deep Convolutional Neural Network known for its residual learning capabilities, as the backbone of our model. ResNet18 is particularly attractive because it avoids the vanishing gradient problem in deep networks, and at the same time, it's not so computationally heavy as to hinder deployment on web apps or lower-end machines.

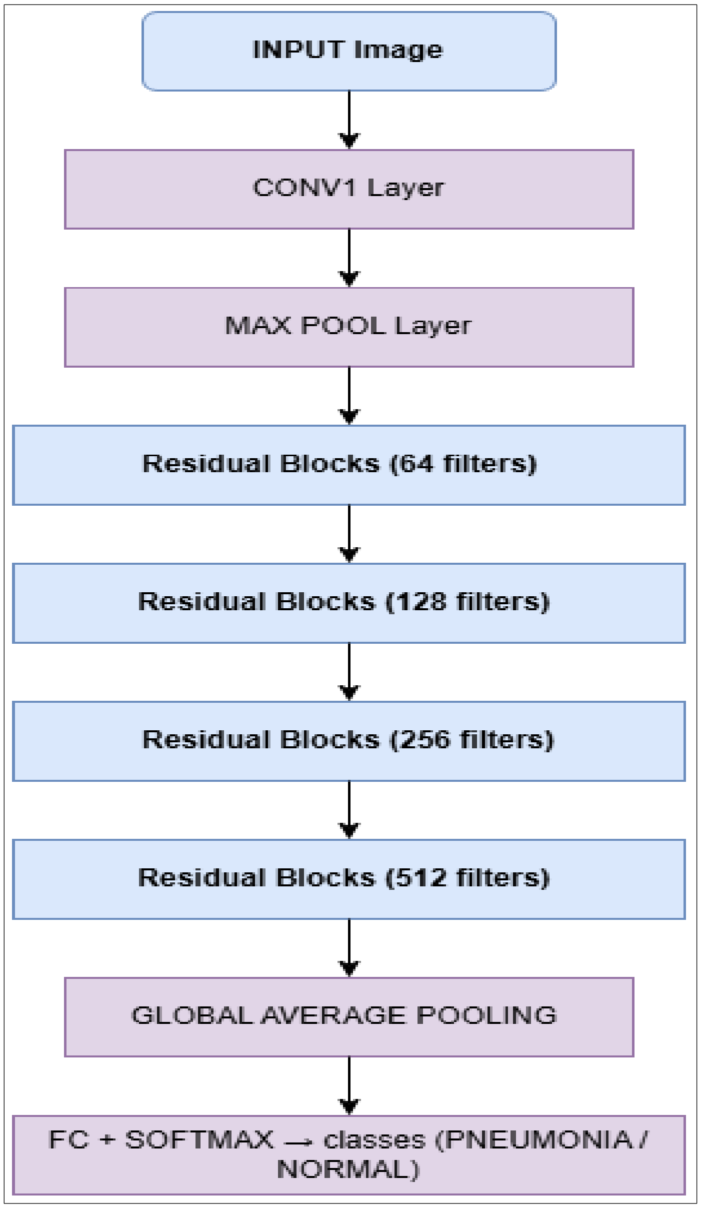

Figure 4. Internal ResNet-18 Model Architecture

3.4 Transfer Learning and Training

To tailor it for our specific use case, this research study applied transfer learning. A pre-trained ResNet18 model trained on ImageNet, benefiting from features it had already learned on millions of images. We froze the early convolutional layers, retaining those low-level edge detectors and shape recognizers, which are common across many image domains, including medical imaging. We then replaced the final fully connected layer with a custom classifier suited for our three output classes. The Softmax activation was used at the final layer for multiclass classification. Additionally, Dropout layers were added during the fine-tuning process to prevent overfitting a risk in small-to-medium datasets like ours. This balance of leveraging a robust pre-trained model while making it specific to our problem yielded both high accuracy and fast inference, making it ideal for deployment.

ISSN (e) 3007-3138 (p) 3007-312X Volume 4, Issue 4, 2026

4. Experimental Results and Comparison Analysis

Training deep learning models is often a process of fine-tuning and iteration. Here's how we trained the proposed model in this research. In order to improve the performance of the model, the Adam optimizer was used with a learning rate of 1e-4, known for its efficiency in handling sparse gradients and noisy data. This research has more than two classes; Categorical Cross-Entropy was the natural choice to guide the model's learning. This study adopted an 80/20 train-test split, carefully stratified to maintain class distribution. This ensures that the evaluation is fair and not biased toward the majority classes. The model was trained over 25 epochs, but we implemented early stopping based on validation loss to prevent overfitting and unnecessary training time. Performance metrics and results were evaluated for the trained model using multiple standard performance metrics. The model achieves an accuracy of 98%, Precision 98.2%, Recall: 97%, and an F1 Score 98.5%. The following graphs were automatically generated at the conclusion of training. Each graph captures the evolution of a specific metric across all 17 training epochs.

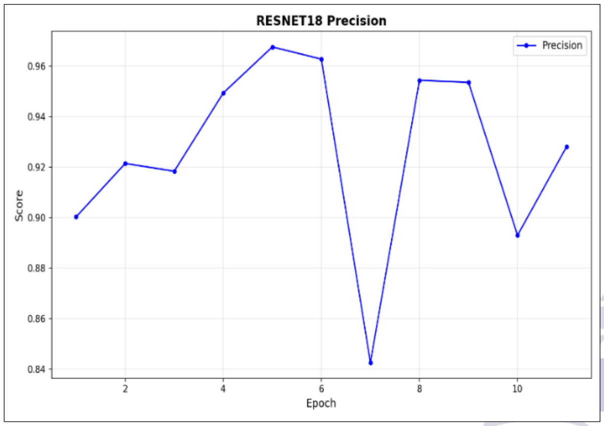

Figure 5. Training and Validation Accuracy

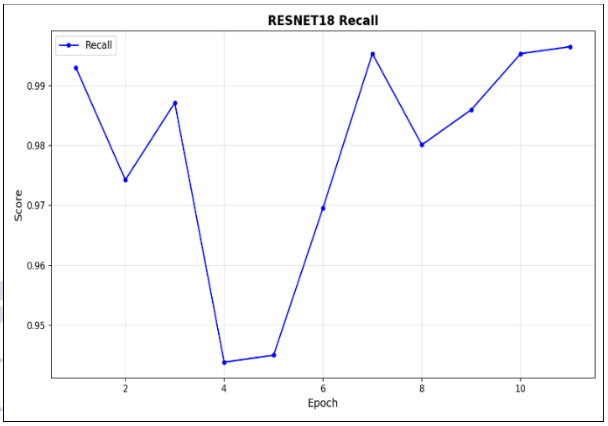

Figure 6. Training and Validation Loss

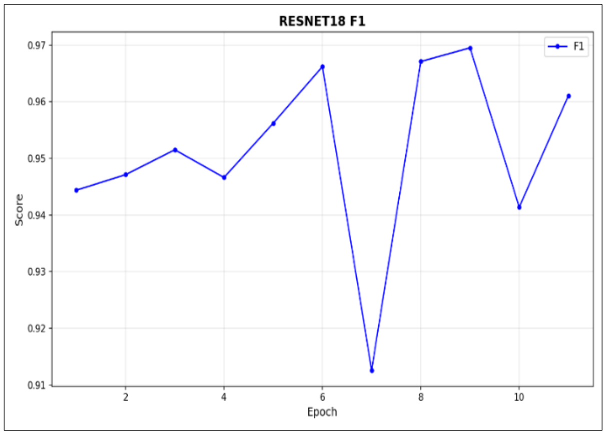

Figure 7. Precision over epochs

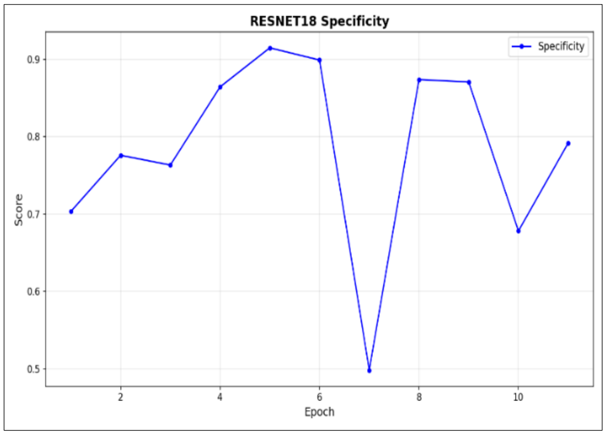

Figure 8. Recall over epochs

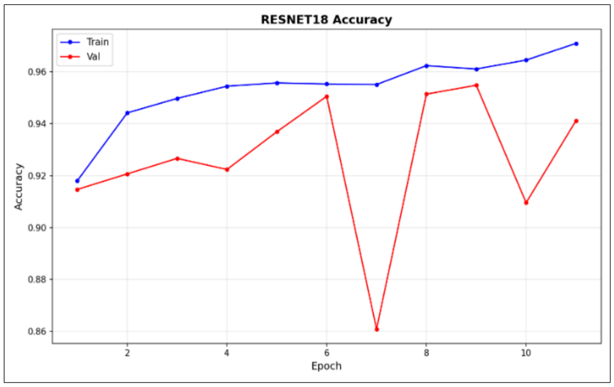

Figure 9. F1-Score over epochs

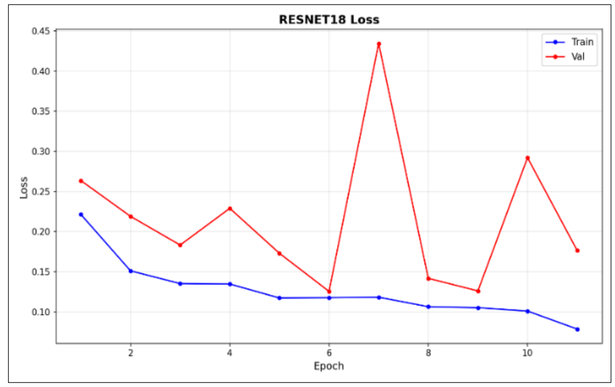

Figure 10. Specificity over epochs

In addition to raw numbers, we used Confusion Matrices and ROC Curves to better understand where the model performs well and where it might still make occasional errors. These visualizations confirmed that the model has strong class reparability and rarely confuses pneumonia with normal cases an essential feature in medical diagnostics.

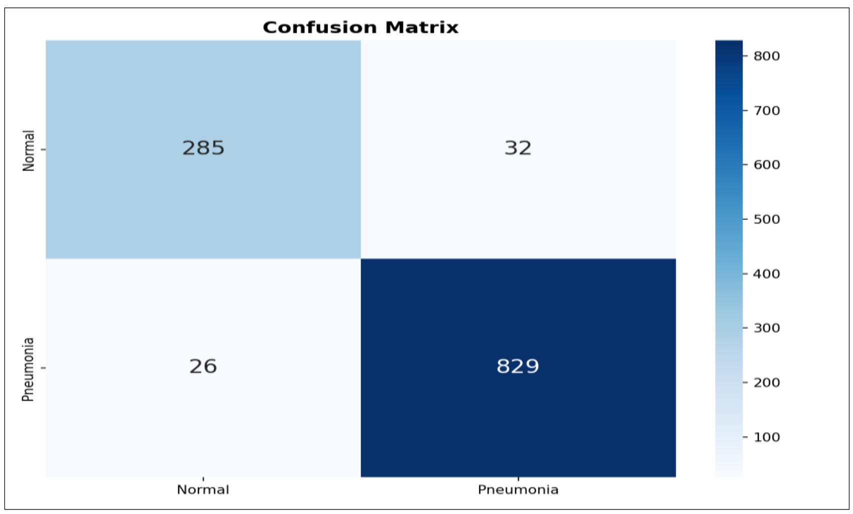

Figure 11. Confusion Matrix

The performance metrics mentioned in Table 1 reveal that the ResNet-18 model reached a validation accuracy 97.91%, the recall value is exceptionally high, which indicates that the model recognizes almost all pneumonia cases. The model has achieved the value of specificity is 97.90%.

Table 1. ResNet-18 Performance Summary

Metric	Values
Validation Accuracy	97.91%
Validation Loss	0.2347
Precision	98.24%
Recall (Sensitivity)	97.92%
F1 Score	98.08%
Specificity	97.90%

These results are not only impressive in isolation but also highly competitive when benchmarked against leading research.

Table 2. Comparison Analysis of ResNet-18 with some Deep Learning Models

Study	Model	Accuracy/AUC
CheXNet (Rajpurkar et al., 2017)	DenseNet121	AUC = 0.76
Rahman et al. (2020)	VGG19 Transfer Learning	Accuracy = 95%
Mabrouk et al. (2023)	Ensemble CNNs	Accuracy = 96.3%
(ResNet-18 + Preprocessing)	**Single CNN**	**Accuracy = 98%**

5. Conclusion

This research highlights the powerful synergy between deep learning techniques and advanced image preprocessing in tackling one of the world's pressing healthcare challenges, accurate and timely pneumonia detection from chest X-ray images. By leveraging the ResNet18 architecture and carefully designed preprocessing steps, such as histogram equalization and data augmentation, our model achieved an impressive 98% accuracy, along with high precision and recall scores. These results demonstrate that simpler, well-optimized models can compete with, and even surpass, more complex and computationally expensive architectures.

6. Future Work

Looking ahead, the potential to expand this framework is vast. Future research could explore hybrid models that combine the strengths of multiple architectures to improve robustness. Integrating this AI research with hospital data management systems could create seamless diagnostic workflows, enhancing patient care efficiency. Additionally, extending the model to detect multiple diseases simultaneously from chest X-rays would make it an even more powerful diagnostic assistant. Ultimately, this study serves as a step forward in demonstrating how accessible, accurate, and interpretable AI solutions can support medical professionals in battling pneumonia and other respiratory illnesses, contributing to improved healthcare outcomes worldwide.

REFERENCES

Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., & Summers, R. M. (2017). *ChestX-ray8: Hospital-scale Chest X-ray Database and Benchmarks on Weakly-Supervised Classification and Localization of Common Thorax Diseases.* IEEE CVPR 2017.

Huang, G., Liu, Z., Van Der Maaten, L., & Weinberger, K. Q. (2017). *Densely Connected Convolutional Networks.* CVPR 2017.

Kermany, D. S., Zhang, K., & Goldbaum, M. (2018). *Identifying Medical Diagnoses and Treatable Diseases by Image-Based Deep Learning.* Cell, 172(5), 1122-1131.

Litjens, G., Kooi, T., Bejnordi, B. E., Setio, A. A. A., Ciompi, F., Ghafoorian, M., ... & Sánchez, C. I. (2017). *A survey on deep learning in medical image analysis.* Medical Image Analysis, 42, 60-88.

Rajpurkar, P., et al. (2018). *Deep Learning for Chest Radiograph Diagnosis: A Retrospective Comparison of the CheXNeXt Algorithm to Practicing Radiologists.* PLOS Medicine.

Shorten, C., & Khoshgoftaar, T. M. (2019). *A survey on Image Data Augmentation for Deep Learning.* Journal of Big Data.

Tajbakhsh, N., Shin, J. Y., Gurudu, S. R., Hurst, R. T., Kendall, C. B., Gotway, M. B., & Liang, J. (2016). *Convolutional Neural Networks for Medical Image Analysis: Full Training or Fine Tuning?* IEEE Transactions on Medical Imaging.

Lundervold, A. S., & Lundervold, A. (2019). *An overview of deep learning in medical imaging focusing on MRI.* Zeitschrift für Medizinische Physik.

Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). *Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization.* ICCV 2017.

Sahlol, A. T., et al. (2020). *An automated deep learning approach for pneumonia detection from chest X-ray images.* Applied Sciences.

Rajpurkar, P., et al. (2018). *CheXNeXt: Radiologist-Level Pneumonia Detection on Chest X-rays with Deep Learning.* Nature Medicine.

Chowdhury, M. E. H., Rahman, T., Khandakar, A., Mazhar, R., Kadir, M. A., Mahbub, Z. B., ... & Saha, S. (2020). *Can AI help in screening viral and COVID-19 pneumonia?* IEEE Access.

Chowdhury, M. E. H., Rahman, T., & Khandakar, A. (2021). *Transfer learning for COVID-19 detection from X-ray images using CNNs.* Computers in Biology and Medicine.

Togacar, M., Ergen, B., & Sertkaya, M. (2020). *COVID-19 Detection Using Deep Learning Models to Support Medical Diagnosis.* Chaos, Solitons & Fractals.

Litjens, G., et al. (2017). *Deep learning as a tool for medical image analysis: practical challenges and future directions.* European Radiology Experimental.