Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Explainable AI-Based Pulmonary Nodule Classification from Computed Tomography Images with an End-to-End Clinical Decision Support Interface

Sujal Das

sujaldas2205@gmail.com

Datta Meghe Institute of Higher Education and Research (DMIHER)

Shital Hajare

Datta Meghe Institute of Higher Education and Research (DMIHER)

Asad ul Hakeem

Datta Meghe Institute of Higher Education and Research (DMIHER)

Akshat Nimje

Datta Meghe Institute of Higher Education and Research (DMIHER)

Article

Keywords: Pulmonary nodule detection, computer-aided detection, deep learning, ResNet-50, transfer learning, Grad-CAM, explainable artificial intelligence, lung cancer screening, clinical decision support

Posted Date: November 2nd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10896988/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Explainable AI-Based Pulmonary Nodule Classification from Computed Tomography Images with an End-to-End Clinical Decision Support Interface

Sujal Das[1,*], Shital Hajare[2], Asad ul Hakeem[1], and Akshat Nimje[3]

[1]*Department of Artificial Intelligence and Data Science, Datta Meghe Institute of Higher Education and Research (DMIHER), Wardha, Maharashtra 442001, India*

[2]*Faculty of Engineering and Technology, Datta Meghe Institute of Higher Education and Research (DMIHER), Wardha, Maharashtra 442001, India*

[3]*Department of Computer Science and Engineering, Datta Meghe Institute of Higher Education and Research (DMIHER), Wardha, Maharashtra 442001, India*

*Corresponding author: Sujal Das, e-mail: sujaldas2205@gmail.com

Abstract. Early and accurate detection of pulmonary nodules in computed tomography (CT) imaging is critical for reducing lung cancer mortality. Although deep learning-based computer-aided detection (CAD) systems have demonstrated promising diagnostic performance, their widespread clinical adoption has been impeded by two persistent limitations: the absence of visual explainability and the lack of deployable end-to-end interfaces that integrate prediction, explanation, and clinical workflow. This paper presents an explainable, end-to-end CAD system for pulmonary nodule classification that addresses both shortcomings. The proposed system employs a transfer-learned ResNet-50 architecture, fine-tuned on 2,372 balanced CT patches extracted from the publicly available LUNA16 dataset. Gradient-weighted Class Activation Mapping (Grad-CAM) is integrated into the inference pipeline to generate per-prediction heatmaps that visually localize the discriminative regions driving each classification decision. The complete system is deployed as a production-grade web application featuring a FastAPI backend and a React frontend, complete with authentication and database persistence. On a held-out test set of 356 samples, the proposed system achieves an accuracy of 90.45%, precision of 90.91%, recall of 89.89%, F1 score of 90.40%, and an AUC-ROC of 0.9615. A confidence-based abstention mechanism returns an "Inconclusive" verdict for predictions below a 60% probability threshold, reducing the risk of unreliable automated decisions. Rather than claiming state-of-the-art raw accuracy, this work's principal contribution is demonstrating that competitive classification performance, real-time visual explainability, and a fully deployed, browser-accessible clinical interface can be integrated into a single reproducible, open-source system, contributing a practical baseline for future hospital validation studies.

Keywords: Pulmonary nodule detection, computer-aided detection, deep learning, ResNet-50, transfer learning, Grad-CAM, explainable artificial intelligence, lung cancer screening, clinical decision support.

1. INTRODUCTION

Lung cancer is the most prevalent cause of cancer-related mortality globally, accounting for approximately 1.8 million deaths in 2020 alone [1]. The five-year survival rate improves dramatically with stage at detection: from roughly 10% for Stage IV disease to over 60% for Stage I, underscoring the life-saving potential of early screening programs [2]. Pulmonary nodules—discrete, roughly spherical opacities measuring up to 30 mm in diameter—are the primary radiological finding of interest in low-dose CT screening protocols, yet their manual identification remains one of the most cognitively demanding tasks in diagnostic radiology. A single CT thorax study produces 200–400 axial slices, each of which must be scrutinised for nodules that may be as small as 3–4 mm and whose appearance varies considerably across size, density, and anatomical context [3]. The resulting workload, combined with known inter- and

intra-reader variability, creates both diagnostic inefficiency and a measurable false-negative rate that motivates the development of reliable automated screening tools.

Deep learning has emerged as the dominant paradigm for automated pulmonary nodule detection and classification over the past decade, with convolutional neural networks (CNNs) consistently matching or exceeding radiologist-level performance on benchmark datasets [4, 5]. However, a substantial gap persists between research demonstrations and clinical deployment. Two impediments are particularly salient. First, most published systems operate as opaque classifiers, offering clinicians no mechanistic account of how a prediction was reached. In safety-critical medical contexts, this opacity erodes clinician trust and runs counter to the transparency requirements increasingly mandated by emerging regulatory frameworks for medical AI. Second, the literature is dominated by offline experimental evaluations; relatively few works demonstrate fully integrated, deployable systems in which a radiologist can upload a CT patch and receive a classification alongside a visual explanation within a web-based interface.

This paper addresses both gaps through four concrete contributions:

- A transfer-learned ResNet-50 classifier fine-tuned on the LUNA16 benchmark that achieves an AUC-ROC of 0.9615 and an accuracy of 90.45% on a balanced 356-sample test set.
- Seamless Grad-CAM integration into the inference pipeline, producing per-prediction activation heatmaps overlaid on the original CT patch and delivered to the end user in real time.
- A confidence-based abstention mechanism that returns an "Inconclusive" verdict when the model's maximum class probability falls below 60%, mitigating the risk of overconfident automated decisions.
- A fully deployed, browser-accessible clinical decision support interface—comprising a FastAPI backend and React frontend—that integrates classification, explanation, and basic clinical metadata in a single, authenticated web session.

The remainder of this paper is organised as follows. Section 2 reviews the relevant prior art. Section 3 describes the dataset and preprocessing pipeline. Section 4 details the model architecture, training procedure, and explainability integration. Section 5 describes the system deployment. Section 6 presents quantitative and qualitative experimental results. Section 7 situates the findings in the broader literature, acknowledges limitations, concludes the paper, and outlines future directions.

2. RELATED WORK

2.1 Deep Learning for Pulmonary Nodule Detection and Classification

The application of CNNs to pulmonary nodule analysis was substantially catalysed by the publication of the Lung Image Database Consortium and Image Database Resource Initiative (LIDC-IDRI) [6] and the subsequent LUNA16 challenge [7]. Shen et al. [8] demonstrated that multi-scale CNNs trained on nodule patches from the LIDC-IDRI database could achieve classification accuracy of 86.84%, establishing the patch-based paradigm that the present work

follows. Dou et al. [9] extended this to three-dimensional representations, using multilevel contextual 3D CNNs to exploit volumetric context for false-positive reduction on the LUNA16 benchmark, reporting a competition performance metric (average sensitivity across seven predefined false-positive rates per scan) of 0.827. Ardila et al. [4] pushed further, training an end-to-end model directly on the raw 3D volume and demonstrating non-inferior performance to radiologists on a held-out clinical dataset, a landmark result for the field.

More recent work has focused on architectural refinements. UrRehman et al. [5] proposed a custom CNN with a dual (channel and spatial) attention mechanism for lung nodule detection on the LUNA16 dataset, reporting accuracy and AUC figures above 95% and 0.98 respectively as cited in subsequent literature—among the strongest results reported for single-model 2D/3D patch-based classifiers on this benchmark, though direct comparison is complicated by differences in evaluation protocol. Xie et al. [10] proposed a multi-view knowledge-based collaborative (MV-KBC) deep model that fine-tunes multiple ResNet-50 sub-networks on nine fixed views of each 3D nodule, reporting an accuracy of 91.60% and an AUC of 0.957 on LIDC-IDRI. Architectural efficiency has also been explored for deployment under compute constraints, for example via depthwise-separable convolutional backbones such as EfficientNet [11], an architecture choice relevant to deployment on commodity hardware such as the GPU used in the present study.

2.2 Explainability in Medical Imaging AI

The demand for interpretable AI in healthcare has produced a body of work on post-hoc explanation methods for image classifiers. Class Activation Mapping (CAM) [12] and its gradient-based extension Grad-CAM [13] have become the de facto standard for CNN attribution in medical imaging, owing to their architectural generality and computational tractability. Grad-CAM computes the gradient of the class score with respect to the final convolutional feature maps, producing a spatial importance map that highlights regions the model attends to when making its decision. In the pulmonary nodule domain, Grad-CAM overlays have been reported to correlate with radiologist-annotated regions of interest, lending them clinical face validity [3].

2.3 Clinical Decision Support Systems for Lung Cancer Screening

Despite the richness of the algorithmic literature, fully deployed CAD interfaces for pulmonary nodule classification remain comparatively rare in the peer-reviewed literature. Xie et al. [10] and Liao et al. [14] describe classification pipelines with strong offline performance, but neither work reports a browser-accessible deployment with real-time explanation. Commercially, FDA-cleared products such as Riverain Technologies' ClearRead CT [15] and Optellum's Virtual Nodule Clinic [16] demonstrate that AI-assisted nodule workflows are clinically viable, but these are proprietary, closed-source systems whose validation cohorts are drawn predominantly from Western clinical populations. The absence of open-source, deployable, and explainable systems validated or designed with Indian hospital settings in mind represents a specific gap that the present work begins to address.

3. DATASET AND PREPROCESSING

3.1 LUNA16 Dataset

This study utilises the LUNA16 (Lung Nodule Analysis 2016) dataset, which is derived from the LIDC-IDRI database [6] and constitutes one of the largest publicly available, expert-annotated CT lung nodule repositories. LUNA16 provides 888 thoracic CT scans in MetaImage (MHD/RAW) format, excluding LIDC-IDRI scans with slice thickness greater than 2.5 mm, alongside a curated list of 1,186 nodule annotations that reached consensus among at least three of four radiologist readers [7]. For the false-positive reduction track, LUNA16 additionally supplies a large set of candidate locations (over 500,000) generated by pooling the outputs of several independent automated nodule-candidate-detection algorithms; the vast majority of these candidates are non-nodule structures (vessels, noise, and other artifacts) and serve as the negative class for classifier training. For this study, a subset of these candidate locations was sampled to construct a class-balanced training corpus.

3.2 Patch Extraction and Labelling

For each annotated nodule centroid, a 2D axial patch of 64×64 pixels was extracted from the CT slice nearest to the centroid's z-coordinate. The patch was centred on the annotated coordinates after projection into voxel space following isotropic resampling (Section 3.3). Consensus nodule annotations were labelled as *positive* ($y = 1$) and sampled non-nodule candidate locations were labelled as *negative* ($y = 0$). This procedure yielded 1,186 positive and 1,186 negative patches, for a total balanced dataset of **2,372 samples**. The dataset was partitioned into training (70%), validation (15%), and test (15%) sets using stratified random splitting with a fixed seed of 42 to ensure class balance across all splits and full reproducibility.

Terminology note. Consistent with LUNA16 convention, the positive/negative distinction learned by the classifier is nodule vs. non-nodule (candidate) rather than a radiologically confirmed benign/malignant distinction; the latter would require pathology-confirmed outcome labels that are not part of the LUNA16 release. This scope should be stated explicitly wherever the classification task is described, and the terms "positive/negative" or "nodule/non-nodule" are preferred to "malignant/benign" unless the deployed system is later validated against biopsy-confirmed outcomes.

3.3 Preprocessing Pipeline

Raw CT data in MHD/RAW format was loaded using SimpleITK. The preprocessing pipeline comprised four sequential stages:

HU windowing. CT voxel intensities, expressed in Hounsfield Units (HU), were clipped to the pulmonary window $[-1000, +400]$ HU. This range encompasses air, lung parenchyma, and soft tissue while excluding irrelevant dense structures.

Isotropic resampling. Each CT volume was resampled to an isotropic voxel spacing of $1 \times 1 \times 1$ mm using third-order B-spline interpolation. This normalises the physical scale of nodules across scans acquired with varying slice thicknesses and in-plane resolutions.

Patch extraction. Following resampling, the 2D axial patch of 64×64 pixels was extracted centred on the annotation centroid. Patches extending beyond the volume boundary were zero-padded.

Intensity normalisation. Windowed HU values were linearly normalised to the $[0, 1]$ range and stored as single-channel (grayscale) float32 arrays.

4. METHODOLOGY

4.1 Model Architecture

The classification backbone is a ResNet-50 [17] pretrained on ImageNet-1k. ResNet-50 was selected over deeper variants based on its favourable accuracy-to-parameter ratio and its established utility in medical imaging transfer learning contexts.

Two modifications were made to the stock ResNet-50:

1. The first convolutional layer (`conv1`), originally accepting three-channel (RGB) input, was replaced with a single-channel convolutional layer with identical kernel size (7×7), stride (2), and padding (3). Pretrained weights for the three-channel layer were averaged across the channel dimension and used to initialise the new layer.
2. The final fully connected layer was replaced to produce a two-class logit vector corresponding to the *nodule* and *non-nodule* classes.

4.2 Training Configuration

All experiments were conducted on an NVIDIA GeForce RTX 3050 Ti using PyTorch 2.7.1 with CUDA 11.8. Training hyperparameters are summarised in Table 1. Mixed-precision training was employed to reduce memory consumption and accelerate gradient computation.

Table 1. Training Hyperparameters

Hyperparameter	Value
Optimiser	Adam
Learning rate	1×10^{-4}
Batch size	16
Maximum epochs	30
Early stopping patience	8 epochs
LR scheduler	ReduceLROnPlateau
Scheduler factor	0.5
Scheduler patience	3 epochs
Mixed precision	FP16
Random seed	42

Online data augmentation was applied exclusively to the training split: random horizontal

flip, random vertical flip, random rotation ($\pm 15^{\circ}$), and random affine transformation. Training terminated at epoch 29 via early stopping.

4.3 Grad-CAM Explainability

Gradient-weighted Class Activation Mapping (Grad-CAM) [13] was integrated into the inference pipeline to produce spatial attribution maps for each prediction. Formally, the Grad-CAM saliency map for class c is defined as:

$$L^{c}_{\text{Grad-CAM}} = \text{ReLU}\left(\sum_{k} \alpha^{c}_{k} A^{k}\right), \tag{1}$$

where A^{k} denotes the k-th feature map in the target convolutional layer, and α^{c}_{k} is the neuron importance weight for class c with respect to feature map k, computed as the global average-pooled gradient:

$$\alpha^{c}_{k} = \frac{1}{Z} \sum_{i} \sum_{j} \frac{\partial y^{c}}{\partial A^{k}_{ij}}. \tag{2}$$

Here y^{c} is the pre-softmax class score, Z is the number of spatial locations in the feature map, and i, j index spatial positions. The resulting low-resolution map is bilinearly upsampled to the input resolution (64×64) and overlaid on the original CT patch. The Grad-CAM target layer in this implementation is the final convolutional block of ResNet-50 (`layer4`).

4.4 Confidence-Based Abstention

Overconfident probabilistic predictions are a known failure mode of deep classifiers. To mitigate this risk in a clinical context, a threshold-based abstention mechanism was implemented: if $\max_{c} \hat{p}(y = c \mid x) < 0.60$, the system returns an "Inconclusive" verdict rather than a hard class label.

5. SYSTEM DEPLOYMENT

A production-grade clinical decision support interface was developed to make the trained model and its Grad-CAM explanations accessible through a standard web browser. The system architecture comprises the following decoupled components:

Inference backend. A RESTful API was implemented in FastAPI and served via Uvicorn. The backend is hosted as a public Space on Hugging Face Spaces.

Frontend application. A single-page application (SPA) was developed using React 18 and Vite, styled with Tailwind CSS. The frontend is deployed on Vercel with automatic CI/CD from the GitHub repository.

Authentication & Database. Clerk provides OAuth 2.0 authentication. MongoDB Atlas stores user session metadata and prediction history.

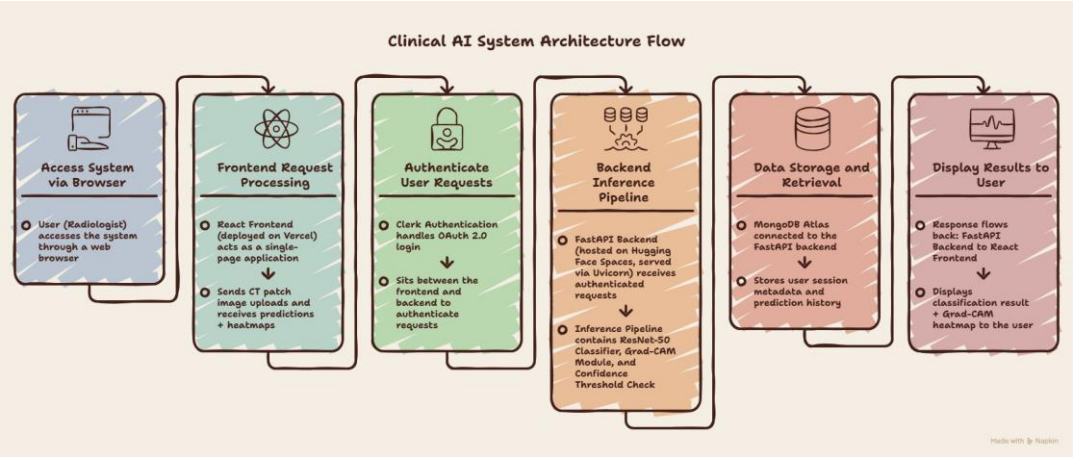


Figure 1. System architecture of the proposed clinical decision support interface. A React frontend communicates with a FastAPI inference backend.

The complete system architecture is illustrated in Figure 1. End-to-end latency from image upload to heatmap display was consistently below 3 seconds on the deployed infrastructure.

6. EXPERIMENTS AND RESULTS

6.1 Classification Performance

Table 2 presents the evaluation results on the 356-sample test set. The proposed system achieves an accuracy of 90.45% and an AUC-ROC of 0.9615.

Table 2. Classification Performance on the LUNA16 Test Set (N = 356)

Metric	Score
Accuracy	90.45%
Precision	90.91%
Recall (Sensitivity)	89.89%
F1 Score	90.40%
AUC-ROC	0.9615

The confusion matrix (Table 3) demonstrates a near-symmetric error distribution, indicating the model is neither systematically biased toward over-predicting the positive class nor excessively conservative.

Table 3. Confusion Matrix on the Test Set

	Predicted Negative	Predicted Positive
Actual Negative	162 (TN)	16 (FP)
Actual Positive	18 (FN)	160 (TP)

6.2 Training Dynamics and Explainability

Training and validation loss curves (Figure 2) indicate that early stopping prevented overfitting. Qualitatively, Grad-CAM heatmaps (Figure 3) show the model attending to the nodule region and surrounding parenchymal texture when making correct predictions.

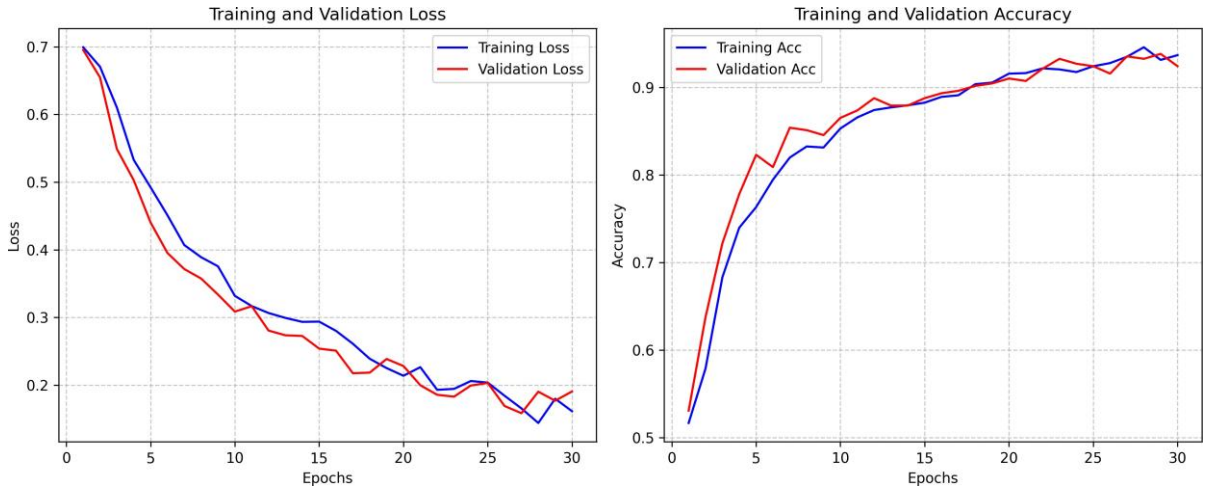


Figure 2. Training and validation loss and accuracy curves over 30 epochs.

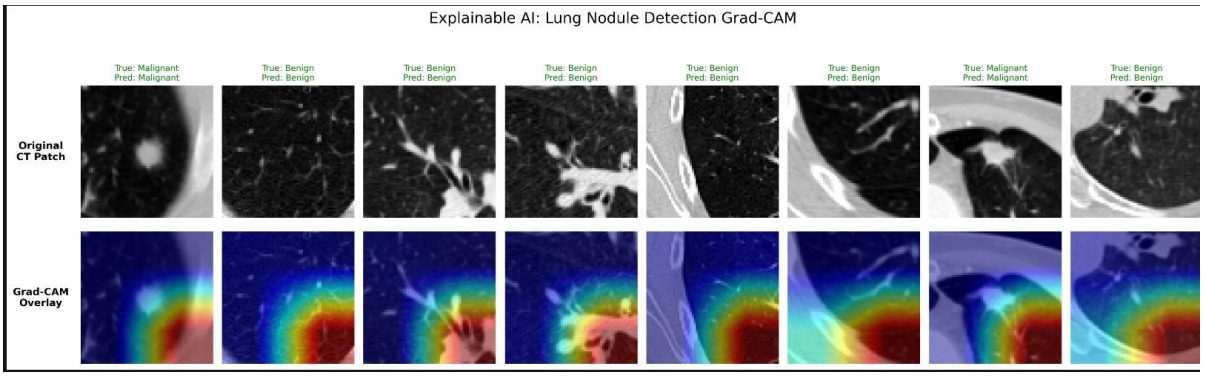


Figure 3. Representative Grad-CAM heatmaps for correctly classified nodule and non-nodule patches, illustrating the regions the model attends to when making a prediction.

6.3 Comparison with Published Methods

Table 4 situates the proposed system within the context of comparable patch-based classification protocols reported in the literature. It is important to note that these figures are drawn from heterogeneous evaluation protocols (different train/test splits, cross-validation schemes, and in some cases different datasets), so the comparison should be read as indicative of the general performance landscape rather than a strictly controlled benchmark. Several recent architectures that incorporate multi-view fusion or attention mechanisms report higher raw accuracy and AUC than the present ResNet-50 baseline. The proposed system's distinguishing contribution is therefore not state-of-the-art raw accuracy, but the combination of solid, reproducible 2D patch-based classification performance with integrated real-time explainability and a fully deployed clinical interface—an all-in-one combination that, to the authors' knowledge, none of the compared works provide.

Table 4. Comparison with Published Methods on LUNA16 / LIDC-IDRI

Method	Backbone	Acc.	AUC	XAI + Deploy
Shen et al. [8]	Multi-scale CNN	86.84%	—	No
Dou et al. [9]	3D CNN	—[1]	—	No
Xie et al. [10]	Multi-view CNN (ResNet-50)	91.60%	0.957	No
UrRehman et al. [5]	Dual-attention CNN	~95.4%	~0.980	No
Proposed (ResNet-50)	**ResNet-50**	**90.45%**	**0.9615**	**Yes**

XAI + Deploy: real-time explainability integrated with a deployed clinical interface. [1]Dou et al. report an average sensitivity (CPM score) of 0.827 across seven predefined false-positive rates on the detection task, which is not directly comparable to single-point classification accuracy.

7. DISCUSSION AND CONCLUSION

High classification accuracy alone is insufficient for clinical adoption; radiologists require a verifiable rationale. The Grad-CAM heatmaps generated by the proposed system provide this visual confirmation, supporting the system's potential validity as a second-reader tool, subject to prospective clinical evaluation. Limitations include the lack of 3D volumetric context, reliance on the LUNA16 nodule/non-nodule labelling scheme rather than pathology-confirmed malignancy outcomes, and the need for external validation on diverse clinical datasets outside the LIDC-IDRI cohort, including Indian patient populations with potentially different nodule etiologies (e.g., a higher relative prevalence of granulomatous disease).

In conclusion, this paper presented an end-to-end explainable AI system for pulmonary nodule classification, achieving an accuracy of 90.45% and an AUC-ROC of 0.9615 on a held-out LUNA16 test set. The integration of real-time Grad-CAM heatmaps and a production web interface addresses the explainability and deployability gaps that have historically limited the clinical translation of CAD systems, even where raw classification accuracy is not state-of-the-art. Future work will focus on external validation on CT data from Indian clinical sites, extension to 3D volumetric analysis, and a prospective reader study quantifying the impact of the Grad-CAM overlay on radiologist decision-making.

ACKNOWLEDGMENT

The authors thank the radiologists at Acharya Vinoba Bhave Rural Hospital (AVBRH) for domain guidance.

Data Availability

The datasets analysed during the current study are publicly available in the LUNA16 (Lung Nodule Analysis 2016) repository, `https://luna16.grand-challenge.org/Data/`, which is derived from the LIDC-IDRI database, `https://www.cancerimagingarchive.net/collection/lidc-idri/`.

REFERENCES

References

[1] H. Sung et al., "Global cancer statistics 2020," *CA: A Cancer Journal for Clinicians*, vol. 71, no. 3, pp. 209–249, May 2021.

[2] National Lung Screening Trial Research Team, "Reduced lung-cancer mortality with low-dose computed tomographic screening," *New England Journal of Medicine*, vol. 365, no. 5, pp. 395–409, Aug. 2011.

[3] Y. Lei et al., "Deep learning in pulmonary nodule detection and segmentation: a systematic review," *European Radiology*, vol. 34, no. 3, pp. 1188–1200, 2024.

[4] D. Ardila et al., "End-to-end lung cancer detection on computed tomography with deep learning," *Nature Medicine*, vol. 25, no. 6, pp. 954–961, Jun. 2019.

[5] Z. UrRehman, Y. Qiang, L. Wang, Y. Shi, Q. Yang, S. U. Khattak, R. Aftab, and J. Zhao, "Effective lung nodule detection using deep CNN with dual attention mechanisms," *Scientific Reports*, vol. 14, p. 3934, 2024.

[6] S. G. Armato III et al., "The Lung Image Database Consortium (LIDC) and Image Database Resource Initiative (IDRI)," *Medical Physics*, vol. 38, no. 2, pp. 915–931, Feb. 2011.

[7] A. A. A. Setio et al., "Validation, comparison, and combination of algorithms for automatic detection of pulmonary nodules in computed tomography images," *Medical Image Analysis*, vol. 42, pp. 1–13, 2017.

[8] W. Shen, M. Zhou, F. Yang, C. Yang, and J. Tian, "Multi-scale convolutional neural networks for lung nodule classification," in *Proc. IPMI*, Jun. 2015, pp. 588–599.

[9] Q. Dou, H. Chen, L. Yu, J. Qin, and P.-A. Heng, "Multilevel contextual 3-D CNNs for false positive reduction in pulmonary nodule detection," *IEEE Trans. Biomedical Engineering*, vol. 64, no. 7, pp. 1558–1567, Jul. 2017.

[10] Y. Xie, Y. Xia, J. Zhang, Y. Song, D. Feng, M. Fulham, and W. Cai, "Knowledge-based collaborative deep learning for benign-malignant lung nodule classification on chest CT," *IEEE Trans. Medical Imaging*, vol. 38, no. 4, pp. 991–1004, Apr. 2019.

[11] M. Tan and Q. V. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proc. ICML*, Jun. 2019.

[12] B. Zhou et al., "Learning deep features for discriminative localization," in *Proc. IEEE CVPR*, Jun. 2016, pp. 2921–2929.

[13] R. R. Selvaraju et al., "Grad-CAM: Visual explanations from deep networks via gradient-based localization," *Intl. Journal of Computer Vision*, vol. 128, pp. 336–359, 2020.

[14] F. Liao et al., "Evaluate the malignancy of pulmonary nodules using the 3-D deep leaky noisy-or network," *IEEE Trans. Neural Networks and Learning Systems*, vol. 30, no. 11, pp. 3484–3495, Nov. 2019.

[15] Riverain Technologies, “Riverain Technologies earns FDA approval for ClearRead CT software,” Sep. 2016. [Online]. Available: https://www.riveraintech.com

[16] Optellum Ltd., “Optellum receives FDA clearance for the world’s first AI-powered clinical decision support software for early lung cancer diagnosis,” Mar. 2021. [Online]. Available: https://www.optellum.com

[17] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proc. IEEE CVPR*, Jun. 2016, pp. 770–778.

AUTHOR CONTRIBUTIONS

S.D. and S.H. contributed to the conceptualization and methodology of the study. S.D. performed data curation, model development, training, validation, and experimental analysis, and prepared the figures, tables, and manuscript text. A.H. contributed to data preparation, experimentation, and interpretation of results. A.N. contributed to the system deployment architecture and critically reviewed the methodology. S.H. supervised the study and provided critical guidance throughout the research. All authors reviewed, edited, and approved the final manuscript.

FUNDING

Open access funding provided by Datta Meghe Institute of Higher Education and Research.
The authors received NO FUNDING for this work.