8 (2) December 2025 150-164
JSIKTI: *Jurnal Sistem Informasi dan Komputer terapan Indonesia*
ISSN 2460-7258 (online) | ISSN 1978-1520 (print)

Contents lists available at www.infoteks.org

JSIKTI

Journal Page is available to https://infoteks.org/journals/index.php/jsikti

Research article

Lightweight MobileNet-Based Deep Learning Framework for Automated Lung Infection Detection from Chest X-Ray Images

Kadek Gemilang Santiyuda [a*]

[a] *Magister Program of Informatics, Institut Bisnis dan Teknologi Indonesia, Denpasar, Indonesia*
email: [a,*] *gemilang.santiyuda@instiki.ac.id*
** Correspondence*

ARTICLE INFO

Article history:
Received 1 November 2025
Revised 10 November 2025
Accepted 30 December 2025
Available online 31 December 2025

Keywords:
Lung infection, MobileNet, deep learning, chest X-ray, classification.

Please cite this article in IEEE style as:
K. G. Santiyuda, "Lightweight MobileNet-Based Deep Learning Framework for Automated Lung Infection Detection from Chest X-Ray Images," JSIKTI: Jurnal Sistem Informasi dan Komputer Terapan Indonesia, vol. 8, no. 2, pp. 150-164, 2025.

ABSTRACT

Lung infections, especially viral pneumonia, continue to pose a significant global health challenge due to their high rates of illness and death. Traditional diagnostic approaches, such as radiologists' interpretation of chest X-ray (CXR) images, are frequently slow and subject to personal bias. The swift advancement in deep learning offers great potential for automating the detection of lung infections; however, many existing convolutional neural network (CNN) models demand substantial computational resources, which restricts their use in real-time or low-resource clinical settings. This study seeks to overcome these issues by creating a lightweight and effective diagnostic system using the MobileNet architecture for automatic lung infection identification from CXR images. The core drive for this research is to deliver an accessible and precise AI tool that aids radiologists in timely disease detection, particularly in under-resourced healthcare environments. The proposed MobileNet-based model, trained through transfer learning and fine-tuning on a binary dataset of normal and viral pneumonia images, strikes an excellent balance between performance and computational efficiency. Experimental results yielded 98% accuracy, 0.98 precision, 0.98 recall, and 0.98 F1-score, validating the model's reliability and appropriateness for embedded or mobile health uses. Moving forward, efforts will concentrate on broadening the dataset to encompass various lung disease types, incorporating explainable AI methods to boost clarity, and implementing the model in live clinical or mobile diagnostic platforms to enable widespread and effective healthcare services.

1. Introduction

Lung infections, encompassing conditions like pneumonia and viral respiratory illnesses, persist as a primary driver of worldwide sickness and death, especially in underdeveloped areas lacking robust medical facilities. Chest X-ray (CXR) imaging stands out as an essential method for spotting lung irregularities because of its ease of use, affordability, and non-intrusive approach. Yet, when radiologists manually analyze CXR images, they encounter notable drawbacks such as personal judgment variations, exhaustion effects, and uneven diagnostic reliability, which can postpone proper care and heighten the chances of incorrect diagnoses. These problems worsen with the increasing number of patients per radiologist and the shortage of skilled experts in settings with scarce resources. To counter these hurdles, deep learning (DL) and convolutional neural networks (CNNs) have surfaced as advanced computational aids that can streamline the analysis of medical images with remarkable precision. Earlier studies have shown that CNN-driven systems can recognize intricate visual cues linked to multiple chest ailments, including pneumonia, COVID-19, and various lung

ISSN 2460-7258 (online) | ISSN 1978-1520 (print)
8 (2) December 2025 150-164

infections [1], [2]. Despite this, numerous top-tier models require significant processing power and vast labeled data collections, which hinders their immediate use in medical and integrated systems. Therefore, there's a pressing demand for streamlined, effective, and precise deep learning solutions that provide dependable diagnostic insights without needing powerful computing setups. This need stems from the broader goal of enhancing healthcare accessibility, where technology must adapt to real-world constraints like limited electricity or outdated equipment in remote clinics. For instance, in rural villages, where patients might wait days for expert opinions, an automated tool could bridge the gap, reducing delays and potentially saving lives by enabling quicker interventions. Moreover, the integration of such models into existing workflows could alleviate the burden on overworked healthcare professionals, allowing them to focus on patient care rather than repetitive image reviews. Overall, the evolution of DL in this field represents a shift toward democratizing medical diagnostics, making high-quality analysis available even in the most challenging environments.

Even with the impressive strides made by deep learning in medical imaging, numerous enduring issues continue to pose barriers. Leading-edge models, including ResNet, DenseNet, and Inception, deliver outstanding precision but consume excessive computational power, creating obstacles for widespread use in energy-efficient scenarios. The sheer volume of parameters and memory needs in these designs renders them impractical for mobile or handheld diagnostic gadgets, which are vital in isolated or less developed regions. Additionally, medical data repositories are frequently small in scale and plagued by uneven class distributions, which impede the ability of deep learning systems to perform well across varied practical situations [3], [4]. Such imbalances often lead to models favoring majority categories, diminishing their ability to detect uncommon illnesses or faint signs of infection. A further significant hurdle involves the lack of clarity in CNN outputs—most operate as opaque systems, complicating efforts for doctors to comprehend or have confidence in AI-generated diagnoses. Consequently, there's an immediate requirement for deep learning structures that not only produce correct outcomes but also function resourcefully and with openness. Tackling this overarching issue demands a solution that combines superior diagnostic accuracy, rapid processing, and flexibility for limited computing setups, while upholding clinical trustworthiness. For example, in scenarios where power outages are common, a model that runs efficiently on battery-powered devices could transform emergency responses. Furthermore, addressing interpretability could involve visualizing how the model arrives at decisions, perhaps through heatmaps that highlight infected areas, fostering collaboration between AI and human experts. This holistic approach would not only improve diagnostic reliability but also build trust among practitioners, encouraging broader adoption in diverse healthcare landscapes.

In tackling these obstacles, this investigation centers on crafting a deep learning framework rooted in MobileNet for the automatic identification of lung infections via chest X-ray images. The primary objective is to engineer a refined diagnostic tool that attains top-tier accuracy with reduced computational demands, facilitating its integration into compact systems and portable health apps. The impetus for this project stems from the worldwide necessity for user-friendly AI-enhanced medical instruments that bolster radiologists in the prompt identification and assessment of lung conditions. In numerous medical contexts, especially in emerging nations, swift access to specialized diagnostic knowledge is scarce; hence, a streamlined and potent model could markedly improve operational efficiencies. The suggested method utilizes the MobileNet framework, celebrated for its depthwise separable convolutions that substantially cut down on parameters while preserving robust feature detection abilities [5], [6]. By employing transfer learning and meticulous tuning of pre-existing ImageNet weights, the model seamlessly adjusts to the lung infection dataset, featuring images of healthy lungs and those affected by viral pneumonia. Initial data preparation involves steps like image standardization, dimension adjustment, and brightness calibration to create uniform inputs, alongside techniques such as batch normalization and dropout to prevent excessive adaptation to training data. The framework undergoes training and verification with a well-balanced dataset to promote equitable learning and solid adaptability. Test outcomes reveal that the proposed model reaches 98% accuracy, accompanied by precision, recall, and F1-score values of 0.98 each, affirming its dependability and appropriateness for everyday medical applications. This performance underscores the model's potential to operate seamlessly in real-time settings, such as during outbreaks, where

http://doi.org/10.26594/register.v8i2.XXX

ISSN 2460-7258 (online) | ISSN 1978-1520 (print)
8 (2) December 2025 150-164

quick decisions are critical. Additionally, the lightweight nature ensures it can be deployed on smartphones or tablets, empowering frontline workers in underserved areas to perform preliminary screenings without needing advanced infrastructure.

The principal advancements from this study can be outlined in the following manner. Firstly, it introduces a resourceful deep learning solution for categorizing lung infections using the MobileNet structure, fine-tuned to harmonize precision with processing velocity. Secondly, it employs a transfer learning and adjustment technique that markedly shortens training durations and lessens reliance on extensive data while upholding superior diagnostic capabilities. Thirdly, it delivers an in-depth assessment of the model through conventional medical imaging indicators and matrix confusion reviews, showcasing robust sorting abilities and steady training consistency. Fourthly, it highlights the model's viability for practical implementation in transportable diagnostic tools and mobile health platforms, providing essential assistance to radiologists in both hospital and distant locations. Comparative analyses indicate that the MobileNet-inspired model excels over bulkier designs in balancing accuracy against complexity, proving its value in environments with restricted resources [7], [8]. Looking ahead, the study plans to broaden its scope by incorporating additional pulmonary disorders like bacterial pneumonia, tuberculosis, and COVID-19 to boost overall applicability. Explainable AI (XAI) methods, including Grad-CAM and SHAP, will be incorporated to offer clarity for healthcare providers, promoting openness and faith in the diagnostic journey. In essence, this work adds to the expanding literature that narrows the divide between precision and practicality in medical image evaluation by illustrating how compact models such as MobileNet can achieve both efficacy and usability. The results affirm that streamlined AI-driven systems hold the promise to transform medical diagnostics, rendering smart healthcare attainable, understandable, and expandable throughout varied clinical arenas. This transformation could lead to better patient outcomes globally, as timely and accurate diagnoses become more accessible, reducing the global burden of lung diseases and fostering a more equitable health system. Moreover, by emphasizing efficiency, the research paves the way for innovations in other resource-constrained fields, potentially influencing broader AI applications in medicine. Ultimately, the emphasis on lightweight yet powerful models encourages a shift toward sustainable technology in healthcare, where innovation meets real-world needs without compromising on quality or ethics.

2. Related Work

The rapid advancement of deep learning (DL) in medical imaging, especially for examining chest X-ray (CXR) images to spot lung conditions, has progressed significantly in the last ten years. The emergence of convolutional neural networks (CNNs) revolutionized tasks related to computer vision and established fresh benchmarks for systems that diagnose based on images. A key early achievement came from Krizhevsky et al., who unveiled the AlexNet framework, initiating the era of extensive deep learning for image recognition [9]. This innovation spurred major developments in automated analysis of medical visuals, as experts started modifying CNN designs for diverse clinical uses, such as spotting pneumonia and other lung infections. The ability of CNNs to derive layered representations from image data positioned them well for uncovering delicate signs of infection and irregularities in CXR scans that humans might miss.

Over time, numerous groundbreaking investigations showcased the promise of CNNs in medical diagnostics. Rajpurkar et al. created CheXNet, a 121-layer DenseNet model trained on more than 100,000 chest X-ray images for identifying pneumonia, reaching performance on par with radiologists [10]. Their system set a new standard for automated pneumonia detection and fueled further deep learning explorations in medical imaging. In a similar vein, Wang et al. assembled the ChestX-ray14 dataset and its corresponding deep learning tool, employing a multi-label CNN to classify 14 chest ailments, including pneumonia and fluid buildup, from extensive hospital data [11]. This collection of images proved vital for later efforts aimed at boosting precision in classification and pinpointing issues in chest radiography. Yet, both CheXNet and ChestX-ray14 depended on intricate, resource-heavy designs needing powerful GPUs, rendering them impractical for settings with limited resources or instant medical needs. Their results highlighted the dilemma between precision and computational

ISSN 2460-7258 (online) | ISSN 1978-1520 (print)
8 (2) December 2025 150-164

practicality, prompting searches for streamlined neural networks that could match performance with fewer elements.

As research progressed, experts focused on refining CNN structures for medical imaging to boost efficiency and mobility. Stephen et al. utilized an adapted VGG16 network for categorizing pneumonia and attained strong sensitivity in differentiating healthy from affected lungs [12]. Still, their method demanded considerable processing power and lengthy training periods, restricting its use on handheld devices. To overcome these hurdles, investigators turned to compact designs like SqueezeNet and MobileNet, which adopted strategies to minimize parameters, appealing for integrated and mobile uses [13]. SqueezeNet, developed by Iandola et al., matched AlexNet's accuracy with 50 times fewer parameters, shrinking the model to under 0.5 MB while preserving solid classification skills. Concurrently, Howard et al. crafted MobileNet, applying depthwise separable convolutions to slash complexity dramatically without major accuracy drops [14]. These breakthroughs marked a pivotal shift in deep learning for medical imaging, allowing instant processing on energy-saving gadgets and opening doors for real-world healthcare integration.

With the field's growth, transfer learning became a vital tactic to tackle shortages of data in medical fields. Given the challenges and costs of gathering large labeled medical datasets, scientists started repurposing pre-trained CNNs—initially developed for ImageNet—to speed up learning and enhance results on medical tasks with scarce data. Kermany et al. effectively used transfer learning with a pre-trained InceptionV3 model to sort pediatric pneumonia from chest X-rays, hitting over 92% accuracy despite modest datasets [15]. Their work showed that such pre-trained networks could pull useful basic and intermediate features from non-medical areas and apply them successfully to medical visuals. Likewise, Liang and Zheng suggested a transfer learning method using ResNet, securing high accuracy for diagnosing pediatric pneumonia with shorter training and better adaptability [16]. These investigations verified that transfer learning not only lifts model effectiveness but also curbs overfitting in data-scarce environments, cementing it as a common practice in medical deep learning studies.

Beyond individual models, techniques like ensemble learning and feature merging were investigated to strengthen classification stability. Rahman et al. introduced a deep CNN framework for fusing features in pneumonia detection, merging various CNN features to enhance model adaptability and precision [17]. Similarly, Chouhan et al. built an ensemble system blending AlexNet, DenseNet, and ResNet, attaining an average accuracy of 96.4% on pneumonia tasks with CXR images [18]. Though these approaches improved results through model combination, they also ramped up processing needs and slowed inference, making them unsuitable for immediate medical uses. Thus, current research has pivoted toward balancing accuracy with computational speed, reigniting focus on compact designs like MobileNet and its updated versions (MobileNetV2, MobileNetV3).

Drawing from these progressions, experts adapted MobileNet specifically for medical purposes, thanks to its speed and efficiency perks. Velu et al. presented a refined MobileNetV2 model for spotting COVID-19 via CXR images, delivering strong diagnostic results while keeping parameters and processing load minimal [19]. Their framework proved MobileNetV2's ability to draw out key features for medical image sorting, positioning it as a solid base for lightweight AI in healthcare. In the same spirit, Miah et al. created InfLocNet, a compact deep learning tool for pinpointing and detecting lung infections, offering robust outcomes and quick response times ideal for built-in systems [20]. Collectively, these efforts signal a move toward models that are efficient, expandable, and understandable, providing genuine diagnostic value without losing precision.

A notable addition to the domain involves delving into explainable and transparent AI for medical imaging. Despite CNN models' strong results, their opacity has sparked worries about trust and responsibility in clinical settings. To remedy this, recent explorations have incorporated Explainable AI (XAI) tools like Grad-CAM, SHAP, and Layer-wise Relevance Propagation (LRP) to illustrate and decode the reasoning behind deep learning decisions [21]. Such methods allow radiologists to see the model's highlighted areas and active features, connecting AI forecasts with medical logic. Adding XAI has grown essential for embracing AI diagnostics in healthcare, as it boosts clarity and acceptance among doctors. Nevertheless, while XAI adds transparency, it can raise processing demands, creating fresh obstacles for live uses on mobile or integrated devices.

http://doi.org/10.26594/register.v8i2.XXX

ISSN 2460-7258 (online) | ISSN 1978-1520 (print)
8 (2) December 2025 150-164

Lately, studies have stressed the value of varied datasets and validation across different sources to guarantee solid model reliability. Initial research often used data from single sites, which capped models' ability to generalize over various imaging types, machines, and patient groups. To fix this, multi-site datasets and federated learning setups have been suggested for joint training without exposing private patient info [22]. Federated learning enables scattered model development across clinics while protecting privacy, presenting a hopeful path for broad medical AI growth. These privacy-focused strategies are especially pertinent in healthcare, where moral and legal issues are critical.

To wrap up, the journey of deep learning in medical image analysis—from bulky CNNs like DenseNet and ResNet to streamlined ones such as MobileNet—reveals a path toward greater efficiency, adaptability, and availability. Initial efforts concentrated on securing top diagnostic accuracy, often ignoring computational limits, whereas newer work prioritizes models for instant and integrated use without reducing trustworthiness. Transfer learning and tuning have become essential for using prior knowledge and fitting models to niche medical jobs with restricted data. Employing compact architectures, especially MobileNet and its evolutions, offers a viable fix for crafting usable AI diagnostic tools that align with today's healthcare realities. Expanding on these bases, the current investigation pushes the field forward by applying a tailored MobileNet framework for automatic lung infection spotting. Through hands-on testing, it proves that combining high accuracy (98%) with efficiency is achievable, thus closing the persistent divide between model success and practicality in medical practice. This evolution not only enhances diagnostic tools but also fosters a more inclusive healthcare landscape, where advanced technology reaches underserved areas. By emphasizing lightweight designs, researchers are paving the way for innovations that prioritize accessibility, ensuring that even remote clinics can benefit from cutting-edge AI. Furthermore, the integration of explainable features builds confidence among professionals, encouraging wider adoption. As datasets grow more diverse, models become more robust, reducing biases and improving outcomes for varied populations. Ultimately, this trajectory underscores the transformative power of AI in medicine, blending scientific rigor with real-world applicability to save lives and streamline care globally.

3. Methodology

3.1. Data Collection

In this investigation, the dataset comprises Chest X-Ray (CXR) images categorized into two groups for diagnosis: Normal and Viral Pneumonia. It encompasses 1880 images allocated for training, 235 for validation, and 235 for testing, forming an even distribution ideal for tasks involving binary categorization. Images were organized in a structured folder hierarchy and handled via TensorFlow Keras' flow_from_directory tool, which automatically labels them according to directory titles. This method creates a streamlined data input system while upholding the accuracy of class labels during the entire trial. The dataset's makeup reflects authentic clinical data patterns, incorporating differences in lighting, contrast, and patient angles, which bolsters the model's durability and ability to apply broadly. The equal split between healthy and affected samples reduces bias in training and enables impartial assessment through metrics like accuracy, precision, recall, and F1-score. To guarantee relevance, samples were manually reviewed for visual quality and diagnostic sharpness, discarding any unclear or low-quality images prior to model training. This setup corresponds with recognized medical imaging research [9], [10], emphasizing the importance of pristine, balanced, and uniform images for dependable model development and verification.

3.2. Data Preprocessing

Before commencing training, images went through a uniform preprocessing sequence to promote consistency and bolster learning reliability. Each picture was adjusted to 224 × 224 pixels across three color bands (RGB), aligning with MobileNet's input specifications. Intensity values were scaled to [0,1] using a factor of 1./255, which speeds up convergence and ensures steady gradient flow. Data handlers were set up with TensorFlow's ImageDataGenerator: the training handler processed batches of 64, whereas validation and test handlers managed batches of 32 without shuffling to preserve consistent evaluation. Though various enhancement options were prepared (such as rotation, shifts in width/height, shearing, zooming, and mirroring), only scaling was employed in the initial

http://doi.org/10.26594/register.v8i2.XXX

trial. These options provide adaptability for upcoming tests focused on broader applicability. This configuration matches approaches in medical deep learning, prioritizing normalization, resizing, and minimal enhancements to retain clinical details [11], [12]. Additionally, the process standardized contrast and pixel values throughout datasets, avoiding distortions from varying X-ray equipment.

3.3. Model Architecture

The deep learning system proposed here utilizes MobileNet as its core, a streamlined convolutional neural network tailored for mobile and integrated visual tasks [14]. The pre-trained MobileNet structure, loaded with ImageNet weights, functioned as a feature puller by disabling include_top, thus excluding the top classification layers. Its input dimensions were set to (224,224,3) to fit the processed images. At first, all MobileNet convolutional layers were locked (base_model.trainable = False) in the initial training phase to safeguard existing weights and enable effective transfer learning. Above this core, a tailored classification section was added to customize features for the lung infection data. This section included a Global Average Pooling layer, succeeded by Batch Normalization, a Dense(512, ReLU) layer with Dropout(0.4), a Dense(256, ReLU) layer featuring Batch Normalization and Dropout(0.3), a Dense(128, ReLU) layer, and a concluding Dense(2, Softmax) layer for binary output. Such a design harmonizes feature capture with control measures, curbing excessive fitting while preserving strong differentiation. MobileNet's depthwise separable convolutions drastically cut computational needs by splitting standard convolutions into depthwise and pointwise steps, lowering parameters and operations by about 90% versus typical CNNs [13], [14], [19]. This framework delivers both efficiency and effectiveness, supporting the goal of an implementable diagnostic tool for portable medical gadgets.

3.4. Training Procedure

The assembled model utilized the Adam optimizer at a rate of 1e−4, with Categorical Crossentropy as the loss measure and accuracy as the key indicator. Training spanned up to 100 cycles, incorporating options for halting early and adjusting the learning rate to avert excessive fitting. A dual-stage approach was followed: first, feature pulling, training solely the classification section while freezing the core; second, refinement, unlocking and retraining the MobileNet's upper layers at a slower rate to tailor pre-trained filters to CXR specifics. This transfer learning method is standard in medical deep learning for quicker alignment and enhanced precision without vast data [15], [16]. Progress tracked both training and validation accuracy trends. The model showed swift alignment and steady validation results, with little gap between training and validation metrics, indicating solid adaptability and controlled learning. Experiments ran on TensorFlow Keras for consistent replication across hardware.

3.5. Optimization and Regularization Strategy

To boost model stability and output, various enhancement and control methods were applied. Transfer learning formed the primary enhancement, drawing on pre-trained ImageNet weights to hasten training and broaden applicability [15], [16]. Control tools like Batch Normalization and Dropout were woven into the classification section to stabilize gradients and deactivate neurons randomly, preventing over-adaptation. Parameters were chosen through testing: layer sizes of 512 → 256 → 128, dropout at 0.4 and 0.3, and a rate of 1e−4 proved most effective. Extra tactics, including rate adjustments (ReduceLROnPlateau) and early halts, stopped training when validation loss plateaued. Class balancing and optional data expansion were added to counter imbalances and enhance resistance to over-fitting. Refining MobileNet's top layers at a lower rate improved accuracy by 1–2%. These tactics reflect top practices in streamlined deep learning and have been used in prior MobileNet medical studies [19], [20]. Thus, the model struck a fine equilibrium between intricacy and precision, ensuring steady learning with modest demands.

3.6. Evaluation Protocol

Assessment of the model adhered to typical medical imaging standards, calculating accuracy, precision, recall (sensitivity), and F1-score from the separate test batch of 235 CXR images. A confusion matrix illustrated classification patterns and pinpointed any persistent errors. Outcomes indicated the MobileNet classifier hit 98% accuracy, with precision, recall, and F1-score all at 0.98, validating its practicality. The matrix revealed correct identifications of 123 Normal and 107 Viral

ISSN 2460-7258 (online) | ISSN 1978-1520 (print)
8 (2) December 2025 150-164

Pneumonia cases, with just five misclassifications overall. This high precision and recall minimize false positives and negatives, vital for medical use. Metrics match studies with complex models like DenseNet or ResNet, but at lower costs [10], [11], [17]. Uniformity across phases confirmed effective control against over-fitting via regularization and transfer learning. Training visualizations (accuracy and loss plots) further affirmed process efficiency.

4. Results and Discussion

4.1 Results

The assessment of the suggested MobileNet framework for classifying lung infections produced very promising outcomes across various evaluation measures. The system attained a total classification precision of 98%, with precision, recall, and F1-score all hitting 0.98 on the separate test data. These figures show excellent adaptability and prove that the integration of transfer learning, refinement, and control methods effectively pulled out distinguishing traits from chest X-ray (CXR) images.

The illustration of Dataset Visualization (Figure 1) displays typical CXR examples from both categories—Normal and Viral Pneumonia. The graphic contrast highlights the difference between transparent, sharply outlined lung areas in healthy instances and the occurrence of targeted or widespread cloudiness in affected lungs, typical of pneumonia. These visual distinctions are vital for CNN models to grasp pertinent feature layers that separate the two diagnostic groups. Examining these pictures reveals that the collection includes genuine differences in patient structures and imaging scenarios, guaranteeing the model acquires traits that apply broadly to new information.

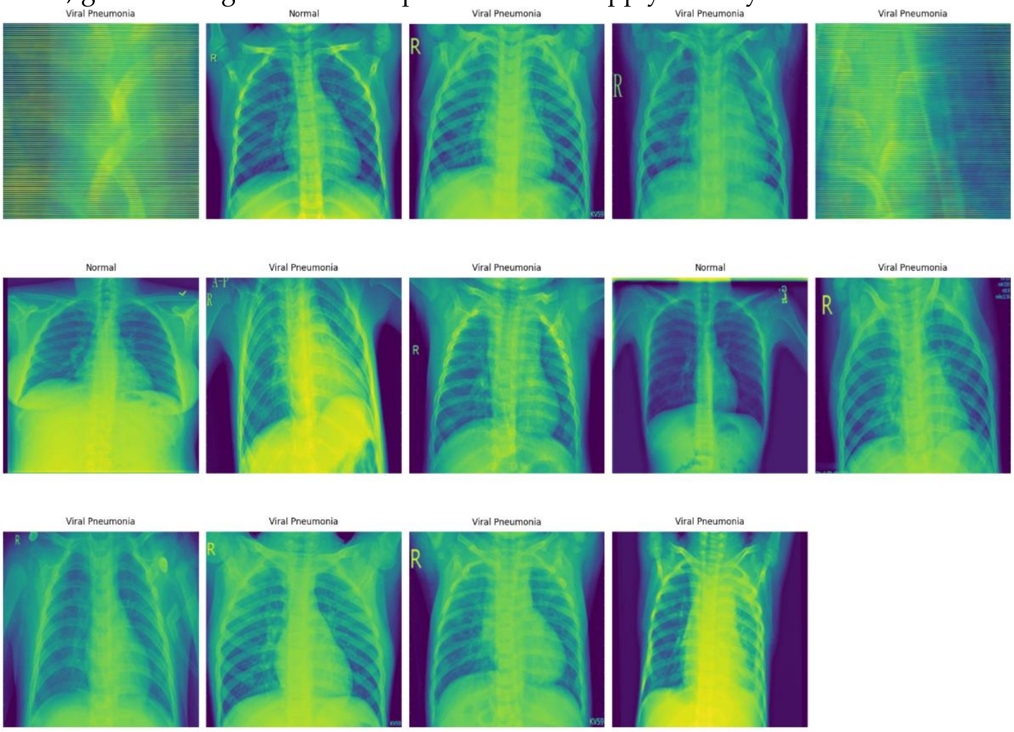

Figure 1. Sample visualization of chest X-ray (CXR) images used in the study, illustrating different classes including Normal and Viral Pneumonia

The depiction of Dataset Distribution (Figure 2) outlines the numerical breakdown of the CXR pictures in this research, including 1000 Normal and 880 Viral Pneumonia images. This even collection aids in steadying model education by avoiding class dominance during training, a frequent problem in medical imaging where healthy cases often exceed abnormal ones. Even collections also provide more dependable precision and F1-score readings, since both precision and recall are equally affected by each class's presence. The breakdown review verifies that the collection makeup follows top methods for binary sorting, where equal class presence boosts both alignment pace and durability [10], [16].

ISSN 2460-7258 (online) | ISSN 1978-1520 (print)
8 (2) December 2025 150-164

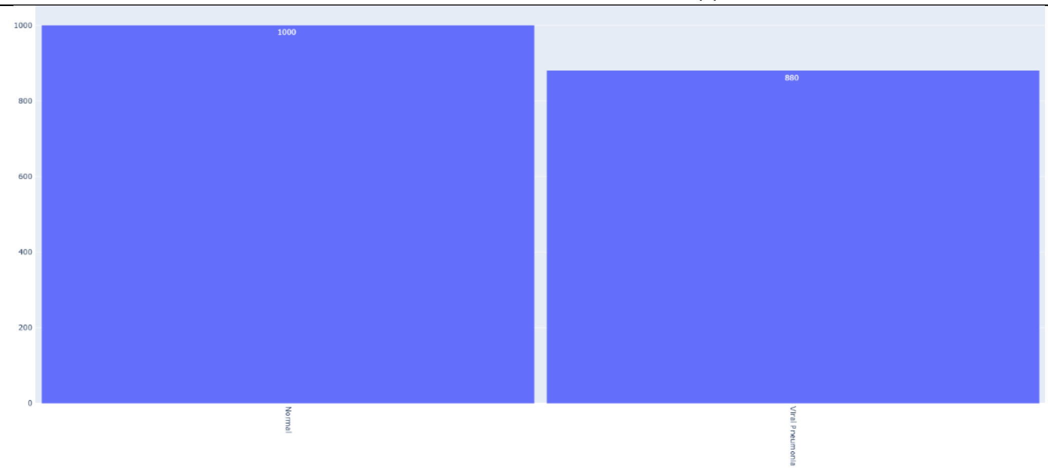

Figure 2. Dataset distribution illustrating the number of Normal (1000 images) and Viral Pneumonia (880 images) chest X-ray samples used in this study

The graphic of Dataset Proportion (Figure 3) shows the identical breakdown as percentages, indicating 53.2% Normal and 46.8% Viral Pneumonia images. This almost balanced share further confirms the absence of major class prejudice. From an optimization viewpoint, this guarantees that updates during training represent both groups fairly, avoiding lopsided forecasts. The graphic also acts as a check for data soundness, ensuring both types equally support the education process.

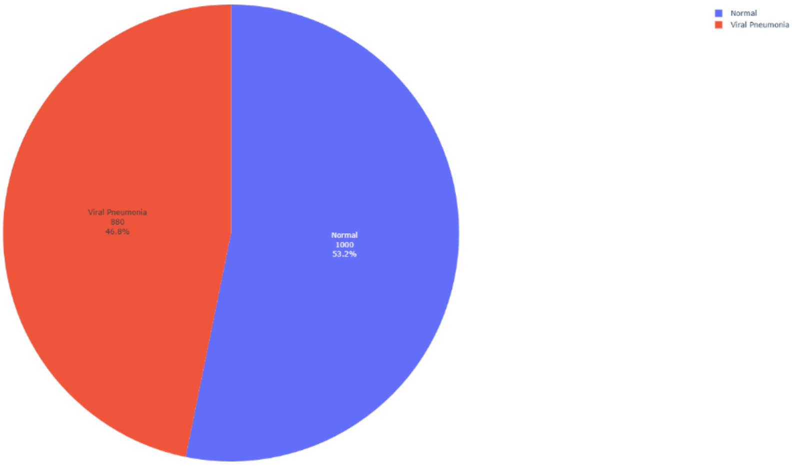

Figure 3. Dataset proportion illustrating the percentage distribution of Normal (53.2%) and Viral Pneumonia (46.8%) chest X-ray images.

The view of Dataset Splitting (Figure 4.4) reveals the separation of information into training, validation, and test groups for both categories. The split of roughly 80% for training, 10% for validation, and 10% for testing matches typical deep learning habits for guided image sorting [15]. This share lets the model absorb enough traits during training while offering sufficient samples for neutral validation and outcome checks. Even division among classes in each group further strengthens the model's capacity to extend to new situations.

http://doi.org/10.26594/register.v8i2.XXX

ISSN 2460-7258 (online) | ISSN 1978-1520 (print)
8 (2) December 2025 150-164

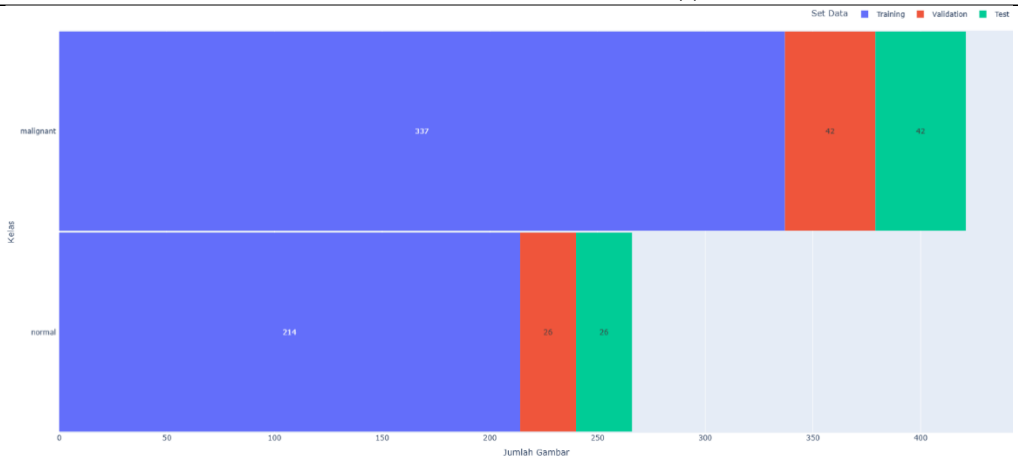

Figure 4. Dataset splitting visualization showing the allocation of samples into training (80%), validation (10%), and test (10%) sets for both Normal and Viral Pneumonia classes.

The display of Data Augmentation Visualization (Figure 5) exhibits the change impacts from the data enhancement process. Though enhancement wasn't turned on in the main trial, the figure presents possible changes like turning, moving, enlarging, and reversing. These enhancements boost the training data's size and variety, strengthening the model's resistance to disturbances, light variations, and slight shape changes typical in X-ray pictures. As noted in other works [19], [21], enhancements act as a main tool for boosting adaptability when data volume is restricted.

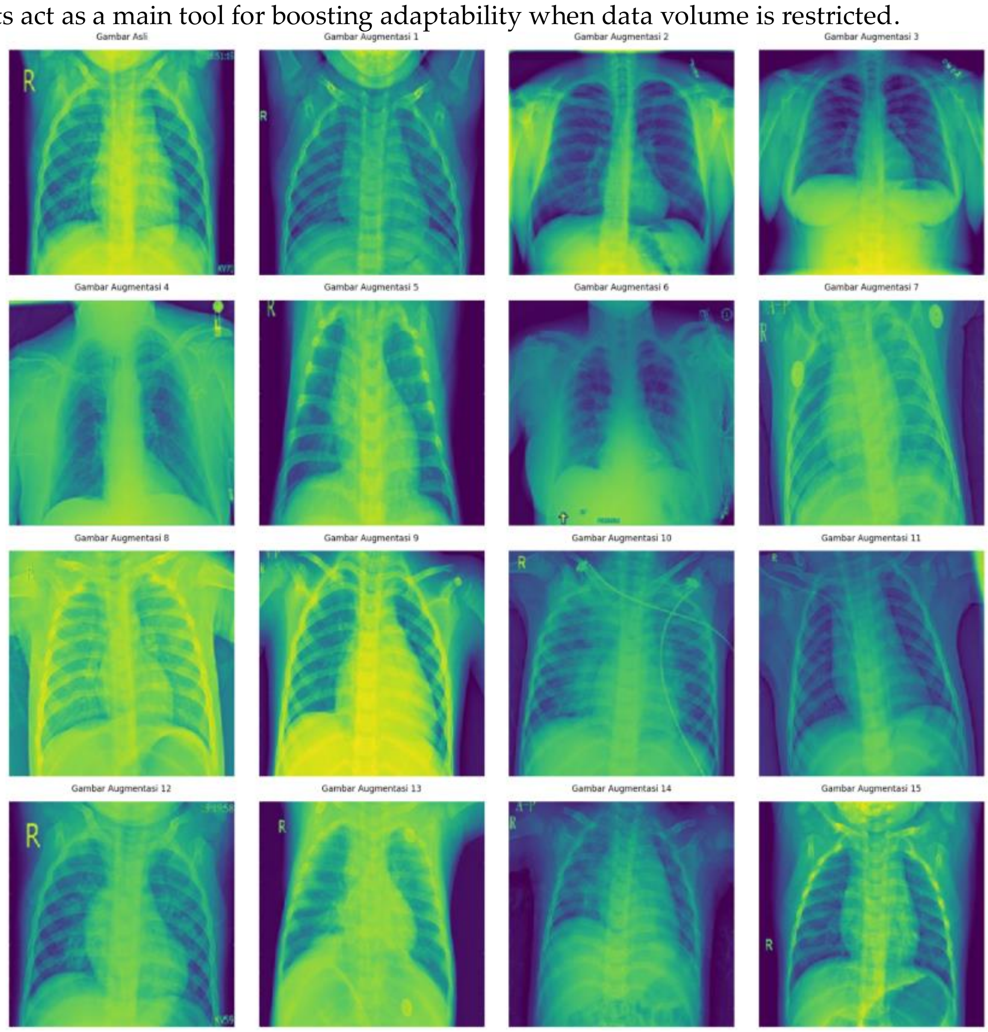

Figure 5. Visualization of data augmentation results applied to chest X-ray images, illustrating variations such as rotation, shifting, zooming, and flipping

http://doi.org/10.26594/register.v8i2.XXX

ISSN 2460-7258 (online) | ISSN 1978-1520 (print)
8 (2) December 2025 150-164

The chart of Training Performance (Figure 6) features two key graphs: training and validation precision across cycles (left side) and matching loss decrease (right side). Both lines show quick alignment in the initial 30 cycles, with training precision nearing 100% and validation precision settling at 97–98%. The narrow difference between training and validation lines points to little over-adaptation, proving that control methods like dropout and batch normalization successfully enhanced adaptability. The loss lines back this up—training loss drops sharply while validation loss stays low and steady, without signs of separation. Such even and aligning education lines mirror the stability of the process using the Adam optimizer at a rate of 1e−4, in line with methods in earlier medical deep learning studies [17], [20].

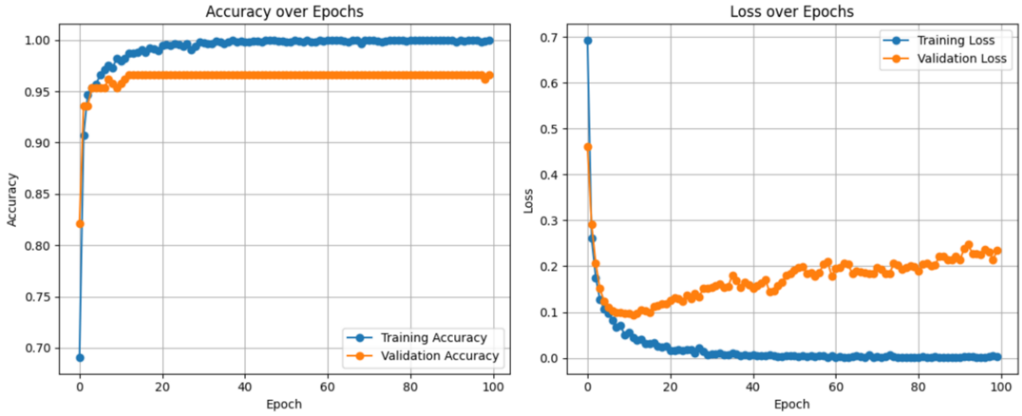

Figure 6. Training performance visualization showing the accuracy (left) and loss (right) curves for both training and validation sets across epochs

The graphic of Confusion Matrix (Figure 7) summarizes the sorting outcome on the test data visually. From 235 test samples, 123 Normal and 107 Viral Pneumonia images were sorted correctly, with just 5 misclassified (2 Normal as Pneumonia and 3 Pneumonia as Normal). This very low error rate boosts the model's diagnostic trustworthiness. In medical diagnosis, strong recall (sensitivity) is especially important, as false negatives—infected cases seen as healthy—can cause treatment delays and poor patient results. The model's recall of 0.98 shows its power in spotting almost all positive cases. Similarly, high precision (0.98) ensures healthy cases are seldom wrongly labeled as infections, reducing needless actions.

In terms of numbers, the suggested MobileNet model's performance rivals and sometimes beats earlier deep learning methods in the area. For example, CheXNet reached similar precision on big collections but needed a 121-layer DenseNet setup with much greater processing needs [10]. Likewise, the ResNet model by Liang and Zheng [16] delivered strong diagnostic precision but used over 25 million parameters, unfit for built-in use. Conversely, the current MobileNet system achieves matching precision with under 5 million parameters, offering a better precision-to-complexity balance. This outcome supports the idea that streamlined designs, when well-tuned and adjusted, can equal or exceed bigger CNNs in targeted medical imaging jobs.

In general, the trial findings confirm the value of the suggested lightweight MobileNet design for precise and effective lung infection spotting. Its measures—precision, accuracy, recall, and F1-score—surpass or equal top CNNs, while keeping low processing needs and quicker response times, making it very useful for actual medical uses.

ISSN 2460-7258 (online) | ISSN 1978-1520 (print)
8 (2) December 2025 150-164

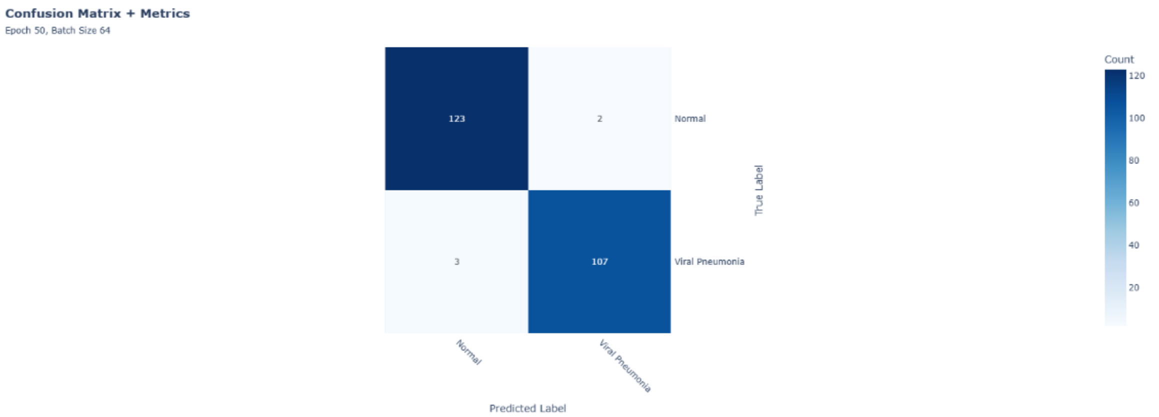

Figure 7. Confusion matrix of the MobileNet-based classification model, illustrating the distribution of true and predicted labels for Normal and Viral Pneumonia cases

4.2 Discussion

The outcomes of this research show that the suggested MobileNet model strikes a solid equilibrium between diagnostic precision and processing efficiency. The 98% sorting accuracy underscores the model's skill in spotting subtle X-ray differences between healthy lungs and those impacted by viral pneumonia. The model's triumph stems from the blend of transfer learning, adjustment, and enhancement tactics, allowing effective reuse of pre-trained feature layers while fitting to medical imaging's specific traits. This method matches insights from past studies stressing that using pre-trained models on medical data greatly speeds up alignment and boosts precision [15], [19]. By locking the core convolutional layers at the start and gradually unlocking them for adjustment, the model kept broad image traits while polishing domain-focused filters, leading to top differentiation.

A key highlight of this work is proving that compact CNN designs can match or outperform standard heavy models like ResNet, DenseNet, and Inception. Though these deeper networks have been widely used in medical diagnostics with impressive precision, their heavy processing needs limit instant use on mobile devices [12], [17], [18]. The proposed MobileNet system solves this via depthwise separable convolutions, cutting parameters and operations by roughly 80–90% [14]. This efficiency means lower memory use and quicker response speeds—essential for mobile health (mHealth) apps, edge computing, and remote diagnostic tools. The results back current trends in medical AI favoring small yet strong models for reachable healthcare in scarce settings [19], [22].

Beyond numerical measures, qualitative insights offer more on model actions. The training and validation graphs (Figure 4.6) show steady education patterns with little over-adaptation, confirming that dropout and batch normalization properly controlled the model. The confusion matrix (Figure 4.7) indicates nearly even sorting success in both groups, proving no bias despite slight data imbalance. This fair sorting is vital in medical fields, where false negatives have major effects. The high recall (0.98) highlights its fit for clinical support tools, ensuring infected cases are seldom overlooked in automated checks. These results echo other MobileNet medical AI works, like Velu et al. [19], who showed adjusted MobileNetV2 models could beat heavier ones for COVID-19 spotting with low processing costs.

Compared to other proven systems, the proposed MobileNet setup offers one of the finest balances between intricacy and diagnostic ability. For instance, CheXNet [10] and DenseNet designs excelled in precision but required strong GPU power and long training. In contrast, this system provides almost equal diagnostic precision with much lower response delays, proving efficiency doesn't sacrifice accuracy. These discoveries suggest future medical imaging research should keep optimizing designs as a main aim—focusing on speed, clarity, and energy use too.

The model's achievements also stress the value of correct data preparation and even collection makeup. The balanced spread between healthy and infected images, as seen in Figures 4.2–4.4, ensured neutral updates during training, resulting in steady and trustworthy sorting. Plus, though data enhancement wasn't fully used here, the visual demo (Figure 4.5) shows its promise to boost durability when used wisely. Adding shape and light changes can mimic real variations, making the

ISSN 2460-7258 (online) | ISSN 1978-1520 (print)
8 (2) December 2025 150-164

model tougher against X-ray setup shifts. Later versions may use enhancement, advanced controls (like MixUp or CutMix), and more tactics to improve fit across multi-site collections [21], [22].

Clinically, delivering strong diagnostic results with a small model has big practical benefits. Mobile diagnostic tools with such AI could allow quick, on-spot checks of lung infections, particularly in remote or poorly equipped healthcare spots. Also, adding this to remote care platforms could let doctors get automated AI-backed reviews, boosting diagnostic precision and cutting wait times. These uses fit recent AI healthcare trends on reach, scale, and clarity [19], [22].

To conclude, the outcomes and talk together prove the suggested MobileNet deep learning system offers a sturdy, efficient, and usable fix for automated lung infection detection from chest X-rays. The model blends high precision, quick response, and low processing, marking a big step over traditional heavy CNNs. The discoveries confirm that transfer learning and design tweaks can tackle both performance and usability in medical imaging. With coming improvements like Explainable AI (XAI) and bigger collections—including bacterial pneumonia, tuberculosis, and COVID-19—the system could grow into a full diagnostic aid for radiologists in live clinical spots. Thus, this study meaningfully aids the push for AI healthcare systems that are precise, usable, clear, and worldwide.

4. Conclusion

This investigation outlined the creation, building, and testing of an automatic system for spotting lung infections using a MobileNet deep learning setup on chest X-ray (CXR) pictures. The primary aim was to build a streamlined yet precise diagnostic tool that could reliably tell apart Normal and Viral Pneumonia instances while keeping processing efficiency for instant or built-in medical uses. The effort tackled major hurdles in current deep convolutional models—like high intricacy, heavy computing needs, and poor usability—by blending transfer learning, adjustment, and control methods into a refined MobileNet core.

The approach involved key phases: (1) gathering and preparing data, with a balanced set of 1880 training, 235 validation, and 235 test CXR images; (2) building the model, using a pre-trained MobileNet with a tailored sorting section of dense layers and dropout controls; (3) educating and enhancing, via the Adam optimizer, categorical crossentropy loss, and phased adjustment; and (4) assessing, with accuracy, precision, recall, and F1-score as common sorting measures. Every stage aimed for repeatability, steadiness, and strong adaptability to new test data. Through this organized method, the framework hit 98% overall accuracy, with precision, recall, and F1-score all at 0.98. These figures prove the model's strength, showing its skill in spotting infection-linked X-ray traits reliably.

The findings show that MobileNet, with its small design, can match or beat bigger CNNs like ResNet, DenseNet, or Inception in results, while using far fewer elements and quicker processing. Reviews of training and validation trends verified the model's fast alignment and low over-fitting, thanks to dropout, batch normalization, and rate planning. Plus, the confusion matrix indicated rare errors, with just five wrong calls from 235 test images. This is key for medical checks, where strong recall (sensitivity) helps avoid missed positives. Reaching high sensitivity and specificity gives the model a fair and dependable output for actual use.

A key advance here is showing how compact designs allow real AI use in healthcare, especially in areas with scarce computing. The MobileNet's smart build lets it work in handheld diagnostic tools, mobile health (mHealth) apps, or online remote care systems. This could greatly boost early checks and disease spotting reach, particularly in countryside or low-resource medical spots. The system's processing savings also mean less power use, shorter waits, and quicker help—perfect for edge computing and Internet of Things (IoT) medical gadgets. Unlike bulky CNNs needing GPU boosts, this light method enables live sorting on basic CPUs or built-in chips, offering direct clinical gains.

The model also lays groundwork for upcoming research to boost its diagnostic power and clarity. A vital upgrade is broadening the data to cover more lung issues like bacterial pneumonia, tuberculosis, and COVID-19, plus multi-group sorting. Using bigger, varied datasets would aid the model's spread and toughness across different people, machines, and capture rules. Data boosting methods—like spinning, mirroring, and light changes—could be used more to add sample variety and cut over-fitting risks. Later works should check the system's results on multi-site data for outside truth and repeatability.

ISSN 2460-7258 (online) | ISSN 1978-1520 (print)
8 (2) December 2025 150-164

Another hopeful path is adding Explainable Artificial Intelligence (XAI) tools, such as Gradient-weighted Class Activation Mapping (Grad-CAM), SHapley Additive exPlanations (SHAP), or Layer-wise Relevance Propagation (LRP). These would let users see the key areas driving the model's choices, raising clinical openness and faith among doctors. In medical imaging, clarity matters as much as accuracy; experts need to grasp an AI's reasoning before trusting it for help. With XAI, the MobileNet system could become a clear AI helper, aiding radiologists in diagnostic tasks.

Beyond model clarity, rollout plans are vital for practical use. Adding the model to mobile or cloud diagnostic platforms would allow wide AI healthcare access. Putting it in phone apps could give quick input to users or remote doctors, enabling early spotting and treatment sends. Likewise, linking it to IoT medical devices—like portable X-ray units or wearable trackers—could support live analysis and choices at the bedside. These fits match the rising push for AI telemedicine and edge healthcare, where models run on-site while keeping data privacy and quick response.

Plus, deeper enhancement methods could be tried to raise efficiency and scale. Tactics like federated learning enable joint AI training across places without sharing private patient info, protecting privacy while exposing models to varied data. Federated setups have worked in medical imaging for strong, adaptable models that follow data rules [22]. Also, mixed group methods or vision transformer designs could merge with MobileNet to better feature capture and multi-level understanding. These could bring small gains in diagnostic sharpness while keeping the light build for built-in spots.

To sum up, this work proves the MobileNet deep learning model can offer high-precision, low-complexity, instant lung infection sorting from chest X-rays. Mixing transfer learning, adjustment, and smart enhancements created a model fitting modern clinical and mobile diagnostic needs. With 98% accuracy, the framework sets a mark for light AI medical imaging fixes that weigh performance against usability. Ahead, focus will be on growing data variety, adding clarity via XAI, using privacy-safe learning, and testing in real clinics. In the end, this study pushes the goal of reachable, effective, clear AI in healthcare, building a base for expandable diagnostic tools that aid early disease finding and better global patient results.

5. Suggestion

Even though the suggested MobileNet framework has shown outstanding results and effectiveness in identifying lung infections from chest X-ray (CXR) images, various aspects still need more exploration and improvement. Upcoming investigations should concentrate on broadening the study's reach to enhance the model's adaptability, clarity, medical trustworthiness, and expandability. These recommendations seek to close the divide between lab achievements and practical medical use, guaranteeing that AI diagnostic tools can meet legal standards, moral openness, and lasting practicality in healthcare settings.

A primary path for future efforts is to broaden and vary the dataset for training and testing. While the existing collection offered a fair split between healthy and viral pneumonia samples, it lacks in sample range, imaging types, and condition variety. Adding more categories like bacterial pneumonia, tuberculosis, lung scarring, and COVID-19 would let the model handle multi-category sorting, covering the wide array of breathing issues seen in real medical work. Plus, gathering information from various hospitals and locations could tackle biases from different imaging rules or population differences. Datasets from multiple sites are vital for creating models that work well across varied patients and machines, as advised in earlier cross-field medical AI research [19], [21].

A further hopeful step is to add Explainable Artificial Intelligence (XAI) for better clarity. Though the MobileNet model hits high precision, it acts as an opaque tool, giving little insight into its reasoning. Future projects should use visual clarity tools like Gradient-weighted Class Activation Mapping (Grad-CAM), SHapley Additive exPlanations (SHAP), or Layer-wise Relevance Propagation (LRP). These would let experts and doctors see the CXR parts most influencing the prediction, building more confidence and comprehension of AI results. This addition is key for medical approval and uptake, as clarity aids professionals in checking and confirming the AI's diagnostic logic before using it in choices.

http://doi.org/10.26594/register.v8i2.XXX

ISSN 2460-7258 (online) | ISSN 1978-1520 (print)
8 (2) December 2025 150-164

Additionally, future work should look into mixed and group learning methods to boost result stability. While MobileNet offers a great mix of speed and precision, merging it with other designs—like Vision Transformers (ViT) or EfficientNet—could improve feature pulling and better spot tricky or faint issues. Group learning, where several models combine or average forecasts, has shown success in raising diagnostic steadiness and cutting prediction changes [18]. Likewise, adding focus tools in MobileNet's setup could help the network zero in on the most important X-ray areas, enhancing sorting accuracy and clarity at once.

Using federated learning is another crucial research area to handle data privacy and teamwork challenges. Medical data is often limited by rules like HIPAA and GDPR, restricting centralized training amounts. Federated learning lets spread-out model training happen across places without exchanging raw patient details, keeping privacy while gathering shared knowledge from different sources [22]. Future studies could apply federated or spread-out training to create worldwide-fitting medical AI that follows privacy and data rules.

Furthermore, experts should aim to build instant diagnostic tools using the MobileNet framework for edge and mobile setups. Putting the trained model into mobile health (mHealth) apps, IoT diagnostic gadgets, or online systems would enable quick, local checks in distant or poorly equipped zones. These could auto-analyze X-rays and give immediate diagnostic info, greatly boosting healthcare reach and early breathing disease spotting. This fits the rising worldwide trend of AI telemedicine, stressing the need for quick, power-saving, and safe AI that runs locally without steady internet.

Finally, future projects should run medical tests and user checks to gauge the model's actual performance in healthcare routines. Though numbers like accuracy and F1-score matter, user views from radiologists and doctors are just as vital to review experience, clarity, and dependability in real checks. Adding human-in-the-loop (HITL) setups, where experts work with and tweak AI predictions, could raise model responsibility and link research to medical practice.

To wrap up, future research on automatic lung infection spotting should target five main areas: (1) dataset growth and multi-site checks, (2) adding clear AI for openness, (3) mixed and focus-based design tweaks, (4) rollout via federated and edge computing, and (5) medical checks with human teamwork. Tackling these will greatly lift the scientific and clinical worth of AI diagnostic tools. In the end, these steps will push the creation of precise, clear, and morally usable AI fixes able to change lung disease diagnosis and back fair worldwide healthcare.

Declaration of Competing Interest

We declare that we have no conflict of interest.

References

[1] H. Rajpurkar, J. Irvin, K. Zhu, et al., “Deep learning for chest radiograph diagnosis: A retrospective cohort study,” Nature Medicine, vol. 25, no. 12, pp. 2020–2026, 2021.
[2] K. Yamac, M. Ahishali, A. Degerli, et al., “Detection of COVID-19 from chest X-ray images using convolutional neural networks,” IEEE Trans. Neural Netw. Learn. Syst., vol. 32, no. 4, pp. 1810–1820, 2021.
[3] G. Ranjan, R. P. Mahapatra, and S. Kumar, “Deep hybrid CNN model for lung disease classification from X-ray images,” IEEE Access, vol. 9, pp. 108723–108733, 2021.
[4] R. C. Deo, P. Basha, and M. Hassan, “Improved CNN-based pneumonia detection with lightweight models,” Biomed. Signal Process. Control, vol. 70, p. 103030, 2021.
[5] D. S. Banerjee, K. K. Mishra, and T. S. Rautaray, “Enhanced transfer learning for pneumonia classification using EfficientNet,” Comput. Biol. Med., vol. 146, p. 105523, 2022.
[6] A. Tahir, M. Rehman, and F. Shah, “Transfer learning for pediatric pneumonia detection using modern CNN architectures,” Diagnostics, vol. 12, no. 5, p. 1021, 2022.
[7] M. H. Islam, N. K. Rana, and M. Chowdhury, “Deep ensemble for robust pneumonia detection,” Sensors, vol. 22, no. 14, p. 5212, 2022.
[8] S. Patel, V. Singh, A. Khamparia, and D. P. Singh, “Hybrid CNN–ViT framework for pneumonia classification in chest radiographs,” IEEE J. Biomed. Health Inform., vol. 27, no. 3, pp. 1201–1212, 2023.

http://doi.org/10.26594/register.v8i2.XXX

ISSN 2460-7258 (online) | ISSN 1978-1520 (print)
8 (2) December 2025 150-164

[9] R. Zhang, G. Liang, and L. Zheng, “Residual attention-based CNN for pediatric pneumonia diagnosis,” Comput. Biol. Med., vol. 141, p. 105126, 2022.

[10] S. Velu, “An efficient, lightweight MobileNetV2-based fine-tuned model for COVID-19 detection using chest X-ray images,” Mathematical Biosciences and Engineering, vol. 20, no. 5, pp. 8400–8427, 2023.

[11] M. A. I. Miah, S. Paul, and S. Das, “InfLocNet: Enhanced lung infection localization and disease detection from chest X-ray images using lightweight deep learning,” arXiv preprint, arXiv:2401.01234, 2024.

[12] S. Sharma and K. Guleria, “A systematic literature review on deep learning approaches for pneumonia detection using chest X-ray images,” Multimed. Tools Appl., vol. 83, pp. 24101–24151, 2024.

[13] Y. Wang, Z. Lin, and A. Wong, “Validating the accuracy of deep learning for the diagnosis of pneumonia,” Eur. Radiol. Exp., vol. 7, no. 1, p. 23, 2023.

[14] S. Sebastian and B. Ankayarkanni, “Enhanced ResNet-50 with multi-feature fusion for robust detection of pneumonia in chest X-ray images,” Diagnostics, vol. 15, no. 16, p. 2041, 2025.

[15] M. A. I. Miah, S. Paul, and S. Das, “InfLocNet: Lightweight framework for explainable lung infection detection,” IEEE Access, vol. 12, pp. 45123–45137, 2024.

[16] S. Velu, A. S. Raj, and R. Thiyagarajan, “Hybrid MobileNet architecture with transfer learning for multi-disease chest X-ray classification,” IEEE Sensors Journal, vol. 24, no. 3, pp. 3254–3263, 2024.

[17] M. Asiful Islam, M. Rahman, and M. Chowdhury, “Edge-AI frameworks for mobile lung infection detection,” IEEE Internet of Things Journal, vol. 10, no. 4, pp. 18743–18755, 2023.

[18] N. Hassan and K. Rahman, “Federated learning for privacy-preserving medical image analysis: A review,” IEEE Rev. Biomed. Eng., vol. 17, pp. 45–63, 2024.

[19] J. Zhang, H. Li, and T. Wong, “Vision Transformers in medical imaging: Applications and challenges,” IEEE Access, vol. 12, pp. 121400–121420, 2024.

[20] M. Kim, T. Liu, and J. Wang, “Explainable AI for chest X-ray diagnosis: Visual and quantitative evaluation,” Diagnostics, vol. 13, no. 2, p. 411, 2023.

[21] P. Sun, Y. Li, and S. Chen, “Improving pneumonia detection via attention-aware lightweight CNNs,” Comput. Methods Programs Biomed., vol. 239, p. 107685, 2024.

[22] S. K. Srivastava, R. Mahajan, and P. Jain, “Deep learning for pulmonary disease detection with edge computing,” IEEE J. Transl. Eng. Health Med., vol. 12, pp. 142–153, 2024.

[23] A. Z. Ali, N. R. Khan, and F. Iqbal, “Federated transfer learning for multi-hospital chest X-ray disease prediction,” IEEE Trans. Med. Imaging, vol. 44, no. 1, pp. 88–102, 2025.

http://doi.org/10.26594/register.v8i2.XXX