



An Optimized Lightweight CNN Architecture for Automated Pneumonia Classification from Chest X-Ray Images

Princi Saini*¹  and Mridula Singh¹ 

¹ Department of Computer Science and Engineering, Quantum University, Roorkee
*princisaini55@gmail.com, mridulaiitr@gmail.com

Abstract. Chest X-ray images are usually interpreted by expert radiologists for the diagnosis of pneumonia, which is a time-consuming and subjective process. We propose a light-weight and computationally efficient convolutional neural network (CNN) for automated pneumonia detection to deal with these challenges. The goal is to achieve high diagnostic accuracy and low computational requirements. This makes the approach suitable for practical use. The dataset of the proposed framework includes 5,863 chest X-ray images of pneumonia and normal cases. Before training, the images are processed with standardization and enhancement techniques and data augmentation methods to better equip the model to generalize over diverse samples. The CNN model is trained on these prepared inputs and is designed to balance performance and efficiency. The experimental results show that the model can reach an accuracy of 95.8% and has a strong ability to distinguish between pneumonia and healthy cases. Such a performance indicates that the model could be used for a reliable and prompt diagnosis. Overall, the study provides an AI based solution which can help healthcare professionals in early detection of pneumonia. The proposed approach is lightweight, making it especially useful in scenarios with limited computational resources, which can enable wider adoption of intelligent diagnostic systems.

Keywords: Pneumonia Detection, Convolutional Neural Network (CNN), Chest X-ray Image, Deep learning, Medical Image analysis.

1 Introduction

Pneumonia is still a major global health problem, especially in young children and older adults, where delays in diagnosis may result in serious complications or even death. Detection using chest X-ray imaging is a widely used technique. However, its interpretation is highly dependent on clinical expertise which may lead to variability and delay in decision making.

The indicators of pneumonia—such as cough, fever, chest discomfort, and difficulty in breathing—are often similar to those of other respiratory conditions. This overlap makes early and accurate diagnosis challenging. Chest radiographs help

reveal lung abnormalities like fluid accumulation and tissue opacities, making them a standard diagnostic tool. However, analyzing these images is typically performed manually by radiologists, which can be time-consuming and prone to variability.

To address these challenges, this study presents a lightweight and efficient convolutional neural network (CNN) designed specifically for pneumonia detection from chest X-ray images. The model uses a compact architecture with four convolutional layers to extract relevant features while keeping computational requirements low. Techniques such as pooling and dropout are incorporated to improve generalization and reduce overfitting. The overall design focuses on achieving a balance between accuracy and efficiency, making it suitable for deployment in real-world and resource-constrained environments.

Several deep learning architectures, including VGG, ResNet, DenseNet, MobileNet, and EfficientNet, have been widely applied to pneumonia detection tasks. These models often achieve high accuracy by leveraging techniques such as transfer learning, data augmentation, and ensemble methods. Transfer learning, in particular, allows models to benefit from knowledge gained on large datasets, improving performance on smaller medical datasets while reducing training time.

Data augmentation methods, including transformations like translating, flipping, and scaling, are commonly used to improve model robustness and handle limited data. Ensemble approaches further enhance prediction correctness by combining multiple models. Despite these improvements, many existing solutions are difficult to implement in clinical practice due to their high computational demands. They often require powerful hardware and may suffer from issues such as overfitting, dataset bias, and limited generalization across diverse patient groups. Additionally, achieving real-time performance remains a challenge.

In this work, a streamlined and effective CNN-based solution is developed to provide accurate and timely pneumonia detection. The proposed model aims to support clinicians in making faster decisions, reduce the workload on radiologists, and enable early identification of the disease. Its lightweight nature allows it to be used effectively even in settings with limited computational resources, potentially improving healthcare accessibility and patient outcomes.

This study offers a practical approach that addresses the trade-off between model correctness and computational efficiency. Unlike many existing deep learning models that prioritize performance at the cost of complexity, the proposed architecture is designed to deliver reliable results while remaining suitable for real-world clinical deployment.

2 Literature Review

Table 1. Comparative Analysis of Recent Deep Learning Methods for Chest X-ray-Based Sickness Detection: Methods, Datasets, Performance, and Key Contributions (2019–2025)

S. No	Author (Year)	Approach Used	Dataset	Performance	Key Contribution / Limitation
1	Mahi et al. (2024)	Lightweight CNN	12,448 X-rays	98% accuracy	High accuracy with low

					computational cost
2	Nallamala et al. (2025)	EfficientNet-B0 vs VGG16	Not specified	91.51%	Improved performance over traditional CNN
3	Gupta et al. (2024)	Efficient Lightweight CNN	Not specified	98.89%	High accuracy but lacks detailed metrics
4	Chatterjee et al. (2021)	EfficientNetB3	Not specified	93%	Supports multi-class classification
5	Suryaa et al. (2021)	Ensemble CNN	5,863 images	95.03%	Improves robustness and accuracy
6	George et al. (2022)	CNN-based model	Not specified	Not given	Lacks experimental validation
7	Singh et al. (2024)	CNN + preprocessing	Not specified	93.75%	Highlights importance of preprocessing
8	Budhya et al.	CNN (VGG, ResNet)	Not specified	High	Outperforms traditional ML
9	Lafraxo & El Ansari (2022)	CNN + regularization	Not specified	97.30%	Better generalization
10	Saha & Rahman (2022)	Modified CNN	Not specified	89%	Moderate performance
11	Varshni et al. (2019)	DenseNet + ResNet	ChestX-ray14	AUC ~0.77	Feature extraction approach
12	Saranya et al. (2020)	CNN	30,059 images	94.3%	High recall (98.2%)
13	Islam & Pathari (2023)	CNN	Not specified	93%	Works on low-quality images
14	Hashmi et al. (2020)	Ensemble / Transfer Learning	5,836 images	98.43%	Very high AUC (99.76%)
15	Rahman et al. (2020)	Transfer Learning	5,247 images	98%	Effective multi-class detection
16	Sharma et al. (2020)	CNN + augmentation	Not specified	Not given	Improves generalization
17	Labhane et al. (2020)	CNN + Transfer Learning	5,964 images	>97%	Strong pediatric dataset results
18	Aljawarneh & Al-Quraan (2025)	Enhanced CNN	5,863 images	92.46%	Limited validation
19	Kumar & Sharma (2025)	MobileNetV2	624 images	92.6%	Lightweight, suitable for mobile
20	Haque et al. (2023)	Deep CNN	Not specified	98.22%	Strong feature extraction
21	El Asnaoui et al. (2021)	Transfer Learning models	Not specified	>96%	Multiple architectures compared
22	Ahmad et al. (2023)	Multiple CNN architectures	8,100 images	98.33%	Hybrid model performs best
23	Eido & Yasin (2025)	CNN + augmentation	Not specified	Up to 99%	Performance varies with dataset
24	Chouhan et al. (2020)	Transfer Learning CNN	Not specified	High	Limited clinical validation
25	Santhoshi & Jyostna (2023)	CNN + RCNN	Not specified	Not given	Adds localization capability
26	Chouhan et al. (2020)	Transfer Learning CNN	Not specified	High	Repeated approach
27	AlGhamdi (2024)	MobileNetV3	>14,000 images	96–99%	Efficient for real-time use
28	Badlani et al. (2022)	Custom CNN	Not specified	89.2%	Balanced metrics, moderate accuracy
29	Kundu et al. (2021)	Ensemble CNN	5,856 images	~98.8%	Strong generalization
30	An et al. (2024)	Hybrid CNN +	5,856 images	95.19%	Advanced but computationally
Attention					heavy

As shown in Table 1, the comparative analysis of past research indicates that existing deep learning approaches for chest X-ray analysis consistently report strong classification performance, often exceeding 90% accuracy. However, more accuracy often comes at the cost of higher computational complexity, making their practical deployment in low-resource settings difficult. More advanced approaches, such as ensembles and hybrid architectures, deliver even higher accuracy, but they come with a significantly higher computational cost than single models. Transfer learning is commonly used to improve the accuracy of models that are primarily trained on small datasets, and lightweight models such as MobileNet have been shown to be suitable for real-time use in resource-constrained environments. Existing studies have many limitations including insufficient reporting of important metrics used to assess performance such as accuracy, recall and area under the curve (AUC), poor transparency and reproducibility of data sets, and minimal validation in clinical practice. Furthermore, there is a trade-off between the complexity of a model and the ability to deploy practically.

On analyzing the research it is clear that most of the researchers build either an accurate system through complicated models or speed up an easy to use architecture but rarely do the researchers find a balance of both methods at once. Thus, the requirement to have an efficient convolutional neural network (CNN) to help in the diagnosis of pneumonia by using chest x-ray images is well justified. The aim of this model will be to develop an efficient CNN architecture that can provide high diagnostic accuracy and also applies a comprehensive set of evaluation metrics, while being able to operate efficiently in real time for the health care system and/or when computational resources are limited. This proposed methodology will provide a novel way of connecting high performance deep learning models with their ability to be used within a clinical environment thus forming an efficient, scalable disease detection system for pneumonia.

3 Methodology

3.1 Dataset

The study uses 5,863 chest x-ray images of 'Pneumonia and Normal'. The dataset is well organized by grouping the images in three folders (training, validation and testing) and each of those folders has the same sub-folder names to facilitate the use of the model train-validation and to keep the model performance evaluation unbiased.

These chest x-ray images (anterior-posterior view) were obtained from retrospective cohorts of pediatric patients aged 1–5 years treated at the Guangzhou Women and Children's Medical Center (Guangzhou, China). These radiographs are representative of typical clinical situations as they were all obtained during routine clinical care and there is inherent variability due to imaging conditions and disease severity.

A large amount of screening was needed to get the data to a reliable and quality state. All chest x-ray images underwent quality control (QC) procedures in the sense of removing those of low quality or with excessive noise or error due to the image being 'unreadable' before the QC labels were assigned. The QC labels were then reviewed by two expert physicians for accuracy and adherence to the Diagnostic Imaging Guidelines (DICG). In order to better improve the accuracy and consistency of QC labels, a third expert physician independently reviewed the evaluation set and made judgements on the accuracy of labels to minimize the possibility of any errors in the QC labels, thus providing a high-quality ground truth for model training and evaluation.

In the present research, an efficient CNN model will be proposed for the application of pneumonia detection through the utilization of chest X-ray images. The aim of this model is to obtain a high diagnostic accuracy and also to be computationally efficient to be deployed in real-world situations.

3.2 Dataset Description

To train and evaluate the proposed model, an available chest X-ray dataset will be adopted in this study. This dataset will contain labelled images in two categories, one of pneumonia and the other of normal, obtained from real clinical environments with varying patient demographics, severity of disease, and imaging conditions (contributing to greater diversity and producer robustness).

3.3 Data Preprocessing

To prepare and standardize input data and improve model performance, a comprehensive image preprocessing pipeline has been employed to resize all images to a standard 224×224 pixel resolution, allowing for uniformity and reduced computational complexity.

The range of pixel intensity values has been normalized between 0 and 1 during image preprocessing to help stabilize the training process by speeding up convergence time during backpropagation.

Denoising and contrast-adjusting techniques have been employed to enhance image quality and bring attention to anatomical features of interest (particularly to highlight the lung area and the opacity of the lungs due to pneumonia).

In order to account for class imbalances and improve generalization, several data augmentation techniques have been used on the images as follows:

- Rotation
- Horizontal flipping
- Zooming
- Width/height shifting

By simulating variations that occur in real life, the transformations performed will help reduce overfitting and allow for the model to learn invariant/robust features.

3.4 Proposed CNN Architecture

A new type of Convolutional Neural Network (CNN) architecture has been created that combines classification accuracy with low complexity, allowing it to function in real time in a clinical setting. The design of the architecture is simple yet elegant and will effectively extract both low and high level features, while also keeping computational costs low.

This model contains four convolutional layers, each of which has 32, 64, 128, and 128 filters, respectively, as given in Table 2. Each of the four convolutional layers has a kernel size of 3×3 , which is an effective size for extracting local features (e.g., edges, textural patterns, and opacity patterns) that relate to the presence of pneumonia in chest X-rays. The convolution operation itself can be represented mathematically as follows:

$$f_{i,j}=(I*K)_{i,j}=\sum_m\sum_n I(i+m,j+n)\cdot K(m,n)$$

where I is the input image and K is the convolution kernel.

Each convolutional layer is followed by a Rectified Linear Unit (ReLU) activation function. The pooling technique takes one maximum value in each pooling window. Therefore, the feature data with the most significant features was preserved.

To improve generalization and prevent overfitting, dropout layers with a rate of 0.5 were used. Dropout works by randomly deactivating some neurons during training, which encourages the model to learn stronger and more reliable features.

After flattening the extracted features' maps, they undergo a fully connected (dense) layer; the 1st Dense layer has 128 neurons, allowing for high-level feature learning. The final layer will output a probability score (0 or 1) using the sigmoid activation function. This is done because we are performing binary classification:

$$\sigma(x)=\frac{1}{1+e^{-x}}$$

Efficient feature extraction and computational complexity are important considerations when designing deep-learning architectures for implementation in environments with low resources. The proposed architecture has been designed and tested according to these two criteria.

Table 2. Design of the proposed Lightweight CNN, layer-wise configuration

Layer No.	Layer Type	Filters / Neurons	Kernel Size	Output Size	Activation	Remarks
1	Input Layer	—	—	$224 \times 224 \times 1$	—	Chest X-ray image
2	Conv2D	32	3×3	$224 \times 224 \times 32$	ReLU	Low-level feature extraction

3	MaxPooling2D	—	2×2	112 × 112 × 32	—	Downsampling
4	Conv2D	64	3×3	112 × 112 × 64	ReLU	Texture features
5	MaxPooling2D	—	2×2	56 × 56 × 64	—	Dimension reduction
6	Conv2D	128	3×3	56 × 56 × 128	ReLU	Deeper feature learning
7	MaxPooling2D	—	2×2	28 × 28 × 128	—	Downsampling
8	Conv2D	128	3×3	28 × 28 × 128	ReLU	High-level feature extraction
9	MaxPooling2D	—	2×2	14 × 14 × 128	—	Final spatial reduction
10	Dropout	—	—	14 × 14 × 128	—	Rate = 0.5
11	Flatten	—	—	25,088	—	Feature vector
12	Dense	128	—	128	ReLU	Fully connected layer
13	Dropout	—	—	128	—	Rate = 0.5
14	Output Dense	1	—	1	Sigmoid	Binary classification

3.5 Training Strategy

In order to achieve efficient training and high performance of the proposed CNN, an optimization strategy has been developed. An optimization algorithm is assigned to train a CNN model using adaptive learning rates for each parameter and combines the advantages of using both momentum and root mean squared propagation (RMSProp).

The parameter update rule for Adam is:

$$\theta_{t+1} = \theta_t - \frac{\eta}{\sqrt{v_t + \epsilon}} \cdot m_t$$

With m_t and v_t indicating the first and second moments, respectively, and η representing the learning rate.

The learning rate is set to 0.001, providing a compromise between convergence speed and stability, making the binary cross-entropy a good fit for use in training:

$$L = - \frac{1}{N} \sum_{i=1}^N [y_i \log(\hat{y}_i) + (1-y_i) \log(1-\hat{y}_i)]$$

Where N = number of samples, y_i = Actual Label, and $\hat{y}_i$ = Predicted Probability.

Trained using a batch size of 32 over 25–30 epochs. Early stopping is used as a method to reduce overfitting – the training is terminated when there has been no improvement in the validation loss for a specific number of epochs.

Multiple metrics from the confusion matrix have been applied to measure the performance of our proposed model. These metrics include accuracy, precision, recall (or sensitivity), specificity, F1-score, and AUC-ROC. This allows for a balanced evaluation. The evaluation metrics are defined as follows:

$$\text{Accuracy: } \frac{TP+TN}{TP+TN+FP+FN}$$

$$\text{Precision: } = \frac{TP}{TP+FP}$$

$$\text{Recall (Sensitivity) : } = \frac{TP}{TP+FN}$$

$$\text{Specificity: } = \frac{TN}{TN+FP}$$

$$\text{F1-Score: } = \frac{2 \times (\text{Precision} \times \text{Recall})}{\text{Precision} + \text{Recall}}$$

The ROC curve helps visualize the balance between correctly identified positives and incorrectly identified. The Area Under the Curve (AUC) is computed as:

$$\text{AUC} = \int_0^1 \text{TPR}(\text{FPR}) d(\text{FPR})$$

Where:

$$\text{TPR} = \frac{TP}{TP+FN} \quad \text{FPR} = \frac{FP}{FP+TN}$$

These metrics provide a detailed understanding of model performance, particularly in medical diagnosis where minimizing false negatives is critical.




Fig. 1. Proposed CNN model with four convolutional layers, pooling, and dropout.

The architecture is specifically designed to balance accuracy and efficiency, making it suitable for deployment in resource-constrained healthcare environments.

3.6 Transfer Learning

Pre-trained models VGG16, ResNet50, and EfficientNet-B0 are used for transfer learning to optimize performance and decrease training times. The rich feature

representations created by these massive training set models can also be utilized in the detection of pneumonia.

There are two strategies used in transfer learning:

- Feature extraction – the use of frozen convolutional layers with only the classifier trained.
- Fine-tuning – some of the deeper layers are re-trained specifically for that domain.

Transfer learning can increase convergence speed, improve accuracy, and help address problems that arise from having a small medical data set.

3.7 Model Training

To use the Adam optimizer to train the model, we implemented an adaptive learning rate to allow for faster convergence. A learning rate of 0.001 has been used to provide stable updates to the model's weights.

A binary cross-entropy loss function was chosen as the loss function for the classification task. Key hyperparameters include: Batch Size: 16–32; Epochs: 20–50; Dropout Rate: 0.3–0.5.

To help reduce the risk of overfitting the model, early stopping based on validation loss was applied and learning rate scheduling was used to adjust the optimization process during training.

3.8 Implementation Details

Use of TensorFlow and Keras to implement the model in an environment equipped with Python. The training process is enhanced by utilizing GPU acceleration, thereby reducing processing time. Batch processing techniques based on asynchronous generation of data and real-time data augmentation have both contributed significantly toward efficiency in the overall pipeline.

To aid in monitoring model development and optimizing the ultimate selection of the best models during training, the use of model checkpoints and various forms of model performance visualization (e.g., use of TensorBoard) will also be an integral component of this effort. Ultimately, utilizing unseen target data for evaluation will allow us to determine how well our model generalizes.

4 Results and Interpretation

4.1 Experimental Results

The experimental results show that the developed CNN model performs consistently well on the test dataset. It achieves an accuracy of 95.8%, along with high precision (96.5%) and recall (94.9%), indicating that the model can correctly identify both pneumonia and normal cases with good reliability. The specificity

(96.2%) and F1-score (95.7%) further demonstrate that the model maintains a good balance between detecting positive cases and avoiding false predictions. An AUC value of 0.97 also reflects its strong ability to distinguish between the two classes.

The training and validation curves indicate smooth and stable learning, with very little sign of overfitting. Both accuracy and loss trends show consistent improvement during training. The confusion matrix confirms that most cases are correctly classified, while the ROC curve further supports the model's effectiveness in separating pneumonia from normal images across different thresholds.

Metric	Value
Accuracy	95.8%
Precision	96.5%
Recall (Sensitivity)	94.9%
Specificity	96.2%
F1-Score	95.7%
AUC-ROC	0.97




Fig. 2. Training and validation accuracy curve showing steady convergence and minimal overfitting.




Fig. 3. Training and validation loss curve indicating stable learning behavior.




Fig. 4. Confusion matrix demonstrating correct classification of most pneumonia and normal cases.




Fig. 5. ROC curve with $AUC \approx 0.97$ indicating excellent discriminative performance.

4.2 Comparative Analysis

The proposed model's performance was evaluated alongside standard models—(VGG16, ResNet50, InceptionV3, and EfficientNetB0)—under identical experimental conditions. The key parameters that will be compared across all models are: (1) accuracy; (2) computational efficiency; and (3) robustness.

The results show that the proposed model delivers a good balance between performance and efficiency by yielding an accuracy that is competitive to other models while also providing a lower computational burden than the other models, Table 3. This makes the proposed model a good candidate for use in real-time clinical applications—especially where resources are limited.

Table 3. Performance comparison of the proposed CNN model with standard deep learning architectures for pneumonia detection

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	AUC	Complexity
Proposed CNN	95.8	96.5	94.9	95.7	0.97	Low
VGG16 [31]	94.2	94.2	92.5	93.3	0.94	High
ResNet50 [32]	95.1	95.1	93.8	94.4	0.95	High
InceptionV3 [33]	95.6	95.3	94.1	94.7	0.96	Very High
EfficientNet-B0 [34]	96.0	96.0	94.5	95.2	0.96–0.97	Moderate

The proposed CNN architecture yields results competitive with those of state-of-the-art deep learning architectures, with significantly less computational complexity.

- The proposed model has an accuracy of 95.8%, which is competitive with deeper architectures such as ResNet50 and InceptionV3.
- Compared to VGG16, which contains a very large number of parameters, the proposed model has a lightweight structure that allows for faster training and inference.
- While EfficientNet-B0 is designed for optimizing scalability, the proposed CNN achieves similar accuracy with a simplified architecture and lower resource requirements.
- The high AUC (0.97), indicating strong discriminatory power, is competitive with modern hybrid models that have been previously reported.

4.3 Interpretation of Results

The results indicate that the developed CNN model performs effectively while remaining computationally efficient, making it suitable for real-world applications. The model achieves an accuracy of 95.8%, showing its ability to reliably differentiate between pneumonia-affected and normal chest X-ray images. Its recall (sensitivity) of 94.9% highlights that most pneumonia cases are correctly identified, which is especially important in medical diagnosis where missing a positive case can have serious consequences.

With a precision of 96.5%, the model also produces very few false positives, reducing the likelihood of unnecessary medical interventions. The F1-score of 95.7% reflects a good balance between precision and recall, indicating stable performance even when class distributions vary. Additionally, the AUC-ROC value of 0.97 demonstrates strong capability in distinguishing between the two classes across different decision thresholds.

When compared to more complex deep learning models such as ensemble or hybrid approaches, the proposed CNN achieves similar accuracy while requiring significantly fewer computational resources. Its lightweight structure allows faster inference and lower hardware requirements, making it well-suited for deployment in real-time clinical environments, particularly where resources are limited.

The training process also shows stable and consistent learning behavior. Both training and validation accuracy improve steadily and converge after around 20 epochs, suggesting effective model learning. The small gap between training and validation accuracy (around 1–2%) indicates good generalization. Likewise, the loss curves decrease smoothly without major fluctuations, confirming that overfitting is well controlled. Minor variations in validation loss during later stages are expected due to data variability and do not affect overall performance. Techniques such as dropout (0.5) and early stopping play an important role in preventing overfitting by ensuring that the model does not simply memorize the training data. Overall, these observations suggest that the model maintains a good balance between learning capacity and generalization, leading to reliable performance on unseen data.

5 Conclusion

In this study, we proposed an optimized lightweight Convolutional Neural Network (CNN) architecture for automated pneumonia classification from chest X-ray images. The proposed model is developed with the aim to achieve a trade-off between high diagnostic accuracy and low computational complexity which can be used for practical healthcare applications particularly in resource-constrained environments. With the help of effective preprocessing techniques, data augmentation, dropout regularization, and a compact CNN structure, the model was able to learn robust features to distinguish pneumonia-affected lungs from normal cases.

The experimental evaluation showed the good performance of the proposed approach with an accuracy of 95.8%, precision of 96.5%, recall of 94.9%, F1-score of 95.7% and an AUC-ROC value of 0.97. The results show that the model can reliably detect pneumonia cases, while also minimizing false positives and false negatives.

Since the architecture is light-weighted, it can be easily deployed in real-time clinical systems like healthcare centers with limited hardware resources and also can provide faster inference. Moreover, the stable training behavior and the reduced overfitting indicate the good generalization capability of the model on unseen data. In conclusion, the proposed CNN framework offers a practical and efficient AI-assisted diagnostic solution that can assist radiologists and healthcare professionals in the early detection of pneumonia.

Therefore, future work can involve extending the model for multi-class classification of lung diseases, including attention mechanisms to improve feature localization, and validating the approach on larger and more diverse clinical datasets to improve robustness and real-world applicability.

References

1. Mahi, A. B. S., Eshita, F. S., & Helaly, T. (2024). An efficient convolutional neural network for pneumonia recognition from chest X-ray images. *Proceedings of the 6th International Conference on Electrical Engineering and Information & Communication Technology (ICEEICT)*, 202–207. <https://doi.org/10.1109/ICEEICT62016.2024.10534356>
2. Nallamala, S. H., Polireddy, Y., Peram, V. N. T., Narra, D., Maddali, P. K. T., & Botlagunta, H. N. K. (2025). Pneumonia detection using CNN: A deep learning approach. *i-manager's Journal on Information Technology*, 14(1), 17–25. <https://doi.org/10.26634/jit.14.1.21936>
3. Gupta, A., Mishra, P., Tiwari, N. K., Bajpai, A., & Yadav, S. (2024). Pneumonia detection in chest X-ray images using efficient and lightweight convolutional neural networks. *Proceedings of the IEEE Students Conference on Engineering and Systems (SCES)*, 1–6. <https://doi.org/10.1109/SCES61914.2024.10652501>
4. Chatterjee, R., Chatterjee, A., & Halder, R. (2021). An efficient pneumonia detection from chest X-ray images. In M. Prateek et al. (Eds.), *Proceedings of International Conference on Machine Intelligence and Data Science Applications*. Springer. https://doi.org/10.1007/978-981-33-4087-9_63

5. Suryaa, V. S., Annie, R. A. X., & Aiswarya, M. S. (2021). Efficient DNN ensemble for pneumonia detection in chest X-ray images. *International Journal of Advanced Computer Science and Applications*, 12(10). <http://dx.doi.org/10.14569/IJACSA.2021.0121084>
6. George, S., N. G., Saifullah, S., & Joe, I. R. P. (2022). An efficient neural network model for pneumonia detection. *Proceedings of the International Conference on Computing, Communication, Security and Intelligent Systems (IC3SIS)*, 1–7. <https://doi.org/10.1109/IC3SIS54991.2022.9885488>
7. Singh, G., Guleria, K., & Sharma, S. (2024). An efficient convolutional neural network model for pneumonia detection in chest X-ray images. *Proceedings of the International Conference on Internet of Everything, Microwave, Embedded, Communication and Networks (IEMECON)*, 1–5. <https://doi.org/10.1109/IEMECON62401.2024.10845994>
8. Budhya, S. M., & Harish, B. G. (n.d.). Detection of pneumonia using convolutional neural networks. <https://doi.org/10.56726/IRJMETS60519>
9. Lafraxo, S., & El Ansari, M. (2022). Regularized convolutional neural network for pneumonia detection through chest X-rays. In J. Kacprzyk et al. (Eds.), *Advanced intelligent systems for sustainable development*. Springer. https://doi.org/10.1007/978-3-030-90633-7_76
10. Saha, A. K., & Rahman, M. (2022). An efficient deep learning approach for detecting pneumonia using convolutional neural network. In S. Shakya et al. (Eds.), *Sentimental analysis and deep learning*. Springer. https://doi.org/10.1007/978-981-16-5157-1_5
11. Varshni, D., Thakral, K., Agarwal, L., Nijhawan, R., & Mittal, A. (2019). Pneumonia detection using CNN-based feature extraction. *Proceedings of the IEEE International Conference on Electrical, Computer and Communication Technologies (ICECCT)*, 1–7. <https://doi.org/10.1109/ICECCT.2019.8869364>
12. Saranya, P., Thapliyal, T., & Bajpai, A. (2020). Detection of pneumonia from X-ray images using convolutional neural networks. IOS Press. <https://doi.org/10.3233/APC200192>
13. Islam, M., & Pathari, F. J. (2023). Detection of pneumonia from chest X-ray images using convolutional neural network (CNN). *Proceedings of the International Conference on Applied Artificial Intelligence (ICAPAI)*, 1–8. <https://doi.org/10.1109/ICAPAI58366.2023.10194027>
14. Hashmi, M. F., Katiyar, S., Keskar, A. G., Bokde, N. D., & Geem, Z. W. (2020). Efficient pneumonia detection in chest X-ray images using deep transfer learning. *Diagnostics*, 10(6), 417. <https://doi.org/10.3390/diagnostics10060417>
15. Rahman, T., Chowdhury, M. E. H., Khandakar, A., Islam, K. R., Islam, K. F., Mahbub, Z. B., Kadir, M. A., & Kashem, S. (2020). Transfer learning with deep convolutional neural network (CNN) for pneumonia detection using chest X-ray. *Applied Sciences*, 10(9), 3233. <https://doi.org/10.3390/app10093233>
16. Sharma, H., Jain, J. S., Bansal, P., & Gupta, S. (2020). Feature extraction and classification of chest X-ray images using CNN to detect pneumonia. *Proceedings of the International Conference on Cloud Computing, Data Science & Engineering*, 227–231. <https://doi.org/10.1109/Confluence47617.2020.9057809>
17. Labhane, G., Pansare, R., Maheshwari, S., Tiwari, R., & Shukla, A. (2020). Detection of pediatric pneumonia from chest X-ray images using CNN and transfer learning. *Proceedings of the International Conference on Emerging Technologies in Computer Engineering*, 85–92. <https://doi.org/10.1109/ICETCE48199.2020.9091755>
18. Aljawarneh, S. A., & Al-Quraan, R. (2025). Pneumonia detection using enhanced convolutional neural network model on chest X-ray images. *Big Data*, 13(1), 16–29. <https://doi.org/10.1089/big.2022.0261>

19. Kumar, K. S., & Sharma, J. (2025). An all-inclusive CNN-based pneumonia detection method. *Proceedings of the International Conference for Emerging Technology (INCET)*, 1–5. <https://doi.org/10.1109/INCET64471.2025.11140921>
20. Haque, A., et al. (2023). Detecting pneumonia from X-ray images using deep convolutional neural network. *Proceedings of the International Conference on Big Data Analytics and Practices (IBDAP)*, 1–6. <https://doi.org/10.1109/IBDAP58581.2023.10271996>
21. El Asnaoui, K., Chawki, Y., & Idri, A. (2021). Automated methods for detection and classification pneumonia based on X-ray images using deep learning. In Y. Maleh et al. (Eds.), *Artificial intelligence and blockchain for future cybersecurity applications*. Springer. https://doi.org/10.1007/978-3-030-74575-2_14
22. Ahmad, N., Qadri, I. A., Naeem, M., Akbar, F., Ahmed, N., & Rai, S. I. (2023). Effective detection of lung disease from X-ray images using convolutional neural networks. *Journal of Population Therapeutics and Clinical Pharmacology*, 30(18), 1975–1992. <https://doi.org/10.53555/jptcp.v30i18.3329>
23. Eido, W. M., & Yasin, H. M. (2025). Pneumonia and COVID-19 classification and detection based on convolutional neural network: A review. *Asian Journal of Research in Computer Science*, 18(1), 174–183. <https://doi.org/10.9734/ajrcos/2025/v18i1556>
24. Chouhan, V., Singh, S. K., Khamparia, A., Gupta, D., Tiwari, P., Moreira, C., Damaševičius, R., & 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. <https://doi.org/10.3390/app10020559>
25. Santhoshi, M., & Jyostna, J. (2023). An efficient approach of pneumonia detection using transfer learning models, RCNN and Fast RCNN. *Proceedings of the International Conference on Advances in Computation, Communication and Information Technology*, 1214–1219. <https://doi.org/10.1109/ICAICCIT60255.2023.10465721>
26. AlGhamdi, A. S. (2024). Efficient deep learning approach for the classification of pneumonia in infants from chest X-ray images. *Traitement du Signal*, 41(3), 1245–1262. <https://doi.org/10.18280/ts.410314>
27. Badlani, K., Sawal, S., Nilkute, M., Belekar, S., & Nilawar, A. (2022). Pneumonia detection through image classification using CNN. *International Journal of Next-Generation Computing*, 13(5). <https://doi.org/10.47164/ijngc.v13i5.931>
28. Kundu, R., Das, R., Geem, Z. W., Han, G.-T., & Sarkar, R. (2021). Pneumonia detection in chest X-ray images using an ensemble of deep learning models. *PLoS ONE*, 16(9), e0256630. <https://doi.org/10.1371/journal.pone.0256630>
29. An, Q., Chen, W., & Shao, W. (2024). A deep convolutional neural network for pneumonia detection in X-ray images with attention ensemble. *Diagnostics*, 14(4), 390. <https://doi.org/10.3390/diagnostics14040390>
30. Iparraguirre-Villanueva, O., Guevara-Ponce, V., Roque Paredes, O., Sierra-Liñan, F., Zapata-Paulini, J., & Cabanillas-Carbonell, M. (2022). Convolutional neural networks with transfer learning for pneumonia detection. *International Journal of Advanced Computer Science and Applications*, 13(9). <http://dx.doi.org/10.14569/IJACSA.2022.0130963>

31. Simonyan, K., & Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556.
32. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
33. Szegedy, C., et al. (2016). Rethinking the inception architecture for computer vision. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2818–2826.
34. Tan, M., & Le, Q. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. *Proceedings of the International Conference on Machine Learning*, 6105–6114.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

