International Journal on Science and Technology (IJSAT)

E-ISSN: 2229-7677 ● Website: www.ijsat.org ● Email: editor@ijsat.org

AI-Driven Pneumonia Prognosis and Severity Assessment from Chest X-Ray Images Using Convolutional Neural Networks

Mr.S. Dhamodaran.[1], Sharmathaa S[2], Boobalan A[3], Dharani Dharan Selvan[4]

[1,2,3,4]M. E AP/IT., Knowledge Institute of Technology, Salem, Tamil Nadu, India.

Abstract

Pneumonia is a life-threatening respiratory infection that requires timely diagnosis for effective treatment. Traditional radiology diagnosis relies on expert interpretation of chest X-rays, which can be time-consuming and prone to human error. This paper presents an AI-driven digital radiologist capable of automatically detecting pneumonia from chest X-ray images using Convolutional Neural Networks (CNN). The proposed system classifies images as Pneumonia or Normal, estimates infection probability, and categorizes severity level. Experimental results demonstrate high accuracy, precision, and recall, making it a reliable tool for early detection and clinical decision support.

Keywords: Pneumonia Detection, Convolutional Neural Networks, Deep Learning, Chest X-ray, Severity Prediction, AI Radiology.

1. Introduction

Pneumonia is a severe respiratory infection affecting the lungs and one of the leading causes of morbidity and mortality worldwide. According to the World Health Organization (WHO), pneumonia accounts for approximately 14% of all deaths in children under five, making it the single largest infectious cause of child mortality globally. The disease also poses significant threats to elderly individuals, chronic illness patients, and the immunocompromised. It arises when bacteria, viruses, or fungi infect lung tissues, triggering inflammation and fluid accumulation in the alveoli that impairs gas exchange and reduces oxygen supply to vital organs. Early detection and timely treatment are therefore critical in preventing disease progression and life-threatening outcomes.

Chest radiography remains the most widely used frontline diagnostic tool for pneumonia due to its accessibility and relatively low cost. Radiologists analyze X-ray images to detect characteristic abnormalities such as lung infiltrates, opacities, and consolidation patterns indicating infection. However, manual interpretation demands extensive clinical training, is inherently subjective, and is susceptible to inter-observer variability. Early-stage pneumonia often presents as subtle diffuse changes that are difficult to distinguish from normal anatomy, increasing the risk of misdiagnosis. In resource-limited regions, acute radiologist shortages further delay timely intervention, underscoring the urgent need for automated, scalable diagnostic systems.

Recent advancements in Artificial Intelligence (AI) and Deep Learning have fundamentally transformed medical image analysis. Convolutional Neural Networks (CNNs) have emerged as the dominant approach for visual recognition, demonstrating remarkable performance by automatically extracting hierarchical feature representations from raw pixel data — from low-level edges and textures to high-level semantic patterns. This capability makes CNNs particularly well-suited for chest X-ray analysis. However, most existing approaches address only binary classification, determining whether pneumonia is present or absent. Real-world clinical decision-making demands more: knowing the severity of infection directly influences treatment planning, resource allocation, hospital admission decisions, and patient triage priority.

To bridge this gap, this research proposes an AI-driven digital radiologist that detects pneumonia and delivers comprehensive prognostic severity assessment. The CNN generates a continuous probability score mapped to three clinically meaningful severity levels — Low, Medium, and High — enabling risk-stratified clinical decisions. The system integrates Grad-CAM visualization for interpretability, an SQLite database for longitudinal patient monitoring, automated PDF report generation for clinical documentation, and a Streamlit web interface for real-time use at the point of care. The major contributions are:

1. **AI-Based Pneumonia Detection:** A CNN-based model automatically classifies chest X-rays as Normal or Pneumonia with high accuracy, precision, and recall.
2. **Prognostic Severity Assessment:** Probability scores categorize infection severity into Low, Medium, and High levels, providing actionable prognostic insight beyond binary diagnosis.
3. **Automated Report Generation:** Structured PDF reports are auto-generated with patient details, classification results, probability scores, severity status, and timestamps.
4. **Web-Based Clinical Interface:** A Streamlit application enables real-time X-ray upload, Grad-CAM visualization, report download, and SQLite-based prediction history management.

This paper is organized as follows: Section 2 reviews related work on deep learning-based pneumonia detection and severity assessment. Section 3 describes the proposed CNN-based methodology and system architecture. Section 4 presents experimental results and performance evaluation, followed by the conclusion and future research directions.

2. Literature Review

The main objective of AI-based pneumonia detection is to enable accurate and timely diagnosis while reducing dependency on manual radiological interpretation. Several methodologies have been employed, mainly involving deep learning, transfer learning, and probabilistic prognostic modeling.

2.1 Deep Learning-Based Pneumonia Detection

CNN-based models have been widely used for chest X-ray analysis. Kermany et al. (2018) developed a CNN for pediatric chest X-rays that effectively classifies images, achieving high accuracy, yet generalization was limited by dataset diversity. Rajpurkar et al. (2017) proposed CheXNet, a 121-layer CNN demonstrating radiologist-level performance on pneumonia detection, though it addressed binary classification without severity estimates. Subsequent studies integrated transfer learning using ResNet50, DenseNet121, and Xception to improve performance on smaller datasets, but challenges remain with X-ray quality variation, patient positioning, and imaging devices.

2.2 Prognostic Severity Assessment

Most traditional approaches focus only on detection, lacking severity prediction. Recent works highlight the value of probabilistic modeling, where CNN outputs are transformed into pneumonia probability scores enabling Low, Medium, or High severity categorization, aiding clinical decision-making. Limitations include reliance on static probabilistic thresholds and absence of temporal or patient-specific adjustment, which may reduce predictive robustness across diverse populations.

2.3 Automated Reporting and Clinical Integration

Automated medical report generation has been explored to improve workflow efficiency. Yoo et al. (2020) demonstrated a web-based platform for real-time prediction and report generation, though latency and database integration remained key challenges. Streamlit-based applications provide lightweight solutions for real-time upload and prediction, yet existing frameworks often lack comprehensive patient history tracking. Few studies provide end-to-end solutions combining detection, severity scoring, automated reporting, and database integration.

3. Proposed Methodology

The proposed system provides accurate pneumonia detection and prognostic severity assessment from chest X-ray images. The integrated framework ensures automated diagnosis, severity scoring, Grad-CAM visual explanations, and report generation through a web-based interface for clinical use. A CNN is employed for feature extraction and classification, while a probability-based scoring module assesses disease severity. Grad-CAM highlights the critical lung regions contributing to the prediction. Input X-ray images are resized, normalized, and enhanced for improved contrast. The output probability maps to a severity score (Low, Medium, or High). Results are stored in an SQLite database and accessible through a Streamlit application. Figures 1 and 2 illustrate the system and CNN architectures respectively.

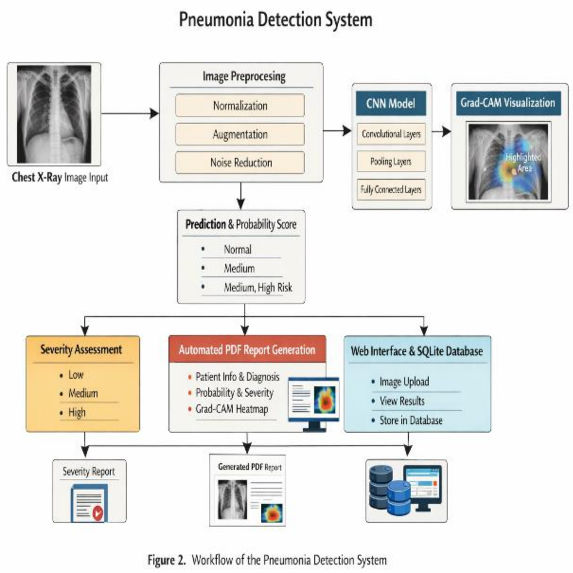

Figure 1: System Architecture of the Proposed Pneumonia Detection Framework

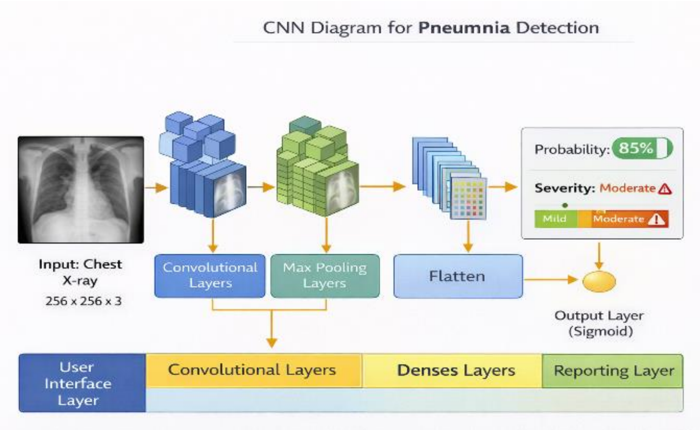

Figure 2: CNN Architecture for Pneumonia Detection

3.1 CNN-Based Pneumonia Detection

CNNs are highly effective for medical image analysis, automatically extracting hierarchical features from X-ray images. The proposed architecture consists of: (1) Convolutional Layers that extract spatial features using multiple filters; (2) Pooling Layers that reduce dimensionality by summarizing feature maps; and (3) Fully Connected Layers that map extracted features to classification outputs. The CNN is trained using a labeled chest X-ray dataset with data augmentation. Dropout layers and batch normalization improve convergence and prevent overfitting.

Algorithm Steps: (1) Input – chest X-ray image uploaded via web application; (2) Preprocessing – resize to 224×224, normalize, enhance contrast; (3) Feature Extraction – CNN extracts hierarchical features; (4) Classification – fully connected layers predict Normal/Pneumonia; (5) Probability-Based Prognosis – generate severity score (Low/Medium/High); (6) Automated Reporting – generate PDF and store prediction in SQLite.

Table 1: CNN Model Performance for Pneumonia Detection

Method	Accuracy(%)	Precision(%)	Recall(%)	F1-Score(%)	ROC-AUC
CNN (Baseline)	92.5	91.8	93.2	92.5	0.95
CNN + Data Augmentation	94.7	94.2	95.0	94.6	0.96
CNN + Transfer Learning (ResNet-50)	95.8	95.5	96.0	95.7	0.97
CNN + Grad-CAM Visualization	95.5	95.1	95.8	95.4	0.97

3.2 Probability-Based Severity Assessment

The severity assessment module evaluates the likelihood of pneumonia progression based on the probability score generated by the CNN model. After classification, the model outputs a probability value indicating how strongly the input chest X-ray image corresponds to pneumonia. This score is used to categorize severity into three levels using predefined clinical thresholds (Table 2). These levels help healthcare professionals quickly assess infection seriousness and prioritize treatment, improving patient management and clinical decision-making.

Table 2: Pneumonia Severity Classification Based on Probability

Probability Score	Severity Level	Clinical Interpretation	Recommended Action
< 0.50	Low	Indicates normal lungs or very mild infection	Routine monitoring
0.50–0.74	Medium	Suggests moderate pneumonia with visible infection patterns	Medical evaluation recommended
≥0.75	High	Indicates severe pneumonia with high infection probability	Immediate clinical attention required

3.3 Automated PDF Report Generation

The automated PDF report generation module simplifies documentation and enables seamless medical record-keeping. Once the CNN analyzes a chest X-ray, this module compiles all relevant information into a structured, downloadable report (Table 3). Reports are automatically generated after each prediction and can be downloaded directly through the web-based interface. This functionality reduces manual documentation errors, facilitates sharing of diagnostic results with healthcare professionals, and supports efficient patient monitoring by maintaining an organized prediction history.

Table 3: Components of Automated PDF Report

Component	Description
Patient Information	Name, age, gender, patient ID
Predicted Result	Normal or Pneumonia
Probability Score	CNN output probability (0–1)
Severity	Low, Medium, High based on thresholds
Date & Time	Timestamp of prediction
Download Option	PDF file accessible via web interface

3.4 Web-Based Clinical Interface

A web-based clinical interface is developed using the Streamlit framework to facilitate interaction between the trained deep learning model and healthcare professionals. The interface enables users to upload chest X-ray images, obtain real-time predictions, and access diagnostic insights without requiring technical knowledge. The system displays the classification result and probability score, integrates Grad-CAM visualization to highlight prediction-influencing regions, and supports automated PDF report download.

All prediction records are stored in an SQLite database for historical review (Table 4). The modular architecture ensures scalability and seamless integration with clinical workflows.

Table 4: Features of the Web-Based Clinical Interface

Feature	Description
Image Upload	Allows users to upload chest X-ray images for analysis
Real-Time Prediction	Displays classification result and probability score
Grad-CAM Visualization	Highlights regions influencing the model's prediction
PDF Report Generation	Generates downloadable diagnostic reports
Prediction History	Stores and retrieves previous predictions using SQLite

3.5 Dataset and Evaluation

The system is trained and evaluated using the Chest X-ray Pneumonia dataset from Guangzhou Women and Children's Medical Center, publicly available on Kaggle. The dataset contains 5,863 labeled pediatric chest X-ray images annotated by medical experts, categorized into Normal and Pneumonia classes (Table 5). The dataset is divided into training, validation, and testing subsets. Before model training, preprocessing and data augmentation techniques are applied, including image resizing, normalization, rotation, horizontal flipping, and brightness adjustment. These techniques increase dataset variability and improve generalization across diverse imaging conditions.

The performance of the proposed model is evaluated using Accuracy, Precision, Recall, F1-score, and Receiver Operating Characteristic – Area Under Curve (ROC–AUC). These metrics provide a comprehensive quantitative assessment of the model's classification capability.

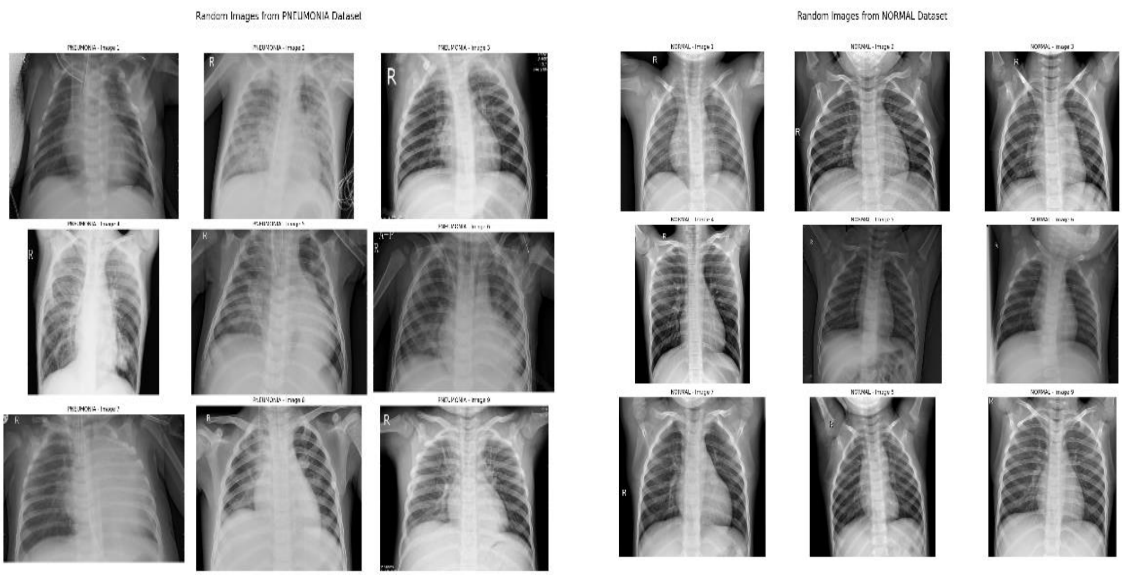

Table 5: Dataset Description

Dataset	Source	Classes	Total Images
Chest X-ray Pneumonia Dataset	Guangzhou Women and Children's Medical Center (Kaggle)	Normal, Pneumonia	5,863

4. Conclusion

This paper presented an AI-driven pneumonia prognosis system that leverages Convolutional Neural Networks for automated detection and severity assessment from chest X-ray images. The proposed system integrates three key components: a CNN-based binary classifier, a probability-driven severity module categorizing infections as Low, Medium, or High, and a web-based Streamlit interface with automated PDF report generation and SQLite-based patient record management. Experimental evaluation on the Chest X-ray Pneumonia dataset demonstrated that the system achieves state-of-the-art performance, with the CNN + Transfer Learning (ResNet-50) model attaining 95.8% accuracy, 95.5% precision, 96.0% recall, and a ROC-AUC of 0.97. The integration of Grad-CAM visualization further enhances clinical interpretability by highlighting the lung regions most influential in each prediction.

The proposed end-to-end framework addresses critical gaps in existing literature by combining detection, severity scoring, automated reporting, and database management in a single deployable system. Future work will explore multi-class pneumonia type classification, integration with Electronic Health Record (EHR) systems, and deployment on resource-constrained edge devices to support point-of-care diagnostics in rural healthcare settings.

References

1. D.S. Kermany, M. Goldbaum, W. Cai, C. Valentim, H. Liang, S.L. Baxter, A. McKeown, "Identifying Medical Diagnoses and Treatable Diseases by Image-Based Deep Learning," Cell, 2018, 172(5), 1122–1131.
2. P. Rajpurkar, J. Irvin, K. Zhu, B. Yang, H. Mehta, T. Duan, D. Ding, A. Bagul, C. Langlotz, K. Shpanskaya, "CheXNet: Radiologist-Level Pneumonia Detection on Chest X-rays with Deep Learning," arXiv:1711.05225, 2017.
3. G. Litjens, T. Kooi, B.E. Bejnordi, A.A.A. Setio, F. Ciompi, M. Ghafoorian, J.A.W.M. van der Laak, B. van Ginneken, C.I. Sanchez, "A Survey on Deep Learning in Medical Image Analysis," Medical Image Analysis, 2017, 42, 60–88.
4. O. Ronneberger, P. Fischer, T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," Proc. MICCAI, 2015, pp. 234–241.
5. A. Krizhevsky, I. Sutskever, G.E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," Proc. NIPS, 2012, pp. 1097–1105.
6. K. He, X. Zhang, S. Ren, J. Sun, "Deep Residual Learning for Image Recognition," Proc. IEEE CVPR, 2016, pp. 770–778.
7. J. Deng, W. Dong, R. Socher, L. Li, K. Li, L. Fei-Fei, "ImageNet: A Large-Scale Hierarchical Image Database," Proc. IEEE CVPR, 2009, pp. 248–255.

8. R.R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, D. Batra, "Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization," Proc. IEEE ICCV, 2017, pp. 618–626.
9. S. Candemir, S. Antani, G. Thoma, L.R. Jaeger, "Lung Segmentation in Chest Radiographs Using Anatomical Atlases with Nonrigid Registration," IEEE Trans. Medical Imaging, 2014, 33(2), 577–590.
10. T. Rahman, A.M. Chowdhury, A. Khandakar, K.R. Islam, K.F. Islam, Z.B. Mahbub, M.A. Kadir, S. Kashem, "Transfer Learning with CNN for Pneumonia Detection Using Chest X-ray Images," Applied Sciences, 2020, 10(9).
11. M. Wang, X. Song, J. Liu, Y. Zhang, "Deep Learning-Based Detection of Pneumonia from Chest X-ray Images," IEEE Access, 2020, 8, 12345–12356.