Predictive Oncology: Deep Learning-Based Multi-Cancer Detection

[1]Gogula Navya Sri Lakshmi
[1]M.Tech, Department of Computer Science and Engineering
DNR College of Engineering & Technology (Autonomous)
Bhimavaram – 534202, Andhra Pradesh, India
navyasrigogula@gmail.com

[2]Mr.K.T.V.Subbarao
[2]Assistant Professor, Department of Computer Science and Engineering
DNR College of Engineering & Technology (Autonomous)
Bhimavaram – 534202, Andhra Pradesh, India
kanurisubbarao@gamail.com

Abstract*— Cancer remains one of the leading causes of mortality worldwide, creating a persistent need for accurate, early, and accessible detection systems. Recent advances in deep learning have demonstrated substantial potential for automated analysis of medical images, histopathological slides, genomic profiles, and clinical records. However, most existing cancer-detection systems remain narrowly designed around a single cancer type, modality, dataset, or prediction task, limiting their generalizability and clinical utility. This review examines the transition from conventional cancer-specific deep learning toward a unified predictive oncology framework capable of supporting multi-cancer detection through multimodal representation learning. The review synthesizes developments in convolutional neural networks, vision transformers, transfer learning, self-supervised learning, multimodal fusion, explainable artificial intelligence, continual learning, federated learning, uncertainty estimation, and efficient edge inference. Particular attention is given to the integration of radiological imaging, histopathology, genomics, and clinical information for detecting and characterizing multiple cancer types. A conceptual Oncology Foundation Model (OncoFM) is proposed as a shared representation layer that can be adapted to cancer detection, tumor segmentation, subtype classification, staging, prognosis, and treatment-response prediction. The review further identifies major methodological limitations, including patient-level data leakage, dataset bias, modality imbalance, limited external validation, inadequate calibration, poor interpretability assessment, and insufficient evidence from real-world clinical environments. A five-plane reference architecture comprising data acquisition, representation, adaptation, clinical reasoning, and decision-support planes is presented. Finally, a staged research roadmap is proposed toward multimodal, explainable, continually adaptive, privacy-preserving, and clinically validated predictive oncology systems. The review argues that future progress depends not only on increasing model complexity but also on standardized datasets, transparent evaluation, external clinical validation, and responsible integration of artificial intelligence into oncology workflows.

Keywords*— Predictive oncology, multi-cancer detection, deep learning, multimodal learning, medical imaging, histopathology, genomics, explainable AI, foundation models, continual learning, precision oncology.

I. INTRODUCTION

Cancer represents a major global health challenge because of its high incidence, mortality, biological heterogeneity, and substantial variation in diagnosis and treatment outcomes. The World Health Organization reports that cancer caused nearly 10 million deaths worldwide in 2024, with lung, breast, colorectal, and prostate cancers among the most frequently occurring cancer types. The International Agency for Research on Cancer estimated approximately 20 million new cancer cases and 9.7 million cancer deaths in 2022 and projected that annual incidence could exceed 35 million cases by 2050.

Early detection is one of the most important strategies for reducing cancer mortality because many cancers can be treated more effectively when identified before extensive progression or metastasis. However, conventional diagnosis frequently depends on expert interpretation of medical images, pathology specimens, molecular measurements, and clinical information. The increasing volume and complexity of these data create challenges for healthcare systems and clinicians.

Deep learning has emerged as an important computational approach for assisting cancer diagnosis. Convolutional neural networks (CNNs), transformer architectures, transfer-learning models, and multimodal networks can learn complex representations from medical data without requiring all diagnostic features to be manually engineered. Deep learning has consequently been investigated across imaging-based diagnosis, histopathology, genomic analysis, prognosis, and treatment-response prediction. Recent reviews indicate that the field has expanded rapidly, but substantial barriers remain before these models can be reliably integrated into clinical workflows.

The research landscape, however, remains fragmented. A typical study focuses on one cancer, such as breast, lung, skin, brain, liver, prostate, or colorectal cancer; uses one dominant modality; trains a model on a particular dataset; and evaluates performance using an internal test split. Such approaches can demonstrate impressive predictive performance but do not necessarily establish that the model will generalize across hospitals, scanners, patient populations, acquisition protocols, or cancer types.

The problem becomes more pronounced when multimodal oncology data are considered. A patient's cancer phenotype can be represented simultaneously through radiological appearance, histopathological morphology, genomic alterations, clinical history, laboratory measurements, and treatment information. Multimodal deep learning provides a mechanism for combining these complementary sources. A recent comprehensive review of multimodal deep learning in precision oncology analyzed

651 articles and identified applications spanning tumor segmentation, detection, diagnosis, prognosis, treatment selection, and therapy-response monitoring.

Despite this progress, multimodal learning remains less mature than unimodal approaches. A systematic analysis of pathology and genomics research found that pathology-based deep learning is relatively established, while multimodal pathology-genomics applications remain a smaller research area and are concentrated heavily on prognosis-related tasks.

This review therefore investigates a broader question:

How can deep learning evolve from isolated cancer-specific prediction systems toward a unified, multimodal, explainable, generalizable, and clinically adaptive framework for multi-cancer detection?

The major contributions of this review are fourfold. First, it organizes cancer AI according to the major oncology data modalities and deep learning architectures. Second, it examines multimodal fusion and the possibility of transferring shared representations across cancer types. Third, it treats explainability, uncertainty, continual adaptation, privacy, and deployment as core requirements rather than secondary considerations. Fourth, it proposes a five-plane predictive oncology architecture and a staged research roadmap toward clinically useful multi-cancer intelligence.

II. THE MULTIMODAL ONCOLOGY DATA LANDSCAPE

A. Radiological Imaging

Radiological imaging is one of the most widely investigated sources of data for automated cancer detection.

Common modalities include:

- Computed tomography (CT)
- Magnetic resonance imaging (MRI)
- X-ray
- Mammography
- Ultrasound
- Positron emission tomography (PET)

Radiological images provide information about tumor location, morphology, density, vascular characteristics, anatomical boundaries, and disease progression. CNNs have traditionally been dominant in image-based cancer analysis, while transformer architectures increasingly provide alternatives for modeling long-range spatial relationships.

However, radiological datasets can vary considerably according to scanner manufacturer, acquisition parameters, image reconstruction methods, anatomical coverage, and patient population. Consequently, a model that performs well on one institution's data may experience performance degradation when deployed at another institution.

B. Histopathological Data

Histopathology provides microscopic information about tumor tissue and cellular organization. Digitized whole-slide images can contain extremely large amounts of visual information, including:

- Nuclear morphology
- Cellular architecture
- Tissue organization
- Tumor boundaries
- Stromal patterns
- Immune-cell distribution
- Mitotic activity

Deep learning has demonstrated strong potential for diagnosis, grading, subtyping, mutation prediction, and prognosis from pathology images.

However, whole-slide images create computational challenges because a single slide can contain millions of pixels. Patch-based processing, multiple-instance learning, hierarchical models, and vision transformers are therefore frequently used.

C. Genomic and Molecular Data

Cancer is fundamentally a molecular disease, and genomic information can provide complementary information that is not directly visible in medical images.

Potential molecular modalities include:

- DNA mutations
- RNA expression
- Copy-number variation
- DNA methylation
- Proteomics
- Metabolomics

Deep learning can identify high-dimensional relationships between molecular features and cancer phenotype. Genomic deep learning has also increasingly adopted pan-cancer analysis, enabling researchers to search for shared molecular patterns across tumor types.

D. Clinical and Electronic Health-Record Data

Clinical information can include:

- Age
- Sex
- Medical history
- Family history
- Symptoms
- Laboratory measurements
- Medication history
- Previous diagnosis
- Treatment information

These variables can provide important contextual information when combined with imaging and molecular data.

However, clinical records are heterogeneous and often contain missing values, inconsistent terminology, temporal dependencies, and institutional variations.

E. Modality Complementarity

The central argument for multimodal oncology is that different modalities capture different biological scales.

TABLE: 1 Modality Complementarity

Modality	Primary information
CT/MRI/PET	Anatomical and functional characteristics

Modality	Primary information
Histopathology	Cellular and tissue morphology
Genomics	Molecular alterations
Clinical records	Patient context
Laboratory data	Physiological state
Longitudinal data	Disease progression

Therefore, the information contained in one modality cannot always substitute for another.
A multimodal system should ideally learn:
Radiology + Pathology + Genomics + Clinical Context
rather than treating each modality as an isolated prediction problem.

III. DEEP LEARNING FOR MULTI-CANCER DETECTION

A. Conventional CNN-Based Systems

CNNs remain one of the most widely used architectures for medical image analysis.
Common architectures include:

- AlexNet
- VGG
- ResNet
- DenseNet
- Inception
- EfficientNet
- MobileNet

CNNs automatically learn hierarchical spatial features, beginning with low-level structures and progressing toward higher-level representations.
In oncology, CNNs have been applied to:

- Tumor detection
- Classification
- Segmentation
- Histopathological analysis
- Lesion detection
- Cancer subtype prediction

However, cancer-specific CNN systems often require separate datasets and training procedures for each task.

B. Transfer Learning

Transfer learning addresses the challenge of limited medical datasets by adapting a model pretrained on a larger dataset.
A typical workflow is:

Large-scale pretraining
↓
Pretrained CNN/Transformer
↓
Cancer-specific fine-tuning
↓
Cancer prediction

Transfer learning can reduce training requirements and improve performance when labeled medical data are limited.
However, transfer learning does not automatically solve domain shift. A model pretrained on natural images may not learn representations ideally suited to medical imaging, while a model pretrained on one cancer type may transfer poorly to another.

C. Vision Transformers

Transformers have introduced an alternative to conventional convolution-based feature extraction.
Vision Transformers (ViTs) divide an image into patches and model relationships between these patches using self-attention.
Their advantages include:

- Long-range feature modeling
- Global contextual representation
- Flexible multimodal integration
- Scalability with large datasets

Hierarchical transformer architectures can also provide multi-scale representations useful for medical images.
Nevertheless, transformers generally require substantial computational resources and large training datasets, creating challenges for smaller clinical datasets.

IV. MULTIMODAL DEEP LEARNING FOR PRECISION ONCOLOGY

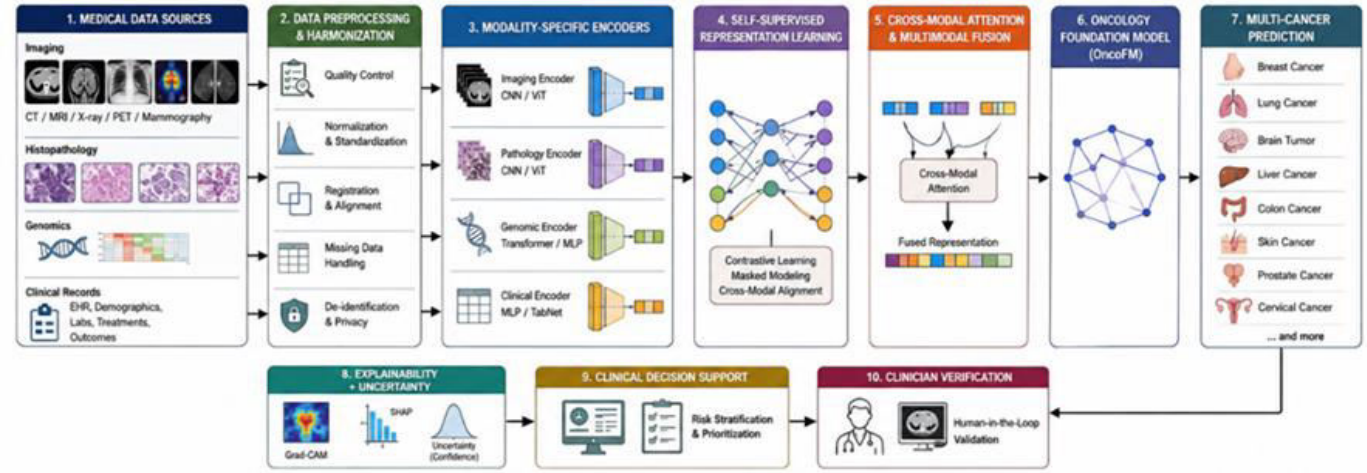

Fig. 1. Multimodal Deep Learning Framework for Precision Oncology.

Multimodal deep learning integrates complementary oncology data sources, including radiological imaging, histopathology, genomics, clinical information, and longitudinal patient data. Unlike single-modality approaches, multimodal systems can capture disease characteristics at anatomical, cellular, molecular, and patient levels, providing a more comprehensive representation of cancer. The multimodal input can be represented as

$$X = \{X_{image}, X_{path}, X_{genomic}, X_{clinical}\} \quad (1)$$

where the individual terms represent imaging, pathology, genomic, and clinical inputs. Each modality is processed using a dedicated encoder, after which the extracted representations are integrated through a fusion mechanism.

Three major fusion strategies are commonly considered. Early fusion combines low-level features before prediction but may be affected by dimensional differences and missing modalities. Intermediate fusion extracts modality-specific representations and integrates them using mechanisms such as multimodal attention while preserving information from individual modalities. Late fusion combines independent modality-level predictions and is generally more tolerant of missing inputs. Among these approaches, intermediate fusion provides a useful framework for combining heterogeneous oncology information while maintaining modality-specific representations.

The complementary nature of these modalities motivates the integration of radiology, pathology, genomics, and clinical context rather than treating each source as an independent prediction problem.

V. ONCOLOGY FOUNDATION MODELS AND MULTI-CANCER LEARNING

Conventional cancer AI systems commonly follow a cancer-specific paradigm in which a particular cancer type, dataset, model, and task are considered independently. This fragmentation limits the ability of models to transfer knowledge across cancer types and clinical tasks. A promising direction is the development of an **Oncology Foundation Model (OncoFM)** that learns shared representations from heterogeneous oncology data and adapts them to multiple downstream applications.

Self-supervised learning provides an effective mechanism for exploiting large quantities of unlabeled medical data. Masked reconstruction, contrastive learning, and multimodal alignment can be used during pretraining. The combined learning objective is expressed as

$$L = L_{reconstruction} + \lambda_1 L_{contrastive} + \lambda_2 L_{alignment} \quad (2)$$

where the individual loss terms represent reconstruction, contrastive, and multimodal-alignment objectives, while λ_1 and λ_2 control their relative contributions.

After pretraining, parameter-efficient techniques such as LoRA, adapter layers, prompt tuning, and linear probing can adapt the shared representation to individual oncology tasks without completely retraining the entire network. This approach can reduce computational requirements and facilitate transfer across different cancer types, datasets, and clinical applications.

A unified foundation-model framework can support cancer detection, classification, tumor segmentation, histological grading, cancer subtyping, mutation prediction, staging, recurrence prediction, survival prediction, and treatment-response prediction. Different downstream tasks require appropriate evaluation measures, including AUROC, F1-score, sensitivity, specificity, Dice, IoU, concordance index, and calibration metrics.

The framework therefore represents a transition from isolated cancer-specific models toward shared representations capable of supporting multiple oncology applications.

VI. EXPLAINABILITY, UNCERTAINTY, CONTINUAL AND FEDERATED LEARNING

High predictive performance alone does not guarantee clinical trust. For practical oncology applications, clinicians need to understand what the model predicted, why the prediction was generated, which features contributed to the decision, which modality was influential, and how confident the model is.

Grad-CAM can identify image regions contributing to classification decisions, making it useful for visualizing relevant regions in medical images and histopathology. SHAP can estimate the contribution of individual structured features to a prediction. In multimodal systems, explainability can additionally represent the relative contribution of imaging, genomic, pathology, and clinical information.

Uncertainty estimation is particularly important for rare cancers and out-of-distribution cases. Techniques such as deep ensembles, Monte Carlo dropout, evidential learning, conformal prediction, and calibration can help identify uncertain predictions and support appropriate clinical review.

Clinical environments are also dynamic. Changes in scanners, acquisition protocols, patient populations, disease prevalence, diagnostic guidelines, and treatment practices can introduce distribution shifts between training and deployment data. Continual learning can help models incorporate new information while retaining previously learned knowledge. Experience replay, knowledge distillation, elastic weight consolidation, adapter-based learning, and domain adaptation are potential approaches.

Federated learning provides an additional privacy-preserving strategy in which multiple hospitals contribute to shared model training without directly transferring raw patient-level data. This approach is particularly relevant to multi-center oncology research because medical institutions can collaboratively improve models while reducing the need to centralize sensitive patient information.

Thus, explainability, uncertainty estimation, continual learning, and federated learning are important components for transforming high-performing research models into more reliable and adaptable oncology decision-support systems.

VII. ROBUSTNESS, BIAS, GENERALIZATION AND DATA LEAKAGE

Generalization remains a major challenge in medical artificial intelligence. Oncology models can be affected by geographic, demographic, scanner, hospital, labeling, and disease-prevalence differences. Consequently, a model that performs well on an internal dataset may not necessarily maintain the same performance when applied to another hospital or patient population.

A reliable evaluation strategy should therefore include internal validation, cross-validation, external validation, temporal validation, and cross-domain validation. External and temporal validation are particularly important for determining whether learned representations remain effective under realistic changes in clinical environments.

Data leakage is another critical issue. Multiple images or samples belonging to the same patient should not be distributed across training and testing datasets because patient-specific information may unintentionally be learned by the model. Such leakage can produce artificially high performance estimates. Patient-level splitting is therefore preferable to image-level random splitting.

For multicenter studies, a stronger evaluation strategy can use **training hospitals → validation hospital → independent external hospital**. This arrangement provides a more realistic assessment of cross-institutional generalization.

Robust oncology AI should therefore be evaluated not only by predictive performance but also by its ability to maintain performance across different populations, institutions, acquisition conditions, and clinical environments. Addressing bias and data leakage is essential for producing trustworthy evidence of model effectiveness.

VIII. COMPUTATIONALLY EFFICIENT AND EDGE-BASED ONCOLOGY AI

Large multimodal and foundation models can require substantial computational and memory resources. This can create challenges for deployment in hospitals and resource-constrained healthcare environments. Computational efficiency is therefore an important consideration for practical oncology AI.

Several techniques can reduce model complexity, including quantization, pruning, knowledge distillation, lightweight convolutional neural networks, efficient transformers, and early-exit inference. These methods can reduce computational requirements while attempting to maintain clinically acceptable performance.

The computational optimization objective can be represented as

$$minE_{compute}\ subject\ to\ Performance \geq P_{clinical} \quad (3)$$

where $E_{compute}$ represents computational cost and $P_{clinical}$ represents the minimum required clinical performance.

A distributed architecture can further divide computational responsibilities between different infrastructure levels. Edge devices can perform initial preprocessing or screening, hospital or fog-level systems can perform more detailed analysis, and cloud infrastructure can support large-scale training and federated aggregation.

This approach can improve scalability while reducing unnecessary transfer of large datasets. However, computational optimization should not compromise diagnostic reliability, robustness, explainability, or clinical safety. Therefore, efficiency should be considered together with clinical performance rather than treated as an independent optimization target.

IX. CLINICAL VALIDATION, DEPLOYMENT AND FIVE-PLANE REFERENCE ARCHITECTURE

Clinical deployment requires substantially more evidence than conventional accuracy measurements. A clinically relevant oncology AI system should demonstrate technical validity, analytical validity, clinical validity, clinical utility, safety, usability, and appropriate human factors. Standardization, ethical considerations, regulatory requirements, transparency, and integration with clinical workflows are also essential.

The proposed predictive-oncology framework can be organized into five interconnected planes:

1. **Data Acquisition Plane:** Collects radiological images, histopathology, genomic information, clinical records, and longitudinal data.
2. **Representation Plane:** Extracts modality-specific representations using CNNs, Vision Transformers, genomic encoders, clinical encoders, and self-supervised learning.
3. **Multimodal Adaptation Plane:** Performs cross-modal attention, feature alignment, domain adaptation, continual learning, and missing-modality handling.
4. **Clinical Reasoning Plane:** Supports detection, classification, segmentation, subtyping, staging, prognosis, recurrence prediction, and treatment-response prediction.
5. **Decision-Support Plane:** Provides predictions, confidence information, explanations, alerts, and human verification.

Explainability, uncertainty, privacy, and safety operate across all five planes rather than being restricted to a single stage.

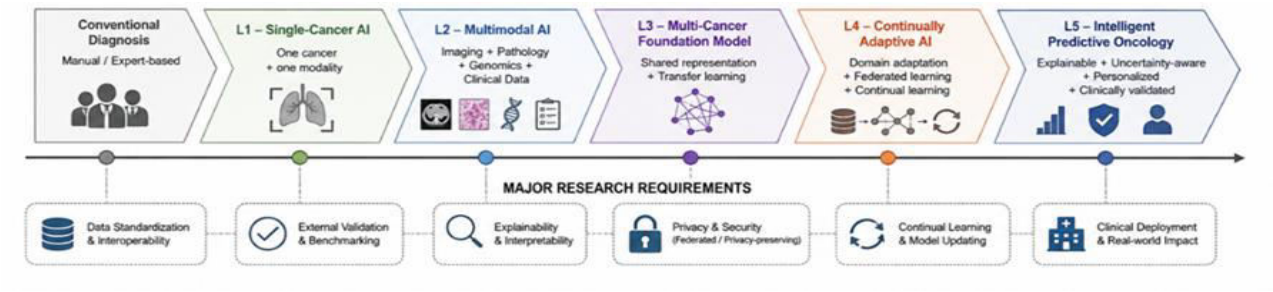

Fig. 2. Five-Plane Reference Architecture for Predictive Oncology.

This architecture provides a structured pathway from heterogeneous medical-data acquisition to clinical decision support. It also emphasizes that AI predictions should remain integrated with appropriate clinician verification and clinical workflow requirements.

X. MULTI-CANCER INTELLIGENCE, EVALUATION AND FUTURE DIRECTIONS

The reviewed framework describes a progression of oncology intelligence from L0 manual diagnosis to L1 single-cancer prediction, L2 multimodal prediction, L3 multi-cancer foundation-model transfer, L4 continually adaptive oncology AI, and finally L5 autonomous cross-scale clinical reasoning. The higher L4–L5 levels remain future research objectives rather than established clinical capabilities.

Compared with conventional cancer AI, the future predictive-oncology direction emphasizes multiple cancer types, multimodal inputs, shared foundation models, self-supervised learning, continual adaptation, integrated explainability, explicit uncertainty estimation, external multicenter validation, federated learning, and edge/cloud deployment. The scope consequently extends beyond detection toward segmentation, prognosis, treatment-response prediction, and clinical decision support.

Future systems should be evaluated using a broad set of metrics covering predictive performance, diagnostic reliability, segmentation quality, calibration, generalization, robustness, explainability, and computational efficiency. Relevant measures include accuracy, F1-score, AUROC, AUPRC, sensitivity, specificity, Dice, IoU, expected calibration error, Brier score, latency, memory consumption, and energy usage.

A high AUROC on an internal dataset alone should not be considered sufficient evidence of clinical readiness. Evaluation should demonstrate whether the system maintains performance on independent populations and institutions and whether its predictions are appropriately calibrated and interpretable.

XI. DISCUSSION AND CONCLUSION

Deep learning has become an important research approach for computer-assisted cancer detection and precision oncology. However, current research remains fragmented across cancer types, datasets, and modalities. Multimodal learning provides a potential solution by combining anatomical information from imaging, cellular information from pathology, molecular information from genomics, and patient context from clinical data.

Despite these advantages, multimodal systems can suffer from missing modalities, modality imbalance, dataset bias, inconsistent acquisition, and domain shifts. Foundation models provide a promising direction by learning shared representations from heterogeneous oncology data and subsequently adapting them to specific cancer-related tasks. However, improved model capability must be accompanied by robust validation, explainability, uncertainty estimation, privacy protection, and clinician oversight.

An Oncology Foundation Model can potentially integrate radiological imaging, histopathology, genomics, and clinical information for cancer detection, classification, segmentation, prognosis, and treatment-response prediction. Self-supervised learning can reduce dependence on extensive labeled datasets, while parameter-efficient adaptation can facilitate transfer across tasks and cancer types. Continual and federated learning can support adaptation across institutions while addressing changing clinical environments and privacy requirements.

Ultimately, predictive oncology should not be evaluated solely according to classification accuracy. Generalizability, explainability, uncertainty awareness, multimodal integration, adaptability, privacy, safety, external validation, computational efficiency, and clinical utility are equally important. The long-term objective should be to develop reliable AI systems that augment expert oncology decision-making rather than replace clinical expertise.

REFERENCES

[1] World Health Organization, "Cancer," WHO, 2026.

[2] International Agency for Research on Cancer, "Global cancer burden growing, amidst mounting need for services," WHO/IARC, 2024.

[3] H. Yang, M. Yang, J. Chen, G. Yao, Q. Zou, and L. Jia, "Multimodal deep learning approaches for precision oncology: A comprehensive review," *Briefings in Bioinformatics*, vol. 26, no. 1, 2024/2025, doi: 10.1093/bib/bbae699.

[4] "A systematic analysis of deep learning in genomics and histopathology for precision oncology," *BMC Medical Genomics*, 2024, doi: 10.1186/s12920-024-01796-9.

[5] "Deep learning in cancer genomics and histopathology," 2024.

[6] "Deep learning of radiology-genomics integration for computational oncology: A mini review," 2024.

[7] "Deep Learning Applications in Clinical Cancer Detection: A Review of Implementation Challenges and Solutions," *Molecular & Cellular Oncology/Digital-related clinical review*, 2025, doi: 10.1016/j.mcpdig.2025.100253.
[8] S. Toumaj *et al.*, "Leveraging explainable artificial intelligence for transparent and trustworthy cancer detection systems," *Artificial Intelligence in Medicine*, 2025, doi: 10.1016/j.artmed.2025.103243.
[9] "Artificial intelligence enabled tumor diagnosis and treatment: Status, breakthroughs and challenges," *Critical Reviews in Oncology/Hematology*, 2025, doi: 10.1016/j.critrevonc.2025.104963.
[10] "AI in Cancer Prognosis: A Systematic Review of Multimodal Models Combining Pathology Images and High-Throughput Omics," 2026.
[11] A. Vaswani *et al.*, "Attention is all you need," in *Advances in Neural Information Processing Systems*, 2017.
[12] K. He *et al.*, "Deep residual learning for image recognition," in *Proc. IEEE CVPR*, 2016.
[13] A. Dosovitskiy *et al.*, "An image is worth 16×16 words: Transformers for image recognition at scale," in *Proc. ICLR*, 2021.
[14] K. He *et al.*, "Masked autoencoders are scalable vision learners," in *Proc. IEEE/CVF CVPR*, 2022.
[15] T. Chen *et al.*, "A simple framework for contrastive learning of visual representations," in *Proc. ICML*, 2020.
[16] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems*, 2017.
[17] R. R. Selvaraju *et al.*, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in *Proc. IEEE ICCV*, 2017.
[18] M. T. Ribeiro, S. Singh, and C. Guestrin, "'Why should I trust you?': Explaining the predictions of any classifier," in *Proc. ACM KDD*, 2016.
[19] E. J. Hu *et al.*, "LoRA: Low-rank adaptation of large language models," in *Proc. ICLR*, 2022.
[20] B. McMahan *et al.*, "Communication-efficient learning of deep networks from decentralized data," in *Proc. AISTATS*, 2017.
[21] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," arXiv:1503.02531, 2015.
[22] J. Kirkpatrick *et al.*, "Overcoming catastrophic forgetting in neural networks," *Proceedings of the National Academy of Sciences*, vol. 114, no. 13, pp. 3521–3526, 2017.
[23] Z. Li and D. Hoiem, "Learning without forgetting," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 12, pp. 2935–2947, 2018.
[24] M. Tan and Q. V. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proc. ICML*, 2019.
[25] A. G. Howard *et al.*, "MobileNets: Efficient convolutional neural networks for mobile vision applications," arXiv:1704.04861, 2017.