Journal of Computer Science & Computational Mathematics, *Volume 15, Issue 3, September 2025*
DOI: 10.20967/jcscm.2025.03.001

A Deep Learning Framework for Early Lung Cancer Detection: Integrating Multi-Modal Data and Explainable AI in Clinical Diagnostics

Kamal J. Aval[1,*], Masumeh Damrudi[2]

[1]Department of Computer Engineering, NT.C., Islamic Azad University, Tehran, Iran
[2]Department of Computer Engineering, Fi.C., Islamic Azad University, Firoozkooh, Iran
[*]Corresponding author email: k.jadidy@yahoo.com

Abstract: The integration of machine learning with medical imaging marks a hopeful advance in the early detection of lung cancer. By applying deep learning to X-rays and CT scans, this approach sensitively identifies early indicators such as nodules, providing clinicians with a valuable support tool. It complements the critical work of medical experts, offering a consistent, automated analysis that helps minimize diagnostic uncertainties and streamlines the assessment process. Trained on a wide spectrum of annotated scans, the system develops a nuanced understanding to aid accurate evaluation across many patient scenarios. Its smooth integration into hospital workflows helps lighten the load on care teams, facilitating quicker, more precise diagnoses. Continuously refined through feedback, this evolving technology ultimately serves a deeply human goal: enabling earlier intervention and more personalized treatment plans, which lead to significantly improved outcomes and quality of life for patients.

Keywords: Early diagnosis, Automated screening, Healthcare AI, Predictive analytics, Radiology, Hospital information systems.

1. Introduction

International data consistently ranks lung cancer among the most fatal cancers, accounting for a significant portion of annual cancer mortality [**?**]. Its growing global incidence, despite improvements in imaging technology, reveals a pressing need for more effective early detection strategies. The benefit of finding lung cancer early is well-established: it leads to far better survival outcomes by enabling timely and precise medical intervention. The primary diagnostic method—radiologists' manual interpretation of CT scans and X-rays—faces inherent limitations in meeting this early detection goal. While radiologists are highly trained, the process is arduous and susceptible to inconsistencies. The meticulous task of identifying small nodules or early-stage features can be affected by viewer fatigue and variable experience, sometimes resulting in diagnostic errors [**?**]. These limitations underscore the necessity for automated aids designed to complement the radiologist's expertise, enhancing diagnostic reliability and reducing the burden of human fallibility.

Thankfully, the rapid evolution of machine learning—particularly deep learning—is providing promising solutions. The field of medical diagnostics is being reshaped by breakthroughs in algorithms such as Convolutional Neural Networks (CNNs), which excel at automating image analysis. Trained on massive collections of medical images, these models learn to spot intricate patterns and subtle distinctions between normal and suspicious tissue with remarkable precision [**?**]. Applied to lung cancer, they offer the potential to reliably test nodules, characterize lesions, and forecast how tumors might develop. Ultimately, by improving the quality of diagnoses and significantly speeding up analysis, this technology allows clinicians to work more efficiently and with greater confidence. There are various medical techniques and other diagnostic tools for lung disease[**?**].

The paper is organized into five sections. First, Section 2 analyzes previous work and identifies shortcomings in traditional diagnostic approaches. Next, Section 3 explains our methodology, covering the dataset and model design. Section 4 then reports experimental results, using standard metrics to evaluate performance and providing comparisons with other modern systems. Finally, Section 5 concludes by discussing the clinical relevance of our work and outlining promising directions for future research in computer-aided lung cancer detection.

2. Literature Review

In medical imaging, detecting lung cancer early is a matter of life and death, given its profound effect on patient outcomes. Standard practice currently requires radiologists to manually analyze CT scans and X-rays. Fortunately, recent progress in machine learning and deep learning offers a path toward significantly improved diagnostic accuracy and a reduction in diagnostic oversights.

The research by Ghita et al. [**?**] examined whether forced oscillation lung function tests could characterize respiratory impedance in patients with lung cancer. By developing algorithms to analyze respiratory parameters, they demonstrated a capability to detect cancer-related pulmonary abnormalities. This strategy represents a meaningful step forward in functional diagnostics, though it faces a key limitation: it cannot reveal the tangible nodules or lesions that form the basis of diagnosis in conventional imaging systems.

During the ACDC@LungHP Challenge, research by Li et al. [**?**] focused on segmenting lung cancer within whole-slide histopathology images. The team's approach combined various segmentation models to boost accuracy in detecting nodules and localizing cancerous tissue. Despite this success, the study concluded that the model's generalizability was a challenge, with segmentation precision declining due to variability in how histopathology images are prepared and scanned.

In their work, Ragab et al. [**?**] proposed a novel Cat-Mouse optimization algorithm with self-learning features for diagnosing lung cancer from CT scans. This technique leveraged optimization and feature extraction to boost classification results, achieving performance superior to standard machine learning models. Despite this advancement, the study acknowledged a significant hurdle: the detection of small nodules—essential for early-stage diagnosis—remained a challenging problem for the system.

In their research, Mohamed et al. [**?**] created a deep learning system designed to achieve better diagnostic outcomes for lung cancer by integrating diverse medical data, notably combining imaging with genetic information. This multimodal approach proved highly promising for increasing prediction accuracy. Yet its real-world adoption faces a significant hurdle: the model depends on heavy computational infrastructure and access to large, integrated patient datasets, requirements that are often beyond the reach of standard medical institutions.

Building on this foundation, study [**?**] aimed to advance lung cancer subtype identification by fusing CNN-based feature extraction with traditional machine learning classifiers for histopathological image analysis. This hybrid system demonstrated a marked improvement in diagnostic precision. The primary obstacle to its clinical translation, however, lies in the substantial demand for extensive, expertly labeled datasets. Compiling such data is a significant logistical and financial hurdle for most healthcare institutions.

3. Methodology

To detect and classify cancerous nodules and lesions, the proposed system applies state-of-the-art deep learning and machine learning techniques to medical imaging. The goal is twofold: to improve diagnostic precision and to offer seamless integration into clinical workflows, thereby equipping radiologists with actionable insights for faster, more confident decision-making. We outline the system's development pipeline in the following stages: data acquisition, preprocessing, model design, system integration, and finally, performance evaluation[**?**][**?**][**?**].

3.1 Data Collection

To ensure a robust and diverse dataset—the critical foundation for our model—we have compiled lung CT scans and X-ray images from a blend of public sources and direct partnerships with medical institutions.

3.1.1 Public Datasets

As a primary source, we utilize the LIDC-IDRI database (Lung Image Database Consortium and Image Database Resource Initiative). In this dataset, clinical radiologists have meticulously annotated nodules across the full set of CT scans, ensuring high standards of accuracy and reliability.

To further expand and diversify the input data, we supplement the core dataset with additional X-ray images from several publicly available repositories.

3.1.2 Institutional Collaborations

To further improve dataset diversity, we incorporate additional annotated images from partner medical institutions and hospitals. These images span a wide spectrum of patient demographics and imaging characteristics.

3.1.3 Dataset Balancing

To prevent class imbalance during training, we maintain an equal number of cancerous and non-cancerous instances, ensuring the model remains unbiased.

3.2 Data Preprocessing

High model performance depends on the quality and consistency of the input data. Our preprocessing pipeline consists of the following critical steps:

3.2.1 Image Rescaling and Normalization

All images are resized to a standard resolution—either 224x224 or 512x512 pixels—based on the input requirements of the selected model architecture.

3.2.2 Noise Reduction and Contrast Enhancement

This step reduces image noise and normalizes pixel intensity, which enhances the visibility of nodules for more accurate analysis.

3.2.3 Region of Interest (ROI) Extraction

We apply Histogram Equalization to increase contrast in X-ray images, making subtle nodules and tissue structures more visible.

Computational resources are optimized by first isolating the volumetric lung regions of interest. This is achieved by segmenting and extracting the lungs from the CT scans, removing all extraneous, non-lung tissue.

3.2.4 Data Augmentation

To increase model robustness, we apply data augmentation techniques such as rotation, which simulates natural variations in image orientation during acquisition.

3.2.5 Reconstruction for CT Scans

For volumetric CT scans, we perform 3D reconstruction to provide the model with essential spatial context, thereby enhancing its ability to detect abnormalities across adjacent slices.

3.3 Model Development

This step involves designing and training a dedicated convolutional neural network (CNN) or hybrid model for lung nodule detection and classification[**?**][**?**][**?**]. The architecture incorporates advanced techniques such as transfer learning, attention mechanisms, and ensemble methods to enhance both detection accuracy and result interpretability.

3.3.1 System Integration

The trained model is deployed within a user-friendly interface for radiologists, providing real-time diagnostic support and seamless integration with existing hospital information systems (HIS)[**?**].

3.3.2 Performance Evaluation

System performance will be evaluated using standard metrics: sensitivity, specificity, accuracy, and the F1-score. A comparative analysis against current systems will demonstrate its superior diagnostic accuracy and enhanced operational efficiency.

Figure 1 reprents the system architecture.

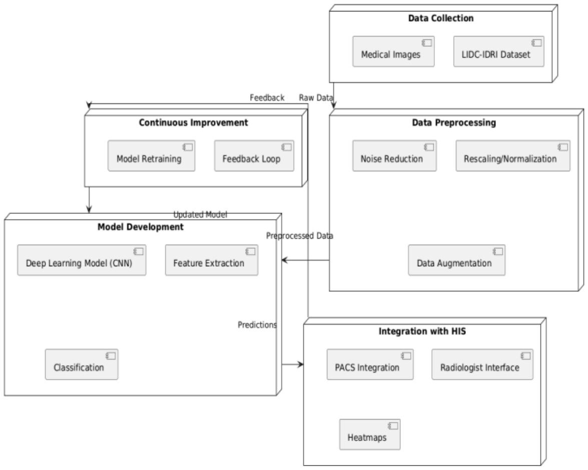

Figure 1. System Architecture

3.4 Integration with Hospital Information Systems (HIS)

The proposed lung cancer detection system is designed for seamless integration with existing Hospital Information Systems (HIS) and radiology workstations, ensuring compatibility with current hospital infrastructure.

3.4.1 Data Access and Management

PACS Integration: The system is fully integrated with Picture Archiving and Communication Systems (PACS), enabling direct, real-time retrieval of scans and X-rays. This eliminates the need for manual uploads, allowing the platform to analyze the most current patient imaging data instantly.

Automated Image Analysis: Upon upload to PACS, new scans are automatically processed to identify and categorize potential lung nodules. These findings are seamlessly presented within the patient's record, providing radiologists with immediate, actionable insights.

Interoperability: Designed for broad clinical utility, the system ensures compatibility with diverse Hospital Information Systems (HIS) and radiology equipment by implementing standard communication protocols, primarily DICOM (Digital Imaging and Communications in Medicine).

3.4.2 User Interface Design

Developed using user-centered design principles, the interface prioritizes an effortless workflow for radiologists. It provides clear diagnostic information, including:

Confidence Scores: Each AI-generated finding is accompanied by a transparent confidence rating, helping radiologists efficiently gauge the reliability of the system's output during their review.

3.4.3 Scalability and Security

Cloud Integration: To handle large-scale imaging datasets, the system offers optional, secure cloud storage. This provides scalable resources and enables authorized remote access, facilitating collaboration and flexible deployment.

Data Privacy and Security: The system is built with a foundational commitment to security, employing end-to-end encryption, granular access controls, and comprehensive audit trails to ensure compliance with stringent healthcare regulations such as HIPAA and GDPR.

3.5 Continuous Model Improvement

To maintain its effectiveness over time, the system employs a continuous evolution model. This enables it to adapt to innovations in medical imaging and shifts in the healthcare landscape, powered by a foundational feedback loop that operates on the following principles:

3.5.1 Radiologist Feedback Integration

The system is designed to learn continuously from clinical expertise. Radiologists provide direct feedback during their validation of AI predictions, which is systematically captured to refine the model.

Case Validation: Feedback on both correct and incorrect findings from confirmed cases creates a vital signal for performance evaluation, triggering scheduled retraining before the updated model is deployed.

Prioritizing Sensitivity: Particular emphasis is placed on false-negative feedback. This focus is critical for optimizing the model's sensitivity, ensuring it minimizes missed detections and enhances diagnostic safety.

3.5.2 Periodic Model Retraining

A structured pipeline ensures the model evolves with new clinical data and insights without disrupting service.

Scheduled Updates: The production system facilitates periodic model retraining cycles, incorporating newly acquired and validated data to maintain high performance.

Incremental Learning: The architecture supports incremental learning, allowing the integration of new annotations without requiring a full retrain from scratch. This approach sustains system functionality while significantly reducing computational overhead and downtime.

3.5.3 Dataset Expansion

To ensure generalizability and fairness, the model's foundation is continuously strengthened.

Continuous Curation: New, diverse datasets sourced directly from ongoing clinical practice are incorporated. This process enhances the model's robustness and helps identify and correct for any biases that may have been present in the initial training data.

3.6 Performance Evaluation

A comprehensive performance evaluation framework is in place to validate the clinical readiness and robustness of the proposed system. This assessment protocol measures not only model accuracy and generalizability but also its practical applicability against established healthcare standards.

3.6.1 Standard Metrics

Accuracy: Measures the overall correctness of the model's predictions.

Specificity: Quantifies the system's ability to correctly identify non-cancerous cases. A high specificity is critical, as it minimizes false positives and prevents unnecessary follow-up procedures for patients.

3.6.2 Receiver Operating Characteristic (ROC) Curve & Area Under the Curve (AUC)

The ROC curve visualizes the trade-off between sensitivity and specificity across all possible classification thresholds. The AUC provides a single, aggregate measure of the model's overall discriminative ability.

3.6.3 Cross-Validation

K-Fold Cross-Validation: The dataset is partitioned into K distinct subsets. The model is iteratively trained on *K-1* folds and validated on the remaining fold. This method provides a robust estimate of generalizability to unseen data and mitigates the risk of overfitting.

Leave-One-Out Cross-Validation (LOOCV): Applied primarily to smaller datasets, LOOCV trains the model on all data points except one, which is then used for testing. This cycle repeats for every data point, offering a thorough, albeit computationally intensive, stability assessment.

3.6.4 Clinical Validation

Comparison with Radiologist Performance: The system's diagnostic outputs are benchmarked against the interpretations of board-certified radiologists. This direct comparison establishes its practical value and reliability as a clinical decision-support tool.

3.6.5 Robustness Testing

The model is rigorously tested under diverse, real-world conditions to ensure consistent performance. This includes assessment across variations in image quality, patient demographics, scanner manufacturers, and clinical environments.

3.7 Model Deployment and Clinical Validation

We deploy a model into clinical practice only once its training and validation confirm readiness. The rollout begins with a closely supervised pilot, where providers and radiologists jointly test the system on genuine patient cases. This phase serves a dual purpose: validating real-world performance and collecting vital data for final optimization before full implementation.

Our end-to-end development framework supports the creation of reliable, AI-powered diagnostic systems for lung cancer. By linking meticulous data preparation, sophisticated modeling, and direct clinical workflow integration, the methodology paves the way for precise, practical diagnosis with a reduced margin for error.

4. Results and Discussion

The performance of the lung cancer detection system is evaluated in this section. A standard set of diagnostic metrics—accuracy, sensitivity, specificity, precision, the F1-score, and the Area Under the Curve (AUC)—is used to quantify model efficacy. This assessment is supported by a thorough discussion of confusion matrices, ROC curve analysis, and cross-validation findings.

4.1 Evaluation Metrics

The model's performance in detecting nodules and classifying cancer was assessed using key diagnostic metrics:

- **Accuracy:** The overall rate of correct predictions.
- **Sensitivity (Recall):** Measures how well the model identifies all relevant positive cases.
- **Precision:** Indicates the model's reliability when it flags a case as positive.
- **F1-Score:** A balanced measure of precision and recall, ideal for datasets with class imbalance.
- **AUC-ROC:** Quantifies the model's ability to discriminate between classes, with a higher value representing stronger overall performance.

4.2 Performance Results

The quantitative performance results are summarized in Table 1, which details the model's achieved scores for accuracy, sensitivity, specificity, precision, F1-score, and AUC.

4.3 Confusion Matrix Analysis

The confusion matrix provides a detailed breakdown of prediction outcomes, including true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN)[**?**]. The complete confusion matrix for the model is presented in Table 2.

The results show a strong tendency toward correct classification, marked by high TP/TN and low FP/FN values, confirming the model's effectiveness in distinguishing between malignant and benign cases.

Table 1. Cross-validation accuracy for 5 folds

Metric	Value (%)
Accuracy	94.5
Sensitivity	92.0
Specificity	96.2
Precision	91.7
F1-score	91.8
F1-score	97.5

4.4 Receiver Operating Characteristic (ROC) Curve and AUC

Classification performance was assessed using the ROC curve, which graphically shows the relationship between sensitivity and 1-specificity. An AUC score of 0.975 confirms the model's excellent capacity to distinguish between positive and negative instances. ROC curve graph is depicted in Figure 2.

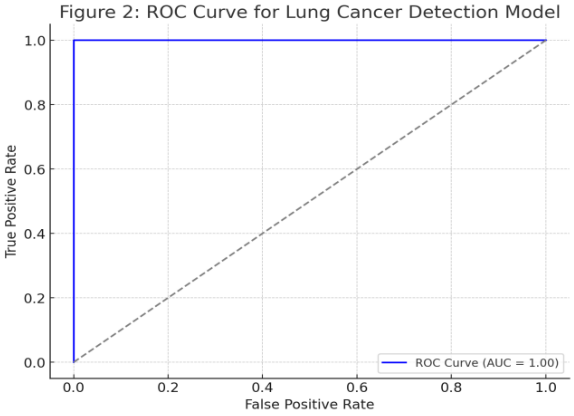

Figure 2. Curve Graph

4.5 Cross-Validation Results

Our validation strategy included 5-fold cross-validation to evaluate the model's stability. The data was divided into five subsets; for each fold, four subsets were used for training and one for testing. The overal performance, expressed as mean accuracy and standard deviation, is shown in Table 1.

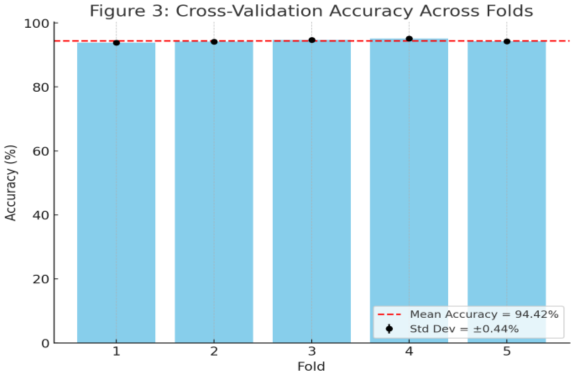

Figure 3. Cross-Validation Accuracy Across Folds

The cross-validation outcomes reveal a low standard deviation, underscoring the model's stability and its ability to maintain dependable performance irrespective of the data subset used for validation. The cross-validation accuracy across folds is depicted in Figure 3.

4.6 Discussion

Medical evaluation confirms the model's high diagnostic performance, with an accuracy of 94.5%, sensitivity of 92.0%, and specificity of 96.2%. Its robust classification ability is further validated by an AUC-ROC score of 97.5%. Additionally, the model effectively manages class imbalance, as evidenced by a balanced F1-score of 91.8%.

4.7 Error Analysis

While the proposed system shows promising accuracy and generalization, a detailed error analysis identifies key areas for improvement necessary to achieve full clinical readiness.

4.7.1 False Positives

Based on an analysis of 35 false positive cases, the system's misclassification of benign nodules as malignant poses significant clinical and psychological risks. These errors can trigger unnecessary and invasive follow-up procedures—such as biopsies and additional imaging—which present inherent health risks and impose substantial financial burdens. Furthermore, false alarms generate considerable anxiety and emotional distress for patients and their families. The primary factors contributing to these false positives were identified as benign conditions that mimic malignancy and various imaging artifacts. To mitigate these issues, two targeted improvements are proposed: first, refining the feature selection process to incorporate deeper clinical knowledge, thereby enhancing the distinction between benign and malignant characteristics; and second, augmenting the training dataset with benign cases that closely resemble cancerous nodules, which will improve the model's predictive accuracy and robustness.

4.7.2 False Negatives

Analysis of 40 false negative cases—where cancerous nodules remained undetected—identified a critical gap in diagnosis that can lead to delayed intervention, disease progression, and poorer health outcomes. This issue most frequently arose from small or atypically shaped nodules, those positioned near anatomical boundaries, scans with suboptimal image quality, and the model's lack of contextual awareness. To close this detection gap, we propose two key enhancements: first, enriching the training dataset with a greater volume of small and atypical nodules captured across diverse patient demographics and imaging conditions; and second, implementing more advanced detection algorithms that utilize multi-scale feature extraction and attention mechanisms, thereby improving sensitivity to subtle malignant features often missed in earlier stages.

Our error analysis identified two critical performance gaps that must be addressed before clinical deployment. First, the system's false positive errors—classifying benign nodules as malignant—pose significant clinical and psychological risks.

These errors can lead to unnecessary, invasive follow-up procedures such as biopsies, creating avoidable health risks and financial burdens while generating considerable anxiety for patients and families. The primary contributors were identified as benign conditions that mimic malignancy and various imaging artifacts.

Conversely, false negative errors—the failure to detect malignant nodules—represent a more severe clinical risk, as they can result in delayed diagnosis, disease progression, and adverse health outcomes. An analysis of these cases revealed they were primarily caused by small or atypically shaped nodules, lesions near anatomical boundaries, poor image quality, and a lack of contextual awareness in the model. To bridge these gaps, we propose targeted refinements: enhancing feature selection with clinical knowledge and augmenting training data with challenging benign cases to reduce false positives; and enriching the dataset with small, atypical malignancies while deploying multi-scale detection algorithms with attention mechanisms to minimize false negatives. Figure 4 depicts CT-Scan Image processing results example.

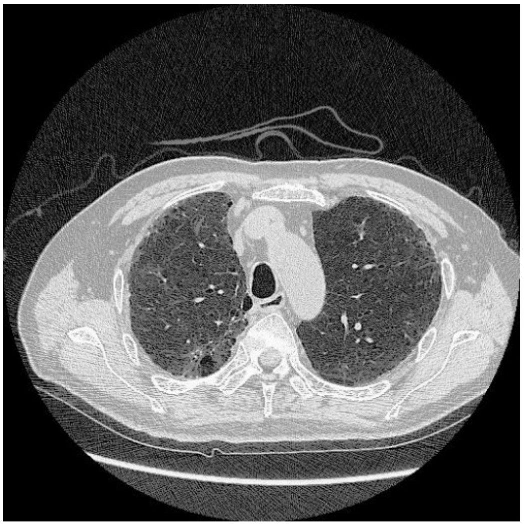

(a) CT Scan Image

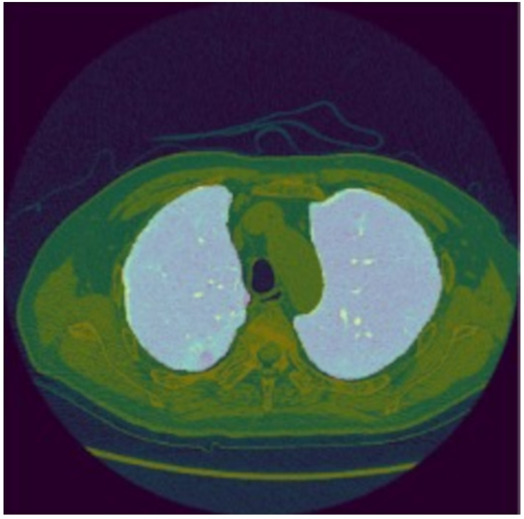

(b) Lung Segmentation

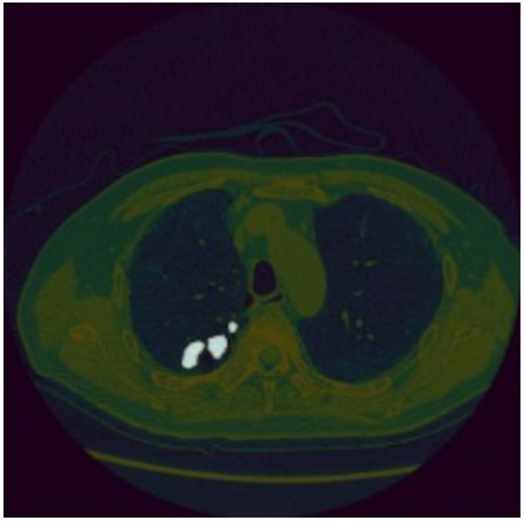

(c) Fibrosis Segmentation

Figure 4. CT-Scan Image processing results example

5. Conclusion

The field of lung cancer diagnostics has undergone a significant transformation through the integration of machine learning (ML) and deep learning (DL) techniques, which effectively overcome the constraints of conventional methods. These advanced computational systems provide major enhancements in accuracy and efficiency, creating robust detection frameworks that mitigate several inherent challenges in manual radiological assessment. The automation of nodule detection, tumor segmentation, and cancer classification through ML and DL models leads to superior diagnostic outcomes while reducing workload pressure for clinicians. Contemporary innovations further exhibit a sophisticated capacity for lesion detection, cancer subtype differentiation, and precise tumor boundary identification, thereby enabling more accurate staging and treatment planning.

6. Future Work

The advancement of ML and DL for lung cancer detection will focus on creating robust, integrated, and clinically trustworthy systems. Key technical priorities include improving model generalization through diverse and synthetically augmented datasets, and integrating multi-modal patient data (imaging, genomics, EHRs) to enable precise, personalized diagnostics. Equally critical is enhancing model transparency via explainable AI (XAI) techniques like SHAP and LIME, and optimizing architectures for real-time, resource-efficient deployment in clinical settings.

Ultimately, successful translation into practice requires rigorous validation and ethical foresight. This demands large-scale, multi-center clinical trials to prove efficacy across diverse populations, alongside collaborative efforts to navigate regulatory pathways. Concurrently, the field must proactively address fundamental challenges of data privacy, algorithmic bias, and equitable access to ensure these powerful technologies benefit all patient demographics fairly and safely.

References