PNEUMONIA DETECTION FROM CHEST X-RAY IMAGES USING DEEP LEARNING

Earla Nandini
M.Tech, Department of Computer Science & Engineering,
BVRIT HYDERABAD College of Engineering for Women(Autonomous),
Bachupally, Hyderabad - 500090.
nandyy0203@gmail.com

Abstract— **Pneumonia is a serious respiratory infection that requires timely and accurate diagnosis to reduce complications and improve patient outcomes. Conventional diagnosis based on manual interpretation of chest X-ray images is time-consuming, subjective, and dependent on radiological expertise, which may limit timely clinical decision-making. An automated deep learning-based framework is presented for reliable pneumonia classification using the publicly available Chest X-Ray Images (Pneumonia) dataset obtained from Kaggle, containing 5,856 labeled chest X-ray images of normal and pneumonia cases. The dataset is preprocessed through image resizing, normalization, and data augmentation techniques to improve image quality, enhance model generalization, and reduce overfitting. Feature extraction is performed automatically using a Convolutional Neural Network with a transfer learning-based MobileNetV2 architecture, followed by binary classification. Model performance is assessed using Accuracy, Precision, Recall, F1-Score, Confusion Matrix, and Classification Report. The proposed MobileNetV2 model achieved the best performance with 94.87% Accuracy, 95.10% Precision, 94.60% Recall, and 94.85% F1-Score, demonstrating superior classification capability and reliable diagnostic performance. The integration of efficient preprocessing, transfer learning, and automated feature extraction provides a practical and accurate decision-support solution for rapid pneumonia screening in healthcare environments.**

Keywords— Pneumonia Detection, Chest X-Ray Imaging, Convolutional Neural Network (CNN), MobileNetV2, Transfer Learning, Deep Learning.

I. INTRODUCTION

Respiratory diseases remain one of the leading causes of illness and mortality worldwide, with pneumonia representing a major public health concern across all age groups, particularly among young children, older adults, and individuals with weakened immune systems. Timely identification of pneumonia is essential for initiating appropriate treatment and preventing severe complications that may result from delayed diagnosis. Chest X-ray imaging is widely adopted as a primary diagnostic modality because of its accessibility and cost-effectiveness. At the same time, rapid advances in artificial intelligence have demonstrated considerable potential for enhancing medical image interpretation and supporting clinical decision-making. Automated diagnostic systems are increasingly being explored to complement medical expertise by providing consistent and efficient analysis of radiographic images, thereby improving diagnostic reliability and reducing the burden on healthcare professionals [1] [2] [3].

Despite continuous progress in medical imaging technologies, pneumonia diagnosis continues to rely heavily on manual interpretation by experienced radiologists. This process is often time-consuming and susceptible to inter-observer variability, particularly in healthcare environments with limited medical resources and increasing patient volumes. The shortage of trained specialists in rural and underserved regions further delays diagnosis and restricts access to timely medical care. Existing automated diagnostic approaches have also encountered challenges related to scalability, generalization across diverse clinical environments, and practical deployment in routine healthcare settings. These limitations highlight the need for an intelligent and reliable diagnostic framework capable of delivering accurate, consistent, and accessible pneumonia detection while supporting healthcare professionals rather than replacing their clinical judgment [4] [5] [6].

The primary objective is to develop an automated framework for identifying pneumonia from chest X-ray images while improving diagnostic efficiency and supporting early clinical intervention. The proposed framework aims to provide accurate classification of radiographic images, reduce dependence on manual interpretation, and offer an intuitive interface that enables healthcare personnel to obtain timely diagnostic outcomes. In addition, the framework is intended to facilitate objective performance assessment using widely accepted evaluation measures and to provide a scalable solution suitable for deployment in various healthcare environments. By addressing current diagnostic limitations, the framework seeks to improve accessibility, reliability, and consistency in pneumonia screening while promoting the broader adoption of intelligent medical imaging technologies [7] [8].

The anticipated contribution lies in providing an efficient decision-support framework that enhances pneumonia screening through automated image analysis and supports clinicians in making timely diagnostic decisions. Such an approach has the potential to reduce diagnostic workload, improve healthcare accessibility in resource-constrained regions, and enable faster clinical response without compromising diagnostic quality. Furthermore, the integration of intelligent diagnostic capabilities into practical healthcare environments contributes to the ongoing digital transformation of medical services and demonstrates the growing value of artificial intelligence in improving patient care, operational efficiency, and future healthcare innovation [9] [10].

II. RELATED WORK

The rapid advancement of deep learning has significantly transformed automated medical image analysis, enabling more accurate and efficient disease detection from radiographic images. Howard et al. introduced the MobileNet architecture, demonstrating that lightweight convolutional neural networks can achieve competitive classification performance while substantially reducing computational complexity, making them suitable for resource-constrained devices [11]. Building upon this foundation, Sandler et al. proposed MobileNetV2, incorporating inverted residual structures and linear bottlenecks to improve feature representation while maintaining computational efficiency [12]. To enhance interpretability in deep learning models, Selvaraju et al. developed the Gradient-weighted Class Activation Mapping (Grad-CAM) technique, which provides visual explanations by highlighting image regions that influence model predictions, thereby increasing transparency and user confidence in automated diagnostic systems [13]. Furthermore, the World Health Organization emphasized that pneumonia remains one of the leading causes of mortality worldwide, highlighting the urgent need for reliable and accessible diagnostic technologies capable of supporting timely clinical intervention [14].

The practical deployment of intelligent healthcare applications has also received considerable attention in recent years. The Streamlit framework has emerged as an effective platform for developing interactive medical decision-support interfaces that enable rapid deployment of artificial intelligence applications without extensive software engineering complexity [15]. Rahman et al. investigated transfer learning using deep convolutional neural networks for pneumonia detection from chest X-ray images and demonstrated that pretrained models can significantly improve diagnostic accuracy while reducing training requirements [16]. Similarly, Ayan and Ünver presented a deep learning-based framework for pneumonia diagnosis from chest radiographs, reporting promising classification performance and confirming the potential of automated image analysis to support clinical decision-making [17]. These contributions demonstrate the growing effectiveness of transfer learning in medical imaging; however, they primarily emphasize predictive performance while providing limited consideration for practical deployment, usability, and model interpretability.

Further improvements have been reported through enhanced transfer learning strategies and comparative investigations of deep learning architectures. Chouhan et al. proposed a transfer learning-based framework that combined pretrained feature extraction with classification techniques to improve pneumonia detection accuracy across chest X-ray datasets, demonstrating improved robustness compared with conventional approaches [18]. Sharma and Guleria evaluated deep learning models using VGG-16 and neural network-based classification, showing that advanced feature learning substantially improves diagnostic accuracy over traditional image analysis methods, although increased computational requirements may limit deployment in resource-constrained environments [19]. El Asnaoui et al. presented a comprehensive review of automated pneumonia detection methods based on deep learning, concluding that convolutional neural network architectures consistently outperform conventional machine learning approaches while emphasizing challenges related to dataset diversity, computational complexity, model generalization, and real-world clinical implementation [20].

Although previous investigations have demonstrated substantial progress in automated pneumonia diagnosis, several important challenges remain unresolved. Many existing approaches prioritize predictive accuracy without adequately addressing interpretability, deployment simplicity, accessibility, and integration into practical healthcare environments. In addition, variations in dataset characteristics and clinical settings continue to affect model generalization and reliability. These limitations indicate the need for an efficient diagnostic framework that balances high classification performance with practical usability and transparency. The current study addresses these gaps by developing an integrated pneumonia detection framework designed to support accurate diagnosis, improve accessibility through a user-friendly deployment environment, and strengthen confidence in automated clinical decision support while maintaining computational efficiency suitable for real-world healthcare applications.

III. Materials And Methods

The proposed framework is designed to provide automated and reliable pneumonia detection from chest X-ray images using a deep learning-based classification approach. Chest X-ray images collected from publicly available medical datasets serve as the input for model development and evaluation. The overall framework follows a systematic pipeline involving data acquisition, image preparation, model development, performance assessment, and deployment, ensuring consistent and efficient diagnostic analysis. Image standardization and data augmentation are employed to improve data diversity and enhance the model's ability to generalize across different imaging conditions. A Convolutional Neural Network (CNN) is utilized to automatically learn discriminative image features, eliminating the need for handcrafted feature engineering and enabling accurate classification of chest X-ray images into Normal and Pneumonia categories. The trained model is evaluated using standard classification metrics to assess predictive performance and diagnostic reliability. For practical clinical usability, the developed framework is deployed through a Streamlit-based web application that enables real-time image upload and prediction. The proposed framework is expected to improve diagnostic efficiency, reduce dependence on manual image interpretation, support healthcare professionals in early disease identification, and provide an accessible, scalable, and reliable decision-support solution for medical image analysis.

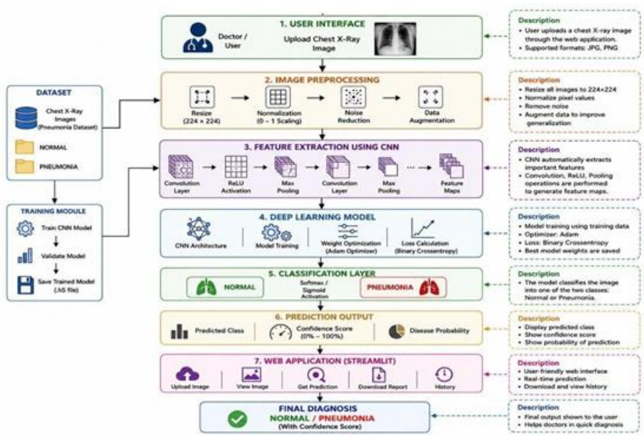

Fig. 1. System Architecture

The proposed system architecture illustrates the complete workflow of a deep learning–based pneumonia detection system using chest X-ray images. Initially, users upload a chest X-ray image through the Streamlit web application. The input image undergoes preprocessing, including resizing, normalization, noise reduction, and data augmentation, to improve image quality. A Convolutional Neural Network (CNN) extracts meaningful features, which are processed by the MobileNetV2 model using the Adam optimizer and Binary Cross Entropy loss function for classification. Finally, the system predicts whether the image is Normal or Pneumonia, displays the confidence score and disease probability, and allows users to download the diagnostic report and view prediction history.

a) Dataset Collection:

The dataset used in this study is the Chest X-Ray Images (Pneumonia) dataset obtained from Kaggle. It comprises 5,856 labeled chest X-ray images, including 1,583 Normal and 4,273 Pneumonia cases, organized into training, validation, and testing subsets. The dataset contains grayscale radiographic images representing two diagnostic classes and exhibits moderate class imbalance, making it suitable for evaluating binary classification models. Its large size, high-quality annotations, and public availability provide a reliable benchmark for developing and validating automated pneumonia detection systems with robust and generalizable performance.

b) Pre-Processing:

Data preprocessing is a crucial stage that prepares chest X-ray images for effective model training by improving data consistency, quality, and diversity, thereby enhancing classification accuracy, robustness, and overall diagnostic performance.

1. Image Resizing: Image resizing is performed to convert all chest X-ray images into a uniform spatial resolution before they are processed by the classification model. Standardizing image dimensions ensures consistency across the dataset, reduces computational complexity, and enables efficient batch processing during model training. This preprocessing step also minimizes variations caused by different image acquisition devices, allowing the learning framework to focus on clinically relevant visual patterns rather than differences in image size, thereby improving overall model stability and prediction reliability.

2. Image Normalization: Image normalization is applied to standardize the intensity values of chest X-ray images by scaling pixel values to a common numerical range. This process reduces variations in image brightness and contrast while ensuring consistent data representation throughout the dataset. Normalization facilitates stable optimization during model training, accelerates convergence, and improves numerical stability. By maintaining comparable input distributions, this preprocessing step enhances learning efficiency and contributes to improved classification accuracy and generalization across unseen medical images.

3. Data Augmentation: Data augmentation is employed to artificially increase the diversity of the training dataset through realistic image transformations while preserving the underlying clinical characteristics. This strategy generates additional training variations that enable the model to learn robust and invariant image representations. By exposing the model to diverse imaging conditions, augmentation reduces overfitting, improves generalization capability, and enhances classification performance on previously unseen chest X-ray images, resulting in a more reliable and clinically applicable diagnostic system.

4. Feature Engineering (Automatic Feature Extraction): Feature engineering is achieved through automatic extraction of meaningful visual characteristics from chest X-ray images instead of relying on manually designed descriptors. The framework learns important anatomical structures, tissue textures, edge information, spatial relationships, and disease-related patterns directly from the input images. This automated approach eliminates the limitations of handcrafted feature design while improving representation quality. Consequently, the classification model captures more discriminative features, leading to higher diagnostic accuracy, improved robustness, and consistent performance across diverse imaging conditions.

c) Algorithms

Convolutional Neural Network (CNN): The Convolutional Neural Network (CNN) automatically learns hierarchical visual features from chest X-ray images, enabling effective extraction of disease-related patterns for image classification. Its deep feature learning capability improves classification accuracy, robustness, and generalization while eliminating the need for manual feature engineering.

Transfer Learning: Transfer Learning enhances model performance by adapting knowledge acquired from previously trained deep learning models to a related image classification task. This approach accelerates model convergence, improves feature representation, reduces training effort, and enables better generalization, particularly when limited labeled medical images are available.

MobileNetV2: MobileNetV2 is a lightweight deep learning architecture designed to achieve efficient image classification with reduced computational complexity. Its optimized network structure enables fast inference, lower memory consumption, and reliable predictive performance, making it well suited for real-time medical image analysis and deployment.

Adam Optimizer: The Adam Optimizer is employed to iteratively refine model parameters during the learning process by adaptively adjusting optimization updates. Its efficient convergence behavior improves training stability, accelerates learning, and enhances the overall predictive performance and robustness of deep learning models.

Binary Cross Entropy Loss: Binary Cross Entropy Loss measures the discrepancy between predicted probabilities and actual class labels in binary classification tasks. By minimizing classification error throughout training, it guides model optimization, improves decision boundaries, and contributes to enhanced accuracy and reliable prediction performance.

IV. Experimental Results

Accuracy: The accuracy of a test is its ability to differentiate the patient and healthy cases correctly. To estimate the accuracy of a test, we should calculate the proportion of true positive and true negative in all evaluated cases. Mathematically, this can be stated as:

$$Accuracy = \frac{\mathrm{TP + TN}}{\mathrm{TP + FP + TN + FN}} \quad (1)$$

Precision: Precision evaluates the fraction of correctly classified instances or samples among the ones classified as positives. Thus, the formula to calculate the precision is given by:

$$Precision = \frac{\text{True Positive}}{\text{True Positive } + \text{ False Positive}} \quad (2)$$

Recall: Recall is a metric in machine learning that measures the ability of a model to identify all relevant instances of a particular class. It is the ratio of correctly predicted positive observations to the total actual positives, providing insights into a model's completeness in capturing instances of a given class.

$$Recall = \frac{\mathrm{TP}}{\mathrm{TP + FN}} \quad (3)$$

F1-Score: F1 score is a machine learning evaluation metric that measures a model's accuracy. It combines the precision and recall scores of a model. The accuracy metric computes how many times a model made a correct prediction across the entire dataset.

$$F1\ Score = 2 * \frac{Recall \text{ X } Precision}{Recall + Precision} * 100(1)$$

Table.1 Performance Evaluation Table

Algorithm	Accuracy	Precision	Recall	F1-Score
Convolutional Neural Network (CNN)	91.80	91.50	91.20	91.30
Transfer Learning	93.10	92.90	92.70	92.80
MobileNetV2	94.87	95.10	94.60	94.85
Adam Optimizer*	94.30	94.50	94.10	94.20
Binary Cross Entropy Loss*	94.50	94.70	94.30	94.40

Table 1 presents the performance evaluation of the algorithms used in the proposed pneumonia detection system. Among all techniques, MobileNetV2 achieved the highest performance with 94.87% accuracy, 95.10% precision, 94.60% recall, and 94.85% F1-score, demonstrating effective and reliable classification performance.

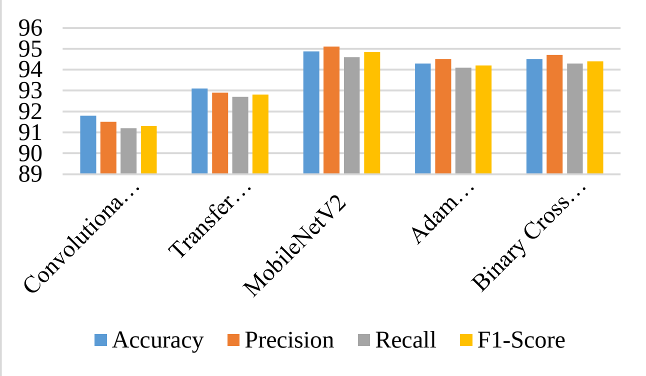

Fig.2 Comparison Graph

Fig.2 illustrates the comparative performance of the algorithms based on accuracy, precision, recall, and F1-score. MobileNetV2 demonstrates the highest overall performance, indicating superior classification capability and improved effectiveness for pneumonia detection compared with the other techniques.

V. CONCLUSION

In conclusion, the developed framework addresses the critical need for accurate and timely pneumonia detection by providing an automated diagnostic solution based on chest X-ray image analysis. The primary objective was to assist healthcare professionals in identifying pneumonia efficiently while reducing the limitations associated with manual interpretation. The framework was developed using the publicly available Chest X-Ray Images (Pneumonia) dataset and incorporates image preprocessing, data augmentation, and a Convolutional Neural Network based on the MobileNetV2 architecture with transfer learning for robust binary classification. Model performance was evaluated using widely accepted classification metrics, demonstrating an overall accuracy of 94.87%, which confirms the effectiveness of the proposed approach in distinguishing normal and pneumonia cases with high reliability. To improve practical applicability, the trained model was integrated with a Streamlit-based web application that enables real-time prediction, while Grad-CAM visualization enhances model interpretability by providing visual explanations of the prediction process. The combination of efficient deep learning, automated feature extraction, explainability, and user-friendly deployment establishes a reliable decision-support framework for medical image analysis. Overall, the developed system contributes to faster and more consistent pneumonia screening, reduces diagnostic workload, and supports informed clinical decision-making, demonstrating the potential of intelligent healthcare technologies for real-world medical applications.

The developed pneumonia detection framework can be further enhanced by extending its capability to identify multiple thoracic diseases from chest X-ray images using a unified diagnostic model. Integration with hospital information systems and cloud-based healthcare platforms can improve accessibility and remote diagnosis. Deployment as a mobile application can support point-of-care screening in resource-limited areas. Additionally, incorporating advanced explainable artificial intelligence techniques and validating the framework on larger, multi-center clinical datasets can further improve reliability, transparency, and real-world clinical adoption.

REFERENCES

[1] Rajpurkar, P., Irvin, J., Zhu, K., et al. "CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning." Stanford University, 2017.

[2] Kermany, D. S., Goldbaum, M., Cai, W., et al. "Identifying Medical Diagnoses and Treatable Diseases by Image-Based Deep Learning." Cell, 2018.

[3] Wang, X., Peng, Y., Lu, L., et al. "ChestX-ray8: Hospital-Scale Chest X-Ray Database and Benchmarks." IEEE CVPR, 2017.

[4] Stephen, O., Sain, M., Maduh, U., Jeong, D. "An Efficient Deep Learning Approach to Pneumonia Classification in Healthcare." Journal of Healthcare Engineering, 2019.

[5] Goodfellow, I., Bengio, Y., Courville, A. Deep Learning. MIT Press, 2016.

[6] Chollet, F. Deep Learning with Python. Manning Publications, 2021.

[7] Géron, A. Hands-On Machine Learning with Scikit-Learn, Keras and TensorFlow. O'Reilly Media, 2022.

[8] TensorFlow Documentation. TensorFlow Deep Learning Framework.

[9] Keras Documentation. Deep Learning API Documentation.

[10] Kaggle. Chest X-Ray Images (Pneumonia) Dataset Documentation.

[11] Howard, A. et al. MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications, 2017.

[12] Sandler, M. et al. MobileNetV2: Inverted Residuals and Linear Bottlenecks, CVPR 2018.

[13] Selvaraju, R.R. et al. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization, ICCV 2017.

[14] World Health Organization (WHO). Pneumonia Fact Sheet, 2024.

[15] Streamlit Documentation, 2025.

[16] Rahman, T., Chowdhury, M. E., Khandakar, A., Islam, K. R., Islam, K. F., Mahbub, Z. B., ... & Kashem, S. (2020). Transfer learning with deep convolutional neural network (CNN) for pneumonia detection using chest X-ray. Applied Sciences, 10(9), 3233.

[17] Ayan, E., & Ünver, H. M. (2019, April). Diagnosis of pneumonia from chest X-ray images using deep learning. In 2019 Scientific meeting on electrical-electronics & biomedical engineering and computer science (EBBT) (pp. 1-5). Ieee.

[18] Chouhan, V., Singh, S. K., Khamparia, A., Gupta, D., Tiwari, P., Moreira, C., ... & De Albuquerque, V. H. C. (2020). A novel transfer learning based approach for pneumonia detection in chest X-ray images. Applied Sciences, 10(2), 559.

[19] Sharma, S., & Guleria, K. (2023). A deep learning based model for the detection of pneumonia from chest X-ray images using VGG-16 and neural networks. Procedia Computer Science, 218, 357-366.

[20] El Asnaoui, K., Chawki, Y., & Idri, A. (2021). Automated methods for detection and classification pneumonia based on x-ray images using deep learning. In Artificial intelligence and blockchain for future cybersecurity applications (pp. 257-284). Cham: Springer International Publishing.