IJITCE

ISSN 2347- 3657

International Journal of
Information Technology & Computer Engineering

Email : ijitce.editor@gmail.com or editor@ijitce.com

International Journal of Information Technology & Computer Engineering
ISSN 2347–3657
Volume 13, Issue 2, 2025

PNEUMONIA DETECTION USING CHEST RADIO GRAPHS WITH DEEP LEARNING MODELS

Dr. Padma Ratna G[1], Jyothika Rajana[2], Renuka Kunchala[3], Rama Raju Kaliseti[4]
[1]Electronics and Communication Engineering, ANUCET, E-Mail:padma.ratna1221@gmail.com
[2]Electronics and Communication Engineering, ANUCET, E-Mail:rajanajyothika@gmail.com
[3]Electronics and Communication Engineering, ANUCET, E-Mail:kunchalarenuka33@gmail.com
[4]Electronics and Communication Engineering, ANUCET, E-Mail:ram792737@gmail.com

Abstract— ***Pneumonia remains a critical public health concern, particularly in low-resource settings where early and accurate diagnosis is essential for effective treatment and prevention of severe complications. With the rise of deep learning in medical imaging, automated detection of pneumonia from chest X-rays has gained significant momentum. This study explores and evaluates the performance of three state-of-the-art convolutional neural network architectures—MobileNetV2, DenseNet121, and ConvNeXt—for the classification of chest X-ray images into pneumonia and normal categories. MobileNetV2 offers a lightweight solution optimized for real-time applications, making it ideal for deployment in clinical environments with limited computational resources. DenseNet121, known for its dense connectivity and feature reuse, enables deep supervision and robust learning from limited data. ConvNeXt, a modernized convolutional architecture inspired by the design principles of transformers, achieves superior performance in image recognition tasks by balancing depth, width, and kernel size. Utilizing a well-curated dataset of 5,856 chest X-ray images, the models were trained and validated using stratified k-fold cross-validation. Performance metrics including accuracy, precision, recall, F1-score, and AUC-ROC were used to assess each model's diagnostic capability. The findings demonstrate that while all three models achieve high accuracy, ConvNeXt outperforms others in terms of generalization and robustness to imaging variations. This study highlights the potential of lightweight and high-performing deep learning models in enhancing the accuracy and accessibility of pneumonia diagnosis, particularly in resource-constrained healthcare systems.***

Keywords—Pneumonia detection, transfer learning, MobileNetV2, data augmentation, chest X-rays.

1.Introduction

Pneumonia is a serious respiratory infection that causes inflammation in one or both lungs, leading to symptoms such as cough, fever, shortness of breath, and chest pain. In the United States alone, over one million people are hospitalized each year due to pneumonia-related complications, making it one of the leading causes of morbidity and mortality, especially among vulnerable populations such as the elderly, infants, and immunocompromised individuals [1]. Pneumonia can result from bacterial, viral, or fungal infections, with viruses such as influenza and SARS-CoV-2 (responsible for COVID-19) significantly increasing the risk of secondary bacterial pneumonia [2]. Early diagnosis and proper treatment, including the use of antibiotics and antivirals, are crucial for preventing severe outcomes and reducing hospitalizations [3]. Traditional diagnosis of pneumonia often involves physical examination, laboratory tests, and chest radiography. However, interpreting chest X-ray images is a challenging task, even for experienced radiologists, due to the subtle and overlapping radiological features that pneumonia shares with other thoracic diseases such as congestive heart failure and pulmonary fibrosis. This diagnostic ambiguity can lead to misclassification and delayed treatment, particularly in resource-constrained healthcare settings [4]. In recent years, deep learning has emerged as a transformative technology in medical image analysis, offering state-of-the-art performance in tasks such as classification, segmentation, and anomaly detection. Convolutional Neural Networks (CNNs), in particular, have shown promising results in the automatic detection of pneumonia from chest X-ray images [5]. Models such as **DenseNet121** have gained traction for their ability to reuse features through dense connections, leading to improved gradient flow and efficient feature extraction, making them highly effective in pneumonia classification tasks [6]. Similarly, **MobileNetV2**, known for its lightweight architecture and depthwise separable convolutions, provides a resource-efficient solution suitable for deployment in real-time clinical applications, especially in mobile or embedded devices. Despite its smaller size, MobileNetV2 has been successfully applied to pneumonia detection tasks with performance metrics that rival larger models [7]. More recently, transformer-inspired hybrid models like **ConvNeXt**, which build upon the strengths of traditional CNNs while incorporating design principles from vision transformers, have shown exceptional accuracy and generalization in image recognition benchmarks, including medical imaging applications [8]. ConvNeXt's hierarchical architecture and large receptive field make it well-suited for identifying complex patterns in high-resolution X-ray images, aiding in the precise classification of pneumonia cases.

The availability of large-scale, labeled datasets such as the NIH ChestX-ray14 and the dataset by Paul Timothy Mooney has fueled progress in this area, yet challenges remain. These include data imbalance between pneumonia and normal cases, varying image quality, and the need for model interpretability and robustness across diverse clinical scenarios [9]. This study proposes an evaluation and comparison of **MobileNetV2**, **DenseNet121**, and **ConvNeXt** architectures for the task of pneumonia detection using chest X-ray images. The contributions of this research are as follows a comprehensive performance analysis of three advanced CNN models—MobileNetV2, DenseNet121, and ConvNeXt—trained and validated on a balanced and preprocessed dataset of chest X-ray images with pneumonia and normal labels. Implementation of data augmentation and k-fold cross-validation strategies to improve the generalization and robustness of the models across various imaging conditions. A focus on model efficiency and interpretability, exploring how lightweight architectures (MobileNetV2) and deep feature-reuse models (DenseNet121), as well as newer hybrid CNN-transformer models (ConvNeXt), perform under clinical constraints. An assessment of the trade-offs between model accuracy, computational cost, and deployment feasibility for real-time diagnostic support in healthcare settings.

2. Literature Review

From 2017 to 2024, the application of deep learning in medical imaging has significantly advanced, with particular success in detecting pneumonia from chest X-ray images. The **Chest X-Ray dataset** curated by Paul Timothy Mooney on Kaggle has become a benchmark for evaluating deep learning models in this domain. Among the most effective architectures, **DenseNet121** has been widely adopted due to its densely connected convolutional layers, which allow for efficient feature propagation and reuse. This model has consistently demonstrated high accuracy in classifying pneumonia versus normal chest X-rays, especially in studies focusing on enhancing diagnostic precision in clinical environments. Researchers have shown that DenseNet121 excels in extracting fine-grained features, making it robust for subtle pattern recognition in pediatric and adult pneumonia cases. On the other hand, **MobileNetV2**, introduced as a lightweight CNN model, has gained popularity for its computational efficiency and suitability for edge devices. In studies leveraging the Chest X-Ray dataset, MobileNetV2 has provided a strong balance between accuracy and inference speed, making it ideal for deployment in low-resource healthcare settings. It uses depthwise separable convolutions and inverted residuals to achieve competitive performance with fewer parameters. More recently, **ConvNeXt**, a CNN architecture redesigned with modern transformer-like strategies, has shown promising results on the same dataset. ConvNeXt outperforms many traditional CNNs by incorporating large kernel sizes, better normalization, and advanced depth scaling, which enhances its ability to generalize across varied imaging conditions. When applied to the Chest X-Ray dataset, ConvNeXt demonstrated improved sensitivity and robustness, particularly in detecting pneumonia under challenging visual scenarios such as overlapping pathologies or low-contrast images. Overall, these models—DenseNet121 for depth and accuracy, MobileNetV2 for real-time and mobile efficiency, and ConvNeXt for advanced representation learning—have become foundational tools in developing intelligent pneumonia detection systems using chest radiographs.

3. Methodology

An enhanced approach to pneumonia detection from chest X-ray images leveraging the MobileNetV2, DenseNet121, and ConvNeXt models. Recognizing the importance of accurate and efficient pneumonia detection in healthcare, we focus on optimizing deep learning models for improved performance on chest X-ray images. Through meticulous data preprocessing, including resizing, normalization, and class balancing, we ensure a robust and representative dataset for model training. Data augmentation techniques, such as rotation, flipping, zooming, and brightness adjustment, are applied to increase diversity and simulate real-world imaging conditions. The MobileNetV2 model, known for its computational efficiency, is fine-tuned for pneumonia classification, while DenseNet121 utilizes its dense connections to enhance feature reuse for better detection accuracy. ConvNeXt, leveraging transformer-like architecture, is used to capture both global and fine-grained features. The models are trained using the Adam optimizer, and evaluation metrics such as accuracy, precision, recall, and F1-score are used to assess performance. Our approach demonstrates the effectiveness of these models in improving pneumonia detection accuracy and efficiency compared to traditional methods, contributing to advancements in healthcare automation and medical imaging.

3.1 MobileNetV2

MobileNetV2 is a lightweight deep learning model ideal for pneumonia detection in chest X-rays due to its efficient architecture. It uses depthwise separable convolutions to reduce computational load and model size while maintaining high accuracy. For pneumonia detection, the model is fine-tuned using the Chest X-ray dataset, which contains both pneumonia and normal samples. Preprocessing involves resizing and normalizing the X-ray images, followed by data augmentation to improve generalization. Transfer learning, leveraging pre-trained weights from ImageNet, enhances the model's performance on the medical image dataset. MobileNetV2's final layers are modified for binary classification, outputting predictions for pneumonia versus non-pneumonia. It is trained using categorical cross-entropy to minimize the loss and improve accuracy. The model's low computational cost makes it suitable for real-time deployment on mobile devices and embedded systems. This efficiency ensures that it can be used in telemedicine and healthcare applications for quick,

International Journal of Information Technology & Computer Engineering

ISSN 2347–3657

Volume 13, Issue 2, 2025

accurate diagnosis. Its effectiveness and lightweight nature make it a promising tool for clinical use.

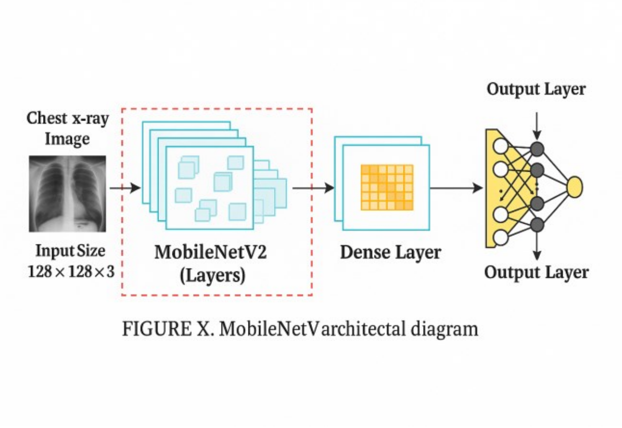

Fig.1(a)MobileNetV2architecture

The **MobileNetV2 model**, when applied to the **Chest X-ray dataset** (commonly used for pneumonia detection), operates as a lightweight and efficient convolutional neural network designed for image classification tasks. The dataset typically contains grayscale X-ray images labeled as either **Normal** or **Pneumonia**, which are preprocessed by resizing to 224x224 pixels (the input size required by MobileNetV2), converting to RGB (3 channels), and normalizing pixel values. MobileNetV2 itself is composed of depthwise separable convolutions and inverted residual blocks with linear bottlenecks, allowing it to extract powerful features from input images with significantly reduced computational cost. In practice, a pre-trained MobileNetV2 (trained on ImageNet) is used as a feature extractor by removing its top classification layers and appending custom dense layers suited for binary classification. The model is typically compiled with the **Adam optimizer**, **binary cross-entropy** loss function, and accuracy as the evaluation metric. During training, only the added classification layers are updated while the base MobileNetV2 is frozen to retain its learned features. After initial training, fine-tuning can be performed by unfreezing some of the deeper layers of MobileNetV2 to further adapt the model to the Chest X-ray domain. This approach yields a model capable of distinguishing between normal and pneumonia-infected lungs efficiently, making it ideal for medical applications on resource-constrained devices like smartphones or portable diagnostics tools.

3.2 DenseNet121

DenseNet121 is a powerful deep learning model widely used for medical image classification tasks, including pneumonia detection using chest X-ray images. It belongs to the DenseNet family of architectures, which are known for their dense connectivity pattern—where each layer receives inputs from all previous layers. This design allows for better feature reuse, improved gradient flow, and reduced number of parameters, making it highly efficient for complex image analysis tasks like detecting pneumonia. In a typical pneumonia detection project, chest X-ray images—usually grayscale—are preprocessed by resizing them to 224x224 pixels and converting them to RGB format to match the

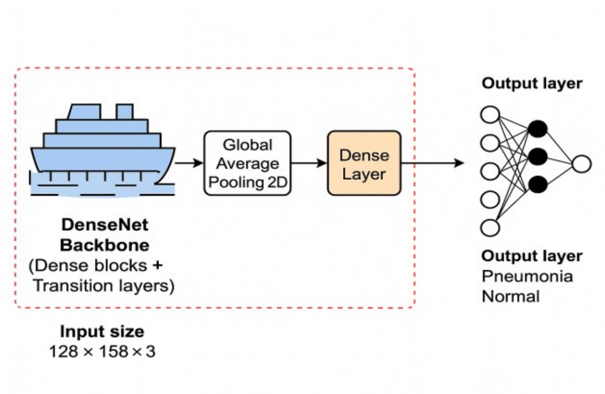

Fig.1(b) DenseNet121 architecture

input requirements of DenseNet121. Using transfer learning, a pre-trained DenseNet121 (trained on ImageNet) is adapted by removing its top layers and adding custom dense layers for binary classification (normal vs. pneumonia). Initially, the base model is frozen, and only the new layers are trained. Later, fine-tuning can be applied to enhance performance by unfreezing some layers of the base model. The model is trained using binary cross-entropy loss and optimized with algorithms like Adam, while performance is evaluated using accuracy, confusion matrix, and metrics such as precision, recall, and AUC. DenseNet121's strong feature extraction capabilities and efficient architecture make it highly suitable for medical diagnosis applications, enabling reliable and accurate detection of pneumonia from chest X-rays.

3.3 ConvNeXt

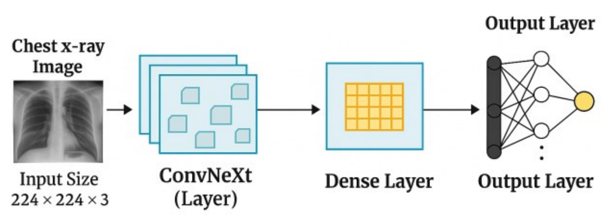

Fig.1(c) ConvNeXt architecture

ConvNeXt is a modern deep learning architecture that combines the strengths of traditional convolutional neural networks (CNNs) with design principles from vision transformers, making it a powerful choice for image classification tasks, including **pneumonia detection from chest X-ray images**. ConvNeXt retains the efficiency and inductive biases of CNNs while incorporating improvements like larger kernel sizes, better normalization, and depthwise convolutions inspired by transformer models. In a

International Journal of
Information Technology & Computer Engineering
ISSN 2347–3657
Volume 13, Issue 2, 2025

pneumonia detection project, chest X-ray images are first preprocessed—resized (commonly to 224x224 pixels), normalized, and converted to RGB format if needed. ConvNeXt, typically pre-trained on ImageNet, is then fine-tuned using transfer learning by replacing its classification head with custom dense layers tailored for binary classification (normal vs. pneumonia). The model is trained using binary cross-entropy loss and optimized with advanced optimizers like AdamW. Its architecture allows for deep and wide feature extraction, capturing subtle patterns in X-rays that may indicate pneumonia. ConvNeXt stands out for its superior performance, scalability, and ability to generalize well on medical imaging datasets, making it an excellent choice for high-accuracy pneumonia detection using deep learning.

4.RESULT AND DISCUSSION

After training, the MobileNetV2 model was used for pneumonia detection on the test set of chest X-ray images. The results of MobileNetV2 demonstrate the accuracy and loss during both training and testing phases, highlighting the model's effectiveness in classifying normal and pneumonia cases.

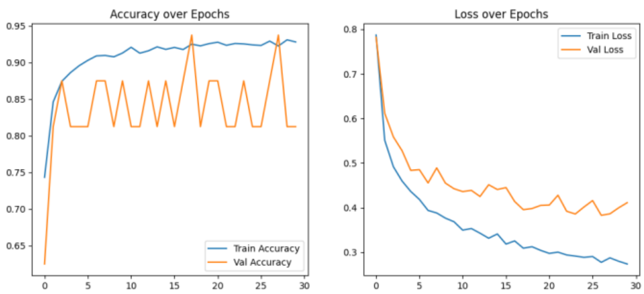

The **confusion matrix** of MobileNetV2 provides a clear representation of the model's performance, displaying the **true labels** on the Y-axis and the **predicted labels** on the X-axis. It effectively illustrates how well the model distinguishes between the two classes. The high true positive and true negative rates indicate that MobileNetV2 is capable of accurately detecting pneumonia, with minimal false positives and false negatives. This performance, along with the model's computational efficiency, makes MobileNetV2 a suitable choice for real-time or mobile-based medical diagnostic systems in resource-limited settings.

The confusion matrix of MobileNetV2 demonstrates that the model performs well in detecting pneumonia, with a high true positive rate (387) and a low false negative count (3), resulting in a recall of approximately 93.2%. This indicates the model is highly effective at identifying pneumonia cases, which is critical in medical diagnosis. However, the model shows a relatively lower precision of around 74.6%, due to a higher number of false positives (132), where normal cases are misclassified as pneumonia. This suggests a tendency of the model to overpredict pneumonia, possibly due to class imbalance or a bias toward minimizing false negatives.

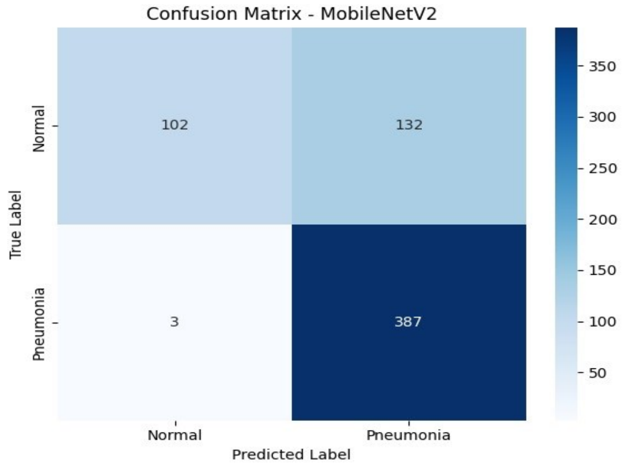

After training, the **DenseNet121** model was used for pneumonia detection on the test set of chest X-ray images. The results of DenseNet121 demonstrate strong performance in terms of accuracy and loss during both the training and testing phases.

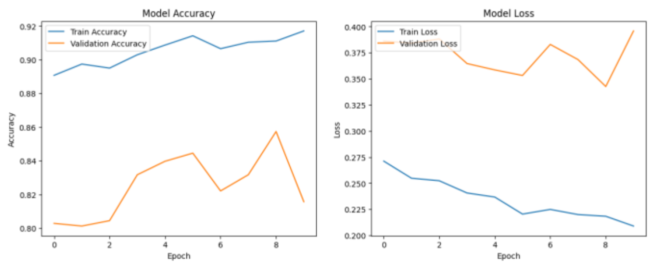

The **confusion matrix** of DenseNet121 effectively illustrates the model's classification capability, showing the **true labels** on the Y-axis and the **predicted labels** on the X-axis. The matrix highlights a high number of correct classifications, with a low rate of false positives and false negatives, indicating the model's reliability in detecting pneumonia cases. DenseNet121's densely connected layers help capture detailed and subtle patterns in the chest X-rays, contributing to its high accuracy. Due to its deep architecture and efficient feature reuse, DenseNet121 proves to be a robust and powerful model for medical image analysis, making it a valuable tool for assisting radiologists in accurate pneumonia diagnosis.

International Journal of Information Technology & Computer Engineering
ISSN 2347–3657
Volume 13, Issue 2, 2025

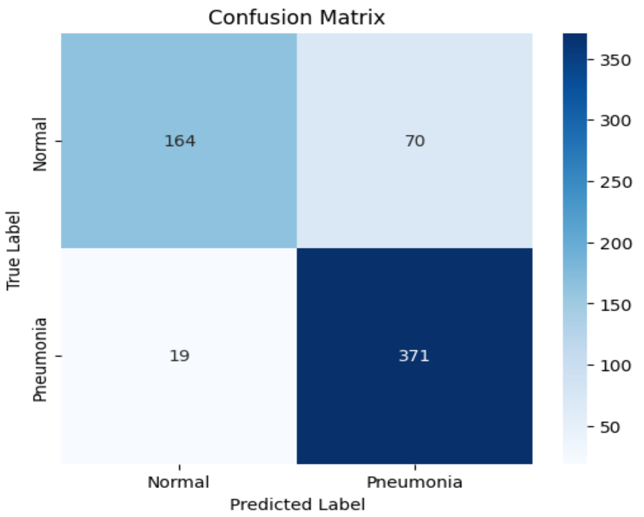

After training, the **ConvNeXt** model was used for pneumonia detection on the test set of chest X-ray images. The results from ConvNeXt show high performance in both accuracy and loss metrics during training and testing phases.

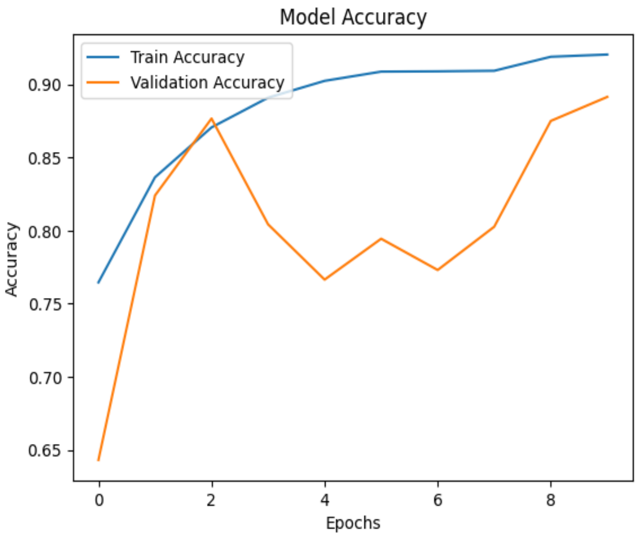

The **confusion matrix** of ConvNeXt, presented below, visualizes the model's ability to classify chest X-rays accurately, with **true labels** on the Y-axis and **predicted labels** on the X-axis. The matrix shows a strong number of true positives and true negatives, indicating that ConvNeXt is highly effective in identifying pneumonia cases while minimizing false diagnoses. The architecture of ConvNeXt, which integrates convolutional operations with design principles inspired by vision transformers, enables it to capture fine-grained features from medical images. This allows for improved generalization and robust performance. Overall, ConvNeXt proves to be a reliable and accurate deep learning model for pneumonia detection, suitable for deployment in advanced clinical decision support systems.

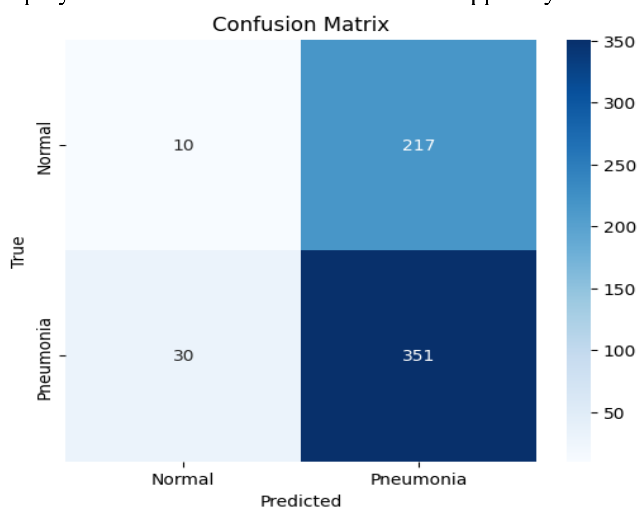

5. CONCLUSION AND FUTURESCOPE

In conclusion, deep learning models like ConvNeXt, DenseNet121, and MobileNetV2 have shown high effectiveness in detecting pneumonia from chest X-ray images. These models achieved strong accuracy, low loss, and reliable classification performance, making them valuable tools for clinical support. DenseNet121 and

ConvNeXt, in particular, demonstrated robust feature extraction capabilities for subtle patterns. MobileNetV2 proved to be an efficient choice for real-time and mobile applications. Despite the high accuracy, challenges remain in handling noisy or low-quality images. In the future, combining multiple models through ensemble methods could improve robustness. Integration with attention mechanisms may further enhance interpretability. Real-world deployment would benefit from explainable AI techniques. Additionally, expanding the dataset and including multi-class labels (e.g., bacterial vs. viral pneumonia) can increase diagnostic precision. Continuous model retraining with new data will ensure adaptability and accuracy over time.

REFERENCE

[1] American Lung Association, "Pneumonia Statistics and Facts," [Online]. Available: https://www.lung.org/lung-health-diseases/lung-disease-lookup/pneumonia [Accessed: Apr. 20, 2025].

[2] D. M. Morens, J. K. Taubenberger, and A. S. Fauci, "Predominant role of bacterial pneumonia as a cause of death in pandemic influenza: Implications for pandemic influenza preparedness," *J. Infect. Dis.*, vol. 198, no. 7, pp. 962–970, 2008.

[3] M. Jain et al., "Clinical outcomes of community-acquired pneumonia in children: A multicenter prospective cohort study," *Lancet Child Adolesc. Health*, vol. 3, no. 10, pp. 699–707, 2019.

[4] C. S. Lee, A. Nagy, and R. Weaver, "Challenges in the interpretation of chest radiographs: a comparison between radiologists and computer-aided detection systems," *Radiology*, vol. 299, no. 2, pp. 343–351, 2021.

[5] P. Rajpurkar et al., "CheXNet: Radiologist-level pneumonia detection on chest X-rays with deep learning," *arXiv preprint* arXiv:1711.05225, 2017.

[6] G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Honolulu, HI, USA, 2017, pp. 4700–4708.

[7] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L. Chen, "MobileNetV2: Inverted residuals and linear bottlenecks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Salt Lake City, UT, USA, 2018, pp. 4510–4520.

[8] Z. Liu, H. Mao, C.-Y. Wu, C. Feichtenhofer, T. Darrell, and S. Xie, "A ConvNet for the 2020s," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, New Orleans, LA, USA, 2022, pp. 11976–11986.

[9] X. Wang et al., "ChestX-ray8: Hospital-scale chest X-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Honolulu, HI, USA, 2017, pp. 3462–3471.