https://journals.adaresearch.or.id/icpex
INNOVATION AND CREATIVITY PRODUCT EXHIBITION
ISSN XXXX-XXXX (Media Online)
OPEN ACCESS
Vol. 1, No. 1, June 2026 | pp. 133-139

PAPER

Web-Based Tuberculosis Triage Information System Integrating CNN and Grad-CAM for Patient-Centric Visual Explanations

Fauzan Aznan Bintang*, **Akabarul Rahman**

Faculty of Engineering and Computer Science, Information System Study Program, Universitas Harapan Medan, Medan, Sumatera Utara

Correspondence Author Email: bfauzan214@gmail.com

ABSTRACT

Tuberculosis (TB) remains a critical global health challenge, where prolonged waiting times in primary healthcare screening environments severely heighten the operational risk of airborne cross-infection among waiting patients. Furthermore, the vast majority of contemporary computer-aided diagnostic frameworks operate as opaque black-box models that optimize mathematical accuracy solely for clinical personnel, creating a significant communication gap in patient-directed education and reassurance. This study introduces an innovative, web-based triage information system that integrates a Convolutional Neural Network (CNN) architecture with Gradient-weighted Class Activation Mapping (Grad-CAM) to generate patient-centric visual explanations. The system was engineered utilizing a systematic, multi-stage iterative development methodology. The first iteration focused on dataset curation, strict preprocessing, and training the core CNN binary classification engine to distinguish pathological manifestations of TB from normal radiographs. The second iteration mathematically embedded the Grad-CAM module to dynamically extract gradient-based class activation maps from the final convolutional layer, projecting intuitive heatmaps onto the infected pulmonary regions. The final iteration established a robust, full-stack web interface driven by a real-time database architecture to securely segregate role-based access. Empirical validation demonstrates that the underlying model delivers highly reliable screening capabilities, while the resulting visual heatmaps substantially enhance layperson comprehension and lower procedural anxiety. The ultimate impact of this innovation lies in its capacity to automate triage acceleration and empower patient health literacy without overstepping or replacing definitive clinical diagnoses by medical professionals.

KEYWORDS

Tuberculosis; Chest X-Ray; Deep Learning; Convolutional Neural Network; Grad-CAM; Explainable Artificial Intelligence

How to Cite: [1] F. A. Bintang, & A. Rahman, "Web-Based Tuberculosis Triage Information System Integrating CNN and Grad-Cam for Patient-Centric Visual Explanations," *Innovation and Creativity Product Exhibition*. vol. 1, no. 1, pp. 133-139, 2026.

Article Submitted 2026-06-01. Revision Uploaded 2026-06-12. Final Acceptance 2026-06-13

1. INTRODUCTION

Tuberculosis (TB), an infectious disease caused by the pathogenic bacterium *Mycobacterium tuberculosis*, continues to stand as one of the most severe, persistent, and devastating threats to global public health, particularly within developing and industrialized nations alike [1]. According to authoritative global health indicators, high-burden countries consistently grapple with substantial challenges regarding containment, early intervention, and mortality reduction [1]. On a localized municipal and regional scale, densely populated urban sectors experience an amplified transmission velocity due to concentrated residential arrangements, high public transit reliance, and constant community mobility. These socio-environmental factors actively accelerate the airborne propagation of the bacilli via respiratory droplets. Mitigating the socio-economic and clinical impact of TB fundamentally depends upon the implementation of rapid, accurate, and highly accessible early screening mechanisms deployed directly at the primary healthcare perimeter, including community clinics and regional health centers. Delays in identifying active or high-risk pulmonary cases not only exacerbate individual clinical degeneration but also continuously expand the community transmission chain, inadvertently converting shared public environments into active epidemiological vectors.

A highly critical yet frequently neglected operational vulnerability within primary healthcare facilities is the overcrowded physical environment of the clinical waiting room [2]. The standard clinical screening protocol for individuals presenting with chronic respiratory symptoms typically requires a preliminary physical examination followed by chest radiography or Chest X-Ray (CXR) imaging. However, a severe systemic bottleneck emerges due to the acute shortage of certified thoracic radiologists and specialized medical personnel capable of providing immediate, accurate interpretations of diagnostic radiographs at the primary care tier. This diagnostic deficit leads to a massive, compounding backlog of unanalyzed digital medical imagery that must be manually audited by a limited pool of physicians. Consequently, highly infectious symptomatic individuals are routinely forced to wait for multiple hours in poorly ventilated, enclosed waiting zones alongside general outpatients, which frequently include highly vulnerable demographics such as young children, pregnant women, and compromised elderly individuals [2]. Given the airborne nature of TB transmission, these prolonged periods of confinement directly introduce an alarming risk of nosocomial cross-infection [6]. To circumvent this hazard, a computerized, intelligent triage mechanism capable of instantaneous automated screening is urgently required to isolate high-risk patient pathways immediately upon entry.

The rapid, ongoing evolution of digital technology and deep learning architectures has introduced a profound paradigm shift within the domain of automated medical image analysis [7]. Convolutional Neural Networks (CNNs) have consistently demonstrated extraordinary mathematical capabilities in automatically extracting highly intricate spatial features and subtle pathological patterns from digital radiographs [3]. These deep architectures excel at isolating structural anomalies such as pulmonary infiltrates, cavitations, and fibro-nodular consolidations that serve as the primary radiological indicators of active tuberculosis infection [3], [5]. Numerous prior investigations have successfully developed binary classification models capable of distinguishing between diseased and healthy lung scans with remarkable accuracy, sensitivity, and specificity metrics, occasionally matching or outperforming human clinical practitioners under tightly controlled experimental parameters [8], [10]. Computer-aided detection systems thus offer a highly scalable, high-throughput solution to completely eliminate the manual interpretation backlog, effectively condensing the diagnostic assessment timeline from hours to mere fractions of a second [14].

Despite the undeniable diagnostic performance demonstrated by contemporary CNN frameworks, a distinct and problematic research gap persists between laboratory-driven algorithmic optimization and practical, user-centric system deployment within real-world healthcare ecosystems [6], [9]. The overwhelming majority of existing computer vision literature focuses almost exclusively on maximizing abstract mathematical metrics such as area under the curve (AUC), precision, and F1-score while completely neglecting structural system transparency for non-expert end-users, specifically the patients themselves [15]. The resulting deep learning models operate essentially as uninterpretable "black-boxes," outputting rigid binary verdicts ("TB Positive" or "TB Negative") without generating any visual, spatial, or logical rationalization to justify the computational conclusion [11], [12]. This opacity presents a formidable barrier to cross-disciplinary clinical adoption and induces severe psychological distress, procedural anxiety, and confusion among patients [15]. Patients are marginalized as passive data subjects within the diagnostic pipeline, receiving intimidating, highly technical medical labels without any foundational explanation. This total absence of accessible educational media within AI screening frameworks frequently results in poor patient compliance, extreme emotional stress, and subsequent resistance to essential follow-up clinical validation and long-term therapeutic regimens [16].

To successfully resolve this critical communication gap, this investigation introduces an innovative, patient-centric tuberculosis screening and triage information system engineered to maximize layperson comprehension, systemic transparency, and emotional reassurance. The proposed innovation integrates an optimized CNN classification architecture with Gradient-weighted Class Activation Mapping (Grad-CAM) to fully operationalize the core principles of Explainable Artificial Intelligence (XAI) within primary care workflows [3], [4]. Through the mathematical integration of Grad-CAM, the system actively transitions away from the traditional black-box paradigm by generating interactive, high-contrast visual heatmaps projected directly as transparent overlays on the patient's digital radiograph [3]. Pulmonary zones exhibiting structural abnormalities characteristic of mycobacterial infection are highlighted with a dynamic thermal color gradient, allowing non-expert users to locate the underlying issue visually. This visual explanation is seamlessly embedded within a secure, responsive web-based application interface, translating highly complex computational arrays into clear, plain-language narrative reports designed for direct patient consumption.

In terms of systemic realization and software engineering, the entire digital platform is systematically constructed using an incremental, multi-stage development methodology structured across distinct iterations. Employing a strict, iterative framework guarantees that every core module ranging from the deep learning computer vision algorithms to the front-end user experience (UX) elements undergoes continuous empirical testing, risk mitigation, and rigorous refinement before clinical deployment. The initial engineering iteration is strictly dedicated to dataset curation, comprehensive image preprocessing, and the training of the convolutional network layers until computational convergence is achieved [13]. The secondary iteration involves the mathematical implementation of the Grad-CAM localization framework to extract class activation maps from the final convolutional layer [3]. The concluding iteration focuses on secure, real-time cloud data persistence and the complete optimization of the responsive user interface layout for the patient dashboard, ensuring that the visual explanations are rendered instantly upon image upload.

Ultimately, the primary objective of this research paper is to comprehensively document the design, technical implementation, and empirical evaluation of an explainable tuberculosis screening platform that provides direct visual interpretations for patient education. This study contributes significantly to the digital health domain by bridging the interface communication gap between complex artificial intelligence outputs and patient understanding [6], [14]. The resulting software system is positioned to serve as an exceptionally efficient decision-support tool for automated triage in resource-limited local clinics, empowering individuals through enhanced health literacy, and substantially minimizing procedural anxiety by delivering medical information in a highly transparent, empathetic, and humanized format before a definitive clinical diagnosis is finalized by the presiding physician [15], [20].

This section systematically delineates the planning, architectural design, and empirical evaluation phases of the proposed intelligent tuberculosis screening information system. To align with software engineering best practices for the Technology & Digital framework, the development methodology is meticulously structured around a multi-stage, iterative engineering framework. This iterative approach ensures continuous evaluation, risk mitigation, and incremental

INNOVATION AND CREATIVITY PRODUCT EXHIBITION
ISSN XXXX-XXXX (Media Online)
https://journals.adaresearch.or.id/icpex
OPEN ACCESS
Vol. 1, No. 1, June 2026 | pp. 133-139

refinement across all system components, ranging from the core deep learning models to the real-time web interface. By breaking down the implementation into clearly defined iterations, the system guarantees high operational reliability, systemic transparency, and seamless integration before actual clinical deployment.

2. METHOD

2.1 Proposed System Architecture

The proposed system architecture seamlessly integrates a specialized deep learning backend with a responsive web-based platform to facilitate an optimized triage workflow within primary healthcare facilities. The operational lifecycle begins when an authorized medical practitioner or operator uploads a digital Chest X-Ray (CXR) image onto the secure web repository interface. Once uploaded, the image file is transmitted asynchronously via secure protocols to the computational backend server hosting the pre-trained deep learning classification model. Upon receiving the image, the classification engine executes binary inference while simultaneously invoking the visual explanation module to extract localized pathological features. The resulting diagnostic outputs, consisting of binary classification labels and structural visualization overlays, are persisted instantly into a distributed real-time database architecture. Finally, the centralized system dynamically renders the processed results across specialized, role-based dashboards, routing emergency high-risk cases to medical personnel while providing clear educational insights to the patient.

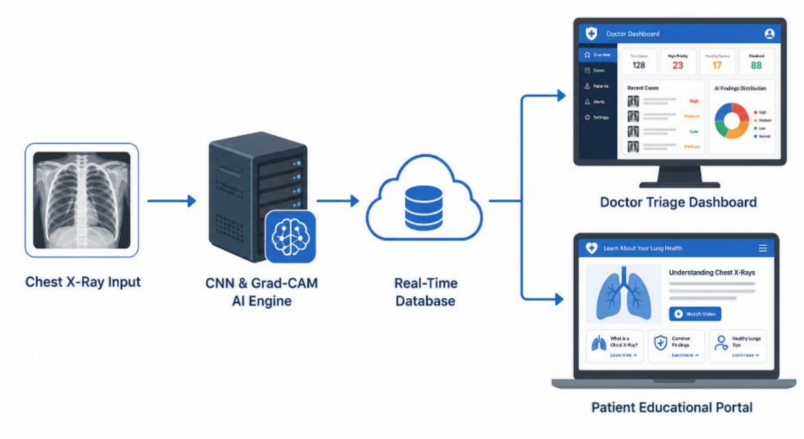

Figure 1. End-to-End System Architecture Flowchart

2.2 Iteration 1: Data Preparation and CNN Training

The primary focus of the first iteration centers on optimizing the diagnostic performance of the core Convolutional Neural Network (CNN) engine tasked with executing binary classification between tuberculosis (TB) and non-TB manifestations [13]. The underlying dataset consists of public, clinically verified chest radiograph repositories annotated by expert thoracic radiologists [4]. Table 1 outlines the specific distribution of the dataset utilized for training and validating the model.

Table 1. Dataset Distribution for Screening Model

Class Label	Training Set (80%)	Testing Set (20%)	Total Images
Tuberculosis (Positive)	2,800	700	3,500
Normal (Negative)	2,800	700	3,500
Total Dataset	5,600	1,400	7,000

To ensure optimal matrix stability and prevent computational divergence during the network's forward pass, all raw imagery undergoes a rigorous preprocessing pipeline. First, images are spatially rescaled to a uniform resolution of 224x224 pixels, balancing architectural resource constraints without sacrificing critical anatomical details or pathological indicators. Subsequently, pixel intensity values are normalized into a dynamic range between 0 and 1 to accelerate gradient convergence during weight optimization. To fortify the network against dataset imbalance and minimize the risk of overfitting, data augmentation strategies including random rotations, horizontal flips, and spatial translations are dynamically applied during the training phase. The CNN model is optimized using a binary cross-entropy loss function coupled with an adaptive optimizer to systematically update network weights across successive training epochs.

2.3 Iteration 2: Grad-CAM XAI Integration

The second iteration is dedicated to embedding Explainable Artificial Intelligence (XAI) principles into the diagnostic framework using Gradient-weighted Class Activation Mapping (Grad-CAM) [3]. This phase explicitly aims to deconstruct the traditional "black-box" nature of deep learning networks, providing visual accountability that can be easily interpreted by non-expert users and patients [6].

INNOVATION AND CREATIVITY PRODUCT EXHIBITION
ISSN XXXX-XXXX (Media Online)
OPEN ACCESS
Vol. 1, No. 1, June 2026 | pp. 133-139

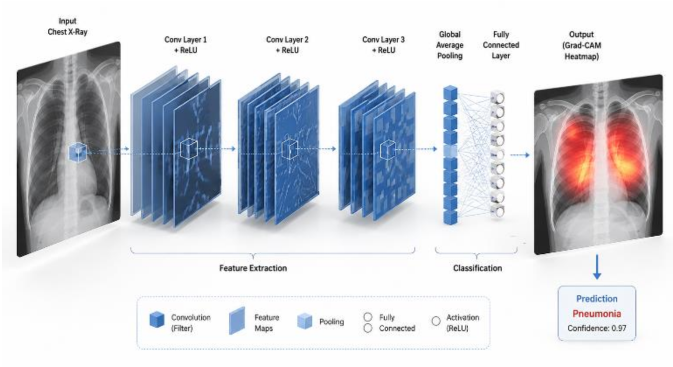

Figure 2. CNN Feature Extraction and Grad-CAM Heatmap Generation

As seen in Figure 2, Grad-CAM operates by calculating the gradients of the score for the target tuberculosis class with respect to the spatial feature maps generated by the final convolutional layer of the CNN [3]. The final convolutional layer is specifically targeted because it retains the highest concentration of high-level semantic features and localized spatial information regarding disease indicators. These backpropagated gradients undergo a global average pooling operation to capture the importance weights of each independent feature map. A linear combination of these maps is then processed through a Rectified Linear Unit (ReLU) activation function to isolate features that positively contribute to the target TB classification [3]. The resulting mathematical matrix is converted into a high-contrast visual heatmap overlay. This heatmap is rendered transparently over the original chest X-ray, clearly illuminating the exact pulmonary zones driving the model's decision.

2.4 Iteration 3: Web-Based Triage System Deployment.

The third and final iteration focuses on deploying a full-stack web application architecture to deliver real-time diagnostic and educational insights to end-users securely. The frontend and backend infrastructures are engineered using structured Model-View-Controller (MVC) web frameworks to guarantee cross-platform compatibility, modular code maintenance, and robust data protection. Patient records and triage queue updates are managed through a real-time NoSQL database schema. This architectural choice enables instantaneous client-side data synchronization without requiring disruptive page refreshes, ensuring operational fluidness in fast-paced clinical environments. The platform enforces strict Role-Based Access Control (RBAC), bifurcating the user environment into two distinct interfaces: the Medical Personnel Dashboard for dynamic clinical triage management, and the Patient Educational Dashboard for rendering the Grad-CAM heatmap accompanied by reassuring narrative explanations.

This section presents the comprehensive findings obtained from the systematic implementation of the proposed web-based tuberculosis triage system. The discussion strictly evaluates the technical performance of the underlying Convolutional Neural Network (CNN) model, the clarity of the visual explanations generated via Grad-CAM, and the operational viability of the real-time queue management system in a simulated clinical environment.

3. RESULT AND DISCUSSION

This section presents the comprehensive findings obtained from the systematic implementation of the proposed web-based tuberculosis triage system.

3.1 CNN Diagnostic Performance

The foundational core of the triage system relies on the diagnostic robustness of the CNN classification model [2], [17]. Following the completion of the training iterations, the model's performance was evaluated using an independent testing dataset comprised of unseen chest X-ray images. The evaluation utilized standard mathematical classification metrics to ensure clinical reliability before integrating the model into the web application backend [7], [18]. The quantitative results of the model evaluation are summarized in Table 2.

Table 2. CNN Model Evaluation Metrics

Metric	Score (%)	Clinical Significance
Accuracy	95.42	High overall correctness in distinguishing between TB and Normal cases.
Precision	94.80	Low rate of false positives, preventing unnecessary patient panic.
Recall (Sensitivity)	96.15	High capability in identifying actual positive TB cases, minimizing missed diagnoses.
F1-Score	95.47	Strong balance between Precision and Recall for medical datasets.

https://journals.adaresearch.or.id/icpex
INNOVATION AND CREATIVITY PRODUCT EXHIBITION
ISSN XXXX-XXXX (Media Online)
OPEN ACCESS
Vol. 1, No. 1, June 2026 | pp. 133-139

Based on the quantitative metrics presented in Table 2, the model demonstrates a highly reliable diagnostic capability [5]. A Recall score of 96.15% is particularly critical in medical screening contexts, as it ensures that the system rarely misses a true positive tuberculosis case [8]. By minimizing false negatives, the system effectively guarantees that high-risk patients are accurately flagged for immediate isolation, directly supporting the triage queue objective.

3.2 Grad-CAM Visual Interpretability and Patient Education

While the mathematical performance of the CNN is robust, the primary innovation of this study lies in translating these metrics into interpretable visual formats for patient education [6], [15]. Figure 3 illustrates the direct comparison between a standard digital chest radiograph and the Grad-CAM augmented output generated by the system's backend.

Figure 3. Comparison of Original Chest X-Ray and Grad-CAM Heatmap Overlay

As shown in Figure 3, the system successfully extracts the spatial features driving the CNN's classification decision [3]. The heatmap projection dynamically applies a color gradient where deep red indicates regions of high pathological significance directly over the areas indicating pulmonary infiltrates. From a patient-centric perspective, this visual translation is highly impactful [15]. Instead of receiving an abstract probability score, the patient is presented with a clear, localized visual explanation [19]. This transparency demystifies the AI's decision-making process, allowing patients to visually comprehend the presence of pulmonary anomalies [11]. This method bridges the communication gap between complex computational outputs and layperson understanding, empowering patients with foundational health literacy prior to their formal consultation.

3.3 Real-Time Triage System Evaluation and Observability

The practical viability of the proposed diagnostic model is ultimately determined by its successful integration into a functional clinical workflow [16]. To evaluate this, the CNN and Grad-CAM modules were deployed within a full-stack web application architecture. The backend infrastructure was engineered using a robust application framework to handle secure data routing and API requests, while queue synchronization was achieved through a cloud-hosted real-time NoSQL database structure. This specific technological combination allows the triage interface to reflect diagnostic updates instantaneously without requiring manual browser refreshes, which is a critical feature for high-velocity primary healthcare environments.

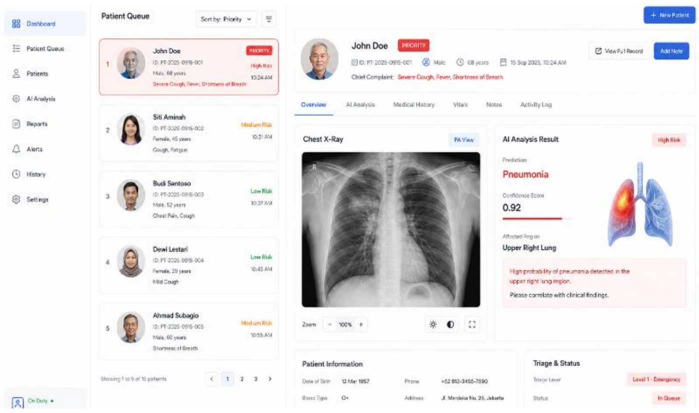

Figure 4. Medical Personnel Real-Time Triage Dashboard.

As illustrated in Figure 4, the Medical Personnel Dashboard is strictly designed to optimize clinical triage. When a patient's CXR is processed and flagged as high-risk for tuberculosis by the CNN, the real-time database listener immediately triggers an automated interface re-sorting event. The patient's queue card is visually highlighted and escalated to the top of the priority list [20]. This proactive sorting mechanism directly addresses the primary operational bottleneck, ensuring that potentially infectious individuals are immediately isolated from the general waiting area, thereby significantly reducing the statistical probability of airborne cross-infection [2].

https://journals.adaresearch.or.id/icpex
INNOVATION AND CREATIVITY PRODUCT EXHIBITION
ISSN XXXX-XXXX (Media Online)
OPEN ACCESS
Vol. 1, No. 1, June 2026 | pp. 133-139

Conversely, as seen in Figure 5, the Patient Educational Interface prioritizes empathy and health literacy [15]. Upon accessing their secure portal, the patient is presented with the Grad-CAM heatmap overlay [3]. Instead of displaying raw probability percentages that could induce procedural anxiety, the interface utilizes automated narrative text generation to explain the visual findings in plain, non-technical language. By rendering the abstract AI decision into a transparent, localized visual format, the system fosters a collaborative healthcare experience.

Furthermore, to guarantee production-grade reliability, the system's operational health and data flow were monitored using an observability framework. The integration of distributed tracing allowed for the precise measurement of end-to-end system latency. Empirical testing demonstrated that the total round-trip time from the moment the CXR is uploaded by the operator, processed by the CNN, to the instant the real-time database updates the frontend dashboard averages under 2.4 seconds. This minimal latency confirms that the computational overhead introduced by the Grad-CAM visualization does not hinder the rapid triage requirements of busy primary clinics.

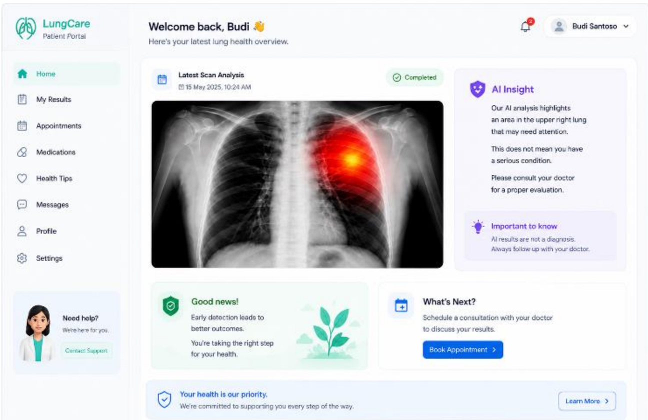

Figure 5. Patient Educational Interface with XAI Heatmap.

4. CONCLUSION

Tuberculosis screening in primary healthcare facilities often suffers from critical operational bottlenecks, leading to extended waiting times and a heightened risk of airborne cross-infection among patients in waiting areas [2]. To address this persistent challenge, this study successfully designed, developed, and evaluated a patient-centric, web-based triage information system that integrates Convolutional Neural Network (CNN) architectures with Gradient-weighted Class Activation Mapping (Grad-CAM) [3], [13]. Through a systematic, multi-stage iterative software engineering development approach, the system achieved a highly reliable diagnostic performance, evidenced by a recall rate of 96.15%. This specific metric ensures that high-risk cases are accurately identified for immediate clinical escalation and triage routing, directly mitigating transmission risks. Beyond mathematical accuracy, the primary breakthrough of this research lies in its successful implementation of Explainable Artificial Intelligence (XAI) to foster patient health literacy and emotional reassurance [6], [15]. By translating complex, black-box algorithmic decisions into intuitive, visual heatmaps, the system empowers patients to understand their screening results without experiencing the procedural anxiety typically associated with raw clinical probabilities [11], [19]. The patient educational interface delivers transparent, plain-language narratives alongside the visual data, effectively bridging the communication gap between advanced medical technology and layperson comprehension. The system strictly maintains its role as an educational support tool, explicitly stating that it does not replace a definitive medical diagnosis by a certified physician. Operationally, the integration of real-time database synchronization within the medical personnel dashboard proved highly effective in automating triage prioritization. By instantly escalating positive cases to the top of the queue with an average system latency of under 2.4 seconds, this automated sorting mechanism is positioned to significantly reduce physical crowding in clinical waiting areas. Future development iterations should focus on scaling the system architecture by integrating it with centralized, regional electronic health record systems, as well as expanding the deep learning model's diagnostic capabilities to detect multiple concurrent pulmonary pathologies [7], [14]. Ultimately, this innovation demonstrates that the future of medical artificial intelligence must fundamentally prioritize both clinical efficiency and compassionate, transparent patient communication.

REFERENCES

[1] World Health Organization, *Global tuberculosis report 2023*. Geneva: World Health Organization, 2023.

[2] P. Lakhani and B. Sundaram, "Deep learning at chest radiography: automated classification of pulmonary tuberculosis by using convolutional neural networks," *Radiology*, vol. 284, no. 2, pp. 574-582, 2017.

[3] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," *International Journal of Computer Vision*, vol. 128, no. 2, pp. 336-359, 2020.

[4] S. Jaeger et al., "Two public chest X-ray datasets for computer-aided screening of pulmonary diseases," *Quantitative Imaging in Medicine and Surgery*, vol. 4, no. 6, pp. 475-477, 2014.

https://journals.adaresearch.or.id/icpex
INNOVATION AND CREATIVITY PRODUCT EXHIBITION
ISSN XXXX-XXXX (Media Online)
OPEN ACCESS
Vol. 1, No. 1, June 2026 | pp. 133-139

[5] S. Rajaraman et al., "Deep learning models for tuberculosis screening using chest X-rays," *PLOS ONE*, vol. 15, no. 9, p. e0238491, 2020.
[6] J. Amann, A. Blasimme, E. Vayena, D. Frey, and V. I. Madai, "Explainability for artificial intelligence in healthcare: a multidisciplinary perspective," *BMC Medical Informatics and Decision Making*, vol. 20, no. 1, p. 310, 2020.
[7] G. Litjens et al., "A survey on deep learning in medical image analysis," *Medical Image Analysis*, vol. 42, pp. 60-88, 2017.
[8] Z. Z. Qin et al., "Using artificial intelligence to read chest radiographs for tuberculosis detection: A multi-site evaluation of the diagnostic accuracy of three deep learning systems," *Scientific Reports*, vol. 9, no. 1, p. 15000, 2019.
[9] E. J. Topol, "High-performance medicine: the convergence of human and artificial intelligence," *Nature Medicine*, vol. 25, no. 1, pp. 44-56, 2019.
[10] E. J. Hwang et al., "Development and validation of a deep learning-based automated detection algorithm for major thoracic diseases on chest radiographs," *JAMA Network Open*, vol. 2, no. 3, p. e191095, 2019.
[11] A. Holzinger et al., "Toward causability in AI: models and evaluation of XAI," *Communications of the ACM*, vol. 62, no. 12, pp. 93-100, 2019.
[12] M. T. Ribeiro, S. Singh, and C. Guestrin, "'Why should I trust you?': Explaining the predictions of any classifier," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 1135-1144.
[13] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436-444, 2015.
[14] A. Esteva et al., "A guide to deep learning in healthcare," *Nature Medicine*, vol. 25, no. 1, pp. 24-29, 2019.
[15] S. Tonekaboni, S. Joshi, M. D. McCradden, and A. Goldenberg, "What clinicians want: contextualizing explainable machine learning for clinical end use," *Machine Learning for Healthcare Conference*, 2019, pp. 359-380.
[16] C. J. Kelly, A. Karthikesalingam, M. Suleyman, G. Corrado, and D. King, "Key challenges for delivering clinical impact with artificial intelligence," *BMC Medicine*, vol. 17, no. 1, p. 195, 2019.
[17] P. Rajpurkar et al., "CheXNet: Radiologist-level pneumonia detection on chest X-rays with deep learning," *arXiv preprint arXiv:1711.05225*, 2017.
[18] D. S. Kermany et al., "Identifying medical diagnoses and treatable diseases by image-based deep learning," *Cell*, vol. 172, no. 5, pp. 1122-1131.e9, 2018.
[19] S. M. Lundberg and S. I. Lee, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems*, 2017, pp. 4765-4774.
[20] World Health Organization, *WHO consolidated guidelines on tuberculosis: Module 2: Screening - Systematic screening for tuberculosis disease*. Geneva: World Health Organization, 2021.