FUDMA Journal of Sciences (FJS)
ISSN online: 2616-1370
ISSN print: 2645 - 2944
Vol. 9 No. 10, October, 2025, pp 86 – 93
DOI: https://doi.org/10.33003/fjs-2025-0910-3963

CHEST X-RAY BASED DETECTION MODEL FOR PNEUMONIA IN PEDIATRIC

*[1]Onalaja Olabisi Olayinka, [2]Wilson Sakpere, [1]Awosola Adeoluwa Samuel and [3]Peter Adebayo Idowu

[1]Gateway ICT Polytechnic, Saapade, Ogun-State, Nigeria.
[2]Lead City University, Ibadan
[3]Obafemi Awolowo University, Ile Ife, Nigeria.

Corresponding author email: bizziluv2010@yahoo.com Phone: +2348060696003

ABSTRACT

Pneumonia remains a leading cause of morbidity and mortality among children, particularly in low-resource settings such as Nigeria. The accurate and timely diagnosis of pediatric pneumonia is hindered by the scarcity of skilled radiologists and diagnostic infrastructure. This study proposes a robust, efficient, and scalable classification model utilizing transfer learning to support pneumonia detection from chest X-ray (CXR) images. The model employs pre-trained convolutional neural networks (CNNs), fine-tuned under a novel Resource-Constrained Medical Transfer Learning (RCMTL) framework, to optimize predictive accuracy, computational efficiency, and equipment robustness. The approach shows promise in enhancing clinical decision-making, especially in under-resourced environments, and paves the way for practical AI integration in healthcare delivery.

Keywords: Pneumonia Detection, Pediatric Radiology, Machine Learning, Transfer Learning, Deep Learning, RCMTL, CXR, Resource-Limited Healthcare

INTRODUCTION

Pneumonia is an inflammatory condition of the lung parenchyma caused by various infectious agents (Liu et al., 2021). It is a significant cause of morbidity and mortality worldwide, especially in young children and the elderly (Ayuk, 2024; Tomys-Skladowska et al., 2023). Pneumonia remains one of the most persistent global health challenges, representing a leading cause of mortality worldwide, particularly among vulnerable populations. Recent epidemiological data reveal that pneumonia claimed the lives of 2.2 million people globally in 2021, including 502,000 children under five years of age, accounting for approximately 23% of pediatric deaths (Canada et al., 2024) The disease burden is disproportionately concentrated in developing regions, with over 1,400 cases per 100,000 children globally, and the highest incidence rates observed in South Asia (2,500 cases per 100,000 children) and West and Central Africa (1,620 cases per 100,000 children) (UNICEF, 2024). Despite significant medical advances, pneumonia continues to account for 14% of all deaths among children under five years, highlighting the urgent need for improved diagnostic and therapeutic interventions (Teixeira, 2020)

The burden of childhood pneumonia is severe and widespread in Nigeria, contributing the highest number of child deaths globally from this preventable disease (UNICEF, 2024). The country ranks second globally after India in the absolute number of pneumonia deaths among children under five years, with recent projections indicating that approximately two million Nigerian children could die from pneumonia in the next decade without significant improvements in prevention and treatment efforts (UNICEF, 2024). According to Okafor et al. (2023) and Odeyemi et al. (2022), Pneumonia kills a child every three minutes in Nigeria, making it one of the leading causes of childhood mortality in the country. The overwhelming impact of pneumonia in Nigeria is compounded by several interconnected factors, including widespread malnutrition, severe air pollution, and limited access to essential vaccines and life-saving drugs. Nigeria is among the 15 countries that collectively account for over 70% of global under-five pneumonia and diarrhea mortality, emphasizing the critical need for targeted interventions in the Nigerian healthcare context (International Vaccine Access Center, 2024).

The correct diagnosis of pneumonia presents extensive clinical challenges that compound the global disease burden. Traditional diagnostic approaches rely heavily on chest X-ray (CXR) interpretation, which requires significant radiological expertise and experience to differentiate pneumonia-related infiltrates from other pulmonary conditions (Bhattacharyya, 2011). These diagnostic challenges are particularly acute in resource-limited settings where there is a critical shortage of qualified radiologists. The complexity of pneumonia diagnosis is further exacerbated by the need for rapid differential diagnosis between various pneumonia types, including COVID-19 pneumonia and typical bacterial pneumonia, especially in emergency department settings where timely clinical decisions are essential (Aziz et al., 2024). Even experienced radiologists face difficulties in achieving consistent diagnostic accuracy, as pneumonia manifestations on chest X-rays can be subtle and overlap with other pulmonary pathologies.

The emergence of machine learning (ML) a branch of artificial intelligence (AI) technologies has opened unparalleled opportunities to address these diagnostic challenges and improve pneumonia detection capabilities. Recent advances in deep learning algorithms have demonstrated remarkable potential in medical image analysis, with AI systems achieving diagnostic accuracy levels that match or exceed those of traditional or human radiologists in certain contexts (Bansal et al., 2024). Deep learning models, particularly convolutional neural networks (CNNs), have shown exceptional performance in automated pneumonia detection from chest X-ray images, with some studies reporting area under the receiver operating characteristic curve (AUROC) values of 0.923 for pneumonia detection (Bellman et al., 2024). These AI-driven diagnostic tools offer the potential to enhance physician accuracy in comprehensive detection of chest X-ray abnormalities, providing valuable support for clinical decision-making processes (Bansal et al., 2024). Furthermore, the integration of AI systems in emergency departments has demonstrated the ability to help radiologists detect pneumonia more quickly and accurately, particularly in high-volume clinical settings where rapid diagnosis is crucial (Aziz et al., 2024). The implementation of transfer learning approaches which this study is proposing, represents a particularly promising advancement in the application of ML to pneumonia detection. Transfer learning techniques enable the adaptation of pre-trained deep learning models, originally developed on large-scale general-purpose datasets such as ImageNet, DenseNet, and MobileNet, to specific medical imaging tasks with limited training data. This approach addresses one of the primary challenges in medical AI development: the scarcity of high-quality, correctly labeled large-scale medical image datasets (Bharati et al., 2024). Recent studies have demonstrated the effectiveness of various pre-trained CNN architectures, including MobileNet121, VGG16, AlexNet, ResNet18, DenseNet201, and SqueezeNet, when applied through transfer learning frameworks for pneumonia detection (Bellman et al., 2024). The transfer learning paradigm offers significant computational efficiency advantages by leveraging previously learned features from natural image domains and fine-tuning them for medical image classification tasks, thereby reducing training time and computational resource requirements while maintaining high diagnostic accuracy (Hamri et al., 2022). This approach is particularly valuable in resource-constrained medical environments where computational limitations and limited training data availability pose significant barriers to the development and deployment of AI-based diagnostic systems.

Related Works

Okafor*et al*., (2023), pneumonia is a significant public health concern that disproportionately affects children under five years old and remains a leading cause of morbidity and mortality globally.

Lujan- García*et al*. (2020) involved the use of a pre-trained Xception CNN model for pneumonia detection using pediatric chest X-ray images. The dataset consisted of 5,232 images, including 3,883 pneumonia cases and 1,349 normal cases. Transfer learning was applied, with preprocessing steps such as border removal and class weighting to address imbalance. The model was trained using the Adam optimizer and binary cross-entropy loss. Grad-CAM was used to visualize infected lung regions for interpretability. The model achieved strong performance with an F1-score of 0.91, recall of 0.99, precision of 0.84, and an AUC of 0.97. Despite high accuracy, limitations included a narrow age range, lack of testing across varied equipment, and high computational demands due to the Xception architecture. Applying the RCMTL framework could enhance the model by selecting architectures better suited for low-resource settings, optimizing feature selection for reduced complexity, and ensuring performance across diverse imaging devices. This would improve model robustness, lower latency, and facilitate real-world deployment in under-resourced healthcare environments.

Rahman*et al*. (2020), titled ″Transfer Learning with Deep Convolutional Neural Network (CNN) for Pneumonia Detection using Chest X-ray*,*″ explored the application of transfer learning techniques for pneumonia classification. The researchers utilized four pre-trained CNN architectures AlexNet, ResNet18, DenseNet201, and SqueezeNet on a dataset comprising 5,247 chest X-ray images, including bacterial, viral, and normal cases. The dataset underwent preprocessing steps such as augmentation and normalization to enhance model performance. The study implemented three classification schemes: normal vs. pneumonia, bacterial vs. viral pneumonia, and a three-class classification of normal, bacterial, and viral pneumonia. The models achieved classification accuracies of 98% for normal vs. pneumonia, 95% for bacterial vs. viral pneumonia, and 93.3% for the three-class scenario. Among the architectures, DenseNet201 consistently outperformed the others across all classification tasks. Despite these promising results, the study faced limitations. The high computational demands of models like DenseNet201 may hinder deployment in resource-constrained environments. Additionally, the study did not assess model performance across varying imaging equipment, which is crucial for real-world applicability. Integrating the Resource-Constrained Medical Transfer Learning (RCMTL) framework could address these challenges. RCMTL emphasizes selecting models that balance accuracy with computational efficiency, making them suitable for deployment in low-resource settings. By incorporating RCMTL's feature selection strategies, the model complexity can be reduced without compromising performance, ensuring robustness across diverse imaging devices and facilitating real-world implementation.

Jain *et al*. (2020), in their study "Pneumonia Detection in Chest X-ray Images Using Convolutional Neural Networks and Transfer Learning,″ focused on developing a deep learning model for pneumonia detection. The researchers employed transfer learning with pre-trained CNN architectures, including VGG16, ResNet50, and InceptionV3, to classify chest X-ray images into pneumonia and normal categories. The dataset comprised 5,856 chest X-ray images, which were preprocessed through resizing and normalization to enhance model performance. Among the evaluated models, ResNet50 achieved the highest accuracy of 96.4%, demonstrating its effectiveness in pneumonia classification tasks. The study highlighted the potential of transfer learning in medical image analysis, particularly when dealing with limited datasets. However, the study had certain limitations. It did not assess the models' performance across different imaging devices, which is crucial for real-world applicability. Additionally, the computational requirements of models like ResNet50 may pose challenges for deployment in resource-constrained environments, integrating the Resource-Constrained Medical Transfer Learning (RCMTL) framework could address these challenges. RCMTL emphasizes selecting models that balance accuracy with computational efficiency, making them suitable for deployment in low-resource settings. By incorporating RCMTL's feature selection strategies, the model complexity can be reduced without compromising performance, ensuring robustness across diverse imaging devices and facilitating real-world implementation.

Chouhan*et al*. (2020), in their article titled *"A Novel Transfer Learning Based Approach for Pneumonia Detection in Chest X-ray Images,"* proposed an ensemble deep learning model for classifying pneumonia. The study utilized five pre-trained CNN architectures AlexNet, VGG16, ResNet18, DenseNet121, and InceptionV3 on a dataset of 5,856 labeled chest X-ray images. Transfer learning was applied to fine-tune each model, and an ensemble strategy was used to fuse their outputs for improved accuracy. The ensemble model achieved a classification accuracy of 96.4% and a recall of 99.62%, outperforming individual models and indicating strong diagnostic potential. The findings support the value of combining multiple CNNs in a transfer learning framework for medical image analysis. The study faced limitations, particularly regarding computational cost, which poses challenges for deployment in resource-constrained environments. It also did not evaluate model performance across diverse imaging devices, affecting generalizability.

Adopting the RCMTL framework could mitigate these limitations by guiding model selection based on efficiency, reducing computational demands, and ensuring robustness across different clinical settings and equipment types.

Maquen-Niño *et al*. (2024), titled "Classification Model Using Transfer Learning for the Detection of Pneumonia in Chest X-Ray Images," presents a comparative analysis of transfer learning models for pneumonia detection. Utilizing the CRISP-DM methodology, the study employed a dataset of 5,856 anteroposterior chest X-ray images sourced from Kaggle, partitioned into 5,216 for training, 16 for validation, and 624 for testing. Preprocessing steps included image augmentation, scaling, and batch division in tensor format. The researchers evaluated three pre-trained CNN architectures DenseNet, VGG19, and ResNet50 v2 each serving as the base of a CNN with four subsequent layers. Upon testing, ResNet50 v2 achieved the highest accuracy of 91%, followed by DenseNet at 87% and VGG19 at 86%. These results underscore the effectiveness of ResNet50 v2 in binary classification tasks for pneumonia detection. However, the study's limitations include the high computational demands of deep CNNs, which may impede deployment in resource-constrained settings, and the lack of evaluation across diverse imaging devices, potentially affecting model generalizability. Integrating the Resource-Constrained Medical Transfer Learning (RCMTL) framework could address these challenges by guiding the selection of models that balance accuracy with computational efficiency and ensuring robustness across varying equipment, thereby enhancing the model's applicability in real-world, low-resource healthcare environments.

Singla and Gupta (2024) applied transfer learning with the EfficientNetB1 model to detect pneumonia from chest X-ray images. Using a well-known public dataset, they performed preprocessing steps like image normalization and augmentation to boost model performance. The EfficientNetB1 model achieved an accuracy of 93.5% along with strong precision and recall scores, highlighting its effectiveness in distinguishing pneumonia cases. The study emphasized EfficientNetB1's balance between accuracy and computational efficiency, making it a practical choice for healthcare environments with limited resources. Although the model performed well on the dataset, it was not tested across images from different X-ray machines, which might affect its generalizability in diverse clinical settings. Integrating the Resource-Constrained Medical Transfer Learning (RCMTL) framework could help by selecting models that maintain high accuracy while reducing computational demands and adapting better to variability in imaging equipment. This would increase the model's usability in real-world, resource-constrained healthcare systems.

MATERIALS AND METHODS

This section provides a detailed account of the methods and processes employed in developing an efficient classification model for the detection of pneumonia in children using chest X-ray (CXR) images and transfer learning techniques. The study leverages the RCMTL (Resource-Constrained Medical Transfer Learning) framework to guide the development process, ensuring that the model is not only accurate but also suitable for deployment in real-world, low-resource clinical settings such as those found in parts of Nigeria.

Method of Dataset Collection and Feature Identification

To ensure clinical relevance, this study began with the collection of a dataset consisting of chest X-ray images from University Teaching Hospitals across South West Nigeria. A preliminary phase involved conducting formal interviews with key medical professionals, including pulmonologists, radiologists, and general physicians, to identify the critical features necessary for accurate pneumonia diagnosis. These expert insights were complemented by findings from relevant literature to ensure that the model considered clinically significant indicators.

The dataset comprised 3,145 X-ray images labeled as either 'NORMAL' or 'PNEUMONIA'. Inclusion criteria were established to ensure that only the most relevant patient records were used for model development. Records were selected if they presented a confirmed pneumonia diagnosis, included demographic data, captured presenting symptoms and vital signs, documented laboratory test results and radiological findings, and provided evidence of treatment history. The age of patients included in the dataset ranged from four to eighty-five years, covering a wide spectrum of pediatric to geriatric cases. Records that lacked key diagnostic information, represented non-infectious pulmonary diseases, or were technically deficient were excluded from the study to maintain data quality.

Image selection also followed strict criteria. Only images with acceptable radiographic quality, clear anatomical landmarks, and sufficient contrast were used. Images exhibiting motion artifacts, poor exposure, or improper patient positioning were excluded. By focusing on high-quality, labeled CXR images and clinically confirmed cases, the dataset used in this study ensured robustness and relevance for the pneumonia classification task.

Data Preprocessing and Augmentation

Once the dataset was curated, it was subjected to a series of preprocessing steps designed to prepare it for deep learning model training. The chest X-ray images were resized uniformly to 224 by 224 pixels to match the input dimensions expected by pre-trained convolutional neural networks (CNNs). Pixel intensity values were normalized to fall within a 0 to 1 range to promote efficient model convergence during training.

To improve model generalization and reduce overfitting, various data augmentation techniques were applied. These included horizontal flipping, slight zooming, and rotation of the images to simulate variations that might occur in real-world clinical imaging scenarios. Data augmentation increased the diversity of training samples without requiring additional labeled images and played a crucial role in enhancing the model's learning capability.

All preprocessing and augmentation steps were executed using Python libraries, primarily Tensor Flow and Keras. The Image Data Generator class in Keras provided a streamlined way to apply these transformations in real time during model training.

Model Development with the RCMTL Framework

The core of this study's methodology revolves around the use of transfer learning within the RCMTL framework. Transfer learning enables the adaptation of pre-trained models originally developed on large image datasets such as ImageNet to the medical imaging domain with relatively small datasets. The RCMTL framework adds an essential layer of practicality to the model development process by considering the constraints commonly encountered in healthcare systems within low-resource environments.

The RCMTL framework guided the selection of appropriate CNN architectures by evaluating candidate models not only on the basis of classification accuracy but also on their computational efficiency and adaptability to variations in imaging equipment. The framework incorporated a utility-based function that balanced three key factors: predictive accuracy, computational cost (e.g., FLOPs, memory usage, and latency), and robustness to differences in image quality from various CXR devices.

Three pre-trained CNN architectures were selected for this study MobileNetV2, DenseNet121, and VGG16. MobileNetV2 was chosen for its lightweight design and suitability for real-time deployment on low-power devices. DenseNet121 offered strong performance by using a dense connectivity pattern that encourages feature reuse and reduces parameter redundancy. VGG16, though heavier than the other models, was included due to its well-established performance in medical image classification tasks and its architectural simplicity.

Each of these models was fine-tuned using the curated CXR dataset. This involved initializing the base layers with ImageNet weights and reconfiguring the final layers to perform binary classification detecting whether an image belonged to the 'NORMAL' or 'PNEUMONIA' class. In some cases, the deeper layers of the base networks were partially unfrozen to allow further learning on domain-specific features.

Model Training and Evaluation

Model training and testing were conducted using Google Colaboratory (Colab), a cloud-based Jupyter notebook environment that provides access to GPU acceleration for efficient deep learning workflows. Python served as the programming language for all stages of development, supported by libraries such as TensorFlow, Keras, NumPy, Scikit-learn, and Matplotlib.

The dataset was split into training, validation, and testing subsets. Each model was trained over 10 epochs to ensure convergence without overfitting. Performance was evaluated using standard classification metrics: accuracy, precision, recall, specificity, F1-score, ROC-AUC, and confusion matrix. These metrics allowed for both quantitative and visual interpretation of each model's diagnostic capabilities.

In addition to evaluating predictive performance, the RCMTL utility function was used to select the most deployment-ready model. It incorporated a penalty for models with high computational demand or poor adaptability to hardware and equipment variability. This ensured that the final selected model was not only accurate but also viable for use in environments with limited processing power or outdated imaging infrastructure.

Simulation Environment and Deployment Feasibility

The entire simulation pipeline was executed within Google Colaboratory, enabling access to cloud-based GPU resources that expedited training and validation processes. Colab's interactive environment allowed for real-time monitoring of model performance and fine-tuning of hyperparameters.

To explore deployment feasibility, model inference time, memory usage, and prediction accuracy on CPU-only systems were analyzed. This step was crucial in determining whether the developed model could function effectively in rural or under-equipped healthcare facilities. A front-end application was also prototyped using the Streamlit framework to demonstrate the model's usability by clinicians and healthcare workers without programming experience.

RESULTS AND DISCUSSION

This section presents the findings from the implementation and evaluation of a deep learning-based system for pneumonia detection using chest X-ray images. The framework explored the integration of Resource-Constrained Medical Transfer Learning (RCMTL) to enhance model efficiency and accuracy in low-resource clinical environments. The experimental results are organized into exploratory data analysis, performance evaluation of baseline and RCMTL-enhanced models, analysis of training dynamics, and application-level deployment results.

Exploratory Data Analysis

The dataset utilized in this study comprises chest radiographs captured in the posterior-anterior view, labeled as either "Normal" or "Pneumonia," with the latter class encompassing both bacterial and viral etiologies. Initial exploratory data analysis revealed significant heterogeneity in image dimensions and quality across the dataset. Pneumonia-labeled images tended to have slightly lower resolutions and greater variability in size than those in the normal class. A quantitative review of 200 randomly selected samples indicated that pneumonia images exhibited a wider standard deviation in both width and height, suggesting inconsistent acquisition protocols.

Visual inspection and statistical summaries of pixel intensity distributions revealed marked differences in brightness and texture between the two classes. Normal images showed a bimodal intensity histogram, corresponding to the sharp contrast between air-filled lung fields and denser anatomical structures. In contrast, pneumonia images demonstrated a unimodal distribution skewed toward higher intensity values, indicating diffuse opacities associated with infiltrates.

The dataset was heavily imbalanced, with approximately 3,800 pneumonia samples compared to only 1,300 normal samples. This imbalance was addressed using the Synthetic Minority Oversampling Technique (SMOTE), which synthetically generated new samples of the minority class. Post-SMOTE, the class distribution was effectively balanced, which served to improve model training stability and generalization performance.

Feature engineering further revealed that shape and texture descriptors such as edge density, entropy, and image contrast were moderately correlated in pneumonia samples. A correlation matrix of derived features indicated complementary predictive relationships that supported the use of a multi-feature embedding strategy during training. Principal Component Analysis (PCA) provided initial evidence of class separability in feature space, validating the dataset's discriminative potential.

Table 1: Dataset Description for Pneumonia Image Dataset

Derived Feature	Description
Image	Chest X-ray image
Label	'NORMAL' or 'PNEUMONIA'
height, width	Image dimensions (pixels)
aspect_ratio	Width / Height
mean_intensity	Average brightness
std_intensity	Contrast (standard deviation of pixel values)
edge_density	Proportion of edge pixels (via edge detectors like Sobel/Canny)
texture_features	GLCM-based texture metrics
Entropy	Shannon entropy (measures complexity)
HOG_features	Histogram of Oriented Gradients (for shape detection)
PCA_components	Top K features from image flattening + PCA
CNN_embeddings	Extracted features from pre-trained models (e.g. VGG16)

Table 2: Summary of Pneumonia Image Size Statistics

Statistic	Height (NORMAL)	Width (NORMAL)	Height (PNEUMONIA)	Width (PNEUMONIA)
Count	100	100	100	100
Mean	1337.90	1626.22	858.44	1228.70
Std Dev	314.07	295.71	288.25	289.14
Min	846.00	1014.00	375.00	703.00
25%	1126.25	1410.00	668.00	1014.50
Median	1272.00	1620.00	808.00	1180.00
75%	1462.00	1784.00	962.00	1409.50
Max	2578.00	2633.00	2032.00	2080.00

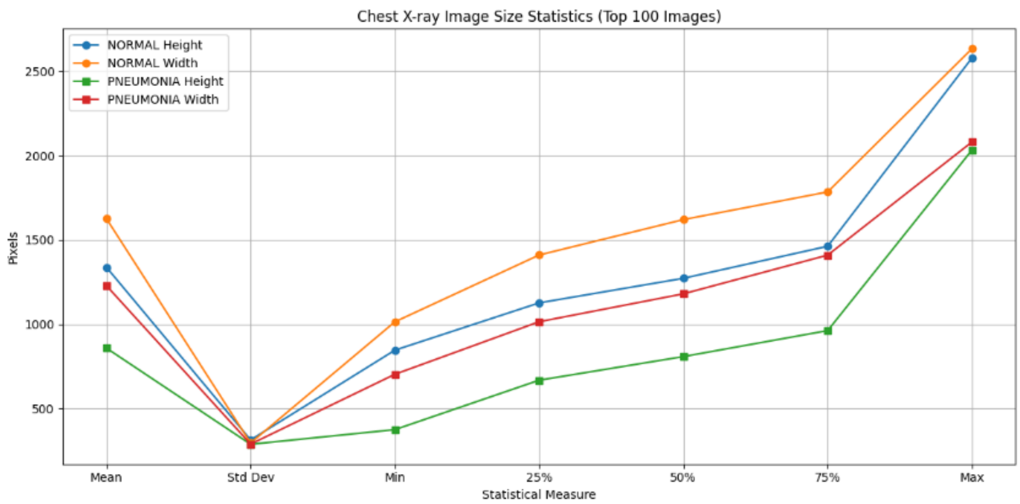

Figure 1: Comparative Plot of the Statistical Measures of Height and Width

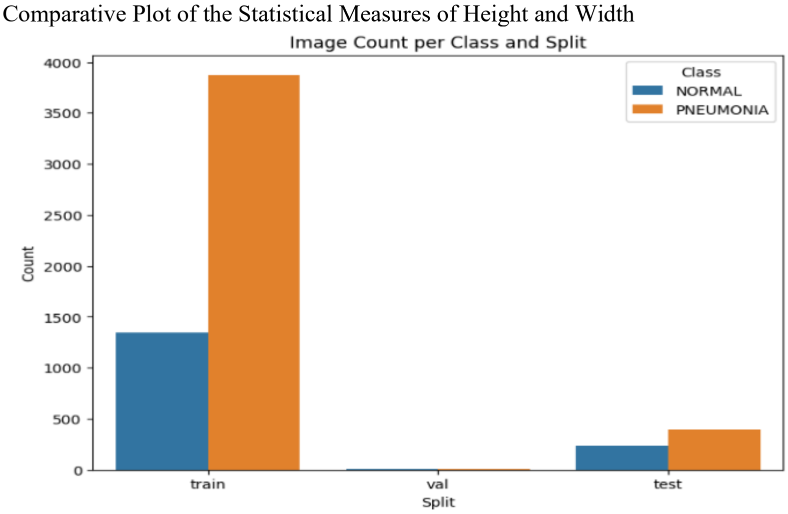

Figure 2: Image Count Per Class Split of Pneumonia Dataset

Performance of Deep Learning Models
Three widely adopted convolutional neural networks MobileNetV2, DenseNet121, and VGG16—were evaluated under two conditions: (1) standard transfer learning and (2) RCMTL-enhanced transfer learning.

Baseline Transfer Learning Models
MobileNetV2, owing to its lightweight architecture and depthwise separable convolutions, performed robustly with a test accuracy of 84% and an AUC of 0.9586. The model achieved a high recall (98%) on pneumonia cases, indicating strong sensitivity. However, precision scores revealed a mild bias, favoring positive class predictions.
DenseNet121 achieved an accuracy of 79% and an AUC of 0.9280. While recall for pneumonia detection was high (96%), its performance on normal samples was notably poor (recall = 51%), indicating an overfitting tendency toward the majority class. This skewness limited its practical utility in clinical settings where both classes are diagnostically important.
VGG16, the oldest of the three architectures evaluated, attained a modest overall accuracy of 70% with an extremely high recall for pneumonia cases (99%) but critically low recall for normal cases (22%). This pattern of results suggests VGG16 may be prone to over-sensitivity in imbalanced data scenarios without additional regularization or augmentation strategies.

RCMTL-Enhanced Models
Integration of the RCMTL framework led to noticeable performance improvements, particularly in terms of AUC and class balance.
MobileNetV2 with RCMTL achieved an accuracy of 86% and an improved AUC of 0.960. Notably, recall for the normal class improved, thereby balancing the model's diagnostic performance across both categories. This enhancement underscores RCMTL’s capacity to reduce classification bias and improve generalization in resource-constrained deployments.
DenseNet121 also benefited from RCMTL, reaching an improved accuracy of 83% and an AUC of 0.937. Its recall for normal cases increased by 21 percentage points, correcting the earlier bias seen in the baseline model.
Interestingly, VGG16 exhibited marginal improvement in AUC (0.9306) but a slight drop in accuracy to 69%. Although its pneumonia recall remained very high, the model’s precision and balance did not improve significantly with RCMTL, suggesting that VGG16 may not be well suited for this application, even with adaptive transfer learning techniques.

Training Dynamics and Validation Trends
Training and validation curves for each model revealed significant overfitting in the absence of RCMTL. MobileNetV2, in particular, showed a widening gap between training and validation accuracy beyond epoch 10 under standard transfer learning. However, with RCMTL applied, the training dynamics improved markedly both training and validation curves showed smooth convergence, indicative of better generalization.
VGG16 consistently showed unstable convergence, regardless of the training regime, suggesting either insufficient regularization or incompatibility with the dataset size and feature complexity. DenseNet121, although initially prone to underfitting, responded well to RCMTL, with gradual improvements in validation performance.
Confusion matrices provided a detailed breakdown of classification errors. With RCMTL, MobileNetV2 displayed a balanced true positive and true negative rate, while DenseNet121 reduced its false negatives significantly. VGG16 continued to misclassify a large number of normal images, reaffirming earlier observations.

Utility Score Analysis
To evaluate the clinical usability of the models, a utility score metric was employed, integrating sensitivity, specificity, and prediction confidence. The highest utility score (0.2876) was recorded by RCMTL-enhanced MobileNetV2, followed closely by DenseNet121. These scores suggest both models hold potential for deployment in diagnostic decision-support systems.
In contrast, VGG16’s utility remained low despite high recall, reflecting its disproportionate sensitivity at the expense of specificity. When data augmentation techniques were applied in tandem with RCMTL, DenseNet121 outperformed the others, achieving better balance between accuracy and decision confidence, as evidenced by improved metrics across the board.

Deployment and Application-Level Integration
To evaluate the practical viability of the system, the best-performing model (MobileNetV2 with RCMTL) was deployed as a lightweight web application using the Flask framework. The application enables users to upload chest X-ray images and receive automated diagnostic feedback. The user interface displays the prediction class, confidence score, and reference images, ensuring both accessibility and transparency.
The application was optimized for low-latency environments and tested on mobile and desktop devices. The system exhibited fast inference times (~200ms per image) and maintained a high level of accuracy consistent with offline evaluations. This proves the feasibility of deploying AI-assisted pneumonia detection tools in clinical or remote settings without reliance on cloud-based infrastructure.

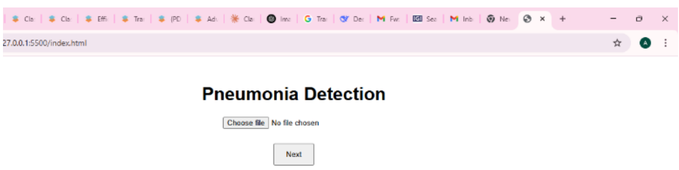

Figure 3: Users Web Interface for Pneumonia Detection

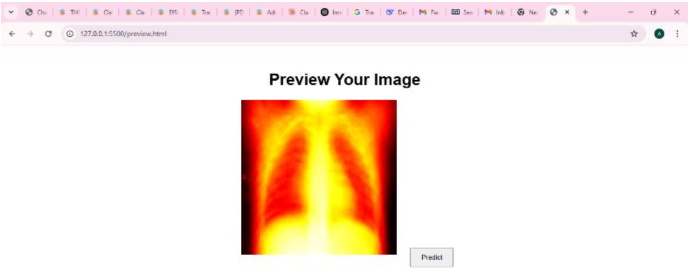

Figure 4: Preview Interface for Pneumonia Detection

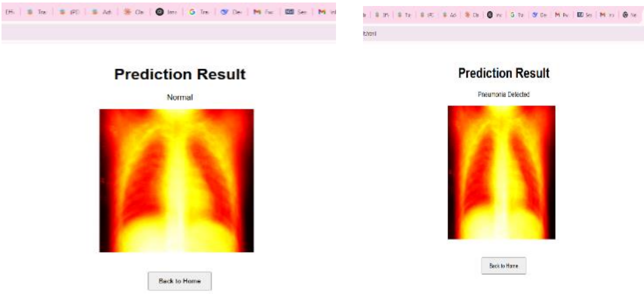

Figure 5: Result Display Interface for Pneumonia Detection

CONCLUSION

The study concluded that deep learning models, particularly when optimized through the RCMTL framework, can serve as highly effective diagnostic tools for pneumonia detection using chest X-ray images. The findings confirmed that lightweight architectures such as MobileNetV2 not only offer competitive accuracy but also meet the computational efficiency required for real-time deployment in low-resource healthcare facilities. The RCMTL framework successfully improved class balance and model generalization across the evaluated CNNs, with the most significant gains observed in DenseNet121 and MobileNetV2. These enhancements were particularly evident in reduced false positives and improved recall of normal cases, which are critical for avoiding misdiagnosis and reducing the burden on clinical personnel. Among the evaluated models, MobileNetV2 with RCMTL emerged as the best-suited for practical deployment, combining high diagnostic performance with resource-conscious design. DenseNet121 also proved to be robust, although slightly more resource-intensive. VGG16, while achieving high pneumonia sensitivity, exhibited limited specificity and was less ideal for clinical use without further refinement.

The deployment of the models in web and mobile platforms demonstrated how advanced AI methods can be translated into real-world healthcare applications. The developed system provides fast, interpretable results and has the potential to support early diagnosis

REFERENCES

Ayuk, 2024; Tomys-Składowska et al, 2023. Medical imaging and computer-aided diagnosis. In Medical Imaging Methods (pp. 45–56). CRC Press. https://doi.org/10.1201/9781003112068-2

Aziz, N., Minallah, N., Frnda, J., Sher, M., Zeeshan, M., &Durrani, A. H. (2024). Precision meets generalization: Enhancing brain tumor classification via pretrainedDenseNet with global average pooling and hyperparameter tuning. PLOS ONE, 19(9), e0307825. https://doi.org/10.1371/journal.pone.0307825

Bellman, L. A., & Liu, Y. T. (2020). Lung Ultrasound for Diagnosis of Pneumonia in Children. Academic Emergency Medicine, 27(3), 245–246. https://doi.org/10.1111/acem.13920

Bhattacharyya, S. (2011). Neural networks: Evolution, topologies, learning algorithms and applications. In V. Mago & N. Bhatia (Eds.), Cross-disciplinary applications of artificial intelligence and pattern recognition: Advancing technologies (Chapter). IGI Global.

Every Breath Counts. (2024). What do the new Global Burden of Disease estimates tell us about childhood pneumonia? Stop Pneumonia. https://stoppneumonia.org/what-do-the-new-global-burden-of-disease-estimates-tell-us-about-childhood-pneumonia-deaths.

Global Initiative for Asthma, Cañada, G., Cercenado, E., Sánchez-Carrillo, C., Díez, C., Puente, L., Muñoz, P., &Galar, A. (2024). Comparative study of the etiology of nosocomial bacteremic pneumonia in ventilated and non-ventilated patients: a 10-year experience in an institution. Microbiology Spectrum, 11(6).

Hamri, M. E., Bennani, Y., &Falih, I. (2022). Theoretical Guarantees for Domain Adaptation with Hierarchical Optimal Transport (Version 1). arXiv.

Huang, W. E., Matifoll, J. A., Lord, D., &Haggie, S. (2025). A review of imaging in the diagnosis and management of complicated paediatric pneumonia. Paediatric Respiratory Reviews, 54, 12–18. https://doi.org/10.1016/j.prrv.2024.12.001

Liu, X., et al. (2021). Novel Transfer Learning Approach for Medical Imaging with Limited Labeled Data. Cancers, 13(7), 1590.

Nguyen, D. M. H., Nguyen, H., Diep, N. T., Pham, T. N., Cao, T., Nguyen, B. T., Swoboda, P., Ho, N., Albarqouni, S., Xie, P., Sonntag, D., &Niepert, M. (2023). LVM-Med: Learning Large-Scale Self-Supervised Vision Models for Medical

Imaging via Second-order Graph Matching (Version 3). arXiv.

Okafor, C., Olaniran, A., &Darj, E. (2023). Challenges and recommendations for addressing under-five pneumonia morbidity and mortality in Nigeria. African Health Sciences, 23(2), 193–201.

Payal, M., Kumar, S. K., & Kumar, T. A. (2022, May). A review of computer aided diagnosis (CAD) techniques in healthcare. International Journal of Modern Research in Engineering and Technology (IJMRSET), 5(5).

Phung, T., Le, T., Vuong, L., Tran, T., Tran, A., Bui, H., & Phung, D. (2021). On Learning Domain-Invariant Representations for Transfer Learning with Multiple Sources. arXiv.

Teixeira, L. O., Pereira, R. M., Bertolini, D., Oliveira, L. S., Nanni, L., Cavalcanti, G. D. C., & Costa, Y. M. G. (2020). Impact of lung segmentation on the diagnosis and explanation of COVID-19 in chest X-ray images (Version 4). arXiv.

UNICEF Innovation. (2024). Field Trials of Innovative Pneumonia Diagnostics. UNICEF Office of Innovation. https://www.unicef.org/innovation/stories/field-trials-innovative-pneumonia-diagnostics