REVIEW ARTICLE

open access

Artificial Intelligence and Multi-Agent Systems in Oncology: Clinical Translation, Infrastructure, and Future Directions

Abhishek Gulia[1], Rita Singh[1], Urja Khanna[3], Raghul RJ[1]

Received: 14 March 2026 | Accepted: 16 May 2026 | Published: 14 June 2026

Abstract

Background: Artificial intelligence (AI) and multi-omics integration have rapidly advanced oncology by enabling analysis of genomics, proteomics, and other biological data. Deep learning techniques now support automated tumour segmentation and drug response prediction using large cancer databases.

Objective: To review key developments, clinical applications, and challenges of AI-driven precision oncology.

Methods: A narrative overview of current literature on AI in imaging, radiomics, pharmacogenomics, federated systems, and interoperability standards.

Results: Radiomics enables quantitative tumour characterization from imaging, while machine learning improves anticancer drug sensitivity modelling. Emerging technologies include autonomous AI agents and federated multi-site systems that enable decentralized, privacy-preserving clinical implementation.

Conclusion: AI-driven multi-omics and multi-agent systems are transforming precision oncology, though challenges in **infrastructu**re, regulation, and clinical translation remain.

Keywords: Artificial intelligence • Multi-agent systems • Oncology • Radiomics • Pharmacogenomics • Federated learning • Nuclear medicine • Precision oncology • Multi-omics integration • Clinical translation

Introduction

Cancer remains a leading cause of global morbidity and mortality (as per Globocan 2022), demanding increasingly precise and scalable solutions for early detection, therapeutic stratification, and longitudinal monitoring. Traditional oncology relies heavily on clinician interpretation of imaging, histopathology, and genomic testing. However, heterogeneity in tumour biology and inter-observer variability limit reproducibility (1–4).

Artificial intelligence has emerged as a transformative tool capable of extracting high-dimensional features from imaging, genomic, and clinical datasets (5–7). Deep convolutional neural networks first demonstrated strong performance in medical image analysis, particularly with the development of U-Net architectures for biomedical segmentation (6,8–12). Radiomics further expanded this paradigm by converting images into quantitative descriptors correlated with tumour phenotype and microenvironment (5).

Simultaneously, pharmacogenomic datasets such as CCLE and GDSC enabled machine learning models to predict drug response based on genomic alterations (1–4). These approaches laid the foundation for precision oncology powered by computational modelling.

Oncology artificial intelligence research has more recently moved toward distributed intelligence and autonomous agents that can coordinate diagnostic and treatment planning. Federated learning approaches allow collaboration across hospitals without sharing patient data while respecting privacy regulations (13–15). Multi-agent AI systems can connect imaging, genomic, and clinical components within integrated decision-making ecosystems. This review aims to consolidate foundational research and emerging frameworks to provide a roadmap for AI and multi-agent systems in oncology.

Deep learning in oncologic imaging

Convolutional neural networks for tumour detection and segmentation

U-Net introduced by Ronneberger et al. was a breakthrough in image segmentation of biomedical images through the implementation of end-to-end training on small, labelled samples (6). Various applications of U-Net and its derivatives in tumour segmentation on CT, MRI, and PET images have appeared since.

Men et al. showed that deep learning could automatically segment clinical target volumes and organs at risk for rectal cancer radiotherapy treatment, allowing substantial reductions in manual contouring time with equivalent accuracy to that achieved by expert radiation oncologists (7). Sheller et al. further applied segmentation in federated settings, illustrating that brain tumour segmentation across multiple institutions could be done without any exchange of actual patient data (13). For nuclear oncology imaging applications, deep learning techniques have advanced PET attenuation and scatter correction. Gong et al. found that a CNN model could directly correct for attenuation correction without requiring a CT scan (16). Chen et al. further showed that ultra-low-dose PET image reconstruction was possible with deep learning, leading to a significant decrease in radiation exposure without affecting image quality (17).

Radiomics and quantitative imaging biomarkers

Radiomics involves the use of medical imaging to generate high-dimensional feature vectors of tumour heterogeneity, morphology, texture, and intensity patterns (5). Radiomics was introduced by Lambin et al. as an interface between medical imaging and precision medicine (5).

[1]Department of Radiation Oncology, Max Hospital, Sec 128, Noida, Uttar Pradesh, India
Correspondence: Abhishek Gulia, Department of Radiation Oncology, Max Hospital Sec 128, Noida, Uttar Pradesh, India
Email: drabhishekgulia@gmail.com
How to cite: Gulia A, Singh R, Khanna U, et al. Artificial intelligence and multi-agent systems in oncology: Clinical translation, infrastructure, and future directions. Ann Compr Oncol Biomed Sci. 2026;1(1):9–16. doi:10.66765/acobs.2026.003
DOI: https://doi.org/10.66765/acobs.2026.003

Table 1: Major Pharmacogenomic and Oncology Datasets

Dataset	*Year*	*Data Type*	*Application*	*Key Reference*
CCLE	2012	Genomic + Drug Response	Drug sensitivity prediction	1
GDSC	2013	Genomic + Pharmacologic	Biomarker discovery	2,4
TCGA	2013	Multi-omics	Molecular tumor profiling	36
AACR GENIE	2017	Clinical-genomic registry	Precision oncology validation	35

It was shown by Sollini et al. that features based on PET radiomics in NSCLC can serve as predictors for response and survival outcome (18). In another study, Werner et al. found that PRRT response in neuroendocrine tumours can be predicted using radiomics and thus guide theranostic decision-making (19).

Methodologic considerations of feature reproducibility, standardization, and external validation have been outlined in detail by Van Timmeren et al., which are important aspects in obtaining approval from regulatory authorities (20). This approach is an essential component of AI-enabled agents in oncology today.

AI IN PHARMACOGENOMICS AND DRUG RESPONSE PREDICTION

Large-scale pharmacogenomic datasets

The Cancer Cell Line Encyclopaedia (CCLE) facilitated predictive modelling of drug sensitivities using genomic information of hundreds of cancer cell lines (1).

Similarly, the Genomics of Drug Sensitivity in Cancer (GDSC) project systematically identified genomic markers associated with drug response (2, 4). Iorio et al. expanded pharmacogenomic interaction landscapes, integrating genomic features with drug response matrices to identify actionable vulnerabilities (3). These datasets catalysed precision oncology modelling (Table 1).

Machine learning models for drug sensitivity

Menden et al. applied machine learning algorithms to predict cancer cell sensitivity using genomic and chemical descriptors (10). Costello et al. organized a community-wide challenge to benchmark drug sensitivity prediction algorithms, revealing variability in model performance and emphasizing reproducibility (12).

Ding et al. demonstrated that integrating multi-omics data significantly improves predictive accuracy over single-omics approaches (11). These findings underscore the importance of multi-layered biological integration in oncology AI.

Functional Genomics and CRISPR Validation

Genome-scale CRISPR-Cas9 screening transformed target validation in oncology. Shalem et al. and Wang et al. independently demonstrated the feasibility of genome-wide knockout screens in human cells (21, 22). Hart et al. further identified genotype-specific vulnerabilities, supporting synthetic lethality approaches (23). AI-driven analysis of CRISPR datasets now accelerates identification of druggable targets.

AI agents in nuclear medicine and theranostics

Deep Learning–Enhanced PET Reconstruction and Quantification

Nuclear medicine has been one of the earliest adopters of AI-driven reconstruction methods due to its inherently noisy and low-count imaging characteristics. Gong et al. demonstrated that deep convolutional neural networks could perform direct attenuation and scatter correction in PET imaging without the need for CT-based attenuation maps, maintaining standardized uptake value (SUV) accuracy while reducing dependency on hybrid imaging (16). Similarly, Yang et al. applied deep learning frameworks for PET attenuation and scatter correction, improving image fidelity and reducing artifacts (24).

Chen et al. showed that ultra-low-dose PET imaging reconstructed via deep learning preserved lesion detectability while significantly reducing radiation exposure (17). These findings are particularly relevant in longitudinal oncology monitoring, where cumulative radiation dose is a concern. Collectively, these studies indicate that AI agents embedded within PET workflows can autonomously optimize reconstruction parameters, detect anomalies, and flag suboptimal acquisitions in real time (Table 2).

Radiomics-guided theranostics

Radiomics has shown particular promise in theranostic nuclear medicine applications. Werner et al. demonstrated that PET-derived radiomic features could predict response to peptide receptor radionuclide therapy (PRRT) in neuroendocrine tumors (19). Such predictive modelling supports individualized radionuclide dosing and patient selection.

Sollini et al. reported that PET radiomic signatures in NSCLC were associated with survival outcomes and treatment response, reinforcing the role of AI in prognostic stratification (18).

AI agents can integrate PET radiomics with genomic and clinical parameters, creating closed-loop theranostic systems capable of adaptive treatment planning.

Autonomous AI agents in clinical nuclear medicine

Going beyond mere predictive modelling to agents, the shift to AI agents requires systems that can perceive, reason, and act. Concepts related to multi-agent systems by Russell and Norvig and Wooldridge form the basis of decision-making agents in distributed settings (25, 26). The large language model-based clinical reasoning systems exemplified by GPT-4 architecture highlight capabilities of contextual reasoning in a multi-step approach suitable for cancer decision support (27). Although such systems need extensive validation, their application suggests the possibility of developing AI agents that can integrate information from imaging, genomics, and guideline recommendations.

Table 2: Key AI Applications in Oncologic Imaging

Application	*Modality*	*AI Method*	*Clinical Impact*	*Reference*
Tumor segmentation	MRI/CT	U-Net CNN	Automated contouring	6,7
PET attenuation correction	PET Deep	CNN	Quantitative accuracy	16,24
Ultra-Low Dose PET	PET	Deep Learning	Dose reduction	17
Radiomics Prognostication	PET/CT	Feature Modeling	Survival prediction	18,19

In nuclear oncology, AI agents might:

- Recommend optimal radiotracer selection
- Suggest individualized radionuclide dosing
- Monitor post-therapy imaging response
- Trigger alerts for atypical biodistribution patterns

This represents a shift from passive AI tools to active clinical collaborators.

Multi-Agent Architectures in Oncology

Conceptual framework of multi-agent oncology ecosystems

The development of AI applications in the field of oncology has been moving more towards integrated systems of intelligent agents rather than standalone prediction models. A Multi-Agent System (MAS) involves intelligent agents working within an environment in order to perform tasks that could not be completed by one single agent and has its theoretical foundations based on distributed AI systems (8, 9). In oncology, a multi-agent system could be conceived as a hierarchical framework of intelligent agents, where each would represent different aspects of cancer intelligence.

The Imaging Agent combines methods for segmentation using convolutional neural networks (6, 7) with radiomic feature extraction workflows (5, 18, 19). Such an agent executes perceptual tasks of lesion detection, volumetric segmentation, measurement of heterogeneity, and longitudinal follow-up of response. Notably, radiomics obtained using PET and CT scans have shown to be significant in predicting outcomes in various cancers (18, 19); hence, they contribute to the phenotypic layer in the MAS framework.

Molecular Interpretation Layer comprises the actions of the Genomic Agent. Variant annotation tools such as ANNOVAR (28), SIFT (29), PolyPhen-2 (30), FATHMM (31), and CADD (32) assess the pathogenicity of variants on the sequence level. They analyse amino acid conservation, structural disruption ability, evolutionary constraint, and scoring of deleteriousness. Within the MAS paradigm, these computational approaches act as sub-agents converting the results of sequencing experiments into actionable information on the functional effects of genetic variants.

The Pharmacogenomic Agent combines large-scale response matrices based on data from platforms like CCLE (1) and GDSC (2, 4) with predictive algorithms for machine learning. Using genomic changes to map out landscapes of therapeutic sensitivities (3), the agent predicts probabilistic responses to various treatment strategies including targeted drugs, immunotherapy, or cytotoxic approaches. Crucially, multi-modal prediction has been shown to be more accurate than unimodal prediction (11), making integration mandatory.

The Clinical Decision Agent combines information from the imaging, genomic, and pharmacologic agents, and correlates these with guideline-based oncology recommendations. The Clinical Decision Agent operates on a reasoning level, carrying out risk stratification and therapy ranking and prediction of toxicity. Reasoning in clinical decision agents is akin to reasoning performed by state-of-the-art language models (25).

Lastly, the Infrastructure Agent manages orchestration, security, and federated coordination. Through federated learning techniques (13, 14), along with methods to ensure privacy (41–42), the layer facilitates learning across institutions without centralizing the data. The layer also enables interoperability through standards such as FHIR (40, 41).

Whereas conventional AI systems have been monolithic, multi-agent systems allow for scalable modules. These agents can be updated, validated, and audited independently yet in a cohesive manner. This is important for regulatory compliance, because the validation process may apply to each individual agent but the decision-making logic will still be transparent.

Systems-level integration and biological convergence

The pathology of cancer is multifaceted with features including genomic instability, aberrant epigenetics, transcriptomic rewiring, metabolism, and microenvironmental diversity. The use of AI methods with single-modalities does not address this multidimensional disease architecture. Therefore, integration of multi-agents is not only a computational solution but also biologically driven.

Hasin et al. noted that multi-omics integration would offer better disease architecture (33). Ritchie et al. explained that there were computational tools that could analyse genotypic and phenotypic connections in different data domains (35). This knowledge underlies the development of multi-agent oncology systems.

In a multi-agent system, multiple biological signals are analysed simultaneously: Radiomic signatures represent spatial heterogeneity and have been linked to hypoxia, necrosis, and proliferation gradients (5, 18).

Transcriptomic signatures reveal oncogenic signalling pathways and states of the immune microenvironment (33). CRISPR-based functional genomics approaches confirm gene essentiality and synthetic lethal interactions (21–23). Large-scale pharmacogenomic data can construct drug response landscapes (1–3).

Through computational fusion, the system shifts from predictive modelling based on correlation to that which is based on mechanism of action. For instance, a cancerous growth with high radiomic heterogeneity and confirmed DNA repair pathway dependency using CRISPR along with pharmacogenomic evidence of PARP sensitivity would suggest a more informed suggestion regarding treatment.

This type of fusion demands significant computational power. Concepts such as cloud computing proposed by Armbrust et al. (36) allow flexible use of computing resources and large scale modelling, while frameworks such as MapReduce (37) allow large-scale parallel analysis of multi-institutional data sets. These advances facilitate real time or near-real time adaptive learning through federated oncology networks.

Dynamic agent communication and adaptive learning

Communication among agents is an important attribute of multi-agent systems. In MAS models in oncology, output from one agent can influence the decision criterion of the other. For example,

- If the genomic agent discovers high tumour mutation burden, the decision criterion for clinical decisions will be adjusted toward immunotherapy. A radiomic progression signal may trigger re-evaluation of drug sensitivity modelling.
- Federated updates may recalibrate prediction coefficients based on newly aggregated global evidence.

The architecture driven by feedback allows for the development of adaptive intelligence, not merely predictions. Federated learning protocols (13–15) permit independent organizations to iteratively improve a common model without revealing sensitive patient data. Mechanisms for preserving privacy include differential privacy (41) and secure aggregation (42), which adhere to regulatory standards and maintain statistical reliability.

Consequently, the oncology multi-agent system transitions from a sequential pipeline process to a sustainable ecosystem that can be updated continually.

Clinical and regulatory implications of multi-agent systems

Multi-Agent Architectures from the standpoint of translational research have several benefits:

1. Improved explainability with modular outputs (imaging score, genomic score, pharmacologic score).

2. Verification of agent performance independently.
3. Scalable deployment in a federated manner across multiple institutions.
4. Reduction of single model bias.

The regulatory authorities have begun to focus on accountability and life-cycle management of devices powered by AI technologies (43, 44). The modular MAS architecture fulfils such requirements, with independent verification of each agent's performance metrics.

- Nonetheless, there are challenges associated with the added complexity, which include: Inter-agent dependency risk
- Computational overhead
- Synchronization latency in federated settings
- Standardization of cross-agent data ontologies

Addressing these issues will require collaborative efforts across oncology, computer science, bioinformatics, and regulatory sciences.

Toward autonomous oncology intelligence

The future course for such multi-agent systems in oncology will be that of semi-autonomous support in clinical settings. Although human intervention is still essential, AI agents will possibly assume roles in:

- Simulating treatment strategies
- Predicting outcomes over time
- Dynamic modification of treatments
- Re-calibrating models worldwide

This would mark a complete transition from analytics assistance to intelligent ecosystems.

Federated learning and distributed oncology intelligence

Foundational federated learning models

McMahan et al. introduced federated averaging algorithms enabling decentralized training of neural networks without centralized data pooling (13). Konečný et al. expanded communication-efficient strategies for distributed model training (14).

Sheller et al. demonstrated multi-institutional brain tumour segmentation using federated learning, achieving comparable accuracy to centralized training without sharing patient data (15). Rieke et al. described the transformative potential of federated learning in digital health ecosystems (38). These frameworks are particularly relevant in oncology, where data privacy regulations and institutional policies restrict data exchange (Table 3).

Oncology genomic networks

Large collaborative networks such as The Cancer Genome Atlas (TCGA) and AACR Project GENIE have demonstrated the power of multi-centre genomic aggregation (39, 40). However, traditional models rely on centralized repositories. Federated AI extends this paradigm by allowing model parameter sharing rather than raw genomic data, enabling scalable precision oncology while maintaining compliance with data protection regulations.

Privacy-Preserving AI

Differential privacy frameworks introduced by Dwork and later adapted to deep learning by Abadi et al. provide mathematical guarantees against patient re-identification (41, 42). Kaissis et al. further described secure federated machine learning approaches in medical imaging (39). Privacy-preserving AI is essential for real-world deployment in oncology.

Interoperability and clinical integration

Clinical AI deployment requires interoperability with hospital information systems. SMART on FHIR enables standards-based application integration within electronic health records (43). HL7 FHIR provides RESTful architectures for healthcare data exchange (44, 45). Warner et al. demonstrated integration of genomic data into clinical workflows using standardized frameworks (46). Without interoperability, AI systems remain siloed experimental tools rather than deployable clinical infrastructure.

Clinical translation and regulatory landscape

From algorithmic performance to clinical utility

The process of transferring artificial intelligence technology for oncology from experimental modelling to actual clinical application is arguably among the greatest challenges facing translational medicine in the present day. While there are many instances where an AI tool displays excellent performance when measured through retrospective data sets, clinical translation of the technology entails more than simply good AUC scores.

In fact, the clinical translation process demands comprehensive validation, with assurances that the model's performance remains robust when measured on independent data sets generated using varying imaging and sequencing technologies. Radiomics models, for instance, have proven vulnerable to changes in acquisition and reconstruction parameters (5, 20), highlighting the need for harmonization prior to clinical implementation.

Just as important, however, is multi-centre reproducibility, which assesses whether algorithm performance is consistent among disparate geographical and institutional populations. In this regard, federated validation methods have demonstrated potential through their capacity for collaborative study without data aggregation (13–15). Indeed, such methods are especially pertinent to oncology because of the effects of tumour heterogeneity and population diversity on treatment efficacy.

The gold standard of validation is prospective validation. In contrast to retrospective validation, prospective methods measure not only predictive accuracy but also clinical engagement, integration into clinical workflow, impact on decision-making, and influence on patient outcomes. Otherwise, AI technology runs the risk of remaining purely academic in its application.

Human–AI synergy in oncology practice

Incorporating AI in oncology should be viewed as an enhancement rather than a replacement for human expertise in oncology. According to Topol, collaboration between human intelligence and AI is the basis of high-performance medicine

Table 3: Federated Learning in Oncology

Study Modality	*Use case*	*Multicenter*	*Privacy preserved*	*Reference*
McMahan et al.	FL algorithm	Yes	Yes	13
Sheller et al.	Brain tumor segmentation	Yes	Yes	15
Rieke et al.	Digital health framework	Yes	Yes	34
Kaissis et al.	Secure FL imaging	Yes	Yes	39

(43). Decisions in oncology involve intricate ethical questions, patients' preferences and multidisciplinary collaboration. Thus, the use of AI should be seen as assistance and not a substitute for human decisions.

Human-AI collaboration includes:

- Clear visual interpretation
- Confidence scores
- Contextual recommendation rankings
- Transparent uncertainty quantifications

In a review of obstacles to effective implementation of clinical AI, Kelly et al. highlighted bias amplification, inadequate generalization, insufficient incorporation in practice and mistrust among stakeholders (44). The latter is especially relevant for oncology, where the lack of equal representation of genomics or imaging data may lead to discriminatory treatment decisions. Therefore, explainability and transparency become ethical considerations.

Regulatory requirements and AI as a medical device

AI-powered oncology technologies have moved into the category of Software as a Medical Device (SaMD) in regulatory terms. The emerging international consensus around AI regulation is based on the following guiding principles:

1. Model architecture transparency
2. Transparency about training data
3. Bias detection and evaluation
4. Explanations of AI models
5. Life cycle performance assessment

Contrary to classical, static medical devices, AI technologies that implement federated learning and continuous updates are dynamic systems. This makes them susceptible to issues like "algorithm drift," wherein model parameters change over time. In consequence, post-market surveillance is crucial to identifying the decline in performance and potential bias development. Predictive radiomics and pharmacogenomics require more attention because of the involvement of complex high-dimensional features in the analysis. Proposed digital health governance frameworks stress the importance of life cycle regulation, taking into account validation, deployment, and real-time auditing of AI systems (43,44).

Federated learning and regulatory scalability

Federated learning provides an exciting potential for regulatory science. By allowing validation of models across multiple centres without exchanging raw data (13–15), federated learning systems increase their generalizability testing capacity. Additionally, increased statistical power is achieved through distributional validation while maintaining data privacy. Data privacy methods like differential privacy (41) and secure aggregation protocols (42) enable federated learning technologies to become fully compliant with data protection laws. It should also be noted that the implementation of federated validation might decrease overfitting and improve performance indicators by testing the model's performance on various data samples. The application of multi-agent system technology provides even more opportunities for the field. The modularity of such a system makes it possible to test different agents separately.

For instance:

- The imaging agent validation process could be focused on segmentation performance;
- The genomic agent performance indicator could be concordance of pathogenicity predictions; and
- The pharmacogenomic agent performance metric could include drug response prediction correlation.

Ethical, legal, and implementation challenges

Although technology may be ready for translation, several barriers still remain, including:

- Ownership issues of data among different organizations
- Liability in decision-making supported by AI systems
- Cybersecurity issues in distributed networks
- Necessity of workforce training
- Resistance to process disruptions

Ethical aspects of autonomous or semi-autonomous agents in oncology need careful consideration. Systems of decision support should retain clinician autonomy and patient-focused approach. Moreover, algorithmic biases based on underrepresented genomics data sets of particular ethnic groups may create an unequal impact on treatment decisions. As such, implementation of science is as important as algorithm development, and successful translation requires collaboration between many experts.

Toward evidence-based AI oncology

For the future development of AI in clinical oncology, it is essential to develop the following standards for reporting:

- Documented datasets
- Benchmarking for reproducibility
- Performance registries in real-world settings
- Benchmarking for comparative effectiveness studies

Eventually, federated and multi-agent systems could act as the infrastructure for global networks of AI benchmarking, facilitating ongoing evidence generation. Here, oncology AI changes its status from that of isolated predictive algorithms into a controlled system of continuous learning, based on transparency and reproducibility.

Challenges and Limitations

Dataset heterogeneity and biological variability

In spite of advancements in AI for oncology, one of the most recurrent issues in applying it to clinical practice is the problem of data heterogeneity. Datasets differ in terms of scanner manufacturers, methods used for image reconstruction, and acquisition settings as well as contrast protocols used for image capture. In turn, radiomic feature extraction depends greatly on voxel resolution, segmentation, and preprocessing pipeline (5,20). Genomic datasets created on various sequencers, tools for variant detection, and annotation schemes can introduce some bias in mutation identification. In vitro datasets available at public repositories like CCLE and GDSC (1–4) only provide information on cells cultured outside a patient's body. Such experiments cannot fully reproduce tumour microenvironment, immune system interaction, or pharmacodynamics in vivo. Heterogeneity within biology is also a contributing factor that adds to computational heterogeneity. The intratumor clonal variability, evolution through therapy pressures, and microenvironment interactions create a problem for models that work only at a static snapshot level. Such AI algorithms trained on a static snapshot level will not be able to represent the dynamic evolution of tumours.

Imaging protocol variability and reproducibility constraints

Reproducibility is still an essential problem for imaging artificial intelligence. Differences in slice thickness, reconstruction kernel types, PET uptake period, and attenuation correction procedures affect quantization significantly (5,16,20). Artificial intelligence models trained using one protocol could be sensitive to changes in acquisition protocols. Furthermore, segmentation procedures are prone to user variation, especially if contours are generated manually (6,7). Although automated segmentations help minimize interobserver differences, such approaches depend on

the characteristics of training datasets. Multicentre evaluation and federated training approaches (13–15) address some challenges associated with image variability by exposing models to various input types during the training process. Still, institutional heterogeneity could result in gradients' divergence and model instabilities in cases where training data isn't IID (13,14). In other words, methodological robustness should take priority over optimal performance at single centres.

Model interpretability and the "black box" problem

Interpretability continues to be an ethical issue within oncology AI research. Deep learning models, especially ones combining imaging and multi-omics features, are frequently high-dimensional nonlinear functions which cannot be easily explained. While genomic variant annotations, such as ANNOVAR (28), SIFT (29), PolyPhen-2 (30), FATHMM (31), and CADD (32), offer pathogenicity scores that are interpretable, multi-agent models may complicate the interpretation of input-output relationships. The lack of interpretability can be detrimental for clinicians and regulatory agencies' trust in the system. In the field of oncology, where recommendations could affect the course of patients' lives, such recommendations are unethical. Techniques for Explainable Artificial Intelligence (XAI), including feature attribution and uncertainty analysis, should thus be part of the initial design process and not an afterthought.

Algorithmic bias and health disparities

The structure of AI models is inevitably influenced by the datasets used for training. Inadequate representation of specific ethnicity, geographic regions, or socioeconomic groups in genomics and radiologic data can result in unfair therapeutic suggestions. Databases of pharmacogenomics (1–4) include information that is mostly obtained from limited ancestry samples. In a similar way, datasets used for imaging (13–15) tend to under-represent low-resource areas. Such issues may aggravate the problem of inequality within the field of precision oncology. A solution can be sought through implementing federated machine learning techniques that allow collecting geographically distributed data. Yet, at the same time, federated models may perpetuate institutional biases if properly audited. Metrics that measure algorithmic fairness, as well as demographic stratifications, should become a standard component in oncological AI studies.

Cybersecurity and infrastructure vulnerabilities

Distributed computing and federated learning create cybersecurity weaknesses. Inversion attacks, data reconstruction attacks, and adversarial attacks present possible threats that might impact patient confidentiality. Differential privacy techniques (41) and secure aggregation algorithms (42) offer cybersecurity through mathematics, although there may be trade-offs related to computational efficiency. Moreover, adversarial attacks on AI for imaging systems may cause changes in the lesions' detection and/or segmentation. As AI systems in oncology move towards becoming interconnected multiagent systems, cybersecurity will become just as important as efficiency.

Ethical concerns in autonomous decision systems

As the role of artificial intelligence agents evolves from that of helpers to semi-independent decision-making systems, the ethical situation becomes more complicated. Some of the issues include:

- Responsibility in cases of treatment mistakes involving artificial intelligence
- Informed consent for therapy choice via artificial intelligence
- Affording patient autonomy using computer assistance

Human involvement is essential. An artificial intelligence oncology framework should incorporate human supremacy and computational power.

Benchmarking and standardization

Strong benchmarking paradigms will continue to be indispensable for addressing fragmentation. Challenges driven by community initiatives, like those developed by Costello et al. in drug sensitivity prediction (12), demonstrate the need for comparison against a set of standards.

Future benchmarking exercises should include:

- Federated evaluation between multiple institutions
- Multi-omics analysis tasks
- Prognostic predictions
- Fairness audits

Without standardized validation pipelines, there is a risk that AI in oncology will become a fragmented field full of disconnected prototypes.

Future directions

The future developments in AI for oncology will probably be defined not by predictive standalone algorithms but by intelligent AI ecosystems working together. There appear new trends in developing autonomous theranostic AI, which would allow for integrating radiomic phenotypes of tumors (18,19) into pharmacogenomic landscapes of response (1–3) and gene dependencies confirmed with CRISPR technology (21–23). This way, AI can contribute to dynamically optimizing treatments depending on the changes in imaging and genetic data. Within the domain of radiation oncology, innovations in automated segmentation (6,7) and dose modelling technologies seem to offer the possibility of adapting treatments according to tumour regression and toxicity profiles. These can be globally optimized within federated infrastructures (13–15). Furthermore, artificial intelligence holds promise for enhancing functional genomics discovery. The integration of CRISPR screening datasets with multi-omics modelling platforms (33,35) may identify context-specific vulnerabilities and synthetic lethalities, accelerating the process of target identification and drug repurposing. A more radical approach would be to develop continuous federated learning networks. Instead of using static algorithms, oncology AI systems of the future can continuously learn by updating their models with contributions from multiple organizations around the globe. These AI agents could function as global oncology intelligence platforms akin to genomics consortia (36,35) but capable of running in near real time. The most revolutionary development would be the fusion of radiomics, pharmacogenomics, digital pathology, and federated learning capabilities. By combining phenotypic imaging biomarkers (5,18) with molecular pathways derived from transcriptomics and genomics models (33), oncology AI can get closer to mechanistic precision medicine. The ultimate goal is not the displacement of physicians by AI algorithms but rather the formation of a continuously learning, collaboration-driven system where AI agents provide actionable insight, while humans provide interpretation and ethical guidance.

Conclusion

AI has been quickly altering the world of oncology, transitioning from fragmented predictive algorithms into comprehensive, multiple layer intelligent systems that can support complicated decision-making processes. All improvements achieved in terms of deep learning-based image analysis, radiomics,

pharmacogenomic predictions, CRISPR-based functional genomics, and federation frameworks altogether suggest that a new era of cancer diagnosis, characterization, and treatment approaches has already come upon us. Throughout this manuscript, we have emphasized that the real revolution brought about by AI will not be the result of individual single-modality prediction methods, but rather in the integration of multiple data streams into a multi-agent intelligent system framework. An imaging phenotype, a genetic alteration, a transcriptional activity in certain pathways, and sensitivity to a variety of drugs all represent different aspects of tumour biology. The fusion of these components in multi-agent frameworks creates an intelligent decision-making ecosystem that enables personalized oncology. Equally important, though, translation into clinical practice needs science, reproducibility, transparency, and appropriate governance. Data heterogeneity, algorithmic bias, interpretation limitations, and regulation continue to be obstacles. Federated machine learning and privacy-preserving approaches provide exciting possibilities for geographical diversity in model validation along with data protection for patients. Still, the responsible application will need joint efforts of oncologists, computational biologists, ethicists, and regulators. The future of AI in oncology is expected to involve continuously improving multi-agent systems combining the advantages of radiomics, pharmacogenomics, and decentralized computing into the daily workflow of clinicians. Instead of displacing the expertise of practitioners, multi-agent AI technology should operate as an integral part of cognitive work assisting clinicians' judgement through high-dimensional analytics. It is worth mentioning that in this era when we see the confluence of AI technology and oncology, what is really happening is that there occurs a paradigm shift in oncological practice towards adaptive data-driven precision cancer medicine. The success of this paradigm shift will hinge on methodology, ethics, and international cooperation. Provided all the conditions are met, multi-agent AI systems could lay the foundation for next-generation cancer precision treatment.

Ethical Approval

Conflicts of interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Ethical approval statement

This article is a review study based exclusively on previously published literature. Ethical approval was therefore not required, as no human participants, animals, or identifiable personal data were involved.

Informed consent statement

Informed consent was not required for this study because it does not involve human participants, patient data, or identifiable personal information.

Author contributions

Abhishek Gulia: Conceptualization, Writing – Original Draft Preparation, Supervision; Urja Khanna: Literature Review; Raghul RJ: Formal Analysis; Rita Singh: Writing – Review & Editing.

All authors have read and approved the final version of the manuscript.

Data availability statement

No new data was generated or analyzed in this study. Data sharing is therefore not applicable to this article.

Funding statement

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Acknowledgements

The authors declare that no funding, technical, administrative, or institutional support was received for this work. No acknowledgements apply.

AI & AI-assisted tool disclosure

The authors declare that AI-assisted tools, specifically ChatGPT, were used only for language editing and grammatical improvement of the manuscript. The authors take full responsibility for the content, interpretation, and originality of the work.

Orcid ID

- **Abhishek Gulia:** https://orcid.org/0000-0002-6635-785X
- **Rita Singh:** https://orcid.org/0009-0004-6428-963X
- **Urja Khanna:** https://orcid.org/0009-0001-8019-7768
- **Raghul RJ:** https://orcid.org/0009-0008-1354-1195

References

1. Barretina J, Caponigro G, Stransky N, et al. The Cancer Cell Line Encyclopedia enables predictive modelling of anticancer drug sensitivity. Nature. 2012;483:603–607.
2. Garnett MJ, Edelman EJ, Heidorn SJ, et al. Systematic identification of genomic markers of drug sensitivity in cancer cells. Nature. 2012;483:570–575.
3. Iorio F, Knijnenburg TA, Vis DJ, et al. A landscape of pharmacogenomic interactions in cancer. Cell. 2016;166:740–754.
4. Yang W, Soares J, Greninger P, et al. Genomics of Drug Sensitivity in Cancer (GDSC): A resource for therapeutic biomarker discovery. Nucleic Acids Res. 2013;41:D955–D961.
5. Lambin P, Rios-Velazquez E, Leijenaar R, et al. Radiomics: Extracting more information from medical images using advanced feature analysis. Eur J Cancer. 2012;48:441–446.
6. Ronneberger O, Fischer P, Brox T. U-Net: Convolutional networks for biomedical image segmentation. In: MICCAI 2015. 2015. p. 234–241.
7. Men K, Dai J, Li Y. Automatic segmentation of the clinical target volume and organs at risk for rectal cancer using deep learning. Radiother Oncol. 2017;124:542–548.
8. Russell S, Norvig P. Artificial Intelligence: A Modern Approach. 3rd ed. Pearson; 2010.
9. Wooldridge M. An Introduction to MultiAgent Systems. 2nd ed. Wiley; 2009.
10. Menden MP, Iorio F, Garnett M, et al. Machine learning prediction of cancer cell sensitivity to drugs based on genomic and chemical properties. PLoS One. 2013;8:e61318.
11. Ding MQ, Chen L, Cooper GF, et al. Precision oncology beyond targeted therapy: Combining omics data with machine learning. Sci Rep. 2018;8:1197.
12. Costello JC, Heiser LM, Georgii E, et al. A community effort to assess and improve drug sensitivity prediction algorithms. Nat Biotechnol. 2014;32:1202–1212.
13. McMahan HB, Moore E, Ramage D, et al. Communication-efficient learning of deep networks from decentralized data. In: Proc AISTATS. 2017;54:1273–1282.
14. Konečný J, McMahan HB, Yu FX, et al. Federated learning: Strategies for improving communication efficiency. arXiv. 2016;1602.05629.
15. Sheller MJ, Reina GA, Edwards B, et al. Multi-institutional deep learning modeling without sharing patient data: A feasibility study on brain tumor segmentation. In: Brainlesion Workshop MICCAI. 2018;11139:92–104.
16. Gong K, Guan J, Kim K, et al. Direct attenuation and scatter correction of PET images using deep learning. Phys Med Biol. 2018;63:225017.

17. Chen KT, Gong E, de Carvalho Macruz FB, et al. Ultra-low-dose PET imaging using deep learning. Eur J Nucl Med Mol Imaging. 2019;46:1807–1818.
18. Sollini M, Cozzi L, Antunovic L, et al. PET radiomics in NSCLC: Clinical implications and future directions. Eur J Nucl Med Mol Imaging. 2017;44:2103–2114.
19. Puranik AD, Dromain C, Fleshner N, et al. Target heterogeneity in oncology: the best predictor for differential response to radioligand therapy in neuroendocrine tumors and prostate cancer. Cancers (Basel). 2021;13(14):3607.
20. van Timmeren JE, Cester D, Tanadini-Lang S, et al. Radiomics in medical imaging—how-to guide and critical reflection. Insights Imaging. 2020;11:91.
21. Shalem O, Sanjana NE, Hartenian E, et al. Genome-scale CRISPR-Cas9 knockout screening in human cells. Science. 2014;343:84–87.
22. Wang T, Wei JJ, Sabatini DM, et al. Genetic screens in human cells using the CRISPR-Cas9 system. Science. 2014;343:80–84.
23. Hart T, Brown KR, Sircoulomb F, et al. High-resolution CRISPR screens reveal fitness genes and genotype-specific vulnerabilities. Cell. 2015;163:1515–1526.
24. Yang J, Sohn JH, Behr SC, et al. CT-less direct correction of attenuation and scatter in the image space using deep learning for whole-body FDG PET. Radiol Artif Intell. 2020;3(2):e200137.
25. OpenAI. GPT-4 technical report. 2023.
26. Hasin Y, Seldin M, Lusis A. Multi-omics approaches to disease. Genome Biol. 2017;18:83.
27. Ritchie MD, Holzinger ER, Li R, et al. Methods of integrating data to uncover genotype–phenotype interactions. Nat Rev Genet. 2015;16:85–97.
28. Wang K, Li M, Hakonarson H. ANNOVAR: Functional annotation of genetic variants. Nucleic Acids Res. 2010;38:e164.
29. Ng PC, Henikoff S. SIFT: Predicting amino acid changes that affect protein function. Nucleic Acids Res. 2003;31:3812–3814.
30. Adzhubei IA, Schmidt S, Peshkin L, et al. PolyPhen-2: Predicting damaging missense mutations. Nat Methods. 2010;7:248–249.
31. Shihab HA, Gough J, Cooper DN, et al. Predicting the functional, molecular, and phenotypic consequences of amino acid substitutions using hidden Markov models. Hum Mutat. 2013;34(1):57–65.
32. Kircher M, Witten DM, Jain P, et al. CADD: A framework for estimating pathogenicity of human genetic variants. Nat Genet. 2014;46:310–315.
33. Huang S, Chaudhary K, Garmire LX. More is better: Recent progress in multi-omics data integration. Front Genet. 2017;8:84.
34. Rieke N, Hancox J, Li W, et al. The future of digital health with federated learning. NPJ Digit Med. 2020;3:119.
35. AACR Project GENIE Consortium. AACR Project GENIE: Powering precision medicine through an international consortium. Cancer Discov. 2017;7:818–831.
36. The Cancer Genome Atlas Research Network. Comprehensive molecular portraits of human tumors. Nature. 2013;500:61–68.
37. Dwork C. Differential privacy. In: ICALP 2006. 2006;4052:1–12.
38. Abadi M, Chu A, Goodfellow I, et al. Deep learning with differential privacy. In: ACM CCS. 2016:308–318.
39. Kaissis GA, Makowski MR, Rückert D, et al. Secure, privacy-preserving and federated machine learning in medical imaging. Nat Mach Intell. 2020;2:305–311.
40. Mandel JC, Kreda DA, Mandl KD, et al. SMART on FHIR: A standards-based platform for substitutable apps in healthcare. J Am Med Inform Assoc. 2016;23:899–908.
41. Bender D, Sartipi K. HL7 FHIR: An agile and RESTful approach to healthcare information exchange. In: IEEE CBMS. 2013:326–331.
42. Warner JL, Jain SK, Levy MA. Integrating cancer genomic data into clinical workflows using standards-based approaches. J Am Med Inform Assoc. 2016;23:1223–1227.
43. Topol EJ. High-performance medicine: The convergence of human and artificial intelligence. Nat Med. 2019;25:44–56.
44. Kelly CJ, Karthikesalingam A, Suleyman M, et al. Key challenges for delivering clinical impact with AI. BMC Med. 2019;17:195.
45. Armbrust M, Fox A, Griffith R, et al. A view of cloud computing. Commun ACM. 2010;53:50–58.
46. Dean J, Ghemawat S. MapReduce: Simplified data processing on large clusters. Commun ACM. 2008;51:107–113.