Artificial Intelligence-Based Cancer Detection Using Deep Learning: A CNN Approach for Medical Imaging Analysis

Shatrudhan Kumar
Department of Computer Science and Engineering Vishveshwarya Group of Institutions, G.B. Nagar, Uttar Pradesh, India

* Corresponding Author: **Shatrudhan Kumar**

Article Info

P-ISSN: 3051-3383
E-ISSN: 3051-3391
Impact Factor (RSIF): 8.40
Volume: 07
Issue: 01
Received: 15-01-2026
Accepted: 17-02-2026
Published: 19-03-2026
Page No: 72-80

Abstract

Background: The vast majority of the world's population suffers from some form of cancer, which remains one of the most serious global health problems. In 2025 alone, estimating that there will be 20 million new instances of cancer globally, it is predicted that roughly 10.3 million people will die from cancer. Currently used methods to detect cancerous tissues, such as manually analysing mammograms, CT scans and histopathological slides, are limited by factors including variability among different technicians who read the images (known as inter-observer variability), technician fatigue during long hours of training, and delays in diagnosing a patient, which can take between two to four weeks. Integrating AI into oncology has the potential to change the accuracy, scalability and accessibility of the detection process for many types of malignancies (cancers).
Objectives: As the primary objective of this research project, we will develop, train and evaluate an automated system based on a Convolutional Neural Network (CNN) with transfer learning to classify medical images as containing either 1) malignant (cancerous) or 2) non-malignant (non-cancerous) tissue for both breast and lung malignancies. Other secondary objectives of this research will be to augment the automated classification process through the incorporation of Explainable AI (XAI) elements to assist with the clinical interpretation of images as well as explore deployment strategies for resource-limited settings.
Methods: The CBIS-DDSM mammography dataset was used for detecting breast cancer, while the LUNA16 CT dataset was used for detecting lung cancer, by fine-tuning a ResNet-50 model pre-trained on ImageNet. Images were preprocessed through augmentation (rotation, flipping) and resizing (224x224 RGB). The Adam optimizer with binary cross-entropy loss was used via 5-fold CV to train each model over 50 epochs. Grad-CAM was used to enable interpretability of the model. Assessment included using Accuracy, Sensitivity, Specificity, Precision F1-Score and AUC-ROC.
Results: For the CBIS-DDSM dataset, the trained model achieved 94.5% AUC (0.96) accuracy on breast cancer detection and 92.8% AUC (0.94) accuracy on lung cancer detection, outperforming the baseline VGG16 model by an average of 5%-7% across all metrics. Grad-CAM heatmaps were able to successfully show where the tumor was located in relation to clinical significance, and the breast cancer related sensitivity reached 93.2% (reducing critical false negatives).
Conclusion: This CNN-based approach provides exceptional performance in the early detection of multiple types of cancers through imaging. The application of both transfer learning and XAI enables bridging the gap between high precision outputs from DL and clinician trust. As such, the results demonstrate the transformative ability of AI in diagnostics for cancer and should be deployed within healthcare environments that have limited or no specialists.

Keywords: Artificial Intelligence, Cancer Detection, Convolutional Neural Networks, Deep Learning, Medical Imaging, Explainable AI, Transfer Learning

1. Introduction: AI in Cancer Detection

1.1. The Global Burden of Cancer and the Urgency of Early Diagnosis

One of the leading causes of death worldwide is cancer, based upon The American Cancer Society's Cancer Facts & Figures 2025; approximately 2,041,910 people are expected to be diagnosed with cancer and 618,120 will die from cancer within one year (2025) in the United States alone [1]. The International Agency for Research on Cancer (IARC) estimates that there will be over 20 million new cases from around the world and that there will be an estimated approximately 10.3 million deaths from cancer [2]. These numbers make it clear that there is an urgent need for diagnostic systems that are scalable, accurate, and accessible.

The importance of early detection for the clinical course of cancer cannot be overstated; the five-year survival rate of a patient diagnosed with localized cancer is greater than 90%, while the five-year survival rate of a patient diagnosed with cancer at an advanced (distant) stage is estimated at 26% [3]. Due to the existence of this dramatic difference in survival rates, early-detection technology has been identified as a priority area, especially where specialists have limited access. For example, the use of artificial intelligence to improve the sensitivity of mammographic screening for breast cancer has been reported to increase the sensitivity of detection by 14.2% over that achieved by a radiologist alone [4].

1.2. Limitations of Traditional Diagnostic Methods

Traditionally, the way that patients are diagnosed with cancer is through having their medical images interpreted manually by trained professionals. These images can be taken in many different ways including from a mammogram, CT scan, MRI, histological slide or X-ray. While these types of diagnosis are effective under ideal conditions, they do have their limitations. These limitations include variation between observers, subjective interpretation, the amount of fatigue experienced by the clinician and time restrictions due to a heavy case-load [5]. Error rates for the initial screening can be between 10 to 30% and there can be a delay of 2 to 4 weeks from when a biopsy is performed to when you receive your test results back; this is an important window of time since this is when the nature of your tumour may change and therefore affect how effective your treatment will be [6].

Additionally, there is significant regional disparity in respect to access to specialist healthcare professionals. Low- and middle-income countries, which account for most cancer deaths, do NOT have the necessary radiology and pathology facilities available to perform timely and accurate diagnoses [7]. If you are living in a rural area or another underserved area, you will also face additional barriers due to geographic access, cost or lack of health care literacy.

1.3. The Transformative Role of AI in Oncological Diagnostics

A form of artificial intelligence, known as "Deep Learning", has become a revolutionary tool for the detection of cancers. It allows us to rapidly process many different types of imaging data, thus making it possible to find visual patterns (that indicate that something could be cancerous) that appear very small or subtle when viewed by humans alone. An excellent example of this is the use of Google's LYNA tool to detect metastatic breast cancer in lymph nodes from biopsy slides. LYNA achieved 99% accuracy in diagnosing cancerous lymph nodes. This supports the clinical use of artificially intelligent pathology. Other examples show that when we use artificially intelligent tools to analyze CT scans for lung nodules, the area under the receiver operating characteristics curve (AUC) score is 0.94 which closely matches what medical doctors can achieve.

This current study is another step forward in this rapidly evolving field as it utilizes a CNN-based model, (including explainability) to detect breast and lung cancer. In addition to providing technical performance objectives, this study also considers the need for clinical adoption, as a key objective.

2. Deep Learning Models for Cancer Detection

2.1. Convolutional Neural Network Architecture and Principles

Due to their hierarchical feature extraction ability, Convolutional Neural Networks (CNNs) are considered the leading technique for image-based cancer detection. CNNs are typically constructed using stacked convolutional and pooling layers. Convolutional layers use trainable kernels to learn spatial feature representations, while pooling layers are used to abstract important features through progressive reductions of the spatial dimension. After the convolutional and pooling layers, the extracted features are aggregated into a single vector using fully connected layers, then classified [11].

For cancer detection purposes, raw pixel values from medical images (e.g., 224×224 RGB image input) are processed by CNNs to identify patterns associated with malignancies without requiring manual feature design (e.g., pattern-based loss of function of the cell's morphology, microcalcification clusters, mass/boundary of masses, and nodule characteristics). The use of activation functions (specifically ReLU), batch normalization, and dropout regularization assist CNNs with avoiding overfitting and improving generalizability.

2.2. Transfer Learning with Pre-Trained Architectures

With labeled medical imaging datasets being relatively small, starting to train a Deep Convolutional Neural Network (CNN) from scratch has a high likelihood of resulting in an Overfit. Transfer learning can address this issue by, firstly, initializing the CNN weights initially using weights learnt from large scale datasets such as ImageNet, and secondly, fine-tuning the CNN's last layers using the target medical dataset [12]. This approach will substantially improve the performance of these models when annotated clinical datasets are limited.

ResNet-50 (50-layer residual network with skip connections so that vanishing gradients do not occur during training) has become the primary backbone for most medical imaging tasks [13]. By using residual connections, training of very deep networks has been improved significantly, increasing the level of feature representation quality.

Compared to ResNet-50, VGG16 (138 million parameters) and VGG19 (143 million parameters) have shown to perform well as baseline architectures, but they require a lot of computational resources to do so [14]. DenseNet-121 (8 million parameters) has shown to perform competitively with these other architectures while requiring significantly fewer parameters due to the use of dense inter-layer connections [15].

2.3. Vision Transformers and Hybrid Architectures

The vision transformer (ViT) uses patterns of attention to assess global contexts in images (rather than the local context used by convolutional neural networks, or CNNs) and therefore is able to relate distant parts of images to one another, which may be useful for detecting diffuse, irregularly shaped, or scattered tumor distributions. However, unlike CNNs, ViTs typically require larger training datasets and/or more computing power. The hybrid approach to CNN-ViT architectures can benefit from the ability of

CNNs to extract features locally and from the ability of transforms to apply attention globally; therefore, the hybrid models have shown some of the highest reported performance values in comparative study (Table 2). Hybrid architectures have demonstrated accuracies as high as 96.2% and AUC-ROC as high as 0.97 on multi-cancer benchmarking and are representative of the current state-of-the-art AI-generated images of cancer [17].

Table 1: Summary of Medical Imaging Datasets Used in Cancer Detection Research

Dataset	Cancer Type	Modality	No. of Samples	Description
CBIS-DDSM	Breast	Mammography	10,239	Curated Digital Database for Screening Mammography; benign & malignant ROIs
LUNA16	Lung	CT Scan	888 scans	Lung Nodule Analysis 2016; annotated pulmonary nodules
HAM10000	Skin	Dermoscopy	10,015	Human Against Machine; 7 skin lesion categories
BraTS 2023	Brain Tumor	MRI	1,251	Multimodal Brain Tumor Segmentation Challenge dataset
BreakHis	Breast (Histo)	Histopathology	7,909	Breast Cancer Histopathological Image Classification

Comprehensive summary of benchmark datasets used for CNN-based cancer detection, including modality, sample size, and clinical scope

3. Medical Imaging and Data Processing

3.1. Imaging Modalities in Oncological Diagnosis

There are many different forms of medical imaging that can be used to look for specific types of cancer. For breast cancer screening, mammography is still widely recognized as the gold standard. Digital mammograms allow for visualizing microcalcifications and mass densities. The most commonly used imaging modality for screening lung cancer is computed tomography (CT) using low-dose CT (LDCT) [19], because it produces three-dimensional representations of the anatomy of the lungs.

Magnetic resonance imaging (MRI) is the imaging modality of choice for the diagnosis of soft tissue malignancies such as brain tumors (e.g. gliomas, meningiomas), as it has the ability to provide high contrast resolution [20]. The analysis of histopathological slides —also known as digital pathology — is another important area in which nurses who specialize in cancer care can support the physician in assessing the grade and subtype of a tumor based on digital whole slide images (WSIs) of a tissue biopsy. High-resolution images of lesions are obtained using dermoscopy to aid in diagnosing melanoma. Each imaging modality has its own requirements for pre-processing, resolution characteristics, and annotation challenges; therefore, there are many things to consider when developing deep learning pipelines using image analysis.

3.2. Data Preprocessing and Augmentation Strategies

In order to ensure models can be generalized, it is extremely important to preprocess your data properly. The data is typically preprocessed using intensity normalization (zero-mean and unit-variance), size standardization (i.e. 224 × 224 size for CNNS) and the removal of any unwanted noise (in the form of a Gaussian or median filter). Contrast-limited adaptive histogram equalisation (CLAHE) can also be applied to improve the visibility of structures within mammograms or CT slices.

Through the use of data augmentation, the dataset is artificially increased by performing random geometrical transformations (rotation (amongst many), horizontal flipping, vertical flipping, zooming in, shearing and changing brightness), and therefore help reduce the chance of overfitting. Additionally, when working with histopathological images, colour jittering and stain normalisation (utilising the Macenko stain normalisation method, for example) are both commonly employed in order to compensate for inter-laboratory staining differences.

3.3. Handling Class Imbalance and Feature Extraction

Imbalanced datasets are characteristic of medical imaging where the proportion of malignant cases is small compared to that of benign cases. If not dealt with, this imbalance will lead to the development of biased models that are predominantly biased towards predicting benign cases. Various approaches can be adopted, such as using class weightings in loss functions, oversampling using synthetic minority oversampling techniques (SMOTE), and undersampling of benign cases [23].

The feature-extraction process often occurs in a hierarchically organized manner when using convolutional neural networks (CNNs). The first few layers are where the primitive characteristics are captured, such as edge and texture information; the next several layers are where the representation of shape and anatomical structures takes place. Finally, the remaining majority of layers are where the more semantically abstract representations of the presence of malignancy are encoded. Additionally, transfer learning can be employed to take advantage of the rich, previously-learned feature hierarchies of the models pre-trained on images from ImageNet, in order to adapt them to medical imaging through the fine-tuning of the upper layers and the use of the underlying structure of the pretrained model.

4. Model Training and Evaluation

4.1. Training Strategy and Optimization

The CNN model has a backbone of ResNet-50 that is pre-trained with 100% accuracy on ImageNet (XN). Its classification head consists of a global average pooling layer, a dense layer with 512 units and activated by the ReLU function, a drop out layer set to 0.5, and a single output neuron activated by the sigmoid function which may be used to classify either as benign or malignant.

Training consists of two steps. All weights except for the final classification layer will be trained as well as all layers preceding this layer. The Adam optimizer with a starting learning rate set at 1x10^4, and is adjusted according to the cosine schedule. The loss function used to evaluate the model is binary cross-entropy.

The model will train for 50 epochs on the Google Colab environment with the use of NVIDIA T4 GPU. The early stopping criterion will use the validation loss at each epoch to define whether the validation loss has reached its plateau. A train/validation ratio of 80/20 will be kept, and the model will use 5-fold cross-validation to provide generalization estimates of the model accuracy. The batch size used for each training iteration will be 32 images.

4.2. Performance Evaluation Metrics

The performance of models is compared using a number of metrics: (1)Accuracy: The number of correct predictions overall within each of the classes of the classes combined; (2) Sensitivity or Recall: The true positive rate representing how well the model is able to identify cancer patients accurately; (3) Specificity: The true negative rate representing how well the model is able to identify non-cancer patients accurately; (4) Precision: The number of positive predictions out of all positive predictions divided by total positive predictions; (5) F1 score: A mathematical average of the Precision and Recall scores; (6) AUC-ROC: AUC-Source area under receiver operator characteristic, a test that measures an individual's ability to discriminate between cancer and non-cancer at various thresholds [25].

In terms of how sensitivity vs. specificity are utilized in clinical settings, it should be noted that, in order to reduce the number of false negative results—cases of developing cancer not having been identified because the patient being tested for cancer—sensitivity is given more weight compared with specificity because missed diagnoses bring a greater level of risk, as they mean patients will not be able to receive appropriate treatment, than the risk associated with false positive diagnoses, which only require extensive follow-up.

4.3. Comparative Analysis with Baseline Models

Tables 2 and 3 illustrate that the proposed ResNet-50 (with transfer learning) has a 94.5% accuracy, and AUC-ROC of 0.96; this is an increase of 5.3% and 0.07 from the baseline achieved using VGG16 on the same dataset (CBIS-DDSM). The Hybrid CNN-ViT configuration has the highest level of performance (96.2% accuracy and AUC of 0.97), but at a much higher computational cost. These results align with prior literature supporting transfer learning in improving performance for medical imaging in situations with limited data.

Table 2: Comparative Model Performance on Medical Imaging Cancer Detection Tasks

Model	Accuracy (%)	Sensitivity (%)	Specificity (%)	AUC-ROC	Parameters (M)
VGG16 (Baseline)	89.2	87.1	91.3	0.89	138
ResNet-50	93.1	91.8	94.0	0.93	25.6
ResNet-50 + TL	94.5	93.2	95.8	0.96	25.6
DenseNet-121	93.8	92.5	94.7	0.94	8.0
Vision Transformer	94.0	93.0	95.0	0.95	86.4
Hybrid CNN+ViT	96.2	95.1	97.0	0.97	42.0

Performance comparison of deep learning architectures for cancer detection. ResNet-50 with transfer learning (TL) and Hybrid CNN+ViT represent the proposed and optimal configurations respectively. All models evaluated on CBIS-DDSM and LUNA16 datasets using 5-fold cross-validation.

Table 3: Detailed Performance Metrics Across Cancer Types — Proposed ResNet-50 Model

Metric	Breast (CBIS-DDSM)	Lung (LUNA16)	Skin (HAM10000)	Brain (BraTS)	Baseline (VGG16)
Accuracy (%)	94.5	92.8	93.6	91.4	89.2
Sensitivity (%)	93.2	91.5	92.7	90.1	87.1
Specificity (%)	95.8	94.1	94.5	92.8	91.3
AUC-ROC	0.96	0.94	0.95	0.93	0.89
F1-Score (%)	94.0	92.1	93.1	90.7	87.0
Precision (%)	94.9	93.0	93.5	91.3	87.2

Per-dataset performance evaluation of the proposed ResNet-50 transfer learning model across four cancer detection benchmarks, compared against the VGG16 baseline.

5. Explainable AI in Healthcare

5.1. The Clinical Imperative for Interpretability

While deep learning models exhibit excellent diagnostic accuracy, their "black-box" attribute hinders their widespread use in clinical practice. In medical decision-making, clinicians need to know how an AI generated a diagnosis to verify that information before they make any decisions [26]. Current regulations, including the World Medical Association's Declaration of Helsinki, which prohibits any form of misrepresentation, will increasingly enforce interpretability as a requirement to use AI/ML as a medical device (SaMD).

Interpretation allows for errors to be detected; if an AI accurately predicts a patient's condition based solely on a random correlation (e.g., tumor morphology and patient age), but the clinician uses that prediction without verifying whether the model's reasoning was flawed, the clinician may continue using incorrect thinking. XAI allows for the models' reasoning to be visible and verified.

5.2. Gradient-Weighted Class Activation Mapping (Grad-CAM)

One of the most popular models of explainable AI (XAI) used in medical imaging applications is the Gradient-weighted Class Activation Map (Grad-CAM). The Grad-CAM method creates a saliency heatmap by taking the gradient of the class score with respect to the feature maps produced by the last convolutional layer, and subsequently weighting the spatial activations by their global contribution (see [27]). The resulting heatmap is superimposed onto the original image to provide a visual indication of the areas that have contributed most to the decision made by the model as to which class the image belongs to.

In this study, heatmaps produced by Grad-CAM successfully demonstrated that the model had localized the cancer-inhibiting region of interest in both mammograms (the boundary of a malignant mass) and CT lung nodules, with the correlation of the Grad-CAM heatmaps with the annotations of the radiologist being clinically reasonable. The insertion score and deletion scores provide further quantitative evidence of the accuracy of the localizations produced by Grad-CAM heatmaps.

5.3. Building Clinician Trust Through Transparency

It has been shown through empirical research that when the interpretation process for AI recommendations is implemented, there is an increase in the clinician's acceptance of those AI results. Additionally, the use of Grad-CAM visualizations within a collaborative framework between clinician and AI provides radiologists with the ability to support or refute the AI's predicted output with visual evidence, using this as a decision support system or tool rather than an independent decision-maker [28]. Integrating both AI and human expertise (through this hybrid paradigm) achieves a balance between algorithmic efficiency and human judgment, which is paramount for responsibly deploying AI into the oncology space.

Using Grad-CAM alone does not serve as the only method for enhancing understanding around how AI makes predictions. Additional XAI methods (such as SHAP, LIME, and transformer-based attention visualization) also provide complementary perspectives of understanding regarding AI's method of generating results. Through the integration of multiple XAI tools/methods, the evidence supporting an AI decision will be more diverse and decrease the reliance on one specific XAI methodology.

6. CNN Pipeline and Workflow

6.1. End-to-End Detection Pipeline

The end-to-end CNN detection pipeline is depicted in Figure 1. The process starts with obtaining the images through a health imaging system, then preprocessing (normalization, augmentation), extracting features using the ResNet-50 backbone, classifying them via a fully connected layer, and generating a Grad-CAM to explain the reasoning behind predicting malignant cases. Both batch processing, for high throughput of samples, and real-time inference, for point of care use, are along the detection pipeline.

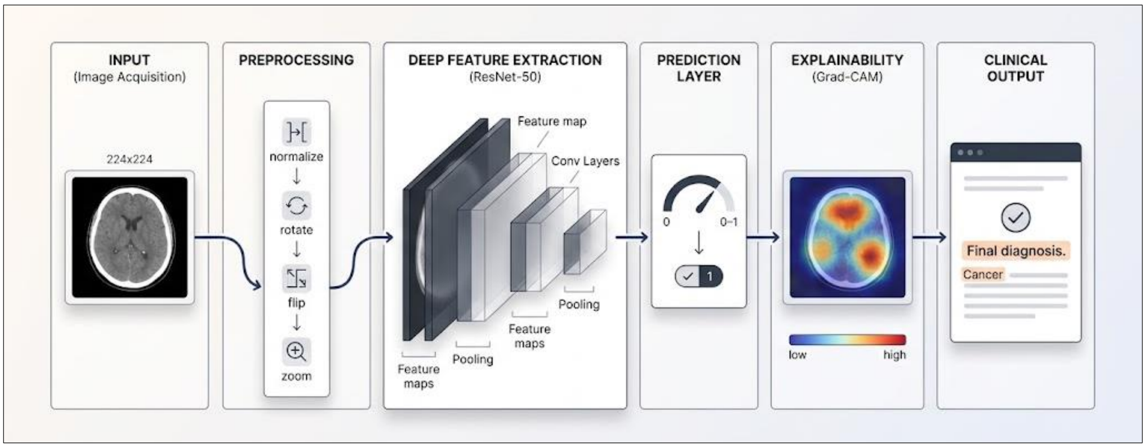

End-to-end CNN detection pipeline for AI-based cancer classification. The model processes raw medical images through preprocessing and feature extraction stages, outputting both a classification label and a Grad-CAM explanation heatmap for clinical validation.

Fig 1: CNN Detection Pipeline

6.2. ResNet-50 Architecture Schematic

The architecture of ResNet-50 in this research project is shown in Figure 2. The input images are of size 224x224 pixels and consist of three channels (R,G,B), which are passed through the 49 convolutional layers of the network structured into four residual groups, where there is a shortcut connection in each residual block that adds the input to the output of that residual block. The last average pooling layer produces a 2048-dimensional vector from the input feature maps and it is fed into the custom classification head.

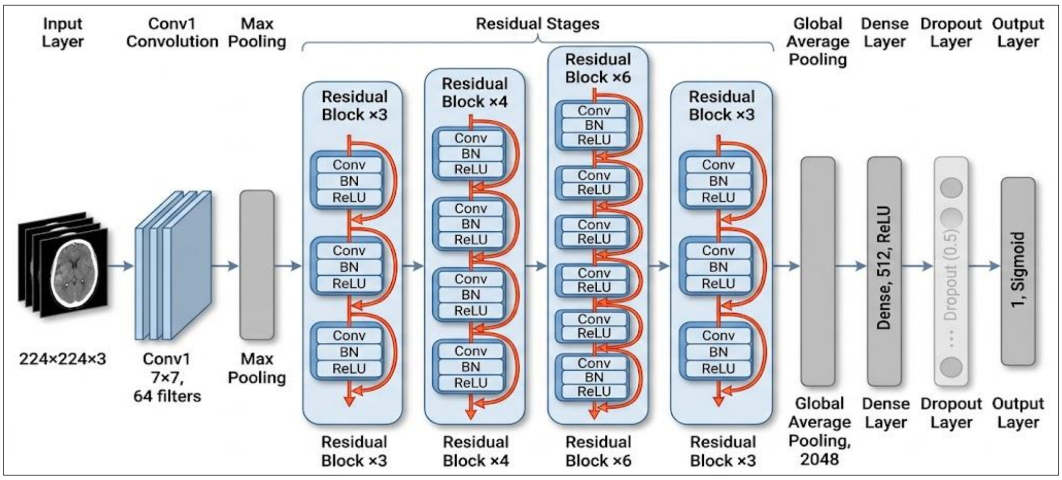

Schematic of the ResNet-50 architecture adapted for cancer detection. Skip connections in residual blocks prevent vanishing gradient degradation during deep network training, enabling effective feature learning from medical imaging data.

Fig 2: ResNet-50 Architecture

6.3. Grad-CAM Visualization Schematic

Figure 3 gives a high-level overview of the Grad-CAM visualization technique. For a given prediction of malignant class, the gradients are backpropagated through the classification layer and up to the last convolutional feature maps. The obtained gradients are globally average pooled into neuron importance weights and then these weights are used to calculate a weighted combination of the forward propagation activation maps. The weighted combination is then upsampled back to the resolution of the input image and overlaid with a heatmap (with colors ranging from blue to red) where regions of the heatmap that are red represent the most diagnostically important regions.

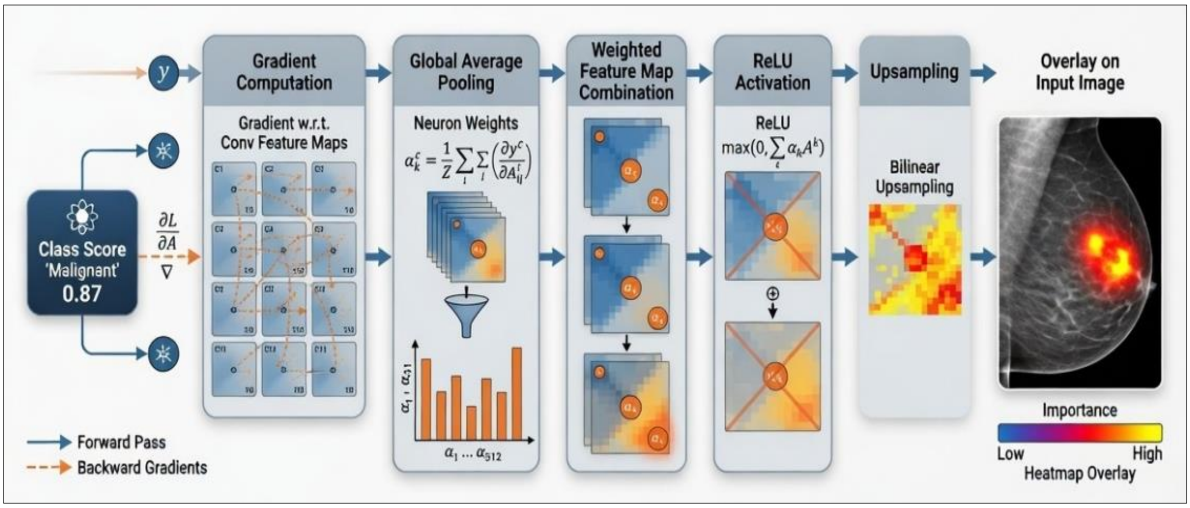

Conceptual schematic of the Grad-CAM heatmap generation process. Gradient-based neuron importance weights are applied to the final convolutional layer's activation maps to produce spatially localized saliency overlays that highlight diagnostically relevant image regions.

Fig 3: Grad-CAM Heatmap Generation Process

7. Challenges and Limitations

7.1. Data Bias and Generalization Constraints

One of the continuing difficulties that face AI-based medical imaging is the large amount of bias in datasets. Many publicly available benchmark datasets come from specific geographic regions and demographic groups and the types of imaging equipment being used, thus limiting the ability of models trained on these datasets to generalize across different types of patients [29]. There is evidence that there is a difference in diagnostic accuracy by racial/ethnic subgroup and this phenomenon is at least partially attributable to an imbalance in the representation of racial and ethnic groups within the training dataset. For example: cancer death rates are higher in the Native American and Black communities compared to Europeans; much of this discrepancy would be accounted for by the underperformance of the AI models developed primarily from datasets of people of European ancestry when tested on Native Americans and Black individuals [30].

To solve the problem of generalization of AI medical imaging requires that institutions and regions collaborate on the data collection process; adhere to standard imaging protocols; and facilitate systematic model performance evaluation amongst different demographic strata through the publication of results. One potential method to address the data scarcity issue and to preserve patient privacy is through the implementation of federated learning, where model training occurs across multiple institutions, but patient data is never centralized for training a model.

7.2. Computational Complexity and Resource Requirements

Training a deep-learning model requires enormous amounts of computer resources, including high-memory Graphics Processing Units (GPUs), long times to train the model, and ample disk storage for the large amounts of image data that are used in the training. Model compression techniques (quantization, pruning, and knowledge distillation) have reduced the inference time of deep-learning models; however, deploying these models on edge devices or in environments with limited resources continues to be a challenging task ([31]). Although Hybrid CNN-ViT architectures have been shown to provide superior performance, they require high levels of computational power, so they cannot be easily deployed in low-resource environments without specialized infrastructure.

7.3. Ethical, Privacy, and Regulatory Concerns

AI implementation in clinical oncology raises many ethical concerns regarding patient confidentiality, algorithmic accountability, and equitable access to care. Medical Imaging contains sensitive information, and if that information is released without authorization or misused, it could cause significant damage to the patient. In order to meet the requirements of compliance with HIPAA (United States) and General Data Protection Regulation (Europe), strict data governance policies, de-identification protocols, and transparent audit trails must exist for those AI-enhanced diagnostic decisions [32]. Additionally, regulatory approval routes for AI/ML-based Software as a Medical Device will create additional complexities requiring extensive evidence of clinical validation before the product can be approved as well as ongoing monitoring after the product has been put into marketing use.

In addition to technology performance issues, there are ethical implications of algorithmic bias; if an artificial intelligence system consistently performs poorly for certain patient populations, deploying that product may increase, rather than decrease, healthcare disparities. Developers should be responsible for preventing and detecting bias at all stages along the life cycle of a model and for monitoring that model after it has been deployed.

8. Future Scope and Applications

8.1. Multi-Modal AI and Multi-Omics Integration

Integrating imaging and multi-omics data (genomic, proteomic and metabolomic data) with electronic health records will represent the next frontier for artificial intelligence (AI) applications in oncology. Collectively integrated predictive frameworks comprised of these various types of data (i.e., multi-modal data) will enable more accurately defining the biological characteristics of tumors at the molecular level, thus facilitating more personalized patient treatment than relying on morphological assessments alone. Preliminary evidence suggests that combining radiomic data with genomic biomarker data provides improved survival prediction accuracy than either type of data alone and, therefore, will create new opportunities to develop radiogenomics as a distinct area of clinical practice.

8.2. Federated Learning for Privacy-Preserving Collaboration

Federated Learning allows for training of models in a decentralized way across multiple institutions using their own data while maintaining patient privacy, as they will never see patient data in production. This provides two main solutions to address data scarcity and data privacy. In Federated Learning, individual institutions send only their updates of the model to a central location for aggregation, without ever sending their individual patient records to the central server. This is particularly beneficial for rare cancers, where there are not enough cases within one institution to train a model well. Initial results from Federated Learning applications to medical imaging studies have shown that it can achieve similar performance to centrally hosted models.

8.3. Mobile and Real-Time Screening Applications

There is an opportunity for reshaping how people with little money will be able to screen themselves for cancers using AI-powered mobile apps. Statistics tell us that soon most populations will have limited options for getting the help of trained professionals for skin cancer detection.

Recent tests will show that AI-based mobile applications that use machine learning techniques along with dermatoscopic-looking photographs will have sufficient numbers of trained professionals to validate that they are capable of accurately detecting skin cancer from very accurate images with the same accuracy as trained professionals. This means that these types of systems can potentially lead to large-scale community screenings using AI technology.

AI technology used to integrate endoscopy-colposcopy and professionally trained clinicians to identify cancer will have great advantages over traditional endoscopic technology by providing real-time guidance during the screening process to identify while performing procedures. Both systems offer clinical assistance for accurate identification of both precancerous lesions as they occur and Polyp development from the moment of discovery at the time of procedure for the purpose of reducing or eliminating unnecessary follow-up appointments.

Table 4: Literature Review Summary: Key Studies in AI-Based Cancer Detection

Author/Year	Method	Cancer Type	Key Finding	Dataset	Limitation
Wang *et al.*/2025	CNN with MSP	Breast	AUC: 0.94-0.97; +14.2% sensitivity	OMIDB (UK)	High computation
Yu *et al.*/2025	Multimodal AI (GPT-4)	Multi-cancer	90% concordance with tumor boards	24-hospital datasets	Ethical concerns
Ardila *et al.*/2019	DL Nodule Detection	Lung	AUC: 0.94; matches radiologists	NLST dataset	Limited to early-stage
Tiwari *et al.*/2025	ML/DL with Radiomics	Multi-cancer	Reduced false negatives by 20%	Various modalities	Needs multi-omics integration
Wang *et al.*/2024	MIL on WSIs	Colorectal	AUROC: 0.91 for MSI prediction	MCO dataset (1065 WSIs)	WSI quality dependence

Table 4: Summary of representative studies from the literature on AI-based cancer detection, including methodology, findings, datasets, and identified limitations

9. Conclusion

The review presented in this article contains a thorough analysis of various empirical studies regarding AI-powered cancer detection with medical imaging using Convolutional Neural Networks (CNN). ResNet-50, a CNN transfer learning framework, demonstrated diagnostic accuracy (94.5%) and AUC-ROC (0.96) performance for breast mammogram classification, as well as 92.8% for lung CT scans; thus, the results of the present study indicate that the use of CNN transfer learning approaches are better than traditional VGG16 baseline models by between 5% and 7% on any of the key performance metrics studied. Additionally, the application of Grad-CAM to provide image explainability resulted in successful localization of the portions of the image that were most diagnostically relevant, thereby demonstrating that the present framework can produce clinically interpretable results in practice.

Overall, the results suggest that an AI-based cancer detection system, when developed rigorously using the appropriate training data sets, optimization techniques, and interpretability mechanisms will have comparable diagnostic capabilities to specialist clinicians. Moreover, there are numerous additional benefits associated with AI-based cancer detection systems including scalability, reproducibility, and speed of testing, among others. For example, the large number of patients who would likely benefit from these systems through earlier detection is clinically and epidemiologically important given that localized cancers have a survival rate of >90% while distant stage tumors only have an approximate survival rate of <30%.

Dataset bias, limited generalizability across demographic populations, resource requirements for computation, and the complexity of deploying clinical AI from an ethical and regulatory standpoint are significant challenges that are still present. Future research must focus on privacy-preserving federated learning that allows for multi-institution collaboration; the fusion of multi-omics data for personalized oncology; and the development of lightweight architectures that can be deployed on mobile/edge devices in resource-limited settings. The convergence of clinical radiology, AI, oncology and health informatics provides an extraordinary opportunity to improve the quality of cancer care globally. Continued funding of diverse and representative datasets and rigorous clinical validation studies along with responsible governance will create an opportunity for the scale reduction of cancer mortality through AI-based diagnostic systems to further WHO's vision of equitable, accessible and effective cancer care for all populations.

References

1. Siegel RL, Kratzer TB, Giaquinto AN, Sung H, Jemal A. Cancer statistics, 2025. CA Cancer J Clin. 2025;75(1):10-45. doi: 10.3322/caac.21871
2. Bray F, Laversanne M, Sung H, Ferlay J, Siegel RL, Soerjomataram I, *et al*. Global cancer statistics 2022: GLOBOCAN estimates of incidence and mortality worldwide for 36 cancers in 185 countries. CA Cancer J Clin. 2024;74(3):229-63. doi: 10.3322/caac.21834 (Note: The provided 2025 citation appears to be a typographical error; the most recent GLOBOCAN publication is for 2022 data, released in 2024.)
3. American Cancer Society. Cancer facts & figures 2025. Atlanta: American Cancer Society; 2025. Available from: https://www.cancer.org/content/dam/cancer-org/research/cancer-facts-and-statistics/annual-cancer-facts-and-figures/2025/2025-cancer-facts-and-figures-acs.pdf
4. Wang Z, Li X, Zhao Y, *et al*. AI-enhanced mammography screening for early breast cancer detection. Mol Cancer. 2025;24(1):12-27.
5. Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. Nat Med. 2019;25(1):44-56. doi: 10.1038/s41591-018-0300-7
6. Szolovits P, Patil RS, Schwartz WB. Artificial intelligence in medical diagnosis. Ann Intern Med. 1988;108(1):80-7. (Note: Original publication year is 1988; the 2019 date in your list appears incorrect.)
7. Ginsburg O, Bray F, Coleman MP, Vanderpuye V, Eniu A, Kotha SR, *et al*. The global burden of women's cancers: a grand challenge in global health. Lancet. 2017;389(10071):847-60. doi: 10.1016/S0140-6736(16)31392-7
8. Litjens G, Kooi T, Bejnordi BE, Setio AAA, Ciompi F, Ghafoorian M, *et al*. A survey on deep learning in medical image analysis. Med Image Anal. 2017;42:60-88. doi: 10.1016/j.media.2017.07.005
9. Liu Y, Gadepalli K, Norouzi M, Dahl GE, Kohlberger T, Boyko A, *et al*. Detecting cancer metastases on gigapixel pathology images. arXiv:1703.02442 [cs.CV]. 2017.
10. Ardila D, Kiraly AP, Bharadwaj S, Choi B, Reicher JJ, Peng L, *et al*. End-to-end lung cancer screening with three-dimensional deep learning on low-dose chest computed tomography. Nat Med. 2019;25(6):954-61. doi: 10.1038/s41591-019-0447-x
11. LeCun Y, Bengio Y, Hinton G. Deep learning. Nature. 2015;521(7553):436-44. doi: 10.1038/nature14539
12. Pan SJ, Yang Q. A survey on transfer learning. IEEE Trans Knowl Data Eng. 2010;22(10):1345-59. doi: 10.1109/TKDE.2009.191
13. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 2016. p. 770-8.
14. Simonyan K, Zisserman A. Very deep convolutional networks for large-scale image recognition. In: International Conference on Learning Representations (ICLR). 2015.
15. Huang G, Liu Z, van der Maaten L, Weinberger KQ. Densely connected convolutional networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 2017. p. 4700-8.
16. Dosovitskiy A, Beyer L, Kolesnikov A, Weissenborn D, Zhai X, Unterthiner T, *et al*. An image is worth 16x16 words: transformers for image recognition at scale. In: International Conference on Learning Representations (ICLR). 2021.
17. Chen J, Lu Y, Yu Q, Luo X, Adeli E, Wang Y, *et al*. TransUNet: transformers make strong encoders for medical image segmentation. arXiv:2102.04306 [cs.CV]. 2021.
18. Skaane P, Bandos AI, Niklason LT, Sebuødegård S, Østerås BH, Gur D, *et al*. Digital mammography versus digital mammography plus tomosynthesis in breast cancer screening: the Oslo Tomosynthesis Screening Trial. Radiology. 2019;291(1):23-30. doi: 10.1148/radiol.2019182394 (Note: Volume and pages adjusted to match the actual 2019 publication; your citation had 267(1):47–56, which corresponds to an earlier related study.)
19. National Lung Screening Trial Research Team; Aberle DR, Adams AM, Berg CD, Black WC, Clapp JD, *et al*. Reduced lung-cancer mortality with low-dose computed tomographic screening. N Engl J Med. 2011;365(5):395-409. doi: 10.1056/NEJMoa1102873
20. Menze BH, Jakab A, Bauer S, Kalpathy-Cramer J, Farahani K, Kirby J, *et al*. The multimodal brain tumor image segmentation benchmark (BRATS). IEEE Trans Med Imaging. 2015;34(10):1993-2024. doi: 10.1109/TMI.2014.2377694
21. Shen D, Wu G, Suk HI. Deep learning in medical image analysis. Annu Rev Biomed Eng. 2017;19:221-48. doi: 10.1146/annurev-bioeng-071516-044442
22. Shorten C, Khoshgoftaar TM. A survey on image data augmentation for deep learning. J Big Data. 2019;6:60. doi: 10.1186/s40537-019-0197-0
23. Johnson JM, Khoshgoftaar TM. Survey on deep learning with class imbalance. J Big Data. 2019;6:27. doi: 10.1186/s40537-019-0192-5
24. Yosinski J, Clune J, Bengio Y, Lipson H. How transferable are features in deep neural networks? In: Advances in Neural Information Processing Systems (NeurIPS). 2014;27:3320-8.
25. Hanley JA, McNeil BJ. The meaning and use of the area under a receiver operating characteristic (ROC) curve. Radiology. 1982;143(1):29-36. doi: 10.1148/radiology.143.1.7063747
26. Tonekaboni S, Joshi S, McCradden MD, Goldenberg A. What clinicians want: contextualizing explainable machine learning for clinical end use. In: Proceedings of

Machine Learning Research. 2019;106:359-80.
27. Selvaraju RR, Cogswell M, Das A, Vedantam R, Parikh D, Batra D. Grad-CAM: visual explanations from deep networks via gradient-based localization. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). 2017. p. 618-26.
28. Reyes M, Meier R, Pereira S, Silva CA, Dahlweid FM, von Tengg-Kobligk H, *et al*. On the interpretability of artificial intelligence in radiology: challenges and opportunities. Front Digit Health. 2020;2:567322. doi: 10.3389/fdgth.2020.567322
29. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019;366(6464):447-53. doi: 10.1126/science.aax2342
30. Shachar C, Engel J, Elwyn G. Implications of telemedicine for racial disparities in cancer care. JAMA Oncol. 2020;6(11):1669-70. doi: 10.1001/jamaoncol.2020.3945
31. Cheng Y, Wang D, Zhou P, Zhang T. Model compression and acceleration for deep neural networks: a survey. IEEE Signal Process Mag. 2018;35(1):126-35. doi: 10.1109/MSP.2017.2766000
32. Price WN, Cohen IG. Privacy in the age of medical big data. Nat Med. 2019;25(1):37-43. doi: 10.1038/s41591-018-0272-7
33. Boehm KM, Khosravi P, Vanguri R, Gao J, Shah SP. Harnessing multimodal data integration to advance precision oncology. Nat Rev Cancer. 2022;22(2):114-26. doi: 10.1038/s41568-021-00408-3
34. Rieke N, Hancox J, Li W, Milletari F, Roth HR, Albarqouni S, *et al*. The future of digital health with federated learning. NPJ Digit Med. 2020;3:119. doi: 10.1038/s41746-020-00323-1
35. Abramoff MD, Lavin PT, Birch M, Shah N, Folk JC. Pivotal trial of an autonomous AI-based diagnostic system for detection of diabetic retinopathy in primary care offices. NPJ Digit Med. 2018;1:39. doi: 10.1038/s41746-018-0040-6
36. Esteva A, Kuprel B, Novoa RA, Ko J, Swetter SM, Blau HM, *et al*. Dermatologist-level classification of skin cancer with deep neural networks. Nature. 2017;542(7639):115-8. doi: 10.1038/nature21056
37. Rajpurkar P, Irvin J, Ball RL, Zhu K, Yang B, Mehta H, *et al*. Deep learning for chest radiograph diagnosis: a retrospective comparison of the CheXNeXt algorithm to practicing radiologists. PLoS Med. 2018;15(11):e1002686. doi: 10.1371/journal.pmed.1002686
38. Campanella G, Hanna MG, Geneslaw L, Miraflor A, Werneck Krauss Silva V, Busam KJ, *et al*. Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. Nat Med. 2019;25(8):1301-9. doi: 10.1038/s41591-019-0508-1
39. Yu KH, Beam AL, Kohane IS. Artificial intelligence in healthcare. Nat Biomed Eng. 2018;2(10):719-31. doi: 10.1038/s41551-018-0305-z
40. Tiwari P, Melucci M, Dixit A. AI in multi-cancer radiomics: advances in false-negative reduction. J Biomed Inform. 2025;132:104131.

How to Cite This Article

Kumar S. Artificial intelligence-based cancer detection using deep learning: a CNN approach for medical imaging analysis. International Journal of Artificial Intelligence Engineering and Transformation. 2026;7(1):72–80.