

Automated Pneumonia Diagnosis from Chest X-Ray Images Using CNN

¹Dr T. Seshu Chakravarthy, ²Gadiyam Shanmukha Yaswanth Reddy, ³Gamidi Vamsi Krishna, ⁴Chinnameerapalli Ravi Teja, ⁵Ankipalli Anil Kumar

¹Associate Professor in Department of CSE Vasireddy Venkatadri Institute of Technology.

^{2,3,4,5}UG Scholars in Department of CSE Vasireddy Venkatadri Institute of Technology.

Abstract

Pneumonia is a serious respiratory infection that affects millions of people worldwide and remains a leading cause of mortality, particularly among children and the elderly. Early and accurate diagnosis is essential for effective treatment, yet traditional chest X-ray interpretation requires skilled radiologists and can be time-consuming, especially in resource-limited healthcare settings. To address this challenge, this project presents Pnemo Detect, an automated pneumonia detection system based on deep learning techniques. The system utilizes chest X-ray images as input and employs a Convolutional Neural Network (CNN) to classify images as either normal or pneumonia-infected, automatically extracting relevant features and reducing human error. Developed as a web-based application using Django, the platform enables healthcare professionals to upload images and receive real-time predictions through a user-friendly interface. To enhance transparency and clinical trust, the system incorporates explainable AI techniques such as Grad-CAM to highlight image regions influencing the model's decisions. The trained model and application are deployed on Amazon Web Services (AWS), ensuring scalability, reliability, and high availability. Extensive evaluation demonstrates that Pnemo Detect provides accurate and consistent predictions, serving as a supportive diagnostic tool to assist medical professionals in faster and more efficient decision-making rather than replacing them.

Keywords: Pneumonia Detection, Chest X-ray, Deep Learning, Convolutional Neural Network (CNN), Medical Image Classification, Explainable AI, Grad-CAM, Web-Based Application, Django Framework, Cloud Deployment, Amazon Web Services (AWS), Healthcare Decision Support System.

I INTRODUCTION

Pneumonia is a severe respiratory infection that affects the lungs and disrupts normal breathing by causing inflammation in the air sacs, which may fill with fluid or pus. This condition reduces

oxygen exchange in the body and can result in symptoms such as cough, fever, chest pain, fatigue, and shortness of breath. Pneumonia can be caused by bacteria, viruses, or fungi and affects individuals of all ages; however, young children, elderly individuals, and people with

weakened immune systems are particularly vulnerable. Despite advancements in healthcare and medical technology, pneumonia remains one of the leading causes of morbidity and mortality worldwide, especially in developing countries where access to quality healthcare services may be limited. Therefore, early detection and accurate diagnosis are crucial to ensure timely treatment and improve patient outcomes.

Chest X-ray imaging is one of the most commonly used diagnostic tools for identifying pneumonia. It allows healthcare professionals to observe abnormalities in lung structures, such as opacities and consolidations that indicate infection. However, interpreting chest X-ray images requires specialized expertise and careful examination by trained radiologists. Manual diagnosis can be time-consuming and may sometimes be influenced by factors such as heavy workload, fatigue, and subjective judgment. With the increasing number of patients in hospitals, there is a growing need for automated systems that can assist medical professionals in making faster and more reliable diagnostic decisions.

The rapid advancement of Artificial Intelligence (AI) and Machine Learning (ML) has opened new possibilities in the field of medical imaging. Deep learning techniques, particularly Convolutional Neural Networks (CNNs), have demonstrated remarkable performance in image classification and disease detection tasks. CNNs are capable of automatically extracting

meaningful features from images without the need for manual feature engineering. This makes them highly effective for analyzing chest X-ray images and detecting pneumonia with high accuracy and consistency.

a deep learning-based pneumonia detection system is developed using chest X-ray images. The proposed system employs a Convolutional Neural Network to classify images as either normal or pneumonia-infected. To ensure practical usability, the trained model is integrated into a web-based application developed using Python and the Django framework, enabling users to upload X-ray images and receive instant predictions. The system is designed to support healthcare professionals by providing quick, accurate, and reliable diagnostic assistance. Overall, this work highlights the potential of combining deep learning techniques with web technologies to improve medical diagnosis and enhance healthcare services.

II LITERATURE SURVEY

Pneumonia detection using chest X-ray imaging has been widely explored in medical research due to its importance in reducing mortality through early diagnosis. Traditionally, pneumonia diagnosis relies on radiologists interpreting chest X-ray images to identify abnormalities such as lung opacities and consolidations. Although this approach is clinically effective, it is time-consuming and

subject to inter-observer variability and human error [1]. The increasing patient load in hospitals has further highlighted the need for automated diagnostic systems that can assist medical professionals and improve efficiency.

Early research efforts focused on conventional machine learning techniques for pneumonia classification. These approaches involved manual feature extraction from chest X-ray images, followed by classification using algorithms such as Support Vector Machines (SVM), K-Nearest Neighbors (KNN), and Naïve Bayes classifiers [2]. While these methods showed moderate performance, their success largely depended on the quality of handcrafted features. Moreover, such models often struggled to generalize across different datasets due to variations in imaging conditions and patient characteristics, limiting their reliability in real-world clinical settings [3].

The emergence of deep learning significantly improved performance in medical image analysis. Convolutional Neural Networks (CNNs), introduced for image recognition tasks, demonstrated the ability to automatically learn hierarchical features directly from raw images [4]. In the context of pneumonia detection, several studies applied CNN architectures to chest X-ray datasets and reported substantial improvements in classification accuracy. Notably, Rajpurkar et al. developed CheXNet, a deep learning model based on DenseNet architecture, which achieved performance

comparable to practicing radiologists in detecting pneumonia from chest X-rays [5]. Similarly, other studies utilizing architectures such as VGG, ResNet, and Inception networks demonstrated strong results in identifying pneumonia-related patterns in radiographic images [6].

Transfer learning has also played a crucial role in improving pneumonia detection systems. Pre-trained models originally trained on large image datasets such as ImageNet were fine-tuned using medical X-ray images, enabling improved feature extraction and reduced training time [7]. This approach proved particularly beneficial when dealing with limited medical datasets. However, while transfer learning models achieve high accuracy, they often require significant computational resources, which can be a challenge for deployment in low-resource healthcare environments [8].

Recent research has emphasized the importance of implementing pneumonia detection systems in real-world applications through web-based and cloud-based platforms. Integrating deep learning models into web applications allows healthcare professionals to upload chest X-ray images and obtain rapid diagnostic results [9]. Cloud deployment further enhances system scalability, accessibility, and availability, especially in remote areas. However, issues related to data privacy, computational efficiency, and model interpretability remain critical considerations in practical deployment [10].

From the existing literature, it is evident that deep learning approaches outperform traditional machine learning methods in pneumonia detection tasks. However, balancing model accuracy with computational efficiency and real-world usability remains an ongoing research challenge. Building upon these studies, the present work focuses on developing an efficient CNN-based pneumonia detection system integrated into a web-based application, aiming to provide an accurate, scalable, and accessible diagnostic support tool for healthcare professionals.

III EXISTING SYSTEM

At present, pneumonia diagnosis mainly depends on traditional clinical practices and the manual examination of chest X-ray images by radiologists. In hospitals and diagnostic centers, chest X-rays are taken using standard imaging equipment and then carefully reviewed by trained specialists. The radiologist studies the image to identify visible signs of infection, such as lung opacities, consolidations, or fluid buildup. This conventional method has been followed for many years and is widely accepted in medical practice. However, even though it is clinically reliable, it has certain limitations that affect the speed, consistency, and overall efficiency of the diagnostic process.

One of the major challenges of the existing system is its complete reliance on human expertise. Accurate interpretation of chest X-ray

images requires extensive training and years of practical experience. Even skilled radiologists may sometimes have different opinions while analyzing the same image, as diagnosis can be influenced by individual judgment. This subjectivity can lead to inconsistencies, particularly in early-stage pneumonia cases where the signs are mild and difficult to recognize. In such situations, there is a possibility of misdiagnosis or delayed detection, which can impact timely treatment.

Another limitation is the time-consuming nature of manual analysis. In many healthcare facilities, especially large hospitals, radiologists must evaluate a high volume of X-ray images every day. The continuous workload can cause fatigue and reduce concentration levels, increasing the risk of human error. As patient numbers continue to grow, the pressure on medical professionals also increases, which may slow down the diagnostic process. Delayed diagnosis can be critical in pneumonia cases, as early intervention plays a key role in preventing severe complications.

The existing system also lacks automation and scalability. Since diagnosis is performed manually, it is difficult to process large amounts of imaging data quickly. In rural or under-resourced healthcare settings, the shortage of experienced radiologists further limits access to timely diagnosis. Patients in such areas may face delays or may not receive expert evaluation at all. Additionally, traditional methods do not

provide any computational support to assist doctors in confirming their findings. Radiologists rely solely on visual inspection and clinical knowledge, which may sometimes lead to overlooking subtle patterns in the images.

while the traditional pneumonia diagnosis system has been effective for many years, its dependence on human interpretation, time consumption, lack of scalability, and absence of automated support highlight the need for more advanced and efficient solutions.

IV PROBLEM STATEMENT

Pneumonia continues to be a serious global health problem, affecting millions of people each year and posing a high risk to children, elderly individuals, and those with weakened immune systems. Although chest X-rays are commonly used to diagnose pneumonia, the process of analyzing these images depends entirely on radiologists. Careful examination of X-ray images requires experience, concentration, and time. In busy hospitals, where a large number of cases must be handled daily, radiologists often face heavy workloads. This can lead to delays in diagnosis and, in some cases, human errors due to fatigue or oversight. Early-stage pneumonia can be especially difficult to detect because the visual signs may be subtle, increasing the possibility of misinterpretation. Another important concern is the lack of access to skilled radiologists in rural or under-resourced healthcare centers. Many

such facilities do not have specialists available at all times, which can delay proper diagnosis and treatment. Without any automated support system, doctors must rely solely on manual observation and clinical judgment. This situation creates a gap in timely and accurate healthcare delivery, particularly in regions where medical resources are limited. There is a clear need for an intelligent and automated system that can assist in the early detection of pneumonia from chest X-ray images. The system should be capable of providing quick, accurate, and consistent results while being simple enough to deploy in real-world healthcare environments. Therefore, the main problem addressed in this project is to design and develop a deep learning-based pneumonia detection system that supports medical professionals by reducing workload, minimizing diagnostic errors, and improving accessibility to reliable healthcare services.

Objectives

The primary objective of this research is to develop an efficient and reliable deep learning-based approach for the detection of pneumonia using chest X-ray images. This study aims to design and evaluate a Convolutional Neural Network (CNN) model capable of automatically extracting meaningful features from radiographic images and accurately classifying them as normal or pneumonia-affected, thereby reducing reliance on manual feature engineering and minimizing human error. Another key objective is to improve diagnostic speed and

consistency by providing a computational method that can support early detection and assist healthcare professionals in clinical decision-making. The research also focuses on integrating the trained model into a web-based framework to ensure practical applicability and real-time accessibility in healthcare environments. Furthermore, this study seeks to analyze the performance, scalability, and feasibility of deploying the proposed system in real-world settings, particularly in resource-constrained regions, with the goal of enhancing accessibility, reliability, and efficiency in pneumonia diagnosis.

V PROPOSED SYSTEM

The proposed system presents an intelligent and automated approach for detecting pneumonia from chest X-ray images using deep learning techniques. Unlike the traditional diagnostic method that relies entirely on manual interpretation by radiologists, the proposed approach utilizes a Convolutional Neural Network (CNN) model to analyze radiographic images and identify signs of pneumonia. The main purpose of this system is to provide accurate, fast, and consistent diagnostic support to healthcare professionals. By combining deep learning with a web-based framework, the system ensures that the solution is not only technically effective but also practical for real-world clinical use. Chest X-ray images are uploaded through a simple and user-friendly web interface. Once uploaded, the images

undergo preprocessing steps such as resizing and normalization to ensure consistent input quality for the model. The processed images are then passed to the trained CNN model, which automatically extracts important features related to lung abnormalities and infection patterns. Based on these learned features, the model classifies each image as either normal or pneumonia-affected. The prediction result is generated within a few seconds and displayed to the user, enabling quick and informed clinical decision-making. The system serves as a supportive second opinion, enhancing confidence and consistency in interpretation. The automatic feature extraction capability of the CNN enables the detection of subtle visual patterns that may not be easily noticeable during manual examination, particularly in early stages of pneumonia. Another significant advantage is the reduction in diagnostic time. The automated model processes images rapidly, which is especially valuable in emergency situations where early intervention is crucial. Faster screening can lead to timely treatment and improved patient outcomes. Additionally, by handling routine image analysis tasks, the system helps reduce the workload of radiologists, allowing them to focus on more complex cases. The proposed system is designed to be computationally efficient and capable of operating on standard computer systems without requiring specialized hardware. Its web-based implementation ensures easy accessibility and scalability, making it suitable for hospitals,

clinics, and rural healthcare centers. By offering a practical, reliable, and efficient diagnostic support tool, the proposed system aims to enhance pneumonia detection and improve overall healthcare service delivery.

VI METHODOLOGY

The methodology of this research focuses on designing and developing an automated pneumonia detection system using deep learning techniques applied to chest X-ray images. The overall approach follows a structured pipeline that includes data preparation, image preprocessing, model development, classification, explainability integration, and web-based deployment.

The process begins with dataset preparation. Chest X-ray images are organized into structured categories representing normal and pneumonia cases. Proper labeling and separation of training and testing data are ensured to maintain consistency and reliability during model evaluation. Before feeding the images into the model, preprocessing steps are applied. These include resizing images to a fixed resolution, normalizing pixel intensity values, and converting them into numerical arrays suitable for deep learning input. Preprocessing ensures uniformity and improves model stability.

The core of the methodology is the use of a Convolutional Neural Network (CNN) for feature extraction and classification. The CNN architecture is designed to automatically learn

hierarchical features from chest X-ray images without manual feature engineering. Convolution layers are used to extract spatial patterns such as edges, textures, and lung abnormalities. These layers are followed by activation functions such as Rectified Linear Unit (ReLU) to introduce non-linearity. Max-pooling layers are incorporated to reduce dimensionality while preserving important features, thereby improving computational efficiency and reducing overfitting.

After feature extraction, the output is passed through fully connected layers, and a softmax function is applied to generate probability scores for each class. The final classification determines whether the X-ray image is normal or pneumonia-affected. The model is trained using labeled data and evaluated based on performance metrics such as accuracy and reliability.

To improve transparency and clinical trust, the methodology integrates the Grad-CAM (Gradient-weighted Class Activation Mapping) technique. Grad-CAM generates heatmaps that highlight the regions of the X-ray image that most influenced the model's decision. This explainability component helps clinicians visually verify predictions and enhances the interpretability of the system.

Finally, the trained model is integrated into a web-based application framework. The system is deployed on a cloud platform to ensure

scalability, availability, and real-time accessibility. The overall methodology ensures that the solution is accurate, efficient, interpretable, and suitable for real-world healthcare environments.

VII IMPLEMENTATION

The implementation phase translates the proposed methodology into a functional and deployable system. The entire application is developed using Python, with deep learning components implemented using PyTorch and web functionalities handled through a backend framework.

The dataset is first organized into structured directories representing different classes. During runtime, uploaded images are passed through a preprocessing pipeline that resizes images to a standard dimension (224×224 pixels), converts them into tensors, and applies normalization to maintain consistency with training conditions.

The CNN model is implemented using a pre-trained architecture (ResNet18), which is fine-tuned for pneumonia detection by modifying the final fully connected layer to match the required output classes. The trained model parameters are saved and loaded during prediction time to avoid retraining, thereby improving efficiency and reducing computational overhead.

When a user uploads a chest X-ray image through the web interface, the backend securely stores the file and triggers the prediction

function. The model performs a forward pass on the preprocessed image and generates output probabilities. The class with the highest probability is selected as the final prediction, which may correspond to Normal, Pneumonia, or other defined categories. The prediction result is then returned to the user interface in real time.

The system also generates Grad-CAM visualizations during inference. These heatmaps are overlaid on the original X-ray image to highlight influential regions, providing interpretability alongside the prediction result.

The web application interface is designed to be simple and user-friendly. It allows users to upload images and view results without requiring technical expertise. Backend operations, including request handling, image storage, preprocessing, model inference, and response generation, are managed efficiently to ensure smooth execution.

The complete system is deployed on a cloud platform to ensure high availability and scalability. Cloud deployment enables the application to handle multiple users simultaneously and process real-time prediction requests without requiring specialized local hardware.

IX RESULTS


The experimental evaluation of the proposed pneumonia detection system demonstrates its effectiveness in accurately classifying chest X-

ray images. The model was trained using augmented image data to enhance generalization and reduce overfitting. Data augmentation techniques improved the robustness of the model by exposing it to varied image orientations and intensities. Training was performed using GPU support on Kaggle Kernels, which enabled efficient processing and faster convergence. After multiple experimental trials, the model achieved optimal performance at 100 training epochs, where validation performance stabilized and no significant overfitting was observed.


The performance of the model was thoroughly evaluated using a confusion matrix, which provided a detailed breakdown of classification results. The matrix clearly illustrated the distribution of correctly and incorrectly classified instances across normal and pneumonia categories. From this analysis, key evaluation metrics such as accuracy, precision, recall, and F1-score were computed. The high accuracy value indicates that the model is capable of correctly identifying the majority of cases. Precision results confirm that the system effectively reduces false positive predictions, while recall demonstrates strong capability in detecting pneumonia cases without missing significant instances. The F1-score reflects a balanced trade-off between precision and recall, further validating the reliability of the model.

The Validation Accuracy versus Validation Loss curve provides additional insight into model performance. As training progressed, validation


accuracy steadily increased while validation loss consistently decreased, indicating effective learning behavior. The absence of large fluctuations between training and validation curves suggests good generalization capability and stable optimization.



Validation Accuracy vs
Validation Loss



Confusion Matrix



X CONCLUSION

This research successfully demonstrates the effective application of deep learning techniques for automated pneumonia detection using chest X-ray images. By leveraging Convolutional Neural Networks (CNNs), the study highlights how artificial intelligence can significantly enhance diagnostic accuracy and consistency in medical image analysis. The developed system is designed with a strong emphasis on reliability, efficiency, and practical usability, ensuring that it can function as a supportive clinical decision-making tool rather than replacing medical professionals. The use of a CNN-based architecture enables automatic feature extraction directly from radiographic images, eliminating the need for manual feature engineering and reducing the possibility of human error. The model's ability to accurately differentiate between normal and pneumonia-affected lungs confirms the effectiveness of deep learning approaches in handling complex medical imaging tasks. Performance evaluation using standard metrics such as accuracy, precision, recall, and F1-score further validates the robustness and stability of the proposed framework. The integration of the trained model into a structured web-based application ensures seamless interaction between the user interface, backend processing, and prediction module. This implementation demonstrates how research-oriented models can be transformed into

deployable real-world solutions. Cloud deployment enhances scalability and accessibility, allowing the system to support multiple users efficiently and making it suitable for hospitals, clinics, and resource-limited healthcare centers. An important contribution of this research is the inclusion of explainability through Grad-CAM visualizations, which provide insight into the regions of the X-ray image influencing the model's predictions. This transparency increases clinical trust and supports the responsible adoption of AI-assisted diagnostic systems. Comprehensive validation confirms that the system meets both functional and performance requirements under different operating conditions.

REFERENCES

- [1] R. S. Kumar, A. Sharma, and P. Verma, "Deep learning techniques for pneumonia detection using chest X-ray images," *IEEE Access*, vol. 8, pp. 134920–134930, 2020.
- [2] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [3] A. Rajpurkar et al., "CheXNet: Radiologist-level pneumonia detection on chest X-rays with deep learning," *arXiv preprint arXiv:1711.05225*, 2017.
- [4] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [5] S. Litjens et al., "A survey on deep learning in medical image analysis," *Medical Image Analysis*, vol. 42, pp. 60–88, 2017.
- [6] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proc. CVPR*, 2017, pp. 4700–4708.
- [7] T. Shen, C. Gou, J. Wang, and F. Chen, "A deep learning-based method for pneumonia diagnosis," *IEEE Transactions on Medical Imaging*, vol. 39, no. 5, pp. 1714–1725, 2020.
- [8] S. Minaee et al., "Deep learning-based approaches for COVID-19 detection from chest X-ray images," *IEEE Access*, vol. 8, pp. 181178–181195, 2020.
- [9] Saikumar, B. (2024). Optimizing Crew Scheduling and Absence Management using Microservices: Enhancing Reliability and Efficiency in Crew Management Systems. *International Journal of Enhanced Research in Management & Computer Applications*, 13(11), 50–55. <https://doi.org/10.55948/ijermca.2024.0116>.
- [10] Kalae, U. K. (2025). Optimizing cost-effective cloud data pipeline orchestration across multiple cloud providers. *Journal of Information*

Systems Engineering and Management, 10(63s), e726–e741.

[11] Reddy, S. K. R. (2021). Strengthening the Security of Loyalty Reward Systems: An In-Depth Analysis of Emerging Cyber Threats and Protection Mechanisms. *Journal of Computational Analysis and Applications*, 29(6).

[12] M. Abadi et al., “TensorFlow: A system for large-scale machine learning,” in *Proc. OSDI*, 2016, pp. 265–283.

[13] F. Chollet, *Deep Learning with Python*. New York, NY, USA: Manning Publications, 2018.

[14] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *Proc. ICLR*, 2015.

[15] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.

[16] R. Girshick et al., “Region-based convolutional networks for accurate object detection,” *IEEE TPAMI*, vol. 38, no. 1, pp. 142–158, 2016.