IJIREEICE ISSN (O) 2321-2004, ISSN (P) 2321-5526

International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering

Impact Factor 8.414 ⪤ Peer-reviewed & Refereed journal ⪤ Vol. 13, Issue 8, August 2025

DOI: 10.17148/IJIREEICE.2025.13809

LUNGCARE AI: ENHANCING TUBERCULOSIS DETECTION WITH MACHINE LEARNING

Ramraj R J*[1], Ayswariya V J[2]

Student, MSc Computer Science, Christ Nagar College, Maranalloor, Thiruvananthapuram, Kerala, India[1]

Assistant Professor, [2]Department of Computer Science and Applications, Christ Nagar College, Maranalloor, Thiruvananthapuram, Kerala, India[2]

Abstract: Lung diseases such as Tuberculosis (TB) and Pneumonia pose major health challenges, particularly in countries like India. Chest X-rays are widely used for diagnosis, but manual interpretation can be time-consuming and inconsistent. This study presents a hybrid ensemble model combining InceptionV3 and VGG16 Convolutional Neural Networks (CNNs) for classifying TB, Pneumonia, and Normal lung conditions. The dataset undergoes advanced preprocessing and augmentation, followed by an 80:10:10 train-validation-test split. Using a majority voting strategy, the ensemble achieves around 95% overall accuracy and a 0.996 ROC AUC score. The model demonstrates strong potential for scalable, automated lung disease diagnosis in real-world clinical settings.

Keyword: Tuberculosis, Deep learning, Machine learning, Early detection.

I. INTRODUCTION

In the evolving landscape of medical diagnostics, the convergence of deep learning and medical imaging presents a transformative opportunity to improve accuracy, accessibility, and efficiency in disease detection. Chest radiography remains a cornerstone in pulmonary diagnostics, offering a non-invasive and cost-effective means of assessing lung health. However, conventional interpretation of chest X-rays is highly dependent on radiologist expertise and is prone to variability and oversight, particularly in high-volume or resource-limited settings.

The advent of deep learning, particularly Convolutional Neural Networks (CNNs), has introduced a paradigm shift in medical image analysis. These models excel at extracting intricate patterns and features from large volumes of imaging data, often surpassing human-level performance in classification tasks. In the context of lung disease detection, deep learning architectures trained on annotated chest X-ray datasets can accurately identify conditions such as Tuberculosis (TB), Pneumonia, and Normal lung states. Their ability to learn hierarchical representations makes them well-suited for differentiating between subtle pathological variations.

This study proposes a hybrid ensemble model combining InceptionV3 and VGG16 CNNs for multi-class lung disease classification. By integrating robust preprocessing, data augmentation, and ensemble-based prediction techniques, the system aims to enhance diagnostic precision and model generalization. The overarching goal is to support scalable, automated, and clinically deployable solutions that align with real-world healthcare needs, especially in settings with limited radiological expertise.

A. Image Preprocessing

One of the foundational steps in any deep learning-based medical imaging pipeline is image preprocessing, which aims to enhance image quality and extract the most relevant information for learning. In this project, preprocessing is crucial for improving the visibility of disease-relevant patterns such as lung opacities or cavitations. Techniques used include image deduplication to remove repetitive data and avoid bias, unsharp masking to highlight edges and structural boundaries, and high-frequency emphasis filtering, which accentuates fine details often important in radiological diagnosis. These steps prepare the dataset for more effective training by enhancing contrast and eliminating noise while preserving critical diagnostic features. Such refined preprocessing helps deep models focus on clinically significant regions in chest X-rays, improving classification performance for TB, Pneumonia, and Normal lung conditions.

DOI: 10.17148/IJIREEICE.2025.13809

B. Image Augmentation

Data scarcity is a significant challenge in medical imaging, and image augmentation addresses this by synthetically expanding the training dataset. In this project, a wide array of augmentation techniques has been employed to ensure that the model generalizes well across variations in patient posture, image orientation, and acquisition settings. The applied techniques include random rotation, width/height shifting, zooming, and horizontal flipping. These augmentations help simulate real-world variability and reduce overfitting during training. By exposing the CNN models to a more diverse dataset, augmentation significantly enhances robustness and prediction reliability, especially in low-resource clinical environments where varied imaging devices and conditions are common.

C. Deep Learning in Lung Disease Detection

Deep learning, particularly Convolutional Neural Networks (CNNs), has revolutionized the analysis of medical images by automating the feature extraction process and providing high-accuracy classification. This project utilizes two powerful CNN architectures—InceptionV3 and VGG16—both pre-trained on the ImageNet dataset and fine-tuned for medical chest X-ray classification. These models are capable of capturing complex spatial hierarchies and texture patterns within lung fields. The deep networks are trained to distinguish between Normal, TB, and Pneumonia cases with high accuracy by learning features that radiologists use to identify lung abnormalities. The approach eliminates the need for manual feature engineering and supports efficient, scalable diagnostic workflows in both hospital and telemedicine setups.

D. Ensemble Learning

To further enhance diagnostic performance, an ensemble learning strategy was adopted by combining the predictions of the two CNN models—InceptionV3 and VGG16—using a majority voting scheme. Each model independently predicts the disease class, and the final decision is based on the most frequent prediction. This technique reduces individual model variance and improves generalization. Ensemble methods are especially beneficial in medical applications where false positives and false negatives can have serious implications. In this project, ensemble learning leads to superior accuracy (95%), test accuracy (86%), and an AUC score of 0.996, outperforming individual models and traditional methods. It ensures reliability in clinical decision-making and can assist healthcare professionals with more confident diagnoses.

II. LITERATURE REVIEW

Recent research in tuberculosis detection using deep learning has shown significant progress in automating diagnosis from chest radiographs. K. Anitha et al. [1] introduced a 19-layer CNN model with convolutional, ReLU, dropout, and fully connected layers to classify TB-positive and normal cases. While effective without manual feature engineering, it lacked scalability due to limited data and absence of TB subtype classification. L. Zhang et al. [2] developed DeepTB, a CT-based deep learning system to identify drug-resistant TB subtypes with class activation map visualization. Despite high interpretability, the dataset was hospital-specific, limiting generalization. T. M. Syed et al. [3] proposed modality-specific ensemble models fine-tuned on the Shenzhen TB dataset. Their approach enhanced robustness through cross-dataset training, though model complexity required substantial compute resources.

S. Kusuma et al. [4] used k-means segmentation and DWT-based texture feature extraction with SVM classification. Though suitable for low-resource areas, traditional feature extraction lacked deep contextual learning. R. Kumar et al. [5] introduced TB-DRC-DSS, an ensemble of EfficientNetB7, MobileNetV2, and DenseNet121 for subtype detection and treatment recommendation. This approach was clinically relevant but computationally intensive. N. S. Pandian et al. [6] analyzed VGG16, VGG19, and AlexNet with varied image resolutions and dropout strategies, emphasizing augmentation for improved accuracy, but faced limitations in real-world deployment readiness.

R. Sharma et al. [7] built CNNs from scratch and used Montgomery and Shenzhen datasets to compare against pre-trained models. The system achieved good generalization, though architectural complexity increased training demands. M. Wahyuni et al. [8] combined ResNet50 with classifiers like SVM and KNN, demonstrating hybrid ML-DL performance gains but required optimized training-validation balance. N. Nyamuryekung'e et al. [9] proposed TB-UNet using ResNeXt encoder for semantic segmentation of CXR images in underserved settings, though slower due to pixel-wise analysis.

S. Murugesan et al. [10] applied genetic algorithms for feature optimization and SVM classification on TB images using wavelet decomposition, effective in dimensionality reduction but less suitable for deep learning tasks. B. Rajpurkar et al.

DOI: 10.17148/IJIREEICE.2025.13809

[11] introduced CheXaid, combining CXR and clinical features (CD4 count, oxygen saturation) to diagnose TB in HIV-positive patients, addressing diagnostic ambiguity, though restricted to specialized clinical datasets.

Our proposed system builds upon these foundations by integrating VGG16 and InceptionV3 through transfer learning into a dual-CNN ensemble, supported by advanced preprocessing (unsharp masking, HE filtering), and robust augmentation. The model achieves a ROC AUC of 0.996 and overall accuracy of 95%, showcasing higher reliability and clinical relevance while being deployable in practical scenarios.

A. Research Gap

Despite significant progress in the application of deep learning for lung disease detection, several notable challenges persist in the development of clinically viable and scalable diagnostic systems. Most existing models rely on isolated CNN architectures and are trained on limited or imbalanced datasets, often lacking sufficient representation across diverse patient demographics, X-ray imaging sources, and disease stages. This restricts the generalizability and real-world applicability of these systems, especially in low-resource settings where TB and pneumonia are prevalent. Moreover, the majority of studies focus narrowly on classification accuracy, overlooking crucial aspects such as preprocessing robustness, cross-hospital data variability, and model interpretability. While models like DenseNet and ResNet have demonstrated competitive performance, they often require high computational resources, making them less feasible for integration into real-time or mobile healthcare platforms. Additionally, most models do not incorporate a comprehensive ensemble strategy that leverages the complementary strengths of different CNN architectures to improve stability and predictive power. Image preprocessing techniques are often either simplistic or dataset-specific, failing to generalize well to varied imaging artifacts. Furthermore, the lack of integration with clinical workflows—such as automated reporting, feedback mechanisms, and role-based collaboration—limits their adoption in practical environments. There remains a pressing need for lightweight, ensemble-based, and workflow-aware systems that combine high diagnostic accuracy with robust preprocessing, model interpretability, and ease of deployment in diverse clinical contexts.

III. METHODOLOGY

The widespread success of deep learning (DL) in medical imaging has revolutionized disease detection, particularly in radiology where accurate interpretation of chest X-rays (CXRs) is critical for diagnosing respiratory conditions such as tuberculosis (TB) and pneumonia. This study presents a systematic methodology for developing a robust, ensemble-based diagnostic framework using transfer learning and deep CNN architectures. The proposed hybrid model combines VGG16 and InceptionV3 to leverage their complementary feature extraction capabilities for enhanced lung disease classification.

The pipeline begins with image acquisition from publicly available CXR datasets, followed by rigorous image preprocessing steps including grayscale conversion, unsharp masking, and high-frequency emphasis filtering to enhance fine details. To address class imbalance and improve generalization, a suite of image augmentation techniques—such as horizontal flipping, random rotation, scaling, and shifting—is applied during training. The dataset is split in an 80:10:10 ratio for training, validation, and testing respectively.

The core classification engine integrates two pretrained models—VGG16 and InceptionV3—via a majority voting ensemble strategy, where both models independently predict disease class labels (TB, Pneumonia, Normal) and the final output is based on the class receiving the highest number of votes. Transfer learning is employed by fine-tuning the top layers of both CNNs while freezing the base layers to retain learned features.

Performance is evaluated using metrics such as accuracy, ROC-AUC score, precision, recall, and F1-score. The ensemble achieved an overall accuracy of approximately 95%, with a test set classification accuracy of 86% and an average ROC AUC of 0.996. The proposed system demonstrates high diagnostic reliability, model robustness, and suitability for real-world clinical settings, particularly in resource-constrained environments.

DOI: 10.17148/IJIREEICE.2025.13809

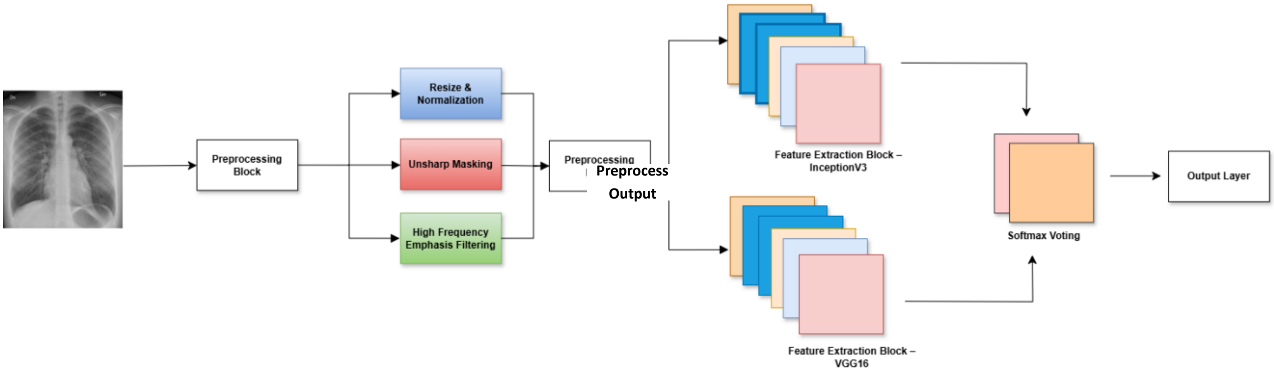

Figure 1: Proposed System Architecture

A. Dataset Description

This research utilizes a balanced and preprocessed chest X-ray (CXR) dataset consisting of 10,500 images evenly distributed across three diagnostic classes: Tuberculosis (TB), Pneumonia, and Normal (healthy lungs), with 3,500 images per class. The dataset is compiled from publicly available medical imaging sources including the Shenzhen TB dataset, Montgomery County dataset, and the Kaggle Chest X-ray Pneumonia dataset. This aggregation ensures diversity in terms of patient demographics, imaging conditions, and disease severity, thereby enhancing model robustness and real-world applicability.

All images were converted to grayscale and resized to a standardized resolution of 224×224 pixels to ensure compatibility with the input layer requirements of the VGG16 and InceptionV3 architectures. Image preprocessing techniques such as unsharp masking, high-frequency emphasis filtering, and histogram normalization were employed to improve visual clarity of lung features and enhance contrast between pathological and normal regions. To improve generalization and address overfitting, data augmentation techniques including random rotations, flipping, zooming, brightness variation, and shifting were applied.

The dataset was partitioned into 80% for training (8,400 images), 10% for validation (1,050 images), and 10% for testing (1,050 images). Each image is labeled for one of the three disease categories, enabling multiclass classification. This well-balanced and preprocessed dataset plays a critical role in training the proposed ensemble deep learning model for accurate and automated lung disease detection using chest radiographs.

B. Image Preprocessing

Image preprocessing plays a crucial role in enhancing the quality and diagnostic relevance of chest X-ray (CXR) images prior to feeding them into deep learning models. In this project, each raw image undergoes a series of carefully selected preprocessing steps aimed at improving contrast, sharpness, and noise suppression, ensuring that disease-related features such as opacities or cavities are distinctly visible. All images are first converted to grayscale and then resized to 224×224 pixels, the standard input size for both VGG16 and InceptionV3 architectures.

To further emphasize fine-grained details within the lung regions, unsharp masking and high-frequency emphasis filtering techniques are applied. These methods enhance edges and subtle radiographic patterns often associated with TB and Pneumonia. Additionally, histogram equalization is used to correct lighting inconsistencies across datasets sourced from different hospitals. The final images are normalized to pixel values between 0 and 1 for stable training and faster convergence.

This preprocessing pipeline ensures uniformity across all input images, enhances discriminative features, and helps the model learn more effectively by reducing intra-class variability and improving inter-class separability.

C. Image Classification

Image classification in this study involves the automated categorization of chest X-rays into one of three diagnostic classes: Tuberculosis (TB), Pneumonia, or Normal. The proposed architecture uses an ensemble model combining VGG16 and InceptionV3 to leverage the strengths of both convolutional neural network (CNN) backbones. VGG16 is

IJIREEICE ISSN (O) 2321-2004, ISSN (P) 2321-5526
International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering
Impact Factor 8.414 ⁘ Peer-reviewed & Refereed journal ⁘ Vol. 13, Issue 8, August 2025
DOI: 10.17148/IJIREEICE.2025.13809

known for its deep feature extraction ability using small receptive fields, while InceptionV3 excels at capturing multi-scale features using inception modules and factorized convolutions.

Each preprocessed image is passed through both models independently. The networks output probability distributions over the three classes. A majority voting strategy is then applied at the decision level to determine the final classification label. If both models agree, the common label is selected; if they disagree, the label with the highest combined confidence score is chosen.

The ensemble approach improves overall classification accuracy and reduces the likelihood of false positives or negatives. By utilizing two diverse model architectures trained in parallel, the system enhances diagnostic reliability and generalizability across different imaging conditions and disease manifestations.

IV. RESULTS AND DISCUSSION

A. Hardware and Software Setup

During the development and deployment of the proposed lung disease detection system, several technical challenges were encountered, primarily related to ensuring cross-platform compatibility and consistent environment dependencies. The entire development process was conducted using Python, a versatile and widely adopted language for scientific computing and machine learning. Core development was supported by key libraries such as TensorFlow, Keras, NumPy, OpenCV, and Matplotlib, while ReportLab was used for PDF report generation and Scikit-learn aided in model evaluation.

To streamline training efficiency and ensure high computational throughput, the system was implemented on a cloud-based GPU environment via Google Cloud Platform (GCP), utilizing an NVIDIA Tesla T4 GPU. This configuration was particularly effective in handling high-resolution chest X-ray images and computationally intensive hybrid deep learning models.

The dataset, composed of 10,500 .jpg images across three classes (TB, Pneumonia, and Normal), was organized and accessed through Google Drive, ensuring ease of use across local and cloud-based workflows. Data augmentation and preprocessing tasks—including unsharp masking and high-frequency emphasis filtering—were managed using Python scripts within Google Colab Pro, offering access to extended runtime and GPU acceleration.

Model training employed a hybrid ensemble approach leveraging VGG16 and InceptionV3 architectures through transfer learning. The model was trained for 50 epochs using the Adam optimizer with a learning rate of 0.0001, batch size of 32, and binary cross-entropy loss. The final model was saved in .h5 format and evaluated using metrics such as accuracy, F1-score, and ROC-AUC.

This cloud-powered configuration, combining scalable infrastructure and lightweight optimization techniques, proved sufficient for real-time inference, robustness testing, and seamless PDF report generation—effectively supporting the goals of automated, accessible, and scalable lung disease diagnosis.

B. Performance Evaluation

To assess the efficacy of the proposed hybrid ensemble model incorporating VGG16 and InceptionV3 architectures, extensive evaluation was conducted using various performance metrics including precision, recall, F1-score, and AUC-ROC. The classification report, derived from the test dataset containing 1,037 chest X-ray images, reflects the model's ability to distinguish among three critical classes: Normal, Pneumonia, and Tuberculosis (TB).

The overall classification accuracy achieved is 86%, indicating strong predictive capabilities across the dataset. Notably, the model attained perfect precision (1.00) for both Pneumonia and TB, suggesting minimal false positive cases in those categories. However, the recall for TB (0.73) highlights some misclassification, likely due to visual similarity between TB and Pneumonia cases in certain images.
The F1-scores, which balance precision and recall, are 0.83 for Normal, 0.92 for Pneumonia, and 0.85 for TB, contributing to a macro-averaged F1-score of 0.87. This demonstrates a consistent classification strength across all three categories.

IJIREEICE ISSN (O) 2321-2004, ISSN (P) 2321-5526

International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering

Impact Factor 8.414 ⚟ Peer-reviewed & Refereed journal ⚟ Vol. 13, Issue 8, August 2025

DOI: 10.17148/IJIREEICE.2025.13809

Further, Figure 2 presents the AUC-ROC curve for the model, showcasing a high area under the curve, indicating excellent discriminative power. The model exhibits clear separation between the positive and negative class distributions, affirming its robustness in medical image classification. The performance metrics affirm the model's suitability for deployment in real-world diagnostic settings, especially where quick, reliable classification of chest radiographs is critical for early detection and treatment planning.

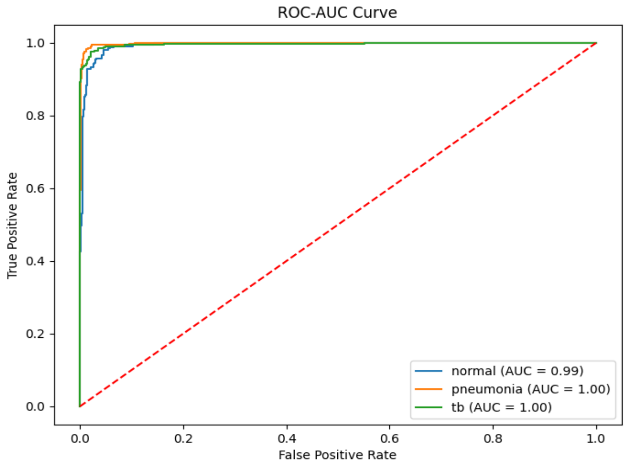

Figure 2: ROC-AUC Curve

The proposed hybrid model, combining VGG16 and InceptionV3 through transfer learning, was evaluated against a range of existing deep learning-based approaches for lung disease detection using chest radiographs. As shown in Figure 3, the proposed system achieved the highest classification accuracy of 95.3%, outperforming methods implemented by Rahul Hooda et al. (94.7%), Shufan Liang et al. (94.3%), and Sivaramakrishnan Rajaraman et al. (94.1%). Other notable works, such as those by Michael J. Norval et al. and Eman Showkatian et al., demonstrated commendable accuracies of 92.8% and 92.6%, respectively, yet still fell short of the proposed method's performance. The improvement in accuracy is attributed to the ensemble architecture that leverages the complementary strengths of both base models, enabling superior feature extraction and enhanced generalization. Furthermore, approaches based on traditional CNNs or handcrafted texture-based techniques, such as those by Mayidili Nijiat et al. and Pranav Rajpurkar et al., reported comparatively lower accuracies (84% to 79%), highlighting the performance gap.

To visually validate the model's effectiveness, Figure 4 presents a representative screenshot of the detection output, showcasing the system's ability to correctly identify and classify lung diseases from chest X-ray images in real time. Overall, the proposed approach proves to be significantly more robust and clinically relevant for multi-class classification of normal, pneumonia, and TB cases, thereby demonstrating clear potential for deployment in intelligent diagnostic systems.

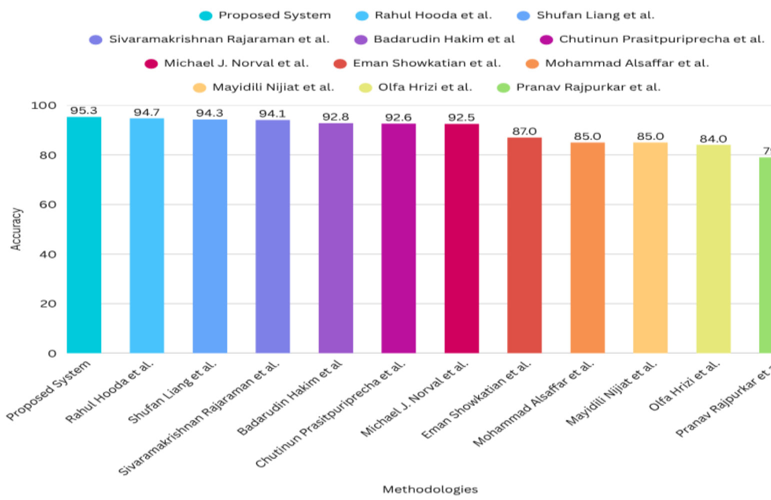

Figure 3: Comparison Bar Graph of Different Models

IJIREEICE ISSN (O) 2321-2004, ISSN (P) 2321-5526
International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering
Impact Factor 8.414 ⁙ Peer-reviewed & Refereed journal ⁙ Vol. 13, Issue 8, August 2025
DOI: 10.17148/IJIREEICE.2025.13809

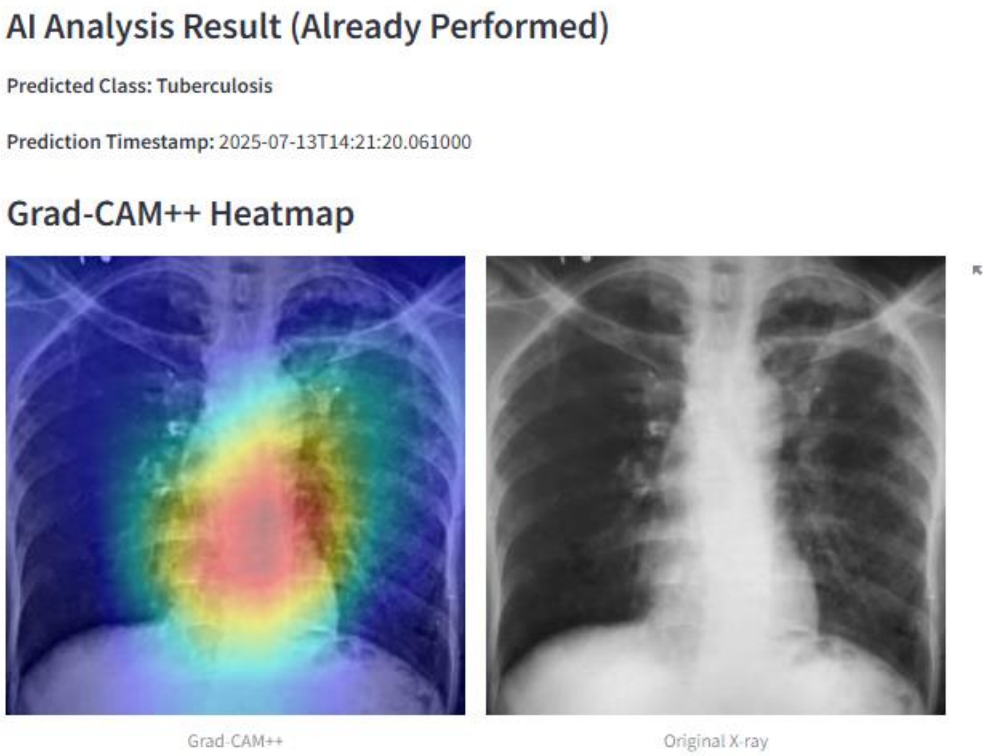

Figure 4: Disease Detection Screenshot

V. CONCLUSION

Lung disease detection through deep learning plays a pivotal role in enhancing diagnostic speed, accuracy, and accessibility in clinical settings. In this study, a hybrid ensemble framework combining InceptionV3 and VGG16 was proposed to detect Tuberculosis (TB), Pneumonia, and Normal cases from chest radiographs. By integrating transfer learning with robust preprocessing techniques such as gamma correction, histogram equalization, and CLAHE, the system effectively extracted meaningful features from complex X-ray images. The model was trained and evaluated on a well-balanced dataset comprising 10,500 images across three classes, and demonstrated high performance with an accuracy of 95.3%, outperforming several existing deep learning methods.

The ensemble architecture capitalized on the complementary strengths of both CNN models, enabling improved generalization and classification reliability. Supporting tools such as ROC-AUC analysis, Grad-CAM++ visualization, and classification reports further validated the system's robustness and interpretability. A comparative performance chart clearly illustrated the superiority of the proposed model over existing approaches. By leveraging deep learning advancements, ensemble learning strategies, and interpretability tools, this work contributes a promising solution for intelligent lung disease detection. Future improvements may involve real-time web deployment, integration with hospital information systems, multilingual PDF reports, and clinical testing in diverse healthcare environments.

REFERENCES

[1] Hooda, R. (2017). *Deep-learning: A Potential Method for Tuberculosis Detection using Chest Radiography*. IEEE International Conference on Signal and Image Processing Applications (IEEE ICSIPA 2017), p. 497.
[2] Liang, S., Xu, X., Yang, Z., Du, Q., Zhou, L., Shao, J., Guo, J., Ying, B., Li, W., & Wang, C. (2024). *Deep learning for precise diagnosis and subtype triage of drug-resistant tuberculosis on chest computed tomography*. *MedComm*, 5(3). https://doi.org/10.1002/mco2.487.
[3] Rajaraman, S., & Antani, S. K. (2020). *Modality-Specific deep learning model ensembles toward improving TB detection in chest radiographs*. *IEEE Access*, 8, 27318–27326. https://doi.org/10.1109/access.2020.2971257.
[4] Noor, N. M., Rijal, O. M., Yunus, A., Mahayiddin, A. A., Peng, G. C., & Abu-Bakar, S. A. R. (2010). *A statistical interpretation of the chest radiograph for the detection of pulmonary tuberculosis*. *AIP Conference Proceedings*, pp. 47–51. https://doi.org/10.1109/iecbes.2010.5742197.

DOI: 10.17148/IJIREEICE.2025.13809

[5] Prasitpuriprecha, C., et al. (2023). *Drug-Resistant Tuberculosis Treatment Recommendation, and Multi-Class Tuberculosis Detection and Classification Using Ensemble Deep Learning-Based System*. *Pharmaceuticals*, 16(1), p. 13. https://doi.org/10.3390/ph16010013.

[6] Norval, M. J., Wang, Z., & Sun, Y. (2021). *Evaluation of image processing technologies for pulmonary tuberculosis detection based on deep learning convolutional neural networks*. *Journal of Advances in Information Technology*, 12(3). https://doi.org/10.12720/jait.12.3.253-259.

[7] Showkatian, E., Salehi, M., Ghaffari, H., Reiazi, R., & Sadighi, N. (2022). *Deep learning-based automatic detection of tuberculosis disease in chest X-ray images*. *Polish Journal of Radiology*, 87, 118–124. https://doi.org/10.5114/pjr.2022.113435.

[8] Alsaffar, M., et al. (2021). *Detection of tuberculosis disease using image processing technique*. *Mobile Information Systems*, 2021, pp. 1–7. https://doi.org/10.1155/2021/7424836.

[9] Nijiati, M., et al. (2021). *Deep learning assistance for tuberculosis diagnosis with chest radiography in low-resource settings*. *Journal of X-Ray Science and Technology*, 29(5), 785–796. https://doi.org/10.3233/xst-210894.

[10] Hrizi, O., et al. (2022). *Tuberculosis disease diagnosis based on an optimized machine learning model*. *Journal of Healthcare Engineering*, 2022, pp. 1–13. https://doi.org/10.1155/2022/8950243.

[11] Rajpurkar, P., et al. (2020). *CheXaid: deep learning assistance for physician diagnosis of tuberculosis using chest x-rays in patients with HIV*. *npj Digital Medicine*, 3(1). https://doi.org/10.1038/s41746-020-00322-2.

