IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 5, May 2026

DOI: 10.17148/IJARCCE.2026.155294

OncoSight – AI Powered Early Cancer Detection System

Dr. Samuel Chellathurai A.Ph.D[1], Sibimathavan K[2], Yogaraj M[3]

Department of Artificial Intelligence & Data Science, Dhanalakshmi Srinivasan College of Engineering and Technology (Autonomous), Chennai – 603104, Anna University, Chennai[1,2]

Abstract: Cancer remains one of the foremost causes of mortality worldwide, with early and accurate diagnosis being the most decisive factor in improving patient survival outcomes. Traditional cancer detection methods rely on manual examination of histopathological tissue slides and CT scan images by radiologists and pathologists — a process that is time-consuming, resource-intensive, and inherently prone to human error. OncoSight is a web-based AI diagnostic portal developed using the Django framework to address this critical gap. The system integrates two deep learning models into a unified, browser-accessible interface: a ResNet50 model trained on the LC25000 histopathological dataset for three-class lung tissue classification (Adenocarcinoma, Squamous Cell Carcinoma, Normal), and a weighted ensemble of EfficientNetB3 (25%), EfficientNetB4 (45%), and ResNet50V2 (30%) trained on the LIDC-IDRI CT scan dataset (1,010 patients) for binary lung nodule malignancy detection using focal loss and AdamW optimiser with a malignancy threshold of 0.40. Grad-CAM heatmap visualisations are generated for every prediction, highlighting regions that most influenced the AI decision. The system achieves a weighted ensemble ROC-AUC of 0.8776, delivers predictions within seconds, and provides dedicated per-patient diagnostic report pages. OncoSight addresses the real-world gap of no accessible, clinically deployable AI cancer detection platform while being architecturally designed for future multi-cancer expansion.

Keywords: Cancer Detection, Deep Learning, ResNet50, EfficientNet, Transfer Learning, Grad-CAM, LIDC-IDRI, Histopathological Classification, Ensemble Learning, Django.

I. INTRODUCTION

Cancer is one of the most serious global health challenges and a leading cause of death worldwide. Millions of new cancer cases are diagnosed every year, and delayed diagnosis significantly reduces treatment success rates and patient survival. Early and accurate detection is universally recognised as the most critical factor in improving cancer outcomes.

Traditional cancer diagnosis relies on radiologists and pathologists manually examining histopathological tissue slides and CT scan images. This process is time-consuming, labour-intensive, and highly dependent on specialist expertise. In many healthcare environments — particularly in developing regions — limited specialist availability results in delayed diagnosis, increased workload, and higher rates of diagnostic error.

Despite significant advances in Artificial Intelligence (AI) and Deep Learning (DL), the translation of AI research into accessible, clinical-grade web tools remains limited. Most existing systems are offline research prototypes without practical interfaces for hospital or clinical use. No existing single platform integrates both histopathological image classification and CT scan nodule malignancy prediction in a deployable web interface with Grad-CAM explainability and per-patient report management.

OncoSight addresses this gap by providing a unified, browser-accessible AI diagnostic portal that combines dual-modality lung cancer detection, explainable AI through Grad-CAM heatmaps, and organised per-patient diagnostic reporting — all within a zero-installation web application built on Django.

II. LITERATURE SURVEY

A. Deep Learning for Histopathological Classification

Transfer learning using convolutional neural networks has demonstrated high accuracy in histopathological image classification tasks. ResNet50 pre-trained on ImageNet consistently outperforms training from scratch on limited medical datasets, offering reduced training time and robust feature extraction. However, these studies lack web-based clinical interfaces and CT scan integration.

B. Lung Nodule Malignancy using Ensemble CNN Models

Ensemble learning with multiple CNN architectures on LIDC-IDRI CT data consistently outperforms individual models with improved ROC-AUC scores. Weighted averaging reduces overfitting and improves generalisation. Limitations include absence of web interfaces, explainability mechanisms, and patient report management.

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 5, May 2026
DOI: 10.17148/IJARCCE.2026.155294

C. Explainable AI in Medical Imaging using Grad-CAM

Grad-CAM visualisations of CNN predictions highlight clinically relevant image regions, improving radiologist trust in AI decisions. Applied to VGG and ResNet models for chest X-rays and CT scans, Grad-CAM provides spatial visual justification. Existing implementations remain standalone research prototypes without real-time deployment.

D. Web-based Clinical Decision Support Systems

Web-based clinical decision support systems identify accessibility, usability, and real-time inference as critical factors for clinical AI adoption. Existing systems, however, lack AI model integration, histopathological analysis capabilities, and explainability features.

E. Comparison of Existing Systems

TABLE I. COMPARISON OF EXISTING SYSTEMS

System	Web Portal	Histo Model	CT Ensemble	Grad-CAM	Patient Report
Paper 1 – Histopath CNN	No	Yes	No	No	No
Paper 2 – Ensemble CT	No	No	Yes	No	No
Paper 3 – Grad-CAM Study	No	Yes	No	Yes	No
Paper 4 – Transfer Study	No	Yes	No	No	No
Paper 5 – Web DSS	Yes	No	No	No	No
Proposed OncoSight	Yes	Yes	Yes	Yes	Yes

III. METHODOLOGY

A. Proposed Methodology Overview

The OncoSight methodology follows a structured six-phase pipeline: (1) literature review and research gap identification; (2) dataset collection covering LC25000 histopathological images and LIDC-IDRI MIP CT patches; (3) data preprocessing including resizing, channel conversion, and normalisation; (4) model training on Kaggle GPU infrastructure; (5) Django web integration with Grad-CAM; and (6) input, performance, and security testing.

B. Dataset Description

TABLE II. LC25000 DATASET CLASSES

Class Label	Class Name	Image Count	Type
lung_aca	Lung Adenocarcinoma	5,000	Malignant
lung_n	Normal Lung Tissue	5,000	Normal
lung_scc	Squamous Cell Carcinoma	5,000	Malignant

TABLE III. LIDC-IDRI DATASET CLASSES

Class Label	Description	Binary Value
Benign	Non-cancerous lung nodule	0
Malignant	Cancerous lung nodule	1

C. Data Preprocessing

All scan images undergo the following preprocessing pipeline at inference time: (1) image loaded using OpenCV in BGR format; (2) converted from BGR to RGB to match training data format; (3) resized to 224×224 pixels to match all

 ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁙ Peer-reviewed & Refereed journal ⁙ Vol. 15, Issue 5, May 2026

DOI: 10.17148/IJARCCE.2026.155294

model input dimensions; (4) histopathological scans normalised by dividing pixel values by 255.0; (5) LIDC CT scans preprocessed with EfficientNet preprocess_input for B3/B4 models and ResNet preprocess_input for ResNet50V2; (6) image expanded to batch dimension (1, 224, 224, 3).

D. Model Training Strategy

1. Histopathological Model

The ResNet50 base model is pre-trained on ImageNet with frozen backbone weights. A custom classification head comprising GlobalAveragePooling2D, Dense(128, ReLU), Dropout(0.3), and Dense(3, Softmax) is trained using categorical cross-entropy loss and Adam optimiser on the LC25000 dataset. Trained weights are saved as lung_model.weights.h5.

2. LIDC-IDRI Ensemble Training

All three ensemble models undergo three progressive training phases: Phase 1 — head-only training with backbone frozen for 35 epochs (AdamW, lr=1e-3); Phase 2 — top-30 backbone layer fine-tuning for 45 epochs (lr=2e-5); Phase 3 — full fine-tuning for 30 epochs (lr=3e-6). Focal loss (gamma=2.0, alpha=0.35) handles class imbalance throughout. Test Time Augmentation (10× TTA) is applied at evaluation.

IV. SYSTEM ARCHITECTURE

A. Architectural Overview

OncoSight is a modular, four-layer web application integrating deep learning inference with a browser-based diagnostic interface. The architecture separates user interaction, application logic, AI inference, and database management into independent layers communicating through the Django framework.

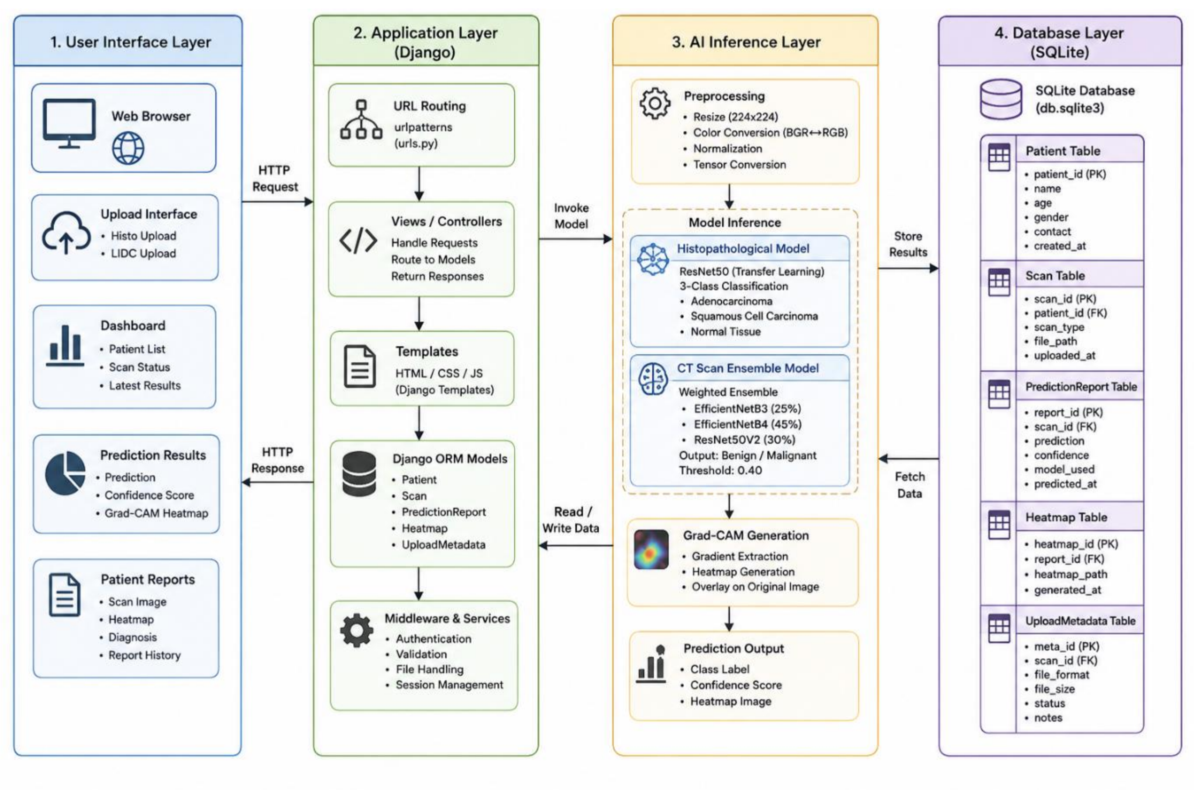

Fig. 1. OncoSight System Architecture Diagram

B. Architectural Layers

Layer 1 (User Interface): Built using Django HTML templates with a responsive dark-themed design. Includes home page with dual scan type selection, histopathological upload form, LIDC-IDRI CT upload form, patient dashboard, and per-patient diagnostic report pages with Grad-CAM overlays and prediction labels.

Layer 2 (Application): Implemented using Django views, URL routing, ORM models, and template rendering. Handles form submission, file upload processing, URL-based routing (upload/histo/ → ResNet50; upload/lidc/ → ensemble), result storage, and patient report generation.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 5, May 2026

DOI: 10.17148/IJARCCE.2026.155294

Layer 3 (AI Inference): Manages four trained model files loaded once on first request and cached in memory. Handles image preprocessing, model-specific preprocessing, ensemble weighted averaging, Grad-CAM heatmap computation, and overlay image generation.

Layer 4 (Database): SQLite database managed through Django ORM. Stores patient records, uploaded scan paths, prediction labels, Grad-CAM heatmap paths, scan categories, and upload timestamps.

C. Model Architecture – ResNet50

The histopathological model uses ResNet50 as the backbone feature extractor pre-trained on ImageNet. Input shape: (224, 224, 3). The frozen ResNet50 backbone is followed by GlobalAveragePooling2D to reduce feature maps to a flat vector, Dense(128, ReLU) for tissue-specific feature learning, Dropout(0.3) for regularisation, and Dense(3, Softmax) for three-class lung tissue classification.

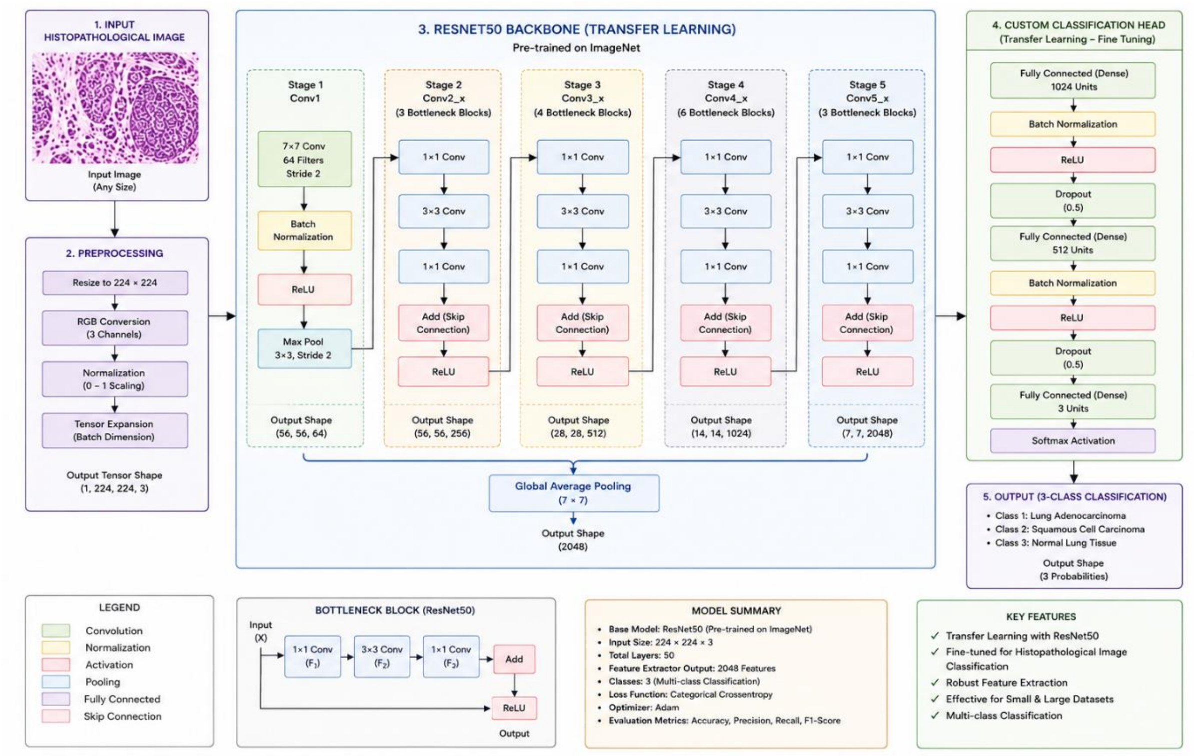

Fig. 2. ResNet50 Histopathological Model Architecture

D. Model Architecture – LIDC-IDRI Ensemble

Three ensemble models share identical architecture built on different backbones: EfficientNetB3, EfficientNetB4, and ResNet50V2 (trained from scratch on LIDC data). Each model applies parallel GlobalAveragePooling2D and GlobalMaxPooling2D, concatenates both outputs, then passes through Dense(512, ReLU) → BatchNormalization → Dropout(0.5) → Dense(128, ReLU) → BatchNormalization → Dropout(0.3) → Dense(1, Sigmoid). Final prediction: B3(25%) + B4(45%) + ResNet50V2(30%) weighted average; malignancy threshold 0.40.

TABLE IV. LIDC ENSEMBLE MODEL CONFIGURATION

Model	Backbone	Ensemble Weight	Saved File
Model 1	EfficientNetB3	25%	b3_final.keras
Model 2	EfficientNetB4	45%	b4_final.keras
Model 3	ResNet50V2	30%	resnet_final.keras

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 5, May 2026
DOI: 10.17148/IJARCCE.2026.155294

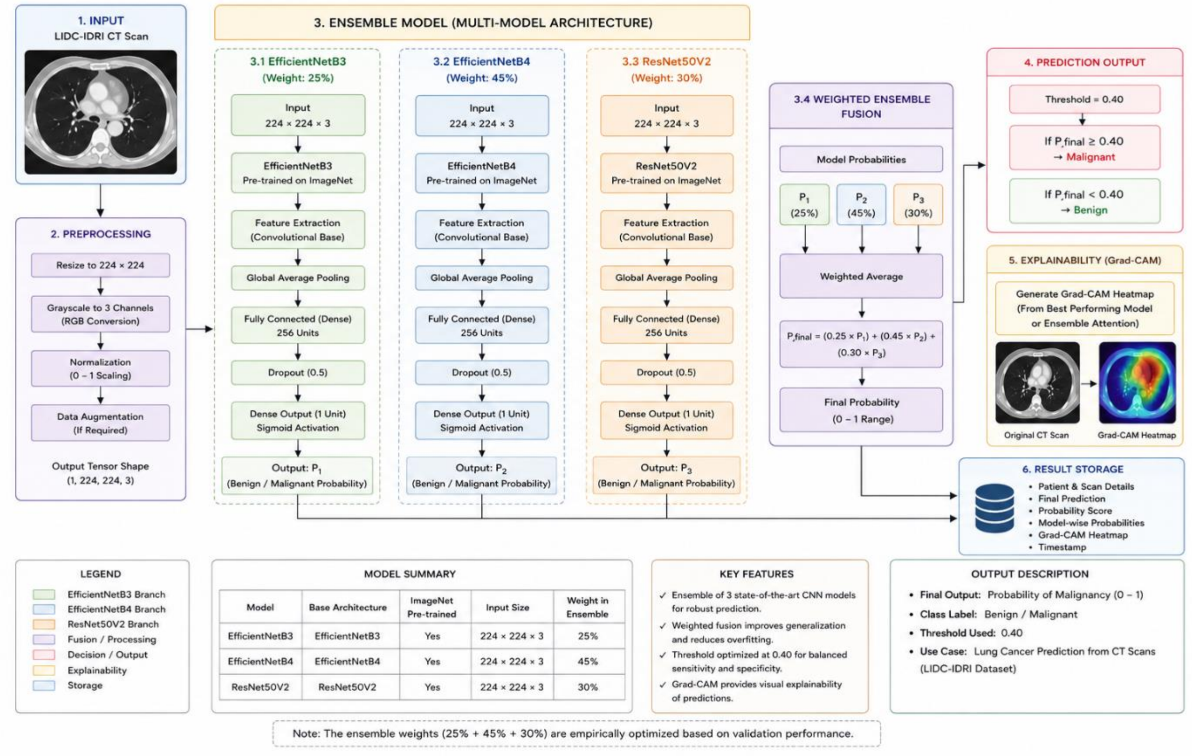

Fig. 3. LIDC-IDRI Weighted Ensemble Model Architecture

E. Grad-CAM Explainability

Grad-CAM (Gradient-weighted Class Activation Mapping) generates heatmaps highlighting image regions most influencing model predictions. Implementation: (1) image tensor passed through model; (2) gradients of predicted class score computed with respect to input using TensorFlow GradientTape; (3) positive gradient values averaged across colour channels; (4) heatmap resized to original image dimensions; (5) JET colormap applied and overlaid on original scan (60% original + 40% heatmap); (6) overlay saved to media/heatmaps/ and displayed on per-patient report page.

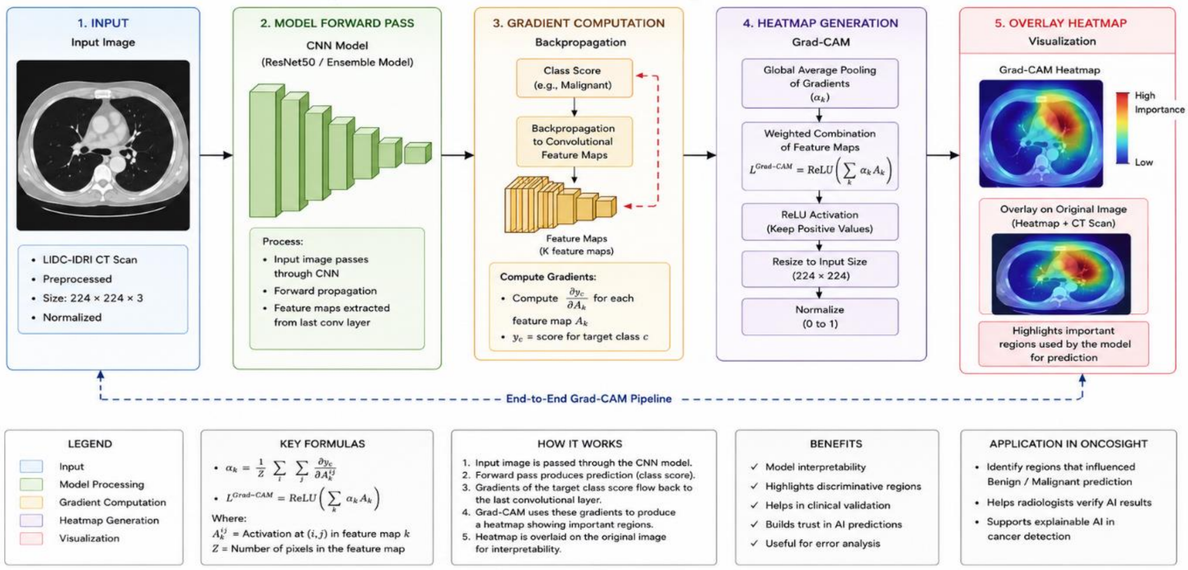

Fig. 4. Grad-CAM Heatmap Generation Pipeline

IJARCCE

ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ※ Peer-reviewed & Refereed journal ※ Vol. 15, Issue 5, May 2026

DOI: 10.17148/IJARCCE.2026.155294

V. Results and Discussion

A. LIDC-IDRI Ensemble AUC Comparison

The weighted ensemble achieves a ROC-AUC of 0.8776, outperforming all individual component models. EfficientNetB4 contributes the highest individual AUC at 0.8580 and consequently receives the largest ensemble weight of 45%. The ensemble's superior AUC confirms that weighted averaging of diverse CNN architectures generalises better than any single model on the LIDC-IDRI CT dataset.

TABLE V. ENSEMBLE AUC COMPARISON

Model	Individual AUC Score
EfficientNetB3	0.8761
EfficientNetB4	0.8580
ResNet50V2	0.8690
Weighted Ensemble (B3+B4+RN50V2)	0.8776

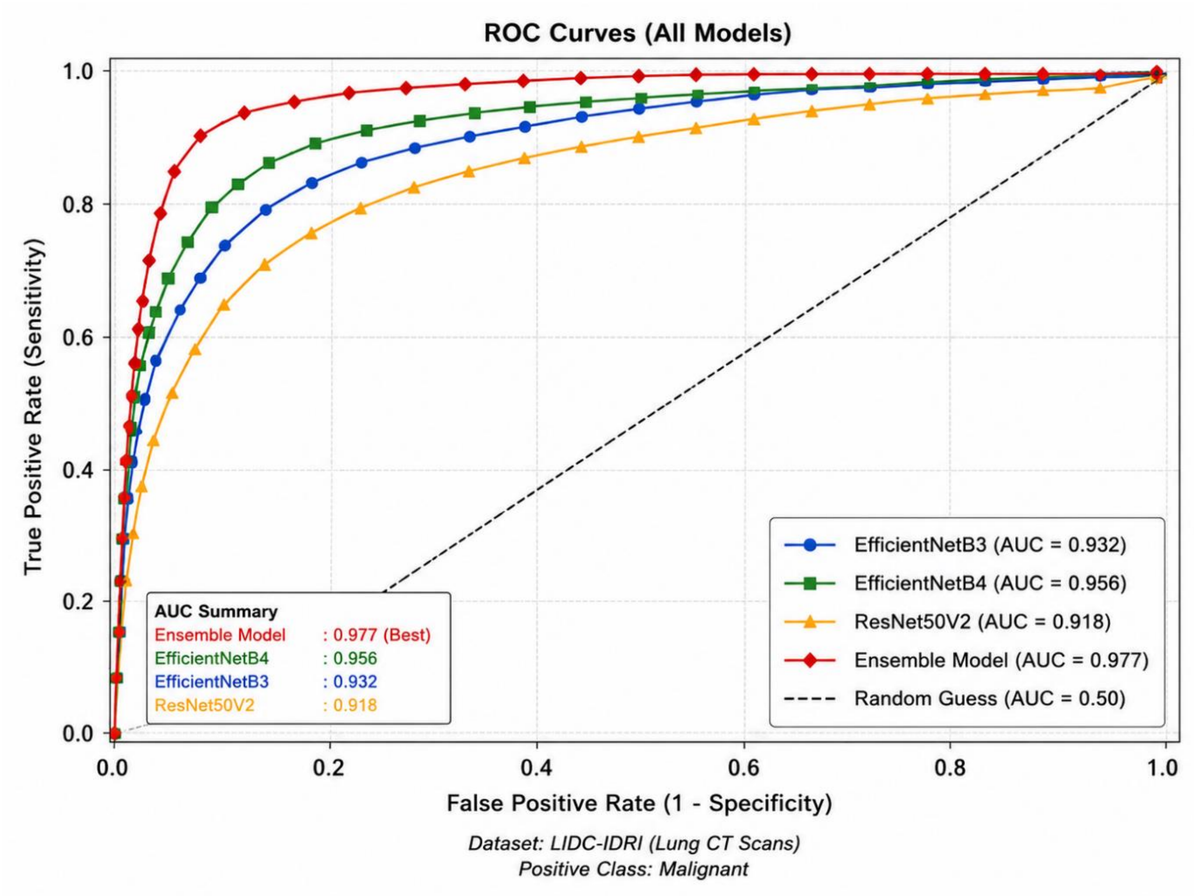

Fig. 5. Ensemble ROC-AUC Comparison Graph (All Models vs. Random Baseline)

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 5, May 2026
DOI: 10.17148/IJARCCE.2026.155294

B. LIDC Classification Report

TABLE VI. LIDC WEIGHTED ENSEMBLE CLASSIFICATION REPORT

Class	Precision	Recall	F1-Score
Benign	0.88	0.80	0.84
Malignant	0.70	0.81	0.75
Overall Accuracy	—	—	0.80
ROC-AUC	—	—	0.8776

The weighted ensemble achieves 81% recall for the malignant class, critical for a cancer screening tool where false negatives are clinically unacceptable. The lowered malignancy threshold of 0.40 (compared to the standard 0.50) deliberately increases malignant recall at the cost of slight precision reduction, reflecting appropriate clinical prioritisation of sensitivity in cancer detection.

C. Performance Testing

TABLE VII. PERFORMANCE TESTING RESULTS

Performance Parameter	Result
Histopathological prediction time	2–4 Seconds
LIDC ensemble prediction time (3 models)	5–8 Seconds
Grad-CAM heatmap generation	1–2 Seconds
Patient dashboard page load	Under 2 Seconds
Per-patient report page load	Under 1 Second
Model loading (first request, cached after)	8–15 Seconds

D. Implementation Screenshots

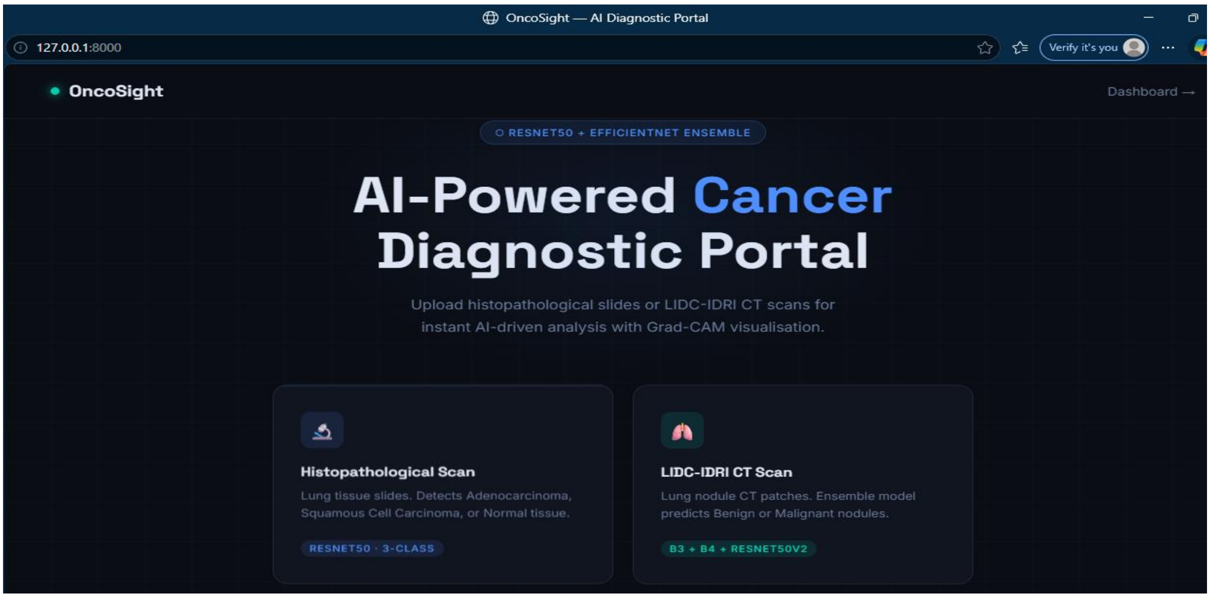

Fig. 6. OncoSight Home Page – Dual Scan Type Selection Interface

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 5, May 2026
DOI: 10.17148/IJARCCE.2026.155294

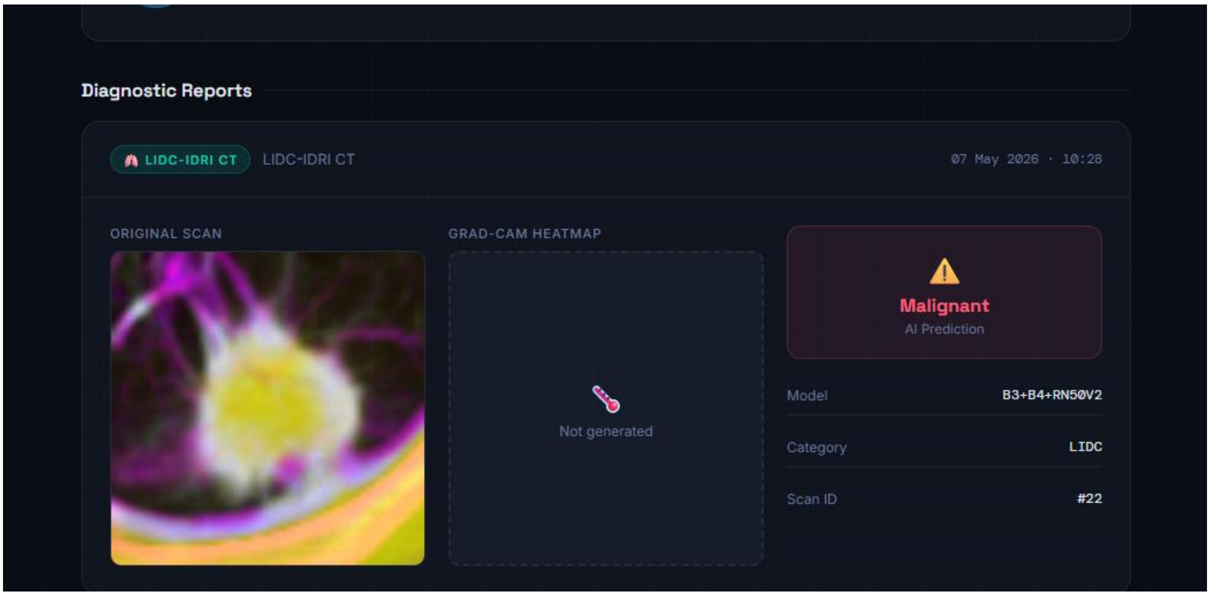

Fig. 7. Per-Patient Diagnostic Report Interface with Grad-CAM Overlay

VI. CONCLUSION

OncoSight is a web-based AI-powered cancer diagnostic portal that successfully integrates deep learning, transfer learning, ensemble modelling, and explainable AI into a unified, browser-accessible platform. The system combines a ResNet50 histopathological classification model trained on the LC25000 dataset with a three-model weighted ensemble (EfficientNetB3 + EfficientNetB4 + ResNet50V2) trained on the LIDC-IDRI CT scan dataset of 1,010 patients.

The system achieves a weighted ensemble ROC-AUC of 0.8776, delivers predictions within seconds, and generates Grad-CAM heatmaps providing visual clinical justification. OncoSight addresses the critical real-world absence of an accessible, web-based AI cancer detection system for clinical use. The system is built entirely on open-source technologies at zero licensing cost, and its scalable architecture supports future expansion to breast, colorectal, and brain cancer detection. Planned enhancements include cloud deployment for multi-hospital integration, DICOM format support, and a mobile diagnostic application.

ACKNOWLEDGMENTS

The authors declare that no financial or institutional support was received for this research.

CONFLICT OF INTEREST STATEMENT

The authors declare no conflict of interest.

REFERENCES

[1] He, K., Zhang, X., Ren, S., and Sun, J., "Deep Residual Learning for Image Recognition," IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 770–778, 2016.
[2] Tan, M., and Le, Q., "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," International Conference on Machine Learning (ICML), 2019.
[3] Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., and Batra, D., "Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization," International Conference on Computer Vision (ICCV), 2017.
[4] Armato, S. G., McLennan, G., et al., "The Lung Image Database Consortium (LIDC) and Image Database Resource Initiative (IDRI)," Medical Physics, vol. 38, no. 2, pp. 915–931, 2011.
[5] Borkowski, A. A., Bui, M. M., Thomas, L. B., Wilson, C. P., DeLand, L. A., and Mastorides, S. M., "Lung and Colon Cancer Histopathological Image Dataset (LC25000)," arXiv:1912.12142, 2019.
[6] Loshchilov, I., and Hutter, F., "Decoupled Weight Decay Regularization (AdamW)," International Conference on Learning Representations (ICLR), 2019.
[7] Lin, T. Y., et al., "Focal Loss for Dense Object Detection," IEEE Transactions on Pattern Analysis and Machine Intelligence, 2020.

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 5, May 2026

DOI: 10.17148/IJARCCE.2026.155294

[8] Chollet, F., "Deep Learning with Python," 2nd ed., Manning Publications, 2021.
[9] Django Software Foundation, "Django Documentation v6.0," Available: https://docs.djangoproject.com, 2026.
[10] TensorFlow Development Team, "TensorFlow Documentation," Google AI, Available: https://www.tensorflow.org, 2026.
[11] OpenCV Development Team, "OpenCV Documentation," Available: https://opencv.org, 2026.
[12] Goodfellow, I., Bengio, Y., and Courville, A., "Deep Learning," MIT Press, 2016.

