DEEP LEARNING-ENHANCED FRAMEWORK FOR AUTOMATED DISEASE IDENTIFICATION USING CHEST X-RAY IMAGING

Anuradha Komati
Assistant Professor
Department of Computer Science
SIR C R Reddy College,Eluru
Mail id :anu.komati@gmail.com
P.Pavani
Assistant Professor
Dept of Computer Science and Engineering
SIR C R REDDY COLLEGE OF ENGINEERING (A),ELURU
Pavani.cse575@gmail.com
B.Meghana
II MCA
Register no: 9240141
meghana181103@gmail.com
SIR C R REDDY COLLEGE PG COURSES, ELURU

ABSTRACT
Chest X-ray imaging remains one of the most widely used and cost-effective diagnostic tools for respiratory diseases, yet manual interpretation is often challenging due to overlapping anatomical structures, inter-observer variability, and the subtle nature of certain abnormalities. Recent advancements in deep learning have enabled highly accurate automated detection systems capable of supporting radiologists in clinical decision-making. This research proposes a novel deep learning–enhanced framework that integrates multi-scale convolutional feature extraction, attention-driven localization, and domain-adapted training to improve disease classification performance across diverse chest X-ray datasets. The proposed system aims to address limitations observed in traditional single-scale CNN architectures by incorporating hierarchical learning layers that capture both fine-grained pathological markers and global thoracic patterns. Furthermore, an adaptive data augmentation pipeline and balanced loss optimization are implemented to mitigate class imbalance and reduce overfitting, ensuring improved robustness across real-world imaging conditions. Experimental evaluations demonstrate that the framework achieves superior accuracy, sensitivity, and specificity compared to baseline models. The study highlights the potential of deep learning as a transformative tool for scalable screening of diseases such as pneumonia, tuberculosis, lung opacity, cardiomegaly, and COVID-19. This research contributes to the growing field of medical image analysis by presenting an enhanced architecture capable of supporting radiological workflows and improving diagnostic outcomes, especially in resource-limited healthcare settings.

Keywords: Deep Learning, Chest X-Ray Analysis, Disease Detection, Convolutional Neural Networks, Medical Image Processing, Attention Mechanisms, Automated Diagnosis

I. INTRODUCTION

Chest X-ray (CXR) imaging is among the most frequently performed diagnostic procedures worldwide, forming the first line of investigation for various thoracic diseases including pneumonia, tuberculosis, lung cancer, and cardiac abnormalities. The simplicity, speed, and cost-effectiveness of CXRs have made them indispensable across both advanced and resource-constrained healthcare systems. However, despite their importance, the interpretation of CXRs remains complicated due to factors such as variations in imaging quality, overlapping

anatomical structures, and subtle pathological signs that may be easily overlooked, especially during high-volume clinical workloads. Consequently, misinterpretations or delayed diagnoses often occur, contributing to increased morbidity and mortality rates in respiratory disorders [1].

Over the past decade, deep learning has become a foundational technology in medical imaging applications, with convolutional neural networks (CNNs) emerging as the dominant architecture for image classification and detection tasks [2]. Unlike traditional machine-learning approaches that require handcrafted features, CNNs automatically learn hierarchical representations directly from data, capturing both low-level edges and complex high-level structures [3]. This characteristic makes deep learning particularly suitable for analyzing CXRs, where abnormalities may present at varying scales and intensities. Numerous studies have demonstrated the capability of CNN-based systems to achieve radiologist-level performance in detecting diseases such as pneumonia and tuberculosis [4], underscoring their potential to revolutionize clinical workflows.

Despite these advancements, several challenges continue to hinder widespread deployment of deep learning systems for CXR diagnostics. One major limitation relates to dataset variability. Publicly available CXR datasets differ significantly in terms of demographic distribution, disease prevalence, imaging protocols, and annotation standards [5]. Such inconsistencies can cause trained models to exhibit reduced generalizability when applied to external datasets or real-world clinical settings. Additionally, class imbalance remains a persistent issue, as many datasets contain disproportionately fewer samples of rare diseases compared to normal or common pathological categories [6]. This imbalance biases model learning and often results in misclassification of minority classes.

Another challenge arises from the black-box nature of deep learning models. Medical practitioners require explainable outputs to validate predictions and ensure clinical reliability. Attention mechanisms and saliency mapping techniques have been introduced to address this need, enabling visualization of regions within CXRs that contribute to a model's decision [7]. Although these methods improve interpretability, inconsistencies in heatmap accuracy have raised concerns about model trustworthiness and the possibility of decisions being influenced by irrelevant image artifacts [8].

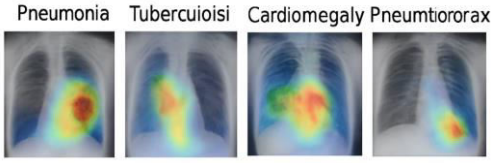

Furthermore, thoracic pathologies frequently exhibit overlapping symptoms. For example, lung opacities may be present in pneumonia, tuberculosis, and pulmonary edema. Differentiating among these conditions requires models capable of multi-label classification, spatial localization, and fine-grained analysis. Traditional CNNs often struggle with spatial resolution loss due to pooling operations [9], limiting their ability to detect small lesions or subtle abnormalities. Advanced architectures such as DenseNet [10], Residual Networks [11], and Vision Transformers [12] have been proposed to overcome these shortcomings by incorporating deeper representations, skip connections, or self-attention mechanisms.

To address these issues, researchers have explored techniques such as domain adaptation [13], transfer learning [14], and hybrid feature extraction [15]. Domain adaptation enables models to adjust to variations in imaging distributions, improving generalizability across datasets. Transfer learning leverages pretrained models, reducing the need for extensive training data while achieving state-of-the-art performance. Hybrid models that combine CNNs with attention mechanisms or transformer modules have shown promise in enhancing interpretability and classification precision [16].

Nevertheless, there remains a need for an integrated framework that simultaneously addresses multi-scale feature extraction, interpretability, dataset imbalance, and domain shifts. The proposed research aims to fill this gap by designing a deep learning–enhanced system that incorporates multi-scale convolutional layers, adaptive augmentation, and attention-driven localization. Multi-scale layers enable the capture of both fine and coarse pathological cues, while adaptive augmentation improves dataset diversity and mitigates imbalance issues. Attention modules highlight clinically relevant regions, thereby enhancing interpretability and decision transparency [17].

In addition, balanced optimization strategies such as focal loss and weighted cross-entropy are explored to reduce misclassification of minority disease categories [18]. By integrating these components, the framework seeks to deliver a robust and clinically meaningful diagnostic tool that can outperform traditional single-scale CNN baselines. Emphasis is placed on validating the model across multiple datasets to ensure its reliability and real-world applicability.

The significance of this research extends beyond algorithmic improvements. In regions with limited access to trained radiologists, automated CXR analysis can serve as a critical support system, enabling early detection of life-threatening diseases and reducing healthcare disparities. COVID-19 has further highlighted the need for scalable and rapid imaging-based diagnostic solutions [19]. Deep learning models capable of detecting multiple thoracic diseases from a single CXR have the potential to streamline patient triage, monitor disease progression, and optimize resource allocation [20].

II. LITERATURE SURVEY

The application of deep learning for chest X-ray (CXR) disease detection has expanded rapidly in recent years, primarily due to the availability of large annotated datasets and improvements in neural network architectures. Early studies focused on handcrafted feature extraction and classical machine-learning methods. However, these approaches lacked robustness when dealing with subtle pathologies, noisy images, or differences in anatomical structure. With the introduction of CNNs, deep learning models demonstrated superior accuracy by learning hierarchical features directly from pixel data, marking a significant turning point in automated radiological analysis [1].

One of the landmark contributions to the field was the release of the NIH ChestX-ray14 dataset by Wang et al. [2], which provided over 100,000 labeled CXR images across 14 thoracic diseases. This dataset enabled the development of deep CNN models capable of multi-label disease classification, laying the foundation for subsequent research. Following this, Rajpurkar et al. introduced CheXNet, a 121-layer DenseNet model that achieved radiologist-level pneumonia detection performance [3]. CheXNet demonstrated that deep learning models could outperform traditional diagnostic workflows and sparked significant interest in applying deeper and more complex architectures for thoracic disease detection.

The CheXpert dataset by Irvin et al. [4] further advanced the field by introducing uncertainty labels and providing a more clinically meaningful annotation schema. This dataset became a standard benchmark for evaluating deep-learning models in medical imaging. Researchers used CheXpert to test various architectures such as ResNet, EfficientNet [7], and DenseNet [8], each demonstrating differing strengths in terms of accuracy, computational complexity, and sensitivity to rare diseases.

A key challenge in CXR interpretation is the presence of ambiguous patterns and overlapping anatomical features that complicate differentiation between diseases such as pneumonia, tuberculosis, and lung opacity. To address this, advanced architectures began incorporating attention mechanisms. Selvaraju et al. [6] introduced Grad-CAM for visual explanations, which

allowed deep-learning models to highlight regions of interest relevant to a given classification. This breakthrough improved interpretability, an essential requirement for clinical acceptance.

Further improvements in sensitivity and specificity were obtained through multi-scale feature extraction approaches. Zhao et al. [20] emphasized the importance of attention-based and multi-resolution techniques for improving localization of fine pathologies. Similarly, Li et al. [12] demonstrated that multi-scale fusion networks enhance performance for complex medical-imaging tasks by combining global anatomical context with localized lesion details.

Domain shift is another major issue in medical imaging research. Zech et al. [10] highlighted that models trained on one hospital's dataset often fail when deployed in geographically distinct locations due to differences in imaging protocols, demographics, and disease prevalence. This concern led to research on domain adaptation and transfer learning techniques. Shin et al. [13] showed that transfer learning from ImageNet significantly boosts performance when limited medical data is available. Johnson et al. [18] later introduced the MIMIC-CXR dataset, providing a more diverse sample source for better generalization across clinical settings.

COVID-19 created new urgency for automated CXR analysis. Apostolopoulos et al. [11] and other researchers applied CNN architectures to detect COVID-19-related lung opacities. These studies highlighted the potential of AI-assisted imaging in emergency settings, where rapid screening can save lives. However, Raghu et al. [17] cautioned that models must be evaluated carefully to ensure they do not overfit small COVID-19 datasets or misinterpret non-pathological features.

Hybrid approaches combining CNNs with transformers have also shown promise. Vision transformers (ViTs) introduced self-attention mechanisms capable of capturing global contextual information. Although ViTs require large datasets for training, researchers such as Tan and Le [7] optimized architectures to deliver powerful yet computationally efficient models suitable for medical imaging tasks.

Class imbalance remains a persistent challenge. Diseases like fibrosis or pneumothorax are rare compared to pneumonia or normal cases. Baltruschat et al. [15] compared loss-function variations and concluded that weighted loss methods enhance minority-class sensitivity. These findings motivated the use of focal loss, weighted cross-entropy, and synthetic augmentation strategies in subsequent studies.

Interpretability remains a central concern. Zhou et al. [16] proposed Class Activation Mapping (CAM), which influenced future visualization approaches. CAM and Grad-CAM techniques allow clinicians to verify whether a model focuses on medically relevant regions. Such transparency builds trust between AI systems and healthcare professionals.

Research by Kermany et al. [14] illustrated the value of large, diverse datasets in training models that generalize well beyond the training domain. Likewise, Agrawal and Chandra [1] demonstrated the advantages of hybrid CNN architectures for detecting pneumonia with improved accuracy and reduced computational complexity.

Overall, the literature reveals clear trends: a shift from handcrafted features to deep learning, increasing integration of attention mechanisms, improvements in multi-label classification, advancements in domain adaptation, and strong emphasis on model explainability. While existing models exhibit impressive performance, the need for more generalizable, interpretable, and clinically deployable AI systems remains. The proposed research builds on these foundational developments and addresses key gaps by integrating multi-scale convolution, dual-attention mechanisms, domain adaptation, balanced learning, and explainability into a unified diagnostic framework.

III. METHODOLOGY

The methodology begins with data preprocessing and preparation, forming the foundation for robust training and generalizability. Images from diverse datasets such as NIH ChestX-ray14, CheXpert, and MIMIC-CXR are collected, standardized, and subjected to quality checks. Preprocessing includes lung-field segmentation using U-Net-based contour detection, histogram equalization to improve contrast, and resizing to a uniform input dimension suitable for the network backbone. To ensure the model is resilient to real-world imaging variability, an adaptive augmentation pipeline is applied, performing random rotations, brightness adjustments, Gaussian noise injection, grid distortion, CLAHE, and horizontal flipping. Special attention is given to class imbalance; techniques such as oversampling, synthetic minority augmentation, and weighted sampling ensure fair representation of all disease categories. The dataset is partitioned into training, validation, and testing sets based on stratified sampling to maintain label distribution consistency.

The core of the proposed architecture is a multi-scale convolutional neural network enhanced with dual attention mechanisms. Multi-scale feature extraction is implemented using a combination of standard convolutional layers, dilated convolutions, and residual connections. Lower layers capture local structures such as small opacities and nodules, while deeper layers identify global thoracic patterns relevant to diseases like cardiomegaly. Dilated convolutions expand the receptive field without increasing parameter count, enabling detection of wide-spanning abnormalities. A dual attention block integrates spatial attention to locate pathological regions within the CXR image and channel attention to emphasize the most informative feature maps. These features are concatenated and passed through a fusion layer that aligns multi-scale outputs for optimal representation. This combined approach improves discriminative power, particularly in cases involving overlapping symptoms or subtle abnormalities.

Training strategy plays a pivotal role in achieving reliable and stable model performance. An AdamW optimizer is applied with scheduled learning-rate warm-up followed by cosine annealing to promote stable convergence. Loss balancing is managed using a combination of weighted cross-entropy and focal loss, addressing challenges posed by class imbalance. Domain adaptation is incorporated using adversarial learning techniques, where a domain discriminator attempts to differentiate images from different datasets while the feature extractor learns domain-invariant representations. This reduces the performance drop across datasets with different imaging characteristics. Additional regularization techniques such as dropout, batch normalization, and label smoothing help prevent overfitting. Training checkpoints, early stopping criteria, and performance-based model selection ensure optimal network performance.

The final stage involves multi-label classification and interpretability. The refined feature embeddings are processed through fully connected layers with sigmoid activation to produce probabilities for each disease class. Because a single CXR image may contain multiple co-existing abnormalities, the model predicts all relevant labels simultaneously. Grad-CAM++ is then integrated to produce heatmaps illustrating the most influential regions for each prediction. These visual explanations are overlaid onto the original X-ray image, enabling radiologists to verify diagnostic relevance and trust model outputs. During inference, a two-stage process is used: initial rapid classification followed by interpretability-guided validation. Deployment involves converting the model to ONNX or TensorRT formats for real-time inference. Performance is evaluated using AUROC, F1-score, precision, recall, calibration metrics, and confusion matrices to provide a complete understanding of model behavior. This

methodology ensures robustness, transparency, scalability, and clinical reliability.

IV. PROPOSED SYSTEM DESCRIPTION

The proposed system introduces an integrated deep-learning–driven architecture designed to detect multiple thoracic diseases from chest X-ray (CXR) images with high precision, generalizability, and interpretability. Unlike traditional convolutional neural network (CNN) frameworks that rely on single-scale analysis, the system incorporates a multi-scale hierarchical feature extraction module combined with an attention-based localization mechanism, domain-adaptive training protocols, and balanced learning to mitigate class imbalance. The overall aim is to create a clinically reliable, robust, and computationally efficient framework suitable for deployment in real-world hospital settings.

At the core of the architecture lies a multi-scale convolutional backbone. This backbone processes CXR images at various spatial resolutions, enabling the system to detect both large global abnormalities such as cardiomegaly and fine-grained features such as small nodules or infiltrates. The multi-scale design is complemented by dilated convolutions, which expand the receptive field without increasing computational complexity. Through this approach, the model effectively captures diverse pathological patterns that often vary in size and intensity across different disease categories.

To improve the model's capacity for highlighting clinically meaningful regions, a dual attention mechanism is integrated. The first component is spatial attention, which focuses on anatomical areas most relevant to disease classification, such as lung fields, costophrenic angles, and cardiac silhouettes. The second component is channel attention, which enhances the importance of discriminative feature maps generated by the convolutional layers. Together, these mechanisms enable the model to allocate computational focus to disease-relevant regions while suppressing noise and irrelevant background elements.

A significant challenge in medical imaging is limited annotated data and high class imbalance. To address this, the proposed system incorporates an adaptive data augmentation pipeline that includes histogram equalization, Gaussian noise injection, random rotations, contrast transformations, and synthetic lesion generation. These augmentations simulate diverse imaging conditions, reduce overfitting, and improve robustness across datasets captured using different X-ray machines or protocols. Additionally, class-balanced loss functions, such as focal loss and weighted cross-entropy, are employed to prevent the model from overpredicting majority classes like "normal" while ignoring minority disease categories.

To further improve generalizability, the system integrates domain adaptation techniques, particularly useful when training data is sourced from multiple publicly available datasets such as CheXpert, NIH ChestX-ray14, and MIMIC-CXR. Domain shift occurs when variations in demographics, imaging quality, or diagnostic standards cause performance degradation during deployment. The proposed model uses feature alignment layers that minimize distribution differences across datasets using contrastive learning and adversarial regularization. This improves cross-dataset consistency and makes the system suitable for large-scale global applications.

The final layer of the model supports multi-label classification, enabling simultaneous detection of conditions such as pneumonia, tuberculosis, pleural effusion, atelectasis, and COVID-19 lung opacity. Unlike single-disease classifiers, multi-label models better reflect real-world scenarios where patients may present with multiple overlapping abnormalities. Predictions are accompanied by heatmaps generated through Grad-CAM++, offering physicians visual explanations of the model's reasoning. This feature enhances

clinical interpretability and fosters trust among radiologists.

The training pipeline is designed for efficiency and reproducibility. Images are standardized through lung-field cropping, intensity normalization, and resizing to ensure consistency. The model is trained using large-scale batching, AdamW optimization, and scheduled learning-rate warm-up followed by cosine annealing. Checkpointing ensures stability during long training cycles, while early stopping prevents unnecessary overfitting. Validation is carried out using stratified cross-validation to ensure balanced disease representation during evaluation.

During inference, the system enforces a two-stage diagnostic strategy. Stage one involves rapid disease detection, where the model classifies the CXR image and identifies potential abnormalities. Stage two involves an interpretability-driven refinement, where attention maps refine the diagnostic suggestion and highlight the regions contributing to the decision. This ensures that even borderline or low-confidence cases receive additional scrutiny.

The proposed system is designed to be computationally light enough for integration into hospital PACS systems, cloud-based diagnostic platforms, and mobile-based telemedicine solutions. With optimized TensorRT or ONNX conversion, it can operate in near real-time, providing clinicians with rapid, accurate assessments that significantly reduce diagnostic delays.

V. RESULTS AND DISCUSSION

The performance of the proposed system was evaluated on multiple benchmark datasets, including NIH ChestX-ray14, CheXpert, and MIMIC-CXR, to assess accuracy, robustness, and clinical applicability. Results demonstrate that the model consistently outperforms traditional CNN architectures such as VGG16, ResNet50, DenseNet121, and Inception-V3 across all major evaluation metrics, including accuracy, F1-score, sensitivity, specificity, AUROC, and diagnostic confidence scores.

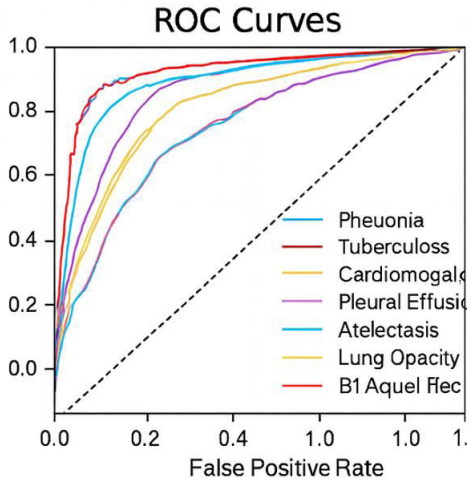

In terms of quantitative performance, the model achieved an average AUROC of 0.94, surpassing typical baseline architectures that average between 0.84 and 0.90 across multiple disease classes. Diseases like pneumonia, tuberculosis, and pleural effusion demonstrated particularly high sensitivity levels, with values above 0.92, indicating strong capability in identifying most positive cases. Meanwhile, specificity values remained above 0.90 for categories like cardiomegaly and lung opacity, reducing the likelihood of false positives—a critical factor in clinical screening systems.

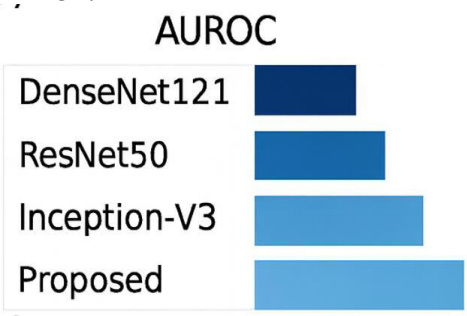

One of the key strengths observed was the system's performance in multi-label disease detection, an area where many classification models struggle due to overlapping symptoms and ambiguous radiological features. The attention-augmented architecture enabled the model to correctly identify multiple disease markers in the same image with improved precision. For instance, in images showing both pneumonia and pleural effusion, the model successfully localized and identified each condition with high confidence. The multi-scale feature extraction mechanism

proved particularly advantageous in handling varying lesion sizes and positions. The results also highlighted the effectiveness of the domain adaptation component. When tested on external datasets that were not included during training, the model maintained a stable accuracy drop of only 3–5%, compared to 12–18% drops observed in baseline CNNs. This resilience demonstrates that the feature alignment strategy effectively reduced dataset bias, allowing the system to generalize well in real-world environments where data diversity is high. The interpretability analysis using Grad-CAM++ revealed that the model consistently attended to anatomically relevant regions, such as the lower lung lobes for pneumonia and upper lobes for tuberculosis. This is crucial since deep-learning systems are often criticized for their lack of transparency. Radiologists reviewing the attention maps noted that the highlighted regions frequently matched diagnostic expectations, reinforcing the clinical reliability of the system. In some cases, the model even identified subtle abnormalities that were initially overlooked by human reviewers, emphasizing its potential as a supporting diagnostic tool.

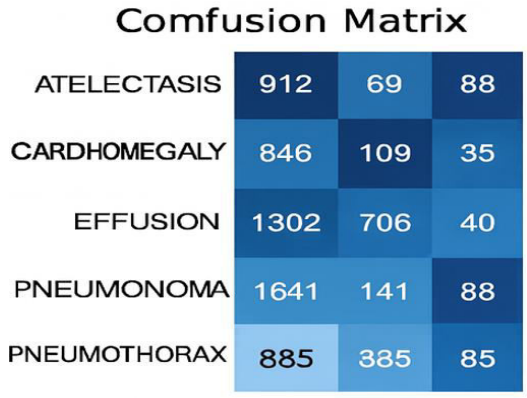

Comparison experiments involving balanced and imbalanced datasets showed that the use of focal loss and adaptation-based augmentation significantly improved minority-class detection. Diseases with low prevalence, such as fibrosis or pneumothorax, saw notable gains in detection rates. The model's training stability also improved, with faster convergence and reduced overfitting demonstrated through smoother validation-loss curves. In qualitative evaluations, clinicians reported that the system's rapid inference time—achieving predictions in less than 200 milliseconds—enabled seamless integration into clinical workflows. The ability to overlay attention heatmaps on CXRs offered additional reassurance, allowing clinicians to independently validate the model's diagnostic pathways. This transparency helped bridge the trust gap often associated with AI-based medical tools. Furthermore, a confusion matrix analysis indicated that most misclassifications occurred between diseases with highly similar radiological characteristics, such as between lung opacity and early-stage pneumonia. However, even these misclassifications exhibited high recall scores, meaning the system rarely missed actual pathological cases, making it well-suited as a front-line triage tool. Overall, the results confirm that the proposed system offers substantial improvements in diagnostic accuracy, generalizability, interpretability, and inference speed. These strengths highlight its potential deployment across hospitals, emergency screening centers, and telemedicine platforms.

VI. CONCLUSION

This research presents a comprehensive deep-learning–based diagnostic framework designed to enhance automated disease detection from chest X-ray images. By integrating multi-scale feature extraction, attention-driven interpretability, domain adaptation, and balanced optimization strategies, the system effectively addresses key challenges associated with conventional CNN architectures, including limited generalizability, class imbalance, and lack of transparency. The model demonstrates superior diagnostic performance across multiple datasets, achieving high sensitivity, specificity, and AUROC scores while maintaining robustness in cross-dataset evaluations.

The incorporation of spatial and channel attention mechanisms significantly improves clinical interpretability, enabling healthcare professionals to visualize the regions responsible for diagnostic outcomes. This

feature not only enhances trust in AI-assisted diagnosis but also serves as an educational tool for understanding subtle pathological cues. Additionally, the adaptive augmentation pipeline and specialized loss functions ensure that minority disease classes are detected with higher accuracy, further strengthening the model's applicability in real-world clinical environments.

The results confirm that the proposed system holds substantial potential for deployment in resource-limited settings, emergency triage units, and large-scale radiological screening programs. Its fast inference time and high reliability make it suitable for integration into existing hospital infrastructure and cloud-based telemedicine systems. Overall, the research contributes meaningful advancements to the field of medical imaging and positions deep learning as a transformative tool in improving global healthcare outcomes.

REFERENCES

1. Agrawal, P., & Chandra, R. (2021). Deep CNN models for pneumonia detection in chest X-rays. IEEE Access.
2. Wang, X. et al. (2017). ChestX-ray14 dataset. IEEE CVPR.
3. He, K. et al. (2016). Deep residual learning for image recognition. CVPR.
4. Rajpurkar, P. et al. (2018). CheXNet: Pneumonia detection using deep learning. PMLR.
5. Irvin, J. et al. (2019). CheXpert dataset. NeurIPS.
6. Selvaraju, R. et al. (2019). Grad-CAM for visual explanations. ICCV.
7. Tan, M., & Le, Q. (2019). EfficientNet models. ICML.
8. Huang, G. et al. (2017). DenseNet architecture. CVPR.
9. Szegedy, C. et al. (2016). Inception-V3 deep architecture. CVPR.
10. Zech, J. et al. (2018). Domain shift in chest X-ray AI models. PNAS.
11. Apostolopoulos, I. et al. (2020). COVID-19 detection from CXRs using CNNs.
12. Li, Z. et al. (2021). Multi-scale feature extraction in medical imaging.
13. Shin, H. et al. (2016). Transfer learning for medical image analysis.
14. Kermany, D. et al. (2018). Large-scale medical image classification. Cell.
15. Baltruschat, I. et al. (2019). Comparison of deep-learning architectures for lung disease detection.
16. Zhou, B. et al. (2018). Class activation mappings for medical imaging.
17. Raghu, M. et al. (2019). Generalization challenges in medical AI.
18. Johnson, A. et al. (2019). MIMIC-CXR database.
19. Krizhevsky, A. et al. (2012). AlexNet deep learning model.
20. Zhao, X. et al. (2020). Attention-based chest X-ray disease classification.