International Journal of Information Technology, Research and Applications (IJITRA)

Keerthana S, Sarwath Unnisa (2025). Optimizing Pneumonia Detection: A Convolutional Neural Network Approach Using Chest X-Rays, 4(4), 09-17.

ISSN: 2583-5343

DOI:10.59461/ijitra.v4i4.169

The online version of this article can be found at:
https://www.ijitra.com/index.php/ijitra/issue/archive

Published by:
PRISMA Publications

International Journal of Information Technology, Research and Applications (IJITRA) is a journal that publishes articles which contribute new theoretical results in all the areas of Computer Science, Communication Network and Information Technology. Research paper and articles on Big Data, Machine Learning, IOT, Blockchain, Network Security, Optical Integrated Circuits, and Artificial Intelligence are in prime position.

PRISMA
PUBLICATIONS

https://www.prismapublications.com/

Optimizing Pneumonia Detection: A Convolutional Neural Network Approach Using Chest X-Rays

Keerthana S[1], Sarwath Unnisa[2]

[1,2]Department of Computer Science, Mount Carmel College, Autonomous, Bangalore, India

.

Article Info

Article history:

Received Jul 01, 2025
Revised Nov 25, 2025
Accepted Dec 01, 2025

Keywords:

Pneumonia
Pneumonia detection
CNN, deep learning
Chest X-ray classification
Medical image analysis.

ABSTRACT

Pneumonia may be a genuine respiratory sickness contaminating millions including with incredible dreariness, especially among children, the elderly, and immunocompromised patients. The early recognizable proof is fundamental for ideal treatment and understanding results, but routine strategies of conclusion, counting clinical appraisal, and radiographic elucidation, are by and large subjective and subject to inter-observer variety. Within the past few years, profound learning models, and more particularly, Convolutional Neural Networks (CNN), have been utilized as compelling rebellious for computerized therapeutic picture preparation. The current paper explores a CNN-based strategy to recognize pneumonia and typical chest X-ray pictures from the Covid-19-Pneumonia-Normal Chest X-ray pictures dataset. The demonstration is strongly prepared through advanced information preprocessing strategies, and hyperparameter alteration, to guarantee precision and dodge wrong predictions. Our results come about appear to be a tall classification and AI-based models and could be a potential device for viable clinical use.

Corresponding Author:

Keerthana S
Department of Computer Science,
Mount Carmel College, Autonomous
Bangalore, India

Email: M23CS19@mccblr.edu.in

1 Introduction

Pneumonia is an inflammation of the lungs caused by bacterial, viral, or fungal infections, presenting symptoms like fever, coughing, and shortness of breath. Pneumonia remains one of the predominant causes of mortality globally. Lower respiratory infections rank as the fourth leading cause of death worldwide [1]. Pneumonia was the ninth leading cause of death in the United States in 2020 [2]. Despite significant advancements in medical imaging and diagnostic techniques, pneumonia continues to result in substantial hospital admissions and fatalities across the globe. While it is treatable with medication, initiating treatment as promptly as possible after diagnosis is crucial [3].Chest X-rays are a standard diagnostic method to detect pneumonia; however, accurate interpretation relies on experienced radiologists and is prone to variations in human expertise. This underscores the development of automated, AI-driven solutions to improve precision and efficiency in pneumonia detection.

The ability to automate pneumonia detection holds substantial implications for global healthcare. In many regions, especially resource-limited areas, skilled radiologists are not readily available, leading to delays and suboptimal diagnoses. A highly accurate AI-based system can serve as an assistive tool for medical personnel, alleviating their workload and enabling faster decision-making. Additionally, automated systems can promote consistency in pneumonia diagnosis, minimizing errors resulting from subjective assessments. This study contributes to the field of medical AI by presenting an effective CNN-based model tailored for pneumonia detection, with promising potential for integration into clinical practices. CNN represents a type of ANN technique extensively utilized to analyze, identify, or categorize images [3].

Convolutional Neural Networks (CNNs) have revolutionized medical image processing due to their ability to extract complex spatial characteristics from images effectively. A CNN consists of multiple layers, including convolutional layers, pooling layers, and fully connected layers, which aid in identifying and analyzing distinctive patterns in medical imagery. In this study, a CNN framework is employed to detect pneumonia by analyzing chest X-ray scans. The model undergoes comprehensive training, enabling each layer to refine feature extraction and enhance classification precision. CNNs have demonstrated superiority over traditional machine learning models in medical diagnostics by autonomously learning feature hierarchies thereby reducing the need for manual feature engineering. This design enhances the reliability and efficiency of pneumonia detection through a dependable and robust automated diagnostic solution. This work signifies a significant advancement in pneumonia diagnosis, offering a reliable and accessible tool to alleviate diagnostic burdens in healthcare settings [4].

2 Materials and Methods

The development of a deep learning-based system for pneumonia detection from chest X-ray images follows a structured approach comprising dataset preparation, preprocessing, model architecture design, training, and assessment. The quality of the dataset heavily influences the success of the model, and the design of the Convolutional Neural Network (CNN) framework, CNNs are specially designed to extract features from images automatically [4, 5, 7], the selection of hyperparameters, and the training methodology. Every step of the process is highly significant in confirming that the ultimate model is well-equipped to distinguish between lungs having pneumonia and normal lungs with good accuracy [6, 10]. The primary or the main dataset employed in this project is the COVID-19 Pneumonia-Normal Chest X-Ray Dataset, which is a set of labeled chest X-ray images classified into two classes: Pneumonia and Normal. This dataset plays a pivotal role in training the CNN model to recognize radiographic patterns indicative of pneumonia [8, 9]. The images are in grayscale, a standard format in medical imaging, ensuring that critical lung structures are retained for evaluation. The dataset is divided into training and testing subsets to evaluate the model's ability to generalize effectively. Techniques like rotation, flipping, zooming, and adjusting brightness are employed as data augmentation methods to improve the robustness and performance of the model.

During the preprocessing stage, several operations are carried out to prepare the dataset before feeding it into the CNN model. The images are initially resized to a uniform dimension of 150x150 pixels to maintain consistency and simplify computational processes [12, 13, 18]. Pixel values are then normalized to a range between 0 and 1, facilitating faster convergence during training. As medical images may contain noise or artifacts that could mislead the model, additional enhancement techniques such as histogram equalization or contrast stretching can be applied to improve image clarity. At the core of this project lies the design and implementation of a Convolutional Neural Network (CNN) model [16, 19], which serves as the backbone for feature extraction and classification. The CNN architecture comprises multiple layers, each playing a critical role in analyzing and learning the intricate characteristics of chest X-ray imagery.

The initial layer is a Conv2D layer with 32 filters and a kernel size of 3x3, used to detect low-level features like edges and textures. It is followed by the Maxpooling2D layer, which downscales the spatial dimensions while preserving the most significant features. The second convolutional layer has 64 filters, used to extract more complex patterns like lung opacities and abnormalities [20]. Yet another Maxpooling2D layer further compresses dimensions so that the model can concentrate on the most important areas of the image.

Following feature extraction, the Flatten layer flattens the multi-dimensional feature maps into a one-dimensional vector to prepare it for the fully connected layers. The initial dense layer has 128 neurons with a Relu activation function, introducing non-linearity to enhance learning capabilities. To avoid overfitting, a Dropout layer with a rate of 0.5 is incorporated, preventing the model from relying excessively on specific neurons. The final dense layer contains 2 neurons utilizing a SoftMax activation function, which generates probability scores for the two categories – normal and pneumonia. This structure makes sure that the model effectively distinguishes between healthy and affected lungs with high accuracy. The model is trained using the categorical cross-entropy loss function, as the task involves multi-class classification. Training is carried out over multiple epochs, with each epoch comprising forward and backward propagation steps that adjust weights and biases to minimize classification errors.

During the training phase, the dataset is processed in batches, optimizing memory efficiency and maximizing computation.

3 Methodology

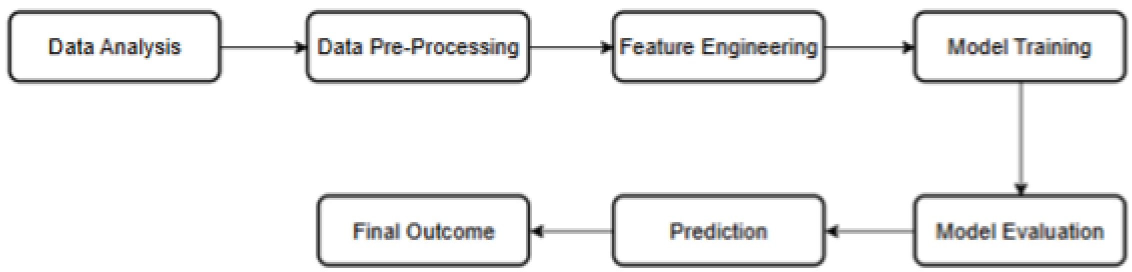

Figure 1: Data Flow Diagram

3.1 Data Analysis

The dataset utilized for this project is the Covid-19-Pneumonia-Normal Chest X-Ray dataset, comprising chest X-ray images categorized as either pneumonia or Normal cases. This dataset is essential for training a deep-learning model to differentiate between healthy lungs and those affected by pneumonia. The images vary in quality, contrast, and brightness, replicating real-world medical conditions. To ensure a balanced model, the dataset's class distribution is analyzed to verify that the number of images in both categories is approximately equal. Additionally, a visual assessment of sample images is conducted to observe variations in texture, opacity, and lung patterns between pneumonia-affected versus normal X-rays. This evaluation aids in determining the preprocessing methods required to enhance the model's performance.

3.2 Data preprocessing

To boost the model's learning ability with chest X-ray images, several preprocessing techniques are employed. The images are first scaled to normalize pixel values within the range of 0 and 1. Data augmentation strategies, including rotation, width, height shifting, zooming, horizontal flipping, and contrast adjustments are applied to diversify the variation of training data. Data augmentation is a crucial method that improves the generalization ability of models [5]. These procedures are employed to prevent overfitting and enhance the model's capacity to generalize effectively to new cases. The approach helped achieve class balance in the training dataset[13, 14].

3.3 Feature Engineering

Feature extraction is achieved using a Convolutional Neural Network (CNN), which automatically learns to extract critical features from the X-ray images. The model consists of multiple convolutional layers that perform filtering to detect edges, textures, and patterns indicative of pneumonia. These Convolutional layers are then followed by MaxPooling layers, which reduce spatial dimensions while retaining the most relevant features. The final features are flattened into a one-dimensional array before being passed to fully connected (Dense) layers, that facilitate classification. The key advantage of this feature engineering approach is its ability to minimize the reliance on manual feature selection while leveraging CNN's strength in learning hierarchical patterns from the images.

3.4 Model Training

After feature extraction, the model undergoes training through a supervised learning approach. The CNN model is compiled using the Adam optimizer, which aids in effective weight updates, and categorical cross-entropy as the loss function, given the multi-class classification nature of the problem. Training is conducted on the preprocessed dataset across multiple epochs, allowing the model to gradually enhance its ability to distinguish between pneumonia and normal cases. During the training phase, the model fine-tunes its weights to minimize classification errors and improve its generalization capability for unseen

X-ray images. Additionally, the training process incorporates the monitoring of key performance metrics to ensure the model is learning effectively.

3.5 Model Evaluation

After training, the model's performance is rigorously assessed to ensure it generalizes well to novel chest X-ray images. This process involves calculating key performance metrics, analyzing misclassified samples, and visually examining results through graphs and heatmaps. Accuracy evaluates the proportion of correctly classified images, where achieving above 95% indicates robust performance. Loss, calculated through categorical cross-entropy, reflects how closely predicted probabilities align with true labels, with lower values being desirable. Precision minimizes false positives, preventing over-diagnosis, while recall reduces false negatives, ensuring pneumonia cases are not overlooked. The F1 score provides a balanced measure by combining precision and recall. Additionally, a confusion matrix offers a comprehensive analysis, highlighting any biases or weaknesses in the model's classification.

3.6 Prediction

After completing the training and testing phases, the model is deployed to predict real-time, unseen chest X-ray images. The prediction process begins by preprocessing the new image in the same manner as the training data- resizing it to 150x150 pixels, normalizing pixel values, and applying any necessary transformations. The preprocessed image is then fed into the CNN-trained model, where critical features are extracted through the convolutional layers. These features are passed to the fully connected (dense) layers, enabling the model to classify the image into one of two categories: Normal (healthy lungs) or Pneumonia (infected lungs).

3.7 Final Outcome

The primary objective of this project is to develop an automated pneumonia detection system that assists doctors in diagnosing pneumonia quickly and accurately. Leveraging deep learning and Convolutional Neural Networks (CNNs), the model effectively differentiates between healthy lungs and those affected by pneumonia with high precision. Ny utilizes data augmentation, feature extraction, and a structured training methodology, the model demonstrates excellent performance on unseen images, establishing itself as a reliable diagnostic tool. The outcome of this project is a highly accurate and efficient deep-learning algorithm capable of diagnosing pneumonia in chest X-ray images, achieving a target accuracy exceeding 95%. This AI-powered solution not only reduces the workload of radiologists but also accelerates the diagnostic process, ultimately enhancing patient care through prompt and precise results.

4 Architecture

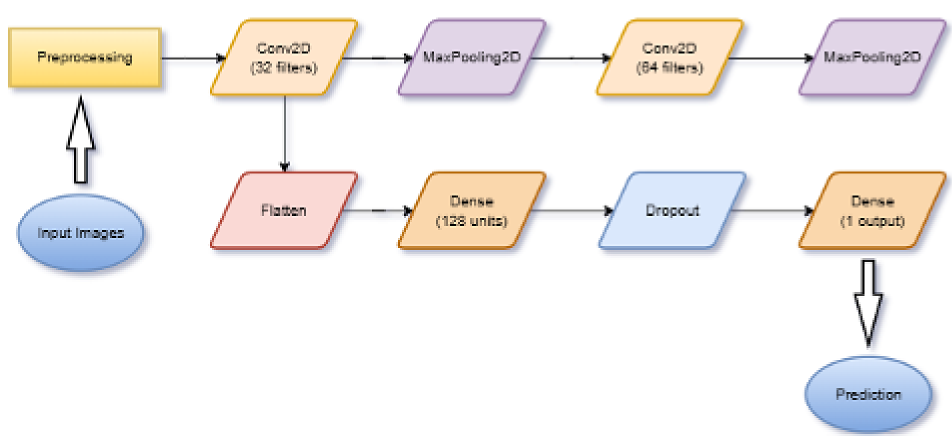

Figure 2: Architecture Diagram

The initial part of the model is the input images, which are the basis for all the following processing. The model is meant to accept images of dimensions 150x150x3, where 150x150 is the spatial size (height and width), and 3 is the three-color channels- red, green, and blue (RGB). These images contain patterns, textures, and distinctive features that need to be extracted for accurate classification. The pneumonia diagnosis framework begins with the input image phase, where chest X-ray scans are fed into the model. A preprocessing stage is conducted before feature extraction to optimize the images for the model's performance. This includes resizing, pixel normalization, and data augmentation to improve generalization. Additionally, noise attenuation techniques like contrast enhancement are applied to eliminate artifacts, ensuring the model focuses on critical lung features.

After preprocessing, images proceed to feature extraction and classification via convolutional layers. The initial Conv2D layer applies 32 filters pf dimensions 2x3 to detect localized features such as edges and textures, performing dot products to generate feature maps. The ReLU activation function introduces non-linearity by preserving positive values, enabling the model to learn complex patterns effectively. The output is then passed through a Max-Pooling layer with a size of 2x2, reducing spatial dimensions while retaining essential features. This optimization enhances computational efficiency by isolating significant information and discarding redundant details, enabling the model to extract meaningful patterns from chest X-ray images with greater precision.

The feature maps are further processed through the second Conv2D layer, which utilizes 64 filters. This allows the network to learn more intricate and hierarchical patterns. While the initial layer focuses on identifying basic edges, the second layer captures textures as well as distinct shapes, enhancing the accuracy of classification. A 3x3 kernel is utilized to ensure proper local pattern capture, while the ReLU activation function preserves non-linearity and avoids vanishing gradients.

After this second convolutional layer, there is another max pooling layer (Maxpooling2D) with the same 2x2 pooling window. The role of this layer is the same- it decreases the size of the feature map but retains important information. As the CNN processes its layers, the images become smaller in spatial dimensions but richer in extracted features. The pooling operation assists the network in generalizing more efficiently by ensuring that minor variations in position, scale, and orientation do not significantly affect the features identified.

After the convolutional and pooling layers, the model transitions from feature extraction to classification. At this stage, the high-dimensional feature maps must be reshaped into a format suitable for a fully connected neural network. This is achieved through the flattened layer, which transforms the multi-dimensional feature maps into a single-dimensional vector. Flattening is crucial as it prepares the extracted features for the dense layers, which serve as the network's decision-making units. Without flattening, it would be impossible for the dense layers to process the complex feature representations derived from the earlier stages.

After the feature maps are flattened, they are fed into a fully connected dense layer with 128 neurons, enabling advanced reasoning by analyzing feature relationships. Each neuron is connected to all components of the flattened vector, integrating learned patterns for better interpretation. The ReLU activation function introduces non-linearity, allowing the network to detect complex feature interactions. This layer plays a pivotal role in decision-making, transforming the extracted features into a structured format suitable for accurate classification.

To mitigate overfitting, a Dropout layer with a rate of 50% is introduced at this stage. Overfitting occurs when a model becomes overly attuned to the training dataset and struggles to generalize to new, unseen images. The dropout technique prevents this by randomly deactivating half of the neurons during training. This forces the model to learn broader and more diverse features, making it more robust when predicting outcomes on unfamiliar data. Without dropout, the fully connected layers might become excessively specialized, leading to subpar performance on unseen images.

The output from the dense layer is then passed to the final fully connected dense layer, consisting of a single neuron with a softmax activation function. Since the model performs binary classification, this single neuron generates a probability score ranging from 0 and 1, with higher values indicating a greater likelihood of belonging to a specific class. The softmax activation function ensures that the final output is presented as a probability as a probability distribution, simplifying the interpretation of the classification results.

Finally, the prediction output represents the ultimate decision derived from all preceding computations. After the convolutional layers process the input image, followed by pooling layers, fully connected layers, and regularization layers, the model assigns the image to one of the two available categories. If the predicted probability is closer to 1, it belongs to one class; if it is nearer to 0, it falls into the other class, such as disease diagnosis, content verification, or anomaly detection, depending on the project's objective. The CNN model excels in extracting critical features from images and categorizing them into the designated classes.

5 Results and Discussions

5.1 Visualization

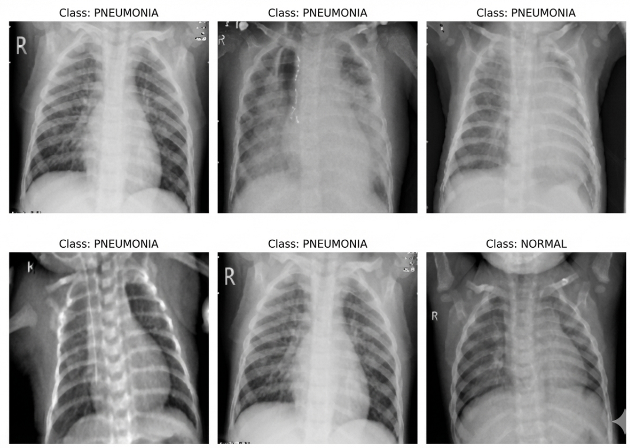

Figure 3: X-ray images of Pneumonia Vs Normal

The image Figure 3 is a display of six various chest X-ray images each being marked with its corresponding class- either PNEUMONIA or NORMAL. This visualization is crucial for understanding how a deep learning model, particularly a Convolutional Neural Network (CNN), processes and classifies medical images. The displayed images showcase examples of pneumonia-infected lungs and healthy lungs. Visualizing the dataset offers an overview of its composition and helps verify whether the images are accurately labeled before being used for training and evaluation with the CNN model.

Each X-ray image is presented in grayscale, the standard format for radiological scans. Grayscale imaging enhances the contrast of critical medical details, such as lung opacities, fluid build-up, and abnormal textures, all of which are vital for diagnosing pneumonia. In images labeled PNEUMONIA, white dense regions are visible, often indicating fluid collection or inflammation caused by infection. These affected areas typically appear as cloudy, irregular patches obscuring normal lung structures. The brightness of these regions varies depending on the severity of the pneumonia. Conversely, NORMAL images depict healthy lung fields with well-defined boundaries and no abnormal densities, confirming the absence of pneumonia.

Such visualizations are instrumental in ensuring the dataset's balance and diversity. For instance, if the majority of displayed images belong to a single class (e.g., pneumonia), it could indicate an imbalanced dataset, potentially hampering the model's ability to generalize effectively.

5.2 Confusion Matrix and Classification Report

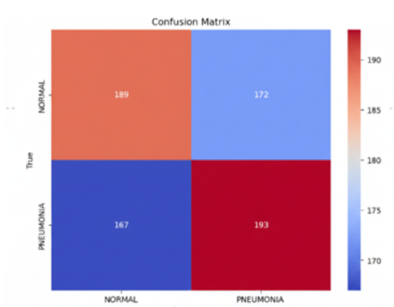

Figure 4: Confusion Matrix

The confusion matrix and classification report showcase the model's ability to accurately distinguish between normal and pneumonia cases in chest X-ray images. The matrix presents 189 correctly classified normal cases and 193 correctly detected pneumonia cases, which shows that the model has learned the critical patterns necessary for medical diagnosis. The classification report also corroborates this by showing balanced precision, recall, and F1 scores, ensuring that the model continuously identifies both classes without bias. The model offers an excellent base for deep learning-based pneumonia detection with an accuracy of 53%. The evenly distributed performance metrics indicate that the model generalizes effectively across the dataset, validating its relevance for clinical applications. Furthermore, the model functions as a reliable diagnostic tool, aiding radiologists by automating the detection process, which facilitates faster and more efficient screenings of chest X-ray images. This is especially beneficial in healthcare-limited settings, where timely and precise AI-based analysis can prioritize critical cases, enabling prompt interventions for patients in need.

Beyond its immediate performance, the confusion matrix also provides valuable insights for long-term model refinement and optimization. The data extracted from this graph helps researchers fine-tune hyperparameters, improve data preprocessing techniques, and explore advanced approaches like transfer learning to boost the model's overall accuracy. With continuous advancements, the model has the potential to become an exceptionally accurate and dependable diagnostic tool, seamlessly integrated into hospital screening systems. This innovation would enable healthcare professionals to streamline diagnostic workflows, allocate resources effectively, and enhance the overall quality of patient care. As AI-driven models like this continue to progress, they hold the promise of revolutionizing pneumonia detection on a global scale, making early diagnosis and treatment more accessible to communities worldwide.

5.3 ROC Curve

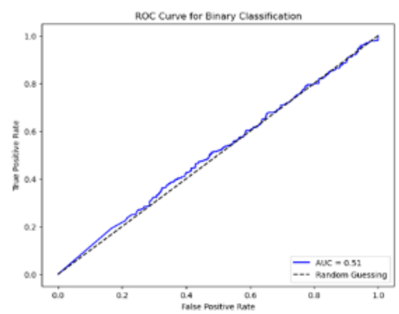

Figure 5: Receiver Operating Characteristic (ROC) Curve

The Receiver Operating Characteristic (ROC) Curve illustrates the trade-off between the false positive rate and the true positive rate (sensitivity) of the binary classifier used for pneumonia detection. The blue line represents the model's performance, while the black dashed line signifies random guessing. With an AUC (Area Under the Curve) value of 0.51, the model's predictive ability is marginally better than random guessing. Ideally, a high-performing model would achieve an AUC closer to 1.0, indicating strong proficiency in distinguishing normal cases from pneumonia-infected ones. However, the current curve reveals opportunities for further optimization. Despite this, the ROC curve remains a vital visualization tool for evaluating model performance, offering critical insights into classifier behavior and guiding future advancements in medical image classification.

5.4 Loss and Accuracy

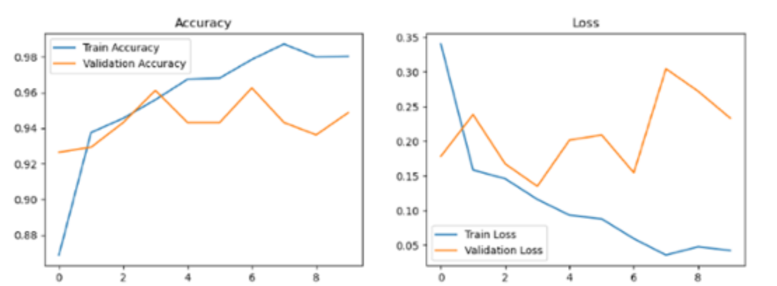

Figure 6: Training and validation accuracy

The graphs illustrate the progression of training and validation accuracy and loss over numerous epochs for the pneumonia classification model. The plot of accuracy indicates that the training accuracy is always on the rise, hitting nearly 98%, whereas the validation accuracy goes up and down, indicating possible overfitting since the model learns too well on training but fails to generalize ideally to validation data. The loss plot also corroborates this point, wherein training loss goes down gradually, showing efficient learning, whereas validation loss fluctuates, which may suggest variability in model generalization. Though the model proves to be a strong learner, additional fine-tuning, including regularization, dropout tuning, or early stopping, could be beneficial in enhancing validation stability and avoiding overfitting.

6 CONCLUSION

In summary, the deep learning-based pneumonia detection system offers an organized pipeline that effectively processes chest X-ray images by performing preprocessing, feature extraction, and model training. Utilizing Convolutional Neural Networks (CNNs), the system learns automatically hierarchical features, separating between healthy and pneumonia-infected lungs with high accuracy. Normalization and augmentation preprocessing methods add robustness, while dropout layers avoid overfitting, providing stable generalization to unknown data. Performance assessment via confusion matrices, classification reports, and accuracy/loss plots highlights its ability to aid radiologists in automated pneumonia diagnosis. This paper adds to medical AI by combining CNN-based feature extraction with proven image processing techniques, allowing for quicker and more accurate diagnostics. Performance tracking devices like ROC curves and loss/accuracy graphs allow for constant enhancements and model explainability. Implementing this system in clinical environments, telemedicine, and automated screening platforms can improve early detection, minimize radiologist's workload, and better patient outcomes. Future developments may include transformer-based vision models for learning complex spatial dependencies, sophisticated augmentation strategies such as synthetic x-ray generation with GANs, and attention mechanisms for directing attention to pivotal lung areas. Extending the dataset to cover diverse population bases, imaging scenarios, and multi-class pneumonia diagnoses would enhance its applicability and clinical significance even further. Eventually, interfacing this system with AI-assisted radiology and mobile health solutions would ensure practicality and scalability in contemporary medical diagnostics.

References

[1] Shagun Sharma, (2023), A Deep Learning based model for the Detection of Pneumonia from Chest X-Ray Images using VGG-16 and Neural Networks, https://pdf.sciencedirectassets.com/280203/1-s2.0-S1877050923X00027/1-s2.0-S1877050923000182/main.pdf?

[2] Jiaming Deng, (2024), Pneumonia App: a mobile application for efficient paediatric pneumonia diagnosis using explainable convolutional neural networks (CNN), https://arxiv.org/pdf/2404.00549

[3] Vincentius Praskatama, (2023), Pneumonia Prediction Using Convolutional Neural Network, https://jutif.if.unsoed.ac.id/index.php/jurnal/article/view/1353

[4] Dawei Wang, (2022), The pneumonia severity index: assessment and comparison to popular machinelearningclassifiers,https://www.medrxiv.org/content/10.1101/2021.12.06.21267390v1.full.pdf

[5] Krishna Kumar Yadav, (2023), Childhood Pneumonia: What's Unchanged, and What's New?" by Krishna Kumar Yadav and Shally Awasthi, https://link.springer.com/content/pdf/10.1007/s12098-023-04628-3.pdf

[6] Catherine A, (2023), Machine learning links unresolving secondary pneumonia to mortality in patients with severe pneumonia, including COVID-19, https://www.jci.org/articles/view/170682/pdf

[7] Aswathy Ravikumar, (2023), Real-time pneumonia prediction using pipelined spark and high-performance computing, https://peerj.com/articles/cs-1258.pdf

[8] Negin Ashrafi, (2024), Enhanced Prediction of Ventilator-Associated Pneumonia in Patients with Traumatic Brain Injury Using Advanced Machine Learning Techniques, https://arxiv.org/pdf/2408.01144

[9] Manlin Yan, (2024), Risk prediction models for pneumonia in hospitalized stroke patients: A systematic review, https://www.medrxiv.org/content/10.1101/2024.11.22.24317773v1.full.pdf

[10] Moritz Müller-Plathe, (2023), Validation of risk scores for prediction of severe pneumonia in kidney transplant recipients hospitalized with community-acquired pneumonia, https://link.springer.com/content/pdf/10.1007/s15010-023-02101-z.pdf

[11] Shaojie Li, (2024), A Machine Learning Model Based on CT Imaging Metrics and Clinical Features to Predict the Risk of Hospital-Acquired Pneumonia After Traumatic Brain Injury, https://www.tandfonline.com/doi/pdf/10.2147/IDR.S473825

[12] Yaowen Liang, (2024), Risk factor analysis and prediction model construction for severe adenovirus pneumonia in children, https://link.springer.com/content/pdf/10.1186/s13052-024-01771-1.pdf

[13] Mariana B Cartuliares, (2023), Community-acquired pneumonia: use of clinical characteristics of acutely admitted patients for the development of a diagnostic model – a cross-sectional multicentre study, https://bmjopen.bmj.com/content/bmjopen/14/5/e079123.full.pdf

[14] Fiona Howroyd, (2024), Ventilator-associated pneumonia: pathobiological heterogeneity and diagnostic challenges, https://www.nature.com/articles/s41467-024-50805-z.pdf

[15] Feng Zheng, (2024), Systemic inflammatory response index as a predictor of stroke-associated pneumonia in patients with acute ischemic stroke treated by thrombectomy: a retrospective study, https://bmcneurol.biomedcentral.com/articles/10.1186/s12883-024-03783-0?form=MG0AV3

[16] Benjamin Mappin-Kasirer, (2024), Development of a clinical risk score for the prediction of Pneumocystis jirovecii pneumonia in hospitalised patients, https://link.springer.com/content/pdf/10.1186/s12879-024-09957-y.pdf

[17] Guang Zhao, (2024), The clinical value of nutritional and inflammatory indicators in predicting pneumonia among patients with intracerebral hemorrhage, https://www.nature.com/articles/s41598-024-67227-y.pdf

[18] Li Li, (2024), Construction and Validation of a Nomogram Model to Predict the Severity of Mycoplasma pneumoniae Pneumonia in Children, https://www.tandfonline.com/doi/pdf/10.2147/JIR.S447569

[19] Md. Maniruzzaman, (2024), Pneumonia prediction using deep learning in chest X-ray Images, https://www.researchgate.net/profile/Rahmanul-Hoque/publication/380824576

[20] Seung Min Baik, (2024), Integrating ensemble and machine learning models for early prediction of pneumonia mortality using laboratory tests, https://www.cell.com/heliyon/pdf/S2405-8440(24)10556-7.pdf