IJIRCCE

ISSN(O): 2320-9801

ISSN(P): 2320-9798

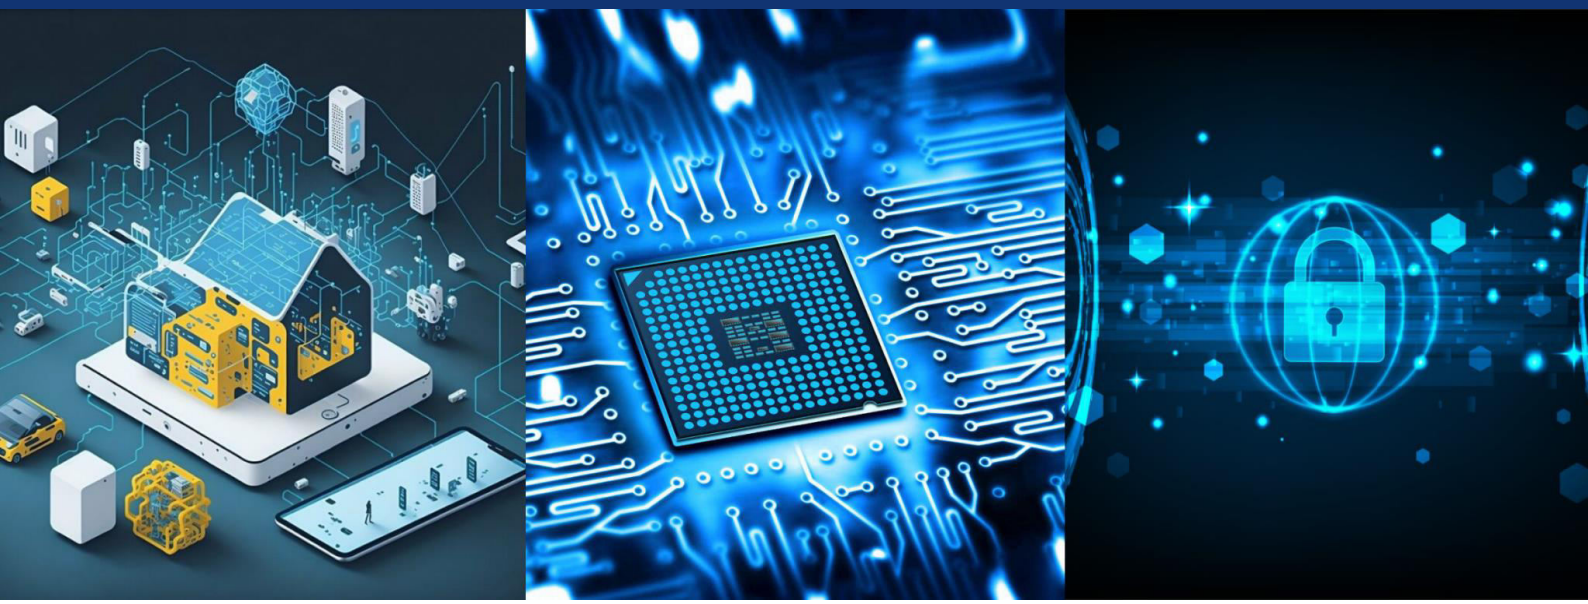

International Journal of Innovative Research in Computer and Communication Engineering

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

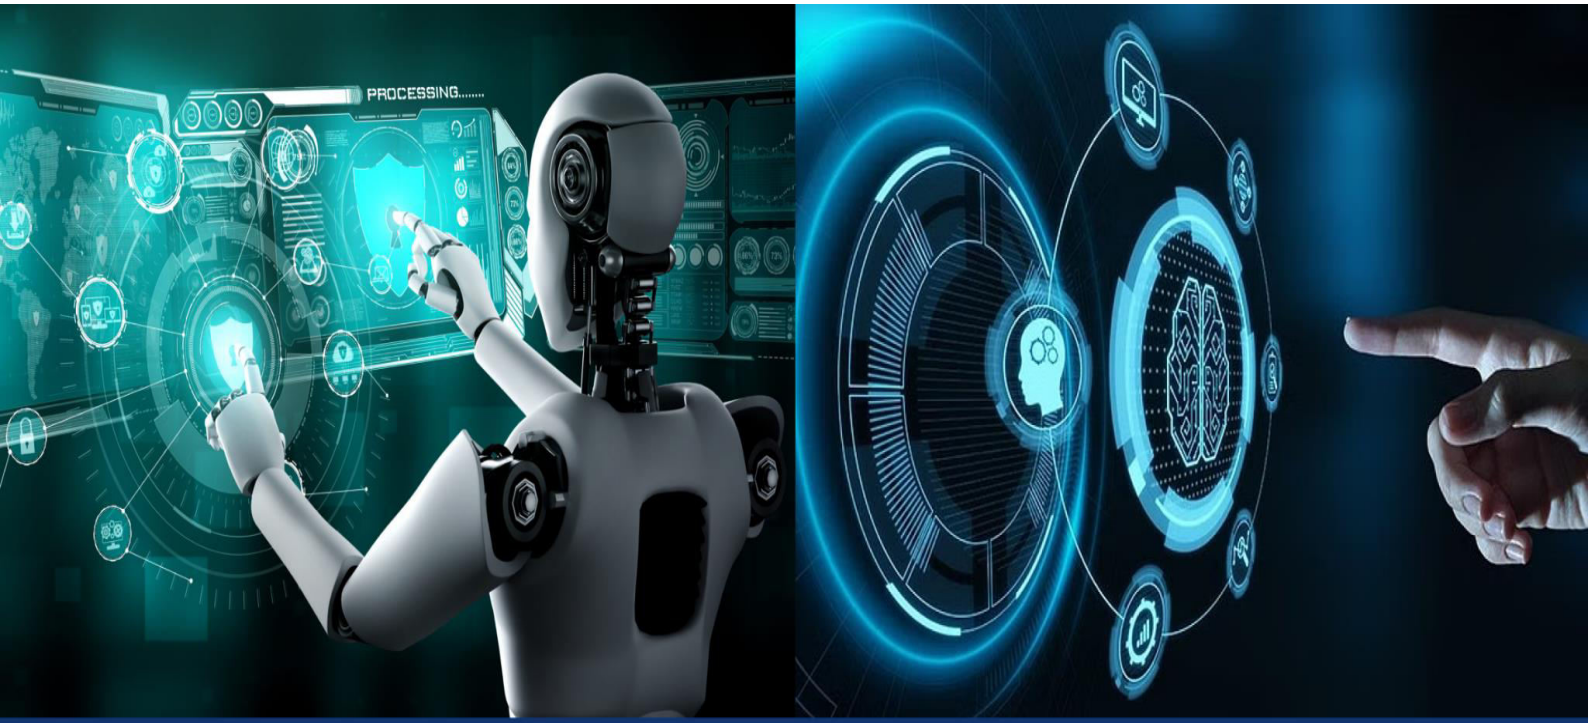

Impact Factor: 9.927

Volume 14, Issue 6, June 2026

© 2026 IJIRCCE | Volume 14, Issue 6, June 2026 DOI: 10.15680/IJIRCCE.2026.1406087

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

A Review on Deep Learning-Based Lung Disease Detection Using Chest X-Ray Images

Joel Jacob Varghese[1], Navin Cholangi[2], Manchit Choudhary[3], Manna Sara Bilu[4], Dr. Yogesh Golhar[5]

Department of Computer Engineering, St. Vincent Pallotti College of Engineering and Technology, Nagpur, India[1-4]

Assistant Professor, Department of Computer Engineering, St. Vincent Pallotti College of Engineering and Technology, Nagpur, India[5]

ABSTRACT: Lung diseases such as pneumonia and tuberculosis, along with occupational conditions like silicosis and pneumoconiosis, remain a major global health burden, particularly in regions with limited access to expert radiologists. Traditional chest X-ray interpretation is time-consuming, depends on specialist availability, and is prone to inter-observer inconsistency. This paper presents a narrative review of how convolutional neural networks (CNNs), strengthened by transfer learning, are being applied to automated multi-class lung disease detection from chest X-ray images. We survey the prominent public datasets, the established architectures (VGGNet, ResNet, DenseNet, and EfficientNet), and the growing role of explainable AI (XAI) through Grad-CAM and LIME. We give particular attention to occupational lung diseases, where data scarcity is the main obstacle to progress. Drawing the threads of the review together, we also outline a modular, web-based reference architecture for deployment in resource-constrained healthcare settings; this is a conceptual design rather than an implemented system. Finally, we discuss the ethical questions around automation, equity, and data privacy. Our analysis suggests that although benchmark performance has reached radiologist-level accuracy for common conditions, the reported numbers should be read with caution, and real-world deployment still demands careful attention to domain shift, regulatory compliance, and clinical workflow integration.

KEYWORDS: deep learning; convolutional neural networks; chest X-ray; lung disease detection; transfer learning; silicosis; occupational lung diseases; explainable AI; clinical decision support; web deployment

I. INTRODUCTION

A. Global Burden and Motivation

Lung diseases including pneumonia, tuberculosis, silicosis, and coal workers' pneumoconiosis continue to pose serious health risks, especially for workers exposed to dust, chemicals, and other hazardous environments [1], [11]. These conditions often develop slowly and go unnoticed until significant damage has already been done. The chest X-ray (CXR) remains the most widely available and cost-effective imaging tool for catching them, yet its clinical value is limited by how many specialists are available to read the films and by the variability between one reader and another.

In many rural and industrial parts of India, access to radiological expertise is severely constrained, with large populations served by relatively few trained specialists. Even where a specialist is on hand, plain radiography struggles with subtle disease: a recent meta-analysis reported that CXR screening for silicosis achieves only about 76% pooled sensitivity against high-resolution CT as the reference standard, with considerable disagreement between readers [1]. That gap is exactly what makes a case for AI-assisted triage tools that support, rather than replace, clinical judgment.

B. Scope and Objectives

This paper is a narrative (non-systematic) review of deep learning approaches to multi-class lung disease classification from chest X-ray images, framed so that the design discussion can later extend to occupational diseases. Our objectives are four. First, to review CNN and transfer learning techniques for sorting chest X-rays into normal, pneumonia, and tuberculosis categories. Second, to examine how such models can be wrapped in a web-based platform that returns real-time predictions with confidence scores. Third, to assess the explainable AI techniques that help clinicians interpret those predictions. And fourth, to sketch a conceptual reference architecture that can be extended to occupational disease classes such as silicosis once suitable datasets exist [2], [3].

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Recent advances in artificial intelligence, and CNNs in particular, have shown strong results in medical image analysis [6], [9]. The growing supply of public medical imaging datasets, together with cheaper computation, has accelerated this progress [14], [15]. Web-based platforms add a further benefit by enabling remote screening and monitoring in areas that are otherwise underserved.

II. DATASETS AND DATA FOUNDATIONS

A. General Lung Disease Datasets

A handful of large public datasets underpin most of the progress in chest X-ray AI. ChestX-ray8/14 (Wang et al., 2017) provides 108,948 frontal-view X-rays labelled with 14 findings extracted from radiology reports using NLP [14]. CheXpert (Irvin et al., 2019) built on this with 224,316 radiographs and added uncertainty labels to handle diagnostic ambiguity [15]. The COVID-19 Image Data Collection (Cohen et al., 2020) showed how quickly the field could mobilise a dataset during an emerging crisis [16].

Systems built on these resources usually draw their training data from such public sources and organise the images into target categories, most often normal, pneumonia, and tuberculosis. How carefully the labelling and category organisation are done at this stage has a direct effect on downstream accuracy and reliability.

B. The Occupational Disease Data Gap

One limitation stands out: publicly available, annotated datasets for occupational lung diseases such as silicosis and pneumoconiosis are almost nonexistent. Akhter et al. (2025) introduced Silicodata, reported as the first publicly annotated CXR benchmark for silicosis, collected from stone workers in Rajasthan, India, which is an important first step [2]. Even so, most occupational disease collections hold only a few hundred images, far too few to train a robust model on their own. Any system meant for this domain should therefore be built to absorb such datasets in later phases without needing to be redesigned.

III. RESEARCH GAPS IN AI-BASED LUNG DISEASE DETECTION

Despite real progress, several challenges remain open. We group them into five recurring themes that surface throughout the literature and that the rest of this paper returns to.

Limited coverage of occupational diseases. Most studies concentrate on common respiratory conditions, pneumonia, tuberculosis, and COVID-19, while occupational diseases such as silicosis and pneumoconiosis receive far less attention. The root cause is the absence of large, well-annotated public datasets for these diseases, which holds back both model development and evaluation [2], [11].

Lack of explainability. Many models reach high accuracy yet are hard to trust, because clinicians are reasonably wary of black-box predictions. XAI techniques such as Grad-CAM and LIME are promising, but their integration into practical clinical systems is still limited [5], [20].

Evaluation under controlled conditions. Most work is tested only on benchmark datasets in the lab. Real deployment brings image-quality variation, domain shift across hospitals, differences in imaging equipment, and diverse patient populations, so models that look excellent in research can quietly lose ground in the clinic [7].

Scarcity of integrated, deployable platforms. Few systems combine prediction, explainability, confidence estimation, and clinical decision support in one place. A great deal of work stops at the algorithm and never addresses deployment, user interaction, or accessibility in low-resource settings.

Privacy, regulation, and fairness. Patient privacy, regulatory compliance, fairness, and ethical adoption all remain barriers to large-scale use. Closing these gaps is what separates an experimental research tool from a dependable clinical decision-support system.

© 2026 IJIRCCE | Volume 14, Issue 6, June 2026 DOI: 10.15680/IJIRCCE.2026.1406087

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

IV. REVIEW OF DEEP LEARNING METHODOLOGIES

A. Convolutional Neural Network (CNN) Based Approaches

CNNs are the backbone of modern medical image analysis. Unlike older methods that depend on hand-crafted features, CNNs learn visual patterns, edges, textures, and abnormalities, directly from the chest X-ray. VGGNet, ResNet, and DenseNet each contributed a key idea: deep stacked filtering [19], residual connections that ease the vanishing-gradient problem, and dense connectivity that encourages feature reuse [17]. Together these have driven strong results in detecting pneumonia, tuberculosis, and COVID-19 [8], [13].

B. Transfer Learning Approaches

Labelled medical data is scarce and expensive, which makes transfer learning a natural fit. Rather than training from scratch, models pre-trained on large datasets such as ImageNet are fine-tuned for a specific medical task, as shown in Fig. 1. Architectures like VGG19, ResNet50, DenseNet121, and EfficientNetB0 deliver high accuracy with far less data and compute [7], [9]. In a field where annotation is the bottleneck, that is invaluable.

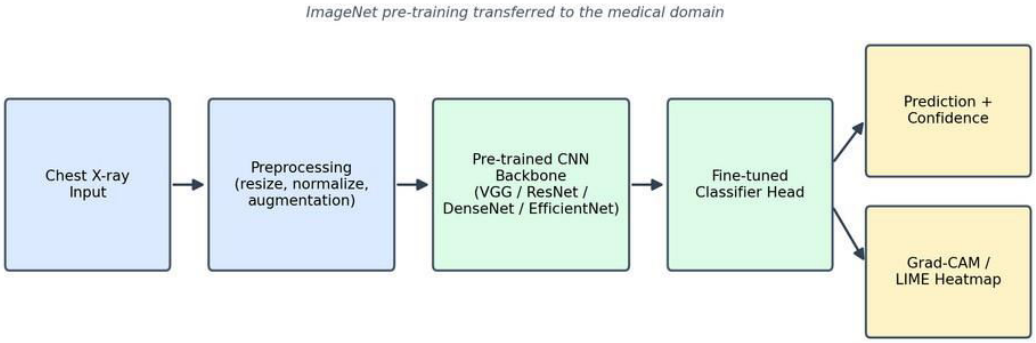

Fig. 1. Typical transfer learning pipeline for chest X-ray classification: an ImageNet pre-trained CNN backbone is fine-tuned on medical data, producing a class prediction with confidence alongside an XAI heatmap.

C. EfficientNet and Model Optimization

EfficientNet rethinks how models are scaled. Where earlier work improved accuracy by increasing depth, width, or input resolution one at a time, EfficientNet uses compound scaling to balance all three together, gaining accuracy without bloating the model [18]. EfficientNetB0 in particular strikes a practical balance between performance and efficiency, which makes it a good fit for clinical settings with limited hardware.

D. Occupational Lung Disease Detection

Work on occupational diseases like silicosis and pneumoconiosis trails behind work on pneumonia and TB, mainly because of the shortage of annotated data [11]. Even so, recent studies show that deep learning can pick up subtle radiographic patterns that a manual read might miss: noise-tolerant staging for pneumoconiosis [3], segmentation-driven pipelines for engineered-stone silicosis [4], and augmentation strategies for small coal workers' pneumoconiosis datasets [10]. Resources like Silicodata are starting to change the picture [2]. The priorities from here are clear: more data, better generalisation, and models that are as reliable for occupational diseases as they already are for common respiratory conditions.

V. COMPARATIVE ANALYSIS OF EXISTING STUDIES

Research in AI-assisted chest X-ray analysis has grown a great deal over the past few years. The dominant approach is transfer learning, fine-tuning ImageNet-pretrained models on medical datasets, which works well even with limited labelled data. Table I summarises the studies most directly relevant to this review, and Fig. 2 compares the classification accuracies reported by representative works.

© 2026 IJIRCCE | Volume 14, Issue 6, June 2026 DOI: 10.15680/IJIRCCE.2026.1406087

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

TABLE I. SUMMARY OF RELEVANT LITERATURE (ordered by year, most recent first)

Year	Author(s)	Contribution	Gap / Limitation
2025	Howlett et al. [1]	Systematic review and meta-analysis of CXR screening for silicosis; pooled sensitivity of ~76% against HRCT as the reference standard.	CXR unreliable as a standalone screen; no AI-based improvement evaluated.
2025	Akhter et al. [2]	Introduced Silicodata, reported as the first publicly annotated benchmark CXR dataset for silicosis, collected from stone workers in Rajasthan, India.	Baseline AI accuracy remains limited; domain adaptation is needed.
2025	Sun et al. [3]	COFINE: a coarse-to-fine, noise-tolerant method for pneumoconiosis staging under label noise; Stage-1 accuracy of ~82.4%.	Dataset noise and stage ambiguity persist; no end-to-end deployment.
2025	Priego-Torres et al. [4]	Automated engineered-stone silicosis screening and staging (rib-cage segmentation + CNN) reporting near-perfect screening performance (ROC AUC up to 1.0).	Reported AUC~1.0 likely reflects a small/single-source test set; staging between simple silicosis and PMF remains difficult; data- and compute-heavy.
2025	Veeramani et al. [5]	XAI-TRANS model combining transfer learning with Grad-CAM and LIME for multi-class lung disease detection; ~97% accuracy.	Not silicosis-focused; no deployment system provided.
2024	Sharma et al. [6]	U-Net + Xception framework for tuberculosis detection with ~99% accuracy and infected-region localisation.	TB-only; not generalised to occupational diseases.
2024	Gu and Lee [7]	Deep transfer learning using real-world image features improved medical image classification, with a pneumonia CXR case study.	Domain gap between natural and medical images remains.
2023	Alshmrani et al. [8]	Multi-class lung disease classification using a VGG19 + CNN architecture; ~96.5% accuracy across five disease classes.	No silicosis focus; lacks explainability.
2023	Priyadarsini et al. [9]	Compared multiple deep learning architectures for lung disease detection, reporting up to ~99% accuracy.	No deployment; limited interpretability.
2022	Dong et al. [10]	Data augmentation with ShuffleNet and ECA-Net for coal workers' pneumoconiosis classification; ~98% accuracy.	Small single-source dataset; limited scope.
2022	Devnath et al. [11]	Systematic literature review of CAD systems for coal workers' pneumoconiosis covering handcrafted features, classical ML, and deep learning.	Highlights the lack of public datasets and practical, deployable systems.
2020	Cohen et al. [16]	COVID-19 image data collection demonstrating rapid mobilisation of CXR datasets during an emerging crisis.	Small initial dataset; emergency-focused rather than general-purpose.
2019	Christe et al. [12]	CAD system for pulmonary fibrosis detection and classification using deep learning on CT images.	Limited to CT; not scalable to CXR-based screening.
2019	Irvin et al. [15]	CheXpert: large chest radiograph dataset (224,316 images) with uncertainty labels for diagnostic ambiguity.	Uncertainty handling adds training complexity; no occupational diseases.
2019	Tan and Le [18]	EfficientNet: compound scaling of depth, width, and resolution, reaching state-of-the-art accuracy with	ImageNet-centric pre-training; computationally intensive at the

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

		fewer parameters.	largest scales.
2017	Rajpurkar et al. [13]	CheXNet: 121-layer DenseNet achieving radiologist-level pneumonia detection on chest X-rays.	Single-disease focus; no multi-class or occupational coverage.
2017	Wang et al. [14]	ChestX-ray8/14: hospital-scale CXR database (108,948 images) with 14 pathology labels extracted via NLP.	Weakly supervised labels; no occupational disease categories.
2017	Huang et al. [17]	DenseNet: densely connected CNNs with feature reuse and improved gradient flow, enabling deeper yet efficient networks.	Higher memory usage; not tailored to medical imaging.
2017	Selvaraju et al. [20]	Grad-CAM: gradient-based visual explanations highlighting the discriminative image regions behind any CNN prediction.	Post-hoc method; not intrinsically interpretable during training.
2015	Simonyan and Zisserman [19]	VGGNet: very deep CNNs (16-19 layers) built from 3x3 filters; a widely used backbone for medical transfer learning.	Large model size; limited efficiency on edge devices.

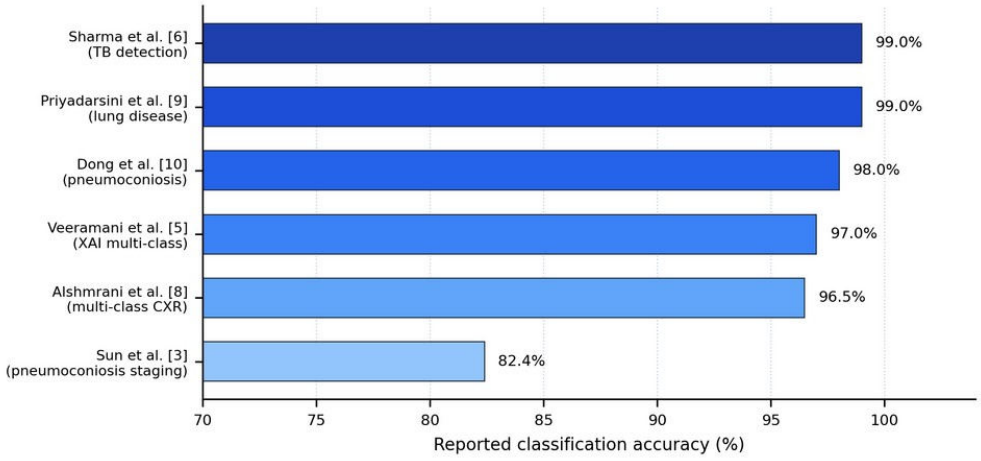

Fig. 2. Reported classification accuracies of representative deep learning studies for lung disease detection on chest X-rays. Direct comparison should be made with caution, as datasets, disease classes, and evaluation protocols differ across studies.

Reported accuracies should be read carefully. The headline figures in Fig. 2 come from different datasets, disease classes, and evaluation protocols, so they are not directly comparable. Values approaching 99-100%, and near-perfect screening metrics such as a reported ROC AUC close to 1.0 [4], are most plausibly explained by small or single-source test sets, favourable class balance, or possible leakage between training and test partitions, rather than by performance that would carry over unchanged to routine clinical use. Such results show strong in-distribution performance but need external validation before any clinical claim.

Across the literature, CNN-based models and transfer learning have steadily improved detection accuracy, and the recent addition of XAI tools such as Grad-CAM is a welcome move toward clinical trust. The structural gaps from Section III nonetheless persist: most models are tested on controlled datasets and never deployed; simple tools for low-resource clinics are scarce; and occupational diseases such as silicosis remain largely overlooked because of data limitations. A deployable web platform with extensible disease coverage speaks directly to the over-dependence on specialist availability.

© 2026 IJIRCCE | Volume 14, Issue 6, June 2026 DOI: 10.15680/IJIRCCE.2026.1406087

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

VI. EXPLAINABILITY AND CLINICAL TRUST

High accuracy on its own is not enough in healthcare; clinicians need to understand why a model reached a decision. Grad-CAM helps by producing heatmaps that show which image regions drove a prediction [20]. LIME complements it by highlighting the features that mattered most for an individual case. Together they make AI outputs more transparent, and recent work shows they can be folded into multi-class lung disease pipelines at little cost to accuracy [5].

That said, saliency maps are a starting point, not a destination. A heatmap confirms that the model looked at plausible lung regions, but it does not explain the reasoning behind a diagnosis in clinical terms. For an AI prediction to genuinely support a radiologist's workflow, the explanation needs to speak the language clinicians use, opacity patterns, distribution, and severity, and it should sit alongside a calibrated confidence estimate so that low-certainty cases are flagged for expert review.

VII. CHALLENGES AND REAL-WORLD DEPLOYMENT CONSIDERATIONS

The structural gaps of Section III become sharpest at the point of deployment. Here we look at how each one turns into a concrete barrier, and sketch a reference architecture that helps mitigate them.

Data scarcity is the most immediate bottleneck. Models need large volumes of accurately labelled images, yet medical data is costly to annotate, privacy-restricted, and hard to access, especially for occupational diseases like silicosis where public datasets are nearly nonexistent [2], [11]. Domain shift compounds the problem: a model trained on one hospital's scans can degrade quietly when moved elsewhere, because equipment, protocols, and patient mix all vary [7], so external validation is essential before any rollout. Interpretability is a sticking point too; Grad-CAM and LIME indicate which regions drive a prediction but do not fully open the black box [20].

Regulation and bias add more complexity. Medical AI must navigate patient-safety laws, data-privacy frameworks, and ethical guidelines, and datasets that under-represent certain groups can bake in bias, producing uneven performance across patient populations, a serious equity concern. Computation matters as well: even efficient architectures like EfficientNet can strain limited hardware [18]. Lightweight models and cloud deployment help, but connectivity and maintenance costs are real obstacles in low-resource settings.

Fig. 3 outlines a modular, web-based reference architecture suited to such environments: a browser-accessible frontend, a lightweight inference service that returns predictions with confidence scores and explanation heatmaps, and a human-in-the-loop review step, with occupational disease classes added later as a module once data allows. We present this as a conceptual synthesis of the requirements drawn from the reviewed literature, not as an implemented or evaluated system.

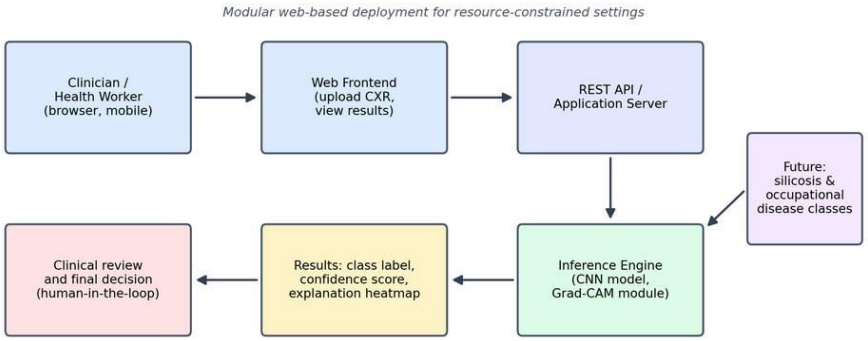

Fig. 3. Proposed modular web-based reference architecture for resource-constrained healthcare settings. The design keeps the clinician in the loop and allows occupational disease classes to be added without architectural redesign.

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

None of this changes a basic point: AI supports clinicians, it does not replace them. A diagnosis draws on patient history, physical examination, laboratory results, and years of judgment, things no image classifier can replicate. The aim is a system that makes capable clinicians more effective, not one that works without them.

VIII. FUTURE RESEARCH DIRECTIONS

The agenda for AI-assisted lung disease detection is promising and still developing. A few directions stand out.
Broader disease coverage. Most models today handle pneumonia, TB, and COVID-19. Extending to silicosis, pneumoconiosis, lung cancer, COPD, pulmonary fibrosis, and pleural effusion would make these systems considerably more useful, and that depends heavily on richer, more diverse datasets becoming available [2].

Multimodal diagnosis. Most systems still look at the chest X-ray in isolation. Combining it with CT scans, laboratory results, electronic health records, and patient history would give models a fuller picture, closer to how a clinician actually reasons [12].

Better explainability. Heatmaps are a start, but future XAI needs to communicate in terms clinicians use, offering explanations that are clinically meaningful rather than only visually intuitive. The more interpretable the system, the more readily it is adopted.

Federated learning. Instead of pooling sensitive data in one place, models can be trained across institutions without the data ever leaving each site, giving better generalisation with less privacy risk.

Foundation models. Large models pre-trained on extensive medical data could serve as powerful starting points for many diagnostic tasks, cutting the time and data needed to build specialised systems.

Conversational AI. Tools that explain results to patients, guide symptom collection, help clinicians with report generation, and improve access in remote areas add another useful dimension.

Real-world validation. Most importantly, validation has to keep pace with ambition. Multi-hospital prospective trials, integration with PACS and EHR systems, and honest evaluation across diverse populations are what separate a promising prototype from a tool clinicians can rely on.

IX. CONCLUSION

This review traced the development of deep learning in chest X-ray analysis, from CNN architectures and transfer learning through to explainable AI and a web-based deployment design. The overall picture is encouraging.

Modern architectures, VGGNet, ResNet, DenseNet, and EfficientNet, have pushed automated classification to the point where, for specific conditions, reported accuracy matches or exceeds expert radiologists [13]. Transfer learning made this possible even where labelled data is thin, and tools like Grad-CAM and LIME have begun to make these models legible to the clinicians who must trust them [5], [20].

The gaps, though, are real. Occupational diseases like silicosis remain data-starved [2], [11]; models trained in one institution do not always travel to another; regulatory hurdles are substantial; and too many studies stop at benchmark accuracy without asking what clinical deployment actually requires. The reported headline numbers, too, deserve a careful eye, since evaluation protocols differ widely from one study to the next.

The more interesting frontier lies ahead, multimodal imaging, federated learning, foundation models, conversational AI, and tighter integration with hospital systems such as PACS and EHRs. These are active directions that could, together, move AI from a laboratory curiosity to a genuine clinical tool.

The central message is simple: deep learning has earned a place in the diagnostic toolkit. What it most needs now is better datasets, more rigorous and honest deployment frameworks, and sustained collaboration between AI researchers and the clinicians who will use these systems. That combination is what turns promising research into technology that helps real patients.

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

REFERENCES

1. P. Howlett, A. Durairaj, M. Lesosky, and J. Feary, "The diagnostic accuracy of chest X-ray screening for silicosis: A systematic review, meta-analysis and modelling study," Occupational and Environmental Medicine, vol. 82, no. 12, pp. 614–620, 2025.
2. Y. Akhter, R. Ranjan, M. Vatsa, R. Singh, S. Chaudhury, A. Agrawal, S. Aggarwal, A. Kalyanpur, and A. Menon, "Silicodata: An annotated benchmark CXR dataset for silicosis detection," Scientific Data, vol. 12, art. no. 1559, 2025, doi: 10.1038/s41597-025-05595-4.
3. W. Sun, D. Wu, J. Shen, et al., "A deep learning method for pneumoconiosis staging on chest X-ray under label noise," IEEE Journal of Biomedical and Health Informatics, 2025, doi: 10.1109/JBHI.2025 (early access).
4. B. Priego-Torres, D. Sanchez-Morillo, E. Khalili, M. Á. Conde-Sánchez, A. García-Gámez, and A. León-Jiménez, "Automated engineered-stone silicosis screening and staging using deep learning with X-rays," Computers in Biology and Medicine, vol. 191, art. no. 110153, 2025, doi: 10.1016/j.compbiomed.2025.110153.
5. N. Veeramani et al., "NextGen lung disease diagnosis with explainable artificial intelligence," Scientific Reports, vol. 15, art. no. 33052, 2025, doi: 10.1038/s41598-025-07603-4.
6. V. Sharma, Nillmani, S. K. Gupta, and K. K. Shukla, "Deep learning models for tuberculosis detection and infected region visualization in chest X-ray images," Intelligent Medicine, vol. 4, no. 2, pp. 104–113, 2024.
7. C. Gu and M. Lee, "Deep transfer learning using real-world image features for medical image classification, with a case study on pneumonia X-ray images," Bioengineering, vol. 11, no. 4, art. no. 406, 2024.
8. G. M. M. Alshmrani, Q. Ni, R. Jiang, H. Pervaiz, and N. M. Elshennawy, "A deep learning architecture for multi-class lung diseases classification using chest X-ray (CXR) images," Alexandria Engineering Journal, vol. 64, pp. 923–935, 2023.
9. M. J. P. Priyadarsini, K. Kotecha, G. K. Rajini, et al., "Lung diseases detection using various deep learning algorithms," Journal of Healthcare Engineering, vol. 2023, art. no. 3563696, 2023.
10. H. Dong, B. Zhu, X. Zhang, and X. Kong, "Use data augmentation for a deep learning classification model with chest X-ray clinical imaging featuring coal workers' pneumoconiosis," BMC Pulmonary Medicine, vol. 22, art. no. 271, 2022.
11. L. Devnath, P. Summons, S. Luo, D. Wang, K. Shaukat, I. A. Hameed, and H. Aljuaid, "Computer-aided diagnosis of coal workers' pneumoconiosis in chest X-ray radiographs using machine learning: A systematic literature review," International Journal of Environmental Research and Public Health, vol. 19, no. 11, art. no. 6439, 2022.
12. A. Christe, A. A. Peters, D. Drakopoulos, et al., "Computer-aided diagnosis of pulmonary fibrosis using deep learning and CT images," Investigative Radiology, vol. 54, no. 10, pp. 627–632, 2019.
13. P. Rajpurkar, J. Irvin, K. Zhu, et al., "CheXNet: Radiologist-level pneumonia detection on chest X-rays with deep learning," arXiv preprint arXiv:1711.05225, 2017.
14. X. Wang, Y. Peng, L. Lu, Z. Lu, M. Bagheri, and R. M. Summers, "ChestX-ray8: Hospital-scale chest X-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases," in Proc. IEEE CVPR, 2017, pp. 2097–2106.
15. J. Irvin, P. Rajpurkar, M. Ko, et al., "CheXpert: A large chest radiograph dataset with uncertainty labels and expert comparison," in Proc. AAAI Conf. Artificial Intelligence, vol. 33, 2019, pp. 590–597.
16. J. P. Cohen, P. Morrison, and L. Dao, "COVID-19 image data collection," arXiv preprint arXiv:2003.11597, 2020.
17. G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in Proc. IEEE CVPR, 2017, pp. 4700–4708.
18. M. Tan and Q. V. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in Proc. 36th Int. Conf. Machine Learning (ICML), 2019, pp. 6105–6114.
19. K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in Proc. Int. Conf. Learning Representations (ICLR), 2015.
20. 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," in Proc. IEEE ICCV, 2017, pp. 618–626.

INTERNATIONAL STANDARD SERIAL NUMBER INDIA

INNO SPACE SJIF Scientific Journal Impact Factor

doi crossref

निस्केयर NISCAIR

INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN COMPUTER & COMMUNICATION ENGINEERING

9940 572 462 6381 907 438 ijircce@gmail.com

www.ijircce.com

Scan to save the contact details