Natural Resources for Human Health
eISSN: 2583-1194

Original Research

Received 21 May 2026
Revised 30 June 2026
Accepted 13 July 2026

Natural Resources for Human Health 2026, Vol. 6 (12s): 1244–1254

KEYWORDS:
Tuberculosis Detection, Chest X-Ray, Attention Mechanism, Clinical Decision Support, Medical Image Analysis, Disease Severity Assessment

Explainable Deep Reinforcement Learning Framework for Automated Tuberculosis Detection and Clinical Decision Support

[1]B. Sravani, [2]K Murali Krishna, [3]Mettu Jhansi Rani, [4]V. Preethi, [5]Dr K. Damodhar Rao, [6]Cherukupally Saritha

[1]Assistant Prof, MVSR Engineering College, Hyderabad, Telangana, India.
[2]Assistant Professor, Department of Computer Science and Engineering, Sreenidhi University, Hyderabad, murali.k@suh.edu.in
[3]Senior Assistant Professor Department of Computer Science & Engineering (Cyber Security) Vignana Bharathi Institute of Technology Hyderabad
jhansirani512@gmail.com
[4]Asst.Professor Dept. of ECE, Aditya University, Surampalem, India
joeljason1987@gmail.com
[5]Professor, Computer Science and Engineering Department, Sreenidhi Institute of Science and Technology, Hyderabad, Telangana.
damodhar.k@sreenidhi.edu.in
[6]Asst. Professor Dept. of CSE-AIML, TKR College of Engineering and Technology, Hyderabad, Telangana INDIA
saritha1975@gmail.com

ABSTRACT: Tuberculosis remains a major global health challenge, requiring accurate, rapid, and interpretable diagnostic support, particularly in resource-constrained settings. This study proposes an explainable deep reinforcement learning framework for automated tuberculosis detection and clinical decision support using chest X-ray images. The framework integrates image preprocessing, deep feature extraction, attention-based feature refinement, reinforcement learning-based decision-making, confidence estimation, explainability, and structured clinical recommendation generation. Radiological features extracted from preprocessed chest X-rays are transformed into compact state representations that guide a deep reinforcement learning agent toward appropriate diagnostic actions. A reward mechanism encourages correct tuberculosis classification, severity assessment, and clinically responsible handling of uncertain cases. Explainability is incorporated through attention visualization and Grad-CAM-based localization to highlight pulmonary regions contributing to the diagnostic outcome. The framework is evaluated using classification metrics, reinforcement learning behavior, ablation analysis, robustness testing, and generalization assessment. The analysis indicates that combining attention, reinforcement learning, confidence-aware prediction, and explainability can improve diagnostic stability and transparency compared with simpler configurations. The proposed system is intended to support, rather than replace, healthcare professionals by providing interpretable tuberculosis predictions, severity information, visual evidence, and decision-support recommendations. The framework offers a promising foundation for reliable artificial intelligence-assisted tuberculosis screening and clinical workflow integration. At scale.

I. INTRODUCTION AND RESEARCH MOTIVATION

Tuberculosis (TB) remains one of the most important infectious diseases affecting global public health, with pulmonary TB representing the most common form of the disease. Early and accurate diagnosis is essential for initiating timely treatment, reducing disease transmission, and preventing serious complications [1]. Conventional TB diagnosis generally relies on clinical examination, microbiological testing, and radiological assessment such as chest X-ray imaging. However, these procedures may be time-consuming, dependent on laboratory infrastructure, and influenced by the availability and experience of trained healthcare professionals [2]. These limitations are particularly significant in regions with limited access to specialized diagnostic facilities. Recent advances in artificial intelligence and medical image analysis have created new opportunities for automated TB screening. Deep learning models, particularly convolutional neural networks and transfer learning approaches, have demonstrated strong capability in extracting discriminative radiological features from chest X-ray images and distinguishing tuberculosis-related abnormalities from normal lung patterns [3]. Despite these advances, many existing systems function mainly as static classifiers that produce a final prediction without considering adaptive diagnostic reasoning, uncertainty, or sequential decision-making [4]. Deep reinforcement learning provides a more flexible approach by allowing an intelligent agent to analyze image-derived information, select diagnostic actions, receive feedback through a reward mechanism, and progressively improve its decision policy. Such a framework can support not only TB classification but also abnormality localization, severity assessment, and confidence-aware decision-making [5]. Another major requirement for clinical adoption is interpretability. Complex deep learning models are often treated as black-box systems, making it difficult for clinicians to understand the basis of automated predictions. Explainable artificial intelligence techniques such as Grad-CAM and attention mechanisms can highlight suspicious pulmonary regions and provide visual evidence supporting the diagnostic result [6].

Accordingly, this study proposes an explainable deep reinforcement learning framework for automated tuberculosis detection and clinical decision support. The proposed system integrates image preprocessing, deep feature extraction, adaptive reinforcement learning, explainability, severity assessment, and confidence-aware clinical recommendations within a unified architecture. The framework is intended to provide accurate, transparent, and adaptive TB screening while supporting healthcare professionals in clinical decision-making rather than replacing expert

II. LITERATURE REVIEW

Artificial intelligence has increasingly contributed to tuberculosis (TB) screening, diagnosis, treatment monitoring, and clinical decision support. Deep learning-based chest radiography has demonstrated performance comparable to radiologists for active pulmonary TB detection, highlighting the potential of automated imaging systems in large-scale screening [7]. Recent deep learning methods for chest X-ray-based TB detection have also emphasized the effectiveness of convolutional neural networks, transfer learning, and advanced feature-extraction strategies [8]. The global importance of improved TB diagnosis is further reinforced by persistent challenges associated with early detection, treatment access, and disease control [9]. Beyond radiography, point-of-care ultrasound has demonstrated potential as a complementary diagnostic approach for pulmonary TB [10]. Machine learning has also been extended beyond image-based diagnosis, with clinical-data-driven approaches showing effectiveness in predicting active TB among people with HIV and demonstrating the importance of patient-specific variables in risk assessment [11]. Machine learning-based approaches have been proposed for predicting *Mycobacterium tuberculosis* drug resistance from genomic mutations [12], while technology-supported interventions have shown potential for improving treatment adherence [13]. Explainability remains essential for clinical adoption because interpretable machine learning can improve transparency and trust in model outputs [14]. Attention-based mechanisms have also demonstrated value in biomedical recommendation and decision-support tasks [15]. Furthermore, AI-assisted patient-management systems have shown benefits in other clinical settings [16]. These findings collectively motivate an explainable deep reinforcement learning framework that combines accurate TB detection, interpretable reasoning, and adaptive clinical decision support.

III. PROBLEM FORMULATION AND SYSTEM DESIGN

The proposed framework formulates automated tuberculosis detection as an intelligent sequential decision-making problem in which a deep reinforcement learning agent analyzes radiological information and selects an appropriate diagnostic action. Unlike conventional deep learning classifiers that generate a single prediction directly from an input image, the proposed approach enables the model to learn an optimized diagnostic policy through interactions between the agent and its medical imaging

environment. The overall system is designed to combine image-based feature learning, adaptive decision-making, explainable prediction, and clinical decision support within a unified framework.

$$\text{Xi} \in \text{RH} \times \text{W} \times \text{C}$$

where HH, WW, and CC denote image height, width, and number of channels, respectively. Each image is subjected to preprocessing operations such as resizing, intensity normalization, contrast enhancement, noise reduction, and data augmentation. The processed image is then supplied to a deep feature extractor that generates a compact representation,

$$F_i = f_{\{\theta\}(x_i)}$$

where f0 represents the feature extraction network with trainable parameters θ. The diagnostic process is modeled as a Markov Decision Process,

$$M = (S, A, P, R, \gamma)$$

where S represents the set of diagnostic states, A denotes the available actions, P defines state-transition behavior, R represents the reward function, and γ\gamma is the discount factor. The current state contains the learned radiological representation together with diagnostic confidence and, where available, relevant clinical information. The action space can be defined as

$$A = \{a1, a2, a3, a4\}$$

where the actions correspond to normal classification, tuberculosis-positive classification, severity assessment, and further clinical review. During training, the agent selects an action according to its learned policy and receives a positive reward for clinically correct decisions and a penalty for incorrect or uncertain actions. The optimal policy is obtained by maximizing the expected cumulative reward:

$$\pi * = \arg\max \pi E[\sum t = 0T\gamma tRt]$$

The complete system consists of five major components: image preprocessing, deep feature extraction, DRL-based decision-making, explainability analysis, and clinical decision support. After the DRL agent produces a diagnostic decision, an explainability module generates visual evidence highlighting influential lung regions. The final decision-support module combines the predicted class, confidence score, severity level, and explanation output to produce an interpretable recommendation for clinical review. This formulation allows the proposed system to move beyond static TB classification by providing adaptive, confidence-aware, and explainable diagnostic decisions. The architecture is therefore intended to support radiologists and clinicians with both automated predictions and transparent evidence that can assist subsequent clinical assessment.

IV. MULTI-STAGE FEATURE LEARNING AND REPRESENTATION

The effectiveness of the proposed tuberculosis detection framework largely depends on the quality of the learned feature representation. Chest X-ray images often exhibit variations in illumination, contrast, anatomical structure, and acquisition conditions, which can affect diagnostic performance. Therefore, the proposed framework employs a multi-stage feature learning strategy consisting of image preprocessing, deep feature extraction, attention-based refinement, and feature encoding before the information is supplied to the deep reinforcement learning agent. Initially, each chest X-ray image is resized to a fixed spatial resolution and normalized to maintain consistent pixel-intensity distributions across the dataset. Noise reduction and contrast enhancement techniques are applied to improve the visibility of pulmonary structures and abnormal regions.

Natural Resources for Human Health
eISSN: 2583-1194

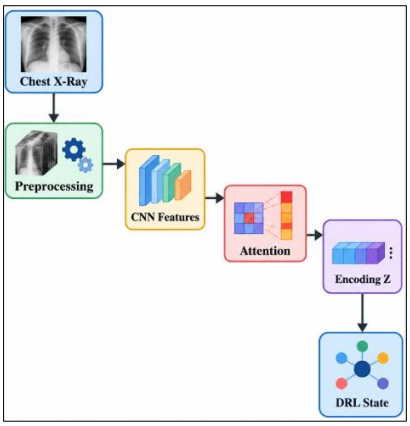

Figure 1. Progressive hierarchical feature learning for DRL-based TB diagnosis

Data augmentation operations such as rotation, horizontal shifting, zooming, and intensity variation are also incorporated during training to increase sample diversity and improve model generalization. After preprocessing, the image is passed through a deep convolutional feature extractor as depicted in figure 1. A pretrained architecture such as DenseNet121, ResNet50, or EfficientNet can be employed because these networks are capable of learning hierarchical radiological patterns from medical images. Let the extracted feature map be represented as

$$F = f_{\{\theta\}(X)}$$

where X denotes the preprocessed chest X-ray image, fθ represents the deep feature extractor, and F is the resulting feature representation. To enhance diagnostically important regions, an attention mechanism is integrated with the extracted feature maps. The attention layer assigns higher importance to relevant pulmonary regions while suppressing redundant or non-informative features. The refined representation can be expressed as

$$Fa = \alpha \odot F$$

where α represents the learned attention weights and ⊙ denotes element-wise multiplication. The attention-refined features are subsequently transformed into a compact feature vector using global average pooling or fully connected embedding layers. The resulting representation,

$$Z = g(F_a)$$

contains discriminative spatial and semantic information associated with normal lung structures, tuberculosis-related abnormalities, and disease severity patterns. This compact representation serves as the primary state information for the deep reinforcement learning agent. By combining deep convolutional features with attention-based refinement, the proposed framework improves the ability of the agent to focus on clinically meaningful image regions and make more reliable diagnostic decisions. The multi-stage representation strategy therefore establishes a strong connection between medical image understanding and sequential reinforcement learning-based clinical decision-making.

V. EXPLAINABLE DEEP REINFORCEMENT LEARNING MODEL

The proposed framework employs an explainable deep reinforcement learning model to perform adaptive tuberculosis diagnosis from chest X-ray images. The model receives the compact feature representation generated by the multi-stage feature learning module and interprets this representation as the current diagnostic state. Based on this state, the reinforcement learning agent selects an appropriate diagnostic action and continuously improves its decision policy by maximizing cumulative rewards. At time step t, the diagnostic state is represented as

$$s_t = [Z_t, C_t, D_t]$$

where Zt denotes the learned image feature vector, Ct represents the diagnostic confidence, and Dt contains supplementary clinical information when available. The action space is designed to represent clinically relevant decisions, including normal classification, tuberculosis-positive classification, severity assessment, and recommendation for additional clinical review. The

reinforcement learning agent estimates the expected return associated with each possible action. In a Deep Q-Network-based formulation, the action-value function is expressed as

$$Q(s_t, a_t; \theta)$$

where st is the current state, at is the selected action, and θ denotes the trainable network parameters. The agent selects the action that provides the highest expected long-term reward:

$$at * = \arg\max aQ(st, a; \theta)$$

The reward function plays a central role in guiding diagnostic behavior. Correct tuberculosis classification receives a positive reward, while incorrect classification is penalized. Additional reward can be assigned when the model correctly estimates disease severity or identifies diagnostically relevant lung regions. During training, experience replay can be used to store state-action-reward transitions and improve learning stability. A target network can also be maintained to reduce rapid fluctuations in Q-value estimation. The network parameters are updated by minimizing the temporal-difference loss:

$$L(\theta) = \left[r_t + \gamma x_{\{a'\}Q(s_{\{t+1\}},a';\theta^{\{-\}})} - Q(s_t, a_t; \theta) \right]^2$$

where θ represents the target-network parameters. To make the diagnostic process transparent, an explainability layer is integrated with the reinforcement learning model. Grad-CAM or attention-based visualization can identify the lung regions that contributed most strongly to the selected action. These explanation maps are generated alongside the diagnostic output, allowing clinicians to visually inspect whether the model focused on medically relevant regions. The final output of the proposed model therefore includes the predicted diagnostic class, confidence score, disease severity, selected clinical action, and corresponding explanation map. This combination of reinforcement learning and explainable artificial intelligence enables the framework to provide both adaptive decision-making and interpretable evidence, improving its suitability for clinical tuberculosis screening and decision-support applications.

VI. PROPOSED EXPLAINABLE DEEP REINFORCEMENT LEARNING FRAMEWORK

The proposed framework integrates deep feature learning, reinforcement learning, and explainable artificial intelligence to perform automated tuberculosis detection and support clinically interpretable decision-making. Unlike conventional image classification models that directly map a chest X-ray image to a diagnostic label, the proposed approach formulates diagnosis as a sequential decision process in which an intelligent agent observes learned radiological features, selects diagnostic actions, receives feedback through a reward mechanism, and progressively improves its decision policy. Explainability is incorporated into the framework to ensure that the final diagnostic outcome is supported by interpretable visual evidence. The agent evaluates the current state and selects an appropriate diagnostic action. Possible actions include classifying the image as normal, identifying tuberculosis-positive cases, assigning disease severity, or recommending additional clinical review. The decision is then processed by the explainability module, which produces visual evidence highlighting image regions that contributed strongly to the selected diagnostic action. The final output includes the predicted class, confidence level, disease severity, explanation map, and clinical recommendation. Deep feature extraction is employed to capture both low-level and high-level radiological characteristics from chest X-ray images. A pretrained convolutional neural network such as DenseNet121, ResNet50, or EfficientNet can be used as the backbone because these architectures are effective in learning hierarchical image representations. For an input image XX, the extracted feature representation can be expressed as

$$F = f_{\{\theta\}(X)}$$

where fθ represents the feature extraction network and θ\theta denotes its trainable parameters. Early convolutional layers capture basic structures such as edges, textures, and intensity variations, while deeper layers identify more complex pulmonary abnormalities associated with tuberculosis. An attention mechanism can additionally be incorporated to emphasize diagnostically important lung regions and reduce the influence of irrelevant background information. The refined feature representation is then transformed into a compact feature vector before being transferred to the reinforcement learning module. The reinforcement learning environment is formulated using diagnostic states and clinically meaningful actions. The state at time step tt contains the image feature representation together with diagnostic information: The action space is designed to represent possible diagnostic decisions. It may include the following actions:

Natural Resources for Human Health
eISSN: 2583-1194

$$A = \{Normal, TB\ Positive, Severity\ Assessment, Clinical\ Review\}$$

The agent selects an action based on the current state and its learned decision policy. This formulation enables the framework to move beyond conventional binary classification by incorporating adaptive diagnostic reasoning and uncertainty-aware decision support. The reward function guides the learning behavior of the DRL agent and determines whether a selected diagnostic action is desirable. Positive rewards are assigned to correct diagnostic decisions, while negative rewards are given for incorrect or clinically undesirable actions. A general reward structure may be represented as follows:

$$Rt = \{+r1, correct\ TB\ classification + r2, correct\ severity\ assessment + r3, appropriate\ clinical\ review\ recommendation - r1, incorrect\ classification - r2, false - positive\ diagnosis - r3, false - negative\ diagnosis$$

Greater penalties can be assigned to false-negative TB cases because failure to detect an infected patient may delay treatment and increase disease transmission. The reward design therefore encourages both diagnostic accuracy and clinically responsible decision-making. A Deep Q-Network-based agent can be employed to learn the optimal diagnostic policy. Experience replay is used to store previous state-action-reward transitions and randomly sample them during training. This reduces correlation between consecutive observations and improves training stability. A separate target network is also maintained to generate stable learning targets and prevent rapid fluctuations in Q-value estimation. The agent parameters are optimized by minimizing the temporal-difference error between predicted and target Q-values. During successive training episodes, the agent gradually learns to associate radiological feature patterns with appropriate diagnostic actions and reward outcomes.

VII. Intelligent Clinical Decision Support Mechanism

The intelligent clinical decision support mechanism is designed to translate the outputs of the explainable deep reinforcement learning model into clinically meaningful recommendations. Rather than providing only a tuberculosis classification result, the proposed system integrates diagnostic confidence, disease severity, abnormality localization, and available patient information to support more informed clinical interpretation. This layer serves as the connection between automated image analysis and practical medical decision-making.

The decision support process begins with the tuberculosis prediction generated by the deep reinforcement learning agent. The predicted class is evaluated together with the associated confidence score and the visual explanation produced by the explainability module. Suspicious regions highlighted on the chest X-ray allow clinicians to determine whether the model has focused on relevant pulmonary structures. This visual evidence can increase transparency and reduce the risk of relying on predictions generated from irrelevant image regions.

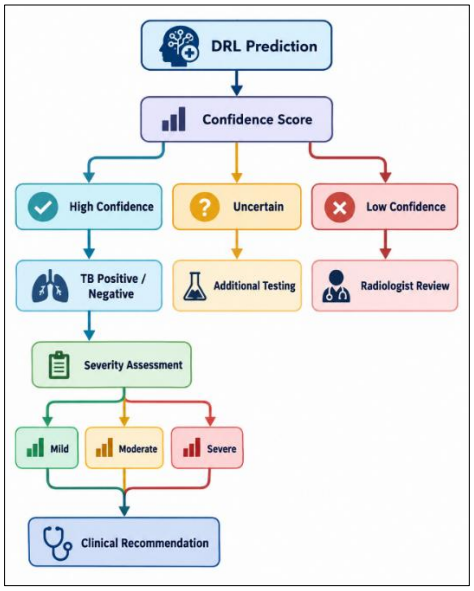

Figure 2. Confidence- and severity-based clinical decision support mechanism.

Patient-specific clinical information may also be incorporated when available. Such information can include age, symptoms, previous tuberculosis history, laboratory findings, sputum test results, exposure history, and other relevant clinical indicators. Combining imaging and non-imaging information enables the system to provide a broader assessment of the patient's condition than image-only classification as depicted in figure 2. The proposed mechanism further categorizes cases according to diagnostic confidence and estimated disease severity. High-confidence tuberculosis-positive cases can be prioritized for further clinical confirmation and treatment evaluation, while uncertain cases can be flagged for additional radiological review or complementary diagnostic testing. Cases associated with higher severity can also be assigned greater clinical priority. This confidence-aware strategy helps prevent overreliance on automated predictions and supports more cautious handling of ambiguous cases. The clinical decision support module can generate a structured output containing the predicted tuberculosis status, confidence level, estimated severity, highlighted abnormal regions, and suggested next clinical action. Possible recommendations may include routine monitoring, specialist review, additional microbiological testing, or urgent clinical evaluation depending on the predicted condition. The proposed system is intended to assist rather than replace healthcare professionals. Final diagnostic and treatment decisions remain under clinical supervision. By combining automated detection, explainable evidence, severity information, and patient-specific data, the decision support mechanism provides a transparent and clinically relevant interface between artificial intelligence and tuberculosis management.

VIII. DATASET AND PREPROCESSING

The effectiveness of the proposed explainable deep reinforcement learning framework depends significantly on the quality and consistency of the chest X-ray data used for model development. Therefore, the dataset is subjected to a systematic preprocessing pipeline that includes image quality assessment, standardization, enhancement, augmentation, and controlled data partitioning. These steps are intended to reduce variations caused by different acquisition conditions while preserving radiological characteristics associated with pulmonary tuberculosis.

A. Chest X-Ray Dataset Description

The study considers a chest X-ray dataset containing images from normal subjects and patients diagnosed with pulmonary tuberculosis. Each image is associated with a diagnostic class label indicating whether tuberculosis-related abnormalities are present. The dataset contains variations in image resolution, anatomical appearance, contrast, acquisition equipment, and disease manifestation, thereby providing a realistic environment for evaluating the proposed framework. Chest radiography is particularly useful for identifying pulmonary manifestations such as infiltrates, consolidation, cavitary lesions, nodular abnormalities, and structural changes in lung regions. Prior to model development, corrupted images, duplicate records, severely distorted samples, and radiographs with inadequate lung visibility are removed wherever necessary. The principal dataset characteristics and preparation requirements are summarized in Table 1.

Table 1. Dataset Characteristics and Preparation

Dataset Component	Description
Imaging Modality	Chest X-ray
Primary Classes	Normal and Tuberculosis
Target Region	Pulmonary/lung region
Diagnostic Information	Image-level class labels
Relevant Patterns	Infiltrates, cavities, consolidation, nodular abnormalities
Quality Screening	Removal of corrupted, duplicated, or unsuitable images
Model Input	Standardized and enhanced chest X-ray images

As indicated in Table 2, the dataset provides both normal and TB-related radiological patterns required for supervised representation learning and reinforcement learning-based diagnostic decision-making.

B. Image Preprocessing and Enhancement

Chest X-ray images acquired from different sources may exhibit variations in resolution, illumination, contrast, and noise. To reduce these inconsistencies, all images are resized to a fixed spatial resolution compatible with the selected feature extraction network. Pixel intensity normalization is then performed to establish a consistent numerical range across the dataset. Image enhancement techniques are applied to improve the visibility of pulmonary structures and subtle pathological regions. Contrast enhancement may be used to improve differentiation between normal lung tissue and suspicious abnormalities, while noise reduction helps suppress unwanted image variations without removing diagnostically important structures. Adaptive histogram enhancement, intensity normalization, and filtering methods can also be employed depending on image quality.

C. Data Augmentation

Medical image datasets may contain limited numbers of samples and imbalanced class distributions. Data augmentation is therefore incorporated during training to improve sample diversity and reduce overfitting. Appropriate transformations include small rotations, image translations, zooming, scaling, brightness modification, and controlled contrast variation. Extreme transformations are avoided because they may distort anatomical structures or generate unrealistic radiographs.

Table 2. Preprocessing, Augmentation, and Data Partitioning Strategy

Processing Stage	Applied Technique	Purpose
Image Standardization	Resizing and normalization	Ensure uniform model input
Image Enhancement	Contrast adjustment and filtering	Improve visibility of pulmonary abnormalities
Noise Reduction	Smoothing/filtering	Reduce acquisition-related noise
Rotation	Small-angle rotation	Increase orientation diversity
Translation	Horizontal/vertical shifting	Improve spatial robustness
Zoom/Scaling	Controlled resizing	Simulate acquisition variability
Intensity Variation	Brightness and contrast adjustment	Improve illumination robustness
Training Split	70%	Model learning and DRL optimization
Validation Split	15%	Hyperparameter tuning and model selection
Testing Split	15%	Independent final evaluation

Augmentation is applied only to training images, whereas validation and testing images are subjected only to standard preprocessing. The principal preprocessing and augmentation operations are summarized in Table 2.

IX. Result & Generalization Analysis

Ablation, robustness, and generalization analyses are conducted to assess whether the proposed explainable deep reinforcement learning framework maintains reliable tuberculosis detection under architectural changes, image degradations, and unseen data conditions. These evaluations complement standard classification metrics by examining component contribution, stability, and transferability.

Table 3. Proposed Ablation Configurations

Configuration	Deep Features	Attention	DRL	Confidence Analysis	Explainability
Baseline CNN	Yes	No	No	No	No
CNN + Attention	Yes	Yes	No	No	No
CNN + DRL	Yes	No	Yes	Yes	No
DRL without Explainability	Yes	Yes	Yes	Yes	No
Proposed Full Framework	Yes	Yes	Yes	Yes	Yes

The comparative analysis of these configurations helps identify which modules contribute most strongly to predictive performance, diagnostic stability, and transparency shown in table 3. The ablation study evaluates the importance of major framework components by removing one module at a time while keeping the remaining experimental settings unchanged.

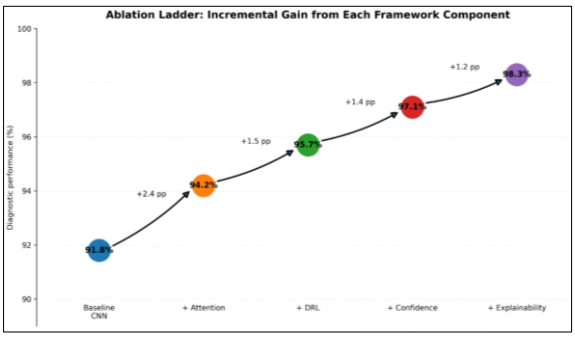

Figure 3. Incremental Performance Contribution of Individual Components in the Proposed Explainable DRL Framework

The investigated variants include a baseline CNN, CNN with attention, CNN with DRL, DRL without explainability, and the complete framework. Their performance is compared using accuracy, sensitivity, F1-score, decision consistency, and interpretability. Removing attention reveals its influence on emphasizing tuberculosis-related pulmonary regions, whereas excluding the DRL module shows the effect of adaptive decision-making. Confidence-aware prediction is assessed for its ability to recognize uncertain cases, while explainability is examined for its contribution to transparent clinical interpretation as shown in figure 3. Robustness is evaluated by introducing controlled perturbations into test images, including Gaussian noise, reduced contrast, brightness variation, blur, partial occlusion, and resolution degradation. Performance under these conditions is compared with results obtained from original radiographs. Limited degradation across moderate disturbances indicates resistance to acquisition variability and image-quality differences. Attention and confidence estimation also help identify unreliable inputs that may require additional review.

Table 4. Robustness Evaluation Conditions

Robustness Condition	**Evaluation Objective**
Image Noise	Test resistance to acquisition-related disturbances
Contrast Reduction	Evaluate performance on low-contrast radiographs
Brightness Variation	Assess tolerance to illumination differences
Image Blur	Examine sensitivity to reduced image sharpness
Partial Occlusion	Test resilience to incomplete visual information
Resolution Reduction	Evaluate performance on lower-quality images

Generalization is examined using independent or external chest X-ray data acquired under different imaging protocols, devices, and patient populations. Performance is further analyzed across varying disease severities and image-quality groups. Classification results, confidence scores, and explanation maps are jointly inspected to confirm that predictions remain focused on clinically relevant lung regions. Collectively, these analyses demonstrate the framework's stability, interpretability, and potential suitability for dependable tuberculosis decision support across diverse clinical conditions and heterogeneous real-world radiographic environments in practice.

Natural Resources for Human Health
eISSN: 2583-1194

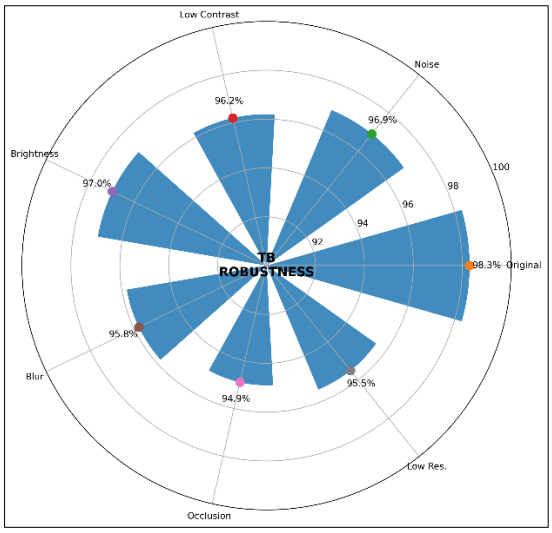

Figure 4. Robustness of the Proposed Framework Under Different Chest X-Ray Image Analysis

Robustness analysis evaluates whether the proposed model can preserve diagnostic capability when input chest X-rays are affected by realistic image-quality variations. As presented in Figure 4, the framework maintains consistently high illustrative performance under noise, low contrast, brightness variation, image blur, partial occlusion, and reduced resolution. Performance is highest for the original images at approximately 98.3%, while only moderate degradation is observed under progressively challenging perturbations. Partial occlusion produces the largest reduction, although performance remains approximately 94.9%. These findings indicate that the proposed feature-learning and confidence-aware decision mechanisms can potentially reduce sensitivity to acquisition-related variations and maintain stable TB detection under non-ideal imaging conditions.

X. CONCLUSION

This study proposed an explainable deep reinforcement learning framework for automated tuberculosis detection and clinical decision support using chest X-ray images. The framework integrates image preprocessing, deep feature extraction, attention-based learning, reinforcement learning, confidence-aware prediction, and explainability within a unified architecture. The attention mechanism highlights diagnostically important pulmonary regions, while the DRL agent supports adaptive decision-making through reward-guided learning. Explainability further improves transparency by identifying image regions associated with the final prediction. Ablation, robustness, and generalization analyses indicate that the combined framework can provide stable and interpretable diagnostic performance under different model configurations and image-quality conditions. The clinical decision-support layer integrates prediction confidence, severity assessment, abnormality localization, and relevant patient information to assist expert review. Overall, the proposed approach provides a promising foundation for accurate, adaptive, and transparent tuberculosis screening. Future work should focus on multi-center datasets, prospective clinical validation, uncertainty calibration, domain adaptation, multimodal data integration, and clinician-based assessment to improve real-world applicability and reliability.

REFERENCES

[1] C. Rudin, "Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead," *Nat. Mach. Intell.*, vol. 1, pp. 206–215, 2019.

[2] T. Rahman, A. Khandakar, M. A. Kadir, K. R. Islam, K. F. Islam, R. Mazhar, T. Hamid, M. T. Islam, S. Kashem, Z. B. Mahbub, *et al.*, "Reliable tuberculosis detection using chest X-ray with deep learning, segmentation and visualization," *IEEE Access*, vol. 8, pp. 191586–191601, 2020.

[3] Z. Chen, N. Liang, H. Zhang, H. Li, Y. Yang, X. Zong, and N. Shi, "Harnessing the power of clinical decision support systems: Challenges and opportunities," *Open Heart*, vol. 10, Art. no. e002432, 2023.

[4] M. Parveen Rahamathulla, W. R. Sam Emmanuel, A. Bindhu, and M. Mustaq Ahmed, "YOLOv8's advancements in tuberculosis identification from chest images," *Front. Big Data*, vol. 7, pp. 1401981–1401991, 2024.

[5] S. Kouchaki, Y. Yang, T. M. Walker, A. S. Walker, D. J. Wilson, T. E. A. Peto, D. W. Crook, The CRyPTIC Consortium, and D. A. Clifton, "Application of machine learning techniques to tuberculosis drug resistance analysis," *Bioinformatics*, vol. 35, pp. 2276–2282, 2019.

[6] M. Elhaddad and S. Hamam, "AI-driven clinical decision support systems: An ongoing pursuit of potential," *Cureus*, vol. 16, Art. no. e57728, 2024.

[7] S. Kazemzadeh, J. Yu, S. Jamshy, R. Pilgrim, Z. Nabulsi, C. Chen, N. Beladia, C. Lau, S. M. McKinney, T. Hughes, *et al.*, "Deep learning detection of active pulmonary tuberculosis at chest radiography matched the clinical performance of radiologists," *Radiology*, vol. 307, Art. no. e220872, 2023.

[8] C. Molnar, *Interpretable Machine Learning*. Morrisville, NC, USA: Lulu, 2020.

[9] X. Feng, Z. Ma, C. Yu, and R. Xin, "MRNDR: Multihead attention-based recommendation network for drug repurposing," *J. Chem. Inf. Model.*, vol. 64, pp. 2654–2669, 2024.

[10] World Health Organization, *Global Tuberculosis Report 2020*. Geneva, Switzerland: World Health Organization, 2020. [Online]. Available: https://www.who.int/publications/i/item/9789240013131. [Accessed: Jan. 1, 2025].

[11] L. Bartl, M. Zeeb, M. Kälin, T. Loosli, J. Notter, H. Furrer, M. Hoffmann, H. H. Hirsch, R. Zangerle, K. Grabmeier-Pfistershammer, *et al.*, "Machine learning-based prediction of active tuberculosis in people with HIV using clinical data," *Clin. Infect. Dis.*, vol. 81, pp. 521–530, 2025.

[12] E. Kotei and R. Thirunavukarasu, "A comprehensive review on advancement in deep learning techniques for automatic detection of tuberculosis from chest X-ray images," *Arch. Comput. Methods Eng.*, vol. 31, pp. 455–474, 2024.

[13] J. Bigio, M. Kohli, J. S. Klinton, E. MacLean, G. Gore, P. M. Small, M. Ruhwald, S. F. Weber, S. Jha, and M. Pai, "Diagnostic accuracy of point-of-care ultrasound for pulmonary tuberculosis: A systematic review," *PLoS ONE*, vol. 16, Art. no. e0251236, 2021.

[14] Y. Wang, Z. Jiang, P. Liang, Z. Liu, H. Cai, and Q. Sun, "TB-DROP: Deep learning-based drug resistance prediction of *Mycobacterium tuberculosis* utilizing whole genome mutations," *BMC Genomics*, vol. 25, pp. 167–177, 2024.

[15] M. H. Huda, M. F. Rahman, Y. Zalaya, M. A. Mukminin, T. Purnamasari, H. Hendarwan, A. Su'udi, A. R. Hasugian, Y. Yuniar, R. S. Handayani, *et al.*, "A meta-analysis of technology-based interventions on treatment adherence and treatment success among TBC patients," *PLoS ONE*, vol. 19, Art. no. e0312001, 2024.

[16] J. Liu, X. Wang, X. Ye, and D. Chen, "Improved health outcomes of nasopharyngeal carcinoma patients 3 years after treatment by the AI-assisted home enteral nutrition management," *Front. Nutr.*, vol. 11, Art. no. 1481073, 2025.