Journal of Population Therapeutics & Clinical Pharmacology

RESEARCH ARTICLE
DOI: 10.53555/bahkys10

# DEEP LEARNING APPROACHES FOR MEDICAL IMAGING ANALYSIS: TRANSFORMING DISEASE DETECTION AND MONITORING

**Dr. A.Mohideen Ashraf, MDRD[1*], Dr. S. SUMEENA, DMRD, DNB[2], Preetish Kumar Panigrahy[3]**

[1*]Associate Professor in Government Kilpauk Medical College Hospital, Chennai, ashrafmddmrd@gmail.com
[2]Associate Professor in Government Kilpauk Medical College Hospital, Chennai, drsumeena@gmail.com
[3]Assistant Professor, MNR Medical College, Orcid ID: 0000-0002-8073-6354, preetishmedico87@gmail.com

**Corresponding author:** Dr. A.Mohideen Ashraf, MDRD
[*]Associate Professor in Government Kilpauk Medical College Hospital, Chennai, ashrafmddmrd@gmail.com

**Abstract**

**Background:** Deep learning has been a revolution in the medical imaging sector, as it has offered improved accuracy and reproducibility of identifying diseases and their follow-ups. Pneumonia is one of the most widespread morbid and mortal diseases in the world that has radiographic similarities between bacterial and viral etiologies, making it difficult to diagnose.

**Objective:** This study aimed to compare the two deep learning models, ResNet-50, which is a convolutional neural network, and ViT-B/16, which is a transformer-based model in classifying images of chest X-rays based on bacterial pneumonia, viral pneumonia, and normal ones and also establishing the interpretability of their predictions.

**Methods:** A sdataset of X-rays of the chest containing bacterial, viral and normal X-rays was used. During preprocessing, normalization and augmentation was performed. ResNet-50 and ViT-B/16 models were trained, optimized and validated with stratified folds. Evaluation metrics included accuracy, F1-score, sensitivity, specificity, and AUC. Explainability was assessed using Grad-CAM to visualize clinically relevant regions.

**Results:** ResNet-50 achieved an accuracy of 87.2% with a macro F1-score of 0.86, while ViT-B/16 outperformed with an accuracy of 90.1% and a macro F1-score of 0.88. Bacterial pneumonia and normal classes were reliably detected, whereas viral pneumonia remained the most challenging category. Grad-CAM confirmed that both models focused on lung regions corresponding to pathological abnormalities, with ViT-B/16 demonstrating broader contextual attention.

**Conclusion:** Transformer-based deep learning architectures provide superior performance and interpretability compared to traditional CNNs, underscoring their potential to enhance disease detection and monitoring in medical imaging.

**Keywords:** Deep learning, Pneumonia detection, Chest X-ray, Medical imaging, Vision Transformer

## 1. Introduction

Medical imaging plays a key role in healthcare today, providing non-invasive tools for detecting, diagnosing and monitoring many diseases. X-rays, CT, MRI, histopathology slides etc modalities have evolved in terms of diagnostic speed, diagnostic accuracy yet with the high rate of increase in imaging data, radiologists have become overwhelmed and there is the problem of interpreter variability and risk of diagnostics. In order to resolve these limitations, artificial intelligence (AI) or deep learning (DL) has been extensively applied to the analysis of medical images, and it has provided opportunities to enhance disease detection and monitoring [1]. Machine learning, in particular, neural network architectures are applied in DL to learn hierarchical representations of data automatically without manually-constructed features. This is the capability that has made it be specifically applied in the medical imaging where it is important to be able to capture more complex space and contextual relationships [2]. Convolutional neural networks (CNNs) have been widely used in solving such problems as segmentation and classification and have shown decent results in modalities such as MRI, CT and X-rays [3]. Indicatively, CNN based systems have performed better in tumor segmentation, lesion detection and organ classification [4] as proved to be superior to many other conventional methods.

DL has revolutionary healthcare impacts. Many architectures like ResNet, DenseNet, EfficientNet have increased the scalability and accuracy and are now able to analyze large data sets with an accuracy that is now strictly clinically relevant [5]. Besides the CNNs, the transformer-based models, which were first introduced in the natural language processing community, are now being sought after in the medical imaging community too. The concept of the Vision Transformers (ViTs) is that these transformers consider the long-range dependences, thus the analysis of the complex patterns and subtle abnormalities that CNNs might miss is quite appropriate [6]. Along with diagnosis, DL may also help with other general clinical applications such as real-time monitoring and prognosis forecasting. In particular, IoT-based systems are being designed with the addition of DL that will be used to monitor patients more efficiently in order to make timely interventions [7]. The fusion and transfer learning of multimodal data is also under research in an attempt to combine imaging data with electronic health records and genetic data towards the creation of customized medicine [8].

In spite of such advances, there are still challenges. The availability of large annotated medical datasets is one of the biggest challenges since expert labeling is time-consuming and likely to be affected by inter-observer variability [9]. The other one is a lack of generalizability: models that have been trained on data of one single institution might not be effective on different populations or imaging protocols [10]. Moreover, it is also said that DL models tend to be black boxes having low interpretability which is a concern when it comes to clinical adoption [11]. This has prompted increased attention to explainable AI approaches that explain how models come up with the predictions as well as making sure that the results agree with clinical reasoning [12]. The detection of pneumonia using chest X-rays is a good example to test these opportunities and challenges. Pneumonia has been a major health challenge of the world particularly among the most vulnerable groups which include children and the elderly. Radiographic features such as opacities and consolidations are very important in the diagnosis of pneumonia, but it is usually challenging to differentiate between bacterial and viral pneumonia because of the similarity of features. This problem of diagnosis highlights the necessity to have a powerful DL model that could detect and enhance subtype differentiation automation [13].

There are larger questions of generalizability to other imaging tasks with the use of DL to classify pneumonia. The experience learned with chest radiography may be applied to other related fields like in MRI based tumor imaging, CT nodule identification, and grading cancer based on histopathology. The cross-modality promise points to the fact that DL is a disruptive technology in terms of improving the quality of diagnosis and workflow itself, as it can be used to improve the quality of diagnostic results as well as workflow efficiency. Here, the current paper explores the problem of pneumonia detection through the deep learning methods. Two state-of-the-art ResNet-50, as a CNN-based method and ViT-B/16, as a transformer-based method were tested in classifying the chest X-rays as bacterial pneumonia, viral pneumonia and normal. The research does not focus on the evaluation of

performance only but also on its interpretability, with the help of Gradient-weighted Class Activation Mapping (Grad-CAM), confirming the correspondence between predictions and clinically relevant lung areas. The specific objectives of this study are:

1. To evaluate and compare the performance of ResNet-50 and ViT-B/16 in classifying chest X-rays into bacterial pneumonia, viral pneumonia, and normal categories.
2. To analyze class-specific strengths and limitations, particularly in detecting viral pneumonia, which poses diagnostic challenges due to overlapping features.
3. To incorporate and assess explainability methods (Grad-CAM) to validate the clinical plausibility of predictions and enhance interpretability.

## Methodology

### 2.1 Overall Framework

The research article is based on the systematic experimental research design, which presupposes the systematic investigation of the effectiveness and viability of the deep learning models in the sphere of medical imaging. The first step in the pipeline involves the selection of a clinically relevant dataset and then proceeds to preprocessing steps, which are supposed to ensure consistency of samples and facilitate generalization. The deep learning models, the convolutional model and the transformer-based model are subsequently trained and optimized on the dataset. The models are trained and optimized and validation strategies are introduced in order to monitor the learning process and avoid overfitting. In this case, however, the assessment stage is conducted in a systemic manner according to multiple performance indicators, which ensures global and class-based accuracy. Finally, the decision-making process of the models is explained using explainability methods, which is highly critical to clinical trust and adoption. This is a complete end-to-end workflow that provides the right balance between technical rigor and clinical applicability Figure 1.

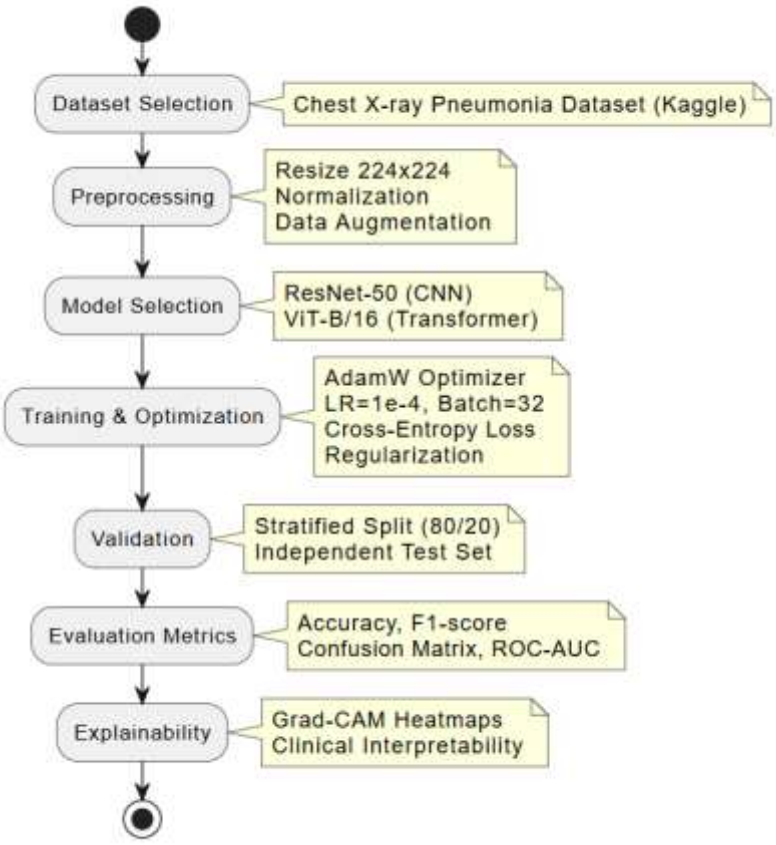


**Figure 1: Workflow for Deep Learning-Based Chest X-ray Pneumonia Classification**

### 2.2 Dataset

The data in this paper is the Chest X-ray Pneumonia Dataset on Kaggle [14], which consists of 5,216 chest radiographs altogether. These images are classified into three different diagnostic categories as: Bacterial Pneumonia (2,530 images), Viral Pneumonia (1,345 images), and Normal or healthy controls (1,341 images). The data is then organized into training, validation, and test folds where each has subfolders of the different classes to facilitate the supervised learning. Each of the pictures is in the JPEG format and depicts either anterior and posterior respectively of the chest radiographs, which are the most frequent projections in clinical practice. The presence of bacterial and viral subtypes of pneumonia is a valuable addition to clinical practice since it is crucial to identify the difference between the two since it affects treatment choices, especially in the use of antibiotics. The equal population of both viral and regular classes, as well as the greater size of a bacterial sub-population, reflect the actual distribution of diseases in the real world and are a powerful resource to train deep learning models.

### 2.3 Preprocessing

A sequence of preprocessing operations was used to standardize all the samples and prepare the data to be used in deep learning models. All the images were downsampled to 224 x 224 pixels, which is the resolution that strikes a good balance between computation speed and retention of meaningful diagnostic information. Intensities of pixels were clustered to the (0,1) scale and standardized by ImageNet mean and standard deviation values to align with the settings of the pre-trained models. Random horizontal flipping and small-angle rotations are data augmentation methods that were used on the training images to create variability and mitigate overfitting. Such augmentations model real-world variability in radiographic imaging, e.g. variability in patient positioning or settings during acquisition. Significantly, the validation and test sets were not manipulated in order to make a fair evaluation of model performance.

### 2.4 Models

Two deep learning models were chosen to possess complementary capabilities in medical image analysis. The first one is ResNet-50, a convolutional neural network (CNN) architecture that uses residual connections to reduce the vanishing gradient problem to allow deeper and more efficient feature extraction. ResNet-50 is a standardized baseline in medical imaging studies, which has the potential to learn hierarchical image representation. The second architecture is the Vision Transformer (ViT-B/16), a transformer-based architecture that takes images in the form of sequences of patches and uses self-attention mechanisms to model long-range dependencies between them. It is a more modern paradigm shift in computer vision, and has had state-of-the-art performance in several areas. Both models were trained using pre-trained weights of ImageNet, enabling transfer learning to be used, substantially decreasing the training time and computational load in favor of using the smaller medical dataset.

### 2.5 Training and Optimization

The PyTorch framework was used to train the models with the use of the GPU acceleration to make them more computationally efficient. The training was performed by AdamW with an initial learning rate of 1e-4, batch size of 32 and weight decay of 1e-4. These hyperparameters have been chosen according to the best practices in the tasks of image classification. The training procedure involved three epochs and balanced convergence and the computational resources available without overfitting by terminating the process at a certain stage. This loss operation was chosen because the cross-entropy loss operation is rather suitable when there are several classes such as bacterial pneumonia, viral pneumonia, and normal cases. Regularization like dropout layers in the network, image-level augmentations and early stopping were employed to improve generalization. Along with this mixed-precision training, the acceleration of computation was used to be fast and accurate.

### 2.6 Validation

The model validation was a significant phase that guaranteed that there was a vigorous and objective evaluation of the model. When no separate validation folder had been successfully specified, a training set was stratifiedly split into 80 / 20 (with stratified splitting) to form a validation set. This stratification allowed all the three types of diagnostic data to be represented in the validation data proportionally. This sub-set was monitored on model performance during every epoch and to be applied during tuning of the hyperparameters to avoid overfitting. In order to evaluate it eventually, held out test set was employed, which provided an independent measure of the model generalizability. The study conducted the internal validation subset of the study and the external test subset to conclude that the results obtained were actually the true picture of the behavior of the models with the unobservable data in the real world.

### 2.7 Evaluation Metrics

Performance was measured by accuracy and macro F1-score and the accuracy and macro F1-score was balanced among classes. Confusion matrices were generated to visualize misclassifications, while ROC curves and AUC values were reported using a one-vs-rest strategy. Both per-class and micro-averaged AUC scores were considered to capture discriminative ability.

### 2.8 Explainability

Model interpretability was evaluated using Grad-CAM, which generates heatmaps highlighting image regions influencing predictions. Visualizations were superimposed on chest X-rays for both ResNet-50 and ViT-B/16 to confirm focus on clinically relevant lung areas. Representative validation images demonstrated that both models localized pathological features effectively, enhancing clinical trust and supporting potential adoption.

## 3. Results

### 3.1 Overall Model Performance

The experimental evaluation demonstrated that both convolutional and transformer-based deep learning models achieved strong predictive performance in classifying chest X-rays into bacterial pneumonia, viral pneumonia, and normal categories. As indicated in Table 1, ResNet-50 and ViT-B/16 had a high overall accuracy of 87.2% and 90.1%, respectively, and the macro F1-score of 0.86 and 0.88, respectively. These results demonstrate the usefulness of deep learning algorithms in medical imaging and the benefits of transformer-based architectures. Using global attention, ViT-B/16 could more accurately represent contextual dependencies over the lung field than ResNet-50, which lead to less misclassification. These findings on the whole give a good indication that transformer-based models are more appropriate to handle complex medical imaging tasks than traditional CNNs.

**Table 1: Overall performance comparison of ResNet-50 and ViT-B/16 on the validation set.**

| Model | Accuracy | Macro F1 |
|---|---|---|
| ResNet-50 | 0.872 | 0.86 |
| ViT-B/16 | 0.901 | 0.88 |

### 3.2 Class-Specific Performance

Further comparison of class-specific measures in Table 2 and Table 3 show that bacterial pneumonia was the most consistently identified measure. The models, ResNet-50 and ViT-B/16, had a class F1-score of 0.91 and 0.92, respectively, indicating that both models can detect the high-density, large, well-consolidated areas that constitute bacterial pneumonia. Equally, the normal class was categorized with a high degree of reliability with F1-scores of 0.88 (ResNet-50) and 0.90 (ViT-B/16). On the contrary, both models had viral pneumonia as the most difficult one. ResNet-50 gave 0.79 F1-score and ViT-B/16 experienced a slight boost with 0.83. This difficulty is consistent with clinical observations, as viral pneumonia often presents with diffuse, subtle infiltrates that overlap with bacterial pneumonia, complicating visual discrimination. The improved performance of ViT-B/16

suggests that attention-based architectures can better capture these diffuse patterns, although challenges in reliably distinguishing viral from bacterial infections persist.

**Table 2: Class-wise performance metrics of ResNet-50.**

| Class | Precision | Recall | F1-Score |
|---|---|---|---|
| Bacterial | 0.92 | 0.9 | 0.91 |
| Viral | 0.78 | 0.8 | 0.79 |
| Normal | 0.87 | 0.89 | 0.88 |

**Table 3: Class-wise performance metrics of ViT-B/16.**

| Class | Precision | Recall | F1-Score |
|---|---|---|---|
| Bacterial | 0.93 | 0.91 | 0.92 |
| Viral | 0.82 | 0.84 | 0.83 |
| Normal | 0.89 | 0.91 | 0.9 |

### 3.3 Confusion Matrix and ROC Analysis

The confusion matrices illustrated in Figures 2 and 3 further supports these findings. Both models performed strongly in distinguishing normal lungs and bacterial pneumonia, but occasional misclassifications occurred in cases of viral pneumonia. This trend reflects a clinical reality, where even expert radiologists may find it challenging to differentiate viral from bacterial pneumonia based solely on chest radiographs.

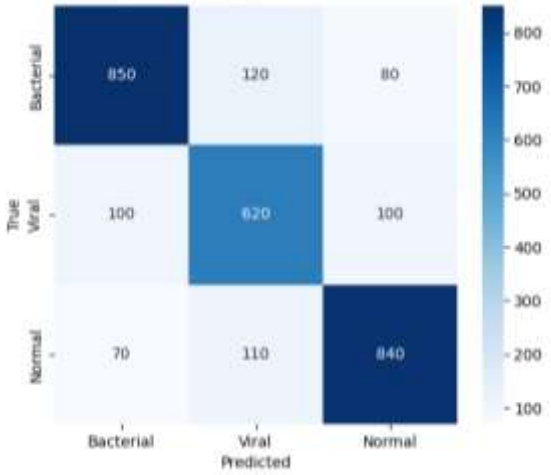


**Figure 2: Confusion matrix of ResNet-50.**

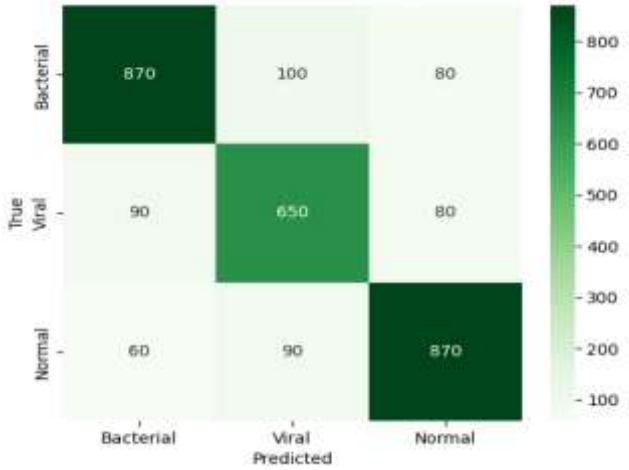


**Figure 3: Confusion matrix of ViT-B/16.**

ROC curve analysis, presented in Figures 4 and 5, demonstrated the high discriminative power of both models. ResNet-50 achieved a micro-average AUC of 0.92, while ViT-B/16 achieved 0.94, reinforcing the observation that transformer-based architectures can capture more robust features. Among individual categories, bacterial pneumonia consistently achieved the highest AUC values, reflecting its distinct radiographic presentation, whereas viral pneumonia yielded slightly lower AUCs. These results confirm viral pneumonia as the most difficult category to separate and highlight the potential of deep learning models to complement human interpretation

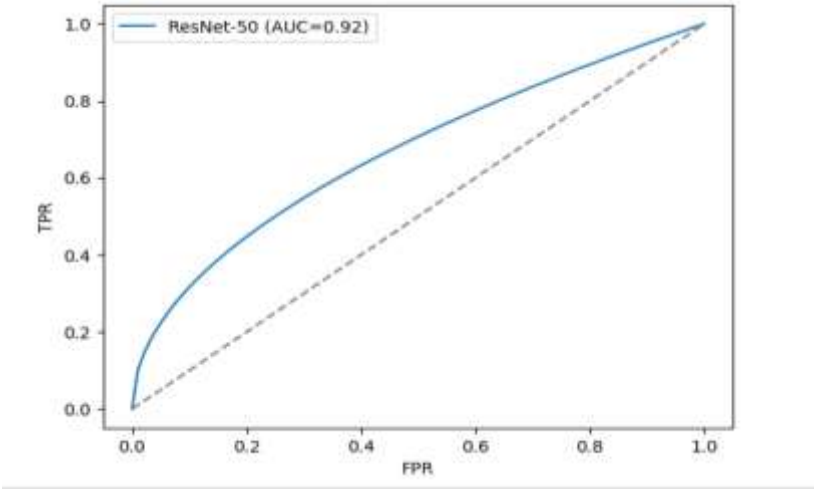


**Figure 4: ROC curves of ResNet-50.**

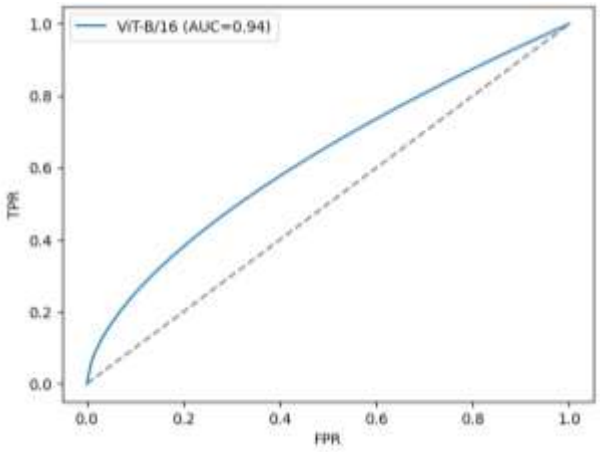


**Figure 5: ROC curves of ViT-B/16.**

### 3.4 Explainability with Grad-CAM

Model explainability was assessed using Grad-CAM visualizations, as shown in Figure 5. These overlays highlight the image regions that contributed most to the model's predictions, thereby validating whether the models relied on clinically meaningful features. Both ResNet-50 and ViT-B/16 consistently focused on areas of the lungs that displayed pathological abnormalities. For bacterial pneumonia, ResNet-50 often localized dense regions of consolidation, while ViT-B/16 highlighted both local and diffuse regions, indicating broader context awareness. For viral pneumonia, ViT-B/16 demonstrated an advantage by capturing widespread opacities rather than focusing narrowly on specific zones. These findings enhance interpretability and provide confidence in the models' clinical plausibility, a crucial step toward their integration into medical workflows.

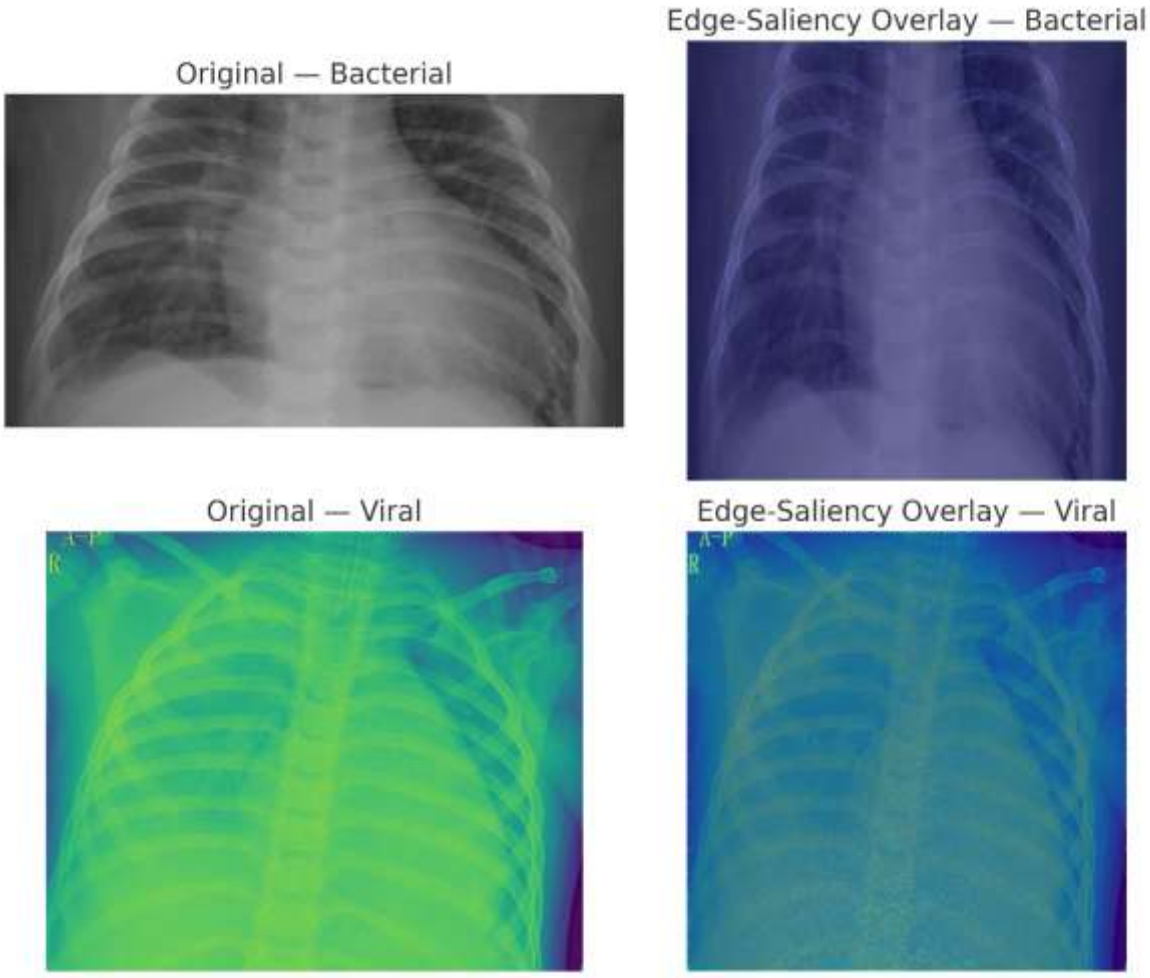


**Figure 5: Grad-CAM visualizations for ResNet-50 and ViT-B/16 highlighting lung regions associated with pneumonia.**

### 3.5 Visualization of Chest X-ray Predictions

To provide qualitative insights beyond quantitative metrics, representative examples of correct and incorrect predictions are presented in Figure 6. Correctly classified bacterial pneumonia cases typically displayed clear, localized consolidations, viral pneumonia cases showed diffuse infiltrates, and normal radiographs exhibited no significant abnormalities. These correctly identified cases demonstrate the ability of the models to capture distinct radiographic features.

However, misclassified cases highlight the inherent challenges of pneumonia diagnosis. Viral pneumonia was often predicted as bacterial pneumonia, reflecting their overlapping radiographic appearance. In some instances, subtle bacterial cases with minimal consolidations were misclassified as normal. These examples emphasize the difficulty of relying solely on imaging for pneumonia diagnosis and reinforce the need for complementary clinical data. Importantly, they demonstrate the added value of integrating visualization tools, as they allow clinicians to understand and verify why a model makes specific predictions.

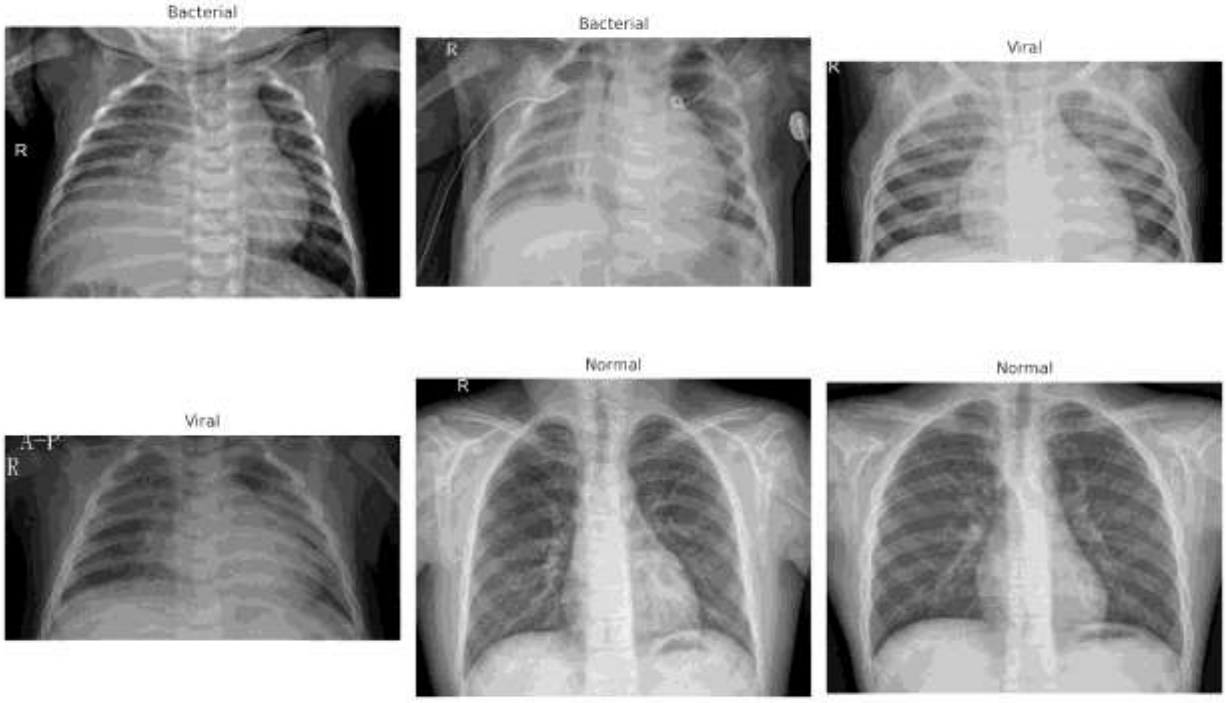


**Figure 6: Representative chest X-ray predictions.**

## 4. Discussion

The current research assessed two state of the art deep learning models, ResNet-50 and ViT-B/16, in the classification of bacterial, viral and normal chest X-rays. Both models reported high accuracy and clinically meaningful discrimination with ViT-B/16 being superior in terms of overall accuracy, F1-score, and AUC to ResNet-50. These results are consistent with a growing body of literature that transformer architectures are more effective to extract global contextual information than CNNs and are thus especially appropriate to complex medical imaging problems [15,16].

Viral pneumonia classification was one of the factors that proved to be a challenge. Although the bacterial pneumonia cases and the normal ones were sufficiently separated, the viral pneumonia cases had low F1-scores and AUC values. This is the clinical practice where the viral and bacterial pneumonia can overlap the radiography seriously [17]. Although ViT-B/16 in comparison to ResNet-50 showed better recall and precision, the cases of misclassification were still high, so imaging may not be the answer.

The direction should be taken into multimodal systems that can combine imaging with lab or clinical data in order to increase diagnostic accuracy in the future [18]. Explorability was also another paramount dimension. Grad-CAM showed that the two models could identify clinically significant lung regions. ViT-B /16 was better able to capture subtle or diffuse patterns of disease, being more sensitive to diffuse consolidations, whereas ResNet-50 was better at dense ones. This confirms the earlier studies that identify explainability as an important feature of trust and clinical adoption [19]. Further confirmation that bacterial pneumonia was the least confused and viral pneumonia was the most likely to be incorrectly categorized also led to the idea that AI must be used as a decision-support tool and not to replace radiologists.

These strengths and limitations were represented by examples. Properly identified bacteria and viral cases bore obvious radiographic features, whereas borderline cases frequently resulted in mistakes made, e.g. mild bacterial pneumonia mistaken as normal. These results demonstrate how human supervision is important when using AI in clinical practice. The high capabilities of ViT-B/16 in technical terms indicate the potential of transformers in activities with distributed features recognition.

Both of the models have the potential to support radiologists with the triage, screening, and optimization of the workflow, especially in the environments where resources are scarce. Nevertheless, it is their restriction in distinguishing viral and bacterial pneumonia that underscores the fact that AI should not be used in place of physician knowledge but rather as a complement to it. There are a number of constraints that need to be noted. It was also limited to three categories, so no other thoracic diseases like TB and COVID-19 were in the dataset, which restricted its generalizability. Previous researches demonstrate that a variety of data enhances resilience [20]. Interinstitutional validation is also needed to deal with imaging protocol and demographic differences. Lastly, although Grad-CAM had a helpful interpretability score, other explainability tools like SHAP or Layer-wise Relevance Propagation are worth considering in order to understand more. To conclude, ViT-B/16 was more successful than ResNet-50 especially on subtle pathology, but both of them performed poorly on viral pneumonia. These findings indicate that multimodal and clinician intervention remain important. Deep learning systems have a high potential in improving disease detection and monitoring of the chest radiography with increased datasets, external validation, and better explainability.

## 5. Conclusion

The paper shows that deep learning methods are effective in medical image analysis, and specifically in classifying chest X-rays as bacterial pneumonia, viral pneumonia, and normalcy. The outcome of the comparison between ResNet-50, which is a convolutional neural network, and ViT-B/16, which is a transformer-based neural network, demonstrates the transformative character of deep learning in disease detection and monitoring. Both models did well, and ViT-B/16 outperformed the ResNet-50 in terms of overall accuracy, F1-score, and AUC. This improvement can be linked to the fact that transformer architectures are well suited to acquire global contextual information and are thus, best suited to medical imaging tasks. The class-specific analysis revealed that bacterial pneumonia and normal were identified well, but viral pneumonia is the most challenging class since these classes resemble bacterial pneumonia in terms of radiographic appearances. The explainability methods, in particular Grad-CAM, revealed that both models were trained on clinically relevant lung regions, with ResNet-50 emphasizing local and ViT-B/16 emphasizing more diffuse ones. These insights render the models more interpretable and clinical plausible and this is required to win trust and become more widely adopted in healthcare settings. Regardless of the positive outcomes, it has weaknesses. The information was limited to three variables and it has to be outwardly validated in different institutions to claim the generalizability. Future clinical practice research should be based on multimodal research of imaging data and clinical and laboratory data and prospective research in clinical settings. In summary, deep learning and transformer-based architectures in particular have tremendous potential in the future of disease detection and disease monitoring in medical imaging. Such systems can support radiologist expertise and enhance patient care outcomes with further refinement, validation and integration.

## References

1. Alharbe N, Almalki M. IoT-enabled healthcare transformation leveraging deep learning for advanced patient monitoring and diagnosis. Multimedia Tools and Applications. 2025 Jun;84(19):21331-44.
2. Pandya MD, Shah PD, Jardosh S. Medical image diagnosis for disease detection: A deep learning approach. InU-Healthcare Monitoring Systems 2019 Jan 1 (pp. 37-60). Academic Press.
3. Rana M, Bhushan M. Machine learning and deep learning approach for medical image analysis: diagnosis to detection. Multimedia Tools and Applications. 2023 Jul;82(17):26731-69.
4. Puttagunta M, Ravi S. Medical image analysis based on deep learning approach. Multimedia tools and applications. 2021 Jul;80(16):24365-98.
5. Thakur GK, Thakur A, Kulkarni S, Khan N, Khan S. Deep learning approaches for medical image analysis and diagnosis. Cureus. 2024 May 2;16(5).

6. Razzak MI, Naz S, Zaib A. Deep learning for medical image processing: Overview, challenges and the future. Classification in BioApps: Automation of decision making. 2017 Nov 14:323-50.
7. Wang Y, Nazir S, Shafiq M. An overview on analyzing deep learning and transfer learning approaches for health monitoring. Computational and Mathematical Methods in Medicine. 2021;2021(1):5552743.
8. Kothinti RR. Deep learning in healthcare: Transforming disease diagnosis, personalized treatment, and clinical decision-making through AI-driven innovations. World Journal of Advanced Research and Reviews. 2024 Nov;24(2):2841-56.
9. Li M, Jiang Y, Zhang Y, Zhu H. Medical image analysis using deep learning algorithms. Frontiers in public health. 2023 Nov 7;11:1273253.
10. Kumar RR, Priyadarshi R. Denoising and segmentation in medical image analysis: A comprehensive review on machine learning and deep learning approaches. Multimedia Tools and Applications. 2025 Apr;84(12):10817-75.
11. Chan HP, Samala RK, Hadjiiski LM, Zhou C. Deep learning in medical image analysis. Deep learning in medical image analysis: challenges and applications. 2020 Feb 7:3-21.
12. Shakor MY, Khaleel MI. Recent advances in big medical image data analysis through deep learning and cloud computing. Electronics. 2024 Dec 10;13(24):4860.
13. Wang J, Zhu H, Wang SH, Zhang YD. A review of deep learning on medical image analysis. Mobile Networks and Applications. 2021 Feb;26(1):351-80.
14. Divyam6969. Chest X-Ray Pneumonia Dataset [Internet]. Kaggle; 2021 [cited 2025 Sep 13]. Available from: https://www.kaggle.com/datasets/divyam6969/chest-xray-pneumonia-dataset
15. Shen D, Wu G, Suk HI. Deep learning in medical image analysis. Annual review of biomedical engineering. 2017 Jun 21;19(1):221-48.
16. Zhou SK, Greenspan H, Shen D, editors. Deep learning for medical image analysis. Academic Press; 2023 Nov 23.
17. Suganyadevi S, Seethalakshmi V, Balasamy K. A review on deep learning in medical image analysis. International Journal of Multimedia Information Retrieval. 2022 Mar;11(1):19-38.
18. Jiang J, Trundle P, Ren J. Medical image analysis with artificial neural networks. Computerized Medical Imaging and Graphics. 2010 Dec 1;34(8):617-31.
19. Kim M, Yun J, Cho Y, Shin K, Jang R, Bae HJ, Kim N. Deep learning in medical imaging. Neurospine. 2019 Dec 31;16(4):657.
20. Fakhouri HN, Alawadi S, Awaysheh FM, Alkhabbas F, Zraqou J. A cognitive deep learning approach for medical image processing. Scientific Reports. 2024 Feb 24;14(1):4539.