Indian Journal of Computer Science and Technology
Volume 5, Issue3 (September-December 2026), PP: 382-388.
www.indjcst.com
ISSN No: 2583-5300

Detecting Pneumonia from Chest X-Rays Using Convolutional Neural Network Model

Arumita Nandi[1], Manaswi[2], Nikhil AE[3], Poornachandra BR[4], P Narashimaraja[*5]
*[1, 2, 3, 4, *5]RV College of Engineering, Bengaluru, Karnataka, India.*
[*]Corresponding Author

To Cite this Article: *Arumita Nandi[1], Manaswi[2], Nikhil AE[3], Poornachandra BR[4], P Narashimaraja[*5], "Detecting Pneumonia from Chest X-Rays Using Convolutional Neural Network Model", Indian Journal of Computer Science and Technology, Volume 05, Issue 03 (September-December 2026), PP: 382-388.*

Copyright: ©2026 This is an open access journal, and articles are distributed under the terms of the Creative Commons Attribution License; Which Permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Abstract: *Pneumonia is a leading cause of mortality all across the globe. The lack of skilled medical expertise in regions where radiology expertise is limited exacerbates the problem. The most common way to detect pneumonia is via atrained medical observation and analysis of chest X-rays. Inconsistencies in interpretation and shortage of skilled radiologists pose a significant challenge. The aim of this study is to present a novel deep learning-based approach for auto detection of pneumonia with a high level of accuracy. In this approach we utilize the DenseNet121 architecture with transfer learning. The proposed methodology employs a systematic multi-stage process encompassing automated data loading, real-time augmentation, class imbalance handling, and network optimization. A publicly available chest X-ray dataset was processed using TensorFlow and Keras. The trained model achieved high predictive performance of 91.0%, an AUC score of 0.9511, a F1 score of 0.9279 at optimal threshold of 0.68, demonstrating its capability to distinguish pneumonia from normal cases with strong generalization on unseen data. Future directions include interface development for real-time diagnostics and expansion to multi-class classification of thoracic diseases.*

Key Word: *Convolution Neural Network (CNN), pneumonia, chest x-rays, DenseNet121, Healthcare, ROC, PIL*

I. INTRODUCTION

Pneumonia remains one of the leading causes of morbidity and mortality worldwide, particularly in low-resource regions where access to skilled radiologists is limited. Chest X-rays continue to be the most widely used diagnostic tool for pneumonia due to their accessibility, cost-effectiveness, and speed. However, the interpretation of these images is often prone to significant inter-observer variability and reduced diagnostic accuracy, especially in overburdened clinical settings. Recently, automated, accurate, and consistent interpretation of radiological data has enabled deep learning to emerge as a promising solution for radiological diagnosis[1,2,3]

This work proposes a deep learning-based framework for automated detection of pneumonia. It leverages on the DenseNet121 [4] convolutional neural network architecture and uses transfer learning to increase learning speed and efficiency. The method used is a systematic multi-stage pipeline that comprises data loading, preprocessing, augmentation, class balancing, model training and hyperparameter optimization. The NIH Chest X-ray database that is publicly available was used for training and validating the data. The data was prepared using TensorFlow and Keras. Image enhancement techniques such as random horizontal flipping, zooming, brightness variation, and contrast adjustment were applied. This was done for generalisation and to increase the robustness of the model.

The inherent class imbalance within the dataset, was addressed via incorporation of dynamically computed class weights, thus ensuring balanced learning across both pneumonia and non-pneumonia cases. Model training was carried out in two phases. In the initial phase, only the classification head was trained while freezing the DenseNet121 backbone. For the next phase, the entire network was fine-tuned to extract deeper feature representations, thereby improving classification performance. Binary cross-entropy was used as the loss function, and the Adam optimizer with adaptive learning rate scheduling was used.

During the initial phase of this project, a survey was conducted with the end-users to understand the gaps in the technology that was currently available, as well as perceptions, needs, and concerns. The survey was conducted amongst patients, care givers and healthcare professionals regarding the extent to which Artificial Intelligence (AI) for pneumonia diagnosis was being used by the healthcare industry. The core objective was to uncover latent expectations, apprehensions and attitudes toward AI-based systems in medical imaging. This survey revealed interesting insights on the use and efficacy of AI in healthcare from amongst a wide range of civil and medical society. We also consulted two medical practitioners regarding the use of AI in healthcare and its positive impacts.

The performance of the model was assessed using widely accepted classification metrics. This included accuracy, area under the receiver operating characteristic curve (AUC), as well as the F1-score. The F1-score was maximised via threshold tuning, and model robustness validated via confusion matrix analysis and comprehensive classification reports. The final model achieved a strong diagnostic performance on the test set. Training and validation curves demonstrated stable convergence and minimal

overfitting.

The study brings out the effectiveness of DenseNet121-based transfer learning towards automating detection of pneumonia. It is done via analysing chest X-rays. Such models can be integrated into clinical decision systems that augment and support medical diagnosis. It can enhance the accuracy of diagnosis in places with limited access to radiology expertise. Future directions include extension of the framework to detection of respiratory diseases beyond pneumonia. Real-world deployment in healthcare environments should be accompanied by developing user friendly operating interfaces[11,8].

II. RELATED WORKS

The detection of pneumonia from Chest X-rays images by using a trained CNN model aims to leverage the use of artificial intelligence to enhance the efficiency and reliability of medical diagnosis. By automating the interpretation of radiographic images, the system minimizes inter-observer variability, and increases the accuracy of clinical decision-making. It will also ensure timely identification of pneumonia; more so in settings where access to trained radiologists are limited. Ultimately, the integration of CNN-based models into clinical workflows has the potential to support physicians, thereby contributing to a reduction of pneumonia-related mortality. The model was developed by furthering earlier work reported by researchers in this field.

Shagum Sharma et al [1] present a systematic literature summary and discussion on CNN models which are pre-trained to detect pneumonia and it provides research gaps in the existing models and Wasif Khan et al [2] reviews all the current available literature on the topic of pneumonia detection using chest-X rays.

N. Shilpa et al [3] focus on using various models such as ResNet50, MobileNetV2, AlexNet, EfficientNetB0, and Xception, showing that the EfficientNetB0 performs consistently much better than the other models, giving 99.78% test accuracy.

Tawsifur Rahman et al [4] use 4 pre-trained deep CNN models, AlexNet, ResNet18, DenseNet201, and SqueezeNet for transfer learning. The classification accuracy of normal and pneumonia images, bacterial and viral pneumonia images, and normal, bacterial, and viral pneumonia were 98%, 95%, and 93.3%, respectively.

Rachna Jain et al [5] compare six models and the validation accuracy of each of these models respectively and Aaryan Chothani et al [6]have implemented the use of a CNN system that detects pneumonia based on pre-processed dataset.

Jovito Colin et al [7] implements the use of deep learning models like ResNet50 and compares four interpretability techniques such as Layer-wise Relevance Propagation (LRP), Adversarial Training, Class Activation Maps (CAMs), and the Spatial Attention Mechanism, and determines which fits best the model, enhancing its transparency with minimal impact on its performance.

Jianpeng Zhang et al[8] have proposed to use the confidence-aware anomaly detection (CAAD) model. The proposed model outperforms binary classification models on the clinical X-VIRAL dataset that contains 5,977 viral pneumonia (no COVID-19) cases, 37,393 non-viral pneumonia or healthy cases, and the model achieved an AUC of 83.61% and sensitivity of 71.70%, which is comparable to the performance of radiologists reported in the literature.

Vikash Chouhan et al [9] have implemented a process to simplify pneumonia detection process for experts as well as for novices, where they approach a novel deep learning framework for the detecting pneumonia with transfer learning. The proposed model reached an accuracy of 96.4% with a recall of 99.62% on unseen data from the Guangzhou Women and Children's Medical Center dataset.

Archit Khatri et al[10] use the concept of Earth Mover's Distance or EMD to identify pneumonia lungs from normal lungs and Rohit Kundu et al [11] proposed an approach where they employed deep transfer learning and the model demonstrated accuracy rates of 98.81% and 86.85% along with sensitivity rates of 98.8% and 87.02% using datasets from Kermany et al and the Radiological Society of North America (RSNA).

Pranav Rajpurkar et al [12] develop an algorithm named CheXNet, a 121 layer convolutional neural network trained on ChestX-ray14 which contains over 100,000 frontal-view X-ray Images. CheXNet performed better than radiologists in detecting atelectasis, with an AUC of 0.862 (95% CI 0.825-0.895), statistically significantly higher than radiologists' AUC of 0.808 (95% CI 0.777-0.838).

Harshvardhan GM et al [13]focus on putting the performances of different CNN models and selecting the best based on minimum loss and maximum accuracy.

Dejun Zhang et al[14] demonstrate the use of a VGG-based model architecture with less layers compared to other models. This gave an accuracy of 96.068%..

Tej Bahadur Chandra et al[15] use a dataset of 412 chest X-ray images containing an equal half of pneumonia and non pneumonia cases from the ChestX-ray14 dataset. The results gave a high accuracy of 95.63%.

Benjamin Antin et al [16] use a dataset from Kaggle containing 112,120 chest X-ray images to plot graphs for pneumonia detection.

Danilo Avola et al [17] performed experiments using 6330 images between train, validation, and test sets, reaching up to 84.46% average f1-score.

Sadman Sadik Khan et al [18] reviews the advances made in medical imaging so far which can be beneficial for detecting pneumonia using chest x-rays.

Rishav Pramanik et al[19]use ResNet50 on a dataset, modifying using a memory-based adaptation parameter and enriched by incorporating an altruistic behavior into the agents.

Mriganka Nath et al [20]use VGG16 model over the dense network to classify images to detect pneumonia x-rays

III. IMPLEMENTATION

Workflow Architecture: Pneumonia Detection System

The proposed system follows a structured pipeline for automated pneumonia detection using deep learning, comprising four main stages: data processing, model training, performance evaluation, and deployment through a web interface.

A. Data Processing Stage

The dataset consists of labeled chest X-ray images categorized as Normal or Pneumonia. Images from train, validation, and test directories were resized to 224×224 pixels to match the DenseNet121 input size. To improve generalization and prevent overfitting, augmentation techniques including horizontal flipping, rotation, zooming, and contrast/brightness adjustments were applied.

To handle class imbalance, class weights were dynamically calculated based on the label distribution. Data pipelines were optimized using TensorFlow's tf.data.AUTOTUNE to enable efficient prefetching and training.

B. Model Training Stage

The backbone model is DenseNet121, pre-trained on ImageNet. Training was executed in two phases:

- Phase 1: Train only the classification head while keeping the DenseNet base frozen.
- Phase 2: Fine-tune the full network with a reduced learning rate.

The model used binary cross-entropy as the loss function during training, and focal loss for better class imbalance handling in deployment scenarios. The Adam optimizer with learning rate scheduling was employed. Model performance was monitored using accuracy, AUC, and F1-score, and an early stopping algorithm was applied so as to avoid overfitting.

C. Evaluation and Threshold Optimization

Evaluation on the test set involved computation and analysis of the confusion matrix, precision, recall, and the F1-score. By analyzing ROC and the precision-recall curves, the optimal threshold (0.68) was selected to maximize the F1-score. The confusion matrix visualizes the pneumonia model's performance during training where 193 cases were true normal and 373 cases were true pneumonia. F1-score of 0.9279 demonstrates it's high sensitivity in healthcare systems.

D. Deployment via Streamlit Web App

A lightweight frontend was developed using **Streamlit**, enabling users to upload images, preprocess and adjust **image enhancements**, and view histograms.

The backend handled preprocessing, model inference, and displays predictions of the **probability and confidence scores.** The system integrates technical rigor and usability, alowing for real-time inference and user-friendly diagnostics in a low-resource setting.

IV. RESULTS AND DISCUSSION

The developed prototype constitutes a backend-optimized deep learning model for automated pneumonia detection from chest X-ray images. It is built upon a fine-tuned DenseNet121 architecture, trained using the TensorFlow and Keras frameworks. Designed for binary classification (Normal vs. Pneumonia), the system was validated on a test set of 624 chest X-ray samples to assess generalizability. Fig 1 represents the precision recall curve of the model.

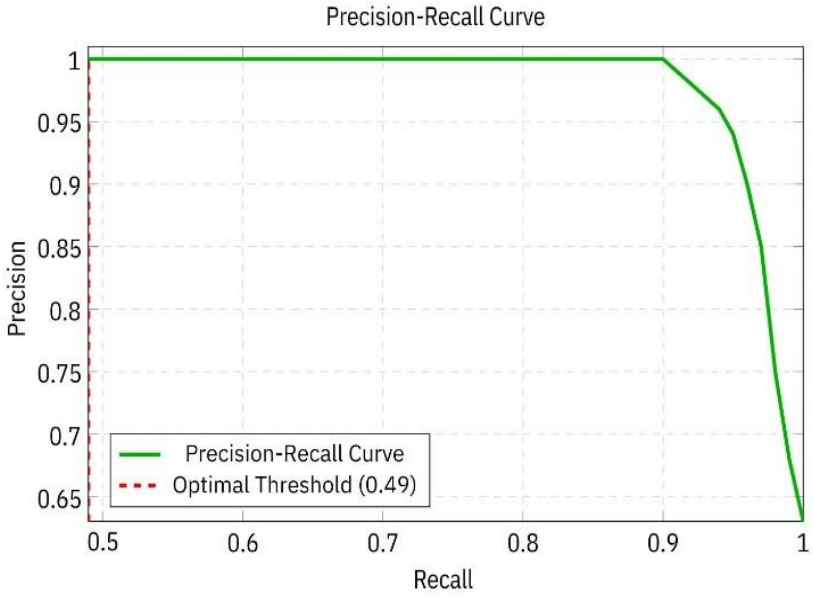


Fig 1. Precision-Recall Curve

The training strategy was designed for computational efficiency and diagnostic performance. Automatic data loading and tf.data prefetching were employed for high GPU utilization. Image augmentation techniques such as horizontal flipping, zooming, and brightness/contrast modulation were deployed for the diversity of training data and thereby improving generalization. During computation of loss, dynamic class weighting was used in order to mitigate class imbalance. Fig 2 shows the ROC curve.

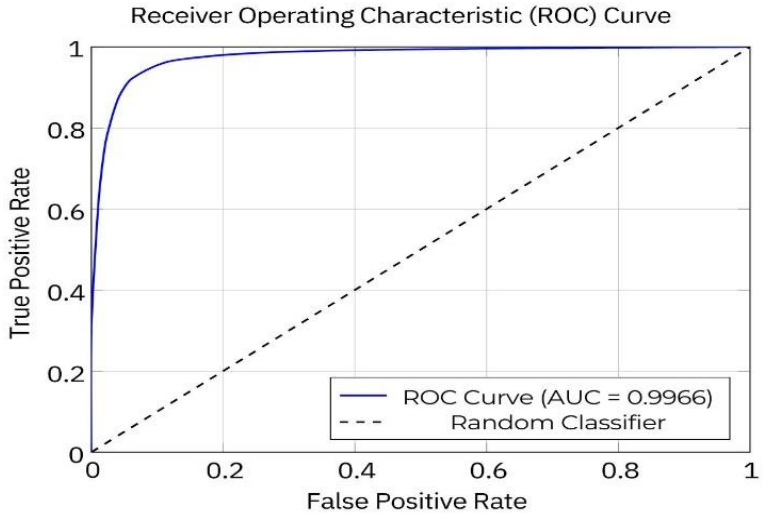


Fig2. ROC Curve

Training was conducted in two phases. During phase 1, only the classification head was trained while the DenseNet base remained frozen. Thereafter, the entire model was fine-tuned. It employed the use of binary cross-entropy as the loss function and used the Adam optimizer with scheduled learning rate decay. Early stopping was integrated so as to prevent overfitting. Fig 3 gives the accuracy and losses obtained during model training.

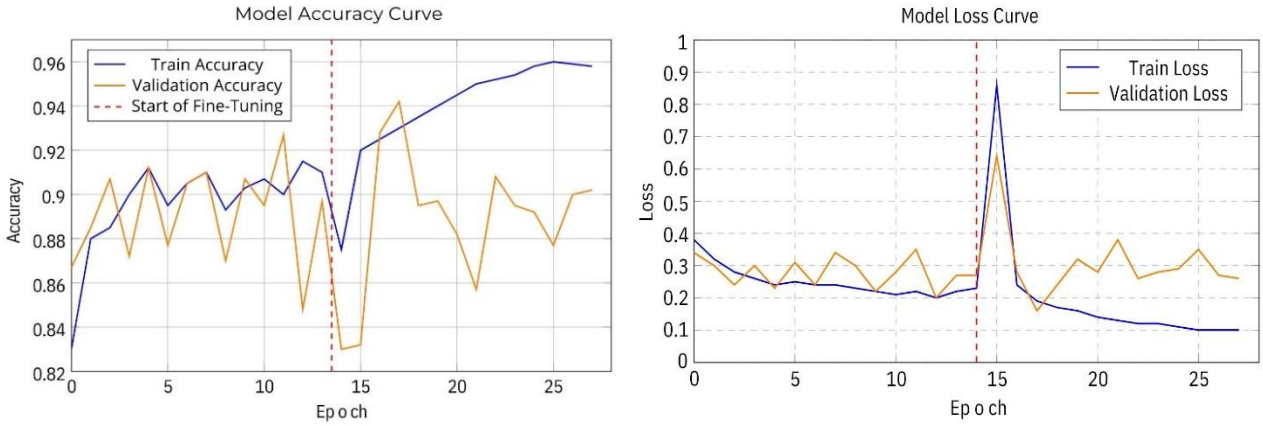


Fig 3. Model loss and Model Accuracy curve

After training, threshold optimization was performed using F1-score maximization. The final model achieved **91% test accuracy,** an **AUC of 0.9511,** and an optimized **F1-score of 0.9279.** This demonstrated high discriminative performance and resilience under real-world conditions. Fig 4 shows the confusion matrix which demonstrates the model has some mismatches, although the overall score overweighs the loss.

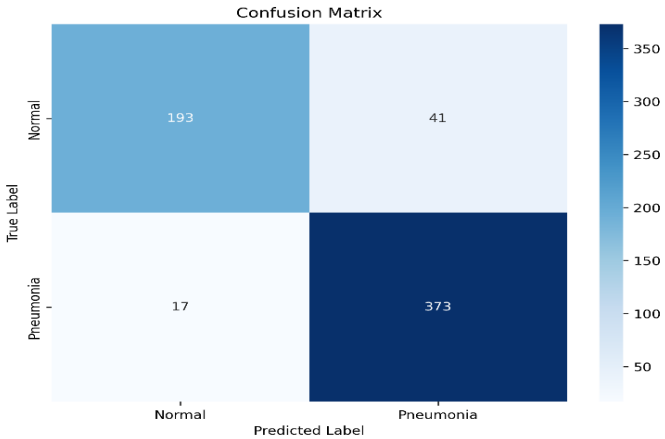


Fig 4. Confusion Matrix

The proposed model exhibits efficient inference performance, with an average prediction time of approximately **313 milliseconds per step**. Its compact deployment in .keras format ensures seamless integration into cross-platform diagnostic tools. This includes applications that are both browser and desktop based. The model is ready to be deployed using frameworks such as **Streamlit**. In this work, we have not created the graphical interface.

Model output consists of binary classification (Normal or Pneumonia), a probability score, and a qualitative interpretation message derived from the predicted class and threshold calibration. For future user-facing deployments, visual aids—including training/validation curves, ROC plots, and confusion matrices—are recommended to enhance interpretability.

While interpretability tools such as Grad-CAM are not currently integrated, transparency is supported through detailed performance metrics and evaluation summaries, ensuring clinical relevance and potential for real-world adoption.

The pixel histogram of the image is based on an 8-bit grayscale system which provides information about the brightness levels in an image from a range of 0 to 255, 0 being black and 255 being white, and in between values are shades of gray.

Our model implements the use of DenseNet121 which has potential benefits such as encouraging feature reuse and better performance [4].

Fig 6 and Fig 7 shows testing of the model for 2 different images and their corresponding histograms were displayed for better enhancement of the images. The **Streamlit** GUI allows manual enhancement along with auto-enhancement of the images. Fig 5 shows the result of the first image the model generated , which works for the 0.68 probability threshold. It displays the confidence in prediction and probability of pneumonia.

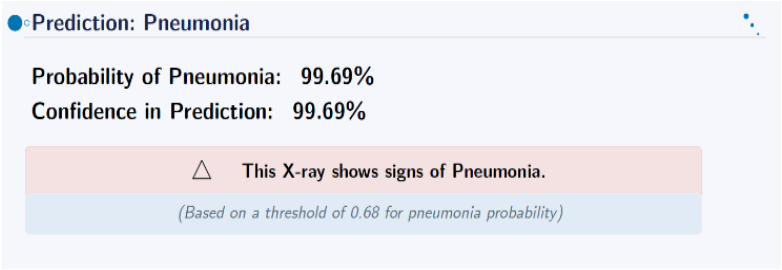


Fig 5. Results of prototype taking first image

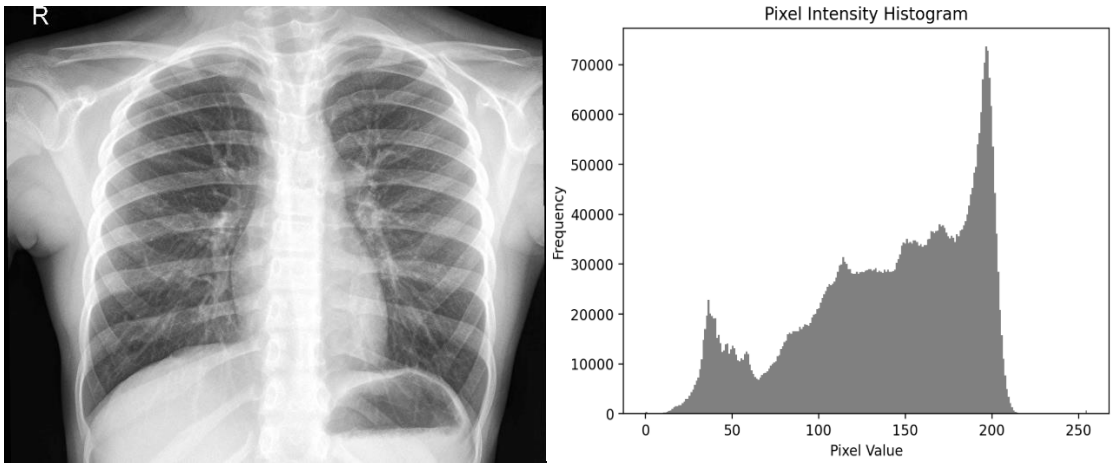


Fig 6. Image and its Pixel Intensity Histogram

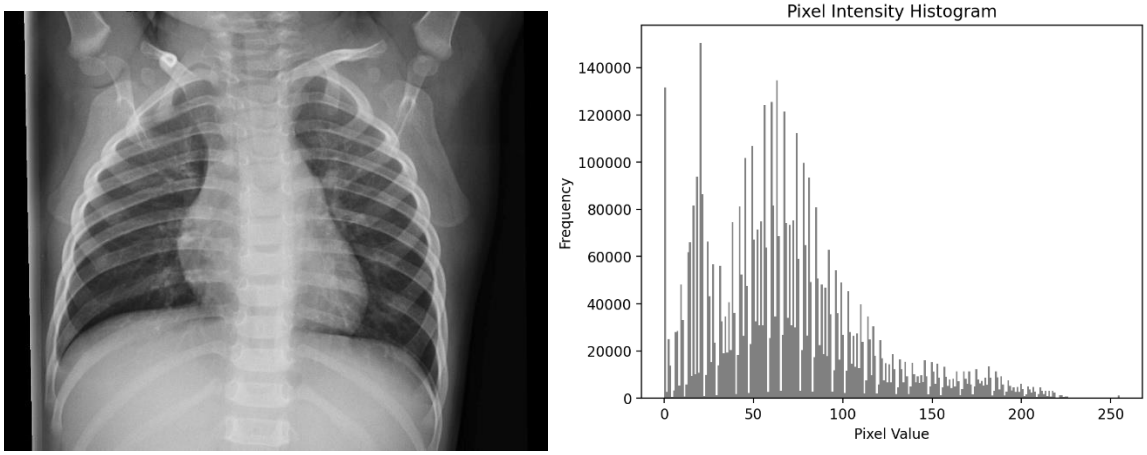


Fig 7. Pixel Intensity Histogram and the Image

V. CONCLUSION

As per the prototype's validation, our model demonstrates the feasibility of deploying a robust and computationally efficient diagnostic tool that can integrate with healthcare infrastructures. The model is shown to achieve a diagnostic accuracy of 91%, thereby confirming its high discriminative performance and resilience under conditions reflective of real-world clinical variability. This model also achieved an AUC of 0.951 and an optimized F1-score of 0.9279. This demonstrates high resilience under actual test conditions. These attributes, such as high reliability paired with operational efficiency makes this system a compelling candidate for integration into clinical decision support systems and telemedicine platforms, where it can assist in improving diagnostic accuracy.

Future work will focus on three development areas to advance this technology for achieving positive impact in healthcare environments. The first will be the development of a clinician-facing interface that is simple and intuitive. This is essential to facilitate seamless workflow integration and user adoption. The second will be to expand the model's capability from binary to multi-class detection for a broader spectrum of pulmonary pathologies. This will increase its clinical utility. Finally, optimization for real-time inference on edge computing devices, such as point-of-care ultrasound machines or mobile platforms, will be pursued. This will enable deployment in resource-constrained and remote environments, thus supporting accurate delivery of diagnosis in decentralized healthcare environments.

In conclusion, this work demonstrates a validated foundation for an AI-assisted diagnostic tool that balances accuracy with practicality. It also offers a tool that will enhance diagnostic precision and scalability in respiratory medicine.

Acknowledgment

We would like to thank Prof P Narashimaraja for his help and guidance in our work. We acknowledge the support by the Department of Electronics, RV College of Engineering for providing the facilities required to initiate and complete this study.

REFERENCES

1. Sharma S. and Guleria K . A systematic literature review on deep learning approaches for pneumonia detection using chest X-ray images. Multimedia Tools and Applications. 2023;83(8): 24101-24151 Available: https://link.springer.com/article/10.1007/s11042-023-16419-1
2. Khan W., Zaki N. and Ali L. Intelligent Pneumonia Identification From Chest X-Rays: A Systematic Literature Review. IEEE Access. 2021. Available: https://doi.org/10.1101/2020.07.09.20150342
3. Shilpa N., Ayeesha Banu W., and Metre P. B. Revolutionizing Pneumonia Diagnosis: AI-Driven Deep Learning Framework for Automated Detection From Chest X-Rays. IEEE Access. 2024;12:171601-171616, Available: https://ieeexplore.ieee.org/document/10753606
4. Rahman T., Chowdhury M. E. H., Khandakar A., Islam K. R., Islam K. F., Mahbub Z. B., Kadir M. A., and Kashem S. Transfer Learning with Deep Convolutional Neural Network (CNN) for Pneumonia Detection Using Chest X-ray. Applied Sciences, 2020;10(9):3233. Available: https://doi.org/10.3390/app10093233
5. Jain R., Nagrath P., Kataria G., Kaushik V. S., and Hemanth D. J. Pneumonia detection in chest X-ray images using convolutional neural networks and transfer learning. Measurement. 2020;165(5):108046. Available: https://www.sciencedirect.com/science/article/abs/pii/S0263224120305844
6. Chothani A., Khemani B., Malave S., Jain M., and Gupta J . Pneumonia Detection in Chest X-rays using CNN. 2024 3rd International Conference for Innovation in Technology (INOCON). Available: DOI:10.1109/INOCON60754.2024.10512113
7. Colin J. and Surantha N. Interpretable Deep Learning for Pneumonia Detection Using Chest X-Ray Images. Information**.** 2025;16(1):53. Available: https://doi.org/10.3390/info16010053
8. Zhang J., Xie Y., Pang G., Liao Z., Verjans J., Li W., Sun Z., He J., Li Y., Shen C., and Xia Y. Viral Pneumonia Screening on Chest X-Rays Using Confidence-Aware Anomaly Detection. IEEE Transactions on Medical Imaging. 2021;40(3): 879-890. Available: https://arxiv.org/pdf/2003.12338v2
9. Chouhan V., Singh S. K., Khamparia A., Gupta, Tiwari P., Moreira C., Damaševičius R. and de Albuquerque. V. H. C. A Novel Transfer Learning Based Approach for Pneumonia Detection in Chest X-Ray Images. Applied Sciences.2020;10(2): 559. Available: https://doi.org/10.3390/app10020559
10. Khatri A., Jain R., Vashista H., and Mittal N. Pneumonia Identification in Chest X-ray Images Using EMD. In book: Trends in Communication, Cloud, and Big Data. 2020: 87-98. Available: https://link.springer.com/chapter/10.1007/978-981-15-1624-5_9
11. Kundu R., Das R., Geem Z. W., Han G.-T., and Sarkar R. Pneumonia detection in chest X-ray images ensemble of deep learning models. PLOS ONE. 2021;16(9): e0256630,. Available: https://doi.org/10.1371/journal.pone.0256630
12. Rajpurkar P., Irvin J., Zhu K., Yang B., Mehta H., Duan T., Ding D., Bagul A., Langlotz C., Shpanskaya K., Lungren M. P., and Ng A. Y. CheXNet: Radiologist-Level Pneumonia Detection on Chest X-rays with Deep Learning. *arXiv preprint arXiv:1711.05225*, 2017. Available: https://arxiv.org/abs/1711.05225
13. Harshvardhan G. M, Gourisaria M. K, Rautaray S. S, and Pandey M. Pneumonia Detection Using CNN Through Chest X-Ray. Journal of Engineering Science and Technology. 2021;16(1):320-332. Available: http://jestec.taylors.edu.my/Vol%2016%20issue%201%20February%202021/16_1_25.pdf
14. Zhang D., Ren F., Li Y., Na L., and Ma Y. Pneumonia Detection from Chest X-ray Images Based on Convolutional Neural Network. Electronics. 2021;10(13): 1512. Available: https://doi.org/10.3390/electronics10131512
15. Chandra T. B. and Verma K.. Pneumonia Detection on Chest X-Ray Using Machine Learning Paradigm. Conference: Proceedings of 3rd international conference on computer vision and image processing. 2019. Available: https://link.springer.com/chapter/10.1007/978-981-32-9088-4_3
16. Antin B., Kravitz J. and Martayan E. Detecting Pneumonia in Chest X-Rays with Supervised Learning. Stanford University *CS229 Course Project Paper*, 2017. Available: cs229.stanford.edu/proj2017/final-reports/5241087.pdf
17. Avola D., Bacciu A., Cinque L., Fagioli A., Marini M. R. and Taiello R. Study on transfer learning capabilities for pneumonia classification in chest x-rays images. Computer Methods and Programs in Biomedicine. 2022;221: 106833. Available: https://www.sciencedirect.com/science/article/pii/S0169260722002152?via%3Dihub

18. Khan S. S., Rupak A. U. H., Faieaz W. W., Jannat S., Prova N. N. I. and Gupta A. K. Advances in Medical Imaging: Deep Learning Strategies for Pneumonia Identification in Chest X-rays. Conference: 15th International Conference on Computing Communication and Networking Technologies (ICCCNT). 2024. Available: DOI:10.1109/ICCCNT61001.2024.10725661
19. Pramanik R., Sarkar S. and Sarkar R. An adaptive and altruistic PSO-based deep feature selection method for Pneumonia detection from Chest X-rays. Applied Soft Computing. 2022;128:10946. Available: DOI:10.1016/j.asoc.2022.109464
20. Nath M. and Choudhury C. Automatic detection of Pneumonia from Chest X-rays using Deep learning. International Conference on Machine Learning, Image Processing, Network Security and Data Sciences(MIND 2022): 175-182. Available: https://link.springer.com/chapter/10.1007/978-981-15-6315-7_14